

Giles Colclough
Magdalen College

Doctor of Philosophy

Oxford Centre for Human Brain Activity
Oxford Centre for the Functional MRI of the Brain
Centre for Doctoral Training in Healthcare Innovation

Supervisors

Professors Mark Woolrich and Stephen Smith

Methods for Modelling Human Functional Brain Networks with MEG and fMRI

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

Trinity Term 2016

COLOPHON

This thesis was typeset by the author with LuaTeX and BibTeX, using the `microtype` and `fontspec` packages.

The main text uses the Linux Libertine typeface. Mono-spaced text is Raph Levien's `Inconsolata`. The format was based on Keith Gillow and John McManigle's `ociamthesis`, André Miede's `classicthesis` and Karsten Tinnefeld's `quotchap`.

© Giles Colclough, October 2016.

To my parents,

Christopher and Sarah

Contents

PREFACE v

A NOTE ON AUTHORSHIP vi

PEER-REVIEWED PUBLICATIONS AND TALKS vii

ABSTRACT x

1 INTRODUCTION 1

2 A SYMMETRIC, MULTIVARIATE LEAKAGE CORRECTION FOR MEG 29

3 HOW RELIABLE ARE MEG RESTING-STATE CONNECTIVITY METRICS? 72

4 THE HERITABILITY OF MULTI-MODAL FUNCTIONAL CONNECTIVITY 102

5 A HIERARCHICAL MODEL FOR FUNCTIONAL NETWORK ANALYSIS 126

6 CONCLUSION 178

A PROBABILITY DISTRIBUTIONS 186

B INFERENCE FOR THE HIERARCHICAL MODEL 188

LIST OF FIGURES 210

LIST OF TABLES 211

LIST OF ABBREVIATIONS 212

REFERENCES 214

Preface

Writing a thesis affords one the rare opportunity to look back over a large body of work, and to draw together the most interesting ideas into a cohesive whole. Such a process encourages reflection on the many people who have, whether directly or indirectly, enabled me to reach this stage, and it is these people I would like to thank now.

It has been an immense privilege to work in the strongly collaborative environments of OHBA and FMRIB. Particular mention should go to Malcolm Proudfoot, Duncan Astle, Angus Stevner, Verena Heise and Ryszard Auksztulewicz, for their enthusiasm in trying my methods and their patience as we got them to work. Special thanks, too, to the analysis team at OHBA, and the MEG team at Nottingham, whom I have used as a sounding-board for all of my ideas. I am also grateful to the useful comments of several anonymous reviewers.

The work in this thesis would not have been possible without financial support from EPSRC's Centre for Doctoral Training in Healthcare Innovation, the Guarantors of Brain, and Magdalen college. I am grateful also to Pedro Valdes-Sosa, Pedro Ariel and all at CNEURO for hosting me in Havana.

My research has been greatly facilitated by the free provision of neuroimaging software and data. Thanks to Elizabeth and Peter Liddle, Sîan Robson, Matthew Brookes, Angus Stevner and Malcolm Proudfoot for sharing their datasets. The Human Connectome Project is a magnificent resource, and has provided much of the data used in this thesis. It is funded by the 16 NIH institutes and centres that support the NIH Blueprint for Neuroscience Research and by the McDonnell Center for Systems Neuroscience at Washington University. My work builds on and adds to the freely-available software packages `SPM12` from the FIL neuroimaging centre at UCL, the HCP's `Workbench`, the Donders Institute's `FieldTrip`, FMRIB's `FSL`, OHBA's `OSL` and Xu Chen's `APACE`.

There are several individuals to whom I owe a great deal of thanks:

Firstly, to my supervisors, Mark and Steve, who have shown exceptional guidance, support, care, criticism, interest and friendship. It is principally because of you that this has been such an enjoyable experience.

To Zita, Andrew, Sam, Malcolm and Adam, for reliably providing wit, antics, punishing cycle races and too much rum.

To Oli—our paths have been so close for so long that, in many ways, it feels like we've done this together.

To Mothusi, for all of your support.

And to my parents, who have always been my mentors and my role models, for giving me the freedom to explore my ideas.

A Note on Authorship

All of the work in this thesis is original and mine, unless otherwise indicated.

Some chapters use substantial material from multi-author papers that have either been published (Colclough *et al.*, 2015, 2016*e*), are under review (Colclough *et al.*, 2016*d*), or were presented at conferences (Colclough *et al.*, 2016*a*,*b*,*c*). In each case I was the first author on the paper; and the entirety of the research presented can be considered as jointly conceived with my supervisors, but led and conducted by me; I wrote the software, performed the analyses and created the figures; and I wrote the text with minor contributions from the other collaborative authors. The only exception is the analysis of the heritability of cortical curvature presented in chapter 4, which was suggested by Stam Sotiropoulos. I have retained the use of the first person plural in these chapters; its use can be understood within this context.

Peer-Reviewed Publications and Talks

Publications

Colclough, G.L., Brookes, M.J., Smith, S.M. & Woolrich, M.W. (2015) 'A symmetric multivariate leakage correction for MEG connectomes.' *NeuroImage*, **117**, 439–448.

Colclough, G.L., Woolrich, M.W., Tewarie, P.K., Brookes, M.J., Quinn, A.J. & Smith, S.M. (2016) 'How reliable are MEG resting-state connectivity metrics?' *NeuroImage*, **138**, 284–293.

Astle, D.E., Barnes, J.J., Baker, K., Colclough, G.L. & Woolrich, M.W. (2015) 'Cognitive training enhances intrinsic brain connectivity in childhood.' *Journal of Neuroscience*, **35**(16), 6277–6283.

Barnes, J.J., Woolrich, M.W., Baker, K., Colclough, G.L. & Astle, D.E. (2016) 'Electrophysiological measures of resting state functional connectivity and their relationship with working memory capacity in childhood.' *Developmental Science*, **19**(1), 19–31.

Cousijn, H., Tunbridge, E., Rolinski, M., Wallis, G., Colclough, G.L., Woolrich, M.W., Nobre, A.G. & Harrison, P.J. (2015) 'Modulation of hippocampal theta and hippocampal-prefrontal cortex function by a schizophrenia risk gene.' *Human Brain Mapping*, **36**, 2387–2395.

O'Neill, G.C., Tewarie, P.K., Colclough, G.L., Gascoyne, L.E., Hunt, B.A., Morris, P.G., Woolrich, M.W. & Brookes, M.J. (2016) 'Measurement of dynamic task related functional networks using MEG.' *NeuroImage*. In Press.

Proudfoot, M., Rokenkohl, G., Quinn, A., Colclough, G.L., Wuu, J., Talbot, K., Woolrich, M.W., Benatar, M., Nobre, A. & Turner, M.R. (2016) 'Altered cortical beta-band oscillations reflect motor system degeneration in amyotrophic lateral sclerosis.' *Human Brain Mapping*. In Press.

Pending Publications

Colclough, G.L., Smith, S.M., Nichols, T.E., Winkler, A.M., Sotiropoulos, S.N., Glasser, M.F., Van Essen, D.C. & Woolrich, M.W. (2016) 'The heritability of multi-modal connectivity in human brain activity.' In revision, *e*LIFE.

Peer-Reviewed Conference Abstracts

Colclough, G.L., Luckhoo, H., Brookes, M.J., Smith, S.M. & Woolrich, M.W. (2014) 'Performing network connectivity analyses between regions of interest using magnetoencephalography.' In *The 28th International Conference on Biomagnetism.* biomag, Halifax, Canada.

Colclough, G.L., Brookes, M.J., Smith, S.M. & Woolrich, M.W. (2015) 'A multivariate leakage correction for meg connectomes enhances between-session consistency of network estimation with HCP data.' In *The 21st annual meeting of the organization for human brain mapping.* ohbm, Honolulu, hi.

Colclough, G.L., Smith, S.M., Nichols, T.E., Glasser, M.F., Van Essen, D.C. & Woolrich, M.W. (2016*a*) 'Genetic factors influence resting-state connectivity more than common environmental factors.' In *The 22nd annual meeting of the organization for human brain mapping.* ohbm, Geneva, Switzerland.

Colclough, G.L., Smith, S.M., Harrison, S.J., Rojas López, P.A., Valdes-Sosa, P. & Woolrich, M.W. (2016*b*) 'A Bayesian approach for hierarchical modelling of sparse functional networks.' In *The 22nd annual meeting of the organization for human brain mapping.* ohbm, Geneva, Switzerland.

Colclough, G.L., Smith, S.M., Nichols, T.E., Tewarie, P.K., Brookes, M.J., Winkler, A.M., Glasser, M.F., Van Essen, D.C. & Woolrich, M.W. (2016*c*) 'The heritability of multi-modal functional connectivity.' In *The 29th International Conference on Biomagnetism.* biomag, Seoul, South Korea.

Liuzzi, L., Colclough, G.L., Quinn, A.J., Woolrich, M.W., Brookes, M.J. & Tewarie, P.K. (2016) 'Can we see dynamic connectivity in resting state meg?' In *The 29th International Conference on Biomagnetism.* biomag, Seoul, South Korea.

O'Neill, G.C., Tewarie, P.K., Colclough, G.L., Gascoyne, L.E., Hunt, B.A., Morris, P.G., Woolrich, M.W. & Brookes, M.J. (2016) 'Measurement of dynamic task related functional networks using meg during a cognitive task.' In *The 29th International Conference on Biomagnetism.* biomag, Seoul, South Korea.

Invited Talks

MEGUK conference, Birmingham, UK (January 2015). 'Performing whole-brain network analyses in MEG.'

Centro de Neurociencias de Cuba, Havana, Cuba (September 2015). 'Resting-state network modelling in fMRI and MEG: improving the characterisation of individual subjects.'

Department of Physics, University of Havana, Cuba (September 2015). 'Modelling brain connectivity at rest.'

Donders Discussions PhD conference, Nijmegen, The Netherlands (November 2015). 'The repeatability and reliability of resting-state connectivity metrics for MEG.'

Cognition and Brain Sciences Unit, Cambridge, UK (November 2015). 'Is our cortical connectivity determined by our genes or our development? Using resting-state MEG recordings to investigate heritability in the functional connectome.'

MEGUK conference, York, UK (March 2016). 'Born this way: the influence of genetics on the connectivity of human brain activity.'

Abstract

Magnetoencephalography (MEG) and functional Magnetic Resonance Imaging (fMRI) offer complementary insights into connected human brain function. Evidence from the use of both techniques in the study of networked activity indicates that functional connectivity reflects almost every measurable aspect of human reality, being indicative of ability and deteriorating with disease. Functional network analyses may offer improved prediction of dysfunction and characterisation of cognition.

Three factors holding back progress are the difficulty in synthesising information from multiple imaging modalities; a need for accurate modelling of connectivity in individual subjects, not just average effects; and a lack of scalable solutions to these problems that are applicable in a big-data setting. I propose two methodological advances that tackle these issues.

A confound to network analysis in MEG, the artificial correlations induced across the brain by the process of source reconstruction, prevents the transfer of connectivity models from fMRI to MEG. The first advance is a fast correction for this confound, allowing comparable analyses to be performed in both modalities. A comparative study demonstrates that this new approach for MEG shows better repeatability for connectivity estimation, both within and between subjects, than a wide range of alternative models in popular use. A case-study analysis uses both fMRI and MEG recordings from a large dataset to determine the genetic basis for functional connectivity in the human brain. Genes account for 20%–65% of the variation in connectivity, and outweigh the influence of the developmental environment.

The second advance is a Bayesian hierarchical model for sparse functional networks that is applicable to both modalities. By sharing information over a group of subjects, more accurate estimates can be constructed for individuals' connectivity patterns. The approach scales to large datasets, outperforms state-of-the-art methods, and can provide a 50% noise reduction in MEG resting-state networks.

1

Introduction

Over the last twenty years, neuroscience has been revising the way it studies and interprets function in the human brain. The brain was long considered as simply a reflexive and responsive organ, a complex computer whose outputs (movement, speech and behaviour) are a function of the sensory inputs it receives and its own internal states. The dominance of this view throughout the 20th Century led to the use of task-based paradigms in nearly all neuroimaging studies, in which the presentation of repetitive stimuli elicits repeatable responses in the cortex, and researchers identify brain regions that 'light up' in response to their experiment as associated in some way with the neural correlates of their task.[1]

An emerging viewpoint considers the brain not as a reflexive machine, but as an active interpreter and predictor of its stimuli and surroundings. The mo-

[1] One of the most lyrical expositions of computational neuroscience and the successes of modelling the brain as a computer is a book called *Computing the Mind,* by Shimon Edelman.

tivation for this view comes from a deep exploration of intrinsic brain activity, in which the brain is unstimulated and functional activity self-generated. This new line of investigation has been growing ever since Shulman *et al.* (1997) demonstrated that a consistent set of cortical regions show elevated activity in the *un*stimulated, as opposed to task-active, parts of an experiment, and Biswal *et al.* (1995) identified correlated patterns of spontaneous activity in the resting phases of a finger-tapping activity that showed 'similar neural organisation [to] task-evoked patterns of activity.' A rich picture of the brain has emerged as an organised, self-stimulated agent that reflects on stored memories, models its environment and predicts future events. These two viewpoints and their change in status are nicely contrasted by Raichle (2009) and further expanded upon by Callard & Margulies (2011), and again by Raichle (2015*b*).

It should be of no surprise to anyone that the brain is active when otherwise unstimulated: we can all relate to experiences of mind wandering, internal monologue or undirected independent thought. What *is* surprising though, is that when many different people are scanned in a *resting state,* in which one person may be thinking about her boss and another about his lunch, a very consistent set of different brain regions show independent, synchronised co-activation (Beckmann *et al.*, 2005; Brookes *et al.*, 2011*b*; Damoiseaux *et al.*, 2006; Smith *et al.*, 2012). Furthermore, these Resting-State Networks (RSNs) are organised in a manner which mirrors the co-ordination of activity triggered by stimulation or cognitive tasks (Smith *et al.*, 2009). In this way, the multifaceted functional responses of the complex computer are all reflected in the intrinsic behaviour of the active brain.

The observation of networks of co-activating regions led naturally to the study of the strength of connections between regions (Craddock *et al.*, 2013;

Greicius *et al.*, 2003; Smith *et al.*, 2013). In contrast to structural connectivity, which is concerned with the number or density of axonal connections across cortex, and to the well-established study of connectivity between individual neuron cells and cell bundles, *functional connectivity* measures synchrony, co-activation, correlation or information transfer between the functional activities in different brain areas.[2] Perhaps the most surprising feature of resting-state activation is the broad range of factors with which functional connectivity is associated, encompassing behaviour, cognition, attention, ageing, development, disease, wealth, health and happiness (Barnes *et al.*, 2016*b*; Craddock *et al.*, 2009; Dennis & Thompson, 2014; Finn *et al.*, 2015; Greicius, 2008; Rosenberg *et al.*, 2016; Smith *et al.*, 2015, 2013; Stam, 2014). Functional connectivity, even in the unstimulated resting state, is clearly integral to the whole gamut of human brain function.

Underpinning the rise of analyses of the brain's resting state, therefore, is a conception of human brain activity as being generated by a macroscopic network of many functional units, a *connectome.* This thesis contributes to current efforts to model the human functional connectome by developing improved network analyses for use in Magnetoencephalography (MEG), Electroencephalography (EEG) and functional Magnetic Resonance Imaging (fMRI), and by using these methods to determine the genetic basis for human functional connectivity.

My work can be equally applied to the analysis of images collected within task paradigms as to resting-state data, but I will concentrate solely on the latter. What, then, are the established features of resting-state brain activity? How do researchers model resting-state connectivity? What are the key directions in

[2] By *brain area* I mean simply a region on the scale of the resolution of whole-brain functional imaging apparatus (MEG, EEG and fMRI), about 0.5–5 cm across.

the field and the barriers to further progress? It is to these questions that I now turn, to paint the landscape within which this thesis should be viewed.

1.1 Functional connectivity and the resting state

CHARACTERISING THE RESTING BRAIN

Shulman *et al.* (1997) used a meta-analysis of PET images to describe a collection of brain regions that are more active in unstimulated or resting states than when the brain is involved in sensory processing or cognitive tasks (shown in figure 1.1 on the following page[3]). These areas, including posterior cingulate cortex, medial prefrontal cortex, bilateral inferior parietal lobule, bilateral temporal cortex and the hippocampal formation (Buckner *et al.*, 2008), are variously associated with reflection, episodic recall, autobiographical memory, imagination and the projection of the self (Gusnard *et al.*, 2001; Gusnard & Raichle, 2001), and were collectively thought to embody a baseline state or default mode of brain function (Raichle *et al.*, 2001). But these regions are not independently active. They exhibit synchronised, correlated activity in fMRI recordings, leading to their collective appellation as the *default mode network* (Buckner *et al.*, 2008; Greicius *et al.*, 2003; Raichle, 2015*a*). (As an example, whole-brain correlations to a seed voxel in the posterior cingulate will consistently highlight co-activation over the whole network.)

And it is not just one network that describes resting-state activity. Independent Component Analysis (ICA; Beckmann & Smith, 2004; Bell & Sejnowski, 1995; Hyvärinen, 1999) is a blind data decomposition technique that uncovers mutually correlated behaviour within certain groupings of regions, but independent

[3] Figure 1.1 is reproduced from Raichle *et al.* (2001); © (2001) National Academy of Sciences.

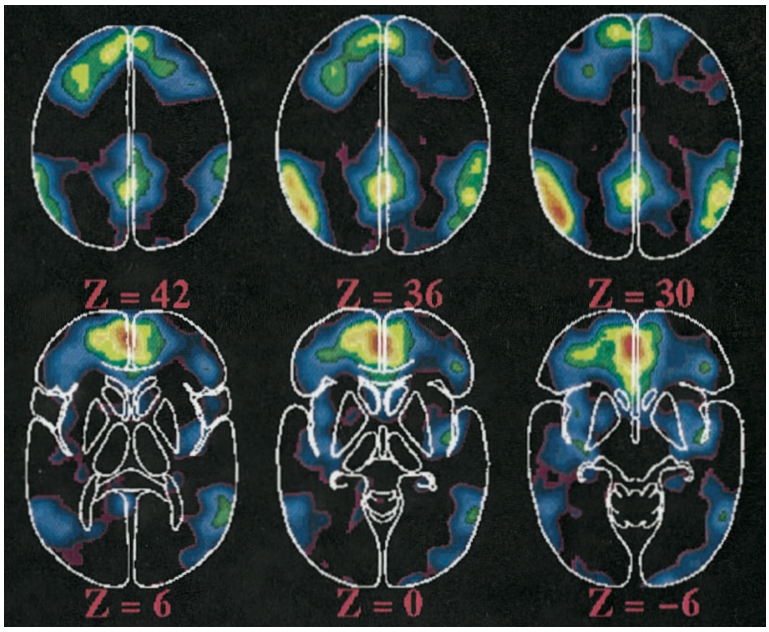


Figure 1.1: THE DEFAULT MODE NETWORK as characterised by Shulman *et al.* (1997). This figure, reproduced from Raichle *et al.* (2001), shows the results of a meta-analysis of PET studies, highlighting regions that are more active during periods of rest than stimulation or cognitive processing, and which form a *baseline state* (Raichle *et al.*, 2001) of human brain activity. The numbers in the image label slices of the z-plane in standard MNI space.

profiles of activity (either temporally or spatially) between the groups. When applied to resting-state fMRI data, a set of around 5–15 resting-state networks are revealed, that explain a large proportion of the variance of the data and are consistent across multiple imaging sessions and multiple subjects (Beckmann *et al.*, 2005; Damoiseaux *et al.*, 2006; Smith *et al.*, 2012). The organisation of these networks reflects the various patterns of co-ordinated function in the brain that are elicited by almost any cognitive task (Smith *et al.*, 2009). Thus, the resting brain segregates into networks of functional activity, which exhibit both local synchrony in cortical areas responsible for sensory processing and motor function, and longer-range communication within higher-order systems that are associated variously with attention, reflection, language, memory and executive control (Buckner *et al.*, 2008; De Luca *et al.*, 2005; Fox *et al.*, 2006;

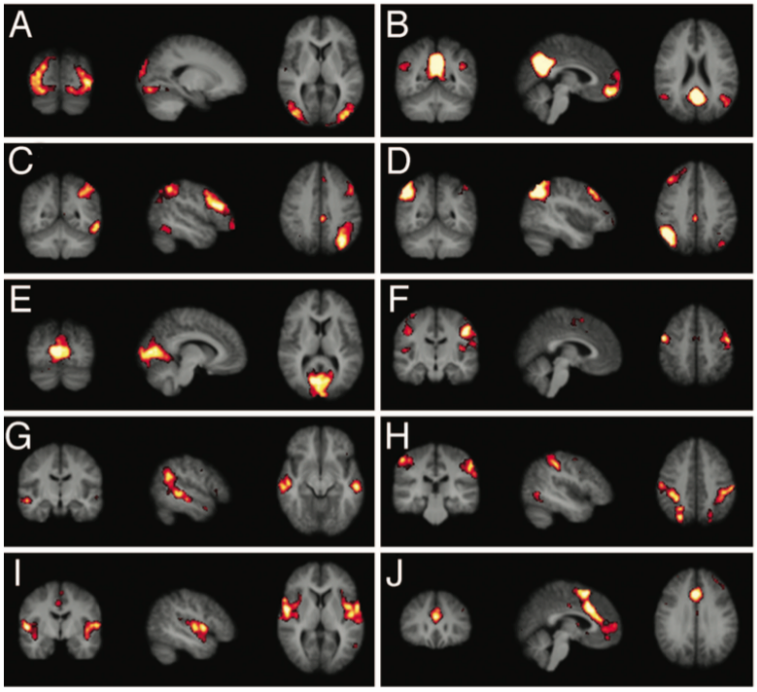


Figure 1.2: RESTING-STATE NETWORKS found by Damoiseaux *et al.* (2006). This figure, reproduced from Damoiseaux *et al.*'s paper in the Proceedings of the National Academy of Sciences in the US, identifies ten consistent resting-state networks in fMRI data using ICA. These can be categorised as visual networks (A and E), the sensorimotor network (F), the ventral stream (G), auditory network (I), the default mode network (B), lateralised fronto-parietal or dorsal attention networks (C and D), and areas of prefrontal cortex and dorsal anterior cingulate (J) that are involved in executive control and working memory. Images are heatmaps of z-statistics overlaid onto the standard MNI152 average brain, thresholded at 2 (cool black) and 5 (hot white).

Raichle, 2015*a*; Seeley *et al.*, 2007; Smith *et al.*, 2012, 2013). A selection of these RSNs is illustrated in figure 1.2.[4]

THE CONTRIBUTION FROM ELECTROPHYSIOLOGY

Independent verification of the existence of resting-state networks of neural activity has come from EEG and MEG. These techniques, in contrast to fMRI that uses the Blood Oxygenation Level Dependent (BOLD) response as a proxy for functional activation, measure the electrical and magnetic fields reaching the scalp from large-scale synchronous firing of bundles of aligned neurons

[4] Figure 1.2 is reproduced from Damoiseaux *et al.* (2006); © (2006) National Academy of Sciences.

(Baillet *et al.*, 2001), forming a direct measure of neural response. Resting-state electrophysiological recordings are dominated by oscillatory signal in the theta (4–8 Hz) band, in the alpha (8–13 Hz) band, particularly in posterior regions, and in the beta (13–30 Hz) band, particularly from dorsal and motor regions. Fluctuations in the power of these oscillations, when decomposed with ICA, show a segregation into resting-state networks that is very similar to that identified in fMRI (Brookes *et al.*, 2011*b*).

Each network can be identified with a particular signature of oscillatory frequencies, so it is clear that many different neural processes are captured within these global synchronous fluctuations in the resting state (Brookes *et al.*, 2011*b*; Hipp *et al.*, 2012; Mantini *et al.*, 2007; Marzetti *et al.*, 2013; de Pasquale *et al.*, 2010). Beta synchronisation and desynchronisation have long been associated with the initiation of movement, and alpha rhythms with visual processing, so it is unsurprising that the motor and visual networks present in these respective frequency ranges. The specific roles and mechanisms by which the dorsal attention network and default mode network operate are still being uncovered. Some authors (Marzetti *et al.*, 2013; de Pasquale *et al.*, 2012) have related their prevalence in the alpha and beta ranges to the growing understanding of alpha power synchrony as a medium for attention to control or suppress distracting activity (Lakatos *et al.*, 2016; Womelsdorf & Fries, 2007), and of beta power as a mechanism for maintaining the *status quo* of internal or external states (Engel & Fries, 2010). While much is still murky, MEG does provide an opportunity to relate specific oscillatory patterns (for which there is an extensive literature describing functional relevance, cognitive mechanisms and distinct neurobiological origins[5]) to the dominant patterns of resting-state activity, clarifying the

[5] See Wang (2010) for a comprehensive review of the relationships between oscillatory frequency,

mechanistic bases and cognitive responsibilities of resting-state networks.

The high temporal resolution of MEG data has also enabled characterisation of the dynamic occupancy of these networks by the resting brain. We now understand that the brain flits between states on a timescale that is much faster than the temporal resolution of fMRI, residing in each for about 200 ms, although with some longer-lived periods (see figure 1.3[6] on the next page, also Baker *et al.*, 2014). Additionally, there are data to suggest that particular areas, such as the posterior cingulate or even the entire default mode network, act as mediators for the integration of information between networks: switching between particular subsets of resting-state networks frequently occurs by passing first through the default mode (de Pasquale *et al.*, 2012, 2015).

A NETWORKED VIEW OF BRAIN FUNCTION

After whole-brain networks of functional activity were identified in the resting state, the move to analyses of the connectivity between regions was natural. A network analysis of the functional connectome can take many forms (see section 1.2 below), but in general provides a description of the strength of connectivity, correlation or synchrony between activation in different regions (Craddock *et al.*, 2013; Friston, 2011; Smith, 2012; Smith *et al.*, 2013). It is now clear that, however it is measured, functional connectivity encodes important information about nearly every topic of interest to human neuroscience.

originating cell population and functional purpose. For more recent efforts to relate connectivity in different parts of the frequency spectrum to directions of information flow and the updating of internally-held beliefs within the cognitive processing systems, as part of the predictive coding framework, see Friston & Kiebel (2009) and Bastos *et al.* (2015).

[6] Figure 1.3 is adapted from Baker *et al.* (2014) under the terms of the Creative Commons attribution license (CC BY 3.0).

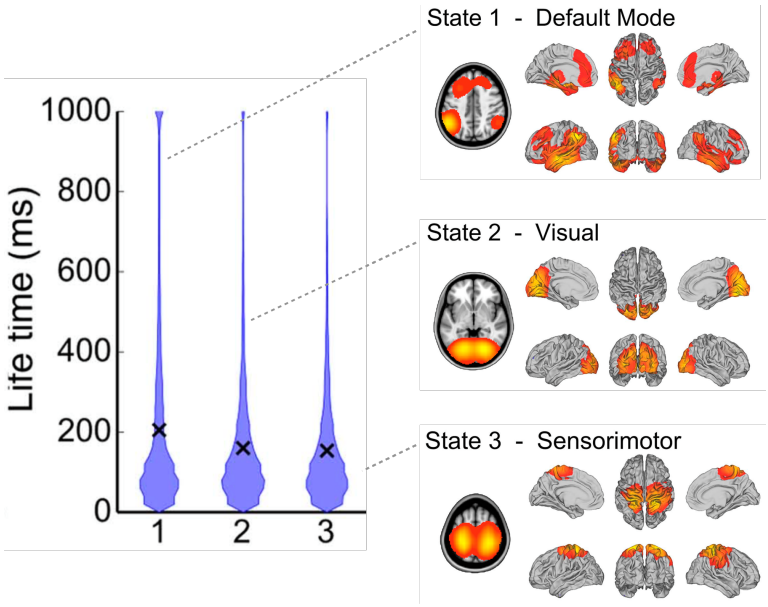


Figure 1.3: SWITCHING BETWEEN NETWORKS IN THE RESTING STATE. Resting-state activity is characterised by fast-timescale switching between different resting-state networks, with only transient occupancy of each. This figure, adapted from Baker *et al.* (2014), shows three out of the eight states identified by a hidden Markov model fitted to resting-state MEG recordings, (A) default mode network, (B) visual network, and (C) sensorimotor network. The distribution of state lifetimes is given on the left (the cross marks the mean): the brain rests in each state for between 50–1000 ms, concentrated below 200 ms with a long tail.

Ageing and development

Connectivity in the brain changes throughout the life cycle. While most resting-state networks can be identified in children, networks associated with attention and control continue to develop in late adolescence, which some authors have related to the improvement in decision-making ability throughout the teenage years (Barnes *et al.*, 2016*b*; de Bie *et al.*, 2012; Fair *et al.*, 2007). These same networks can be enhanced in children and teenagers by the use of cognitive training exercises targeted at working memory ability (Astle *et al.*, 2015; Barnes *et al.*, 2016*a*).

On the other hand, ageing is thought to degrade functional network integrity (Dennis & Thompson, 2014; Dosenbach *et al.*, 2010; Schlee *et al.*, 2012),

particularly affecting longer-range connections in the default mode network and dorsal attention network (Andrews-Hanna *et al.*, 2007; Ferreira & Busatto, 2013; Ferreira *et al.*, 2016; Tomasi & Volkow, 2012), and relates to age-associated cognitive decline (Onoda *et al.*, 2012). Ageing can involve some increases in functional network connectivity, particularly across the hemispheres, thought perhaps to be a method for functional re-organisation as the brain attempts to compensate for the loss of grey matter (Gutchess, 2014; Sala-Llonch *et al.*, 2015)—although this re-organisation cannot always be said to underpin the preservation of language ability or cognitive performance (Shafto & Tyler, 2014).

Findings of hypoconnectivity in fMRI among healthy ageing populations must be interpreted with some caution, as there is growing evidence that these results may be particularly susceptible to uncontrolled movement artefacts (Power *et al.*, 2012, 2015; Siegel *et al.*, 2016).

Disease

A broad range of neurological disorders and diseases have been associated with changes (both positive and negative) in the strength of functional connectivity. In general, while there is some sensitivity to diseases including schizophrenia, Alzheimer's disease and its precursory mild cognitive impairment, Parkinson's disease, motor neurone disease, Huntingdon's disease, autism spectrum disorder and attention deficit disorder, it is fair to say that there is a lack of specificity in prediction: many of these diseases and disorders present partially coincident changes to connectivity[7] (Barkhof *et al.*, 2014; Bullmore, 2012; Craddock *et al.*, 2009; Fornito & Bullmore, 2015*a*,*b*; Greicius, 2008; Harrington *et al.*, 2015; Monk

[7] This is in contrast to disease-induced atrophy of brain tissue. Many different neurodegenerative diseases cause specific atrophy within regions associated with distinct resting-state networks (Seeley *et al.*, 2009).

et al., 2009; Müller *et al.*, 2011; Stam, 2014; Stam & van Straaten, 2012; van Straaten & Stam, 2013). Of this list, the connection between Alzheimer's disease and functional connectivity is perhaps the best established, with consistent findings of hypoconnectivity between the nodes of the default mode network (Hafkemeijer *et al.*, 2012; Liu *et al.*, 2014; Sheline & Raichle, 2013), enabling sensitive classification of patients from healthy controls (Dhillon *et al.*, 2014). This is sensible, given that the regions involved with the network are known to be vulnerable to Alzheimer's, in terms of reduced function, grey matter loss, and the build-up of beta-amyloid plaques (Buckner *et al.*, 2005)—but the disruption to functional connectivity may simply be a direct consequence of these effects, and not an independent contributor to the memory dysfunction and cognitive under-performance associated with the disease (Buckner *et al.*, 2008). Given the low specificity of functional connectivity as a sole predictor, efforts to find biomarkers for disorder and disease will increasingly use connectivity as part of multi-modal feature sets for the prediction of risk and prognosis.

Function, cognition and behaviour

The idiosyncrasy of subjects' functional activations in response to different tasks is encoded in their functional connectivity at rest. Classifiers trained on resting-state connectivity are able to accurately predict the activation map of task response in new subjects, even if the location and topology of the activation pattern is atypical of the population (Tavor *et al.*, 2016).[8] As a natural extension of this property, resting-state functional connectivity can predict aspects of be-

[8] Being able to predict atypical topologies of response is important. A different topological configuration in the task-related activation within a particular subject indicates a different organisation of function in the cortex of that individual; no (sensible) image registration method would be able to align that response to the population average.

havioural response to cognitive tasks, such as reaction time or working memory performance (Jamashita *et al.*, 2015; Sala-Llonch *et al.*, 2012; Smith *et al.*, 2013; Zou *et al.*, 2013). Broader properties of cognition are encoded in the strength of connectivity, such as fluid intelligence (Finn *et al.*, 2015) and the ability to sustain attention (Rosenberg *et al.*, 2016). More broadly still, patterns of functional connectivity are associated with a very wide range of lifestyle factors, such that suppression of a particular configuration of connections correlates (independently) with alcohol intake, smoking, blood levels of tetrahydrocannabinol (the psychoactive component of cannabis), parental drug abuse, stress and anger, whereas expression of the same system is indicative of greater levels of education, improved working memory, fluid intelligence, attention, health, income and life satisfaction (Smith *et al.*, 2015).

Genetic control

It is unsurprising, given this range of associations, that functional connectivity is to some extent genetically determined. Glahn *et al.* (2010) used ICA on fMRI data to show that about 40% of the variation in connectivity to the default mode network is controlled by genes. However, a study of the heritability of functional connection strength, using whole-brain network analyses on either fMRI or MEG data is lacking—see chapter 4.

Improvement of image alignment and region categorisation

Finally, the use of information from the functional connectome can create dramatic improvements to other analyses. The current state of the art in the registration of MR images, Emma Robinson's multimodal surface matching algorithm `MSMall`, performs topology-preserving transformations on data defined on the

cortical surface (Robinson *et al.*, 2014), and employs functional connectivity to gain precise alignment of functional and structural images between different subjects.

Additionally, functional connectivity information has been crucial in the generation of highly repeatable categorisations of the human cortex into discrete functional units that relate to previously understood anatomical boundaries, but which broaden our understanding of functional specification (Baldassano *et al.*, 2015; Craddock *et al.*, 2012; Dhillon *et al.*, 2014; Glasser *et al.*, 2016; Gordon *et al.*, 2016; Harrison *et al.*, 2015; Yeo *et al.*, 2011).

1.2 A crash course in network modelling

A functional network analysis of preprocessed data can be broken down into two steps: parcellation and network modelling. In an important sense, the parcellation defines the functional network that is to be subsequently estimated. Each parcel[9] constitutes a vertex of the network, and is described by its own time course. The process of network modelling uses these time courses to characterise the connectivity strengths between the vertices.

PREPROCESSING

As with all analyses of functional activation, some minimal preprocessing is required. The objectives are to constrain the data to frequency bands of interest, to remove physiological and structured environmental noise, and to bring functional recordings across subjects into a common space, within which comparisons can be made, by a process of image registration. Most of these procedures

[9] Parcel, mode, vertex, node, Region of Interest (ROI)—I shall use these terms interchangeably.

are well-established, and only a couple of points need to be highlighted.

Structured noise is commonly removed both in MEG and fMRI with a combination of ICA and regression (Mantini *et al.*, 2011; Salimi-Khorshidi *et al.*, 2014). The data are decomposed into a set of independent components, and components showing spatial profiles of well-known movement or eye-blink artefacts, spectral profiles characteristic of mains noise or heartbeat, or time courses that correlate with physiological measures,[10] are regressed out of the original data. Of particular concern is movement artefact, which can create structured correlations in functional data that mimic the arrangement of several resting-state networks (Power *et al.*, 2012, 2015). Very recent results suggest that even cleaning data with this combination of ICA and regression, which is the current best practice in fMRI, is insufficient to remove all effects of movement (Siegel *et al.*, 2016). The extent of the problem in MEG is less well-known, although most systems are able to continuously track head position and register the data into a common stationary frame.

Recordings from MEG sensors must additionally be projected onto the cortical surface or points within the brain volume. This process of source reconstruction is important for connectivity analyses, providing noise rejection, interpretability, comparability across datasets and to fMRI results, and some reduction in the magnetic field spread problems that afflict connectivity analyses between sensors (Schoffelen & Gross, 2009). A wide range of tools are available to reconstruct the cortical sources that generate the magnetic signals measured at the scalp. I use beamformers (van Drongelen *et al.*, 1996; Hillebrand *et al.*, 2005; Van Veen *et al.*, 1997) exclusively in this thesis, partly because they can be

[10] These might include the electrocardiogram, electro-occulogram, electromotorgram, estimates of movement in the MRI scanner, or the respiration signal from a chest band.

employed when no good estimate of the strength and covariance of the noise sources are available, partly because they have been shown to have good resolution for resting-state data (Schoffelen & Gross, 2009), and partly because they are the standard tool used within the in-house software and by researchers at OHBA; the methods that I develop should be relatively agnostic to this choice.

PARCELLATION

Performing a connectivity analysis between every brain voxel with functional data is not always sensible. These so-called *dense connectomes* are noisy and computationally-intensive to estimate, difficult to visualise and hard to interpret.[11] In particular, MEG rarely contains information on more than about 100 independent cortical sources, making a dense connectome analysis nonsensical in this modality.

A parcellation is therefore a method for informed dimensionality reduction that reflects the functional activation and connectivity patterns within the data. The goal is to identify distinct functional units that can be well-described by a single time course of functional activity, and which can be associated with a set of functional purposes based on existing knowledge.

Anatomical parcellations define boundaries between regions using the results from histological and cytoarchitectural studies, together with supervised segmentations of structural MR scans. While the use of well-established categorisations of cortex, such as the Brodmann areas (Brodmann, 2006), makes it easy to discuss the function of the regions they describe, and to relate findings to previous work, they produce less repeatable and predictive connectivity

[11] Although dense connectomes have been useful for *creating* parcellations (Baldassano *et al.*, 2015).

matrices than data-driven *functional* parcellations that use the properties of interest within the data themselves to inform the decomposition (see Dhillon *et al.* (2014) for just one example of this difference).

Functional parcellations can be as simple as ICA decompositions of resting-state data. In this approach, the spatial maps (whether of small, contiguous regions or of whole networks such as the default mode) are used as a weighting over voxels, from which the mean time course or principal component is taken. Alternatively, a functionally-informed binary parcellation of cortex can be produced, creating discrete regions of interest, either by using a clustering algorithm on large functional connectivity matrices (Baldassano *et al.*, 2015; Craddock *et al.*, 2012; Yeo *et al.*, 2011), or differentiating areas by identifying changes in functional connectivity profiles at the boundaries between regions (Gordon *et al.*, 2016). A particularly ambitious (and very recent) parcellation scheme using this latter idea on both functional and structural data shows promise both in improving the sensitivity of functional experiments that it informs, and in relating regions to previously-established categorisations (Glasser *et al.*, 2016).

As a final point, there is some evidence that finding subject-specific versions of the global parcellation can improve connectivity estimation (Beckmann *et al.*, 2009; Calhoun *et al.*, 2001; Dhillon *et al.*, 2014; Harrison *et al.*, 2015; Luckhoo, 2013).

NETWORK MODELLING

Having settled upon a set of ROIs to function as nodes in the functional network, the next step is to infer the strength and, potentially, the direction of the influence between each pair of ROIs—to infer the edges of the network. Methods

for performing this inference split broadly into three categories: biophysical network models, and measures of either phase or amplitude synchrony between ROIs. Biophysical network models assume some generative process for the data in each region and model the connectivities between these underlying systems. By contrast, the latter two approaches model what Friston in 2011 called 'effective connectivity,' focussing on the statistical relationships between the ROI time courses.

Amplitude coupling

In electrophysiology, synchronised fluctuations in the power of oscillations in different regions are characterised by correlations in the envelopes of the oscillatory signals (Brookes *et al.*, 2012*b*; Hillebrand *et al.*, 2012; Hipp *et al.*, 2012). The resulting patterns of amplitude coupling relate closely to the slow scale resting-state co-ordinations identified by fMRI (Brookes *et al.*, 2011*a*,*b*; Engel *et al.*, 2013). Within the fMRI literature itself, the most common statistical measure of connectivity is correlation. While linear correlations do not discriminate between direct and indirect connections (when two nodes are connected by one or more intermediaries), and cannot determine the direction of network edges, correlation measures performed well in a large comparison of network analysis methods for use in fMRI (Smith *et al.*, 2011).

When an understanding of solely the direct connectivity of amplitude fluctuations is required, one of two approaches is taken. The first possibility is to posit an undirected Gaussian Graphical Model (GGM), in which the activations of the ROIs are described by a joint multivariate normal distribution, and the connections between nodes estimated with the partial correlation matrix (Fisher, 1924). The partial correlation matrix is simply calculated from the precision matrix of

the data (inverse of the covariance matrix), and zeros in this matrix indicate the absence of a direct edge between two variables (Kunihiro *et al.*, 2004). Good estimates of partial correlation matrices are quick to calculate (Friedman *et al.*, 2008), and perform very well in fMRI data (Duff *et al.*, 2013; Marrelec *et al.*, 2006; Smith *et al.*, 2011). The second approach hypothesises a directed, acyclic graphical model. (These models consider a network with directed influences between nodes, but forbid the formation of loops or cycles in the graph.) Models using Gaussian likelihoods for the data in each node tend to be inferred by ranking the possible network structures using a scoring system, often related to the Bayesian information criterion,[12] and an efficient searching scheme (see, for example, Meek, 1995, and Ramsey *et al.*, 2010). Models using non-Gaussian likelihoods, such as LINGAM (Shimizu, 2014; Shimizu *et al.*, 2006; Tashiro *et al.*, 2014) can exploit moments in the data above second-order statistics (variances) to infer directionality in the communication between nodes, and use fast ICA algorithms for inference. Most of these directed acyclic models seem to perform about as well as partial correlation in fMRI data (Ramsey *et al.*, 2014; Smith *et al.*, 2011), but have not been used in electrophysiology.

Finally, if an undirected connectivity matrix has been inferred between ROIs, whether with correlation, partial correlation, mutual information or some other measure of synchronous amplitude fluctuations, a directionality can be assigned to each network edge using properties of the high-order moments in each node (Bielczyk *et al.*, 2016; Hyvärinen & Smith, 2013; Patel *et al.*, 2006; Ramsey *et al.*, 2011). This *post-hoc* approach seems more successful than LINGAM when tested on simulated data (Bielczyk *et al.*, 2016; Ramsey *et al.*, 2014).

[12] The description of these models as 'Bayes net' methods is therefore something of a misnomer.

Phase relationships

Phase coupling between oscillations is often thought of as a distinct process to amplitude coupling, with the two modes of communication having different functional significance (Engel *et al.*, 2013). Phase coupling might be employed for more transient or meta-stable states of synchrony than amplitude coupling (Palva & Palva, 2011), although the demonstration that modes of amplitude coupling also exhibit short-lived transience has cast shade over this suggestion (Baker *et al.*, 2014).

Whether or not it has a particular and individual functionality, the phase synchrony between signals from different ROIs can be assessed using spectral measures of coherence or narrow-band estimates of phase differences. A wide range of estimators are available, but all perform the same purpose, and will be discussed more extensively in chapter 3. It is only worth highlighting a set of connectivity measures based on multivariate autoregressive models, such as Partial Directed Coherence, that attempt to estimate directed information flow from lagged phase relationships.

All of the spectral and lag-based methods can also be applied to fMRI time series, but they tend to perform poorly (Smith *et al.*, 2011). Attempts to assign directionality, in particular, will suffer as different hæmodynamic delays are present across the brain, presenting a strong confound to any lag-based analyses of BOLD signals.

The spectral richness of MEG has also revealed phase-amplitude coupling between regions as an additional method of information transfer and control during cognitive processing (Canolty *et al.*, 2006) and motor function (de Hemptinne *et al.*, 2015). Phase-amplitude coupling manifests as nested oscillations, in

which the amplitude of one fast oscillatory signal is modulated by the phase of another, slower wave. Methods for elucidating the occurrence of this coupling use a combination of the techniques employed separately to find phase and amplitude coupling, for example by computing the correlation or phase locking between the modulat*ing* signal and the envelope of the modulat*ed* signal (Penny *et al.*, 2008; Tort *et al.*, 2010). Not many of these measures are very robust, and all seem particularly susceptible to false inference when transient bursts of non-sinusoidal oscillations occur (Cole *et al.*, 2016; Jones, 2016; Lozano-Soldevilla *et al.*, 2016). However, this is an area under active development, and emerging methodology for the analysis of phase-amplitude relationships in non-stationary signals may offer significant improvements (Huang *et al.*, 2016). At the present time, there is no framework to build whole-brain connectivity matrices based on phase-amplitude coupling.

Biophysical models

Biophysical models vary widely in scope. They range from simple parametrisations of the relationships between functional volumes, constrained by the structural connections estimated from diffusion MRI or known neuroanatomy (such as the structural equation models of McIntosh & Gonzales-Lima, 1994), to models of abstract network dynamics that are overlaid with neural responses and hæmodynamics (for example, the dynamic causal models for MRI; Friston *et al.*, 2003, 2014, 2011). In their most detailed extent, these models can be highly sophisticated characterisations of the collective phenomena of aggregated neurons, such as the dynamic causal model for MEG, that fits well to recorded data (David *et al.*, 2006; Kiebel *et al.*, 2008).

Parameters in these models are mostly estimated using Bayesian inference

methods,[13] but all are sufficiently complicated that strong prior beliefs about the relevant ROIs and network connections must be used to constrain the parameter space, and undirected explorations of the whole connectome are not yet possible. These models do, however, allow for the interpretation of connectivity in terms of neurobiologically meaningful parameters, such as the concentrations of metabolites and neurotransmitters, or the communication between particular cell populations (Woolrich & Stephan, 2013).

CONTINUED ANALYSES

What happens once connectivity matrices have been estimated? This obviously depends on the goals of the study, but three lines of enquiry are common.

Comparative network analysis Functional networks inferred from separate groups of subjects (patients and controls, for example, or participants taking different dosages of a drug) can be compared using the standard tools for statistical comparison, fitting linear models to look for differences in connection strength on each network edge. The number of edges in a network grows quadratically with the number of nodes, rapidly presenting a large multiple comparisons problem for statistical inference. A number of methods for controlling the false positive rate are available (Benjamini & Hochberg, 1995; Winkler *et al.*, 2016*a*, 2014), but unlike functional activation data, for which the spatial smoothness of the results reduces the amount of independent information in the data, there is no clear relationship between the statistical tests on different network edges

[13] There is an additional literature on computational models of connected neural populations, that aim to understand by simulation the most important influences on connected brain function (Deco *et al.*, 2008). At the moment, however, these neural mass models and neural field models are not invertible, and therefore not applicable for the network analysis of recorded data, except by way of comparison (for example in Cabral *et al.*, 2014, and Deco & Kringelbach, 2016).

that can be exploited to model their interdependencies, and so no direct parallel of cluster-based inference (Nichols & Hayasaka, 2003) is available. Zalesky *et al.* (2012, 2010) have attempted to produce tests based on structural properties of the inferred networks (see also the discussion below), but the null hypotheses they construct are complex, and difficult to relate to the real questions of interest.

Structural analysis of the network The complexity embodied by large functional networks can be described and simplified. Graph theory is the relevant branch of mathematics that seeks to characterise the properties of binary and weighted networks. It provides tools to assess the potential efficiency of information transfer that is possible between the nodes of a graph, and to identify nodes that are highly involved in this process. Many of these graph theoretic measures focus on the structural properties of the functional network, viewing it as a binary graph. As functional connectivity is (nearly always) expressed as a set of connection strengths, these properties become conditional on a thresholding process. Because of this limitation, and because many properties can be so abstract as to make comprehension in terms of the original connectivity unintuitive, only a small proportion of researchers within the field have explored the value of this analysis approach. Among those who have, good summaries of the measures in use, their meaning, their predictive ability for disease, and attempts to understand the organisation of human brain function using these network descriptors, can be found in Bullmore (2012); Bullmore & Sporns (2009, 2012); Fornito & Bullmore (2015*a*); de Pasquale *et al.* (2015); and Rubinov & Sporns (2010).

Multivariate analyses Rather than analysing changes in connectivity on each network edge, the whole functional connectome can be treated in a multivariate fashion. Using the strengths of functional connections as predictors, and each subject as a data point, a model can be built to classify subjects by group or to predict disease severity, behavioural response or other experimental variables of interest (Smith, 2012; Smith *et al.*, 2013). Typical models include regularised regressions (for example, logistic regression with the elastic net), non-linear kernel regressions (support vector machines, relevance vector machines), Gaussian process regressions, or neural network models. Significant classification or regression performance indicates that the whole connectome encodes relevant information about the predicted variable, but extracting which connections were most important in generating this performance is less trivial.[14]

As an alternative, canonical correlation analyses, which identify the principal modes of variation over the whole set of network connections, have proved successful in both MEG and fMRI (O'Neill *et al.*, 2015; Smith *et al.*, 2015).

1.3 The principal challenges to progress

There are three key directions that the field of functional connectivity analysis is exploring, as far as I can see. These are a desire to integrate information from many functional imaging sources, the use of very large datasets, and an increasing focus on the connectivity of individuals, rather than that of the population average. In each direction, there are a number of major barriers in

[14] For good resources covering multivariate prediction that are specific to brain imaging, functional connectivity, and the pitfalls around the interpretation of results from these data, see Dadi *et al.* (2016); Pereira *et al.* (2009); and Varoquaux *et al.* (2016). For tutorials on fitting multivariate models more generally, see Hastie *et al.* (2009); Murphy (2012); and Rasmussen & Williams (2006).

the path of progress that must be overcome.

Multi-modal analyses

The analysis of functional connectivity, particularly in the resting state, was born in fMRI. Section 1.1 discussed the significant impact that MEG has had on our understanding of resting-state connectivity, and the opportunities afforded by its rich temporal and spectral resolution. There is therefore a desire for integrated analyses that can draw on information from multiple spatial and temporal scales simultaneously. The growth of simultaneous EEG and fMRI, and the acquisition of publicly-available datasets such as the Human Connectome Project (HCP; Van Essen *et al.*, 2013) that provide MEG and fMRI recordings from the same subjects, are initial steps in this direction.

On this track, there are numerous hurdles to be jumped. Functional connectivity in fMRI has an established methodology for network analysis, and that is covariance modelling. Correlation and partial correlation models (or GGMs) are cheap to compute, are among the most effective tools for fMRI (Smith *et al.*, 2011) and were used in nearly all the influential studies discussed in section 1.1. This is not the case at all for MEG analysis, where the community has not converged on a source reconstruction method, let alone a set of optimal network modelling tools. More guidance is needed on the choices made as part of a connectivity analysis. Furthermore, the successful methods from fMRI cannot be simply applied to MEG, because a modality-specific confound makes the results uninterpretable. Straightforward comparison of connectivities inferred from the two modalities therefore becomes more difficult, as do any attempts to construct a unified multi-modal inference framework.

Big data

Much can be gained by moving away from small-group comparative studies. Very large public studies offer an increase in statistical power, an opportunity to acquire highly representative samples, the resources to measure a very wide range of attributes and cognitive measures from participants, (frequently) genetic information on subjects, flexibility to investigate many research questions from the same data, and a shared resource that enables the reproduction of results. The downside is, of course, that a few particular samples become very heavily studied, and the results from these analyses over-generalised. Three very large datasets have recently become, or are becoming, available with fMRI data. The Alzheimer's Disease Neuroimaging Initiative holds data from about 1500 subjects, and is targeted at the search for clinical biomarkers for Alzheimer's disease (Mueller *et al.*, 2005). The Human Connectome Project (HCP) has collected high-quality structural, diffusion and functional data from nearly 1000 subjects in a nested family design (Van Essen *et al.*, 2013). Lastly, the UK Biobank is aiming to collect MRI scans from 100 000 people in an effort to provide population-level data for use in applications throughout the medical sciences (Miller *et al.*, 2016).

The obvious difficulty associated with large datasets is computational. Scalable inference methods need to be developed or improved for functional connectivity analyses, to allow the most advanced connectivity estimates to be used on these important resources. Large datasets also throw up the possibility of new modelling challenges, in which the characterisation of the variability over subjects is important, and distributional modelling over the whole population becomes achievable (Marquand *et al.*, 2016).

A focus on the individual

Accurate construction of connectivity models for individual subjects is required in order to answer nearly every important question resulting from functional connectivity analyses. Studies of genetic heritability, investigations into the associations between connectivity and cognitive and behavioural measures, and efforts to construct predictive biomarkers by training multivariate predictive tools on functional connectivity data: all of these analyses are based on connectivity estimates for the individual, not for the group. Network estimation, even using simple covariance models of the data, can be very noisy, particularly with short-acquisition datasets (Friedman *et al.*, 2008; Smith *et al.*, 2011).[15] There is therefore a need to improve the quality of connectivity estimation from single subjects. Upgraded methods intended for use on big datasets would need to be computationally scalable as the number of subjects increases, but could make use of shared information content within the population.

1.4 Summary of the argument

Functional network modelling in neuroimaging is the process of describing the relationships between activation in different distributed regions of cortex. Some of the simplest, but most successful, techniques for describing these relationships are models of the covariance structure of the functional activity. This thesis attacks some of the principal barriers to progress in the use of covariance models for connectivity estimation with fMRI and MEG. These include: the inability to perform comparable network analyses in both MEG and fMRI, the lack of

[15] While the HCP collected an hour of resting-state fMRI from its participants, it only collected six minutes of resting-state MEG; the UK Biobank is collecting six minutes of fMRI at rest.

understanding of the quality of single-subject network estimation in MEG and the absence of convergent good practice in network modelling methods, and the need for improved models of individual functional connectivity that are scalable to very large datasets.

Chapter 2 discusses covariance modelling in MEG. It describes a modality-specific barrier, the leakage of inferred sources between regions, to interpretable modelling methods. It then describes a fast and efficient method for diminishing this artefact, enabling a large range of multivariate analysis techniques to be applied to these data. This permits the use of the same techniques for functional connectivity estimation in MEG as in fMRI.

Chapter 3 compares the performance of covariance modelling to other models of functional connectivity used for electrophysiological data. It seeks to provide additional guidance to researchers making decisions about network modelling methods but warns that the quality of individual connectivity estimates is in general quite poor. It demonstrates that correlation analyses, as set out in chapter 2, provide the most reliable and repeatable inference for single-subject and population-level network estimation, outperforming a wide range of popular techniques.

Chapter 4 provides an example of the use of covariance modelling for an analysis of the heritability of functional connectivity, using a large multi-modal twin dataset from the Human Connectome Project. Because covariance modelling can now be used in both fMRI and MEG, a very similar analysis is able to be performed in both modalities, facilitating the synthesis of results. The data from both modalities support the conclusion that genes play an important role in de-

termining the strength of connectivity within human functional activation, and suggest that this genetic control is stronger than the influence of developmental environments.

Chapter 5 sets out a new method for multi-subject covariance modelling. A hierarchical Bayesian inference system for functional connectivity is described, which has a computational complexity that scales linearly with the number of subjects it models. It is applicable to both fMRI and MEG data, and is shown to improve network estimation in individual subjects, particularly for short-acquisition datasets.

In sum, covariance modelling techniques are brought to a new modality, MEG, and then scalable models are proposed to improve functional network estimation for single subjects within both MEG and fMRI. Both sets of methods are shown to outperform the existing range of techniques, and the results of a case-study analysis provide a new understanding of the genetic origins of human functional brain connectivity.

Chapter 6 looks ahead to the impact of these methods, and to new systems for building hierarchical models for similar applications.

2

A symmetric, multivariate leakage correction for MEG

Ambiguities in the source reconstruction of magnetoencephalographic (MEG) measurements can cause spurious correlations between estimated source time courses. This is a modality-specific barrier to the use of a wide range of multivariate functional network modelling techniques. This chapter sets out a symmetric orthogonalisation method to correct for artificial correlations between a set of multiple ROIs. This process enables the straightforward application of network modelling methods, including common methods from fMRI such as GGMs or partial correlation analyses, and multivariate autoregressive modelling, to infer connectomes, or functional networks, from the corrected ROIs. Here, we apply the correction to simulated MEG recordings of simple networks and to a resting-state dataset collected from eight subjects, before

computing the partial correlations between power envelopes of the corrected ROI time courses. We show accurate reconstruction of our simulated networks, and in the analysis of real MEG resting-state connectivity, we find dense bilateral connections within the motor and visual networks, together with longer-range direct fronto-parietal connections.

2.1 Introduction

MEG has provided independent corroboration of the existence of functional networks in the human brain (Brookes *et al.*, 2012*a*, 2011*b*; Hipp *et al.*, 2012; Luckhoo *et al.*, 2012; Mantini *et al.*, 2011; Marzetti *et al.*, 2013; de Pasquale *et al.*, 2010; Stam & van Straaten, 2012; Wens *et al.*, 2014), but network connectivity analysis in this modality is less mature than in fMRI. Immediate application of the tools used in fMRI to MEG is hampered by the presence of artefactual correlations between the inferred cortical sources. These correlations are a result of the ill-posed inverse problem: the few hundred magnetic field sensors cannot provide sufficient discriminatory information to independently estimate the source activity in thousands of brain voxels. Reconstructions of activity within the brain are coupled over space, with neighbouring regions exhibiting spatially-correlated behaviour. This effect is frequently termed *source leakage,* as the reconstructions of true point dipole sources from the measured signals will be spread over several voxels.

Source leakage therefore creates problems for the interpretation of network analyses. Relationships identified between network nodes could either reveal true functional interactions, or instead reflect the imperfections of electrophysiological source reconstruction. It is clearly important to be able to

perform interpretable whole-brain network analysis in MEG: the use of similar analyses in both MEG and fMRI enables the integration of evidence from multimodal sources; MEG has the potential to offer spectrally resolved information on connectivity, linking different aspects of the connectome to specific neuronal oscillatory frequencies (Brookes *et al.*, 2011*b*; Marzetti *et al.*, 2013); and MEG is a viable tool for exploring short-timescale fluctuations in resting-state connectivity (Baker *et al.*, 2014; O'Neill *et al.*, 2016; de Pasquale *et al.*, 2010, 2015; Vidaurre *et al.*, 2016). Source leakage is therefore a major barrier to mature functional connectivity investigations.

The amount and spatial profile of source leakage is dependent on the strength of the signal, the location in cortex and the assumptions of the inverse solution. Beamformers (Van Veen *et al.*, 1997) seem to reduce field spread more effectively (Schoffelen & Gross, 2009) than the distributed solutions of minimum norm estimates (Sarvas, 1987), even in the presence of correlated sources (Brookes *et al.*, 2007), but neither provides adequate control of the problem. Recent developments of nulling beamformers (Dalal *et al.*, 2006; Hui *et al.*, 2010), and the highly flexible inverse solutions of Hauk & Stenroos (2014) provide better alternatives. A solution can be estimated for each region of interest, and cross-talk between the spatial locations of these regions prevented. If the regions are good characterisations of the true sources, this offers a promising approach for connectivity analyses, although correlations between the inferred sources are not guaranteed to be removed, and this method has yet to be used or evaluated for this purpose.

Irrespective of source reconstruction method, MEG connectivity studies that seek to be completely robust to source leakage artefacts can choose network modelling methods that are unaffected by the induced correlations. Such meth-

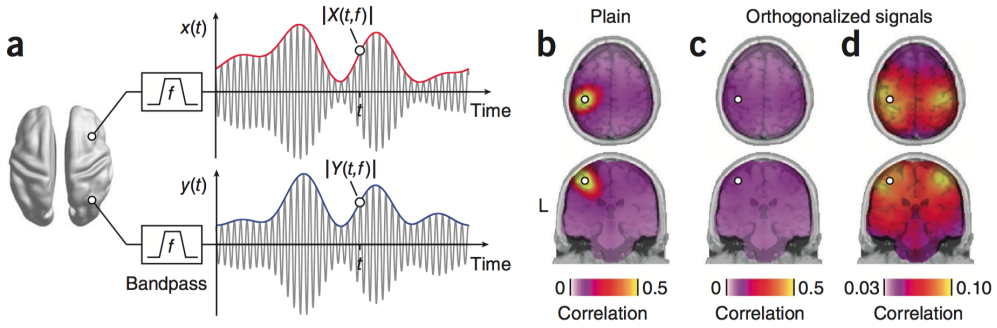


Figure 2.1: SUCCESS OF ORTHOGONALISATION FOR SEED-BASED CONNECTIVITY ANALYSES IN MEG. Figure reproduced from Hipp *et al.* (2012) to illustrate a traditional seed-based connectivity analyses. (A) The reconstructed time courses in two locations, one of which is chosen as the *seed,* are filtered into a specific frequency band. The seed's time course is regressed from the other, then the correlation between the power envelopes of the two computed. (B) Functional connectivity to a seed in the left motor cortex, when no orthogonalisation is performed. The heatmap of correlation is dominated by local correlations induced by the source reconstruction process. (C) Applying an orthogonalisation between the seed and test regions dramatically reduces the strength of correlations, but when the scale is altered (D), sensible bilateral resting-state motor connectivity is revealed.

ods either assess metrics of phase lag between sources, which are insensitive to linear leakage effects (Baccala & Sameshima, 2001; Kaminski & Blinowska, 2014; Nolte *et al.*, 2004; Stam *et al.*, 2007), or attempt to correct for leakage before performing signal amplitude correlations. Authors using the latter technique have so far been limited to pair-wise comparisons of source activities between voxels (Brookes *et al.*, 2012*b*; Hipp *et al.*, 2012). Although this work has recently been extended from the analysis of two voxels to a Canonical Correlation Analysis (CCA) between a pair of ROIs[1] (Brookes *et al.*, 2014; O'Neill *et al.*, 2015), there has so far been no simple technique correcting amplitude correlations between larger collections of brain areas, which limits the scope of whole-brain ROI network analyses.

Source leakage is a linear effect, and it introduces cross-correlations only

[1] See also Soto *et al.* (2016), a statistical thresholding technique to prevent false inference when performing CCA in the presence of signal leakage.

at zero phase lag. Thus, the correlation artefacts caused by source leakage can be corrected (under certain assumptions) by removing zero-lag correlations between the variables of interest, before performing any connectivity analyses. The existing methodology (Brookes *et al.*, 2012*b*; Hipp *et al.*, 2012) is to orthogonalise reconstructed source time courses with respect to a seed voxel (which removes their mutual correlation at zero lag), then to analyse frequency-dependent connectivity by correlating the down-sampled power envelopes of the orthogonalised sources within specific frequency bands. This process has been successful for enabling seed-based connectivity analyses (a clear illustration of the procedure, and its results, is shown in figure 2.1[2]). The approach is, however, limited to pair-wise comparisons between the seed and the rest of the brain, and the correlations are sensitive to which of the two voxels is chosen as the seed. Often a compromise is taken, by averaging the correlations in band-limited power, for each pair of voxels, over the two choices of seed (Hipp *et al.*, 2012).

More subtle leakage corrections that attempt to model the spatial profile of signal leakage have since been developed (Wens, 2014; Wens *et al.*, 2015). These build a model of the leaked signal from a seed voxel, on the assumption that the signal reconstructed at that seed is a good estimate of the underlying source, and subtract this leaked signal estimate from all other voxels, before performing a connectivity analysis. This approach potentially improves longer-range sensitivity compared to the simple, brute-force orthogonalisation method, but the limitation on pair-wise analyses has not been broken.

This chapter presents a symmetric, multivariate, spatial leakage correction

[2] Figure 2.1 is reproduced by permission from Macmillan Publishers Ltd: Nature Neuroscience (Hipp *et al.*, 2012); © (2012) Macmillan Publishers.

that extends the orthogonalisation method to enable corrected multivariate network analyses between multiple ROIs. We remove the effects of signal leakage at zero lag using a symmetric orthogonalisation process over all ROIs which has three valuable properties: firstly, the outcome is not biased by any ordering of the ROI time courses; secondly, the resulting orthogonal vectors are as close as possible to the original time courses; lastly, the computation is based on the singular value decomposition, which is a highly optimised routine on most computational platforms. The number of regions which can be corrected is limited only by the dimensionality of the recorded data. Here, we demonstrate the method by modelling the connected ROIs as a GGM, and estimate the edges of the graph (that is, the connections between ROIs) using regularised partial correlations between the band-limited power envelopes of the corrected ROI time courses. We validate our approach on simulated networks of dipoles within the cortex and on resting-state data from a small sample of healthy subjects.

2.2 Theory

CHARACTERISATION OF THE PROBLEM

Signal leakage

The signals measured over N samples in the P magnetic field sensors of an MEG scanner, $\boldsymbol{Y} \in \mathbb{R}^{P \times N}$, are typically modelled as a linear combination of sources, $\boldsymbol{S} \in \mathbb{R}^{Q \times N}$, at a set of locations inside the brain $\left\{\boldsymbol{x}_q\right\}_{i=q}^{Q}$, and additive Gaussian measurement noise, $\boldsymbol{\epsilon}$, with variance σ_ϵ^2,

$$\boldsymbol{Y} = \boldsymbol{H}\boldsymbol{S} + \boldsymbol{\epsilon} \tag{2.1}$$

where the lead fields $H_{ij}(\boldsymbol{\theta}_j)$ encode the mapping between the field measured at sensor i from a source at location $\boldsymbol{x}_j$ with fixed orientation $\boldsymbol{\theta}_j$, after propagation through the brain and skull. Source reconstruction approaches, such as beamformers (van Drongelen *et al.*, 1996; Van Veen *et al.*, 1997), attempt to infer estimates of source activity, $\hat{\boldsymbol{S}}$, at locations $\left\{\boldsymbol{x}_q\right\}$ using a linear combination of signals from the sensors, defined by a set of weights $\boldsymbol{W} \in \mathbb{R}^{P \times Q}$,

$$\hat{\boldsymbol{S}} = \boldsymbol{W}^{\mathrm{T}}\boldsymbol{Y} = \boldsymbol{W}^{\mathrm{T}}\boldsymbol{H}\boldsymbol{S} + \boldsymbol{W}^{\mathrm{T}}\boldsymbol{\epsilon} \; . \tag{2.2}$$

Each source is therefore estimated as a mix of the true sources within the brain. The signal leakage becomes apparent when we consider the covariance between two sources estimated from a dipole configuration that is truly uncorrelated, $\boldsymbol{S} \sim \mathcal{N}\left(\boldsymbol{0},\, \sigma_s^2\, \boldsymbol{I}_Q\right)$:

$$\begin{aligned} \mathrm{Cov}\left[\hat{\boldsymbol{S}}\right] &= \frac{1}{N}\left\{\boldsymbol{W}^{\mathrm{T}}\boldsymbol{H}\boldsymbol{S}\boldsymbol{S}^{\mathrm{T}}\boldsymbol{H}^{\mathrm{T}}\boldsymbol{W} + \boldsymbol{W}^{\mathrm{T}}\boldsymbol{\epsilon}\boldsymbol{\epsilon}\boldsymbol{W}\right\} \\ &= \left\{\sigma_s^2\boldsymbol{W}^{\mathrm{T}}\boldsymbol{H}\boldsymbol{H}^{\mathrm{T}}\boldsymbol{W} + \sigma_\epsilon^2\boldsymbol{W}^{\mathrm{T}}\boldsymbol{W}\right\} \; , \end{aligned} \tag{2.3}$$

as the covariance between any two estimated sources, $\hat{\boldsymbol{s}}(\boldsymbol{x}_i)$ and $\hat{\boldsymbol{s}}(\boldsymbol{x}_j)$,

$$\begin{aligned} \mathrm{Cov}\left[\hat{\boldsymbol{s}}(\boldsymbol{x}_i), \hat{\boldsymbol{s}}(\boldsymbol{x}_j)\right] &= \frac{1}{N}\left\{\hat{\boldsymbol{s}}(\boldsymbol{x}_i)^{\mathrm{T}}\,\hat{\boldsymbol{s}}(\boldsymbol{x}_i)\right\} \\ &= \left\{\sigma_s^2\boldsymbol{w}(\boldsymbol{x}_i)^{\mathrm{T}}\,\boldsymbol{H}\boldsymbol{H}^{\mathrm{T}}\boldsymbol{w}(\boldsymbol{x}_j) + \sigma_\epsilon^2\boldsymbol{w}(\boldsymbol{x}_i)^{\mathrm{T}}\,\boldsymbol{w}(\boldsymbol{x}_j)\right\} \; , \end{aligned} \tag{2.4}$$

is non-zero, and dependent on the data, lead fields and the noise. Lead fields from neighbouring spatial locations tend to be highly correlated, so this mixing takes the form of local signal leakage, creating correlations with zero temporal lag between neighbouring spatial regions.

Bi-variate orthogonalisation and its failure

We will illustrate the existing approach for leakage correction using orthogonalisation, and its weaknesses, using a simplified example. Consider two regions or voxels, A and B, between which there is little leakage. The inferred signals at A and B, denoted $\hat{\boldsymbol{s}}$, we therefore take to be good estimates of the true behaviour of the region, $\boldsymbol{s}$. In our example, we assume a functional connection exists between A and B such that the activity at B is identical to the activity at A, but time-lagged by n points.

$$\hat{s}_A[t] = s_A[t] \tag{2.5}$$

$$\hat{s}_B[t] = s_B[t] = s_A[t-n] \tag{2.6}$$

where t denotes time in samples. Now consider our estimate of the signal in a region C that is proximate to B, and to which it has no functional connection,

$$\begin{aligned}\hat{s}_C[t] &= bs_B[t] + \epsilon_C[t] \\ &= bs_A[t-n] + \epsilon_C[t]\end{aligned} \tag{2.7}$$

where an amount b of the signal at $\boldsymbol{x}_B$ has leaked into $\hat{\boldsymbol{s}}_C$, and for simplicity we define the residual signal at $\boldsymbol{x}_C$, $\boldsymbol{\epsilon}_C$, to have no cross-correlation with $\boldsymbol{s}_B$ or $\boldsymbol{s}_A$ at any time-lag. It may be intuitive from equation 2.7 that no bivariate linear subtraction or orthogonalisation process will be able to completely remove the influence of region A on C.

To see this explicitly, compare an attempt to correct $\hat{\boldsymbol{s}}_C$ for leakage from regions A and B using a bi-variate orthogonalisation process. We can construct

a version of the inferred signal at $\boldsymbol{x}_C$ that removes the influence of $\boldsymbol{s}_B$ on C,

$$\begin{aligned}\hat{\boldsymbol{s}}_{C\perp B} &= \hat{\boldsymbol{s}}_C - b\hat{\boldsymbol{s}}_B \\ &= \boldsymbol{\epsilon}_C \ .\end{aligned} \tag{2.8}$$

An analysis continuing in the style of Brookes *et al.* (2012*b*) would compute the correlation between the power envelopes of corrected time courses $\hat{\boldsymbol{s}}_B$ and $\hat{\boldsymbol{s}}_{C\perp B}$, repeating the procedure over all pairs of nodes to build up a network matrix. It would not be possible to perform other network analyses on the corrected data unless the connectivity measure could be built up using pairs of ROIs at a time—multivariate network analyses would not be feasible.

If we now attempt to perform the same process between C and A,

$$\begin{aligned}\hat{\boldsymbol{s}}_{C\perp A} &= \hat{\boldsymbol{s}}_C - \frac{\hat{\boldsymbol{s}}_C \cdot \hat{\boldsymbol{s}}_A}{|\hat{\boldsymbol{s}}_A|^2}\hat{\boldsymbol{s}}_A \\ \hat{s}_{C\perp A}[t] &= b s_A[t-n] + \epsilon_C[t] - \frac{b\sum_l s_A[l-n]\, s_A[l]}{|\boldsymbol{s}_A|^2} s_A[t] \ ,\end{aligned} \tag{2.9}$$

we find that even the orthogonalised signals at $\boldsymbol{x}_C$ and $\boldsymbol{x}_A$ can have a non-zero cross correlation for non-zero time lags, when there was no such underlying relationship. The only resolution to the problem is to use a multivariate leakage correction that removes the influence of both A *and* B from C.

A SYMMETRIC, MULTIVARIATE CORRECTION FOR SIGNAL LEAKAGE

Signal leakage is corrected for by removing any correlations with zero temporal lag between all ROIs. A multivariate, symmetric orthogonalisation technique (Everson, 1998; Löwdin, 1950) produces an optimal set of mutually orthogonal time courses for each ROI: the solution is unique, unaffected by any

re-ordering of ROIs and is minimally displaced from the uncorrected ROI time courses (as measured by the least-squares distance).[3]

In the following, we will characterise the behaviour of each ROI by a single time course. The particulars of this definition are unimportant here; common methods in fMRI and MEG include taking the mean time course, or the coefficients of the principal component accounting for the majority of the ROI's variance; some authors in the MEG literature select just the single voxel within the ROI with maximal power (e.g., Hillebrand *et al.*, 2012).

We construct the corrected time courses in two stages, illustrated in figure 2.2. Firstly, we find the closest set of orthonormal time courses, for which there is a simple analytic solution. Secondly, we relax the normality constraint and finesse the solution by iteratively adjusting the lengths and orientations of the corrected vectors until we converge to a solution which is as close as possible to the uncorrected time courses.

We outline the proof from Everson (1998) here. For a set of n ROI time courses containing m time samples $\boldsymbol{Z} = \{\boldsymbol{z}_1, \boldsymbol{z}_2, \cdots, \boldsymbol{z}_n\} \in \mathbb{R}^{m \times n}$, $m \geq n$, we seek the matrix of corrected time courses $\boldsymbol{P}$ that minimises the Frobenius norm

$$\epsilon = ||\boldsymbol{Z} - \boldsymbol{P}||_F^2 \equiv \text{trace}\left[(\boldsymbol{Z} - \boldsymbol{P})^{\mathrm{T}} (\boldsymbol{Z} - \boldsymbol{P})\right] \ . \tag{2.10}$$

By constraining $\boldsymbol{P}$ to contain orthogonal vectors, it can be expressed as

$$\boldsymbol{P} = \boldsymbol{O}\boldsymbol{D} \tag{2.11}$$

where $\boldsymbol{O}^{\mathrm{T}}\boldsymbol{O} = \boldsymbol{I}_n$ and $\boldsymbol{I}_n$ is the $n \times n$ identity matrix, $\boldsymbol{D} = \text{diag}\,(\boldsymbol{d})$ is diagonal,

[3] A broader discussion of methods for producing optimal orthonormal decompositions of data, under various cost functions, can be found in Kessy *et al.* (2015). The authors do not touch on optimal methods for retaining magnitude information in the vectors.

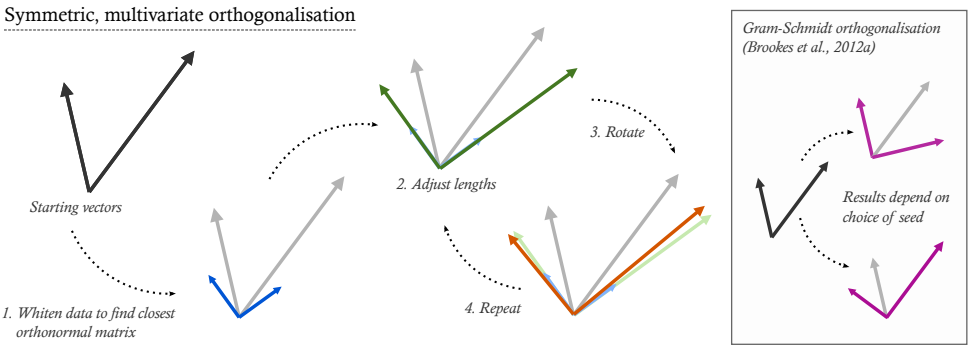


Figure 2.2: A SYMMETRIC, MULTIVARIATE ORTHOGONALISATION PROCESS. The correlations between ROIs that are introduced during source reconstruction can be removed by mutually orthogonalising the ROI time courses, illustrated here for two vectors in two dimensions. An optimal set of corrected time courses can be constructed by iterating towards the closest set of orthogonal vectors to the starting time courses. The process is initialised with the closest orthonormal matrix to the uncorrected vectors, then adjusts in turn the vector magnitudes and orientations to minimise the Euclidean distance between the corrected and uncorrected time courses. The box, right, shows the result of regressing one vector from the other, with differing outcomes depending on the choice of seed, as used for correcting pairs of voxels in Brookes *et al.* (2012*b*).

and $\boldsymbol{P}^{\mathrm{T}}\boldsymbol{P} = \boldsymbol{D}^2$. Hence,

$$\epsilon = \operatorname{trace}\left(\boldsymbol{Z}^{\mathrm{T}}\boldsymbol{Z}\right) - 2\operatorname{trace}\left(\boldsymbol{Z}^{\mathrm{T}}\boldsymbol{O}\boldsymbol{D}\right) + \operatorname{trace}\left(\boldsymbol{D}^2\right) \ . \tag{2.12}$$

If $\boldsymbol{D}$ is known and $\boldsymbol{Z}$ is full rank, ϵ is minimised by maximising trace $\left(\boldsymbol{Z}^{\mathrm{T}}\boldsymbol{O}\boldsymbol{D}\right)$. Using the singular value decomposition of $\boldsymbol{Z}\boldsymbol{D}$,

$$\boldsymbol{Z}\boldsymbol{D} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^{\mathrm{T}} , \tag{2.13}$$

$$\boldsymbol{\Sigma} = \operatorname{diag}\left(\sigma_1, \sigma_2 \ \ldots, \sigma_n\right), \ \sigma_i > 0 \ \ \forall\, i \ ,$$

we can bound expression (2.10) by noting that all elements of the orthonormal

matrix $\boldsymbol{T} = \boldsymbol{U}^{\mathrm{T}}\boldsymbol{O}\boldsymbol{V}$ must have a magnitude $|T_{ij}| \leq 1 \; \forall \; (i, j)$.

$$\operatorname{trace}\left(\boldsymbol{Z}^{\mathrm{T}}\boldsymbol{O}\boldsymbol{D}\right) = \operatorname{trace}\left(\boldsymbol{T}\boldsymbol{\Sigma}\right) \leq \sum_{i=1}^{n} \sigma_i \tag{2.14}$$

is maximal when $\boldsymbol{T} = \boldsymbol{I}_n$; thus, (2.10) can be minimised by setting

$$\boldsymbol{O} = \boldsymbol{U}\boldsymbol{V}^{\mathrm{T}}, \tag{2.15}$$

which is the symmetric or Löwdin orthogonalisation (Löwdin, 1950) of $\boldsymbol{Z}\boldsymbol{D}$.

On the other hand, if $\boldsymbol{O}$ is known, the optimal $\boldsymbol{d}$ is given by

$$\frac{\partial \epsilon}{\partial d_k} = -2 \sum_{i=1}^{m} Z_{ik} O_{ik} + 2 d_k = 0 \tag{2.16}$$

which yields

$$\boldsymbol{d} = \operatorname{diag}\left(\boldsymbol{Z}^{\mathrm{T}}\boldsymbol{O}\right) . \tag{2.17}$$

To find the corrected time courses $\boldsymbol{P}$, we follow Everson's tandem algorithm by starting with $\boldsymbol{D}^{(1)} = \boldsymbol{I}_n$, which gives $\boldsymbol{P}^{(1)}$ as the symmetric orthogonalisation of $\boldsymbol{Z}$ and the unique closest ortho*normal* matrix to the uncorrected ROI time courses; we then allow the vector magnitudes to vary and reduce the error ϵ by iterating (2.13), (2.15) and (2.17) until convergence (illustrated in figure 2.2). Convergence is guaranteed, but not necessarily to a global minimum of ϵ. In our data, this procedure tends to converge within twenty iterations.

This orthogonalisation process to find corrected time courses $\boldsymbol{P}$ by definition removes all correlations between ROIs at zero phase lag. Any remaining correlations between the band-limited power envelopes of these orthogonalised time courses are thought to represent true biological dependencies, at the

expense of true zero-phase-lag connectivity (Brookes *et al.*, 2012*b*; Hipp *et al.*, 2012; Luckhoo *et al.*, 2012). The approach is limited by the rank of the data: you cannot correct more ROIs than there are dimensions in $\boldsymbol{Z}$ as there is no longer a unique solution to (2.14).

NETWORK ANALYSIS

Following Brookes *et al.* (2012*b*) and Luckhoo *et al.* (2012), we compute power envelopes of the corrected ROIs, $\tilde{\boldsymbol{P}}$, as the absolute values of the analytic Hilbert transform of $\boldsymbol{P}$, which are then low-pass filtered to 0.5 Hz and re-sampled at 1 Hz.

We estimate the connections between ROIs by modelling the corrected and down-sampled power envelopes with an undirected GGM, $\tilde{\boldsymbol{P}} \sim \mathcal{N}_n\left(0,\, \boldsymbol{\Omega}^{-1}\right)$, where $\boldsymbol{\Omega}$ is the precision or inverse covariance matrix. The network connections between ROIs—the edges of the graph—are estimated from the partial correlation matrix $\boldsymbol{\rho}_\perp$, which is the conditional correlation between variables with the effect of all other variables removed,

$$\boldsymbol{\rho}_\perp = -\ \mathrm{diag}\left(\boldsymbol{\Omega}\right)^{-1/2} \boldsymbol{\Omega}\ \mathrm{diag}\left(\boldsymbol{\Omega}\right)^{-1/2}\ . \tag{2.18}$$

Zeros in the partial correlation matrix identify variables that are conditionally independent (unconnected by a network edge), thus partial correlations relate only to direct connections in the network graph (Fisher, 1924), whereas full correlations will also identify indirect associations. The estimation of partial correlations from a limited dataset is noisy (Varoquaux & Craddock, 2013; Varoquaux *et al.*, 2010): it is standard practice to regularise the estimate of the precision matrix, $\hat{\boldsymbol{\Omega}}$, and suppress null edges using the graphical LASSO (Fried-

man *et al.*, 2008). This is a sparsity promoting procedure that maximises the log-likelihood of the multivariate Gaussian model, subject to a penalty on the L_1 norm of the precision (the sum of the absolute values in $\boldsymbol{\Omega}$),

$$\hat{\boldsymbol{\Omega}} = \max_{\boldsymbol{\Omega}} \left[\log\det \boldsymbol{\Omega} - \mathrm{tr}\left(\frac{1}{m}\tilde{\boldsymbol{P}}^T\tilde{\boldsymbol{P}}\,\boldsymbol{\Omega}\right) - \lambda||\boldsymbol{\Omega}||_1 \right] . \tag{2.19}$$

The strength of the regularisation penalty is determined by a free parameter λ, which can be chosen by cross-validation.

2.3 Methods

We test the efficacy of our approach on both real and simulated resting-state datasets. Data processing and analyses were performed using in-house Matlab scripts, C routines from the University of Newcastle, Australia (Ninness *et al.*, 2005), `SPM8`, `FieldTrip` (Oostenveld *et al.*, 2011) and `FSL` (Jenkinson *et al.*, 2012). The graphical LASSO was solved using custom implementations of the algorithms in Mazumder & Hastie (2012*a*,*b*). The analysis software is hosted on GitHub and is freely available.[4]

GENERATION OF SIMULATED DATA

Fifteen resting-state MEG experiments of 600 s were simulated. In each experiment, five equivalent current dipoles were active, positioned in the left and right frontal gyrus, the left and right lateral occipital cortex, and the right premotor cortex. Beamforming methods for source reconstruction were used, which have imperfect reconstruction success with sources that are highly linearly-

[4] `github.com/OHBA-analysis/MEG-ROI-nets/`

correlated over the whole acquisition (Gross *et al.*, 2001; Robinson & Vrba, 1999; Van Veen *et al.*, 1997). We therefore simulated each dipole as oscillating with amplitude 1 nAm at a different carrier frequency in the range 8–26 Hz, modulated by a functional activity that was generated by a simple 5-node network model drawn from Smith *et al.* (2011), based on the dynamic causal model for fMRI activity (Friston *et al.*, 2003). In brief, amplitudes or activities $\boldsymbol{a}$ are generated according to

$$\dot{\boldsymbol{a}} = \boldsymbol{A}\boldsymbol{a} + \boldsymbol{u} + \boldsymbol{e}\,, \tag{2.20}$$

where the network edges are non-zero entries on the off-diagonal elements of $\boldsymbol{A}$, with a strength randomly chosen from $\mathcal{N}(0.6,\ 0.1)$; the diagonal elements of $\boldsymbol{A}$ are set to -1, modelling within-node temporal decay; a set of inputs, $\boldsymbol{u}$, drive each node, modelled by a binary Poisson process[5] of strength 0.4, with a mean on-time of 2 seconds and a mean off-time of 7 seconds; and a Gaussian noise input $\boldsymbol{e}$ is drawn from $\mathcal{N}(0,\ 0.02)$. Equation (2.20) was integrated using a fourth-order Runge-Kutta method. The network structure is shown in figure 2.3D on the following page, together with example activity time courses and modulated oscillations for a simpler 2-node network in figures 2.3E and 2.3F.

From these dipole sources, MEG data were simulated for the 306 sensors on an Elekta Neuromag MEG instrument (Stockholm, Sweden), at a sampling rate of 150 Hz, using a multiple local sphere volume conductor (Huang *et al.*, 1999) and the Sarvas (1987) current dipole forward model. Sensor noise was added by sampling from a zero-mean multivariate Gaussian distribution, using a covariance matrix estimated from ten minutes of empty-room recordings. The noise was scaled to produce a unit signal-to-noise power ratio over all the

[5] The probability of waiting $t \geq 0$ seconds for a transition from one state, with mean occupation time τ seconds, to the other is given by the exponential distribution, $p\,(t|\tau) = \frac{1}{\tau}e^{-t/\tau}$.

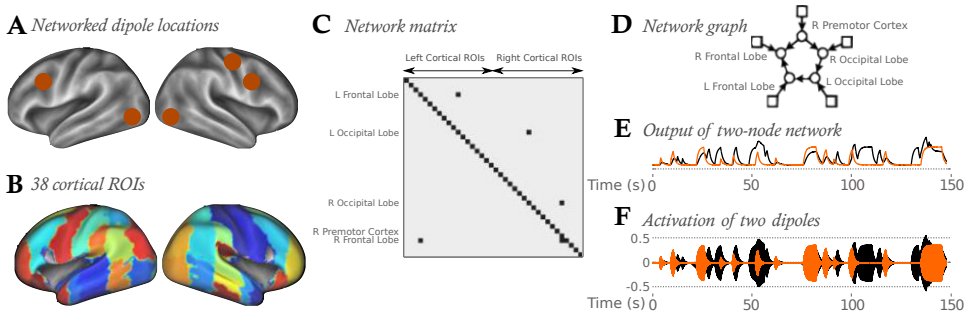


Figure 2.3: SIMULATION OF A FIVE-DIPOLE NETWORK. Five dipoles with differing carrier frequencies in the range 8–26 HZ, modulated in amplitude by a simple five-node network drawn from Smith *et al.* (2011), were placed into five ROIS of a 38-ROI parcellation of the cortex, and the signals recorded in a set of 306 MEG sensors were simulated with additive measurement noise. (A) shows the locations of the five simulated dipoles; (B) a surface rendering of the 38 ROIS drawn from a laterally-split version of the Harvard-Oxford cortical atlas in FSL; (C) shows the network matrix used to simulate the activities in the five dipoles, placed within context of all 38 ROIS considered in later analysis; (D) shows the network as a graph, indicating the directions of the network edges and the Poisson processes stimulating each node independently; (E) shows example amplitude time courses from a simpler two-node network, with one node (orange) driving the other (black), and both stimulated by independent Poisson processes; and (F) displays the the modulation of the activations of two dipoles by the amplitudes in (E).

sensors (SNR, calculated as the ratio of the mean-square signal amplitude in all the sensors to the variance of the noise).

A second dataset of fifty experiments was created using the same methodology, but increasing the number of dipoles to 38, chosen to match the number of ROIS used in later analysis. Only five dipoles were networked, using the same structure as above; the rest were also generated from (2.20) but without any network structure and uncorrelated. Dipoles were randomly located such that each ROI contained one dipole, and the ROIS containing networked dipoles varied between experiments. The simulated SNR in the sensors was 0.4. This batch of experiments was repeated with a denser network structure, in which fifteen dipoles were networked. The network structure is shown in figure 2.7A on page 55.

RESTING-STATE DATA COLLECTION AND PRE-PROCESSING

Ten minutes of MEG data were acquired from ten healthy subjects (7 males, 3 females of which one was left-handed; 27±0.5 years of age) in a resting-state using a 275-channel CTF whole-head system (DC, Canada) at a sampling rate of 600 HZ with a 150 HZ low-pass anti-aliasing filter applied. Subjects were asked to lie in the scanner and view a centrally-presented fixation cross. The study was approved by the University of Nottingham Medical School Research Ethics Committee (approval code F/12/2006). All volunteers received a study information sheet, completed a safety questionnaire, and provided written informed consent, including consent to publish anonymised results.

Head localisation within the MEG helmet was achieved by measuring the locations of three energised electromagnetic coils, taped to the head, using a magnetic dipole fit. The positions of the coils relative to the subject's head shape were measured prior to acquisition using a 3D digitiser (Polhemus Isotrack). A structural MR image was acquired for each subject using an MPRAGE sequence on a Philips Achieva 3T (1 mm^3 resolution, TR = 8.3 ms, TE = 3.9 ms, TI = 960 ms, a shot interval of 3 s, FA = 8°, and a SENSE factor of 2). Each structural scan was registered to the MNI152 standard brain template. The location of the MEG sensors relative to each subject's brain anatomy was found by matching the digitised head-shape to a scalp surface extracted from the MNI template and transformed (with 12 degrees of freedom) into the space of the subject's structural scan. Two subjects' data were discarded at this stage because of poor alignment. All further source-space analysis was carried out in MNI space.

Sensor data were down-sampled to 200 HZ. Any channels or segments of data showing obvious artefacts were removed following visual inspection.

ICA (Hyvärinen, 1999) was used to further remove artefacts from the data as follows: sensor data were decomposed into 150 temporally independent components; artefact components were manually classified as relating to eye blinks (a high kurtosis, with a value over 20, and a repetitive blink structure over time), cardiac sources (strong resemblance to typical electrocardiogram signals and a kurtosis over 20) or mains interference (dominant 50 HZ component in the frequency spectrum); the artefact components were then regressed out of the sensor data (Mantini *et al.*, 2011), which were subsequently band-pass filtered to 4–30 HZ.

SPECIFICATION OF REGIONS OF INTEREST

Separate cortical parcellations were employed for the analysis of the simulated and real data, shown in figures 2.3B and 2.8A. ROIS were defined for the simulated data by selecting a subset of 19 of the ROIS in the Harvard-Oxford cortical brain atlas, available in FSL, and splitting each into two lateral halves to create 38 binary ROIS. For the analysis of the resting-state data, a parcellation based on functional data was desired, which would define regions of interest based on functional specificity. We sought a parcellation which would allow us to investigate connections between and within both hemispheres and between individual, contiguous regions. We took a 100-dimensional group-spatial-ICA decomposition from resting-state fMRI scans of the first two hundred subjects in the HCP database (Van Essen *et al.*, 2013), providing a set of spatial maps with highly local functional areas, although mostly with bilateral distributions. We lateralised each component, and ran the resulting maps through a cluster-identifying algorithm in FSL, thresholding at $z = 3.1$, and retaining the cluster from each map with the maximum peak intensity. We then selected 38 of these clusters, which

provided a largely non-overlapping spanning of the cortex, choosing where possible pairs of ROIs matched in each hemisphere. The z-statistic weightings of the spatial maps that survived the cluster thresholding were retained within the ROIs.

DATA PROCESSING

The same pipeline was used to process both simulated and real datasets.

Dipole sources on an 8 mm uniform grid were reconstructed using a scalar beamformer (Robinson & Vrba, 1999; Van Veen *et al.*, 1997; Woolrich *et al.*, 2011), a spatial filtering technique that uses a minimum-variance constraint to reconstruct signals at the locations of interest, while suppressing power from other locations. The data covariance matrix, used to estimate the beamformer reconstruction weights, was estimated from the entirety of each subject's data, and regularised by reducing the data dimensionality with Principal Component Analysis (PCA), retaining between 260 and 264 components. (The simulated data were not passed through the `maxfilter` data cleaning software from Elekta, which substantially reduces the dimensionality to around 70 components.)

We represent the behaviour of each ROI with a single time course, obtained using PCA. The reconstructed sources within each ROI were first bandpass-filtered: the simulated data were analysed over a 4–30 Hz band, and the real MEG data within the alpha-band (8–13 Hz) and beta-band (13–30 Hz), chosen because many RSNs present well in these frequency ranges (Brookes *et al.*, 2011*b*). Each ROI was normalised to have a positive peak height of unity. The coefficients of the principal component accounting for the majority of the variance of the voxels within the ROI, weighted by the ROI map, were then taken as an appropriate representation of source activity for that region.

Signal leakage with zero temporal lag, induced by the beamforming process, was removed by symmetrically orthogonalising all ROI time courses (section 2.2). Following previous analyses (Brookes *et al.*, 2012*b*; Luckhoo *et al.*, 2012), the power envelopes of these corrected ROI time courses were then found by taking the absolute value of their Hilbert transform, low-pass filtering to 0.5 Hz and downsampling to 1 Hz. The pipeline of ROI time course extraction, band-pass filtering, orthogonalisation and correlation is illustrated in figure 2.4.

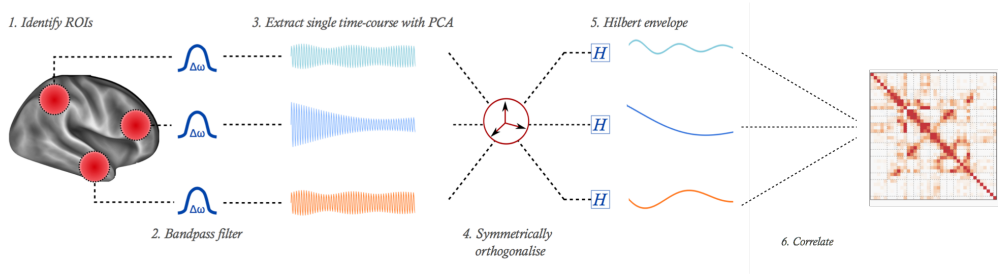


Figure 2.4: ILLUSTRATION OF THE PIPELINE FOR MEG FUNCTIONAL CONNECTIVITY ESTIMATION. From a whole-brain source-reconstruction, a single time course is extracted to represent each ROI. These time courses are bandpass filtered, then orthogonalised to remove shared signal which is potentially attributable to spatial leakage effects. The power envelope of each time course is computed, then correlated to form the network matrix.

CONNECTIVITY ANALYSIS

Linear correlation and partial correlation (equations 2.18 and 2.19) analyses were performed on the down-sampled power envelopes of the corrected ROI time courses of each subject or simulation. We chose the strength of regularisation applied to the partial correlation matrix for each subject using 10-fold within-subject cross-validation, which provides a data-driven assessment of the generalisability of the chosen value, where the optimal λ was that which minimised the corrected Akaike information criterion (Burnham & Anderson, 2002), an information-theoretic model selection criterion which maximises a model's goodness of fit while penalising complexity. After the initial parameter

selection, the search grid for λ was finessed three times.

Correlation values were converted to z-scores using Fisher's transformation.[6] In order to test the null hypotheses that there were no correlations between pairs of variables, the z-scores for each edge were scaled such that the distributions of edge scores would be standard normal under fulfilment of the null. We computed the scaling as the standard deviation of an empirical null correlation distribution, generated by simulating eight iterations of a null, uncorrelated Gaussian dataset, sharing the same number of nodes and temporal smoothness (modelled using the first auto-regressive coefficient) as the ROI time courses, but with no spatial leakage, and taking correlations and unregularised partial correlations between the down-sampled, band-limited power envelopes of these data. This non-parametric approach is designed to provide accurate inference of the correlation structure, even in the presence of temporal autocorrelations in the power time courses.

Group inference on the network structure was performed using a fixed-effects analysis testing the significance of the mean, over all subjects, of the standardised full or partial correlation z-scores between each pair of ROIs. (We use fixed effects because we are not primarily interested here in subject variability, which would likely be poorly estimated from only eight subjects. The fixed effects analysis is telling us about the strength of the group-average effect.) We control for false positives in the group inference at the network level using a False Discovery Rate (FDR) threshold.

[6] Fisher's transformation, $Z(\rho) = \operatorname{artanh}(\rho)$ converts correlations sampled under a null distribution into normally-distributed variables. It is a monotonic transformation, producing greater separation between values as correlations approach unity.

METHOD COMPARISONS

The source-leakage correction method outlined above was compared to two alternatives. Firstly, not employing any correction at all. Secondly, a pair-wise orthogonalisation procedure, which employs the pair-wise orthogonalisation steps previously used for voxel-wise analyses (Brookes *et al.*, 2012*a*; Luckhoo *et al.*, 2012) on the ROI time courses. This works as follows: one ROI time course is regressed from the other, power envelopes are found for the corrected time courses, and their correlation computed. To estimate partial correlations between the same two nodes, the corrected time course for each is regressed from all other ROI time courses; power envelopes are found over all ROIs and the partial correlation between the power envelopes of the two ROIs of interest is computed. This procedure is iterated over all pairs of ROIs. Different correlations are found between each pair depending on which is chosen as the seed (i.e., which ROI time course is regressed from the other). The two correlations are averaged to produce the final estimate. Developing an approach to regularise these partial correlation estimations, using penalised regression approaches, is not trivial and is not relevant here; we restrict our comparison between the multivariate source-leakage correction and this pair-wise orthogonalisation procedure to unregularised estimates of partial correlations.

This comparative pair-wise procedure is significantly more computationally costly: to perform the symmetric orthogonalisation and cross-validated computation of regularised partial correlations between 38 node time courses of one of the resting-state datasets took 70 seconds on a Macbook Pro with a 2.8 GHz Intel Core i7 processor and 16 GB of RAM; performing the orthogonalisation and calculation of (unregularised) partial correlations in a pair-wise fashion took

2000 seconds.

2.4 Results

SIMULATED DATA

The performance of the symmetric source-leakage correction method for ROI network connectivity analyses (section 2.2) was assessed using a set of fifteen simulated experiments, with five networked dipoles, at an SNR of 1.0. The network structure and dipole locations are shown in figure 2.3 on page 44 and the simulations are described in section 2.3.

Figure 2.5 on the following page compares group-average estimates of the networked behaviour of the 38 ROIs in the brain, showing both the correlations and partial correlations between the power envelopes of all ROIs. We compare our proposed symmetric multivariate leakage correction to other forms of analysis: to not using a correction at all, and a pair-wise orthogonalisation procedure of ROI time courses, an extension of a previously-employed voxel-wise correction for seed-based correlation analyses (see section 2.3). We did not attempt to regularise the partial correlations computed for this pair-wise method, which would require a significant increase in algorithmic and computational complexity; in this figure we display only unconstrained partial correlations to compare like with like ($\lambda = 0$).

Using both correction methods, we find much stronger suppression of local source spread, and better identification of the true network structure (as assessed with both full and partial correlation), compared with the results when no correction was applied. The partial correlation matrices, being estimates

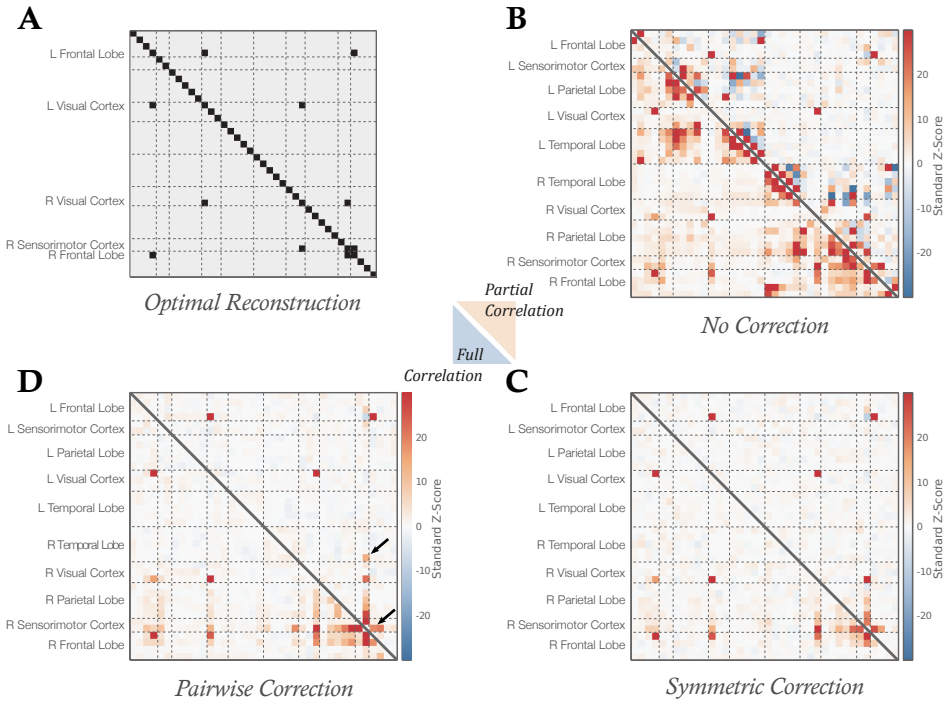


Figure 2.5: RECONSTRUCTION OF A FIVE-DIPOLE NETWORK: COMPARISON OF ABILITY TO INFER A SIMULATED NETWORK STRUCTURE UNDER DIFFERENT CORRECTIONS FOR SOURCE LEAKAGE. Five connected dipoles were simulated in the left and right occipital lobes, the right and left frontal lobes, and the right somatosensory cortex (see figure 2.3), for 15 experiments of duration 10 minutes, at an SNR of 1.0. Group-average z-statistics for the full and partial correlations between the power time courses of 38 cortical ROIs (figure 2.3B) are shown, after the application of corrections for source leakage. We combine two metrics into one heat map with full correlations below the diagonal and (unregularised) partial correlations above. We compare uncorrected data (B) to data adjusted using the symmetric multivariate correction presented in this chapter (C), and to data corrected for zero-lag correlations pairwise between ROIs (D). (C) is the only set of matrices that can be thresholded to achieve perfect reconstruction. The higher arrow highlights an example false positive, and the lower a weak true positive with a neighbouring strong false positive (see text). A symmetrical version of the simulated network matrix is shown (A), indicating the relationships between ROIs which would constitute a perfect network discovery for these (undirected) network analysis methods.

of the direct connections in the network, provide in general much better discrimination of the underlying network than the correlation matrices. (Neither method attempts to assign a direction to the functional connections.) While the spurious correlations present between the envelopes of the uncorrected ROI time courses are nearly eliminated by the orthogonalisation methods, the correlations (marginal and partial) between ROIs corresponding to 'true' edges are also to some extent weakened.

If we compare the symmetric multivariate correction (figure 2.5C) against the pair-wise orthogonalisation (figure 2.5D), we find the latter performs less well in suppressing artificial correlations, particularly among the full correlations. The partial correlations between these pair-wise corrected ROIs identified one additional false-positive network edge (indicated by the upper arrow), and a reduced correlation on a true network edge (lower arrow; with adjacent, stronger, false-positive) when compared to the symmetric correction, given an overall-optimal thresholding.

To understand if this behaviour was reproducible over a range of different dipole locations we extended our simulation. We generated fifty new datasets, with a reduced SNR of 0.4, in which each of the five dipoles was placed randomly in one of the 38 ROIs used for the analysis. We also placed a dipole in each of the remaining ROIs, and stimulated these individually with identically-generated but uncorrelated inputs to those driving the five networked dipoles. We then followed the same analysis procedure to arrive at correlation and partial correlation matrices for each of the fifty experiments, between all 38 ROIs, after applying each of the three source leakage correction methods. We investigated the extent to which each method removed spurious correlations between ROIs by comparing the false positive rates for single-edge detections with those ex-

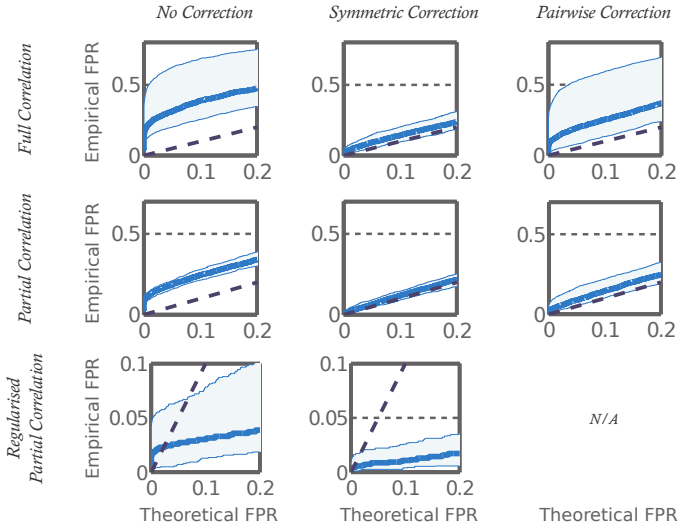


Figure 2.6: CONTROL OF FALSE POSITIVES DURING NETWORK RECONSTRUCTION. In each of fifty datasets, thirty-eight dipoles were simulated, of which five had activities driven by a directed network (figure 2.3D) and the rest were uncorrelated, and placed at random within 38 cortical ROIs (figure 2.3B), one dipole within each ROI. For each experiment, the full and partial correlations (calculated with and without regularisation of the precision matrix) between the band-limited power envelopes of the ROI time courses were computed and used to infer the network structure between ROIs. For each of these metrics, we compare the expected False-Positive Rate (FPR) for single edge detections against the empirical FPR as a measure of how well the spurious correlations between ROIs, introduced during source reconstruction, are removed by three orthogonalisation methods under test: applying no correction for spatial leakage (left column), applying the symmetric multivariate correction presented in this chapter (middle column) and applying a pairwise orthogonalisation approach to the ROI time courses (right column). Each graph plots the empirical FPR against the expected FPR as the threshold for defining network edges is moved. The solid blue line indicates the median behaviour over the fifty runs; the shaded background covers the central 95% of the data. An algorithm to combine the pairwise correction with regularised partial correlation has not been developed.

pected under a perfect correction for source leakage over a range of thresholds (figure 2.6).

Inference on network edges after applying the symmetric multivariate correction (figure 2.6, middle column) closely follows the expected false positive rates, with full correlation performing worse than partial correlation (a one-tailed Wilcoxon signed rank test for a difference in the median observed FPR, conducted at an expected FPR of 0.05 to match conventional thresholds, was sig-

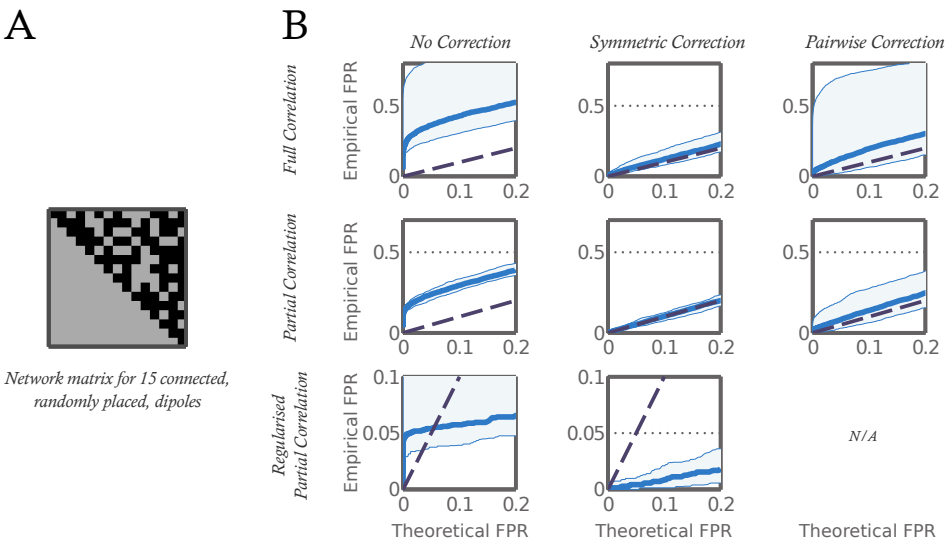


Figure 2.7: CONTROL OF FALSE POSITIVES DURING NETWORK RECONSTRUCTION OF A DENSE NETWORK. In each of fifty datasets, thirty-eight dipoles were simulated, of which fifteen had activities driven by a directed network (A) and the rest were uncorrelated, and placed at random within 38 cortical ROIs (figure 2.3B), one dipole within each ROI. For each experiment, the full and partial correlations (calculated with and without regularisation of the precision matrix) between the band-limited power envelopes of the ROI time courses were computed and used to infer the network structure between ROIs. For each of these metrics, the expected FPR for detection of each edge was compared against the empirical FPR as a measure of how well the spurious correlations between ROIs, introduced during source reconstruction, are removed by three orthogonalisation methods under test: applying no correction for spatial leakage (left column), applying the symmetric multivariate correction presented in this chapter (middle column) and applying a pairwise orthogonalisation approach to the ROI time courses (right column). Each graph plots the empirical FPR against the expected FPR as the threshold for defining network edges is moved. The solid blue line indicates the median behaviour over the fifty runs; the shaded background covers the central 95% of the data. An algorithm to combine the pairwise correction with regularised partial correlation has not been developed.

nificant, $p < 10^{-3}$), and dramatically improves upon these rates when regularisation is applied to the partial correlation matrix. Applying no leakage correction (figure 2.6, left column) leads to a much higher proportion of false positives, although regularisation of the partial correlation matrix is able to suppress many of the edges, reducing the false positive rate dramatically, although not to the same extent as for the corrected data. (Optimised regularisation strengths were similar: $\lambda = 3.1 \pm 1.1 \times 10^{-2}$ and $\lambda = 1.1 \pm 0.6 \times 10^{-2}$, for the uncorrected and

corrected data respectively.) Correcting for source leakage in a pairwise manner (figure 2.6, right column) is less effective at removing false positives than the full multivariate correction, although when the network estimation is based on the partial rather than full correlations, this discrepancy is much reduced (yet still significant: a Wilcoxon signed rank test for the median of the observed pair-wise corrected FPR values being greater than the symmetrically corrected FPR values, conducted at an expected FPR of 0.05, was significant, $p < 10^{-3}$).

The connectivity observed in the resting-state dataset is relatively dense. We repeated our simulations and analyses using a dense network, to ensure comparability between the results of our simulations and experiments. All aspects of this final simulation were identical to those described above, save that fifteen of the 38 dipoles randomly placed within 38 ROIs were densely networked. The network structure for those fifteen is given in figure 2.7A on the preceding page. The false positive rates for detecting network edges under no correction for leakage, a pair-by-pair orthogonalisation process and the multivariate, symmetric orthogonalisation presented in this chapter, are shown in figure 2.7B. We find, for all network inference methods and leakage correction methods under test, that the false positive rates inferred for edge detection exhibit the same behaviour for dense simulated networks as for sparse, with the observations made about figure 2.6 equally applicable to the densely-connected simulated dataset.

RESTING-STATE DATA

Functional connectivity was analysed for alpha- and beta-band oscillations between thirty-eight fMRI-derived ROIs (figure 2.8A on the next page) in resting-state MEG recordings from eight subjects, both with and without the application

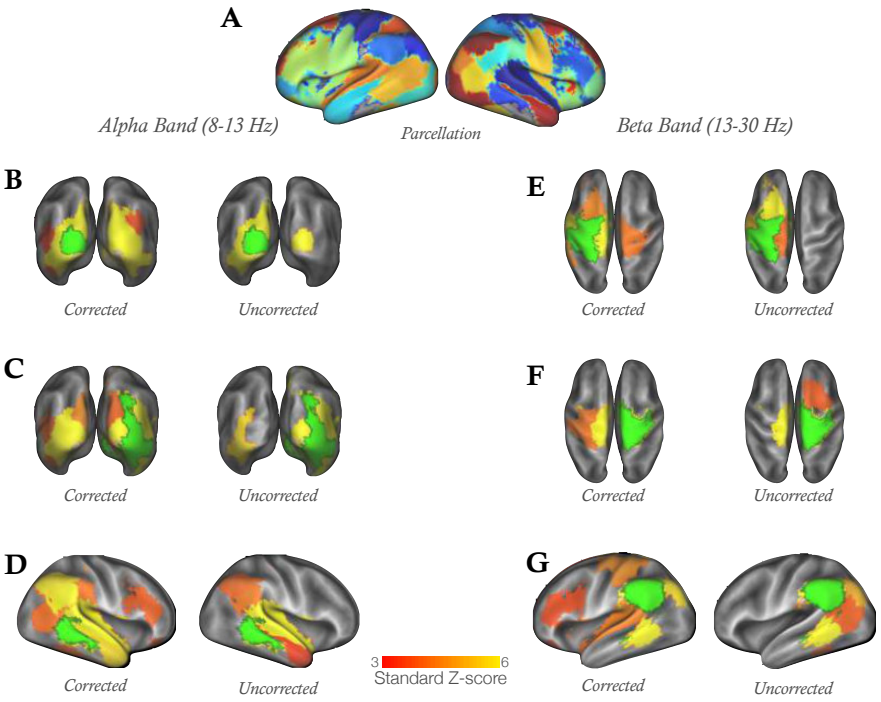


Figure 2.8: GROUP-AVERAGE RESTING-STATE DIRECT FUNCTIONAL CONNECTIONS DERIVED FROM EIGHT SUBJECTS, INFERRED WITH AND WITHOUT THE APPLICATION OF A MULTIVARIATE SPATIAL LEAKAGE CORRECTION BETWEEN ROIS. Surface plots in the alpha-band (8–13 HZ, B–D) and beta-band (13–30 HZ, E–G) show direct functional connections between a single seed ROI (green) and the remaining 37 ROIS derived from an FMRI ICA decomposition. The chosen seed ROIS are in the visual (B and C) and sensori-motor (E and F) cortices, the left parietal (G) and right temporal (D) lobes; these plots correspond to the appropriately indicated columns in the network matrices of figure 2.10. The colour scale indicates inference on the group-average of individual regularised partial correlation z-statistics. All the surface plots have been thresholded at a 5% false discovery rate ($z = 3.4$ for the alpha band and $z = 3.6$ for the beta-band, in both corrected and uncorrected datasets). A surface rendering of the 38 ROIS is given for reference (A).

of the proposed symmetric multivariate source leakage correction. The connectivity between ROIS was modelled using an undirected GGM (equations 2.18 and 2.19), applied to the down-sampled power envelopes of the ROI time courses in the alpha band (8–13 HZ) and beta band (13–30 HZ). We estimated the presence and strength of edges in the network using a group-average analysis of the regularised subject-specific partial correlation matrices between the ROIS. (Op-

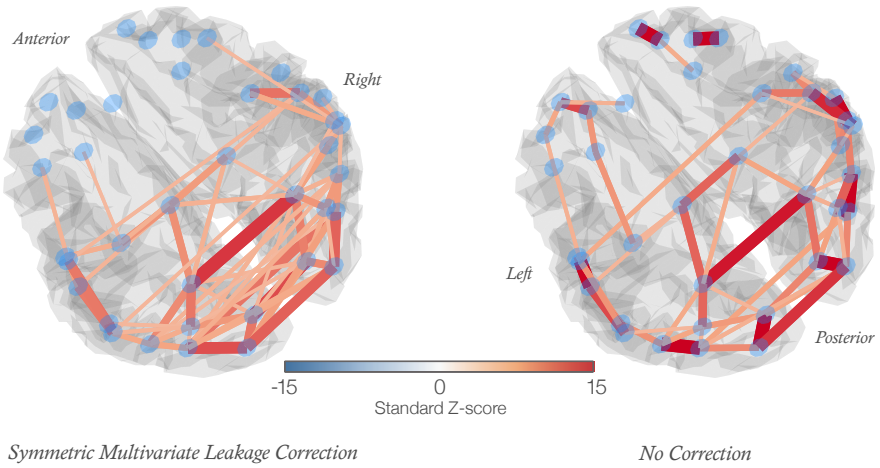


Figure 2.9: ALPHA-BAND RESTING-STATE NETWORK STRUCTURE SURVIVING THE 5% FALSE DISCOVERY RATE CORRECTION FOR MULTIPLE COMPARISONS. A comparison of network edges inferred for alpha-band (8–13 Hz) resting-state oscillations shows a much more densely connected visual network once the correction for source leakage is applied. Edges are shown as joining centres of mass of each ROI, with the colour scale and edge thicknesses indicating the group-average inference on regularised partial correlations between power envelopes of ROI time courses. Only edges above the 5% false discovery rate thresholds are shown ($z = 3.4$ for both the corrected and uncorrected networks).

timised regularisation strengths were of similar magnitude for the uncorrected and corrected data: $\lambda = 3.2 \pm 1.3 \times 10^{-2}$ and $\lambda = 2.0 \pm 1.1 \times 10^{-2}$ respectively for the alpha band; and $\lambda = 2.2 \pm 0.7 \times 10^{-2}$ and $\lambda = 1.6 \pm 0.9 \times 10^{-2}$ for the beta band.)

Estimated networks for functional activity within these two bands, both with and without the correction for source leakage between ROIs, are shown in figures 2.8, 2.9 and 2.10. The surface plots on the preceding page show direct connections to and from individual ROIs, with a 5% False Discovery Rate (FDR) threshold applied ($z = 3.4$ for the alpha band and $z = 3.6$ for the beta-band, in both corrected and uncorrected datasets). These reveal strong clusters of short-range connections that clearly relate to regions well known as participating in independent networks: in the beta-band, the sensorimotor and parietal regions;

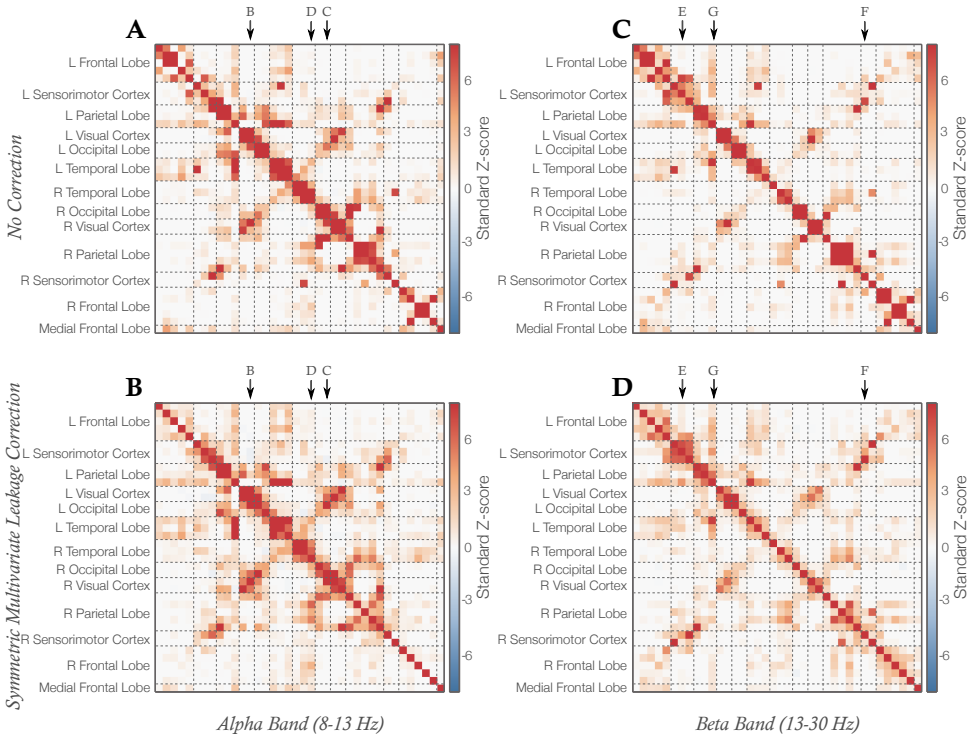


Figure 2.10: GROUP-AVERAGE RESTING-STATE NETWORK MATRICES DERIVED FROM EIGHT SUBJECTS, INFERRED WITH AND WITHOUT THE APPLICATION OF A MULTIVARIATE SPATIAL LEAKAGE CORRECTION BETWEEN ROIS. Alpha-band (8–13 HZ, left column) and beta-band (13–30 HZ, right column) functional connectivity network matrices are inferred from a group average, at each edge, of individual regularised partial correlation z-statistics, between 38 ROIS derived from an FMRI ICA-based parcellation. Network matrices are presented both with (B and D) and without (A and C) a symmetric multivariate correction for source leakage applied to the ROI time courses. Arrows above the matrices indicate the columns (i.e., seed ROIS) that are rendered as surface maps in figure 2.8.

in the alpha band, the visual network and local temporal lobe activity.

With the application of the symmetric source leakage correction between ROIS, the group-average networks, thresholded at a 5% FDR, showed an increased number of bilateral connections within the alpha-band visual network (although the net increase in un-thresholded cross-hemisphere edge-strengths within the 6 ROIS spanning the visual cortex was insignificant, median difference $\Delta z = 1.4$, 95% Confidence Interval (CI) [-0.4, 2.9] , $p = 0.055$, and there was no clear differ-

ence in overall edge-strength between these ROIs, $\Delta z = 1.2$, 95% CI [-4.8, 2.3]). (For these and all estimates of visual alpha-band and sensorimotor beta-band connectivity differences in this section, the median difference with 95% confidence intervals is given by the Hodges-Lehmann estimator, associated with the Wilcoxon signed-rank test, together with a p-value on the same test with no adjustment for multiple comparisons.) At the 5% FDR threshold, the beta-band sensorimotor network emerged more cleanly as an independent cluster, showing in contrast to the uncorrected networks, bilateral motor connectivity from both left and right seed parcels (figure 2.8E and figure 2.8F on page 57), and the removal of some more frontal regions from this motor system (more generally, the median increase in un-thresholded edge strengths between 5 ROIs covering the beta-band sensorimotor network was $\Delta z = 1.6$, 95% CI [0.6, 2.5], $p = 0.004$, with the purely bilateral connections showing an insignificant change, $\Delta z = 1.3$, 95% CI [-0.3, 3.1], $p = 0.063$). Direct connections between frontal and temporal lobes (alpha band, right hemisphere) and within the left fronto-parietal network (beta-band) also emerged at this threshold. By contrast, without the correction, direct connections between ROIs tend to be dominated by links to neighbouring regions, bilateral connections within the motor and visual networks were inhibited, and there were no edges linking ROIs in the frontal and parietal lobes. Some limited connectivity between neighbouring ROIs in the frontal lobe were present in the uncorrected data in both frequency bands; these were largely eliminated when the correction was applied.

The full connectomes of the thresholded network matrices for the alpha band are shown in figure 2.9 on page 58, with edges drawn between the centres of mass of each ROI.[7] Again, for the symmetrically-corrected connectome as

[7] Also available in a video, which makes the network matrices clearer. The video is the first

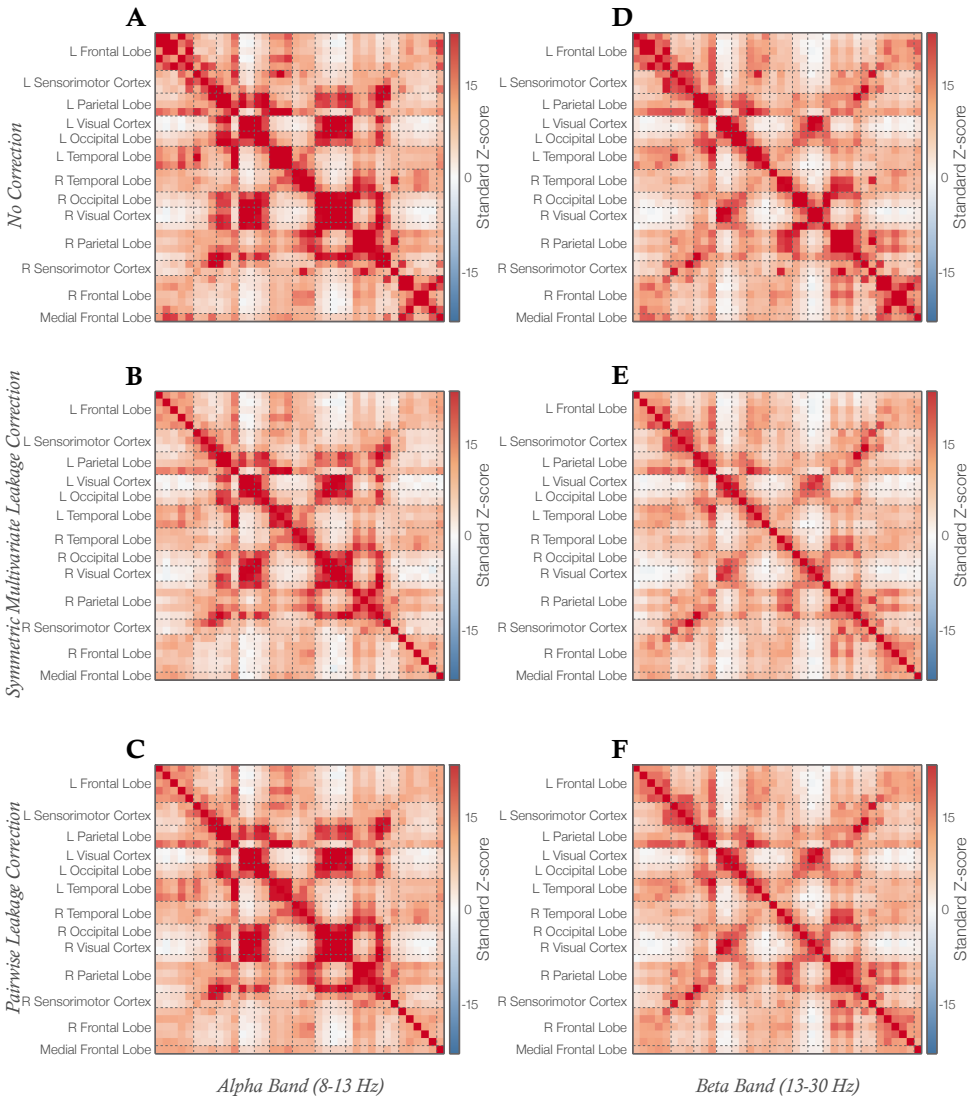


Figure 2.11: GROUP-AVERAGE RESTING-STATE CORRELATION MATRICES DERIVED FROM EIGHT SUBJECTS, INFERRED WITH AND WITHOUT THE APPLICATION OF A SYMMETRIC, MULTIVARIATE SPATIAL LEAKAGE CORRECTION BETWEEN ROIS, AND A PAIR-BY-PAIR LEAKAGE CORRECTION. Alpha-band (8–13 HZ, left column) and beta-band (13–30 HZ, right column) functional connectivity correlation matrices are inferred from a group-average average, at each edge, of individual correlation z-statistics, between 38 ROIS derived from an FMRI ICA-based parcellation. Correlation matrices are presented both with (B and E) and without (A and D) a symmetric multivariate correction for source leakage applied to the ROI time courses. Also presented (C and F) are correlation matrices for ROIS that have been orthogonalised in a pair-by-pair fashion before the computation of correlations.

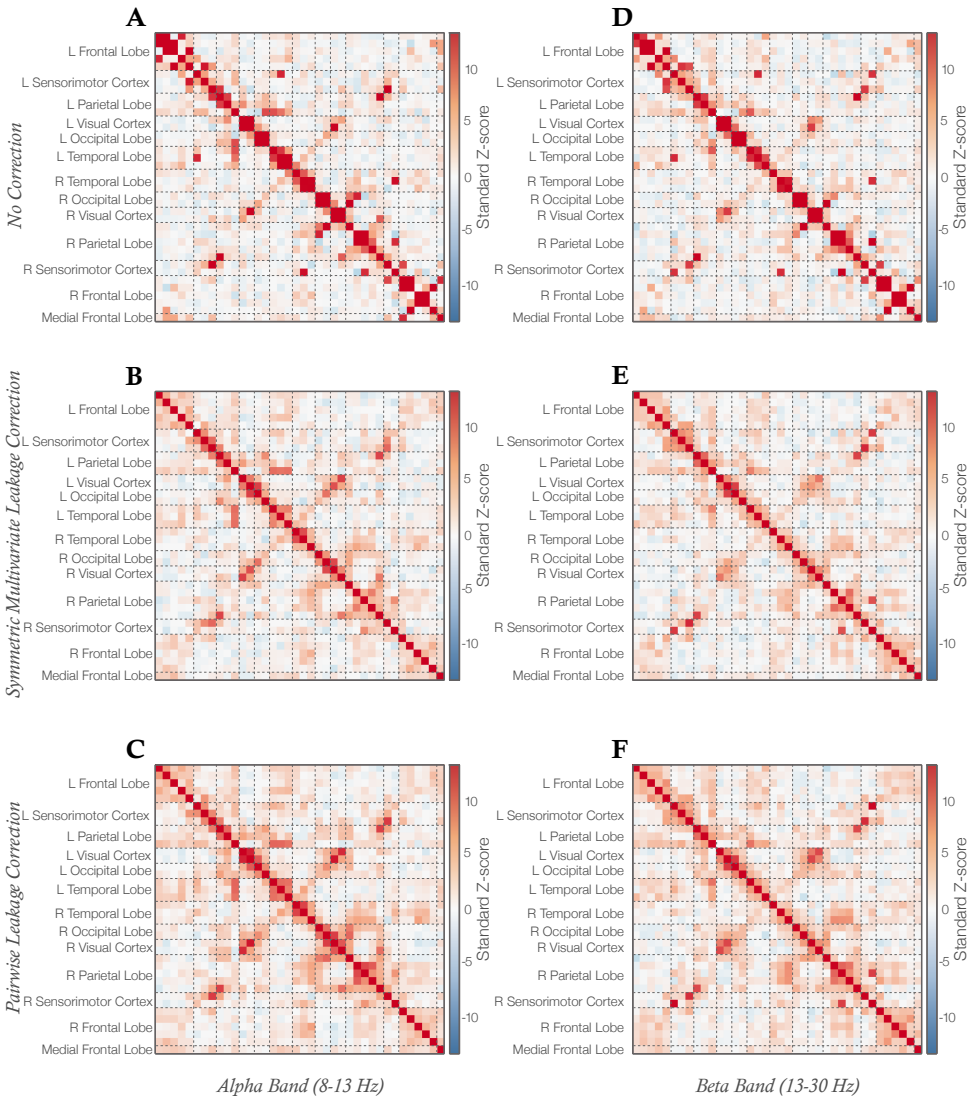


Figure 2.12: GROUP-AVERAGE RESTING-STATE UNREGULARISED PARTIAL CORRELATION MATRICES DERIVED FROM EIGHT SUBJECTS, INFERRED WITH AND WITHOUT THE APPLICATION OF A SYMMETRIC, MULTIVARIATE SPATIAL LEAKAGE CORRECTION BETWEEN ROIS, AND A PAIR-BY-PAIR LEAKAGE CORRECTION. Alpha-band (8–13 Hz, left column) and beta-band (13–30Hz, right column) functional connectivity (unregularised) partial correlation matrices are inferred from a group-average average, at each edge, of individual correlation z-statistics, between 38 ROIs derived from an fMRI ICA-based parcellation. Partial correlation matrices are presented both with (B and E) and without (A and D) a symmetric multivariate correction for source leakage applied to the ROI time courses. Also presented (C and F) are correlation matrices for ROIs that have been orthogonalised in a pair-by-pair fashion before the computation of partial correlations.

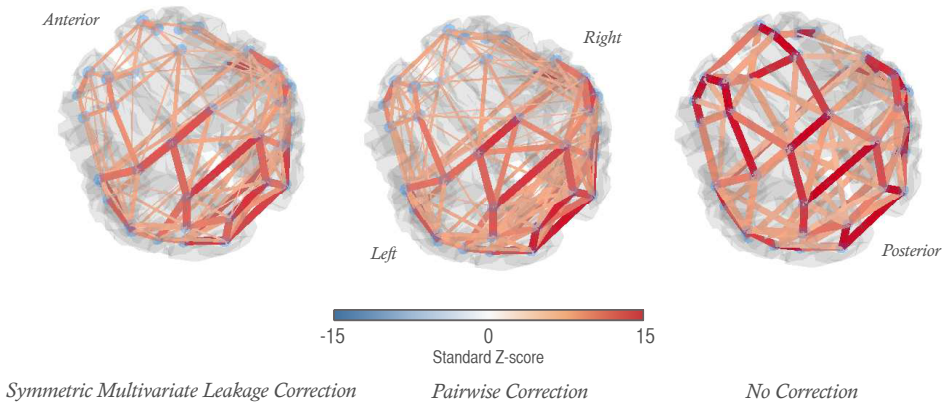


Figure 2.13: ALPHA-BAND (8-13 HZ) RESTING-STATE PARTIAL CORRELATIONS, WITH A COMPARISON BETWEEN ORTHOGONALISATION METHODS. Partial correlation, inferred without any regularisation to sparsity, yields denser connectomes (c.f..figure 2.9) and may be noisier than regularised estimates. In this analysis, there is little to differentiate the two leakage correction methods. Edges are shown as joining centres of mass of each ROI, with the colour scale and edge thicknesses indicating the group-average inference on regularised partial correlations between power envelopes of ROI time courses. Only edges above the 5% false discovery rate thresholds are shown ($z = 3.2$, 3.4 and 3.3 respectively for the symmetrically-corrected, pairwise-corrected and uncorrected networks).

opposed to the uncorrected, there is a significantly higher prevalence of bilateral connections (median difference in edge strength for a bilateral connection within all 38 ROIS $\Delta z = 0.11$, 95% CI [0.08, 0.16], $p < 10^{-3}$). The unthresholded network matrices (both as regularised partial correlations, and as full correlation matrices) are shown in figures 2.10 and 2.11.

Lastly, in figures 2.12, on the previous page, and 2.13, on this page, we present for comparison the unregularised partial correlations between all 38 ROIS, thresholded at a 5% FDR. Estimating direct network connections using partial correlations without regularisation is noisy; here we merely observe the similarity between partial correlations estimated after the pair-wise and

supplementary figure of the published article based on this chapter, and is freely-available at sciencedirect.com/science/article/pii/S1053811915002670/.

symmetric orthogonalisation processes. As in figure 2.11, there is little to differentiate the results after the application of these two different corrections for source leakage.

THE EFFECT OF ORTHOGONALISATION ON POWER ENVELOPES

To what extent does the application of a leakage-correction method alter the power envelopes of the time courses (the data which are passed on for the estimation of correlations between ROIS)? Figure 2.14 compares the down-sampled amplitude envelopes of a time course produced by the functional network of five connected and thirty-three unconnected nodes, used to generate the dipole amplitudes used in the simulations for figure 2.6, after the application of the symmetric, multivariate orthogonalisation of all 38 nodes and after a Gram-Schmidt orthogonalisation to one other, randomly-chosen, node time course. This allows

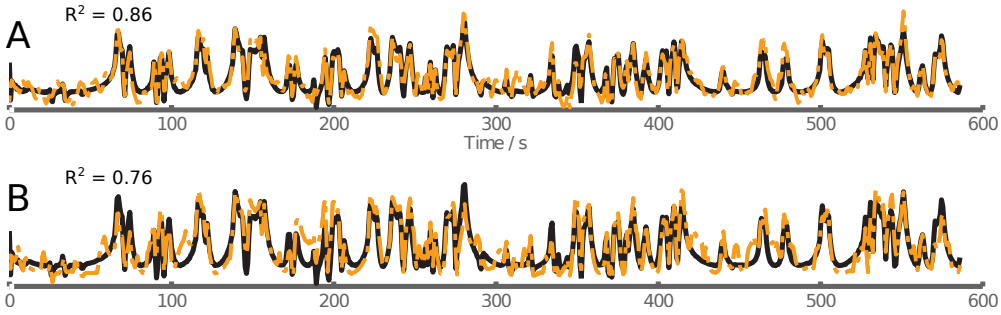


Figure 2.14: ILLUSTRATION OF THE EFFECT OF THE SYMMETRIC ORTHOGONALISATION PROCESS ON A TIME COURSE. An illustration of the effect of the symmetric orthogonalisation process and the pair-by-pair orthogonalisation on the envelopes of two time courses simulated from the functional network of 38 nodes, with 5 connected to each other and the remainder independent, used to generate the data for figure 2.6. In black is the amplitude envelope of the activity time course for one randomly-chosen unconnected node. In (A), in dashed orange, is the envelope of the same node after symmetrically orthogonalising all 38 nodes. In (B), in dashed orange, is the envelope of the same node after orthogonalising using the Gram-Schmidt process to one other, randomly chosen node. For this example, the symmetrically-corrected envelope is closer to the uncorrected envelope than is the Gram-Schimdt corrected version (Pearson's R^2 is 0.86 for the former and 0.76 for the latter).

for comparison of the effects of the orthogonalisation procedures in the absence of any leakage effects. The orthogonalised envelopes are inevitably altered to some extent, appearing as noisy versions of the uncorrected envelope—although the envelopes resulting from the pair-wise correction are more distorted than those after the symmetric, multivariate correction: the average fractional increase in displacement compared to the uncorrected signal envelope (measured as the norm of the difference between the two envelopes), of the pair-by-pair orthogonalisation process over the symmetric orthogonalisation, taken over all node pairings and for all 38 nodes of this simulation draw, is 20.2±0.6% (mean ± std. error).

2.5 Discussion

We have presented a method for removing the confounding effects of source leakage before performing a correlation-based network analysis between regions of interest in source-reconstructed MEG data. This approach corrects the time courses of each ROI using a symmetric, multivariate orthogonalisation step before inferring direct network edges between ROIs from the L_1-regularised partial correlations between the band-limited power envelopes of the corrected time courses. The orthogonalisation step is optimal in the sense that it minimally displaces the set of ROI time courses from their uncorrected forms.

This work extends previous leakage correction methods (Brookes *et al.*, 2012*b*; Hipp *et al.*, 2012) which, with one exception (Brookes *et al.*, 2014), were only appropriate for enabling correlation analyses between pairs of voxels in turn. Any multivariate analysis can now be performed on corrected ROI time courses—possibilities include multivariate auto-regressive models or hidden

Markov models (Baker *et al.*, 2014).

Our approach shares a limitation of these foundation methods, that the statistical testing for the presence of connectivity may be imperfect when the distributional assumptions in the non-parametric tests do not match the distribution of the data. Generating surrogate, null time courses using an auto-regressive model or a phase-randomisation technique produces normally-distributed data (albeit with the appropriate temporal smoothness). If the recorded data are substantially non-Gaussian in distribution, poor control of false positives can result. This was demonstrated by simulation in the appendix of Brookes *et al.* (2014).

More generally, our method is limited by the compartmentalisation of brain function into a set of ROIs with only a single time course each. On the one hand, using too few ROIs to capture the complexity of brain function risks a misrepresentation of the true network dynamics. Here, results were obtained using 38 ROIs from an fMRI group ICA, which is perhaps a coarse rendering of cortical network structure; future studies could explore the effectiveness of this technique using higher-dimensional parcellations. On the other hand, dividing into several ROIs part of the cortex that largely functions as a single coherent functional unit (given the spatial resolution and point spread function of the source reconstruction method), will cause those ROIs to be described by very similar time courses. This situation is problematic for any network inference that partials out nodes, including partial correlation, as the partialling of one node against another with a shared time course prevents the robust discovery of connectivity between their joint cortical unit and the rest of the brain. Under a full correlation analysis, the multivariate leakage correction proposed here could still create similar signal cancellation properties if two nodes are dominated by the same signal with zero-lag. As a result, whether or not a partial network

analysis is performed, using ROIs that are too small or overlapping should be avoided: for example, the application of this method to compute dense or voxelwise connectomes would be inappropriate. One approach for mitigating this behaviour is to select ROIs using data-driven decompositions such as ICA, which differentiate cortical areas based on functional specificity. In particular, ICA decompositions of the same MEG dataset may provide cortical parcellations that reflect the spatial profile and resolution of the particular MEG measuring systems and reconstruction algorithms in use.

We have discussed the use of regularised partial correlation matrices for inference on a simple undirected Gaussian network model between ROIs. In line with the conclusions of many previous authors (Duff *et al.*, 2013; Marrelec *et al.*, 2006; Smith *et al.*, 2011) it seems that partial correlation, whether with or without any imposed sparsity, is more effective at illuminating the true network structure than full correlation, which makes no attempt to discriminate between direct and indirect connections. Compare for example the network structures inferred by correlation and partial correlation methods in figure 2.5 on page 52, and the elevated false-positive rate in the entire top row of figure 2.6 relative to the middle row. Fitting a model that encourages network sparsity in the partial correlation matrix is known to reduce noise and improve the estimation of correlations (Duff *et al.*, 2013; Varoquaux & Craddock, 2013; Varoquaux *et al.*, 2010), and in our simulations significantly improves network detection, with strong suppression of unconnected edges to produce very low false positive rates.

Applying the multivariate leakage correction is effective, in simulation, at removing artefactual correlations between ROIs induced during the source reconstruction. In simulations of network dipoles at random locations through the

cortex, inference on network structure using the partial correlations between the power envelopes of corrected ROI time courses showed, on average, the expected control of false positives, whereas inference using uncorrected ROI time series had a significantly higher rate of false detections. Inspection of the network matrices (figure 2.5) suggests that these are driven by false connections between proximate ROIs.

When we performed our network analyses on a resting-state dataset, applying the symmetric leakage correction to ROIs yielded denser connections within the alpha-band visual and beta-band sensorimotor networks, including a higher incidence of cross-hemispheric connections within these areas and a reduction in connections between the motor cortex and neighbouring regions in the frontal lobe (see figure 2.8 on page 57). By contrast the networks inferred from the uncorrected ROIs are more dominated by local connections. Both of these observations have been previously noted in investigations of pair-wise source leakage corrections at the voxel level (Brookes *et al.*, 2012*a*; Hipp *et al.*, 2012; Maldjian *et al.*, 2014). Applying the correction also elucidated the detection, at a 5% false discovery rate threshold, of longer-range beta-band fronto-parietal and alpha-band fronto-temporal connections, which were less dominant than local connections in the uncorrected data. We expect these direct connections between frontal and parietal regions as the lateral fronto-parietal networks are common and stable independent components in ICA decompositions of both MEG and fMRI data. Direct connections linking all of the regions known to be involved in the default mode network, such as the medial prefrontal cortex and inferior parietal cortex, were not evident in this dataset, either before or after corrections for source leakage. (Some electrophysiological studies identify this network, but with varying degrees of completeness compared to fMRI results.

See, for example, Baker *et al.*, 2014; Brookes *et al.*, 2011*b*; Duncan *et al.*, 2013; Hillebrand *et al.*, 2012; Maldjian *et al.*, 2014; de Pasquale *et al.*, 2010 and Wens *et al.*, 2014.) In line with Hillebrand *et al.* (2012), we find an increased number of connections in regions associated with higher source power (occipital, parietal and temporal cortex in the alpha band; sensorimotor cortex in the beta-band). This may reflect the fact that corrected time courses for ROIs with a higher SNR (and therefore better source specificity) are less rotated from their original positions than those with lower SNR. Lastly, we find it interesting (and encouraging) that the removal of zero-lag correlations from the data has actually increased the number of inferred direct connections (visible in figure 2.9).

In our simulations, we also evaluated an extension of previous bi-variate voxel-wise corrections for source leakage to a pair-by-pair orthogonalisation process between ROIs. While it offered a reduction in the number of false positive edge detections over applying no correction at all, the control of the false positive rate was significantly poorer than desired for a leakage correction method (see figure 2.6). (Although when applied to the resting-state dataset, a comparison by eye of the correlations and *un*regularised partial correlation matrices showed little to differentiate the results of the symmetric and bivariate orthogonalisation processes. It does not follow that this would be true for all datasets.) We hypothesise that the poorer performance of the pair-by-pair orthogonalisation is simply reflecting the failure of a bivariate correction process to account for the entrainment of non-zero lag connectivity between a pair of regions into their respective local neighbourhoods. A mathematical example was provided in section 2.2; a nice discussion can be found in Palva & Palva (2012). We can briefly illustrate this point here. Consider two ROIs, A and B, which are functionally connected in the sense that their time courses exhibit

strong cross-correlations even at non-zero time-lags. The neighbours of B (the most proximate ROIs) may contain significant proportions of the signals in B due to the source leakage effect; this manifests itself as a strong correlation between B and the neighbours of B at purely zero lag. We may be able to prevent the false discovery of a connection between B and the neighbours of B by orthogonalising one with the respect to the other, which removes all zero-lag correlations. However, leakage from B into its neighbours can set up cross-correlations *at non-zero lag* between A and the neighbours of B. This *inherited connection* will not be accounted for by a pair-wise orthogonalisation procedure, but will by a multivariate process that considers the behaviour of all ROIs. Indeed, the *partial* correlations computed under the pairwise correction showed better false-positive control.

Note also that this pairwise approach is not only much more computationally demanding than the simple symmetrical orthogonalisation step we are proposing, but also that a method for finding regularised estimates of partial correlations has not yet been developed, is algorithmically more challenging and would create additional computational overheads, and that it is not straightforward to immediately apply any desired multivariate analysis method to the corrected time courses. We therefore recommend the more effective and more efficient, symmetric approach for removing correlations between ROIs induced by source leakage. The number of ROIs which can effectively be corrected is limited by the rank of the data: our method cannot be applied to voxel-level network analyses, for example. Asking questions about the connectivity of more ROIs than there are independent time series in the sensor data has dubious merit; but if these analyses are pursued, orthogonalisations between pairs of ROI time courses may be the only feasible approach for limiting the effects of source

leakage on amplitude correlation analyses.

Using a decomposition of the cortex into a set of ROIs to investigate functional network properties is common and profitable in fMRI. We believe that our approach for correcting ROI time courses for source leakage correlations will enable a larger range of multivariate ROI analyses to be performed in the MEG community.

3

How reliable are MEG resting-state connectivity metrics?

There are a wide range of network estimation methods available for MEG and EEG, and little in the way of existing guidance on which ones to employ. Chapter 2 augments this problem by introducing an additional method for connectivity analysis. This chapter investigates the extent to which many popular measures of stationary connectivity are suitable for use in resting-state MEG, and how the method of chapter 2 compares to these. We use as empirical criteria that network measures for individual subjects should be repeatable, and that group-average connectivity estimation shows good reproducibility. Using publicly-available data from the Human Connectome Project, we test the reliability of 12 network estimation techniques. We find that the impact of magnetic field spread or spatial leakage artefact is profound, creates a major confound for

many connectivity measures, and can artificially inflate measures of consistency. Among those robust to this effect, we find poor test-retest reliability in phase- or coherence-based metrics such as the phase lag index or the imaginary part of coherency. The most consistent methods for stationary connectivity estimation over all of the tests are simple amplitude envelope correlation and partial correlation measures—the methods introduced in chapter 2.

3.1 Introduction

A plethora of options is available to the researcher who embarks on a functional connectivity analysis using electrophysiological data. To assist in these decisions, much work has gone into establishing an agreed set of *good practices.* These might be summarised by four key points. Firstly, functional connectivity should be estimated between reconstructions of the cortical sources responsible for the sensor recordings (as network modelling on the data in the sensors themselves is hard to interpret; Schoffelen & Gross, 2009). Secondly, the relatively low rank of electrophysiological data places a limit on the number of independent network nodes that can be sensibly specified in a connectivity model (there is little merit in estimating a network matrix over a 2 mm source grid, for example). We can, however, attempt to describe as many of the dominant cortical sources as our data allow using sensible ROIs. Ideally, these regions should be specified using a *functional* parcellation: a decomposition of the data into regions whose boundaries reflect the functional specificity of cortex, as expressed in recordings of functional activity (Craddock *et al.*, 2012; Glasser *et al.*, 2016; Harrison *et al.*, 2015; Smith *et al.*, 2013; Yeo *et al.*, 2011). Thirdly, the spatial leakage of inferred cortical sources into their local neighbourhood can create the semblance of

connectivity between regions that in truth just share components of the same sensor signal (discussed extensively in chapter 2), and methods must be employed that ignore or remove this confound.[1] Lastly, statistical techniques used for inference on the functional networks, or for comparisons of connectivity between groups of subjects, must appropriately control for the large number of comparisons when statistical tests are performed on each network edge. In contrast to spatial activation data, there is no clear relationship between these concurrent tests (although see Zalesky *et al.* (2012, 2010) for hypothesis tests related to properties of network graph structures), and researchers often have to resort to traditional multiple-comparison corrections or combination-testing approaches (Winkler *et al.*, 2016*b*).

In contrast to the prescriptions of these emerging guidelines, the researcher is afforded a lot of flexibility in her choice of network modelling method. The available measures may assess spectral coherence, the presence of a consistent phase delay, or the amplitude coupling of oscillatory patterns. Some of these methods estimate the directionality of information flow between the network nodes, and some reveal solely the direct, and not indirect, connections.[2] Methods for achieving these latter two goals, for example by partialling out data (Marrelec *et al.*, 2006) or using temporal precedence as an indicator of causality (Baccala & Sameshima, 2001; Granger, 1969), are noisier to estimate, and the extent to which they exhibit reliable behaviour in MEG is poorly understood. Indeed, despite the wide range of choices for network modelling, the reliability and consistency of most common network measures have not been fully compared. As a result,

[1] Many authors continue to ignore spatial leakage confounds, despite the increasingly large literature discussing this issue (Brookes *et al.*, 2012*b*; Ewald *et al.*, 2011; Hipp *et al.*, 2012; Maldjian *et al.*, 2014; Nolte *et al.*, 2004; Palva & Palva, 2012; Stam *et al.*, 2007; Vinck *et al.*, 2011).

[2] Full correlation and partial correlation are good examples of methods that reveal all relationships (direct and indirect), and solely direct relationships, respectively.

many researchers may be picking analysis techniques without an understanding of the reproducibility of their results. This chapter seeks to provide guidance to the researcher who is setting out on a resting-state analysis as to which network estimation method to choose. Additionally, having proposed a new method in chapter 2 that enables leakage-corrected multivariate covariance modelling methods for MEG and EEG, it is worthwhile assessing its performance in relation to the existing range of methods in popular usage.

An objective analysis, then, of the quality of the most popular connectivity estimation methods is required. How can we assess and compare their performance in realistic situations, when there is no certain, 'true' network structure for the connectivity of real resting-state data? A key property that any network connectivity measure must satisfy is the repeatability and consistency in estimation of both the single-subject result and the group-average structure. Consistency of estimation is a measure of the noisiness of a network metric, given a set of resting-state scan protocols and source reconstruction methods. Indeed, the sensitivity of inference on the mean effect, by the reduction of noisy sources of variability, can be a good surrogate test in the absence of a known ground truth, even in the presence of smaller (and potentially informative) sources of variation in the network matrices (Smith *et al.*, 2005). Knowledge of estimation consistency is also clinically relevant, as a key requirement in the development of normative population distributions, and an understanding of the run-to-run variability of network estimation is a crucial first step when moving onward from group comparison studies to the classification or characterisation of new subjects.

In this chapter, we compare the consistency and reproducibility of whole-brain network estimation between 39 functionally-defined cortical ROIs both

for group-average networks, and within and between individual subjects' networks, for 12 commonly-employed network estimation metrics. We include in our analysis connectivity measures that are susceptible to spatial leakage confounds, not because they are necessarily to be recommended but because they are still used in the field. We will explore the extent to which the false connectivity induced by these confounds can artificially increase the consistency of network estimation. We concentrate on stationary connectivity, rather than characterising the evolution of networks over time in each subject. We perform our analysis on the resting-state MEG data provided by the HCP, for a number of reasons. The data are freely available, have passed through a stringent pre-processing pipeline, include a large number of young, healthy subjects (61 with resting scans in the MEG2 release), and include three repeated sessions of resting-state recordings for each participant.

As an alternative to the range of methods for connectivity estimation that are robust to spatial leakage artefacts, there are some source reconstruction approaches that are designed to minimise the cross-talk between specific cortical locations (for example, Dalal *et al.* (2006); Hauk & Stenroos (2014); Hui *et al.* (2010), discussed in chapter 2 on page 31). These methods require the multi-source reconstruction to have fewer dipoles than the rank of the sensor data: each ROI would have to be described by a single principle component of its lead fields or a representative voxel. The relative benefits of performing the dimensionality reduction down to the number of ROIs either before or after the source reconstruction (as we do in this chapter) remain to be explored; the reliability of networks inferred using these alternative methods has not been investigated here.

Other authors have essayed similar studies, but without the range of metrics

presented here, or without the focus on source-space connectivity. Astolfi *et al.* (2007) found Partial Directed Coherence (PDC; Baccala & Sameshima, 2001) to reliably infer connectivity in a Stroop task; Hardmeier *et al.* (2014) assessed the Phase Lag Index (PLI; Stam *et al.*, 2007) and weighted Phase Lag Index (WPLI; Vinck *et al.*, 2011) in the resting state, but only analysed sensor-level signals, not cortical sources; David *et al.* (2004) used neural mass models and MEG task data to compare a small number of methods (not including envelope correlations), finding non-linear estimators to be the most successful; Wang *et al.* (2014) undertook a large simulation study to assess accuracy of network estimation, using the outputs from neural mass models and simulated fMRI data—they found PDC and coherence estimated by Fourier transform to outperform non-linear estimators, but they did not look at envelope correlations nor at reproducibility in real datasets; and Jin *et al.* (2011) and Deuker *et al.* (2009) both focussed on the reproducibility of graph metrics, without studying the underlying stationary connectivity estimates.[3]

We present a detailed reproducibility analysis of commonly-used connectivity estimation methods, covering directed, undirected, partial and marginal techniques. The aim is not to propose one individual measure as 'the best,' but to determine, using real data, which of the measures in common use show sufficient reliability to provide a platform for the inference of resting-state connectivity.

[3] Shortly after the work in this chapter was published, Garcés *et al.* (2016) released a very similar study, although they used a much smaller dataset (both in terms of the number of subjects and the length of recordings), tested only three network measures, and used Brookes *et al.*'s approach to leakage correction. Garcés *et al.* came to very similar conclusions as we do here.

3.2 Methods

DATA COLLECTION AND PROCESSING

MEG data

Resting-state MEG recordings were collected on a whole-head Magnes 3600 scanner (4D Neuroimaging, San Diego, CA, USA) from 61 subjects as part of the HCP MEG2 release (Larson-Prior *et al.*, 2013; Van Essen *et al.*, 2013). Of these, 22 are monozygotic twins, 22 are dizygotic twins and the remainder are unrelated. All subjects are young (22-35 years of age) and healthy. (67 subjects were included in the release, but resting-state recordings that passed the quality control checks, which included tests for excessive SQUID jumps, and sensible power spectra and correlations between sensors, were not available from six.[4])

HCP data were acquired using protocols approved by the Washington University institutional review board. Informed consent was obtained from subjects. Anonymised data are publicly available from `ConnectomeDB`[5] (Hodge *et al.*, 2016). Certain parts of the dataset used in this study, such as the family structures of the subjects, are available subject to restricted data usage terms, requiring researchers to ensure that the anonymity of subjects is protected (Van Essen *et al.*, 2013).

Resting-state measurements were taken in three consecutive sessions for each subject with little or no break in between, for 6 minutes each. The data are provided pre-processed, after passing through a pipeline to remove any artefactual segments of time from the recordings, identify any recording channels that

[4] Unavailable subjects by ID: 104012, 125525, 151526, 182840, 200109, 500222.
[5] Accessible at `db.humanconnectome.org`.

are faulty, and to regress out artefacts which appear as independent components in an ICA decomposition with clear artefactual temporal signatures (such as eye-blinks or cardiac interference). Sensor-space data were down-sampled from 509 Hz to 300 Hz to facilitate processing, with the application of a zero-phase anti-aliasing filter.

Source localisation

MEG data from each session were source-reconstructed using a linearly constrained minimum variance scalar beamformer (Robinson & Vrba, 1999; Van Veen *et al.*, 1997; Woolrich *et al.*, 2011), which shows limited susceptibility to field spread (Schoffelen & Gross, 2009) and performs well even with correlated sources (Brookes *et al.*, 2007). Pre-computed single-shell source models are provided by the HCP at multiple resolutions, registered into the standard coordinate space of the Montreal Neurological Institute. Data were beamformed onto a 6mm grid using normalised lead fields and estimates of the data covariance computed separately in the 1–30 Hz and 30–48 Hz bands. Covariance estimation was regularised using PCA rank reduction. The rank was conservatively reduced by five more than the number of ICA components removed during preprocessing (2–20 components, with a mean of 7). Source estimates were normalised by the power of the projected sensor noise. Source-space data were filtered into delta (1–4 Hz), theta (4–8 Hz), alpha (8–13 Hz), beta (13–30 Hz), low beta (13-20 Hz), high beta (20-30 Hz) and low gamma (30–48 Hz) bands.

Parcellation

We were interested in comparing common network estimation methods for whole-brain, region-of-interest connectivity analysis. We re-used the parcella-

Table 3.1: INDEX OF ROI NUMBERS.

ROI number	ROI location
1, 2, 3, 4, 5	Left Frontal Lobe
6	Left Somatosensory Cortex
7, 8	Left Motor Cortex
9, 10, 11	Left Parietal Cortex
12, 13	Left Visual Cortex
14, 15	Left Occipital Lobe
16, 17, 18	Left Temporal Lobe
19, 20, 21	Right Temporal Lobe
22, 23	Right Occipital Lobe
24, 25	Right Visual Cortex
26, 27, 28, 29, 30	Right Parietal Lobe
31	Right Motor Cortex
32	Right Somatosensory Cortex
33, 34, 35, 36, 37	Right Frontal Lobe
38	Posterior Cingulate Cortex
39	Medial Frontal Cortex

tion employed in chapter 2, which comprises clusters identified from a resting-state ICA decomposition of fMRI data from the first 200 subjects of the HCP. An additional ROI was included, defined using the same process on the same data, to cover the posterior cingulate, yielding a total of 39 ROIs, described in table 3.1. A single time course was constructed to represent each node as the first principal component, after weighting the PCA over voxels by the strength of the ICA spatial map. This analysis yielded 39 time courses for each frequency band and session.

Removing contributions from source leakage

We compensate for spatial leakage confounds using the symmetric orthogonalisation procedure introduced in chapter 2 to remove all shared signal at zero lag

between the network nodes. It removes any zero-lag signal overlaps that could be attributed to spatial leakage effects, before performing a multi-variate network analysis. The approach was designed to correct correlation-based analyses on power envelopes for artefactual connections caused by the MEG point-spread function; in this work we test all metrics with and without this correction to understand the extent to which spatial leakage engenders artificial consistency in network estimation.

In the main text, we only apply this leakage correction to metrics that are sensitive to source leakage and do not have an equivalent corrected form. For completeness, results with this correction applied to all metrics are available in the supplementary information of Colclough *et al.* (2016*e*).[6]

The orthogonalisation was applied in the time windows over which the network metrics were computed (O'Neill *et al.*, 2015): for the correlation, spectral and autoregressive model measures, this was the entire time series; for the phase-based metrics, it was applied within each epoch.

CONNECTIVITY ESTIMATION

Connectivity between each region pair was subsequently estimated using the connectivity metrics described below. Generally these metrics fall into marginal (full) and partial network estimation measures, with directed and undirected edge orientation (see also table 3.2 on the next page).

[6] Freely available at www.sciencedirect.com/science/article/pii/S1053811916301914.

Table 3.2: CATEGORISATION OF THE TESTED NETWORK METRICS. We can separate the connectivity estimation methods by the inference process; whether or not only direct network edges are found (partial methods) or direct and indirect connections (marginal methods); whether or not directionality is ascribed to each edge; and whether or not the method is robust against spatial leakage artefacts.

Abbreviation	Connectivity Metric	Type	Direct Associations	Causal Relations	Leakage-Corrected
AEC	Amplitude Envelope Correlation	Amplitude coupling	Marginal	Undirected	Yes, with orthogonalisation
PAEC	Amplitude Envelope Partial Correlation	Amplitude coupling	Partial	Undirected	Yes, with orthogonalisation
COH	Absolute Coherence	Spectral coherence	Marginal	Undirected	No
IMC	Imaginary Coherency	Spectral coherence	Marginal	Undirected	Yes
PCOH	Partial Coherence	Spectral coherence	Partial	Undirected	No
IMPC	Imaginary partial coherency	Spectral coherence	Partial	Undirected	Yes
PLV	Phase-Locking Value	Phase estimation	Marginal	Undirected	No
PLI	Phase Lag Index	Phase estimation	Marginal	Undirected	Yes
WPLI	weighted Phase Lag Index	Phase estimation	Marginal	Undirected	Yes
PSI	Phase Slope Index	Phase estimation	Marginal	Directed	Yes
MI	Mutual Information between phases	Phase estimation	Marginal	Undirected	No
PDC	Partial Directed Coherence	Auto-regressive modelling	Partial	Directed	Yes[1]

[1] PDC may be sensitive to magnetic field spread, see discussion in Methods section 3.2 on page 86.

Phase estimation methods

Phase Lag Index, weighted Phase Lag Index and Phase Locking Value PLI, WPLI and PLV are all measures for assessing phase synchrony between two signals in a particular frequency band, formed from estimates of the instantaneous phase of the signal. In the task and resting-state literature, these measures are commonly computed in trials or artefact-free epochs of 2 s to 25 s, taking average values over these epochs (Hardmeier *et al.*, 2014; Lachaux *et al.*, 2000; Stam *et al.*, 2009, 2007; Tewarie *et al.*, 2014; Vinck *et al.*, 2011). We used window lengths of both 2 and 10 seconds, with little or no change in the results: we therefore present data based on 10 second epochs (results from 2 s epochs are available in the supplementary information of Colclough *et al.*, 2016*e*).

The analytic signal $h(t)$ is derived from a measured signal $s(t)$ as

$$h = s + i\hat{s} = |z(t)|e^{i\phi(t)}, \tag{3.1}$$

where ϕ is the instantaneous phase and

$$\hat{s}(t) = \text{p. v.} \int_{-\infty}^{\infty} d\tau \, \frac{s(\tau)}{t-\tau} \tag{3.2}$$

is the Hilbert transform of signal s within each epoch (p. v. indicates the Cauchy principal value).

The phase-locking value (Lachaux *et al.*, 1999) between two signals is given by

$$PLV = \left| \left\langle e^{i\Delta\phi(t)} \right\rangle \right|, \tag{3.3}$$

where $\langle\dot{}\rangle$ indicates the expectation operator, over time and epochs.

The phase-lag index (Stam *et al.*, 2007) quantifies the asymmetry of the phase difference, rendering it insensitive to shared signals at zero phase lag (see section 3.2 above),

$$PLI = |\langle \text{sign} \sin \Delta\phi \rangle| , \tag{3.4}$$

while the weighted PLI (Vinck *et al.*, 2011) attempts to further weight the metric away from zero-lag contributions,

$$wPLI = \frac{|\langle z_1 z_2 \sin \Delta\phi \rangle|}{\langle |z_1 z_2 \sin \Delta\phi| \rangle} . \tag{3.5}$$

Mutual information of the phase The mutual information shared in the instantaneous phases of two signals is a non-linear measure of their coupling, first proposed by Paluš (1997)

$$I = \int_{-\pi}^{\pi} d\phi_1 d\phi_2 \, p(\phi_1, \phi_2) \log \frac{p(\phi_1, \phi_2)}{p(\phi_1) \, p(\phi_2)} , \tag{3.6}$$

where $p(\phi_1, \phi_2)$ is the joint probability of ϕ_1 and ϕ_2, and $p(\phi)$ the marginal probability of ϕ. As suggested in Wilmer *et al.* (2012), we normalise the mutual information by the joint entropy, E, to get a normalised mutual information estimand ι between 0 and 1,

$$\iota = \frac{I}{E}, \quad E = -\int_{-\pi}^{\pi} d\phi_1 d\phi_2 \, p(\phi_1, \phi_2) \log p(\phi_1, \phi_2) . \tag{3.7}$$

We used Peng *et al.*'s Matlab toolbox for computing the mutual information (Peng *et al.*, 2005).[7]

[7] The `mutual information computation` toolbox is available from `uk.mathworks.com/matlabcentral/fileexchange/14888-mutual-information-computation/`.

Spectral estimation methods

Coherence and Imaginary Coherence The coherency of two signals is their normalised cross-spectral density,

$$\tilde{C}_{ij}(f) = \frac{\tilde{S}_{ij}(f)}{\sqrt{\tilde{S}_{ii}(f)\,\tilde{S}_{jj}(f)}},\ \tilde{S}_{ij} = \left\langle \tilde{s}_i(f)\,\tilde{s}_j^{\star}(f) \right\rangle, \tag{3.8}$$

where the star indicates the complex conjugate and $\tilde{s}$ the Fourier representation of s. This provides a measure of the extent to which they share a constant relative phase. We estimated the cross-spectral density of each pair of ROI time courses, for each session and frequency band, using Welch's method, with windows of length 10 seconds and a Hamming windowing function. We repeated this analysis as for the phase-based metrics, using windows of 2 seconds, again finding little or no change in the results.[8] We estimated the coherency for the complete time series using these cross-spectra. Coherence, the absolute value of coherency, was averaged over the frequency band of interest.

The imaginary part of coherency (Nolte *et al.*, 2004) is insensitive to phase synchronisation with no phase difference, and thus to the effects of signal leakage. Again, we computed coherency as above, and averaged the imaginary part of coherency over the frequency band of interest.

Partial Coherence and Imaginary Partial Coherence The partial coherency of two signals is the normalised inverse of their cross-spectral density,

$$\tilde{PC}_{ij}(f) = \frac{\tilde{S}_{ij}^{-1}(f)}{\sqrt{\tilde{S}_{ii}^{-1}(f)\,\tilde{S}_{jj}^{-1}(f)}}. \tag{3.9}$$

[8] Again, these results can be found in the supplementary information of Colclough *et al.* (2016*e*).

and is to coherency as partial correlation is to correlation: two signals are partially coherent if they retain coherency conditional on the spectra of all other signals in the sample. Partial coherence and partial imaginary coherency are obtained from the magnitude and imaginary part of partial coherency, respectively.

Phase Slope Index The Phase Slope Index (PSI; Nolte *et al.*, 2008) is a directed phase synchronisation index which weights information from different frequencies. It is insensitive to leakage effects, and can handle non-linearities in the phase spectrum. It is defined as

$$PSI_{ij} = \mathfrak{I}\left(\sum_{f} \tilde{C}^{*}_{ij}(f)\,\tilde{C}_{ij}(f+\delta f)\right), \tag{3.10}$$

where $\mathfrak{I}$ indicates the imaginary part and the summation is over frequencies f within the band and δf indicates the frequency resolution of the Fourier transform. We calculated PSI in the same epochs as PLI and PLV, averaging the results over the windows.

Auto-regressive models

Partial Directed Coherence Partial Directed Coherence (PDC) is a directed measure of signal transfer from one site to another (Baccala & Sameshima, 2001). Modelling the source-space signals X as white noise passed through a complex transfer function H (normally inferred by fitting an Auto-Regressive (AR) model),

$$\tilde{X}(f) = H(f)\tilde{E}(f), \tag{3.11}$$

where $\tilde{A}$ denotes the Fourier representation of time-domain signal A, the partial directed coherence is given by

$$PDC_{ij}(f) = \frac{H_{ij}^{-1}(f)}{\sqrt{\sum_l H_{lj}^{-1^*} H_{lj}^{-1}}} . \tag{3.12}$$

PDC is often regarded as insensitive to source leakage, and has been defended as such (Kaminski & Blinowska, 2014), but other authors have thrown doubt on this fact (Leistritz *et al.*, 2013; Omidvarnia *et al.*, 2014).

To compute this metric, we fitted an AR model to the source-reconstructed signals in each ROI, for each band, by first downsampling to 2.5 times the Nyquist frequency of the band of interest, then fitting an AR model with 69 lags and computing the PDC using A. Schlögl's time series analysis toolbox.[9] For the beta and gamma bands, this approach produced poor model fits and PDC estimations, but an AR model with 201 lags applied to the data in each of these bands, maintaining the sample rate at 300 Hz, produced good results.

Amplitude coupling methods

Amplitude envelope correlation and partial correlation We estimated the amplitude coupling between ROIs using linear correlations and partial correlations of the envelopes of the band-pass filtered signals (Brookes *et al.*, 2012*b*; Hipp *et al.*, 2012). Power envelopes were computed as the magnitude of the analytic signal, h, formed from the Hilbert transform of the ROI time courses according to equation (3.2), low-pass filtered at 1 Hz and down-sampled to 0.5 Hz (Luckhoo *et al.*, 2012). We then took the linear correlation between the logarithm of ROI power envelopes, and computed partial correlations by inverting the covariance

[9] Available from pub.ist.ac.at/~schloegl/matlab/tsa/.

matrix of the power envelopes, and regularising using the DP-GLASSO (Mazumder & Hastie, 2012*b*). The regularisation parameter was chosen separately for each session and frequency band as that which maximised the corrected Akaike information criterion under 5-fold cross validation on the data within that session.

RELIABILITY AND CONSISTENCY OF NETWORK MEASURES

We assessed the network measures under test in three ways: (i) in their ability to produce consistent estimates of the population's connectivity patterns, (ii) to show consistent connectivity patterns between separate recording sessions of the same individuals, and (iii) similar connectivity between different subjects. Tests (i) and (iii) are closely related, being driven by the same sources of variation. After computing a network matrix for each recording session, estimates of group connectivity were made by taking the mean value at each network connection, over all sessions and subjects.

For the first test, we assessed the reliability of these group-average estimates by randomly dividing the dataset in half 100 times (keeping all sessions for each subject within a single half), computing group connectivity matrices with each network measure for each half, and correlating the edge strengths of these group matrices between each partition.

For the second, reliability of network estimation within-subject was tested by correlating the edge strengths between network matrices generated from each of the three scans for each subject.

For the third, the consistency of networks estimated from different subjects was tested by correlating edge strengths between network matrices computed from different subjects.

We further investigated the relevance of individual edges to the creation of highly repeatable group-average metrics. On each edge, we calculated the contribution towards the correlation between the networks inferred from each half of the dataset. Thus, if the set of network edges inferred from one half of the data are denoted x and the other set y, the correlation between the two is

$$\begin{aligned}\rho &= \frac{\sum_i \left(x_i - \langle x \rangle\right)\left(y_i - \langle y \rangle\right)}{\sqrt{\sum_i \left(x_i - \langle x \rangle\right)^2}\sqrt{\sum_i \left(y_i - \langle y \rangle\right)^2}} \\ &= \sum_i z_i^x z_i^y \,, \end{aligned} \tag{3.13}$$

for $z^x = \frac{(x_i - \langle x \rangle)}{\sqrt{\sum_i (x_i - \langle x \rangle)^2}}$. The relative contributions from each edge, i, to our measure of reliability, ρ, are then given by the element-wise products $z_i^x z_i^y / \sum_i z_i^x z_i^y$.

Network matrix correlation scores were converted to z-values using Fisher's transformation.[10] This is a monotonic transformation which produces a clearer separation of results when correlations approach unity.

3.3 Results

We focus on presenting results just from the alpha band in this section, because the differences between metrics were smallest for this band, presumably because it exhibits the largest signal to noise ratio. Other bands show exaggerated versions of the plots here, but follow the same patterns. Figures presenting these extended results can be found in the supplementary information of Colclough *et al.* (2016*e*).[11]

[10] $Z(\rho) = \operatorname{artanh}(\rho)$, see also page 49.
[11] Freely-available at sciencedirect.com/science/article/pii/S1053811916301914/.

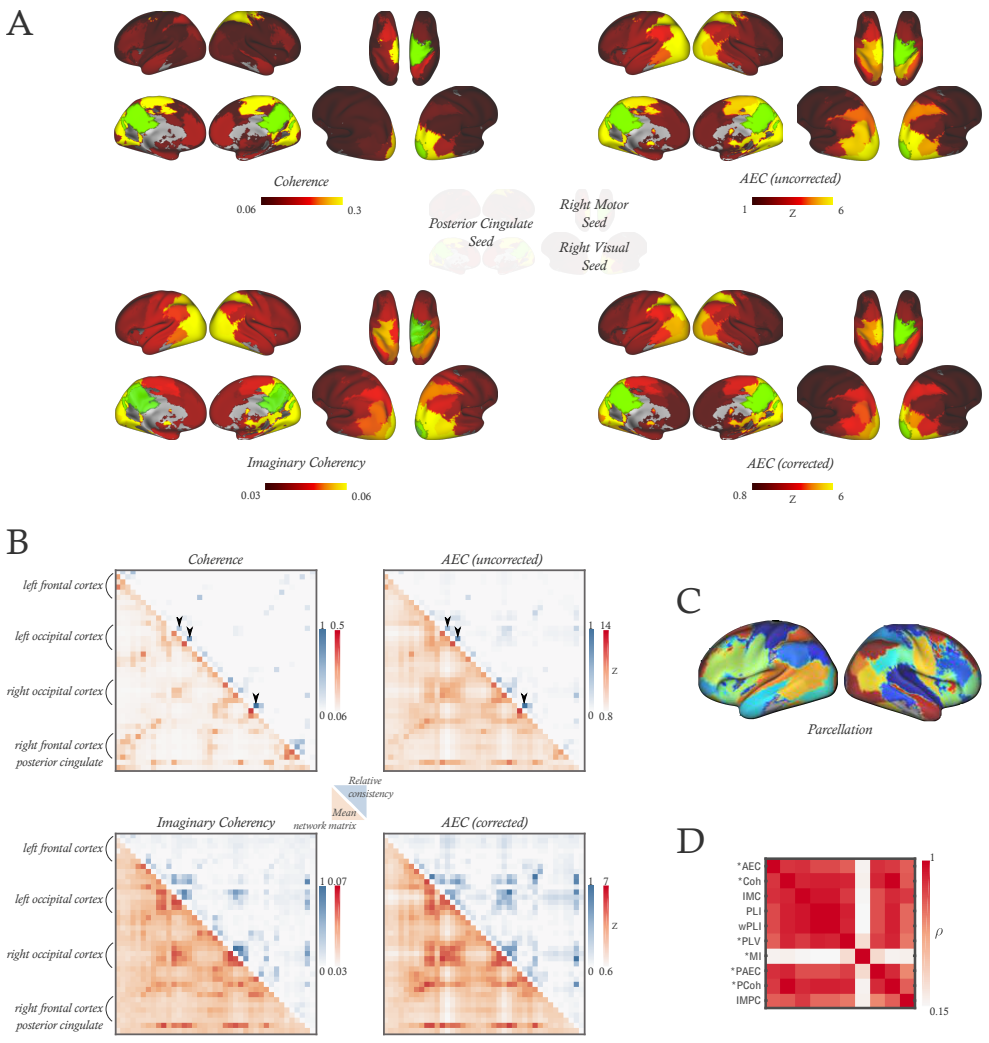


Figure 3.1: MEAN ALPHA-BAND NETWORK MATRICES AT GROUP LEVEL. (A) Views of single rows from the network matrices in (B). Left, connection strength to the posterior cingulate cortex (seed ROI number 38, shown in green). Upper right, connection strengths to the right motor cortex (seed ROI number 24). Lower right, connection strengths to right occipital cortex (seed ROI number 24). (B) Mean network matrices, and the relative edge-level consistency, for a subset of four of the metrics under test: the absolute value of coherency between ROIs, the imaginary part of coherency, and the Amplitude Envelope Correlation (AEC), both with and without a multivariate correction for source leakage. Lower triangles, in red, show the mean network matrix over all 183 sessions in the dataset. The colourbar to the side indicates the scale. Correlations have been converted to z-values before averaging. The upper triangles, in blue, indicate the relative contribution of each edge to the overall consistency of group-average network estimation, computed with equation 3.13 (c.f..figure 3.3A). The two network matrices which are susceptible to source leakage (AEC and coherence, top row) are dominated, even at group level, by a few strong edges (highlighted with arrows), and these drive the high consistency of network estimation. (C) The parcellation of 39 fMRI-derived ROIs used for this connectivity analysis. (D) Similarity between group-average network matrices. The heat map shows the high correlations between network matrices inferred at the group-average using the undirected network measures that are either immune to source leakage bias, or have been corrected for source leakage bias (starred). Expanded acronyms can be found in table 3.2 on page 82 or in the glossary on page 213.

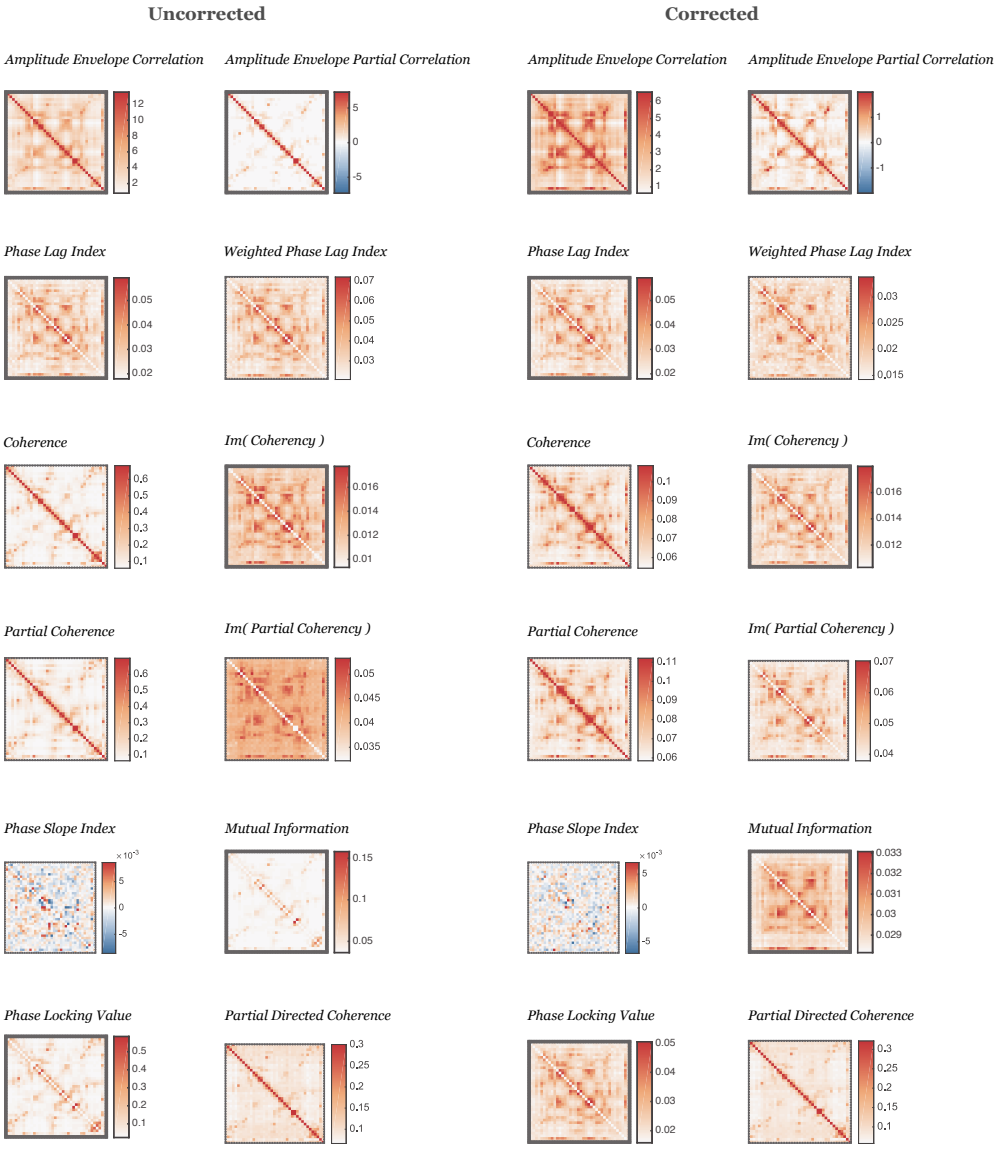


Figure 3.2: MEAN ALPHA-BAND NETWORK MATRICES AT THE GROUP LEVEL inferred for all 183 sessions in the HCP dataset, for all metrics under test.

GROUP CONNECTIVITY STRUCTURE

We find that the group-average connectivity matrices derived from the parcellation of 39 ROIs all exhibit the basic structure of connectivity expected in the resting-state, with segregated sensory networks. The posterior cingulate cortex is highly connected through the parietal and occipital lobes, as would be expected for part of the default mode network (although, as in Maldjian *et al.* (2014), with low connectivity to the anterior cingulate). In figure 3.1A–B on page 90, we illustrate alpha-band connectivity for three of the most common network analysis methods (amplitude envelope correlation, with and without a correction for source leakage; coherence and the imaginary part of coherency). Renderings of these networks in 3D are available in the supplementary information of Colclough *et al.* (2016*e*), and group-average network matrices for all methods under test are presented in figure 3.2 on the previous page. A listing of the locations of the 39 ROIs is given in table 3.1 on page 80.

General features of these network matrices are very similar over all network measures of a given type (partial, marginal, directed). We assessed the extent to which these different network measures are capturing similar information by correlating the mean edge strengths at group level computed with different metrics (shown for undirected metrics immune to, or corrected for, spatial leakage in figure 3.1D). Three-quarters of these between-metric comparisons were correlated with $\rho \geq 0.7$.

CONSISTENCY OF GROUP INFERENCE

We investigated the consistency of group-average inference by repeated bootstrapped division of the entire dataset into two equal halves, and comparing

the networks inferred from each half. Figure 3.3A on the next page gives the distributions of correlations between each of the two dataset halves (about 90 sessions in each), for all metrics studied here, in the alpha band. Those metrics that are affected by spatial leakage effects—amplitude correlation and partial correlation, phase locking value, coherence, mutual information—show strong consistency in group level estimation (with correlations between the two halves around 0.99). Of those metrics designed to be immune to spatial leakage effects, most achieve between-group correlations of $\rho = 0.9$, but solely in the alpha band; PLI, WPLI, IMC and PSI fare poorly ($\rho < 0.9$) in all other bands.[12] Crucially, applying an orthogonalisation correction to remove the impact of spatial leakage effects reduces group consistency for all methods, but uniquely the correlations and partial correlations of corrected power envelopes retain $\rho > 0.95$. Among the directed measures, PDC performs the best (median $\rho = 0.94$).

Among those metrics that are affected by spatial leakage, the same small subset of network edges dominated the high levels of correlation between the split halves. Figure 3.1B on page 90 shows as heat maps the edge-wise contribution to the correlations between network matrices inferred from each half of the data, for networks inferred using coherence, imaginary coherency, correlation and leakage-corrected correlation. When spatial leakage between regions remains uncorrected (top half of figure 3.1B), a small subset of edges close to the diagonal show a very strong contribution to the consistency of group network matrix estimation, and this subset becomes much less important once the leakage-induced connections are accounted for.

[12] Again, for the sake of brevity, these data are freely available online in the supporting information of Colclough *et al.* (2016*e*).

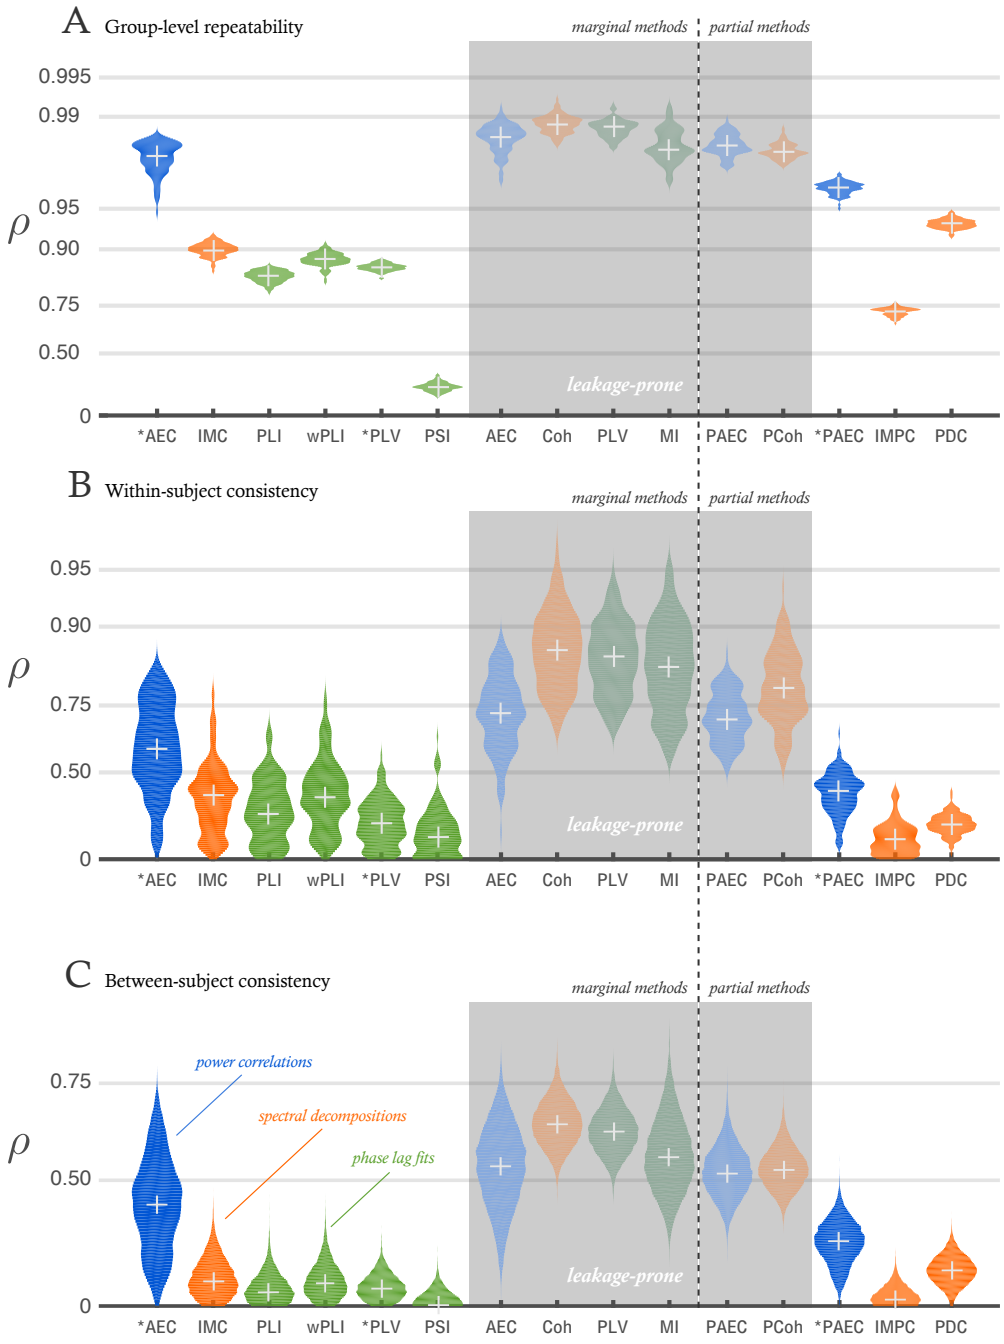


Figure 3.3: CONSISTENCY OF NETWORK MATRIX ESTIMATION. (A) *Stability of group-average inference.* Correlations between network matrices inferred from separate halves of the HCP dataset, using resting-state recordings in the alpha band. The dataset was randomly partitioned in half 100 times, and the distribution of correlations between the network edge strengths in each half was produced over the bootstrapped samples. (B) *Within-subject consistency of network inference.* Correlations between alpha-band network matrices inferred from each of three resting-state sessions from 61 subjects of the HCP dataset. (C) *Between-subject consistency of network inference.* Correlations between alpha-band network matrices inferred from 61 subjects of the HCP dataset. In all cases, the violin plots show smoothed histograms of these distributions, the median marked by a white cross. Correlations have been converted to z-values using Fisher's transformation, but labelled with the original rho value, for clearer display. Darker distributions with background shading identify metrics which are known to suffer from source-leakage confounds. Metrics with a star indicate that a multivariate source leakage correction has been applied before computation of the network measure. Expanded acronyms can be found in table 3.2 on page 82 or in the glossary on page 213.

RELIABILITY OF SINGLE-SUBJECT INFERENCE

The extent to which repeated measurements of resting-state activity within the same subject can produce similar network matrices can be a good test of the reliability of a network measure. We compare correlations of network edge strengths between three consecutive resting-state recordings of each subject, for each of our tested network measures, in figure 3.3B on the previous page. Again, we find high levels of repeatability among those metrics affected by spatial leakage (median correlations for Amplitude Envelope Correlation (AEC), partial correlation, PLV, coherence and Mutual Information (MI) lie between 0.7 and 0.9). This repeatability is destroyed by the orthogonalisation procedure for MI, coherence and PLV. Network measures designed to be immune to zero-lag correlation effects show large amounts of variability over subjects in the extent to which repeated sessions can reproduce the same network structure: correlations vary between 0 and 0.7, with the medians for PLI, WPLI, and the imaginary part of coherency lying around $\rho = 0.3$. Of the measures immune to, or corrected for, spatial leakage, AEC alone (and to some extent, partial correlation), when applied to leakage-corrected data, retains moderate amounts of between session repeatability, although the distributions are broad (the central 95% of within-subject network edge correlations for amplitude correlations between leakage-corrected power envelopes is 0.1–0.8).

CONSISTENCY OF NETWORK MATRICES BETWEEN SUBJECTS

Lastly, the consistency of network matrices inferred from different subjects speaks both to the amount of between-subject variability present in healthy datasets, and to the extent to which the population average of a network measure

is reflected in single-subject measurements. In our assessment of the HCP dataset, we find that network metrics sensitive to artificial correlations induced by spatial leakage retain high consistencies between subjects. Of those which ignore zero-lag effects, or are corrected for leakage by orthogonalisation, AEC (median $\rho = 0.4$) shows the largest amount of between-subject correlation, with the noisier partial correlation measure (median $\rho = 0.3$) a not-too-close second.

3.4 Discussion

We have assessed a range of network measures commonly used to evaluate connectivity in electrophysiological recordings. We were interested in understanding which network measures, when applied to resting-state recordings, are able (i) to produce repeatable inference at group level, and which show consistent patterns of connectivity both (ii) between different healthy subjects, and (iii) on repeated measures of the same subject.

We find that, at a very broad level, the patterning of connectivity described by different network measures is very similar. While this is consistent with the idea that multiple methods of cortical communication are simultaneously employed (for example, co-incident phase synchronisation and coupling of oscillatory power), as reported for shorter-range connections in Mehrkanoon *et al.* (2014), these network measures may simply be picking up different characteristics of the information transfer in the data. Alternatively, the different measures may not be independent estimators of different coupling mechanisms—the dependencies between phase-based metrics and cross-correlation have been discussed by Aydore *et al.* (2013). However, it is also clear that we frequently find strong connections, in all or most network metrics, between ROIs with strong

(high power) signals. Thus, we find dense connections both to and within visual cortex for the alpha band, and strong motor connectivity in beta (figure 3.1A on page 90). None of these observations are unexpected, but a signal strength bias should be recognised as a potential source of similarity between the network models under test here.

The confounds introduced into network estimation due to the spatial leakage of inferred cortical sources, a symptom of the broadness of the EEG and MEG point spread functions, appear highly repeatable over scans and between subjects. We observe the highest levels of repeatability of group-average inference, and the highest consistency of network matrix estimation from single sessions, for those metrics known to be affected by zero-lag phase coupling. Versions of these metrics which regress out the shared signal (leakage-corrected envelope correlation and partial correlation), or ignore coupling with zero phase difference (PLI, PSI, imaginary part of coherency), all exhibit a large relative drop in consistency. The influence of spatial leakage is profound; clearly visible in the inferred group network matrices when corrected and uncorrected metrics are compared (top of figure 3.1B on page 90: compare power envelope correlation with and without orthogonalisation, or coherence and the imaginary part of coherency, for example). This confound generates sufficiently many consistent, strong, spurious network connections over repeated measurement sessions, with similarity over different subjects, that it artificially inflates measures of repeatability between inferred networks. We do not, therefore, recommend the use of any network measure which cannot obviate the spurious connections engendered by spatial leakage effects.

Not all zero-lag connectivity is necessarily spurious. There is a body of experimental and theoretical work that describes the presence and potential

mechanisms for zero-lag synchronisation of neural oscillations (Gollo *et al.*, 2014; Roelfsema *et al.*, 1997). Methods for signal reconstruction that can suppress contributions from other cortical sources (Dalal *et al.*, 2006; Hauk & Stenroos, 2014; Hui *et al.*, 2010) may allow the detection of these effects. However, using standard beamformers or minimum-norm source estimates, it is not possible to disambiguate spatial leakage effects from zero-lag synchrony of biological origin (although see Wens *et al.* (2015) for recent work constructing models of field spread for pair-wise connectivity estimation). Because the impact of the spurious connections is considerable (see above), we (and others, Hipp *et al.*, 2012; Palva & Palva, 2012; Schoffelen & Gross, 2009) argue that interpretable connectivity estimation is only possible when zero-lag connections are removed or otherwise ignored.

The most consistent network measure that is corrected for the biases induced by spatial leakage artefacts, and which we can recommend for use, is the correlation between band-limited power envelopes, after a symmetric orthogonalisation correction to remove shared zero-lag signals (as set out in chapter 2). It is considered good practice (Smith *et al.*, 2013) to use network measures which elucidate only direct connections, if the resulting network matrices are desired to be physiologically interpretable, rather than simply used for discriminative purposes. These partial methods tend to be noisier to estimate, even with regularisation to encode underlying assumptions of sparsity in the connectivity pattern. (Although in high-quality fMRI recordings, partial methods were found by Smith *et al.* (2011, 2013) to be more discriminative than marginal approaches.) Of these, we believe that the partial correlations between leakage-corrected power envelopes show sufficient repeatability, particularly in group-average inference, to be recommended for most resting state experiments. We should

recognise that these recommendations are for methods that rely on orthogonalisation corrections for their removal of source leakage. If this corection badly fails then these methods may have inflated performance for the reasons discussed above. If directional connectivity estimation is required, a large amount of repeatability is sacrificed, but the PDC performed best out of the directed metrics on test. (Some concerns also exist about the robustness of PDC to leakage artefacts, see comment above in Methods section 3.2 on page 86.)

We find, over all frequency bands, phase-based measures and imaginary partial coherence to perform the worst on all assessments, in particular showing very poor within-subject repeatability. Why do these measures exhibit poor performance? We reject the notion that they do not measure anything related to cortical connectivity: there is a wide literature of interpretable results using these measures, and in our data they produce sensible connectivity structure at group level, and show better within-subject consistency than between-subject. Their acceptably-high repeatability for group-average inference, and well-known performance with large task datasets, suggest that these metrics may simply be noisier estimators than the others on test, leading to a reduced ability to extract a consistent basis of functional connectivity from noisy recordings. Longer resting-state scans than are used here, perhaps of fifteen minutes or more, may be required to sufficiently increase the SNR that these measures improve performance at the single-subject level. We note also that Hardmeier *et al.* (2014) find good test-retest reliability with PLI for sensor-level EEG analyses (intraclass correlation coefficients of 0.6-0.8) using recordings of 12 minutes. We have restricted ourselves to a study of repeatability, and so it may be that (despite the strong correlations between group-average results between each metric) some phase-based methods are more sensitive than others to network covariation

with subject attributes or transient connectivity features. Further, we have only performed our analysis using a data-driven parcellation from fMRI, and it may be that some metrics would show improved performance should a reliable MEG parcellation driven by the patterns of phase-based connectivity within the data become available. Korhonen *et al.* have described a parcellation that is designed to maximise the phase synchrony within each ROI. This may work well to ensure that ROIs describe a single, complete, cortical source, and is reported to increase (in simulation) the true positive rate for PLV-based connectivity analyses. Unlike an ICA approach, however, it may not be driven strongly by between-source connectivities. The potential for this type of parcellation to improve phase-based connectivity could be tested in future work. Lastly, it is possible that the beamformer used for source reconstruction shows variability between runs, for example in estimating the source orientations. While this would affect all of our tests equally, it is a concern which is yet to be fully explored in the literature.

In this chapter, we have focussed on the behaviour of connectivity measures commonly employed in MEG and EEG, applied to a large resting-state dataset. We believe that most of our conclusions should carry over to many task paradigms, where cortical responses can be more repeatable between trials and subjects than in the resting state. However, it may be that this inherent repeatability in response boosts the SNR of many connectivity metrics when connections are averaged over trials, such that poorly-performing measures in our study, such as the PLI, become efficacious.

It is becoming established practice that resting-state connectivity analyses in MEG and EEG are carried out in source space, using a set of functionally-defined regions which cover as much of the cortex as possible given the recorded data. We further suggest employing network measures that are inured to the

connectivity confounds introduced by source leakage profiles. To satisfy all these purposes, our results suggest that the most consistent connectivity measure to employ for resting-state studies is the correlation between orthogonalised, band-limited, power envelopes.

4

The heritability of multi-modal functional connectivity

Patterns of intrinsic human brain activity exhibit a profile of functional connectivity that is associated with behaviour and cognitive performance, and deteriorates with disease. This chapter investigates the relative importance of genetic factors and the common environment between twins in determining this functional connectivity profile, using data from two functional imaging techniques, fMRI and MEG. It provides an illustrative application of the MEG analysis methods developed and evaluated in chapters 2 and 3, and demonstrates the utility of performing comparable analyses in two modalities.

Using fMRI data from 477 subjects of the Human Connectome Project, and MEG recordings from a subset, the heritability of connectivity between 39 cortical regions was estimated. On average over all connections, genes account for up

to 20% of the observed variance in fMRI connectivity (20-65% in MEG), which is seven times the contribution from the environment shared between twins (one-and-a-half times in MEG). Therefore, in as far as twins share a common upbringing, it appears that genes, rather than the developmental environment, play a dominant role in determining the coupling of neuronal activity.

4.1 Introduction

Chapters 2 and 3 have outlined and verified a methodology for performing correlation-based multivariate network analyses in MEG. This enables studies using both fMRI and MEG data to use very similar connectivity analysis methods. The power of such a multi-modal approach is that coincident results from both modalities can lend credence to inference on the relationships of interest—functional connectivity, in our case—rather than to the various confounds associated with measurements from either scanner (such as vascular effects or concerns over the quality of source reconstruction). Additionally, we can draw from the strengths of each modality: MEG gives us frequency-specific insight into the coupling of neural oscillations and fMRI can provide improved spatial specificity, and more robust connectivity estimation.[1]

This chapter develops an example multi-modal analysis in this fashion, and takes the opportunity to act as a case-study in functional network analysis. We will perform connectivity analyses between the same set of functionally-defined cortical ROIs in two different imaging modalities. Modality-specific confounds

[1] There are not, to our knowledge, studies comparing the quality or repeatability of network estimation using MEG and fMRI, so this final point is subjective. In our experience, functional networks inferred with fMRI are more repeatable within subjects than those from MEG, and more capable of predicting task responses and activations—but this may all be a consequence of the larger datasets available.

are either removed or subject to control analyses, and we employ appropriate methods for the control of type I error rates over multiple statistical tests.

The Human Connectome Project is a multi-modal twin study design, and the large datasets it provides create an ideal opportunity for the application of our methods to answer important questions on the role of genes in the generation of human functional connectivity. As discussed in chapter 1, profiles of functional connectivity are associated with cognition, health and happiness. We were therefore interested in identifying the relative importance of genetic and shared environmental factors in determining these fundamental patterns of cortical communication.

We employ fMRI recordings to perform a heritability analysis on the strength of functional connectivity within a network of 39 regions. Additionally, we perform the same heritability analysis on source-localised MEG recordings (Van Veen *et al.*, 1997) from a subset of the subjects. This complementary analysis allows us to focus specifically on communication mediated by the coupling (correlation) of oscillatory amplitudes, within particular frequency bands, using MEG as a more direct measure of neuronal activity that is unaffected by vascular confounds.

We present separate analyses of functional networks estimated from the fMRI response (Smith *et al.*, 2013), and of MEG-derived networks in the theta (4–8 Hz), alpha (8–13 Hz) and beta (13–30 Hz) oscillatory bands. These span the frequency range within which the most convincing patterns of resting-state MEG connectivity have been shown to be expressed (Baker *et al.*, 2014; Brookes *et al.*, 2012*b*, 2011*b*; Hipp *et al.*, 2012; Mantini *et al.*, 2007; de Pasquale *et al.*, 2012, 2015). We employ resting-state recordings from 477 subjects from the Human Connectome Project (HCP; Van Essen *et al.*, 2013). Subjects from the

S500 data release with fMRI (MEG) scans comprise 56 (11) monozygotic (MZ) twin pairs, who share 100% of their genetic structure and a common environment, and 57 (11) dizygotic (DZ) twin pairs, who share 50% of their genetic structure and a common environment. Each subject has provided three MEG scans (of six minutes each) and four fMRI scans (of fifteen minutes each). This stratified sample allows the estimation of the relative effects of genetic influence and environmental factors on the variability observed in connectivity structure.

We compare the similarity in overall network structure between pairs of subjects, assessing the extent to which the functional connectivity of two subjects becomes more alike as their proportion of shared background and genetics increases. We fit variance-component models on each network edge, using recently-developed permutation and bootstrap-based methods for fast, accurate, non-parametric statistical inference with family-wise control of type I errors (Chen, 2014). This provides estimates of the mean genetic and shared environmental influences on the observed phenotypic variation in cortical connectivity. We run similar heritability analyses on the oscillatory power and BOLD variance within the ROIs which constitute the network nodes, to determine whether the observed effects of genes and developmental environment on cortical connectivity could be attributable to simple differences in signal strength. Lastly, we analyse the heritability of cortical folding patterns in each ROI, to investigate whether genetic control of anatomical structure is contributing to the heritability of connectivity.

4.2 Methods

SUBJECTS

Four fifteen-minute resting-state fMRI recordings from 461 subjects were collected as part of the HCP500+MEG2 data release by the Human Connectome Project (Larson-Prior *et al.*, 2013; Van Essen *et al.*, 2013). Additionally, 61 subjects provided three six-minute resting-state MEG recordings. 45 subjects are shared across both cohorts. All subjects are young adults (22-35 years of age) and healthy. Of the fMRI (MEG) subjects, 112 (22) are monozygotic twins and 114 (22) are dizygotic twins. Zygosity of twin subjects was self reported. The MEG dataset is the same as that employed in chapter 3.

FMRI ANALYSIS

Resting-state fMRI data were acquired with 2 mm isotropic spatial resolution and a temporal resolution of 0.72 s. The HCP provides comprehensively preprocessed data (Glasser *et al.*, 2013; Robinson *et al.*, 2014) which are registered to a standard cortical surface, and for which structured artefacts have been removed by a combination of ICA and `FIX` (Salimi-Khorshidi *et al.*, 2014), `FSL`'s automated noise component classifier.

We estimated functional connectivity between 39 fMRI-derived cortical ROIs. We used the parcellation employed in chapter 3, which contains contiguous clusters identified from a resting-state group-ICA decomposition on recordings from the first 200 HCP subjects, and re-formed for our fMRI analysis using identical operations but applied to the same ICA maps defined on the cortical surface rather than within the standard volume. The parcellation can be seen

in figure 3.1C on page 90, and ROI labels are shown in table 3.1 on page 80. A single BOLD time course to represent each ROI was constructed using multiple spatial regression in FSLnets.[2] Partial correlation matrices were constructed using the tools in FSLnets, using mild Tikhonov regularisation ($\lambda = 0.01$) and a conversion to z-values with Fisher's transform, scaled to have the same standard deviation as the correlations within eight null datasets, using the methodology described in section 2.3 on page 48. Group-average networks were estimated as the mean z-transformed correlation matrices over all sessions.

MEG ANALYSIS

Analysis of the MEG data proceeded as outlined in section 3.2 on page 78, using only data in the theta (4–8 Hz), alpha (8–13 Hz) and beta (13–30 Hz) bands. The same parcellation was employed as in chapter 3 (and as here for the fMRI data), and power envelopes of the data computed in the same fashion, after using the orthogonalisation correction of chapter 2 to remove confounds to network analyses from spatial leakage effects. A schematic of this analysis pipeline was presented in figure 2.4 on page 48.

Correlations between the power time courses in each band were computed, converted to z-values using Fisher's transform, and standardised with reference to an empirical null distribution, as described above for the fMRI data. Also in the same fashion, group-average networks were estimated as the mean z-transformed correlation matrices over all sessions.

We use correlation measures in MEG and partial correlation measures in fMRI as these are the most robust network construction approaches for each modality (chapter 3; Smith *et al.*, 2011, 2013) while maintaining similarity in the

[2] FSL's network analysis toolbox is available at fsl.fmrib.ox.ac.uk/fsl/fslwiki/FSLNets/.

analysis between the two modalities.

COMPARISON OF ENTIRE NETWORK STRUCTURE

To compute the dissimilarity between network matrices estimated from different recording sessions, we constructed a distance metric by mapping the positive definite correlation matrices to a Euclidean space (Barachant *et al.*, 2013; Ng *et al.*, 2014; ?), then taking the 2-norm of the vector separating the transformed points. A suitable space for each band was found as the Euclidean plane tangent to the Riemann manifold of positive semi-definite matrices, taking the geometric mean of the networks computed from all sessions as the tangent point. The logarithmic mapping between Riemann manifold and tangent Euclidean space can decouple the inter-relations between the elements within each positive-definite network matrix, and can improve the performance of classification algorithms when network matrices are used as discriminative features (?). It also provides for the definition of a true distance metric, where more similar matrices are projected to more proximate locations on the plane.

The separation between each pair of network matrices in each band was computed, forming different distributions of network similarity for pairs of subjects, split by the shared genetics of each pair. To assess whether shared environmental factors and shared genetics were associated with global network structure, we tested the difference in mean between the distributions of the logarithm of network separation for pairs of unrelated subjects and pairs of dizygotic twins, and between pairs of dizygotic twins and monozygotic twins, using a non-parametric t-test based on 5000 permutations of the group labels. In the fMRI data, we were also able to compare unrelated subjects to siblings, and siblings to dizygotic twins.

We applied a False Discovery Rate (FDR) correction to all 21 statistical tests conducted for this chapter (over modalities, frequency bands, etc.) to compensate for the multiple comparisons we perform (Benjamini & Hochberg, 1995). Uncorrected and corrected p-values are available in tables 4.2 and 4.3 on pages 124–125.

THREE-COMPONENT VARIANCE MODELS

Estimates of heritability using linear regression on squared differences

The ACE linear variance-components model splits observed phenotypic variability σ^2 into three factors: (A), additive genetics or heritability (h^2), (C), common environment (c^2), and (E), measurement error or external sources of variability,

$$\sigma^2 = A + C + E\,, \tag{4.1}$$

such that the heritability and common environment factors can be expressed as proportions of the total variability,

$$\begin{aligned} h^2 &= A/\sigma^2 \\ c^2 &= C/\sigma^2\,. \end{aligned} \tag{4.2}$$

The APACE[3] system of fast, accurate heritability estimation with permutation inference (Chen, 2014) was employed to infer on these components. In brief, connectivity strengths on each network edge are assumed to vary about the population mean in a correlated fashion according to family structure. Taking $\boldsymbol{y}$ to be connection strengths for a single edge, for a set of n subjects, and $\boldsymbol{X} \in \mathbb{R}^{n \times p}$

[3] The Advanced Permutation inference for ACE models (APACE) software is available at warwick.ac.uk/tenichols/apace/.

a design matrix with a bias term and $p - 1$ confound regressors,

$$\boldsymbol{y} = \boldsymbol{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}\,, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{V})\ . \tag{4.3}$$

Writing $\boldsymbol{\mu} = (A, C, E)^T$, the covariance matrix describing the family structure can be written as

$$\boldsymbol{V} = \sum_{l=1}^{3} \mu_l \boldsymbol{Q}_l\,, \tag{4.4}$$

and the construction of the $\boldsymbol{Q}$ matrices is best illustrated by example. Assuming we have two MZ twins, two DZ twins and one unrelated subject, arranged in that order, we see easily that $\boldsymbol{V}$ is constructed using

$$\boldsymbol{Q}_1 = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1/2 & 0 \\ 0 & 0 & 1/2 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}, \quad \boldsymbol{Q}_2 = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}, \quad \boldsymbol{Q}_2 = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

Accurate inference of non-negative variance components $\boldsymbol{\mu}$ can be performed in a simple linear regression framework by noting that the expectations of squared differences in connectivity strengths between subjects have simple linear relationships to the variance components A, C and E. Permutation inference on the null hypotheses $h^2 = 0$ on each network edge, and $\bar{h^2} = 0$, taking the average over all edges, is possible by relabelling the MZ / DZ statuses of twin pairs. For full details we refer to Chen (2014).

Application of the ACE model to fMRI and MEG data

An ACE model for the subject-to-subject variability in connection strength was fitted for each network edge, using the mean network matrix (over resting-state

recording sessions) for each subject. On each edge, we regressed out the effect of age, the square of age, sex, an age and sex interaction, the cube root of intracranial volume and of cortical volume (both estimated with `FreeSurfer`), and, for the fMRI data, a measure of subject motion in the scanner (fMRI_motion) and the MR image reconstruction software version. Estimates of heritability were computed on each edge, with family-wise error corrected p-values computed by permutation, randomly shuffling monozygotic and dizygotic twin statuses 15 000 times. Over all edges, mean h^2 and c^2 were computed; the p-value for mean heritability was computed by permutation, as above; confidence intervals for h^2, c^2 and $\left(h^2 - c^2\right)$ were computed by bootstrap re-sampling twin-pairs with replacement 15 000 times.

The ACE model was also fitted to the logarithm of the variance of the corrected power envelopes (and the variance of the BOLD time course) in each ROI. This provided confidence that the conclusions of the ACE model for functional connectivity were not being driven straightforwardly by heritable effects in power or SNR within the network nodes. We additionally re-ran the ACE models for functional connectivity, adding in as an additional regressor the square root of the product of the power of the two nodes which each connection joins; the results were not substantively changed.

Lastly, we sought to understand whether the heritability effects we observe were driven by heritable patterns of cortical folding. The ACE model was fitted within each ROI to estimates of cortical curvature (calculated in `FreeSurfer` and provided by the HCP) for each of the subjects in the fMRI sample. We computed the mean heritability (h^2) over all points within each ROI. In order to compare spatial profiles of cortical folding with connectivity, we averaged the heritabilities of each network connection onto their constituent nodes, to create

a spatial map suggesting the heritability of connectivity by ROI. These maps were correlated with the maps of heritability of cortical curvature in each ROI, and significance assessed by permuting ROIs 5000 times.

4.3 Results

The functional connectivity measured with fMRI was estimated between 39 functionally-defined cortical regions of interest (network nodes) by computing partial correlations between BOLD time courses representing each region, and the functional connectivity measured with MEG, corresponding to amplitude coupling of oscillations in each frequency band, was assessed between the same set of regions by correlating fluctuations in oscillatory power (Engel *et al.*, 2013; O'Neill *et al.*, 2015). The group-averaged functional networks are shown in figure 4.1A on page 113. The fMRI connectome shows strong bilateral connectivity, together with fronto-parietal correlations and coupling between the anterior cingulate and posterior cingulate cortices, reflecting well-known patterns of resting-state networks in fMRI (Beckmann *et al.*, 2005; Smith *et al.*, 2012). In the MEG data, the alpha band shows a strongly-connected visual system, with strong connectivity from visual cortex extending out to the temporal and parietal lobes, and a highly-coupled posterior cingulate. The beta-band exhibits broad connectivity over the whole cortex, with the strongest coupling within the sensorimotor cortices and superior parietal lobules. The theta band also exhibits visual and motor connectivity. These patterns of connectivity are in general agreement with previously reported results in these frequency bands (Baker *et al.*, 2014; Brookes *et al.*, 2011*b*; Hillebrand *et al.*, 2012; Hipp *et al.*, 2012; de Pasquale *et al.*, 2012, 2015; and the results presented in chapter 2).

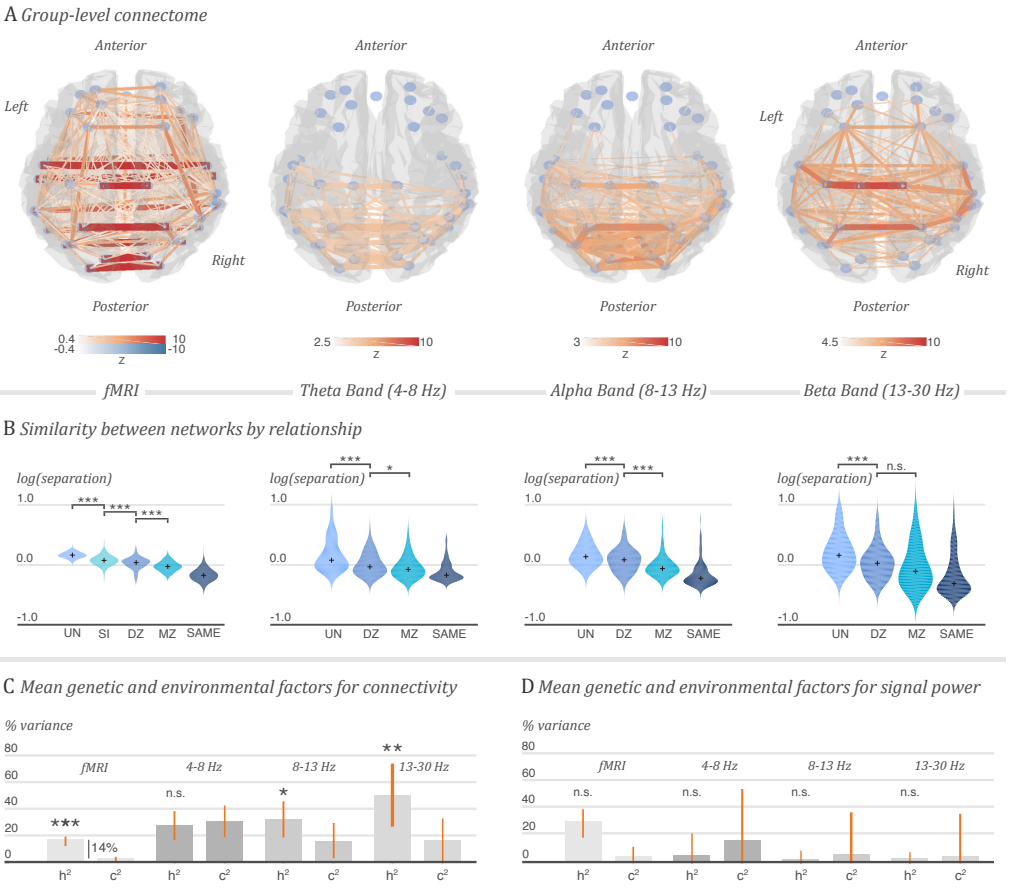


Figure 4.1: CONTRIBUTION OF GENETIC FACTORS AND ENVIRONMENTAL EFFECTS SHARED BETWEEN TWINS TO THE COUPLING OF NEURAL OSCILLATIONS IN THE THETA, ALPHA AND BETA BANDS, AND TO THE CONNECTIVITY OF FUNCTIONAL BEHAVIOUR AS INDEXED BY THE FMRI BOLD RESPONSE. (A) *Group average functional connectome for each oscillatory band and for* FMRI. The coloured edge maps show group-average network matrices for correlations in oscillatory amplitude in each band, and partial correlations in BOLD response, thresholded for visualisation. (B) *Similarity of pairs of network matrices, separated by the relationship status of each pair.* Subjects with a shared environment, and a greater proportion of shared genetics, have a more similar organisation of neuronal coupling. Violin plots show the distribution of distance values between pairs of network matrices estimated from single resting-state recording sessions, on a logarithmic scale and relative to the mean network separation, for pairs of unrelated subjects (UN), siblings (SI), dizygotic twin pairs (DZ), monozygotic twin pairs (MZ) and for repeated sessions with the same subject (SAME). (C) *Average genetic and shared environmental contributions to the variability of functional connectivity.* The bar charts show the mean genetic component (heritability, h^2) and mean common environment component (c^2) from a variance-components model fitted on each network edge, with 95% bootstrapped confidence intervals on the mean. (D) *Average genetic and shared environmental contributions to the variability of oscillatory power, or* BOLD *response, in each* ROI. Stars indicate significant differences in mean value for each of the displayed comparisons: ★★★ : $p < 0.001$; ★★ : $p < 0.01$; ★ : $p < 0.05$; n.s., not significant. The non-parametric p-values were computed by permutation and corrected for multiple comparisons over the twenty-one tests performed. Confidence intervals formed by bootstrap and the corrected permutation-based p-values do not have to align; as a result, not all estimates where the 95% confidence interval is above the zero-axis in C and D are marked as significant.

The similarity of the entire functional network structure, estimated from individual recording sessions, was compared between pairs of unrelated subjects, dizygotic twin pairs and monozygotic twin pairs. Figure 4.1B (left-hand plot, middle row) shows that the structure of fMRI networks of functional activity are progressively more similar as the strength of relationship is increased, from unrelated subjects, through siblings and dizygotic twins to monozygotic twins ($p^{fMRI} < 10^{-3}$ in each case). MEG oscillatory coupling is also significantly more similar between monozygotic twins than between dizygotic twins in the theta and alpha bands ($p^{\theta}_{MZ<DZ} = 0.04$, $p^{\alpha}_{MZ<DZ} < 10^{-3}$, $p^{\beta}_{MZ<DZ} = 0.11$), which is again more alike than between pairs of unrelated subjects, in all three bands ($p_{DZ<UN} < 10^{-3}$ for each band). Also displayed, for comparison, is the distribution of network similarity over repeated recording sessions within the same subject. Tests for significance were performed by assessing the difference in mean of the logarithm of network separation, relative to the mean separation, using a non-parametric permutation-based t-test. All 21 tests performed for this chapter are quoted after a false discovery rate correction for multiple comparisons.

A three-component variance model was fitted for the variability observed in the strength of coupling in individual edges. This model ascribes non-negative proportions of the variance either to additive shared genetics (the heritability, h^2), shared environmental factors (c^2), or to measurement error and other unknown sources of variance. Estimates of the mean heritability and contribution from common environment over these connections were constructed with bootstrapped confidence intervals, and the significance of the mean heritability tested with non-parametric permutation-based methods. Although there is a large amount of variability in the data that is not attributable to either of these

factors, we found that additive genetics accounted for about 20%, on average, of the functional connectivity measured with fMRI. In MEG, about 30% of the variability in oscillatory coupling strengths in the alpha band, and 50% of the variance in the beta band, was determined by additive genetics (although there was no significant heritability on average in the theta band). These estimates and comparisons are presented graphically in figure 4.1C, and tabulated with confidence intervals and *p*-values in tables 4.1 and 4.2 on pages 123–124.

In the absence of a permutation scheme that is invariant under the null hypothesis of equal genetic and shared environmental factors, we additionally estimate $\left(h^2 - c^2\right)$ with bootstrapped confidence intervals to investigate the relative importance of these two components. Only in the fMRI data does our confidence in a differential effect of heritability and shared environment exclude zero (estimate with 95% confidence interval, $\left(h^2 - c^2\right) = 0.14\,[0.09, 0.17]$). Permutation tests for a significant contribution of heritability on each individual functional connection were performed, but no edges were significant after correcting for family-wise error at $\alpha = 0.05$. Parameter estimates for individual fMRI connections are visible in figure 4.2 on page 116.

The three-component model was also fitted to the variability in the logarithm of the signal power in each ROI, in both the fMRI and MEG data, but no mean influence of genetic factors was found (figure 4.1D).

Finally, we fitted the heritability model to the cortical surface curvature of each subject, computing average h^2 within each ROI. Cortical curvature exhibits the highest heritability around the sensorimotor areas and insula (as has been reported previously, in an analysis using a different sample of the HCP subjects; Sotiropoulos *et al.*, 2015). Figure 4.3 on the following page compares this spatial profile of the heritability of curvature to the average heritability of

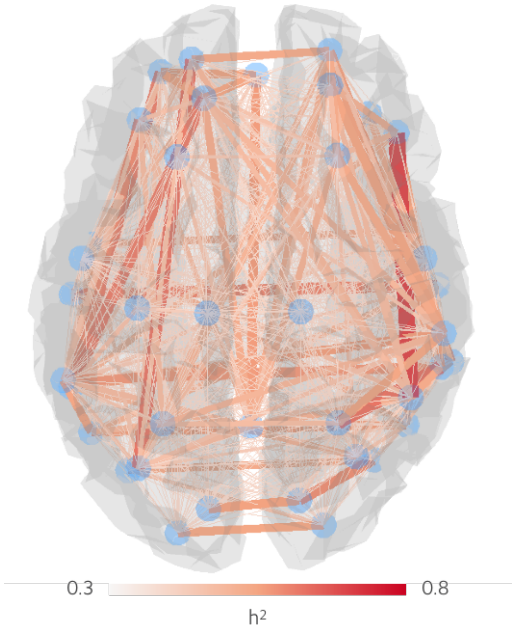


Figure 4.2: HERITABILITIES OF FMRI CONNECTION STRENGTHS. Parameter estimates for the influence of additive genetics, h^2, on each individual network connection, for partial correlations in BOLD time courses derived from each ROI.

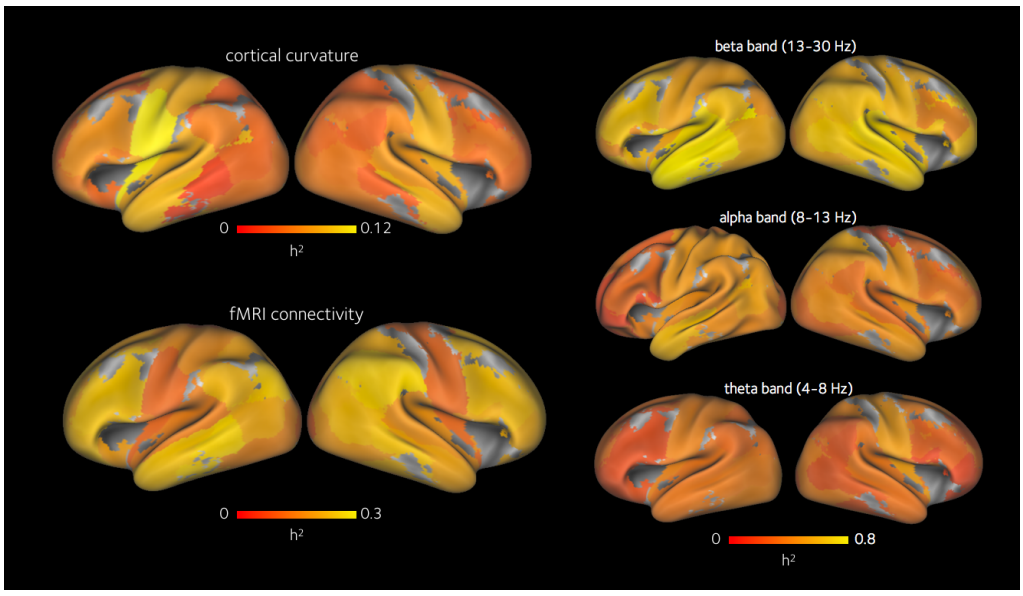


Figure 4.3: HERITABILITY OF CORTICAL CURVATURE, AND THE SPATIAL PROFILE OF THE HERITABILITY OF FUNCTIONAL CONNECTIONS. Top left: mean heritability of cortical curvature within each ROI. Bottom left: heritability of each fMRI network connection, averaged onto each relevant ROI. Right, top to bottom: heritability of each MEG network connection in beta, alpha and theta bands, averaged onto each relevant ROI.

network connections for each ROI within each modality (that is, for each ROI, the average h^2 of all connections involved with that ROI). In general, no significant positive correlations were found between the spatial arrangement of cortical curvature heritability and the heritability of network connections in particular ROIs. The exception was moderate correlation ($\rho = 0.39$) with the MEG theta band connectivity pattern, for which there was no significant heritability of connection strength in the first instance. (These data are presented in table 4.3 on page 125.)

4.4 Discussion

Using resting-state fMRI and MEG recordings released as part of the HCP, we have constructed the functional network structure which expresses the coupling in MEG oscillatory power within three frequency bands, and fMRI partial correlation networks, between 39 ROIs. Based on these analyses, we make two key claims. First, genetic factors help to determine the strength and form of cortical oscillatory communication and functional connectivity. We have shown that the entire functional network structure is found to be more similar for two subjects the more closely they are related, and we estimate that the average heritability of individual connection strengths is about 15-20% for BOLD correlations, and between 20% and 65% in the alpha and beta frequency bands. Secondly, the fMRI data demonstrate that genetic make-up is more important, on average, than the shared environment among twins, when determining the strength of these couplings. A similar trend is visible in the MEG networks.

Our results are drawn from two imaging modalities, fMRI and MEG, with functional connectivities estimated from the same set of ROIs. Our analyses

with these technologies produce complementary assessments of functional connectivity: slow-timescale couplings of functional activation indexed by BOLD response, and faster co-ordinations in oscillatory amplitude measured with MEG. (MEG is a more direct measure of neuronal oscillatory activity, unaffected by vascular confounds.) We found, over both modalities and in two of the three frequency bands, the same patterns of influence on functional connectivity from additive genetics and twins' shared developmental environment. Strong confidence in the relative influence of the former over the latter was only possible in the much larger fMRI sample. Given the differences in connectivity structures that are expressed at these different oscillatory frequencies and in the slow-timescale couplings measured in fMRI, these results suggest that genes have a broad control over functional connectivity in cortex, which outweighs shared environmental factors.

The 61 subjects scanned with MEG under the HCP 500-subject dataset present a relatively small sample for an heritability analysis. As a result, our conclusions from this modality, while coincident with the results in fMRI, are necessarily simple characterisations of the role of genes and environment on cortical oscillatory coupling. The confidence intervals on our parameter estimates are inevitably broad, although our claims are based on significant results from robust non-parametric tests. Various connectivity analyses in EEG (Posthuma *et al.*, 2005; Schutte *et al.*, 2013) have produced estimates in a similar range (with additive genetic heritabilities of 20-75%) as those reported here. However, these studies used different methods for estimating heritability, and have to-date been performed on the signals measured in scalp electrodes rather than on electromagnetic sources localised in cortex, and are thus less interpretable (and less reliable; Schoffelen & Gross, 2009) as measures of cortical coupling.

Our results join the growing evidence for the importance of genetic control on human cortical connectivity. Structural connectivity, the distribution of the white matter tracts across the brain that enable communication, is known to be highly heritable (Jahanshad *et al.*, 2012; Kochunov *et al.*, 2015; Zhu *et al.*, 2015)—and in the HCP dataset, Sotiropoulos *et al.* (2015) estimated the mean heritability of structural connections to be 25%. Previous studies of functional connectivity, while predominantly utilising fMRI rather than electrophysiology, and focussing on network topology or particular sub-networks rather than the profile of whole-brain connectivity, also find significant genetic effects with heritability estimates lying between 20%–60% (Fornito *et al.*, 2011; Glahn *et al.*, 2010; van den Heuvel M. P. *et al.*, 2013; Jansen *et al.*, 2015; Smit *et al.*, 2008; Thompson *et al.*, 2013). Our finding of a differential influence between additive genetic and shared environmental factors is mirrored in a broad range of cognitive, behavioural, physiological and neurological traits measured in adults (Polderman *et al.*, 2015). Our estimates of the heritability of connectivity are lower than the 40–60% often exhibited in these domains, and lower than that of simple neural oscillatory characteristics such as frequency or power (van 't Ent *et al.*, 2009; Posthuma *et al.*, 2001), but this may simply reflect the difficulty in obtaining clean connectivity estimates. (Heritability is expressed as a fraction of the observed variance, which includes measurement noise.) It is therefore interesting that we are able to observe a differential effect of genetics and shared environment in this regime.

Twin studies cannot probe the precise genetic mechanisms of these influences. Recent work (Hawrylycz *et al.*, 2015; Richiardi *et al.*, 2015) has correlated the profiles of gene expression in different regions of cortex, and compared these patterns of gene transcription to resting-state fMRI connectivity profiles,

identifying sets of genes with co-expression patterns that reflect patterns of functional connectivity. However, detailed investigations into the specific genes affecting healthy resting-state connectivity must await large datasets with coincident resting-state functional imaging and genotyping. (Genetic data will soon be available both for the HCP and for the UK Biobank imaging data.)

We have used variance-component models to represent the observed variability in functional connectivity as a set of contributions from additive genetic factors and from shared developmental environment in twin pairs. Our conclusions rest on a number of assumptions implicit in this model (Boomsma *et al.*, 2002). These include an assumption that people choose their partners randomly; that the relevant genetic mechanisms are additive; and that there is no significant interaction between genes and the shared environment—an effect which has been reported to exist in various aspects of cognition (Nisbett *et al.*, 2012). Most importantly, we assume that monozygotic and dizygotic twin pairs will equally share exposure to environmental factors in their upbringing. If this is an accurate assumption, then estimates of c^2 and h^2 are informative about the relative roles of developmental environment and additive genetics in the variation of functional connectivity phenotypes. While there is growing evidence in support of the equal-environment assumption (Conley *et al.*, 2013; Felson, 2014), it has been frequently challenged on the basis that monozygotic twins look more alike, behave more alike, and are treated more similarly than dizygotic twins (Joseph, 1998). Without entering into the lively discussion on this point, we note that if this assumption fails in our sample, it would lead to an overestimation of genetic heritability and an underestimation of the impact of developmental environment, therefore potentially strengthening artificially our conclusions on the differential effect of these two influences. Lastly, the

c^2 term in our model cannot be solely identified with developmental environment. It encompasses all effects which are common between twins, which may include intra-uterine and mitochondrial influences, in addition to the shared environmental factors.

It is commonly observed, in both fMRI and MEG estimates of functional connectivity, that higher coupling strengths are identified for regions with higher signal power. This effect can be caused just by the increased signal to noise ratio in the observed regions, rather than any difference in underlying coupling strength (Friston, 2011). By fitting a genetic and environmental factors model to the power in each ROI of our analysis, and finding no influence of these factors on the measured signal strength, we are able to exclude this signal-to-noise effect as a likely confound in our analysis. Finally, it is possible that our heritability measurements include indirect measures of other effects on functional connectivity, such as motion or breathing patterns (although the data are pre-processed to minimize physiological artefacts), levels of arousal, cortical folding patterns or some similarity in source localisation. Of these concerns, the most substantive is perhaps the influence of cortical folding patterns (which are potentially heritable; see evidence either way in Gómez-Robles *et al.*, 2015; Tallinen *et al.*, 2016; and Ronan & Fletcher, 2015), in particular for the MEG measurements. The lack of strong correlations between the spatial profile of the heritability of cortical curvature and the heritability of edges associated with each ROI (as shown in figure 4.3 on the following page and table 4.3 on page 125) is good evidence, we believe, that this confound has not seriously impacted our findings.

We have made particular emphasis of the increased importance of genetics over developmental environment in determining cortical functional connectiv-

ity. The relevance and implications of this finding are widespread. Functional connectivity profiles are associated with intelligence, which is well-known to be heritable (Bartels *et al.*, 2002; Neisser *et al.*, 1996), and are also implicated with a broad range of behavioural and life-style factors, including earning power, measures of health, assessments of cognitive performance, and self-reported life satisfaction (Smith *et al.*, 2015). As the findings of the Moving To Opportunity experiment in the United States make clear (Chetty *et al.*, 2015; Ludwig *et al.*, 2013), together with the results from studies on the effects of schooling, socio-economic environment and adoption (Nisbett *et al.*, 2012), the environment in which individuals develop can still make an impact on their performance and long-term outcomes. However, while genes are not the only important factor, results such as we present here suggest that nature, rather than nurture, is a major force in determining integrated cognitive function and, by extension, cognitive capability.

Table 4.1: PARAMETER ESTIMATES FOR THE ACE MODEL ON FUNCTIONAL CONNECTIVITY. Parameter estimates and 95% confidence intervals for the mean genetic and shared environmental contributions to the observed phenotypic variability in functional connectivity and signal power.

	Estimated parameter	Value	95% confidence
Connectivity ACE model (θ)	h^2	0.27	[0.17, 0.38]
	c^2	0.31	[0.16, 0.46]
	$(h^2 - c^2)$	−0.03	[−0.26, 0.18]
Connectivity ACE model (α)	h^2	0.32	[0.19, 0.44]
	c^2	0.16	[0.06, 0.30]
	$(h^2 - c^2)$	0.16	[−0.06, 0.32]
Connectivity ACE model (β)	h^2	0.50	[0.27, 0.66]
	c^2	0.16	[0.00, 0.42]
	$(h^2 - c^2)$	0.34	[−0.12, 0.64]
Connectivity ACE model (fMRI)	h^2	0.17	[0.14, 0.19]
	c^2	0.03	[0.02, 0.04]
	$(h^2 - c^2)$	0.14	[0.09, 0.17]
Power ACE model (θ)	h^2	0.05	[0.00, 0.50]
	c^2	0.16	[0.00, 0.54]
Power ACE model (α)	h^2	0.02	[0.00, 0.30]
	c^2	0.06	[0.00, 0.37]
Power ACE model (β)	h^2	0.03	[0.00, 0.36]
	c^2	0.04	[0.00, 0.36]
Power ACE model (fMRI)	h^2	0.30	[0.19, 0.39]
	c^2	0.04	[0.00, 0.11]

Table 4.2: SIGNIFICANCE OF THE HERITABILITY OF FUNCTIONAL CONNECTIVITY. p-values for permutation-based significance tests performed for this chapter, both before and after a false discovery rate correction for multiple comparisons over the 21 tests performed in this chapter.

	Test	p	p_{FDR}
Network Dissimilarity (θ)	$UN > DZ$	3×10^{-8}	1×10^{-7}
	$DZ > MZ$	0.02	0.04
Network Dissimilarity (α)	$UN > DZ$	1×10^{-4}	3×10^{-4}
	$DZ > MZ$	1×10^{-5}	4×10^{-5}
Network Dissimilarity (β)	$UN > DZ$	7×10^{-8}	3×10^{-7}
	$DZ > MZ$	8×10^{-2}	0.12
Network Dissimilarity (fMRI)	$UN > SI$	1×10^{-8}	1×10^{-7}
	$SI > DZ$	1×10^{-8}	1×10^{-7}
	$DZ > MZ$	1×10^{-8}	1×10^{-7}
Connectivity ACE model (θ)	$h^2 > 0$	6×10^{-2}	0.10
Connectivity ACE model (α)	$h^2 > 0$	0.02	0.03
Connectivity ACE model (β)	$h^2 > 0$	5×10^{-3}	0.01
Connectivity ACE model (fMRI)	$h^2 > 0$	1×10^{-4}	3×10^{-4}
Power ACE model (θ)	$h^2 > 0$	0.60	0.67
Power ACE model (α)	$h^2 > 0$	0.66	0.69
Power ACE model (β)	$h^2 > 0$	0.58	0.67
Power ACE model (fMRI)	$h^2 > 0$	0.14	0.20

Table 4.3: CORRELATIONS BETWEEN THE HERITABILITY OF CORTICAL CURVATURE AND THE HERITABILITY OF FUNCTIONAL CONNECTIVITY. Correlations over ROIs, with permutation-based p-values, between the average heritability of cortical curvature in each ROI and the average heritability of connections from each ROI. A lack of strong positive correlations suggests that any heritability in cortical curvature is not driving the heritability observed in functional connection strengths. p-values are given both uncorrected, and after a false discovery rate correction for multiple comparisons over the 21 tests performed in this chapter.

	ρ	p	p_{FDR}
fMRI	−0.57	1.00	1.00
Beta-band	−0.04	0.60	0.67
Alpha-band	0.15	0.18	0.23
Theta-band	0.39	0.01	0.01

5

A Hierarchical Model for Functional Network Analysis

A HIERARCHICAL MODEL for sparse multi-subject functional connectivity in the human brain is described, which can estimate the strength of connectivity between regions at both a single-subject and a population level, and which is applicable to both fMRI and MEG data. This model uses a Bayesian inference scheme to simultaneously infer functional connection strengths over all subjects. Explicit inference on the sparse structure of the underlying network graph is possible, producing a posterior distribution over network structures. The principal objective is to combine a hierarchical inference approach with sparse penalties to improve the quality of connectivity estimation for individual subjects, and we demonstrate these improvements using a range of simulated data and fMRI and MEG recordings from the Human Connectome Project. Finally,

in this dataset, we show that we are able to reduce the proportion of variability in MEG beta-band functional connection strengths that is associated with a range of unmodelled noise sources by up to 20 percentage points, increasing the estimate of genetic control by the same amount. This is equivalent to a noise reduction of about 50%.

5.1 Introduction

Functional connectivity is becoming an increasingly important part of fMRI and MEG studies into disease and cognition in humans. Among the major breakthroughs discussed in chapter 1 (on page 8), a common feature of the analyses is a focus on individual differences in connectivity patterns. Variations in subjects' functional connectomes have been analysed for genetic origins (chapter 4; Glahn *et al.*, 2010); correlated with wealth, health and life satisfaction (Smith *et al.*, 2015); shown to be sensitive to cognitive ability (Finn *et al.*, 2015; Rosenberg *et al.*, 2016) and neurodegeneration (Greicius, 2008; Stam, 2014); used as multivariate predictors of task-evoked activity (Tavor *et al.*, 2016); and depended upon for accurate spatial alignment of activation patterns (Robinson *et al.*, 2014) and categorisation of cortex (Glasser *et al.*, 2016; Yeo *et al.*, 2011). For all of these forms of analysis, accurate estimation of single-subject functional networks is crucial.

Despite the recent explosion of research and high-quality findings, whole-brain functional connectivity estimation is relatively immature. Biophysically-motivated network models are developing (Friston *et al.*, 2003, 2011; Kiebel *et al.*, 2008; Woolrich & Stephan, 2013), but are impractical for whole-brain connectomics. Most of the key developments highlighted above instead use

very simple Gaussian Graphical Models (GGMs) for the covariance of the data, in which the partial correlations between regions indicate the strengths of connections. However, accurate estimation of individual subjects' functional networks using GGMs can be difficult, particularly without long acquisition times. In an effort to reduce the noise in network estimation, sparsity in the networks tends to be encouraged by suppressing weak connections (Duff *et al.*, 2013; Smith *et al.*, 2013; Varoquaux & Craddock, 2013). Imposition of sparsity can also aid interpretation, by explicitly suggesting that certain individual functional connections are absent. There is an entire sub-field that characterises the function and dysfunction of cognitive networks using certain properties of this underlying graph structure (Bullmore & Sporns, 2009; de Pasquale *et al.*, 2012, 2015; Stam & van Straaten, 2012; van Straaten & Stam, 2013).

The immaturity of the field is most obvious in comparison to the modelling of linear activations in neuroimaging. It is relatively common to use hierarchical models for multi-session or multi-subject task and resting-state activations, and implementations are available in the popular software packages `SPM12` and `FSL` (Friston *et al.*, 2005; Harrison *et al.*, 2015; Woolrich, 2008; Woolrich *et al.*, 2004, 2009). In general, hierarchical models can explicitly capture between-subject variation, use information on subject-level measurement error to inform the uncertainty in group-level parameters, and regularise individual estimates using information from the whole group (Gelman *et al.*, 2014, 2008; Woolrich, 2008; Woolrich *et al.*, 2004).[1]

Relatively few efforts have been made towards the joint covariance estimation relevant for multi-subject, whole-brain network inference. (Our most

[1] A hierarchical model structure is one of the few methods that adequately controlled false positive rates in a recent, damning review of statistical analyses in neuroimaging (Eklund *et al.*, 2016).

complete list is Danaher *et al.*, 2015; Guo *et al.*, 2011; Harrison *et al.*, 2015; Lee & Liu, 2015; Liang *et al.*, 2016; Marrelec *et al.*, 2006; Qiu *et al.*, 2015; Varoquaux *et al.*, 2010; and Yang *et al.*, 2015.) None of these methods can simultaneously provide improved single-subject network estimation using sparsity penalties and the sharing of information over the whole dataset with a hierarchical model structure, nor estimates of the population connectivity and a Bayesian characterisation of the posterior uncertainty. The challenge is to achieve all of these goals within a computationally efficient framework that can be applied to large datasets of hundreds of recordings.

We present in this chapter a hierarchical model and scalable inference framework for Bayesian estimation of sparse functional networks, with joint characterisation of subject-level and population-average connectivities. We model functional connectivity simply as an undirected GGM between the network nodes—this model can be applied to MEG data in addition to fMRI (chapter 2), and is among the most successful and repeatable of measures in either modality (Ramsey *et al.*, 2014; Smith *et al.*, 2011; and chapter 3). We estimate covariances between nodes both at subject-level and group- or population-level. We impose a shared sparsity structure on the underlying network graph, and explicitly infer on this network structure. We prove that the inference problem can be reformulated as a series of linked linear regressions, and describe a Markov Chain Monte Carlo (MCMC) sampling approach. We evaluate the performance of our model on a number of simulated network systems of varying size and complexity, comparing to the current state-of-the-art in covariance modelling. We characterise the time for inference with our approach and how it scales with model dimension and the number of subjects. Lastly, we use fMRI and MEG data from the HCP to demonstrate the improvements possible to a range of follow-on

analyses when connectivity is estimated using our framework.

5.2 Existing approaches for sparse variable modelling

Encouraging sparsity in functional networks aids interpretability and improves the quality of estimation. A large body of work has developed for the estimation of sparse GGMs, which will be worthwhile reviewing so that our choice of strategy and its improvements on existing approaches are understood. Sparse precision models[2] have frequently borrowed from the literature for variable selection in linear regression problems, so we will discuss ideas from both areas.

OPTIMISATION METHODS

L_1 and L_2 penalties have a long history of use for the regularisation of unconstrained problems and the imposition of sparsity. Ridge regression (Hoerl & Kennard, 1970*a*,*b*) and the Least Absolute Shrinkage and Selection Operator (LASSO; Tibshirani, 1996) have been particularly successful when combined as the elastic net (Zou & Hastie, 2005), and the equivalents for covariance selection are probably the most widely used methods for estimating precision matrices (these are the graphical LASSO (Friedman *et al.*, 2008), which has been improved with covariance thresholding techniques (Mazumder & Hastie, 2012*a*) and updated to the DP-GLASSO (Mazumder & Hastie, 2012*b*); and the graphical elastic net, which was employed for fMRI data in Ryali *et al.*, 2012).

[2] Also referred to as covariance selection methods or concentration graph models.

CONTINUOUS PRIORS

The penalties employed in optimisation methods have inspired a wide range of prior distributions. L_1 norm regularisers have a direct equivalent in the Laplace or double exponential prior,[3] and the Bayesian LASSO has met with success (Hans, 2009; Park & Casella, 2008). Ishwaran & Rau proposed continuous bimodal priors, with one mode at the origin to encourage sparsity (Ishwaran & Rau, 2005), but a range of priors from the scale mixture of normals family have become more popular because of their flexibility, ability to adaptively impose sparsity on weak signals, and their tractable, conditionally conjugate formulations (notably the Cauchy distribution (Polson & Scott, 2012) and the horseshoe prior (Bhadra *et al.*, 2016; Carvalho *et al.*, 2010; Peltola *et al.*, 2014); for a review of the entire family see Polson & Scott, 2010).

Wang developed the Bayesian graphical LASSO for covariance selection, with Laplace priors on the off-diagonal elements of the precision matrix, and invented a block-Gibbs sampling scheme that operates efficiently across the columns of the matrix (Wang, 2012*a*). The same prior was used in Khondker *et al.* (2013), although with a less efficient Metropolis-Hastings (MH) sampler.

Continuous priors struggle to suppress variables completely to zero. Inference on the structure of a Gaussian graphical model is therefore difficult, but can be performed using an *ad hoc* thresholding procedure (Carvalho *et al.*, 2010; Wang, 2012*a*).

[3] Often employed in the automatic relevance determination literature—see MacKay (1995).

DISCONTINUOUS PRIORS

Methods of approximate Bayesian computation permit inference on models with prior structures that explicitly code for the presence or absence of variables, a stricture that is not possible to frame as a convex optimisation problem. The basic prior structure is that of a *spike and slab* distribution (Mitchell & Beauchamp, 1988):

$$b \sim \left(\mathcal{N}\left(0, \sigma^2\right)\right)^{\gamma} \left((1-a)\,\delta\left(0\right)\right)^{1-\gamma}, \quad \gamma \sim \text{Bernoulli}\left(a\right), \tag{5.1}$$

where the normal distribution represents the broad slab of prior mass, the Dirac delta-function is the spike at zero and γ is a selection variable that chooses between the two.

Apart from the Stochastic Search Variable Selection (SSVS) algorithm of George & McCulloch (1993) (and its counterpart for covariance selection in Wang, 2015), which approximates the spike with a narrow-width Gaussian, a true Dirac delta-function is used in these priors. The benefits of this approach include the explicit selection of variables and the ability to average over different model structures; a reduction of computational complexity for increasingly sparse models; and adaptive shrinkage which can leave untouched variables with strong support in the data, but strongly regularise those of more dubious merit. The downside is the scaling of the problem. For a regression model with k variables, or a network matrix with p nodes and $k = \frac{p(p-1)}{2}$ links, there are 2^k possible models to consider. For example, a network with 30 nodes has about 10^{131} possible models! Variational models do exist for some continuous prior formulations, but will be inapplicable to the non-conjugate hierarchies we consider. MCMC samplers, then, are the tool of choice for approximate inference, but come

with a health warning. The scale of the inference problem will likely be so large that 'the sampler is used to search for promising models, rather than compute the whole posterior' (George & McCulloch, 1997). Converging rapidly to a set of high-probability states is, however, a task at which MCMC chains can excel. We could not find a review that adequately covers past work on this problem, but Hahn & Carvalho (2015), O'Hara & Sillanpää (2009) and Mallick & Yi (2014) are good places to start.

Early samplers used Gibbs updates to test for the inclusion of each variable in turn (Geweke, 1996; Kuo & Mallick, 1998), but it was rapidly shown that Metropolis proposals are more efficient (George & McCulloch, 1997; Raftery *et al.*, 1997). A range of conditionally-conjugate prior structures are available that permit marginalisation over the regression coefficients, which allows an MH sampler to explore just the model space $p(\boldsymbol{\gamma}|\boldsymbol{Y})$, with no subsidiary variables. This does seem to be the most efficient approach. Of the choice of jumping rules to move over the set of indicator variables $\boldsymbol{\gamma}$, multi-step moves or variable swaps outperform single-variable updates (Guan & Stephens, 2011; Kohn *et al.*, 2001; Peltola *et al.*, 2012). Further improvements are possible if the proposal distribution for $\boldsymbol{\gamma}$ can learn from its past. Various adaptive samplers exploit this property (Casella & Moreno, 2006; Nott & Kohn, 2005), and some of the most advanced methods are discussed in Peltola *et al.* (2012), which targets massive regression problems on genome-wide association study data. If marginalisation of lower-level variables is not possible, reversible-jump MCMC samplers can be used to move through spaces of changing dimensionality (Dellaportas & Forster, 1999; Denison *et al.*, 1998; Green, 1995).

While difficult to apply to the problem of sparse precision modelling, it is worth noting that schemes for sampling without replacement (Clyde *et al.*,

2011) and adaptations of the renowned Swendsen-Wang algorithm for Potts models (Edwards & Sokal, 1988; Nott & Green, 2004; Swendsen & Wang, 1987) have also been successful for linear regression. Approaches for performing exact Hamiltonian Monte Carlo (HMC) (without the need for leapfrog integration) are available, but are inefficient unless there are strong additional constraints in the prior structure (Pakman & Paninski, 2013, 2014).

It is good to bear in mind that with sufficient data, spike and slab models (and indeed, any prior) can become dominated by the likelihood and infer very small, but non-zero, effects for some parameters. The rejection of models where variables are exactly zero is not necessarily problematic—and there is no need to introduce priors that scale with the quantity of data, as in Ishwaran & Rau (2005).

Sparse modelling methods for precision matrices have focussed on the *G*-Wishart prior (Letac & Massam, 2007; Roverato, 2002). The Wishart distribution is a multivariate extension of the gamma distribution, and is the conjugate prior for the precision matrix in the multivariate normal. The *G*-Wishart constrains the elements of the precision matrix to be supported by a graph, $\mathcal{G}$, which encodes the sparsity structure of the matrix or network. Monte-Carlo solutions for joint inference of the distribution and its graph $\mathcal{G}$ do exist (Atay-Kayis & Massam, 2005; Mohammadi & Wit, 2015; Wang, 2012*b*), and the distribution itself can be sampled from relatively efficiently for a given graph structure (Lenkoski, 2013). The most efficient inference implementation for the graphical search uses Gibbs sampling, testing each off-diagonal element in turn for inclusion in the model (Hinne *et al.*, 2014). This procedure has been applied to low-dimensional subcortical fMRI data (Hinne *et al.*, 2015), but becomes excessively slow for large model sizes (greater than 30 nodes), even for one subject. The major diffi-

culty in building hierarchical priors over several subjects using the *G*-Wishart distribution is that the normalising constant for the distribution is difficult to approximate and depends on the graph, $\mathcal{G}$.[4] This makes a hierarchical formulation *at least* doubly intractable. (Intractable distributions, for which the normalising constant is not known, can be sampled from using conventional Monte Carlo methods. Distributions for which the denominator is not constant, but depends on other variables in the hierarchy, are called *doubly intractable,* and much harder to tackle. The *G*-Wishart distribution itself falls into this category, and the various methods discussed above use tailored versions of trans-dimensional MCMC algorithms to sample from the distribution. See Green (2003) and Murray *et al.* (2006) for general discussions of how to tackle this problem.)

COVARIANCE ESTIMATION FOR GROUPS OF SUBJECTS

As the Wishart distribution is conjugate to the precision matrix in a multivariate normal distribution, it is unsurprising that Bayesian models with hierarchies of Wishart and Inverse-Wishart distributions have been designed. However, they give up on any attempt to impose sparsity in the system (Harrison *et al.*, 2015; Marrelec *et al.*, 2006).

A number of multiple-class models have been formulated as optimisation problems (Danaher *et al.*, 2015; Guo *et al.*, 2011; Lee & Liu, 2015; Yang *et al.*, 2015), and have been applied to fMRI data (Liang *et al.*, 2016; Varoquaux *et al.*, 2010). They generally use a mixture of L_1 and L_2 norms to suppress all edges in the model, over all subjects, and to encourage similarity between subjects' network matrices. However, all suffer from the usual criticisms of the need for nested

[4] Alex Lenkoski *et al.* have recently released, in preprint, an exact method for computing the normalisation constant, although it is computationally expensive (Uhler *et al.*, 2016); a Laplace approximation does exist as an alternative (Lenkoski & Dobra, 2011).

cross-validation loops to tune the optimisation parameters, and an inability to characterise the uncertainty in the model inference.

5.3 Building a hierarchical covariance model

We will discuss a broad class of priors for precision matrix estimation for a single subject, and demonstrate that it can be reformulated as a linear regression problem, with access to all of the existing priors and inference engines from the Bayesian linear regression literature. The extension to multi-subject models will be described by specifically outlining our proposed prior structure. The objective of the prior is to provide a simultaneous description of the connectivities in individual subjects and their population average, under a set of sparse penalties. The structure is designed to allow efficient inference even on large problems. Full details of our inference program and MCMC algorithm for this model, denoted HIPPO (Hierarchical Inference of Posterior Precisions in OSL), are given in appendix B. We start by setting out the noise model for activation data within a particular ROI.

GENERATIVE MODEL FOR RESTING-STATE ACTIVATION

Description of activation within regions of interest

We are interested in describing the connectivity between several regions of interest over the brain. Given temporally demeaned activations $\boldsymbol{X} \in \mathbb{R}^{v \times n_s}$ in each of N subjects s, defined over a grid of v voxels or points on a cortical surface, and sampled at n_s time points, we can estimate a single time course for

each of p ROIs using a spatial basis set, a linear weighting over voxels,

$$\boldsymbol{X}_s = \boldsymbol{A}\boldsymbol{Y}_s + \boldsymbol{E}; \quad \boldsymbol{E} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma}_e). \tag{5.2}$$

Here, $\boldsymbol{Y}_s = \left[\boldsymbol{y}_1^s, \boldsymbol{y}_2^s, \ldots \boldsymbol{y}_{n_s}^s\right] \in \mathbb{R}^{p \times n_s}$ represents the activation within ROIs, $\boldsymbol{A} \in \mathbb{R}^{v \times p}$ is the spatial basis set (which may simply be binary region allocations, or, for example, a spatial map derived from ICA), and $\boldsymbol{E}$ represents additive noise with covariance structure $\boldsymbol{\Sigma}_e$ over voxels.

We use the general term *activation* to encompass changes in BOLD response over time in fMRI, or fluctuations in the power envelope of oscillatory activity measured with MEG.

Likelihood for the connectivity model

We describe the (temporally demeaned) data within ROIs as being drawn independently from a multivariate Gaussian distribution with zero mean and precision matrix $\boldsymbol{\Omega}_s$

$$\boldsymbol{Y}_s \sim \mathcal{N}\left(\boldsymbol{0}, \boldsymbol{\Omega}_s^{-1}\right), \ \boldsymbol{\Omega}_s \in \mathbb{P}_G, \tag{5.3}$$

where $\mathbb{P}_G$ is the cone of positive definite $p \times p$ matrices restricted to the graph $G = (V, E)$ such that an absence of an edge from set E implies conditional independence of the two relevant variables in all subjects,

$$(i, j) \notin E \implies \omega_{ij}^s = 0 \ \forall \, s.$$

PRECISION MODELLING AS LINKED LINEAR REGRESSION

Wang describes, in 2012*a* and 2015, two different prior structures that allow, after a variable substitution, simple block-Gibbs sampling along the columns of precision matrices to draw from the posterior of particular sparse models for covariance in single datasets. We build on his approach and demonstrate that a very broad range of priors from the linear regression literature can be applied to the elements of a precision matrix, with a simple restriction on the prior for the diagonal elements.

It is worth introducing a bit of notation. Without loss of generality, we take just the final column of precision matrix $\boldsymbol{\Omega}_s$, $\left[\boldsymbol{\omega}_{12}^s, \omega_{22}^s\right]^T$. Let $\boldsymbol{\Omega}_{11}^s$ represent the first principal minor (the block matrix without the final row or column) and $\boldsymbol{S} = \boldsymbol{Y}_s^\dagger \boldsymbol{Y}_s = n_s \boldsymbol{\Sigma}_s$ be the sample inner product matrices. Similar subscripts indicate identical partitions of other matrices, so S_{22} is the diagonal element of the selected column and $\boldsymbol{S}_{12}$ the off-diagonal elements of the column. This is the same convention employed in Friedman *et al.* (2008) and Wang (2012*a*).

We define independent exponential prior distributions on the diagonal elements of the precision matrices, and require the priors on the off-diagonal elements to factorise over the elements (although they may share some common hyper-parameters, $\boldsymbol{\theta}$),

$$\pi\left(\boldsymbol{\Omega}_s\right) = \prod_{i=1}^{p} \operatorname{Exp}\left(\omega_{ii}^s;\ \frac{\lambda^s}{2}\right) \prod_{i<j} \pi\left(\omega_{ij};\ \boldsymbol{\theta}\right), \tag{5.4}$$

using $\pi\left(\cdot\right)$ to denote a prior probability density. Combining (5.3) and (5.4), we

can extract the conditional posterior for a column of the precision matrix,

$$p\left(\boldsymbol{\omega}_{12}^{s}, \omega_{22}^{s}|-\right) \propto \frac{1}{(2\pi)^{n_s/2}} \left(\omega_{22}^{s} - \boldsymbol{\omega}_{12}^{s\,\dagger} \boldsymbol{\Omega}_{11}^{s\,-1} \boldsymbol{\omega}_{12}^{s}\right)^{n_s/2} |\boldsymbol{\Omega}_{11}^{s\,-1}|^{-n_s/2} \tag{5.5}$$
$$\times \exp\left(-\frac{1}{2}\left[S_{22}\omega_{22}^{s} + 2\boldsymbol{S}_{12}^{\dagger}\boldsymbol{\omega}_{12}^{s}\right]\right) \text{Exp}\left(\omega_{22}^{s};\, \frac{\lambda^s}{2}\right) \pi\left(\boldsymbol{\omega}_{12}^{s};\, \boldsymbol{\theta}\right) .$$

Performing the variable substitution (Wang, 2012*a*)

$$(\boldsymbol{u}^s,\, v^s) = \left(\boldsymbol{\omega}_{12}^{s},\, \omega_{22}^{s} - \boldsymbol{\omega}_{12}^{s\,\dagger} \boldsymbol{\Omega}_{11}^{s\,-1} \boldsymbol{\omega}_{12}^{s}\right) , \tag{5.6}$$

for which the Jacobian is the identity matrix, we obtain

$$\log p\left(\boldsymbol{u}^s, v^s|-\right) = \frac{n_s}{2} \log v^s - \frac{1}{2}\left(S_{22} + \lambda^s\right) v^s \tag{5.7}$$
$$- \frac{1}{2}\boldsymbol{u}^{s\dagger}\left[\left(S_{22} + \lambda^s\right)\boldsymbol{\Omega}_{11}^{s\,-1}\right]\boldsymbol{u}^s - \boldsymbol{S}_{12}^{\dagger}\boldsymbol{u}^s + \log \pi\left(\boldsymbol{u}^s;\, \boldsymbol{\theta}\right) + const.$$

We can tidy up with the substitution $\boldsymbol{\Upsilon}^s = \left(s_{22} + \lambda^s\right)\boldsymbol{\Omega}_{11}^{s\,-1}$ to give a normal form for $\boldsymbol{u}^s$ and a Gamma distribution on v^s,

$$\log p\left(\boldsymbol{u}^s|-\right) = -\frac{1}{2}\boldsymbol{u}^{s\dagger}\boldsymbol{\Upsilon}^s\boldsymbol{u}^s - \boldsymbol{S}_{12}^{\dagger}\boldsymbol{u}^s + \log \pi\left(\boldsymbol{u}^s;\, \boldsymbol{\theta}\right) - \frac{n_s}{2}\log \pi + const. \tag{5.8}$$

$$\log p\left(v^s|-\right) = \frac{n_s}{2}\log v^s - \frac{1}{2}\left(S_{22} + \lambda^s\right) v^s + const. \tag{5.9}$$

We draw the comparison to conventional linear regression, $\boldsymbol{y} = \boldsymbol{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}$, $\boldsymbol{\epsilon} \sim$

$\mathcal{N}\left(0, \sigma^2\right)$, conditional on σ^2,

$$
\begin{aligned}
p\left(\boldsymbol{\beta}|\boldsymbol{y}, X, \sigma^2\right) &\propto \frac{1}{\left(2\pi\sigma^2\right)^{nsp/2}} e^{-\frac{1}{2}\boldsymbol{\beta}^\dagger \boldsymbol{\Upsilon} \boldsymbol{\beta} + \boldsymbol{r}^\dagger \boldsymbol{\beta}} \; \pi\left(\boldsymbol{\beta};\, \boldsymbol{\theta}\right) \\
\boldsymbol{\Upsilon} &= X^\dagger X / \sigma^2 \\
\boldsymbol{r} &= \boldsymbol{y}^\dagger X / \sigma^2,
\end{aligned}
$$

to see that this factorisation into column-conditionals leads to inference as a set of p linked regressions on one variable and its interactions, given all the other variables. We are free to choose any factorisable prior on the off-diagonal elements of the precision matrix and retain the simple block-Gibbs sampling scheme (5.8) and (5.9). Choices of priors from the scale mixture of normals family (Gelman, 2006; Gelman *et al.*, 2008; Polson & Scott, 2012) will allow simple draws from multivariate normal distributions. Moreover, this framework allows for easy expansion to a prior structure that defines a hierarchy over subjects, as we shall see below.

HIERARCHICAL PRIOR FOR PRECISION MATRICES

For each subject, we place an exponential prior on the diagonal elements of the precision matrix. This choice allows us to implement the column-wise sampling scheme described in equations (5.4) to (5.9). The free parameter in this distribution is given (for each subject) the 'neutral' Gamma conjugate hyperprior (Kerman, 2011). There is normally plenty of information with which to estimate the diagonal elements (inverse variances), and so an uninformative prior is appropriate. The neutral Gamma prior is relatively uninformative on $\log \lambda^s$ and is

claimed to perform better than traditional $\mathrm{Ga}(\epsilon, \epsilon)$ priors (*ibid.*),

$$\lambda^s \sim \mathrm{Ga}(1/3, 0)\,. \tag{5.10}$$

The full prior on the off-diagonal elements is given in equation (5.11), discussed in detail below, and is illustrated in figure 5.1 on page 142. In essence, it is a spike-and-slab prior with shared sparsity over subjects, normally-distributed connection strengths about the population mean, and mild regularisation on the mean effect.

$$\begin{aligned}
\left(\omega_{ij}^s | \sigma_{ij}, \mu_{ij}, z_{ij} = 1\right) &\sim \mathcal{N}\left(\mu_{ij}, \sigma_{ij}^2\right) \\
\left(\omega_{ij}^s | z_{ij} = 0\right) &\sim \delta_0 \\
\log \sigma_{ij} &\sim \mathcal{N}\left(\log m_\sigma, s_\sigma^2\right) \\
\mu_{ij} &\sim \mathcal{N}\left(0, \chi^2\right) \\
\chi &\sim \mathcal{C}^+(0, A) \\
z_{ij} &\sim \mathrm{Bernoulli}(a) \\
a &\sim \mathrm{Beta}(a_\pi, b_\pi)
\end{aligned} \tag{5.11}$$

Each off-diagonal element of the precision matrices is given a spike and non-central slab prior. The spike, a Dirac delta-function at zero, imposes a shared sparsity structure over all subjects, using edge inclusion variables $\boldsymbol{Z}$. For those edges that are included, the non-central slab is a normal distribution, whose variance σ_{ij} characterises the between-subject variability of that particular connection strength, and whose mean μ_{ij} captures the group-level behaviour.

Following advice in Gelman (2006) and Polson & Scott (2012) on the infer-

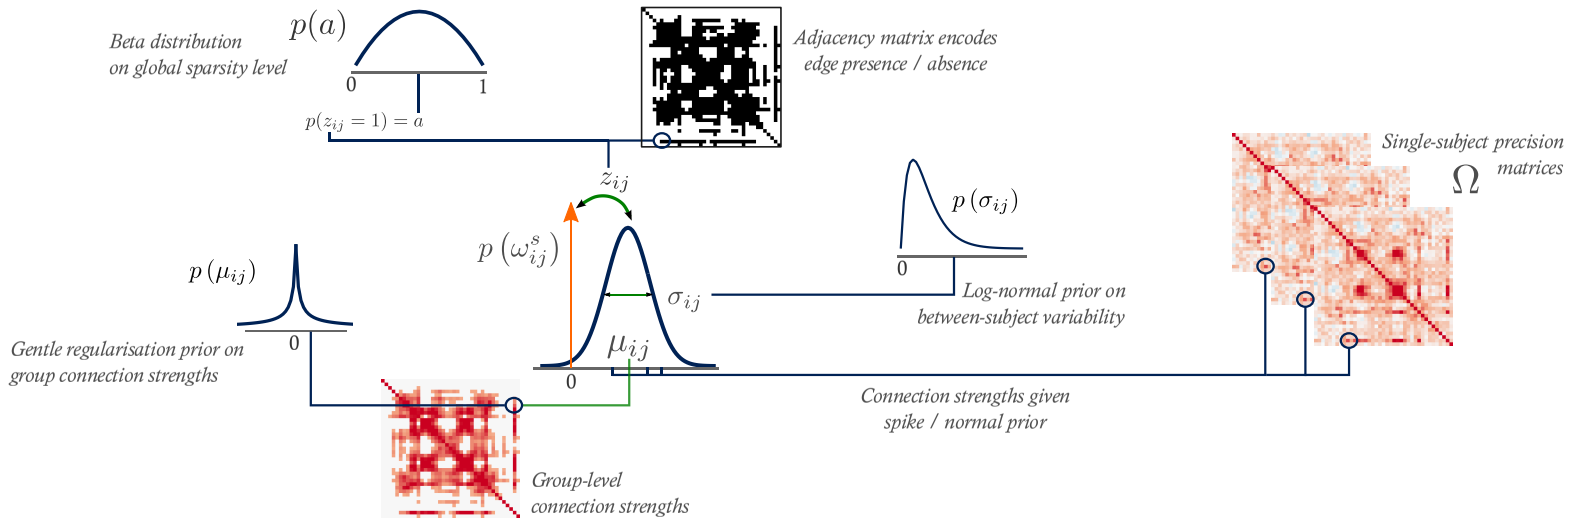


Figure 5.1: HIERARCHICAL PRIOR ON PRECISION MATRICES. A spike and non-central slab prior is placed on each off-diagonal element of the precision matrices. Selection of the slab or spike (presence or absence of a network connection) is controlled by an adjacency matrix, with a learnt sparsity level. The slab is modelled as a normal distribution describing the mean and variance over connection strengths. The mean is regularised by a sparsity-inducing prior, and the variance by a weakly-informative log-normal prior. We call this model and its inference scheme Hierarchical Inference of Posterior Precisions in OSL (HIPPO).

ence of higher-level group parameters in regression, we apply gentle regularisation to the mean edge strengths towards zero using a normal distribution, learning the rough scale of these connectivities, χ^2, from the data and pooling this information over all the edges. This learnt variance parameter should be constrained with a proper, weakly-informative prior, because such parameters can be difficult to estimate with few subjects or when they are close to zero. Gelman and Polson & Scott recommend the half-Cauchy distribution for this application, denoted $\mathcal{C}^+$, because it has a finite mass at zero, is undefined for negative values, drops off near a scaling point A, has a heavy tail which can allow the likelihood to dominate, and can be expressed in a convenient conditionally-conjugate fashion. The scale of the Cauchy distribution can be set sensibly based on the data: we use $A = 0.7$ as an appropriate value for variance-scaled data where correlations and partial correlations do not frequently exceed this number. The resultant prior on the group connection strengths $\boldsymbol{\mu}$ encourages greater sparsity (greater regularisation towards zero) than a simply-parametrised normal distribution or normal-inverse-gamma distribution (Carvalho *et al.*, 2010; Peltola *et al.*, 2012).

A broad normal prior is placed on the logarithm of σ_{ij}, centred on $m_\sigma = 0.5$ and with a standard deviation $s_\sigma = 1$ to allow order-of-magnitude deviations from this value. Finally, a Beta-Bernoulli conjugate prior is placed on the edge inclusion variables, with a shared sparsity parameter a inferred from the data. The hyperparameters in the top-level Beta distribution can be set to weakly encourage levels of sparsity encountered in functional networks. Using values of $(a_\pi, b_\pi) = 6$ places most of the prior mass between 0.3 and 0.7.

The values of the hyperparameters we use are set out in table 5.1.

Table 5.1: VALUES OF HYPERPARAMETERS employed for functional network modelling.

Parameter	Value
m_σ	0.5
s_σ	1
A	0.7
a_π	6
b_π	6

A weakly-sparse prior

The explicit sparsity can be removed from the model by setting $a = 1$ and $z_{ij} = 1 \, \forall \, i, j$ to create a *weakly-sparse* prior. Inference will be performed in exactly the same fashion as for the fully-sparse model, but without the need for updates on $\boldsymbol{Z}$. Sparsity is encouraged through the distribution on the group mean terms $\boldsymbol{\mu}$, and the modelling of a distribution for subject-level connection strengths about these mean terms.

Expanding the Cauchy distribution

The prior on μ_{ij} is complex, but can be re-written as a scale mixture of normals using a simple parameter expansion. Usefully, this change of variables creates a conditionally-conjugate prior structure, simplifying the inference procedure and permitting marginalisation over μ_{ij} to collapse part of the Gibbs sampler (van Dyk & Park, 2008). Following the steps in Gelman (2006); Gelman *et al.*

(2014, 2008); and Polson & Scott (2012), we obtain

$$\begin{aligned} \mu_{ij} &= \xi \eta_{ij} \\ \eta_{ij} &\sim \mathcal{N}\left(0, \sigma_\eta^2\right) \\ \sigma_\eta^2 &\sim \text{Inv-}\chi^2\left(1, A\right) \\ \xi &\sim \mathcal{N}\left(0, 1\right) \\ \chi &= |\xi| \sigma_\eta \, . \end{aligned} \tag{5.12}$$

Redundant parameter expansion in this manner was originally designed to improve the convergence of the Expectation Maximisation algorithm (Liu *et al.*, 1998), then used to increase the movement available to Gibbs samplers (Gelman *et al.*, 2008). It has since been used to create a wide range of flexible priors, in addition to that used here (Carvalho *et al.*, 2010; Gelman *et al.*, 2014; Peltola *et al.*, 2012). It is not necessary, when sampling, for ξ and $\boldsymbol{\eta}$ to converge, only for the meaningful parameters χ and $\boldsymbol{\mu}$.

JOINT PROBABILITY DISTRIBUTION

We are now in a position to write out the full joint probability distribution of the parameters and data,

$$
\begin{aligned}
p\left(\{\boldsymbol{\Omega}_s\}, \boldsymbol{Z}, \boldsymbol{\mu}, \{\lambda^s\}, \boldsymbol{\sigma}, \chi, a, \{\boldsymbol{Y}_s\}\right) \propto \prod_s^N \Bigg\{ & \frac{|\boldsymbol{\Omega}_s|^{\frac{n_s}{2}}}{(2\pi)^{\frac{pn_s}{2}}} e^{-\frac{1}{2}\,\mathrm{trace}[\boldsymbol{S}\boldsymbol{\Omega}_s]} \prod_{i=1}^p \frac{\lambda^s}{2} e^{-\frac{\lambda^s}{2}\omega_{ii}^s} \\
& \times \prod_{i<j} \Bigg\{ \left(\frac{1}{\left(2\pi\sigma_{ij}^2\right)^{1/2}} e^{-\frac{\left(\omega_{ij}^s - \xi\eta_{ij}\right)^2}{2\sigma_{ij}^2}} \right)^{z_{ij}} \\
& \quad \times \delta\left(\omega_{ij}^s\right)^{(1-z_{ij})} \Bigg\} \\
& \times \frac{{\lambda^s}^{1/3-1}}{\Gamma\left(1/3\right)} \Bigg\} \\
& \times \prod_{i<j} \left\{ a^{z_{ij}} \left(1-a\right)^{(1-z_{ij})} \frac{1}{\left(2\pi\sigma_\eta^2\right)^{\frac{1}{2}}} e^{-\frac{\eta_{ij}^2}{2\sigma_\eta^2}} \right\} \\
& \times \prod_{i<j} \left\{ \frac{1}{\sigma_{ij}\left(2\pi s_\sigma^2\right)^{\frac{1}{2}}} e^{-\frac{\left(\log\sigma_{ij} - \log m_\sigma\right)^2}{2s_\sigma^2}} \right\} \\
& \times \frac{a^{a_\pi - 1}\left(1-a\right)^{b_\pi - 1}}{\mathrm{B}\left(a_\pi, b_\pi\right)} \frac{1}{\sqrt{2\pi}} e^{-\xi^2} \frac{\left(A/2\right)^{1/2}}{\Gamma\left(1/2\right)} \frac{e^{-\frac{A}{2\sigma_\eta^2}}}{{\sigma_\eta^2}^{1+\frac{1}{2}}}\,.
\end{aligned}
\tag{5.13}
$$

MODEL INFERENCE

The procedure taken for inference on the HIPPO model, and the Metropolis-within-collapsed-Gibbs (MWCG) sampling scheme, is described in full in appendix B. First, time courses are estimated for each ROI. Second, the hierarchical model is fitted to these data to model the sparse covariance structure simultaneously over all subjects. The sampling scheme builds on the ideas of Wang (2012*a*) and its extension in equation (5.8), sampling linked linear regressions

in each consecutive column of the precision matrices. Within each column, the sampler exploits the progress made in MCMC algorithms for variable selection in linear regression, building largely on Peltola *et al.* (2012). An MH sampler checks for variables to add or remove by the testing model evidence of each proposal: this is a Bayesian model averaging scheme at the top level. Parameters describing group-average and subject-level connectivities are conditionally conjugate on the network structure, and can be sampled directly.

5.4 Methods

SIMULATED DATA

Data generation

Ten simulated datasets were created to test a range of sparse network structures, model sizes, quantities of data, and amount of subject variability in connectivity. Except for some minor differences detailed below, each dataset consists of a number of subjects, with an individual precision matrix to represent each subject's functional connectivity; the data for each subject is a draw of samples from a zero-mean multivariate normal distribution, using the appropriate precision matrix. The general properties of each simulation are summarised in table 5.2 on page 148.

Circle models Simulations 1–4, 7 and 8 use a simple circular network structure. This was set out in Wang (2012*b*), and used in Hinne *et al.* (2015); Mohammadi & Wit (2015); and Wang (2012*a*). A precision matrix of any dimension p is constructed as $\omega_{ii} = 1$, $\omega_{i,i+1} = 0.5$, and $\omega_{1,p} = 0.4$, with the lower diagonal

Table 5.2: DESCRIPTION OF SIMULATED DATASETS used in figure 5.3 on page 160. Datasets are characterised by their size (number of subjects, network nodes and links, and data samples), the amount of subject variability (expressed as the standard deviation of connection strengths (over subjects) divided by the mean connection strength; we take the mean of this coefficient of variation over all connections present in the network), the sparsity of the network, and the type of network structure. We use simple circle models of varying sizes, first set out in Wang (2012*b*) and used in Wang (2012*a*) and Hinne *et al.* (2015), together with the largest simulation (netsim 4) from Smith *et al.* (2011), a random graph structure and connection matrices built on estimates of the networks in cat cortex (Scannell *et al.*, 1999) and macaque visual cortex (Felleman & Van Essen, 1991).

ID	Subjects	Nodes	Links	Samples	Variability	Sparsity	Model Structure
1	5	6	15	18	0	0.60	Circle
2	25	6	15	18	0	0.60	Circle
3	25	6	15	50	0	0.60	Circle
4	25	6	15	100	0	0.60	Circle
5	50	50	1225	200	0	0.95	Netsim 4
6	25	25	300	500	0.2	0.28	Random
7	25	6	15	25	0.5	0.60	Circle
8	25	30	435	100	1.1	0.93	Circle
9	30	52	1326	500	1.8	0.67	Cat cortex
10	25	30	435	100	2.7	0.63	Macaque Visual cortex

elements matched to ensure symmetry.

Simulations 1–4 use the same precision matrix for each subject. Simulations 7 and 8 treat the circle model as the group mean matrix, assign random signs to the connection strengths, and draw single-subject connection strengths from a normal distribution about the group mean with a standard deviation of 0.05 and 0.15 respectively. Subjects' connectivity matrices were then adjusted to be positive semi-definite (Higham, 1988)[5] and rescaled to unit variance.

Netsim Dataset 5 is the 4^{th} network simulation from Smith *et al.* (2011), chosen because the 4^{th} simulation was the largest model in that work. Smith *et al.*

[5] The code used to perform this adjustment, `nearestSPM.m` was written by John D'Errico and is available from `uk.mathworks.com/matlabcentral/fileexchange/42885-nearestspd`.

generated autocorrelated fMRI-like data from an asymmetric network model. As we are estimating precision matrices, we sought instead a symmetric ground truth against which to compare results, and as the individual network matrices used in the simulation are no longer available (each subject had a slightly different set of network strengths), we make the simplifying assumption that all subjects share the same symmetric precision matrix. We estimate this as the unregularised partial correlation computed from the entire concatenated dataset, masked by the symmetrised adjacency matrix of the original simulation.

Random Dataset 6 is designed to mimic the prior structure. Each network edge is randomly assigned a probability of 0, 0.5 or 1 of being included in the network. Subjects' connection strengths on each edge are normally distributed about 0.25 with a standard deviation of 0.05. As each data point is drawn for each subject, edges in the model are turned on or off in accordance with their probability of edge inclusion. The 'true' matrix for each subject is the product of the edge inclusion probability matrix and the subject's connection strengths.

Cat cortex and macaque visual cortex Simulations 9 and 10 use the structures of mammalian cortical networks as their basis. Simulation 9 employs the cat cortical network from Scannell *et al.* (1999) and simulation 10 that of the macaque visual cortex from Felleman & Van Essen (1991).[6] In each case, the network connection matrix is binary and asymmetrical. We used only the upper triangular part, symmetrising to the lower half. For each simulation, the group mean connection matrix was drawn from a *G*-Wishart distribution (Lenkoski, 2013) with identity scale matrix and degrees of freedom equal to one less than the

[6] Both of these matrices are available from the Brain Connectivity Toolbox, at `sites.google.com/site/bctnet/datasets`.

Table 5.3: CHARACTERISATION OF METHODS used on simulated data in figure 5.3 on page 160. We classify the methods under test by their inference method, the style of sparsity imposition, and whether they are fitted to individual subjects, the concatenated dataset, or infer individual connectivity matrices using information from the whole group. HIPPO is our acronym for the hierarchical sparse Bayesian model.

Name	Fitted to	Sparsity	Inference
Partial correlation	individuals	none	analytic
Tikhonov	individuals	continuous	optimisation
GLASSO	individuals	continuous	optimisation
Group GLASSO	group	continuous	optimisation
Wishart	individuals	continuous	analytic
Bayesian GLASSO	individuals	continuous	MCMC
SSVS	individuals	normal-mixture	MCMC
G-Wishart	concatenation	spike & slab	MCMC
single-subject HIPPO	concatenation	spike & slab	MCMC
weakly-sparse HIPPO	group	continuous	MCMC
HIPPO	group	spike & slab	MCMC

number of network nodes. Individual subjects' network matrices were drawn from a *G*-Wishart distribution using the mean connection matrix as the scale, and degrees of freedom set to 50 for the simulations from macaque visual cortex and 400 for those from cat cortex. The *G*-Wishart distributions were constrained using the relevant network matrices (cat, macaque) as the underlying graph.

Models tested

Twelve different covariance models were fitted to the test datasets, representing the range of methods in use from the most naïve to the most advanced. Their basic properties are set out in table 5.3 on page 150.

Partial correlation The sample covariance matrix for each subject is inverted, using the Cholesky algorithm, and normalised to produce the unregularised partial correlation matrix.

Tikhonov A Tikhonov-regularised estimate of the precision matrix is constructed by slightly augmenting the diagonal of the sample covariance matrix,

$$\hat{\boldsymbol{\Omega}} = (\boldsymbol{\Sigma} + \lambda \boldsymbol{I})^{-1} \ . \tag{5.14}$$

The regularisation parameter λ was chosen to minimise the RMS distance between the subjects' matrices and their unregularised group average.

Glasso The graphical LASSO algorithm of Friedman *et al.* (2008), with modifications for computational efficiency (Mazumder & Hastie, 2012*a*,*b*), solves the optimisation problem

$$\hat{\boldsymbol{\Omega}} = \arg\max_{\boldsymbol{\Omega}} \left(\log\det \boldsymbol{\Omega} - \operatorname{trace}\left(\boldsymbol{\Sigma}\boldsymbol{\Omega}\right) - \lambda ||\boldsymbol{\Omega}||_1 \right) \ , \tag{5.15}$$

where $||\cdot||_1$ indicates the L_1 norm (sum of the absolute values of the elements of the matrix). The regularisation parameter λ was chosen to minimise the RMS distance between the subjects' matrices and their unregularised group average. It is safe to say that GLASSO is the most common method for estimating partial correlation brain networks, and one of the most popular of any functional network estimation method in FMRI, making it a good benchmark for our work.

Group glasso The group graphical LASSO of Danaher *et al.* (2015) is fitted to all subjects at once, and imposes sparsity on the mean connectivity matrix. It

solves the optimisation problem,[7]

$$\left\{\hat{\boldsymbol{\Omega}}^s\right\} = \underset{\{\boldsymbol{\Omega}\}}{\arg\max}\left(\sum_s \left[n_s \log\det \boldsymbol{\Omega}^s - \text{trace}\left(\boldsymbol{\Sigma}^s\boldsymbol{\Omega}^s\right)\right] - \lambda\sum_s\sum_{i\neq j}||\omega_{ij}^s||_1\right). \tag{5.16}$$

The regularisation parameter was chosen to minimise the Bayesian information criterion associated with this likelihood.

Fused glasso The fused graphical LASSO of Danaher *et al.* (2015) is also fitted to all subjects at once, and seeks to impose collective sparsity on all subjects' network elements, while encouraging networks from different subjects to be alike. It is therefore the closest existing non-Bayesian approach to the models proposed in this chapter, and the regularisation terms could be viewed as a type of hierarchical prior. Inference is set up as an optimisation problem with two penalty terms,[8]

$$\left\{\hat{\boldsymbol{\Omega}}^s\right\} = \underset{\{\boldsymbol{\Omega}\}}{\arg\max}\left(\sum_s \left[n_s \log\det \boldsymbol{\Omega}^s - \text{trace}\left(\boldsymbol{\Sigma}^s\boldsymbol{\Omega}^s\right)\right] - \lambda_1\sum_s\sum_{i\neq j}||\omega_{ij}^s||_1 + \lambda_2\sum_{s<s'}\sum_{i,j}||\omega_{ij}^s - \omega_{ij}^{s'}||_1\right). \tag{5.17}$$

The regularisation parameters were chosen to minimise the difference between the group mean connectivity inferred using half of the dataset and the unregularised mean of the full dataset.

[7] The code for solving this problem using the alternating direction method of multipliers algorithm was provided by Pedro Ariel Rojas López, from CNEURO, Havana.

[8] R code for solving this problem using the alternating direction method of multipliers algorithm is available as the package `JGL`, from `cran.r-project.org/web/packages/JGL/index.html`.

Wishart Following Gelman *et al.* (2014), a simple Wishart prior distribution is placed independently over each subject's precision matrix, $\boldsymbol{\Omega} \sim \mathcal{W}_p\left(p+1,\, \frac{1}{p+1}\boldsymbol{I}\right)$. This leads analytically to the posterior for each subject, which we summarise by its expectation.[9] The similarity to Tikhonov regularisation is clear.

$$p\left(\boldsymbol{\Omega}^s|\boldsymbol{S}\right) = \mathcal{W}_p\left(\boldsymbol{\Omega}^s;\, p+1+n_s, \left(\left(p+1\right)\boldsymbol{I}+\boldsymbol{S}\right)^{-1}\right) . \tag{5.18}$$

Bayesian glasso The Bayesian graphical LASSO of Wang (2012*a*) places a Laplace or double-exponential prior on the off-diagonal elements of the precision matrix,

$$\pi\left(\omega_{ij}\right) = \frac{\lambda}{2}\exp\left(-\lambda|\omega_{ij}|\right) . \tag{5.19}$$

For each subject, 3000 samples were drawn using Wang's algorithm,[10] after discarding 1000 as warm-up.

SSVS The Bayesian Stochastic Search Variable Selection algorithm (Wang, 2015) for covariance selection places a mixture of normal priors on the off-diagonal elements of the precision matrix,

$$\pi\left(\omega_{ij}\right) = \left(1-a\right)\mathcal{N}\left(\omega_{ij};\, 0, v_0^2\right) + a\,\mathcal{N}\left(\omega_{ij};\, 0, v_1^2\right) , \tag{5.20}$$

where v_0 is chosen to be much smaller than v_1 and sampling proceeds using Gibbs sampling along the columns. We follow Wang's recommendations and set v_0 to 0.05, v_1 to 2.5 and a to 0.5. For each subject, 3000 samples were drawn after 1000 warm-up samples.[11]

[9] If $\boldsymbol{\Omega}$ is (G-) Wishart distributed, $\boldsymbol{\Omega} \sim \mathcal{W}_\mathcal{G}\left(\delta, \boldsymbol{V}\right)$, then its expectation is $\langle\boldsymbol{\Omega}\rangle = \delta\boldsymbol{V}$.
[10] `BayesGlasso` is available from `msu.edu/~haowang/RESEARCH/Bglasso/bglasso.html`.
[11] `BayesGGM_SSVS` is available from `msu.edu/~haowang/`.

G-Wishart The *G*-Wishart distribution is the conjugate prior on the multivariate normal, describing both the precision matrix and the graph $\mathcal{G}$ on which it is supported,

$$\begin{aligned} \pi\left(\boldsymbol{\Omega}|\mathcal{G}, \delta, \boldsymbol{V}\right) &= \mathcal{W}_{\mathcal{G}}\left(\delta, \boldsymbol{V}\right) \\ &= \frac{|\boldsymbol{\Omega}|^{\delta-p-1/2}}{Z_{\mathcal{G}}\left(\delta, \boldsymbol{V}\right)} \exp\left[-\frac{1}{2}\operatorname{trace}\left(\boldsymbol{V}^{-1}\boldsymbol{\Omega}\right)\right] \mathbb{1}_{\boldsymbol{\Omega}\in\mathbb{P}_{\mathcal{G}}}, \end{aligned} \tag{5.21}$$

where $\boldsymbol{V}$ is the scale matrix, δ indicates the degrees of freedom, $Z_{\mathcal{G}}$ is the intractable normalising constant, and $\boldsymbol{\Omega}$ is constrained to live on the cone $\mathbb{P}_{\mathcal{G}}$ of positive definite $p \times p$ matrices with zeros indicated by the graph $\mathcal{G}$.

We fitted this model to the entire dataset concatenated over subjects, using software provided by Hinne *et al.*,[12] using 5000 warm-up samples and 10 000 draws from the distribution. Fitting the model to the concatenated data provides a useful comparison to the hierarchical models, using all of the data for inference, but assuming that each subject shares the same network matrix. It would not be computationally feasible to run the *G*-Wishart algorithm separately on individual subjects for the larger network models.

Single-subject HIPPO A model based on the sparse hierarchical prior presented here (equation 5.11), but simplified for single-subject inference, was designed as a comparison to the performance of the *G*-Wishart model. The prior can be expressed as

[12] `ggm_gwish_cbf_direct.m` is available from `github.com/ccnlab/BaCon/tree/master/ggm`.

$$
\begin{aligned}
(\omega_{ii}|\lambda) &\sim \text{Exp}\left(\frac{\lambda}{2}\right) \\
(\omega_{ij}|z_{ij}=1) &\sim \mathcal{N}\left(0, 0.7^2\right) \\
(\omega_{ij}|z_{ij}=0) &\sim \delta_0 \\
z_{ij} &\sim \text{Bernoulli}\,(a) \\
a &\sim \text{Beta}\,(6,6) \\
\lambda &\sim \text{Ga}\,(1/3, 0)\ .
\end{aligned} \tag{5.22}
$$

Inference follows the format above. We use a single chain for inference, drawing 5000 warm-up samples and 10 000 samples from the distribution.

Weakly-sparse HIPPO The hierarchical model is used, conditional on a full graph: all edge inclusion variables z_{ij} are set to 1 (as described in section 5.3 on page 144). We draw 30 000 samples, with 10 000 as warm-up.

HIPPO The full sparse hierarchical model set out in equation (5.13). We draw 30 000 samples in a single chain, with an additional 10 000 as warm-up.

Analysis

After fitting each model, we compute the Root Mean Square (RMS) error (over edges) between each subject's simulated connection strengths and the inferred partial correlation matrices. We compare the mean and standard deviation of this metric over subjects. We also computed the *KL* divergence between the simulated and estimated precision matrices (not shown); the results show the same pattern as those for RMS error.

RESTING-STATE DATA

We use fMRI data from the first 200 subjects of the Human Connectome Project's HCP900 data release (Van Essen *et al.*, 2013). All subjects provided four fifteen-minute resting-state fMRI scans. We also use the 61 subjects from the MEG2 data release (Larson-Prior *et al.*, 2013), who provided three resting-state MEG scans of six minutes' duration. This dataset has been discussed in detail in chapters 3 and 4. The preprocessing steps taken to remove artefacts from the data, and to source-localise and bandpass filter the MEG data, are identical to those described previously.

fMRI analysis

We modelled connectivity between the 25 non-contiguous spatial components, computed by group ICA, that are released by the HCP. For simplicity, we fitted our models to the concatenated data over all four scans. We fitted both the fully sparse and weakly sparse hierarchical models, running three sampling chains for 40 000 samples in the sparse model, with 20 000 needed for convergence in the weakly-sparse model, using an additional 10 000 samples as warm-up. An example of the convergence of the log probabilities is shown in figure 5.2. Additionally, we estimated Tikhonov-regularised precision matrices for each subject. We followed the procedure used for the connectomes released from the HCP, applying only gentle regularisation with λ set to 0.02.

Having computed precision matrices for each subject with these three methods, and converted into partial correlations (equation B.32), we fitted linear predictive models to two traits recorded as part of the HCP: sex and the number of correct scores on a picture vocabulary test. We used the partial correlations

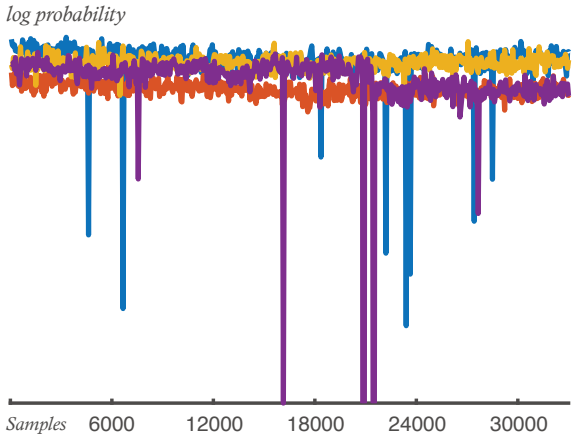


Figure 5.2: CONVERGENCE OF MCMC CHAINS FOR THE SPARSE HIERARCHICAL MODEL. Example of the convergence of four chains of inference for the fMRI data using the sparse hierarchical model. Chains are very stable in slightly different local maxima, showing occasional large divergences as an unlikely model is periodically sampled. The slow exploration of the wells of high probability is illustrated by the purple trace, which moves between two regions of high probability at about the 20 000 sample mark.

on each network edge, for each subject, as the predictors, after regressing out the confounding effects of age, the square of age, sex and an age–sex interaction term,[13] the cube root of cortical volume and of intra-cranial volume, both computed with `Freesurfer`, the software version for image reconstruction, and an estimate of each subjects' motion in the scanner (`rfmri_motion`). Sex was converted to a binary variable, and predicted using logistic regression with elastic net regularisation (Zou & Hastie, 2005). Scores on the picture vocabulary test were demeaned and standardised, and predicted with linear regression using elastic net regularisation. Parameters for the elastic net were tuned in both cases by two nested loops of 5-fold cross-validation. Performance of the models was assessed by computing accuracy (sex) or correlation between predicted scores and real performance (picture vocabulary task), using a 5-fold cross-validation loop for training and prediction. The stratification of subjects into the cross-validation folds took into account the family relationships between

[13] These two confounds are excluded when sex is the predicted variable.

the subjects (Winkler *et al.*, 2015). Prediction and cross-validation were carried out using `nets_predict`.[14]

MEG analysis

We focussed our connectivity modelling on the MEG beta band (13–30 Hz), which is associated with a range of resting-state network profiles (Baker *et al.*, 2014; Brookes *et al.*, 2012*b*, 2011*b*; Hipp *et al.*, 2012; Mantini *et al.*, 2007; de Pasquale *et al.*, 2012, 2015) and exhibits strong heritability in its functional connectivity profile (chapter 4). We employed the parcellation from chapter 4, constructing time courses for each of the 39 nodes, correcting for spatial leakage (chapter 4), and computing down-sampled power envelopes, in the same manner as previous chapters. Time courses were concatenated over sessions for the purpose of functional connectivity estimation.

We estimated functional connectivity in the same manner as for the fMRI data, using both the strong and weakly-sparse hierarchical models (with an identical inference procedure), and lightly Tikhonov-regularised inversion of the sample covariance matrices. The mean heritability of functional connectivity was computed using `APACE` (Chen *et al.*, 2014).[15] ACE models that ascribe a portion of the variability in each phenotype (functional network connection) to additive genetics, environmental factors shared between twins and other unmodelled noise sources, were described in detail in chapter 4. The mean heritability (h^2) was computed over the network connections, with 95% bootstrapped confidence intervals estimated using 5000 sub-samples of the data, and permutation-based p-values computed using 1000 relabellings of the twin

[14] Written by Diego Vidaurre and Stephen Smith; available as part of `FSLNets`, at `fsl.fmrib.ox.ac.uk/fsl/fslwiki/FSLNets`.

[15] Available from `warwick.ac.uk/tenichols/apace`

pairs.

Restricted data analysis

Lastly, we estimated functional connectivities with the fully- and weakly-sparse hierarchical models (HIPPO) for each MEG and fMRI subject using only a small portion of the available data: the first resting-state session (of six minutes) in MEG and the first five minutes of recording for fMRI. We compared these estimates to Tikhonov estimates from the full datasets (18 minutes and 60 minutes, respectively), using the HCP's standard setting of the regularisation parameter λ to 0.02. Additionally on these restricted data samples, we tested network estimation using the same Tikhonov regularisation approach, naïve covariance inversion, and GLASSO, where the parameter λ was set so as to minimise the distance between regularised networks and the unregularised group average network of the same data.

5.5 Results

SIMULATED DATA

The performance of the sparse and weakly sparse hierarchical models presented in section 5.3 on page 136 is compared with that of 10 additional models, over 10 different simulated datasets, in figure 5.3 on page 160. The models are summarised in table 5.3 on page 150 and the datasets in table 5.2 on page 148. It is worth drawing out some general trends. Firstly, as the amount of data increases for inference, covariance estimation becomes less noisy and the error in reconstruction goes down (figure 5.3, simulations 1–4). Of particular note is the difference

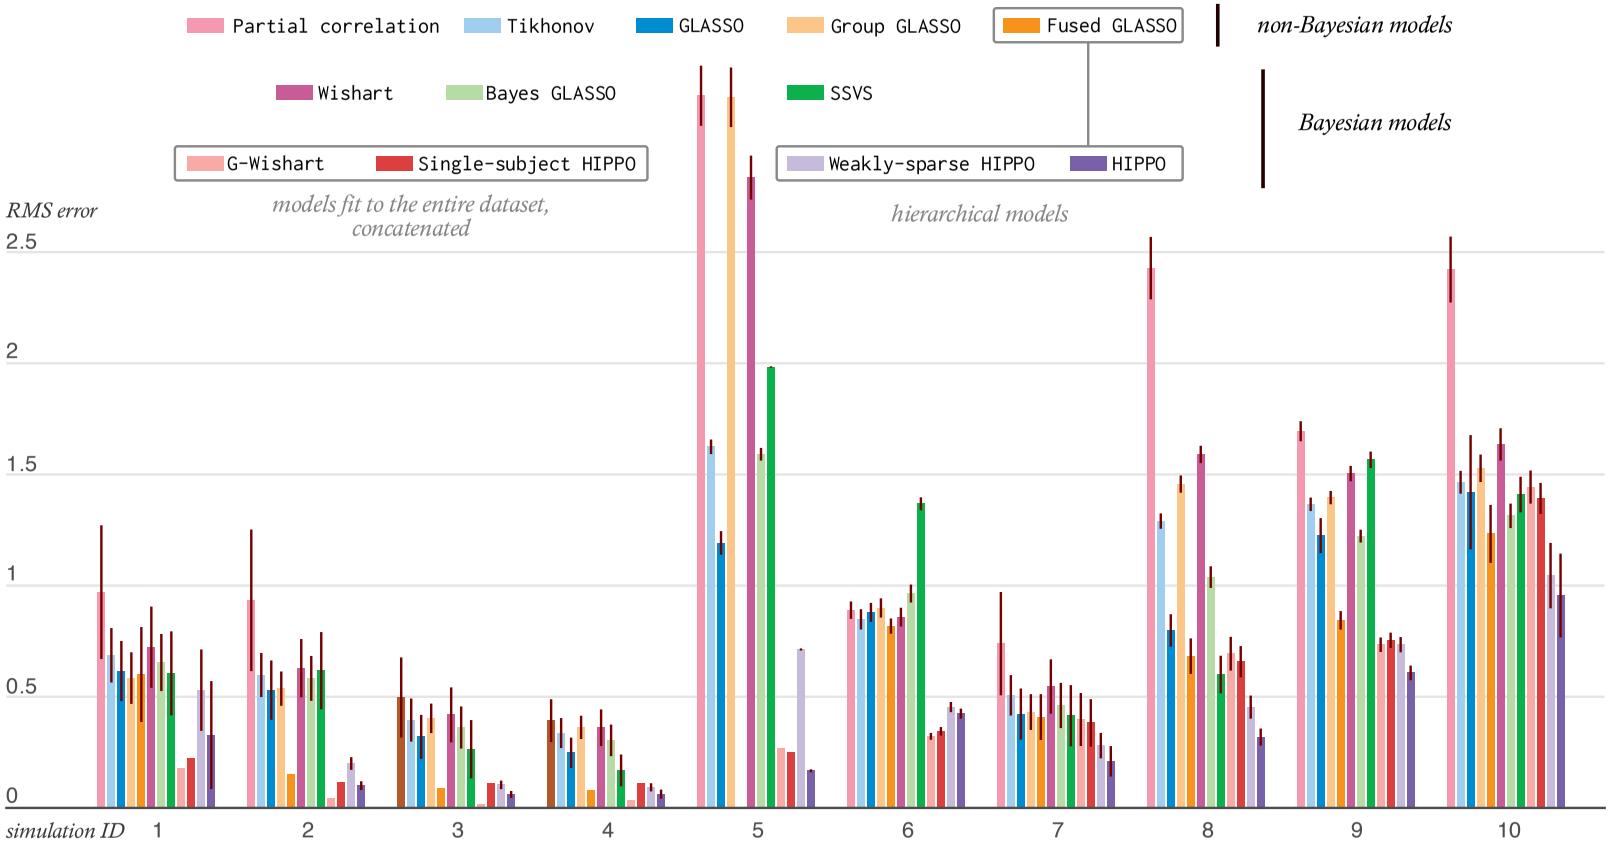


Figure 5.3: COMPARISON OF SPARSE NETWORK MODELLING METHODS ON SIMULATED DATASETS. The RMS error between the simulated precision matrices and estimated precision matrices is shown for 10 multi-subject artificial datasets, using 12 different models for inference (indexed by colour). Bars indicate the mean error over subjects, with the standard deviation over subjects given by the associated dark red line. Datasets 1–5 have no subject-variability in the simulated networks; subject variability increases through datasets 6–10. With limited subject variability, models fitted to the concatenated data perform the best. As this variation increases, the Bayesian hierarchical models win out. There is no result for the fused GLASSO on simulation 5, because the model would not run in a feasible time-frame.

between datasets 1 and 2, where the number of subjects increases but the amount of data per subject is constant. The hierarchical Bayesian models are able to use this increase in information to reduce reconstruction error, whereas models fitted to individual subjects self-evidently are not. Secondly, attempting to perform inference on precision matrices without any form of regularisation is in general a bad idea: all methods tested outperform the simple inversion of the sample covariance matrix. Thirdly, in these simulations of sparse networks, models which build in explicit sparsity priors show improved reconstruction compared to models with continuous priors for sparsity. Fourthly, in simulations with little-to-no between-subject variability, models fitted to the concatenated group data perform the best (unsurprisingly). However, as the variability between subjects increases, the hierarchical models that allow individual subject network estimation win out.

The best-performing non-Bayesian model is Danaher *et al.*'s fused GLASSO, which uses a regularisation term to encourage similarity between subjects' networks. It performs well on the small datasets with limited subject variability, although the sparse hierarchical model consistently produces better estimates (albeit by a very small margin). On datasets 6–10, it is beaten by both hierarchical Bayesian models, and on dataset 5, the largest on test, the fused GLASSO took so long to run (many days, even without attempting cross-validation) that no result could be found.

In summary, these simulations demonstrate that more accurate individual-subject connectivity estimation is possible with our hierarchical Bayesian framework than with the existing state of the art approaches.

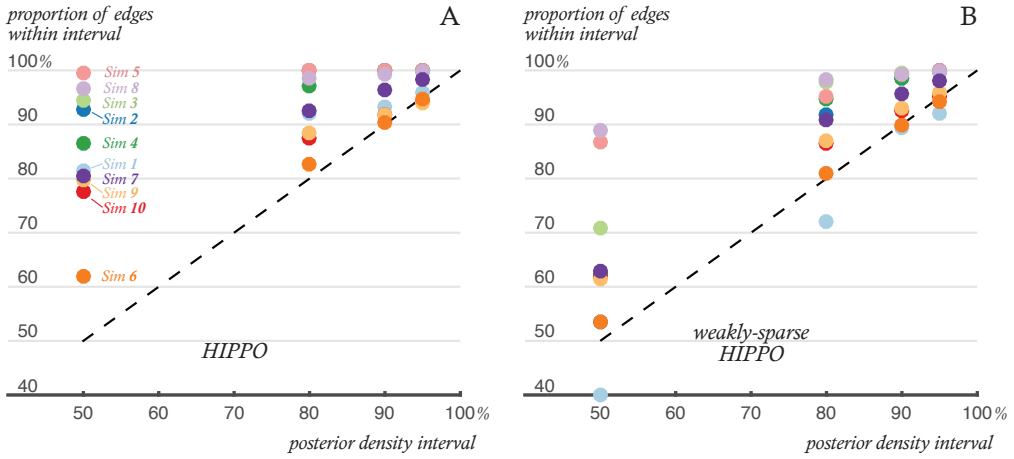


Figure 5.4: COVERAGE OF POSTERIOR INTERVALS ON SIMULATED DATASETS. Proportion of connection strengths included within the 50%, 80%, 90% and 95% highest posterior density intervals, split by simulation, for the sparse (A) and weakly-sparse (B) hierarchical models (HIPPO). The dashed line indicates perfect coverage of the posterior intervals. The models are, in general, conservative, producing wider posterior intervals than is necessary.

CHARACTERISATION OF THE HIERARCHICAL MODELS

Capturing uncertainty

The Bayesian models produce a posterior distribution for each subject's connectivity strengths, summarised by the set of MCMC samples, which could be used to make inference on, for example, differences in connectivity between groups of subjects. Figure 5.4 on page 162 shows that, on the simulated datasets, the hierarchical models tend to be conservative in their posteriors: the 90% highest density interval contains at least 90% of the true edge strengths. In neuroimaging, where control of false positives is the reigning dictum, erring on the side of caution in this way is acceptable. The most accurate density intervals occur for simulation 6, which was designed to mimic the structure of the prior.

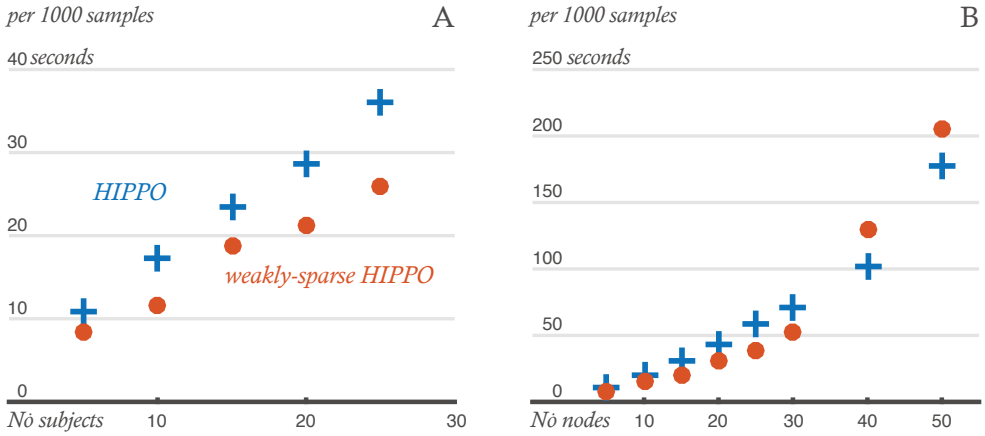


Figure 5.5: SCALING OF COMPUTATIONAL COST WITH CHANGING MODEL DIMENSION. The time for 1000 samples to be drawn using the weakly- and fully-sparse hierarchical model (HIPPO) as the dimensions of the problem are altered. The scaling is linear as the number of subjects are increased (A) and can be quartic as the number of network nodes increases (B). Estimates are the average of 10 runs on 5-node simulation 2 and 5 subjects of simulation 9, respectively.

Table 5.4: COMPARATIVE COMPUTATIONAL COST OF SPARSE BAYESIAN MODELS. The time taken to draw 1000 MCMC samples, or to move through 1000 different model structures, for the three sparse Bayesian samplers tested on simulated data. Of these, only HIPPO was fitted to all subjects simultaneously, the other two were fitted to the first subject only. Timings are given using an average of ten runs on dataset 10, which has 25 subjects, 30 nodes and 100 time points, using a MacBook Pro with 2.8 GHZ processor and 16 GB of RAM. All times are given in seconds.

25 subjects, 30 nodes	HIPPO	Single-subject HIPPO	*G*-Wishart
Time for 1000 samples *(seconds)*	250	11	6600
Time for 1000 samples per subject	10	11	6600
Time for 1000 models	4200	12	6600
Time for 1000 models per subject	170	12	6600

Computational efficiency

Both the weakly-sparse hierarchical model and the fully-sparse model have a computational complexity that scales linearly in the number of subjects, and is roughly quartic in the number of network nodes, on a per-sample basis.[16]

[16] The main computational burden in each sample is the Cholesky decomposition of $p\,(N+1)$ square matrices of size $(p-1)$ (for each of N subjects, we update each column in turn within a single sample), and a series of rank-1 updates to these decompositions that scale with the number of edges proposed for addition or removal in the MH step. The Cholesky decomposition has $\mathcal{O}\left(n^3\right)$ complexity in general for square matrices of size n, although this can be reduced

(Larger models may take more time to converge.) These scalings are illustrated in figure 5.5 on page 163. Sparsity in the model reduces the dimensions of the necessary Cholesky factorisations, which enables the fully-sparse model to outperform the weakly-sparse model as the network size increases in figure 5.5B.

Both methods improve upon the complexity of the best (single-subject) *G*-Wishart algorithms, which show quintic scaling with the number of nodes (Hinne *et al.*, 2015). A comparison of the speed of inference using our hierarchical approach, a single-subject version of the prior (5.22) and Hinne *et al.*'s *G*-Wishart sampler is given in table 5.4 on the following page. Exploration of the sparse model space is slow using the hierarchical prior—but still faster than the *G*-Wishart model, which becomes infeasible to run, even for a single subject, on networks with more than 50 nodes.

SMALL DATA SAMPLES

A useful metric for assessing improvements to network estimation using real data is the ability of a model to estimate connectivity from a small chunk of the recording. We compare beta-band network matrices inferred from resting-state MEG recordings from the HCP, using either all three sessions of data, or only a single 6-minute session of data. Additionally, we compare fMRI network matrices inferred from resting-state HCP data, using either all four sessions of 15 minutes, or only the first 5 minutes of recording. Treating networks inferred from the full datasets using mild Tikhonov regularisation ($\lambda = 0.02$) as a good approximation of 'the truth,' figure 5.6 on the next page compares the RMS differences to these estimates from five inference methods, which only had access to the first portion of the data. The hierarchical models significantly outperform

for sparse matrix computations.

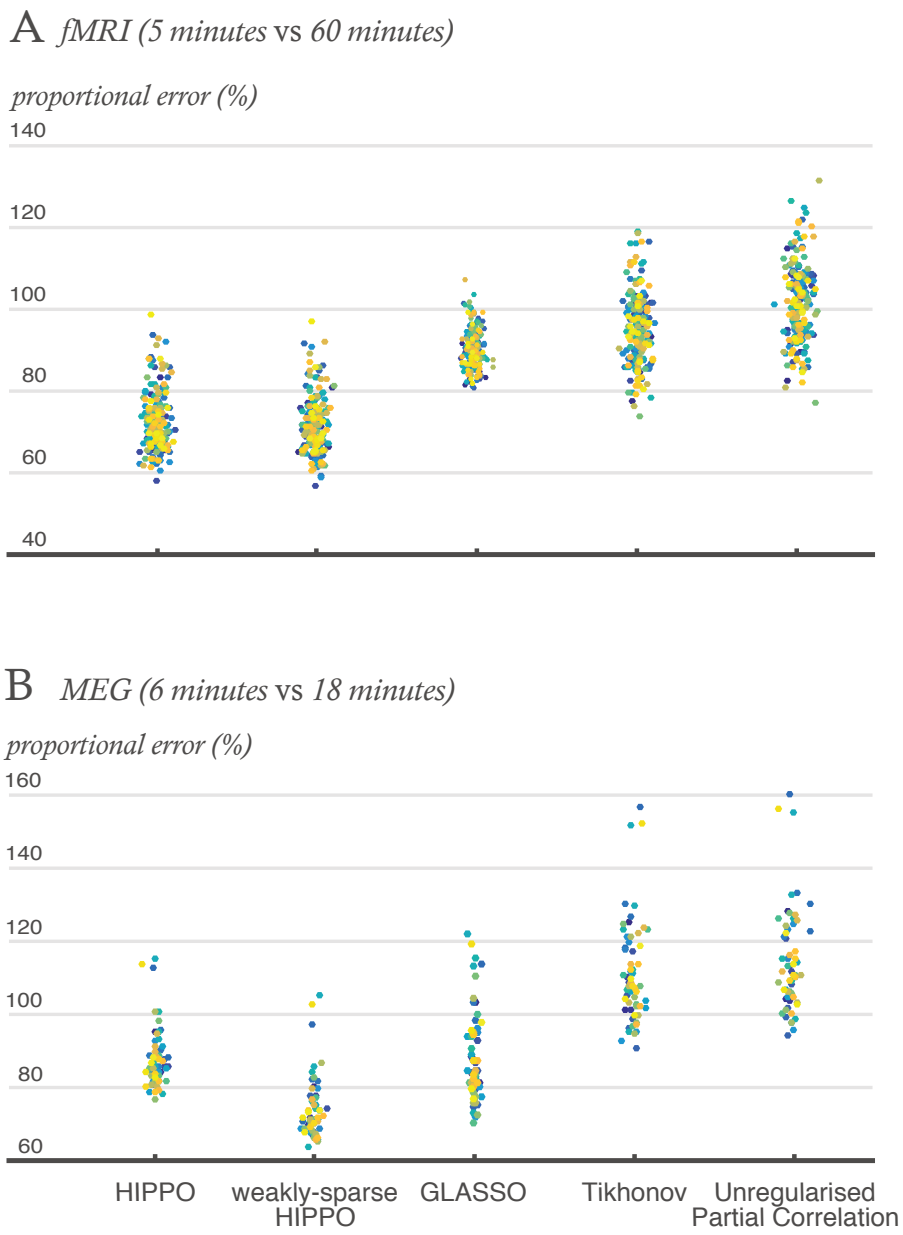


Figure 5.6: ESTIMATION OF NETWORK MATRICES FROM SMALL SAMPLES OF DATA. Single-subject networks estimated from the first 5 minutes of resting-state fMRI data (A), and single-subject networks estimated from the first resting-state session of MEG recordings in the beta band (B), were compared to the average of each subject's complete data, for each modality. Networks were estimated using the hierarchical model HIPPO and its non-sparse version, with the graphical LASSO, with Tikhonov regularisation and without any regularisation. The results were compared in each subject to a mildly Tikhonov-regularised estimate from all three sessions' concatenated data; the RMS error from this estimate is displayed as a percentage of the mean connectivity of each subject's network matrix. Coloured dots identify individual subjects.

the standard regularised solutions, producing a reduction in error which is on the same order of magnitude as the subject-to-subject variation in this metric. The only exception is the sparse hierarchical model in MEG, which does not outperform the Graphical LASSO (GLASSO). (A paired t-test for a difference in mean performance, conducted non-parametrically using 5000 sign flips of the difference between pairs, gave p=0.45 for this exception, and p <0.001 for the other cases.) These results equate to a mean improvement in estimation (with standard deviation over subjects), compared to GLASSO, of 0.7±7% (14±6%) of the average connection strength for the MEG data, and 18±6% (19±5%) for the fMRI data, when using the (weakly-) sparse hierarchical model. It is worth noting that the weakly-sparse hierarchical model can beat GLASSO, in both modalities, even when the regularisation parameter for the latter is chosen *with knowledge of the correct solution*—when it is allowed to cheat. If λ is chosen for each subject so as to minimise the difference between the estimated network and the solution used here as truth, the proportional error with standard deviation over subjects for the GLASSO estimate is 82±8% for fMRI and 83±7% for MEG. The hierarchical model can therefore reduce network estimation error to an extent that is better than GLASSO would ever be able to achieve.

The fused GLASSO of (Danaher *et al.*, 2015), which performed quite well on simulations, did not approach convergence on these short datasets even after five days of computation on a MacBook Pro (with a 2.8 GHz processor and 16 GB of RAM). Cross-validation of the parameters and computation of a solution was therefore unachievable in a sensible time-frame, and results from this method are not available for comparison.

TRAIT PREDICTION USING FMRI FUNCTIONAL CONNECTIVITY

The HCP has released an analysis of the ability of functional connections to predict a wide range of biological and behavioural traits,[17] using all 841 subjects with resting-state fMRI recordings, suggesting that in some cases there is discriminative information embedded within the functional connectome. We chose two of the better-predicted traits, scores on a picture vocabulary test (a measure of crystallised intelligence[18]), and sex. Figure 5.9A on page 170 presents a comparison of predictive performance on these two measures using partial correlation networks estimated using Tikhonov regularisation (the algorithm employed in the HCP's published results), and the fully-sparse and weakly-sparse hierarchical models. In both cases, the differences in predictive ability between the models is smaller than the error on the cross-validated estimate (although there is no difference in sampling or random variation between the methods). The correlation between scores on the picture vocabulary test and the predicted responses are slightly worse (by a few percent) for the hierarchical models, whereas there is a slight improvement in accuracy for the prediction of biological sex (although the accuracies are so close to 0.5 that it is difficult to have confidence in the performance of any model). We were principally interested in the differential prediction ability of the models, so permutation tests to look for significant classification have not been performed.

The group average functional connectivities inferred by the three models are very similar (too close to see major differences when connection strengths

[17] Available at db.humanconnectome.org/megatrawl/.

[18] *Crystallised intelligence* is defined by Wikipedia (on the 21st August 2016) as 'the ability to use skills, knowledge, and experience. It does not equate to memory, but it does rely on accessing information from long-term memory.' This is in contrast to *fluid intelligence* which characterises an ability to reason, deduce and to solve novel problems.

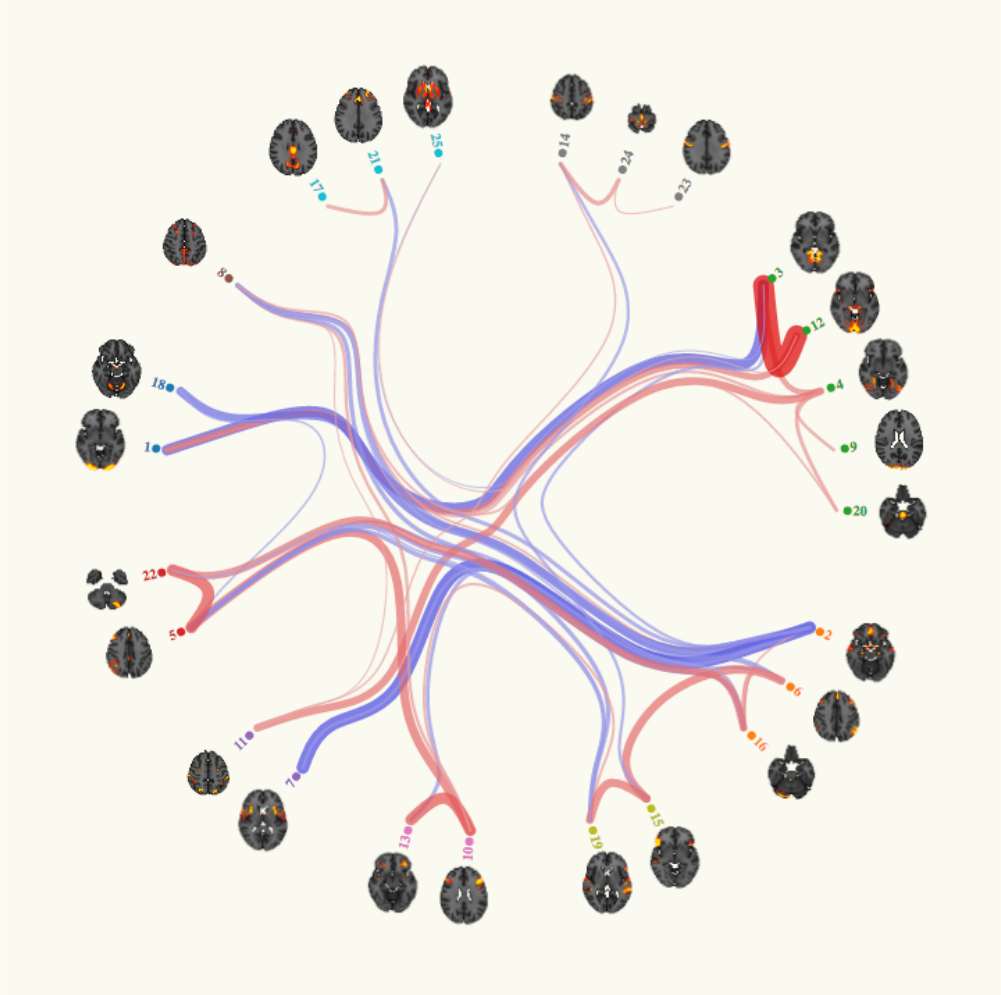


Figure 5.7: GROUP AVERAGE FUNCTIONAL NETWORK FOR THE HCP FMRI DATA. Group average partial correlation network, computed using the hierarchical model (HIPPO). The results obtained using the traditional (Tikhonov-regularised) or weakly-sparse hierarchical models are visually identical. Numbers and brain slices indicate the ICA components which act as network nodes.

are displayed on a heat map). The average partial correlation network for the hierarchical model is shown in figure 5.7 on page 168.[19] The posterior for the edge inclusion variables gave very high probabilities (over 99%) for all connections, presumably because the quantity of data in an hour's total recording time is sufficient to provide evidence for connectivity between all nodes, even if this connectivity is small in some cases. Further support for this suggestion

[19] Figure produced using Paul McCarthy's visualisation tool in `FSLNets`.

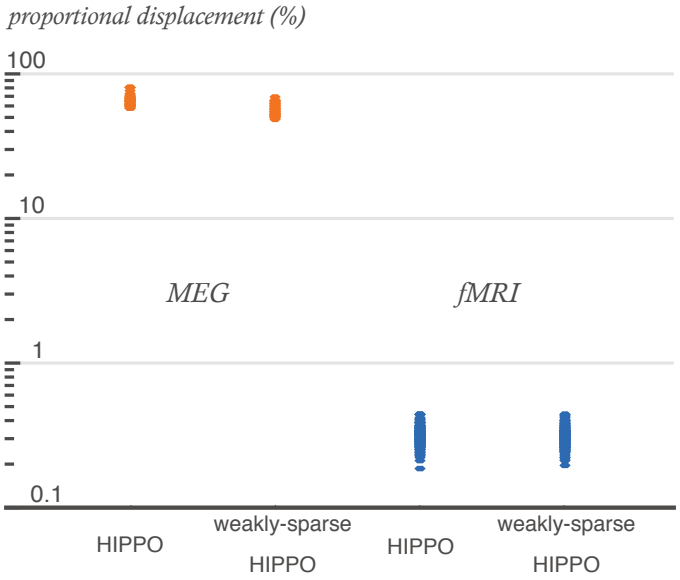


Figure 5.8: DISPLACEMENT OF NETWORK MATRICES FROM THEIR UNREGULARISED FORMS after estimation with the hierarchical model. The RMS distance between partial correlation matrices, inferred with the strongly- and weakly-sparse hierarchical models, and the standard Tikhonov-regularised estimates, is displayed as a percentage fraction (on a log scale) of the mean connection strength in each subject. Each dot represents a different subject. Much more adjustment was made to the MEG matrices than to the fMRI networks.

comes from the individual subjects' networks, which are hardly changed under the hierarchical models from the conventional estimates; this is particularly apparent in comparison to the adjustment made to the MEG networks, shown in figure 5.8 on page 169.

HERITABILITY OF CONNECTIVITY

The HCP dataset is a twin study, and the variability within the functional connectomes of the subjects is determined, in part, by genetic and shared environmental effects (chapter 4). Heritability, h^2 (A), is the proportion of variance in a phenotype which can be explained by additive genetic factors. It is estimated using linear decompositions of the variance into heritability, the environmental effect shared between twins, c^2 (C), and any other unmodelled variance sources and noise, e^2 (E). Improving the quality of network matrix estimation, therefore,

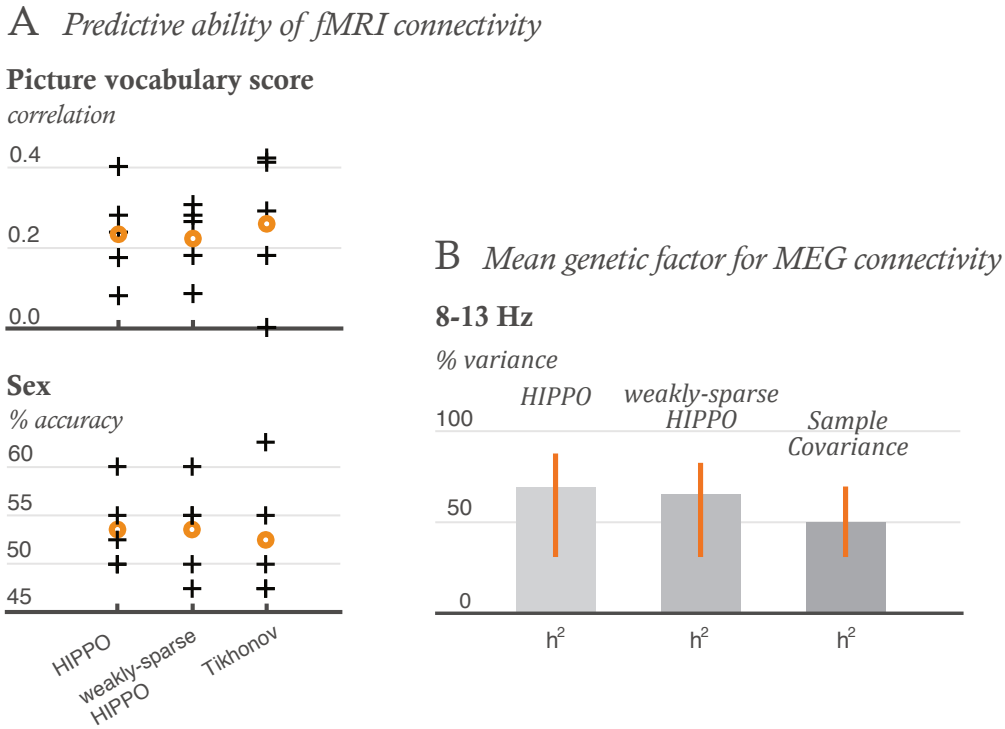


Figure 5.9: PERFORMANCE OF THE HIERARCHICAL MODEL. (A) *Prediction of picture vocabulary scores (top) and sex (bottom) from* fMRI *functional connectomes.* The correlation between scores on a picture vocabulary test and the predicted scores after training a regression model using resting-state fMRI functional connection strengths is shown, together with the accuracy in prediction of subjects' sex using a similar logistic regression model. Results are presented for standard Tikhonov-regularised network inference, and for networks inferred using the weakly-sparse and fully sparse versions of the hierarchical model (HIPPO) presented in this chapter. Black crosses indicate results from each of five cross-validation folds; the orange circle indicates the mean of these scores. (B) *Estimation of the mean heritability of functional connections in the* MEG *beta band.* Heritability estimates are compared between the two versions of the hierarchical model and estimation using sample covariance matrices, as performed for chapter 4. Bars give the estimated effect, on average over all connections, and the error lines denote 95% bootstrapped confidence intervals.

with a hierarchical model that is blind to the twin structure of the data, should reduce estimates of e^2 and increase estimates of heritability.

Replicating the analysis in chapter 4, we fitted ACE models on each edge, and analysed the average heritability over the edges, computing bootstrapped confidence intervals and permutation-based tests for significance (results shown in figure 5.9B). To allow easy comparison to the original study, we fitted the models to correlation matrices: those estimated from the sample covariance matrix, and

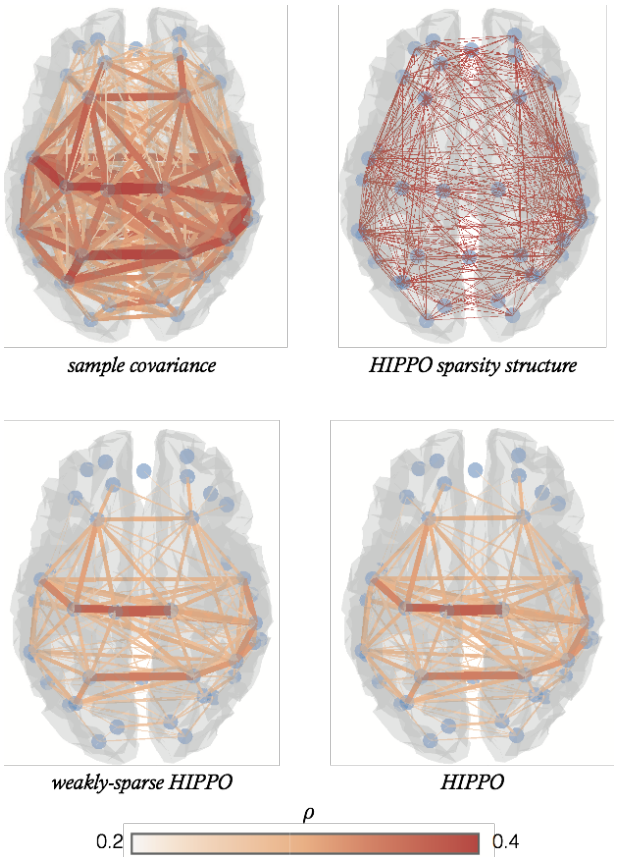


Figure 5.10: GROUP AVERAGE NETWORKS FOR THE MEG BETA BAND. Group average functional correlation networks in the MEG beta-band (13–30 Hz), computed using three estimation methods: the sample correlation matrix (the same result as presented in chapter 4), the hierarchical model HIPPO and the weakly-sparse version of HIPPO. Also shown is the median probability model of direct connections in the group partial correlation matrix, computed using HIPPO.

correlation matrices estimated from inversion of precision matrices regularised using the hierarchical inference procedure. Using the (weakly-) sparse hierarchical inference procedure, the estimates of heritability (with 95% confidence intervals in square brackets) increased from 0.5 [0.3, 0.7] to 0.7 [0.3, 0.9] (0.7 [0.3, 0.8]). (Permutation-based p-values computed for each respective model are 0.005, 0.001 and 0.002.) This increase in heritability is related to a reduction in the variance that is *unexplained* by genetic factors, shared environmental, maternal or household effects, age, sex or brain size, of 0.19 and 0.15 percentage points respectively for the sparse and weakly-sparse models (from an original

estimate of 39% for e^2). This corresponds to a noise reduction of about 50% and 40% using the sparse and weakly-sparse models. These differences are not explained by any random variation in noise or sampling, as exactly the same data were passed into each of the covariance models. Group mean networks for the three models are compared in figure 5.10 on page 171.

5.6 Discussion

We have presented two hierarchical models for the functional connectivity measured with MEG or fMRI, which characterise connectivity by the partial correlation between activations in ROIs, and jointly infer connection strengths for individual subjects and the population average. This ability to perform joint inference at multiple levels of the hierarchy, sharing information between subjects and regularising connection strengths towards the group mean, is an innovation in covariance modelling, previously only possible for functions of linear effects (Gelman *et al.*, 2014). A sparsity structure in the underlying GGM, which is shared over subjects, can be explicitly encouraged and estimated. A Bayesian modelling framework has been described, which allows scalable inference in large models over many subjects, and which provides an approximate posterior distribution over all variables.

Accurate estimation of precision and covariance matrices is difficult and noisy, particularly if the number of data points available is not much greater (or indeed, is less than) the number of free parameters in the matrices. A wide range of techniques have been developed to tackle this problem, most of which regularise weak elements of the matrices towards zero with some sparsity-promoting scheme. The importance of estimating precision matrices in this

fashion is clear in figure 5.3 on page 160, where over many different datasets, the estimation of precision matrices using inversion of the sample covariance matrix leads to a large amount of reconstruction error. In their ability to draw information from many subjects at once, the hierarchical covariance models outperform a wide range of methods, fitted to individual subjects, that represent the current state of the art in covariance modelling. In simulations with large amounts of between-subject variability, the hierarchical models continue to succeed, additionally outperforming models fitted to the entire concatenated dataset.

Hierarchical model structures and the partial pooling of information over subjects can be most useful when limited data are available within each subject (Gelman *et al.*, 2014). Compare, for example, the improvement in matrix reconstruction for simulations 2–4, in figure 5.3, for which the number of subjects and network nodes remains constant, but the amount of data available within each subject increases. The hierarchical models do well in each case, but the differential improvement over more basic models is largest for the case with the least data. We observe the same effect in our studies of real data. Network estimation with a limited subset of both fMRI and MEG recordings can be greatly improved using the hierarchical models (see figure 5.6 on page 165). The key point is that the quality of single-subject connectivity inference is enhanced, using commonalities between subjects to reduce the noise within each. Thus, estimates of the heritability of functional connectivity using MEG beta-band data are increased (figure 5.9B on page 170), because a portion of the noisy variability within the dataset is reduced.

Why, then, is there no improvement to speak of in the quality of biological and behavioural trait prediction using functional networks estimated from fMRI

data? Part of the answer may be that, in fact, functional connectivity encodes very little information that can be extracted by a linear model about subjects' sex or their scores on picture vocabulary tests. However, we suggest that the same line of reasoning can be applied. In the HCP dataset, there is vastly more data available for connectivity inference in fMRI than for MEG: the total recording time for fMRI is one hour, for MEG 18 minutes, and in our analyses we estimate networks with 25 nodes in the former case and 39 in the latter. (Sampling rates are comparable across the two modalities because we apply our network models to the down-sampled power envelopes of MEG recordings.) The quantity of data may therefore be dominating the prior, such that the hierarchical model provides little improvement over simple estimates. An assessment of the difference between networks inferred using the hierarchical models and the standard inference approach confirms that the connectivity matrices inferred from fMRI data were indeed manipulated much less than those inferred from MEG data (figure 5.8 on page 169). If we were to tackle larger networks in fMRI, increasing the number of parameters to be inferred from the data, our hierarchical approach might become more useful.

Sparsity in functional connectivity matrices provides not only a mechanism to improve noisy estimates, but can also improve the interpretability of the networks. An assessment of whether a connection is present in one group of subjects (for example, healthy controls), but absent in another (patients, to complete the analogy), appears to be a pleasingly simple categorisation to many authors, although this may only reflect attitudes promulgated by the prevailing null hypothesis testing framework. Our sparse hierarchical model offers an analyst the ability to draw samples from the approximate posterior distribution of the sparse graph representing the network structure of her dataset. She would

then be able to construct posterior summaries of any function of that graph, $f(\boldsymbol{Z})$, such as posterior confidence intervals on a range of graph theory metrics. One can envisage performing a Bayesian comparison of the degree centrality in functional networks inferred from schizophrenic patients as compared to controls, for example. Using a fully probabilistic description of the network connections would be preferable to testing degree centrality (or another metric) over many binary network matrices created with a sliding scale of thresholds, as is currently best practice. However, while simply using MCMC chains to average over possible models can provide effective regularisation of the parameter estimates, making inferences about graph theoretic functions of the network structure requires two conditions to be met. The first is basic, in that the analyst must be confident that she has run the sampling chains for long enough that she has obtained a good characterisation of the posterior. The second condition is that she must believe that a network model with a shared sparsity structure across subjects is a good representation of the data. If it is not, then any further analysis based on this graph structure is uninterpretable, whether a posterior distribution can be constructed or not. We turn to this second assumption now.

The sparse hierarchical model we present expressly shares the sparsity structure over all subjects: the network structure is therefore considered a property of the entire population, about which no subject is considered to deviate. This may be plausible, particularly if analyses are restricted to sub-populations in which this assumption holds (fitting the model separately to patients and healthy controls, for example). However, we also made this choice because it enables tractable sparse inference within the hierarchical structure, and allows efficient computation within the MCMC sampler. Although we can conceive of vastly more complicated prior structures, with individual network structures for each

subject, linked perhaps by a Markov Random Field, we have not attempted such an inference program. Can we evaluate to what extent our prior assumption holds water? This is difficult without comparison to more general structures. We note, however, that inference using the large fMRI dataset gave evidence overwhelmingly in favour of the full model—that is, the model with all connections present. (Although, with sufficient data, there can be enough support for very small, but significant, effects.) Additionally, the performance of the sparse hierarchical model was not as strong as the weakly-sparse prior structure when used to estimate networks from a small subset of the MEG data (figure 5.6 on page 165). As a result, we suggest that the sparse model structure can be used for effective regularisation of connectivity estimation, and that further interpretation of the network structure be performed with care.

Explicit sparsity was only one flavour of the hierarchical prior that we introduced. The other keeps all connections in the model but weakly promotes sparsity by regularising the population-level connectivity estimates, and encouraging connectivity strengths in individual subjects to be distributed about this point. Our framework for the prior and the inference scheme will allow further customisation: stricter continuous distributions that more strongly encourage sparsity, such as the Laplace or horseshoe priors, could easily be incorporated. In our tests on real data, the weakly-sparse hierarchical model performs similarly to the stricter spike-and-slab prior (see figures 5.6 and 5.9). However, the need for convergence of the posterior over the edge inclusion variables Z (which exhibit highly autocorrelated behaviour) is removed, so the total run time need not be as long. The fMRI and MEG results for this chapter were produced using parallel sampling chains that took about 14 hours to run for the weakly-sparse model, and 20 hours for the sparse model.

As well as providing a Bayesian framework for characterising uncertainty about connection strengths and network structures, the hierarchical models outperform the best competition in a number of additional ways. The popular GLASSO algorithm cannot offer the same quality of network reconstruction. An advanced multi-subject extension of this program, the fused graphical LASSO, cannot match the sparse HIPPO algorithm in our simulations, and is too slow to be useful on real datasets. The best Bayesian sparse model for connectivity, the *G*-Wishart distribution, offers precise estimation for individual subjects, but is impractical for use in any study with moderate numbers of subjects, or even functional networks with 50 or more nodes.

In conclusion, we have presented an advance in functional connectivity and covariance modelling by designing sparse and weakly-sparse hierarchical priors for the distribution of connection strengths over subjects. The inference program, HIPPO, is sufficiently scalable to allow it to be applied to conventional neuroimaging datasets. Improvements to the quality of single-subject network estimates are possible, particularly in small or noisy datasets, with concomitant increases in sensitivity to properties of interest (such as heritability) in the functional connectomes. The models are applicable both to fMRI and to MEG data, and we hope they will enable improved inference for studies in both modalities.

6

Conclusion

I HAVE SET OUT new network modelling methods for use in MEG and fMRI. In a case-study multi-modal analysis, I have found new evidence for the genetic basis of functional connectivity in humans. I will briefly highlight the novel contributions and impact of this thesis, before looking ahead to how these new techniques can continue to be improved and applied.

6.1 Contributions of the thesis

NETWORK MODELLING METHODS

Leakage-corrected connectivity analysis in MEG

Chapters 2 and 3 developed and tested a new method for network modelling in MEG that is robust to spatial leakage artefacts. This method removes any shared signal between the nodes of the network, enabling any multivariate con-

nectivity model to be applied to the data. I focussed in particular on correlation analyses and GGMs, the use of which enable direct parallels to be drawn to network estimation in fMRI, facilitating the synthesis of evidence from multi-modal datasets. Chapter 3 demonstrated that the repeatability of network estimation in MEG, using short-acquisition resting-state scans, is in general quite poor, with the variance shared between networks inferred from successive scanning sessions rarely greater than 0.5. Nevertheless, correlations and partial correlations between the envelopes of leakage-corrected time series were shown to outperform a wide range of alternative, popular methods in this regard.

Outside of the work presented in this thesis, leakage-corrected envelope correlations estimated using my methodology have already shown sensitive differences in connectivity between patients suffering from motor neurone disease, when compared to healthy age-matched controls (Proudfoot *et al.*, in preparation), and revealed marked differences in functional co-activation exhibited during the different stages of sleep (Stevner, in preparation).

The software developed for leakage correction and network modelling in MEG is freely available from `github.com/OHBA-analysis/MEG-ROI-nets/`, and is being incorporated into routine analysis pipelines at the Universities of Cardiff, Toronto and Nottingham, as well as in Oxford.

A Bayesian multi-subject network model

Chapter 5 sets out a new hierarchical Bayesian model and inference procedure for the group analysis of functional networks, with application to both MEG and fMRI data. It offers improved inference of single-subject functional connectivities, particularly for short-acquisition datasets, by leveraging information flexibly from the whole group. It also provides the opportunity for inference on the

underlying graphical network, and a posterior distribution over any graph theoretic function of this structure.

The `HIPPO` package for inference of this model is hosted at `gitlab.com/giles-colclough/OHBA-HIPPO/`.

A GENETIC BASIS FOR FUNCTIONAL CONNECTIVITY

Using data from twins, collected with two imaging modalities, chapter 4 presented an analysis of the contribution of genetic factors to the phenotypic variability of functional brain connectivity. I was able to show that genetic factors, rather than the developmental environment, are of greater importance in determining how the human cortex is functionally connected.

These findings are an important contribution to an active field within imaging genetics that explores the links between genetics, disease, cognition and patterns of brain activity. In addition, the results of the investigation into genetic and environmental influences feed into the wider 'nature *vs* nurture' public debate, by characterising the difference in contribution from these two factors to the strength of intrinsic cortical communication, which is a fingerprint of cognition and behaviour. This point is one of the more interesting outcomes of the analysis, and has been largely ignored by the field so far.

Lastly, as a case study in multi-modal functional connectivity analysis, it is intended to provide an example of how such studies should be carried out. It promotes new, robust, non-parametric inference methods for heritability analysis and encourages the incorporation of careful control analyses that past work has lacked.

6.2 Looking ahead

IMPROVING THE EXISTING MODELS

A number of opportunities are available for enhancing the performance and utility of the multi-subject network model presented in chapter 5. These can be roughly split into enhancements to the sampling efficiency or the quality of estimation on the network structure, and a look towards integration of the method within a wider-scope analysis of the resting-state.

A more integrated analysis of resting-state network behaviour

Functional network estimation is predicated upon the underlying parcellation. Indeed, it is becoming increasingly clear that the information content of a functional connectome can be strongly affected by the quality and form of the parcellation used. Separating the inference of network and nodes, then, is perhaps a bad idea.

While my work has focussed on the second part of this process, many others have improved the methods for categorisation and parcellation of cortex. A joint inference framework for both parcellation and functional connectome has been suggested for fMRI (Harrison *et al.*, 2015), but much of the modelling effort focussed on the spatial variability of nodes over subjects. The hierarchical network modelling method of chapter 5 will fit nicely into this framework, providing a more structured inference system—although for this union of variational and sampling inference methods to succeed, there will be some hurdles associated with computational efficiency to overcome.

Efficiency of sampling

The hierarchical model can explicitly infer on the structure of a Gaussian graphical network model, by sampling the binary elements of the adjacency matrix that codes for the edges in the network. However, these Monte-Carlo estimates of the probability of edge presence,

$$p\left(Z_i = 1\right) = \frac{1}{N}\sum_{s=1}^{N} Z_i^s\,, \tag{6.1}$$

formed by averaging the N sampled values of Z for each edge i, can be noisy. More robust estimates can be obtained by estimating the presence or absence of an edge, conditional on all other variables in the model,

$$p\left(Z_i = 1\right) = \frac{1}{N}\sum_{s=1}^{N} p\left(Z_i^s = 1|\{Z_{j\neq i}^s\}, \{\boldsymbol{\mu}_{j\neq i}^s\}, \boldsymbol{\sigma}^s\right)\,. \tag{6.2}$$

This process of Rao-Blackwellization (Blackwell, 1947; Ghosh & Clyde, 2011; Guan & Stephens, 2011) can be particularly useful for the estimation of binary variables, and is an obvious next-step improvement to the sampling estimators employed in chapter 5.

Broader work could be done to improve the efficiency of the sampler. For example, better use could be made of parallel computational processes. Following Gelman *et al.*'s advice, multiple sampling chains were used for inference and analysis of convergence. This provides opportunities to use recent methods that pass information on-line between different chains, allowing large jumps between regions of parameter space and the sharing of information between the parallel streams (Craiu *et al.*, 2009; Laskey & Myers, 2003; Murray, 2010). Implementations of the algorithm that performs the necessary sequence of Cholesky

decompositions on GPU cards might also offer performance improvements.

The methods developed for inference were specific to the model structure. While Gibbs sampling is an old technique, it was appropriate for this application. The model is not conjugate, and so not amenable to classic variational inference (Beal, 2003) or expectation propagation (Minka, 2001, 2004) approaches. The parameter space was too large and the datasets too big even for the advanced multi-purpose sampling packages using Gibbs sampling (`JAGS`), Hamiltonian Monte Carlo (`STAN`) or nested sampling (`MultiNest`) algorithms. A tailored solution seemed the best fit.

However, the process of designing specific inference schemes, whether for variational or Monte Carlo inference, both in this application and in general, is very time-consuming. There is an emerging methodology for approximate Bayesian inference that looks set to dramatically cut this development timeframe, offering opportunities for model development in three weeks rather than three years.

AN IMMINENT REVOLUTION IN BAYESIAN INFERENCE

No-one working in any field of data analysis could fail to have missed the massive success of deep neural network models, which are now at the heart of nearly every artificial intelligence, robotic vision and on-line advertising platform. The Bayesian inference community has, however, largely kept itself isolated from the explosion of deep learning, fearing perhaps to dip its toe too deeply into the pool of black-box methods, golden rules and dirty hacks that are often needed to get adequate performance from deep neural net architectures. During the course of my DPhil research, a bridge has been established between the two camps. Kingma & Welling, in 2014, made two conceptual leaps. The first

was that the low-dimensional layers of a neural network can be understood as a parametrised *model* for the data, and that the decoding and encoding networks of an autoencoder are therefore synonymous with a non-linear forward model and associated approximate inference procedure.[1] The second was to employ the variational bound for probabilistic inference as the optimisation target within a neural net framework, demonstrating that very noisy Monte Carlo approximations of the intractable integrals[2] that comprise this bound can be tackled using established tools for stochastic gradient descent. Development of these ideas by Yarin Gal has led to the realisation that almost any deep learning model can be recast as a Bayesian inference problem (Gal & Ghahramani, 2015*a*,*b*). *Vice versa,* almost any Bayesian model can be fitted using an approximate variational solution (Gal, 2015), exploiting the scalable and efficient solvers and automatic differentiation engines that are part of current deep learning packages.[3] This represents a major advance in Bayesian inference, with greater flexibility and performance than other attempts in this direction (such as Kucukelbir *et al.*, 2015).

In the new paradigm, a model and its approximate posterior can be expressed in a high-level coding language. The variational objective, or lower bound on the model's evidence, is evaluated using one or more Monte Carlo samples. Symbolic gradients of this objective are automatically calculated, and stochastic optimisers used to fit the model. By eliminating many arduous analytical steps and by-hand integrations, and making the software implementation easier to relate to the overall model, much more time can be spent investigating model specification and performance, rather than dealing with implementation.

[1] The variational autoencoder was simultaneously discovered by Rezende *et al.* (2014).
[2] Even using only one sample of the integrand as an estimator!
[3] The most popular are perhaps `Theano` and Google's `TensorFlow` engine.

This approach is very new, and yet to be used for neuroimaging analyses. However, I believe it holds great promise for bringing the power of large Bayesian models to bear on the wealth and complexity of data that we collect. It will surely be of particular use to the network modelling community as the scope and quality of models continue to develop.

6.3 Final remarks

All of the complexity of human function and dysfunction is bound up within, and expressed through, the functional connectome. By identifying a strong genetic basis for these patterns of communication, I have added to the explorations of this system that are currently scratching the surface of the multifaceted role of connectivity for cognition. Our understanding will continue to grow, by continuing to model the relationships between connectomes, behaviour and health. Accurate modelling of individuals' functional networks, in multiple imaging modalities, will be the basis of this growth. It is my hope that the methods developed in this thesis will contribute to this effort, and, with use, will themselves grow and improve within it.

A

Probability distributions

Several definitions of the gamma, Wishart and G-Wishart distributions exist. This apppendix lays out the conventions used for this thesis, which align with Wikipedia and Gelman *et al.* (2014).

Gamma distribution

The gamma distribution with shape α and scale β is

$$\mathrm{Ga}\left(x;\, \alpha, \beta\right) = \frac{\beta^{\alpha}}{\Gamma\left(\alpha\right)} x^{\alpha-1} e^{-\beta x} \tag{A.1}$$

where $\Gamma\left(\cdot\right)$ is the gamma function.

Wishart distribution

The G-Wishart distribution is

$$\mathcal{W}_{\mathcal{G}}\left(\boldsymbol{X};\, n, \boldsymbol{V}\right) = \frac{|\boldsymbol{X}|^{(n-p-1)/2}}{Z_{\mathcal{G}}\left(n, \boldsymbol{V}\right)} \exp\left[-\frac{1}{2}\operatorname{trace}\left(\boldsymbol{V}^{-1}\boldsymbol{X}\right)\right] \mathbb{1}_{\boldsymbol{\Omega}\in\mathbb{P}_{\mathcal{G}}} \tag{A.2}$$

where $\boldsymbol{V}$ is the scale matrix, n indicates the degrees of freedom, $Z_{\mathcal{G}}$ is the intractable normalising constant, and $\boldsymbol{X}$ is constrained to live on the cone $\mathbb{P}_{\mathcal{G}}$ of positive definite $p \times p$ matrices with zeros indicated by the graph $\mathcal{G}$.

When the graph $\mathcal{G}$ is full (all edges are present), the distribution is simply the Wishart distribution, and the normalising constant is

$$Z\left(n, \boldsymbol{V}\right) = 2^{\frac{np}{2}} |\boldsymbol{V}|^{\frac{n}{2}} \Gamma_p\left(\frac{n}{2}\right), \tag{A.3}$$

where $\Gamma_p\left(\cdot\right)$ is the multivariate gamma function.

B

Inference for the hierarchical model

Full details of the inference procedure and MCMC sampling scheme for the hierarchical model of functional connectivity, set out in chapter 5, are given in this appendix. The conditional distributions for the sampler are derived at the end of this appendix in section B.6 on page 206 and the full sampling algorithm is given in algorithm B.2 on page 201.

Inference occurs in two stages. First, estimation of a time course within each ROI. Second, inferring the precision matrices for each subject. Welding these two stages together into a joint inference procedure is a level of complexity that we have not yet tackled. However, uncertainty about the ROI time courses in the first phase can be easily passed up to the second by drawing consecutive samples of the sum-of-squares matrix $\boldsymbol{S}$ for each subject, and using these as the input for each step of the MCMC algorithm (Gelman *et al.*, 2014).

B.1 Activation within ROIs

To illustrate this point, we describe a simple linear regression model to infer the activation within each ROI (5.2). Assuming for simplicity that the measurement noise in the voxels is isotropic and uncorrelated, sharing a common variance σ^2, and using simple uninformative priors,

$$\left(\boldsymbol{Y}, \sigma^2\right) \sim 1/\sigma^2 \tag{B.1}$$

(Gelman *et al.*, 2014, ch. 14), the posterior is easily found as

$$\begin{aligned}
\boldsymbol{Y}|\sigma^2, \boldsymbol{X} &\sim \mathcal{N}\left(\hat{\boldsymbol{Y}};\, \boldsymbol{V}\sigma^2\right) \\
\sigma^2|\boldsymbol{X} &\sim \text{Inv-}\chi^2\left((v-p)\, n_s,\, s^2\right) \\
\hat{\boldsymbol{Y}} &= \left(\boldsymbol{A}'\boldsymbol{A}\right)^{-1}\boldsymbol{A}'\boldsymbol{X} \\
\boldsymbol{V} &= \left(\boldsymbol{A}'\boldsymbol{A}\right)^{-1} \\
s^2 &= \frac{1}{(v-p)\, n_s}\,\text{trace}\left(\left(\boldsymbol{X}-\boldsymbol{A}\hat{\boldsymbol{Y}}\right)\left(\boldsymbol{X}-\boldsymbol{A}\hat{\boldsymbol{Y}}\right)'\right)\,,
\end{aligned} \tag{B.2}$$

and samples for the sum-of-squares matrix for demeaned $\boldsymbol{Y}$ can be easily drawn, conditional on draws of σ, from a Wishart distribution,

$$\boldsymbol{S} \sim \mathcal{W}_p\left(n_s, \boldsymbol{V}\sigma^2\right)\,. \tag{B.3}$$

It would also be possible to use inference procedures that can account for more complicated or correlated noise models (e.g..Chappell *et al.*, 2009; Penny *et al.*, 2007). However, as the best approach for constructing representative time courses for ROIs is still a matter of much debate (Beckmann *et al.*, 2009;

Calhoun *et al.*, 2001; Filippini *et al.*, 2009; Harrison *et al.*, 2015), and we prefer to focus here on the precision modelling, we opt to use a simple linear regression point-estimate, $\hat{Y}_s = \left(A^\dagger A\right)^{-1} A^\dagger X_s$.

B.2 Estimation of temporal degrees of freedom

Our model for functional connectivity ignores the temporal complexity of the data, assuming that all samples are independently drawn from the GGM structure. However, we cannot ignore the fact that the data we describe have strong temporal autocorrelations, such that not every sample provides new information, independent from its predecessors. Using a poor (over-) estimate of the number of independent samples in the data will cause the Bayesian model to inadequately capture the appropriate level of uncertainty within the posterior distribution.

One approach to dealing with autocorrelated input data would be to temporally whiten the data, but the merits of fitting an auto-regressive model then modelling the covariance structure of the residuals have not been investigated. Motivated partly by the estimation of spatial degrees of freedom in functional neuroimaging studies (Groves *et al.*, 2011; Worsley *et al.*, 1992) and partly by the estimation of the effective sample size of an MCMC chain, we estimate the number of independent data samples in each subject as

$$n_s = \frac{n}{1 + 2\sum_{t=1}^{\infty} \rho_t}, \tag{B.4}$$

where n indicates the number of data points and ρ_t is the temporal autocorrelation of an ROI time course at lag t. We estimate this effective sample size in

each ROI, then simply average over ROIs to get the final estimate. See Gelman *et al.* (2014), ch. 11, for a full discussion in the context of Monte-Carlo sampling and computation of the autocorrelations.

B.3 MCMC sampling scheme

DESCRIPTION OF APPROACH

We propose a Metropolis-within-collapsed-Gibbs (MWCG) sampler for inference of the precision matrices. The basic structure is of conditional Gibbs draws of all parameters that can be identified with particular columns of the matrices. A sample consists of drawing from each column in turn, in a random order. Because of the symmetry within the matrices, this means that these parameters are actually drawn twice within one sample. Some parameters are common to all columns, and are drawn from their conditional distributions separately.

Within each Gibbs draw for a column, the model is largely conditionally conjugate, and many parameters can be marginalised out, collapsing the Gibbs sampler for this subset of parameters (van Dyk & Park, 2008). At the top level, an MH algorithm is used to sample $\boldsymbol{Z}$, the variables identifying which edges are included in the model. The MH sampler is able to explore a conditional distribution which marginalises over all subject-level edge strengths, and the group mean edge strength, substantially reducing the correlations between successive draws. It can be considered as a Bayesian model averaging procedure, in which the sampler accepts or rejects proposed models based on the model evidence (within a column) associated with their particular structure. The group-average connectivity strengths $\boldsymbol{\mu}$ can be drawn conditional on $\boldsymbol{Z}$, marginalised

over the subject-level strengths, which themselves fall out as Gibbs steps at the bottom of the hierarchy.

EDGE-ALLOCATION VECTORS

We infer edge inclusion, for each column of the precision matrices, using a Metropolis-Hastings chain. Following ideas in Peltola *et al.* (2012), we propose changing k edges on each proposal, drawing k from a truncated geometric distribution—a geometric distribution has only one tunable parameter, simplifying the optimisation of the sampler. The proposal consists of deciding in turn, for each of the k edges, whether to add or remove a variable, then picking the altered variable at random. The acceptance probability corresponds to the ratio of model evidences for the posterior with the set of new variables to the set of old variables. The model evidence is computed by averaging out $\boldsymbol{u}^s$ and $\boldsymbol{\mu}_{12}$.

Move proposal

A move of k edges is proposed from a geometric distribution with parameter $l \in (0, 1]$, $p(k) = (1 - l)^{k-1} l$, which has a mode at 1, and which is truncated at either 20 or the number of edges in the column, whichever is smaller. Variables are proposed to be added or removed from the model with a probability of $1/2$, until there are no more variables to add or remove, when the probability becomes 1 in the remaining direction, until k proposals have been made. Taking $\boldsymbol{\gamma}$ to be the edge inclusion variables for a particular column, $\boldsymbol{\gamma} = \boldsymbol{z}_{12}$, we call the proposed model $\boldsymbol{\gamma}^\star$. Borrowing the notation from Peltola *et al.* (2012), we call the sequence of variable changes in the proposal $\boldsymbol{\zeta}^\star = [\zeta_1^\star \dots \zeta_k^\star]$, such that $\boldsymbol{\gamma}^\star$ is produced by applying to $\boldsymbol{\gamma}$ flips of the variables contained in $\boldsymbol{\zeta}^\star$. The proposal distribution for each variable to add, conditional on the previous changes, is

given by the normalised edge inclusion probabilities of all variables not in $\boldsymbol{\gamma}$ or already proposed in $\boldsymbol{\zeta}^\star$. We estimate edge inclusion probabilities m_j from MCMC samples in the warm-up phase, bounding from below at 0.1 and from above at 0.9. The proposal distribution for variables to remove uses the same process, with normalised estimates of the edge exclusion probabilities $1 - m_j$ over the variables in $\boldsymbol{\gamma}$ and not already proposed in $\boldsymbol{\zeta}^\star$. The full proposal distribution is a sequential product of the individual proposals. We further define $\boldsymbol{\zeta}$ to be the set of edge changes which exactly undoes the processes in $\boldsymbol{\zeta}^\star$, to take $\boldsymbol{\gamma}^\star$ back to γ.

$$
\begin{aligned}
q_1\left(\boldsymbol{\gamma}^\star | \boldsymbol{\gamma}, \boldsymbol{\zeta}^\star\right) &= \begin{cases} 1 & \text{if } \boldsymbol{\gamma}^\star \text{ incorporates changes } \boldsymbol{\zeta}^\star \text{ to } \boldsymbol{\gamma} \\ 0 & \text{otherwise} \end{cases} \\
q_1\left(\boldsymbol{\zeta}^\star | \boldsymbol{\gamma}\right) &= \prod_{r=1}^{k} q_1\left(\zeta_r^\star | \boldsymbol{\gamma}, \zeta_{1:r-1}^\star\right) \\
q_1\left(\zeta_r^\star = +j | \boldsymbol{\gamma}, \zeta_{1:r-1}^\star\right) &= \frac{1}{2} \frac{m_j}{\sum_{s \notin \left(\boldsymbol{\gamma} \cup \zeta_{1:r-1}^\star\right)} m_s} \\
q_1\left(\zeta_r^\star = -j | \boldsymbol{\gamma}, \zeta_{1:r-1}^\star\right) &= \frac{1}{2} \frac{1 - m_j}{\sum_{s \in \left(\boldsymbol{\gamma} \setminus \zeta_{1:r-1}^\star\right)} 1 - m_s} \\
h_1\left(\boldsymbol{\zeta} | \boldsymbol{\zeta}^\star, \boldsymbol{\gamma}, \boldsymbol{\gamma}^\star\right) &= \begin{cases} 1 & \text{if } \boldsymbol{\zeta} \text{ is constructed using algorithm B.1} \\ 0 & \text{otherwise} \end{cases}
\end{aligned}
$$

Algorithm B.1: PROCEDURE FOR GENERATING REVERSE PROPOSALS. This algorithm, designed by Peltola *et al.*, is used to construct ζ and $\zeta^{\star\prime}$, the set of steps to undo proposals $\gamma^\star$ and γ'. The algorithm is designed such that it is reversible (generating ζ, then reversing to generate $\zeta^\star$, returns to the initial proposal).

```
def form_backward_move_proposal(zetaS, gamma,  gammaS):
    ## 1.
    # Construct intermediate
    zeta_b = zetaS
    zeta_b = reverse_order_of_removals(zeta_b)
    zeta_b = set_all_operations_to_additions(zeta_b)

    ## 2.
    # convert elements to removals to make
    # a valid proposal from gammaS
    zeta = make_valid(zeta_b, gammaS)
    zeta = reverse_order_of_removals(zeta)

    return zeta
```

Acceptance probability

The jump ratio and acceptance probability for a proposal $(\gamma^\star, \zeta^\star, \zeta)$ are

$$\begin{aligned} r &= \frac{p\left(\gamma^\star|\left\{Y_s\right\}, \chi, \sigma, a\right) q_1\left(\gamma, \zeta|\gamma^\star\right) p\left(k\right) h_1\left(\zeta^\star|\zeta, \gamma^\star, \gamma\right)}{p\left(\gamma|\left\{Y_s\right\}, \chi, \sigma, a\right) q_1\left(\gamma^\star, \zeta^\star|\gamma\right) p\left(k\right) h_1\left(\zeta|\zeta^\star, \gamma, \gamma^\star\right)} \\ &= \frac{p\left(\gamma^\star|\left\{Y_s\right\}, \chi, \sigma, a\right) q_1\left(\gamma|\zeta, \gamma^\star\right) q_1\left(\zeta|\gamma^\star\right) h_1\left(\zeta^\star|\zeta, \gamma^\star, \gamma\right)}{p\left(\gamma|\left\{Y_s\right\}, \chi, \sigma, a\right) q_1\left(\gamma^\star|\gamma, \zeta^\star\right) q_1\left(\zeta^\star|\gamma\right) h_1\left(\zeta|\zeta^\star, \gamma, \gamma^\star\right)}, \end{aligned}$$

$$a_1\left(\gamma^\star|\zeta^\star, \zeta, \gamma\right) = \min\left(1, r\right), \tag{B.5}$$

and the evidence for model γ is derived in section B.6 on page 206 as

$$\begin{aligned} p\left(\gamma|\left\{Y_s\right\}, \chi, \sigma_{12}, a\right) \propto & \frac{1}{\left(2\pi\right)^{\sum_s n_s/2}} |D_\tau^+|^{N/2} a^{|\gamma_+|} \left(1-a\right)^{|\gamma_-|} \\ & \times \prod_s \left\{ |\Delta_+^s|^{1/2} e^{\frac{1}{2} s_{12}^{+\dagger} \Delta_+^s s_{12}^+} \right\} \frac{|\Xi|^{1/2}}{\chi^{2^{|\gamma_+|}/2}} e^{\frac{1}{2}\epsilon_+^\dagger \Xi \epsilon_+} \end{aligned} \tag{B.6}$$

where symbols + and − denote the subsets of variables within the column which are included or excluded, respectively, in the model; $|\gamma_+|$ is the number of included variables in the column; and the remaining variables defined as

$$\boldsymbol{\tau} = \boldsymbol{\sigma}_{12}^{-2} \tag{B.7}$$

$$D_\tau = \text{diag}\left(\boldsymbol{\sigma}_{12}^{-2}\right) \tag{B.8}$$

$$\boldsymbol{\rho}_+^s = \left(D_\tau^+ \boldsymbol{\mu}_+ - \boldsymbol{S}_{12}^+\right) \tag{B.9}$$

$$\boldsymbol{\Delta}_+^s = \left(\boldsymbol{\Upsilon}_+^s + D_\tau^+\right)^{-1} \tag{B.10}$$

$$\boldsymbol{\Xi} = \left(\frac{1}{\chi^2} + ND_\tau^+ - \sum_s D_\tau^+ \boldsymbol{\Delta}_+^s D_\tau^+\right)^{-1} \tag{B.11}$$

$$\boldsymbol{\epsilon}_+^\dagger = \sum_s \boldsymbol{S}_{12}^{+\dagger} \boldsymbol{\Delta}_+^s D_\tau^+ \,. \tag{B.12}$$

An MH update with this acceptance probability will converge to the target distribution (Peltola *et al.*, 2012; Storvik, 2011).

Delayed rejection scheme

The sampler outlined above can be very inefficient in exploring the parameter space, with only 1-10 edge changes in 100 samples on real data. To improve on this, we use an additional technique suggested for sparse linear regression by Peltola *et al.* (2012): the ability to create a second proposal, after the first is rejected, which is in some way conditional on the first but which uses fewer computations to propose and test for acceptance than the creation of an entirely new proposal. Delayed rejection (Green & Mira, 2001; Mira, 2001; Trias *et al.*, 2009) is this process of proposing a secondary move, in such a fashion that detailed balance of the whole chain is preserved. Peltola *et al.* demonstrate that a secondary move can be constructed which is guaranteed to be accepted. The

proposal distribution for this second move picks from a small number of models (possible choices for the vector $\boldsymbol{\gamma}'$), which include the original model $\boldsymbol{\gamma}$. We describe this process now.

If the first proposal $\boldsymbol{\gamma}^\star$ is rejected, which occurs with probability $1 - a_1$, a second proposal $\boldsymbol{\gamma}'$ is created. If k edges were first proposed to be altered, by applying $\boldsymbol{\zeta}^\star$ to $\boldsymbol{\gamma}$, $\boldsymbol{\gamma}'$ is chosen from the 2^k models that consist of applying subsets of $\boldsymbol{\zeta}^\star$ to $\boldsymbol{\gamma}$. Thus if $\boldsymbol{\gamma} = [1100]$ and the first proposal was $\boldsymbol{\zeta}^\star = \{+3, +4, -1\}$ to form $\boldsymbol{\gamma}^\star = [0111]$, the second proposal must be formed by applying one of $\boldsymbol{\zeta}' : \{$ $\{\}$, $\{+3\}$, $\{+4\}$, $\{-1\}$, $\{+3, +4\}$, $\{+3, -1\}$, $\{+4, -1\}$, $\{+3, +4, -1\}$ $\}$ to $\boldsymbol{\gamma}$: $\boldsymbol{\gamma}' = \{$ $[1100]$, $[1110]$, $[1101]$, $[0100]$, $[1111]$, $[0110]$, $[0101]$, $[0111]$ $\}$. This set of models *includes* the original model $\boldsymbol{\gamma}$ and the rejected first proposal $\boldsymbol{\gamma}^\star$.

As the delayed rejection proposal is a *two-step* move, in order to satisfy reversibility we must be able to construct a two-step backwards move, with proposal distribution $q_2\left(\boldsymbol{\gamma}|\boldsymbol{\gamma}', \boldsymbol{\zeta}^{\star\prime}\right)$. This requires its own intermediate proposal, $\boldsymbol{\gamma}^{\star\prime}$, which need not be the same as $\boldsymbol{\gamma}^\star$, but should require the same number of edge flips to reach, k, as the initial proposal.

The proposal distribution q_2 for choosing between this set of models is constructed such that the acceptance probability of this delayed step is unity, $a_2 = 1$. We can see that the full expression for the acceptance probability after a two-stage proposal is

$$a_2\left(\boldsymbol{\gamma}^\star|\boldsymbol{\zeta}^\star, \boldsymbol{\zeta}, \boldsymbol{\gamma}\right) = \min\left(1,\ \frac{p\left(\boldsymbol{\gamma}'|\left\{Y_s\right\}, \chi, \boldsymbol{\sigma}, a\right) q_1\left(\boldsymbol{\gamma}^{\star\prime}, \boldsymbol{\zeta}^{\star\prime}|\boldsymbol{\gamma}'\right) p\left(k\right)}{p\left(\boldsymbol{\gamma}|\left\{Y_s\right\}, \chi, \boldsymbol{\sigma}, a\right) q_1\left(\boldsymbol{\gamma}^\star, \boldsymbol{\zeta}^\star|\boldsymbol{\gamma}\right) p\left(k\right)}\right. \tag{B.13}$$
$$\left. \times \frac{\left(1 - a_1\left(\boldsymbol{\gamma}^{\star\prime}|\boldsymbol{\zeta}^{\star\prime}, \boldsymbol{\zeta}', \boldsymbol{\gamma}'\right)\right) q_2\left(\boldsymbol{\gamma}|\boldsymbol{\gamma}', \boldsymbol{\zeta}^{\star\prime}\right)}{\left(1 - a_1\left(\boldsymbol{\gamma}^\star|\boldsymbol{\zeta}^\star, \boldsymbol{\zeta}, \boldsymbol{\gamma}\right)\right) q_2\left(\boldsymbol{\gamma}'|\boldsymbol{\gamma}, \boldsymbol{\zeta}^\star\right)} \right),$$

where as before, $\boldsymbol{\zeta}^{\star\prime}$ is constructed using algorithm B.1. The secondary proposal

is therefore guaranteed to be accepted for

$$q_2\left(\boldsymbol{\gamma}'|\boldsymbol{\gamma},\boldsymbol{\zeta}^{\star}\right) \propto p\left(\boldsymbol{\gamma}'|\left\{Y_s\right\},\chi,\boldsymbol{\sigma},a\right) q_1\left(\boldsymbol{\gamma}^{\star\prime},\boldsymbol{\zeta}^{\star\prime}|\boldsymbol{\gamma}'\right) \times p\left(k\right)\left(1-a_1\left(\boldsymbol{\gamma}^{\star\prime}|\boldsymbol{\zeta}^{\star\prime},\boldsymbol{\zeta}',\boldsymbol{\gamma}'\right)\right), \tag{B.14}$$

which is the product of the probability of the target model, $\boldsymbol{\gamma}'$, the jumping probability to the intermediate proposal in the reverse direction, and the probability of rejecting that proposal. The distribution is normalised over the range of models considered for $\{\boldsymbol{\gamma}'\}$. Sampling from q_2 involves re-computing (B.6) with sequential updates to a set of Cholesky decompositions as variables are added or removed. It is possible to do this with an efficient implementation, in which the computational complexity scales as $\mathcal{O}\left(N2^k k\right)$; see Clarke (1981); Peltola *et al.* (2012); and Smith (1989) for details.

Utilising this delayed rejection scheme can greatly increase the efficiency on real data, realising around 25 edge changes in 100 samples, an improvement of two- to ten-fold on performance without it.

Construction of reverse move

We use the construction of a backward proposal set out in Peltola *et al.* (2012) to build $\boldsymbol{\zeta}^{\star\prime}$ and, for consistency, $\boldsymbol{\zeta}$. The algorithm (B.1 on page 194) is designed such that it is reversible (generating $\boldsymbol{\zeta}$, then reversing to generate $\boldsymbol{\zeta}^{\star}$, returns to the initial proposal), and such that $q_2\left(\boldsymbol{\gamma}'|\boldsymbol{\gamma},\boldsymbol{\zeta}^{\star}\right) = q_2\left(\boldsymbol{\gamma}'|\boldsymbol{\gamma}',\boldsymbol{\zeta}^{\star\prime}\right)$, which ensures that the probabilities in the distribution of the delayed proposal are independent of which primary model they are conditioned upon.

BETWEEN-SUBJECT VARIANCE

Within each column, we sample the between-subject variance parameters $\boldsymbol{\sigma}_{12}$ using an MH kernel. A Gibbs draw conditional on the subject-level connection strengths is possible, but we found that this led to highly correlated behaviour.

The target density is derived in section B.6 on page 206,

$$\begin{aligned} p\left(\boldsymbol{\sigma}_{12}|\chi,\boldsymbol{\gamma},\boldsymbol{\mu}_{12},\{Y_s\}\right) \propto & \frac{1}{(2\pi)^{\sum_s n_s/2}}|D_\tau^+|^{N/2}e^{-\frac{N}{2}\boldsymbol{\mu}_+^\dagger D_\tau^+\boldsymbol{\mu}_+} \\ & \times\prod_s\left\{|\boldsymbol{\Delta}_+^s|^{1/2}e^{\frac{1}{2}\boldsymbol{\rho}_+^{s\dagger}\boldsymbol{\Delta}_+^s\boldsymbol{\rho}_+^s}\right\} \\ & \times\prod_i\left\{\frac{1}{\sigma_{i\cdot}\left(2\pi s_\sigma^2\right)^{\frac{1}{2}}}e^{-\frac{(\log\sigma_{i\cdot}-\log m_\sigma)^2}{2s_\sigma^2}}\right\}. \end{aligned} \tag{B.15}$$

We use a proposal density that is multivariate normal on $\log\boldsymbol{\sigma}_{12}$,

$$\begin{aligned} q\left(\boldsymbol{\sigma}'|\boldsymbol{\sigma}\right) = & \frac{1}{\prod_i\sigma_{i\cdot}\left(2\pi\right)^{q/2}}|c^2\boldsymbol{\Sigma}_{\log\sigma}|^{-1/2} \\ & \times e^{-\frac{1}{2}(\log\boldsymbol{\sigma}-\log\boldsymbol{\sigma}')^T\left(c^2\boldsymbol{\Sigma}_{\log\sigma}\right)^{-1}(\log\boldsymbol{\sigma}-\log\boldsymbol{\sigma}')}, \end{aligned} \tag{B.16}$$

with a scaling on the covariance matrix of c^2. This gives an acceptance probability of

$$a\left(\boldsymbol{\sigma}_{12}{}'|\boldsymbol{\sigma}_{12}\right) = \min\left(1, \frac{p\left(\boldsymbol{\sigma}_{12}{}'|\chi,\boldsymbol{\gamma},\boldsymbol{\mu}_{12},\{Y_s\}\right)}{p\left(\boldsymbol{\sigma}_{12}|\chi,\boldsymbol{\gamma},\boldsymbol{\mu}_{12},\{Y_s\}\right)}\frac{\prod_i\sigma_{i\cdot}}{\prod_i\sigma_{i\cdot}{}'}\right). \tag{B.17}$$

ADAPTIVE SAMPLING PHASE

We use a *warm-up period* (an initial set of samples, often the first half) to automatically adapt the properties of our sampler. At the end of the warm-up, we fix all values and let the sampler continue into the *sampling phase* proper. Adjusting

the properties of the kernel during the warm-up breaks the detailed balance conditions of the MCMC chain, but this is not problematic if the samples from the warm-up period are not used for inference (Gelman *et al.*, 2014). Four parameters are adapted. For the proposals on $\boldsymbol{Z}$, the edge inclusion probabilities $\boldsymbol{M}$ and the parameter for the geometric distribution which controls how many edges are proposed, l, are updated. For the proposals on the between-subject variance parameters, $\boldsymbol{\sigma}$, the covariance matrix and scale of the kernel are adjusted.

The covariance structure of $\log \boldsymbol{\sigma}$ and the edge inclusion probabilities can be periodically updated using the history of samples in the warm-up period. To tune l and c, however, a criterion for the optimal acceptance rate of the respective MH samplers is needed. Various authors propose different heuristics for this target. David MacKay (MacKay, 2003) motivates a rule-of-thumb acceptance rate of 0.5 as maximising the rate of information transfer about the sampled distribution, viewing the sampler as a noisy communication channel. More theoretical work by Gelman *et al.* suggests an acceptance rate of 0.44 for unidimensional systems, dropping to 0.23 for dimensionalities greater than 5 (Gelman *et al.*, 2014, 1996; Roberts *et al.*, 1997); this may be achieved for Gaussian proposal distributions and approximately normal posteriors on i.i.d. variables by choosing a jumping kernel which has a covariance matrix scaled relative to the target distribution by a factor of $2.38/\sqrt{d}$ (for dimensionality d). Alternatively, for highly non-Gaussian systems, the expected squared jump distance for the Markov chain can be maximised (Pasarica & Gelman, 2010). Peltola *et al.* take this last approach, deriving the optimisation target for the binary variables in their sparse linear regression problem. On simulated and real data, our sampler never approaches optimal acceptance rates (presumably because of the high correlations induced over samples by the Gibbs structure over columns), and it suffices to maximise

the number of edges changed per sample.

In total, we divide sampling into five phases.

1. Sample from the full model to allow all parameters to converge without the sparsity constraints. Set $c = 0.5$, $\boldsymbol{\Sigma}_{\log \sigma} = \boldsymbol{I}$ and $m_{ij} = 0.5$.

2. Introduce the MH sampler on $\boldsymbol{Z}$, moving one edge at a time (for example, by setting the geometric distribution parameter l to 0.999). This allows the sampler to move towards a likely model structure. Set $\boldsymbol{\Sigma}_{\log \sigma}$ to match the covariance of the sampled values of $\boldsymbol{\sigma}_{12}$ about the posterior mean, and set the scaling for its MH kernel $c = {}^{2.38}/_{\sqrt{q}}$. Periodically update the covariance structure for the sampler on $\boldsymbol{\Sigma}_{\log \sigma}$ and adjust c with proportional increases or decreases to achieve an acceptance rate for samples in $\log \boldsymbol{\sigma}$ of about 0.23.

3. Start adjustment of l. A bracketing system is used to identify the optimal value, which maximises the number of edges accepted per sample. Periodically update the estimate of edge inclusion probabilities $\boldsymbol{M}$, using the history of samples of $\boldsymbol{Z}$. Fix $\boldsymbol{\Sigma}_{\log \sigma}$ and c in this phase.

4. Last phase of the warm-up. Fix l. Continue adjustment of edge inclusion probabilities $\boldsymbol{M}$ and the covariance of $\log \boldsymbol{\sigma}$, $\boldsymbol{\Sigma}_{\log \sigma}$.

5. Sampling phase. Set all sampling hyperparameters ($\boldsymbol{M}$, l, $\boldsymbol{\Sigma}_{\log \sigma}$ and c). Only these samples are used for inference.

SAMPLING ALGORITHM

Algorithm B.2: MCMC SAMPLING ALGORITHM for Hierarchical Inference of Posterior Precisions in OSL (HIPPO).

1. Draw top-level probabilities of edge inclusion from a Beta distribution, incorporating the relative proportions of edges present or absent in the model.

$$a|\boldsymbol{Z}, a_\pi, b_\pi \sim \text{Beta}\left(a_\pi + \sum_{i<j} z_{ij}, b_\pi + n_E - \sum_{i<j} z_{ij}\right) \tag{B.18}$$

   where $n_E = \frac{p(p-1)}{2}$ is the total number of edges.

2. Sample auxiliary variable σ_η as a variance parameter for edges of strength η_{ij}.

$$\sigma_\eta^2|\boldsymbol{\eta}, A \sim \text{Inv-}\chi^2\left(\sum_{i<j} z_{ij} + 1, \frac{A + \sum_{i<j} z_{ij}\eta_{ij}^2}{1 + \sum_{i<j} z_{ij}}\right) \tag{B.19}$$

3. Sample auxiliary variable ξ, which takes the form of a posterior from linear regression over all edges.

$$\xi|\left\{\boldsymbol{\Omega}_s\right\}, \boldsymbol{\sigma}, \boldsymbol{\eta} \sim \mathcal{N}\left(\frac{\sum_{i<j} \frac{z_{ij}}{\sigma_{ij}^2} \sum_s \omega_{ij}^s \eta_{ij}}{1 + \sum_{i<j} \frac{z_{ij} N \eta_{ij}^2}{\sigma_{ij}^2}}, \frac{1}{1 + \sum_{i<j} \frac{z_{ij} N \eta_{ij}^2}{\sigma_{ij}^2}}\right) \tag{B.20}$$

4. Sample the regularisation on the diagonal for each subject, which is

Gamma distributed under its conjugate prior.

$$\lambda_s|\mathbf{\Omega}_s \sim \text{Ga}\left(\frac{1}{3}+p, \sum_i \frac{\omega_{ii}^s}{2}\right) \tag{B.21}$$

5. Reform $\chi = |\xi|\sigma_\eta$

6. Choose column j at random from $\{1 \dots p\}$, without replacement.

   (a) Re-arrange rows and columns of each set of matrices, so that column j is at the end.

   (b) Compute the principal minor $\mathbf{\Omega}_{11}^s$, its inverse

$$\mathbf{\Omega}_{11}^s{}^{-1} = \mathbf{\Sigma}_{11}^s - \mathbf{\Sigma}_{12}^s \mathbf{\Sigma}_{12}^s{}^\dagger / \Sigma_{22}^s \,, \tag{B.22}$$

   the diagonal matrix of inverse-variances $D_\tau = \text{diag}\left(\sigma_{ij}^{-2}\right)$ and, for those edges in the model,

$$\mathbf{\Upsilon}^s = (s_{22} + \lambda^s)\, \mathbf{\Omega}_{11}^s{}^{-1} \tag{B.23}$$

$$\mathbf{\Delta}_+^s = \left(\mathbf{\Upsilon}_+^s + D_\tau^+\right)^{-1} . \tag{B.24}$$

   (c) Sample edge inclusion variables, $\boldsymbol{\gamma} = \boldsymbol{z}_{12}$, drawing values of σ_{ij}, η_{ij} from their priors (for proposed additional edges) when necessary to compute the acceptance ratio'

   Generate a proposal move $(\boldsymbol{\gamma}^\star, \boldsymbol{\zeta}^\star, \boldsymbol{\zeta})$ using proposal distribution

$$q_1\left(\boldsymbol{\gamma}^\star, \boldsymbol{\zeta}^\star, \boldsymbol{\zeta}, k|\boldsymbol{\gamma}\right) = p\,(k)\, q_1\left(\boldsymbol{\gamma}^\star|\boldsymbol{\gamma}, \boldsymbol{\zeta}^\star\right) q_1\left(\boldsymbol{\zeta}^\star|\boldsymbol{\gamma}\right) \tag{B.25}$$

   and accept with probability $a_1\,(\boldsymbol{\gamma}^\star|\boldsymbol{\zeta}^\star, \boldsymbol{\zeta}, \boldsymbol{\gamma})$ (B.5). Perform a low-rank

update to $\boldsymbol{\Delta}_+^s$ if necessary.

If this first proposal is rejected, perform a delayed rejection step, accepting with certainty one of the models formed by applying a subset of elements of $\boldsymbol{\zeta}^\star$ to $\boldsymbol{\gamma}$. The accepted model is drawn from the proposal distribution $q_2\left(\boldsymbol{\gamma}'|\boldsymbol{\gamma}, \boldsymbol{\zeta}^\star\right)$ (B.14).

(d) Sample contributions to group mean edge strengths, integrating out subject-level strengths.

$$\boldsymbol{\eta}_{12}^+|\sigma_\eta, \xi, \{Y_s\}, \boldsymbol{\sigma} \sim \mathcal{N}\left(-\boldsymbol{\Xi}^+ \sum_s \left(D_\tau^+ \boldsymbol{\Delta}_+^s S_{12}^+\right)/\xi,\ \boldsymbol{\Xi}^+/\xi^2\right) \tag{B.26}$$

for $\boldsymbol{\Xi}^{+-1} = \frac{1}{\chi^2} + ND_\tau^+ - \sum_s \left(D_\tau^+ \boldsymbol{\Delta}_+^s D_\tau^+\right)$.

(e) Reform $\mu_{ij} = \xi\eta_{ij}$.

(f) Sample substituted variable v^s for each subject.

$$v^s|\boldsymbol{Y}_s, \lambda_s \sim \text{Ga}\left(\frac{n_s}{2} + 1, \frac{S_{22} + \lambda_s}{2}\right) \tag{B.27}$$

(g) Sample edge strengths for each subject.

$$\boldsymbol{u}_s| \sim \mathcal{N}\left(\boldsymbol{\Delta}_+^s \left(D_\tau^+ \boldsymbol{\mu}_+ - \boldsymbol{S}_{12}^+\right), \boldsymbol{\Delta}_+^s\right) \tag{B.28}$$

(h) Undo parametrisations $\boldsymbol{\omega}_{12}^s = \boldsymbol{u}^s$ and $\omega_{22}^s = v^s + \boldsymbol{u}^{s\dagger} {\boldsymbol{\Omega}_{11}^s}^{-1} \boldsymbol{u}^s$.

(i) Update the subject covariance matrices in block fashion.

$$\boldsymbol{\Sigma}^s_{12} = -\boldsymbol{\Omega}^s_{11}{}^{-1}\boldsymbol{u}^s/v^s \tag{B.29}$$

$$\Sigma^s_{22} = 1/v^s \tag{B.30}$$

$$\boldsymbol{\Sigma}^s_{11} = \boldsymbol{\Omega}^s_{11}{}^{-1} + \left(\boldsymbol{\Omega}^s_{11}{}^{-1}\boldsymbol{u}^s\right)\left(\boldsymbol{\Omega}^s_{11}{}^{-1}\boldsymbol{u}^s\right)^\dagger/v^s \tag{B.31}$$

(j) Sample the subject-level variance, for those edges for which $z_{ij} = 1$, using an MH sampler with a Gaussian proposal distribution on $\log \boldsymbol{\sigma}_{12}$ to draw from $p\left(\boldsymbol{\sigma}_{12}|\chi, \boldsymbol{\gamma}, \boldsymbol{\mu}_{12}, \{Y_s\}\right)$.

(k) Return to loop over columns until complete.

B.4 Initialisation and assessment of convergence

It is good practice in Bayesian computation to run several MCMC chains from different starting points, then assess the quality of the inference through a process of calculating heuristics and inspecting the chains (Gelman *et al.*, 2014). Common heuristics include assessing the effective number of samples in the chains n_{eff} (using the same formula that we employ to estimate temporal degrees of freedom in our data, equation B.4), and requiring that the Gelman-Rubin scale reduction statistic $\hat{R}$ falls below 1.1 for all parameters (Gelman & Rubin, 1992). However, both of these estimates are only valid for approximately normally-distributed parameters. Furthermore, we have discussed in chapter 5 the impossibility of completely exploring the parameter space of the sparse model. The Metropolis sampler is able to converge very rapidly to plausible models, but in our experience tends to get stuck in local minima.

We therefore take a two-pronged, pragmatic approach. Firstly, we initialise

the sampler using a quick exploration of the target space. We run a large number of model chains (say 10) for a short amount of time (say 5000 samples). Typically, one chain will find a set of models with consistently higher probabilities than the other chains. We seed the full inference run using a subset of models from this first bout as initialisation points for the network structure, picking model structures using an importance sampling algorithm on the exploration samples. This enables us to initialise inference from a region of high probability.

Secondly, we perform inference using several sampling chains (perhaps 4), for as long as we can afford computationally. Afterwards, we inspect the chains visually, and compute n_{eff} and $\hat{R}$ for the non-binary variables. If the chains have explored very different parts of the probability space, we can either run the simulations for longer, restart the inference with new initialisation points, or make do with a poor characterisation of the posterior. This method is similar in style to the practical attitude taken in the variable selection literature, searching and averaging over a subset of plausible models (George & McCulloch, 1993, 1997; Raftery *et al.*, 1997). While not characterising in full the highly multi-modal posteriors, if it is able to improve on the performance of previous algorithms for covariance modelling and provide interpretable inference, then we feel that our approach is justified.

B.5 Posterior summaries

With the caveats mentioned above held in mind, the set of samples from the MCMC chains can provide (approximately) a complete description of the posterior. Typically, the posterior over particular variables can be summarised with the mean and variance of the relevant samples. Two small subtleties pertain.

Firstly, we summarise the graphical model structure not with the mode of the distribution (the highest probability model), but with the median probability model (constructed by taking the mean of Z across all samples), for which there is greater theoretical support (Barbieri & Berger, 2004; Guan & Stephens, 2011). Secondly, we have drawn samples of the precision and covariance matrices for each subject, but we are more interested in the correlation and partial correlation matrices as expressions of functional connectivity. Following Marrelec *et al.* (2006), we simply rescale each sample, computing partial correlations as

$$\rho_{ij} = -\frac{\omega_{ij}}{\sqrt{\omega_{ii}\omega_{jj}}} \tag{B.32}$$

and construct samples of population-level effects by averaging the partial correlation matrices across subjects.

B.6 Conditional distributions

We derive the conditional distributions needed for the Gibbs draws in the MCMC scheme.

We start with the conditional distribution for all variables grouped into columns, taking the last column without loss of generality. We perform variable substitutions (5.6)

$$\begin{aligned}(\boldsymbol{u}^s,\, v^s) &= \left(\boldsymbol{\omega}_{12}^s,\, \omega_{22}^s - \boldsymbol{\omega}_{12}^{s\,\dagger} {\boldsymbol{\Omega}_{11}^s}^{-1} \boldsymbol{\omega}_{12}^s\right) \\ \boldsymbol{\Upsilon}^s &= (s_{22} + \lambda^s)\, {\boldsymbol{\Omega}_{11}^s}^{-1},\end{aligned} \tag{B.33}$$

to yield

$$p\left(\boldsymbol{\gamma}, \boldsymbol{\mu}_{12}, \boldsymbol{\sigma}_{12}, \boldsymbol{u} | \left\{Y_s\right\}, -\right) \propto \prod_s \left\{ \frac{1}{(2\pi)^{n_s/2}} e^{-\frac{1}{2}\boldsymbol{u}_+^{s\dagger}\boldsymbol{\Upsilon}_+^s\boldsymbol{u}_+^s - S_{12}^{+\dagger}\boldsymbol{u}_+^s} \frac{|D_\tau^+|^{1/2}}{(2\pi)^{|\gamma_+|/2}} e^{-\frac{1}{2}\left(\boldsymbol{u}_+^s - \boldsymbol{\mu}_+\right)^\dagger D_\tau^+ \left(\boldsymbol{u}_+^s - \boldsymbol{\mu}_+\right)} \right.$$
$$\left. \times e^{-\frac{1}{2}\boldsymbol{u}_-^{s\dagger}\boldsymbol{\Upsilon}_-^s\boldsymbol{u}_-^s - S_{12}^{-\dagger}\boldsymbol{u}_-^s} \delta\left(\boldsymbol{u}_-^s\right) \right\}$$
$$\times a^{|\gamma_+|}(1-a)^{|\gamma_-|}\, \pi\left(\boldsymbol{\mu}_{12}\right) \pi\left(\boldsymbol{\sigma}_{12}\right), \tag{B.34}$$

where $\boldsymbol{\gamma} = \boldsymbol{z}_{12}$, the edge inclusion variables for the column; symbols $+$ and $-$ denote the subsets of variables within the column that are included or excluded, respectively, in the model; $|\gamma_+|$ is the number of included variables in the column; $D_\tau = \operatorname{diag}\left(\boldsymbol{\sigma}_{12}^{-2}\right)$; and $\pi(x)$ indicates the prior on variable x. Complete the square to find

$$p\left(\boldsymbol{\gamma}, \boldsymbol{\mu}_{12}, \boldsymbol{\sigma}_{12}, \boldsymbol{u} | \left\{Y_s\right\}, -\right) \propto \frac{1}{(2\pi)^{\sum_s n_s/2}} \frac{1}{(2\pi)^{N|\gamma_+|/2}} |D_\tau^+|^{N/2} a^{|\gamma_+|}(1-a)^{|\gamma_-|}\, e^{-\frac{N}{2}\boldsymbol{\mu}_+^\dagger D_\tau^+ \boldsymbol{\mu}_+}$$
$$\times \prod_s \left\{ e^{-\frac{1}{2}\boldsymbol{u}_+^{s\dagger}\boldsymbol{\Delta}_+^{s\,-1}\boldsymbol{u}_+^s + \boldsymbol{\rho}_+^{s\dagger}\boldsymbol{u}_+^s} \right\} \pi\left(\boldsymbol{\mu}_{12}\right) \pi\left(\boldsymbol{\sigma}_{12}\right),$$

where $\boldsymbol{\rho}_+^s = \left(D_\tau^+ \boldsymbol{\mu}_+ - S_{12}^+\right)$ and $\boldsymbol{\Delta}_+^s = \left(\boldsymbol{\Upsilon}_+^s + D_\tau^+\right)^{-1}$.

This immediately gives the conditional for $\boldsymbol{u}$,

$$\boldsymbol{u}_+^s \sim \mathcal{N}\left(\boldsymbol{\Delta}_+^s \left(D_\tau^+ \boldsymbol{\mu}_+ - S_{12}^+\right), \boldsymbol{\Delta}_+^s\right) \tag{B.35}$$
$$\boldsymbol{u}_-^s = 0\,. \tag{B.36}$$

The subject-level variables $\boldsymbol{u}$ can be marginalised over,

$$\begin{aligned} p\left(\boldsymbol{\gamma}, \boldsymbol{\mu}_{12}, \boldsymbol{\sigma}_{12} | \left\{Y_s\right\}, -\right) \propto & \frac{1}{\left(2\pi\right)^{\sum_s n_s/2}} \frac{1}{\left(2\pi\right)^{N|\gamma_+|/2}} |D_\tau^+|^{N/2} a^{|\gamma_+|} \left(1-a\right)^{|\gamma_-|} e^{-\frac{N}{2}\boldsymbol{\mu}_+^\dagger D_\tau^+ \boldsymbol{\mu}_+} \\ & \times \prod_s \left\{ \int \mathrm{d}\boldsymbol{u}_s^+ \, e^{-\frac{1}{2}\boldsymbol{u}_+^{s\dagger} {\boldsymbol{\Delta}_+^s}^{-1} \boldsymbol{u}_+^s + \boldsymbol{\rho}_+^{s\dagger} \boldsymbol{u}_+^s} \right\} \pi\left(\boldsymbol{\mu}_{12}\right) \pi\left(\boldsymbol{\sigma}_{12}\right) \\ \propto & \frac{1}{\left(2\pi\right)^{\sum_s n_s/2}} |D_\tau^+|^{N/2} a^{|\gamma_+|} \left(1-a\right)^{|\gamma_-|} e^{-\frac{N}{2}\boldsymbol{\mu}_+^\dagger D_\tau^+ \boldsymbol{\mu}_+} \\ & \times \prod_s \left\{ |\boldsymbol{\Delta}_+^s|^{1/2} e^{\frac{1}{2}\boldsymbol{\rho}_+^{s\dagger} \boldsymbol{\Delta}_+^s \boldsymbol{\rho}_+^s} \right\} \pi\left(\boldsymbol{\mu}_{12}\right) \pi\left(\boldsymbol{\sigma}_{12}\right) , \end{aligned} \tag{B.37}$$

giving the conditional distribution for $\boldsymbol{\mu}_{12}$, marginal over $\boldsymbol{u}_s$,

$$p\left(\boldsymbol{\mu}_{12} | \chi, \boldsymbol{\gamma}, \boldsymbol{\sigma}_{12}, \left\{Y_s\right\}\right) \propto \frac{1}{\left(2\pi\chi^2\right)^{\frac{|\gamma_-|}{2}}} e^{-\frac{1}{2}\boldsymbol{\mu}_-\boldsymbol{\mu}_-/\chi^2} \frac{1}{\left(2\pi\chi^2\right)^{\frac{|\gamma_+|}{2}}} e^{-\frac{1}{2}\boldsymbol{\mu}_+ \boldsymbol{\Xi}^{-1} \boldsymbol{\mu}_+ - \boldsymbol{\epsilon}_+^\dagger \boldsymbol{\mu}_+} \tag{B.38}$$

$$\begin{aligned} p\left(\boldsymbol{\eta} | \xi, \sigma_\eta, \boldsymbol{\gamma}, \boldsymbol{\sigma}_{12}, \left\{Y_s\right\}\right) & \propto \xi^q \, p\left(\xi\boldsymbol{\eta} | \chi, \boldsymbol{\gamma}, \boldsymbol{\sigma}_{12}, \left\{Y_s\right\}\right) \\ & \propto \xi^q \frac{1}{\left(2\pi\chi^2\right)^{\frac{|\gamma_-|}{2}}} e^{-\frac{1}{2}\xi^2\boldsymbol{\eta}_-\boldsymbol{\eta}_-/\chi^2} \frac{1}{\left(2\pi\chi^2\right)^{\frac{|\gamma_+|}{2}}} e^{-\frac{1}{2}\xi^2\boldsymbol{\eta}_+ \boldsymbol{\Xi}^{-1} \boldsymbol{\eta}_+ - \xi\boldsymbol{\epsilon}_+^\dagger \boldsymbol{\eta}_+} , \end{aligned} \tag{B.39}$$

where $q = p-1$, the number of edges in a column, $\boldsymbol{\Xi} = \left(\frac{1}{\chi^2} + N D_\tau^+ - \sum_s D_\tau^+ \boldsymbol{\Delta}_+^s D_\tau^+\right)^{-1}$ and $\boldsymbol{\epsilon}_+^\dagger = \sum_s \boldsymbol{S}_{12}^{+\dagger} \boldsymbol{\Delta}_+^s D_\tau^+$.

The conditional distribution for $\boldsymbol{\sigma}_{12}$ can also be drawn out here,

$$\begin{aligned} p\left(\boldsymbol{\sigma}_{12}|\chi,\boldsymbol{\gamma},\boldsymbol{\mu}_{12},\{Y_s\}\right) \propto & \frac{1}{(2\pi)^{\sum_s n_s/2}}|D_\tau^+|^{N/2}e^{-\frac{N}{2}\boldsymbol{\mu}_+^\dagger D_\tau^+\boldsymbol{\mu}_+} \\ & \times\prod_s\left\{|\boldsymbol{\Delta}_+^s|^{1/2}e^{\frac{1}{2}\boldsymbol{\rho}_+^{s\dagger}\boldsymbol{\Delta}_+^s\boldsymbol{\rho}_+^s}\right\} \\ & \times\prod_i\left\{\frac{1}{\sigma_{i.}\left(2\pi s_\sigma^2\right)^{\frac{1}{2}}}e^{-\frac{(\log\sigma_{i.}-\log m_\sigma)^2}{2s_\sigma^2}}\right\}. \end{aligned} \tag{B.40}$$

For the edge inclusion variables, we can further integrate out $\boldsymbol{\mu}_{12}$, in the same manner,

$$\begin{aligned} p\left(\boldsymbol{\gamma}|\{Y_s\},\chi,\boldsymbol{\sigma}_{12},a\right) \propto & \frac{1}{(2\pi)^{\sum_s n_s/2}}|D_\tau^+|^{N/2}a^{|\gamma_+|}(1-a)^{|\gamma_-|}\prod_s|\boldsymbol{\Delta}_+^s|^{1/2} \\ & \times\int \mathrm{d}\boldsymbol{\mu}_{12}\, e^{-\frac{N}{2}\boldsymbol{\mu}_+^\dagger D_\tau^+\boldsymbol{\mu}_+}\prod_s\left\{e^{\frac{1}{2}\boldsymbol{\rho}_+^{s\dagger}\boldsymbol{\Delta}_+^s\boldsymbol{\rho}_+^s}\right\}\frac{1}{(2\pi\chi^2)^{\frac{q}{2}}}e^{-\mu_{12}^2/2\chi^2} \\ \propto & \frac{1}{(2\pi)^{\sum_s n_s/2}}|D_\tau^+|^{N/2}a^{|\gamma_+|}(1-a)^{|\gamma_-|}\prod_s\left\{|\boldsymbol{\Delta}_+^s|^{1/2}e^{\frac{1}{2}s_{12}^{+\dagger}\boldsymbol{\Delta}_+^s s_{12}^+}\right\} \\ & \times\int \mathrm{d}\boldsymbol{\mu}_+\, e^{-\frac{1}{2}\boldsymbol{\mu}_+^\dagger\boldsymbol{\Xi}^{-1}\boldsymbol{\mu}_+-\sum_s s_{12}^{+\dagger}\boldsymbol{\Delta}_+^s D_\tau^+\boldsymbol{\mu}_+}\frac{1}{(2\pi\chi^2)^{\frac{|\gamma_+|}{2}}} \\ \propto & \frac{1}{(2\pi)^{\sum_s n_s/2}}|D_\tau^+|^{N/2}a^{|\gamma_+|}(1-a)^{|\gamma_-|} \\ & \times\prod_s\left\{|\boldsymbol{\Delta}_+^s|^{1/2}e^{\frac{1}{2}s_{12}^{+\dagger}\boldsymbol{\Delta}_+^s s_{12}^+}\right\}\frac{|\boldsymbol{\Xi}|^{1/2}}{\chi^{2^{|\gamma_+|/2}}}e^{\frac{1}{2}\boldsymbol{\epsilon}_+^\dagger\boldsymbol{\Xi}\boldsymbol{\epsilon}_+}. \end{aligned} \tag{B.41}$$

List of Figures

1.1 The default mode network 5
1.2 Resting-state networks 6
1.3 Switching between networks in the resting state 9

2.1 Success of orthogonalisation for seed-based connectivity analysis in MEG 32
2.2 A symmetric, multivariate orthogonalisation process 39
2.3 Simulation of a five-dipole network 44
2.4 Illustration of the pipeline for MEG functional connectivity estimation 48
2.5 Reconstruction of a five-dipole network: comparison of ability to infer a simulated network structure under different corrections for source leakage 52
2.6 Control of false positives during network reconstruction 54
2.7 Control of false positives during reconstruction of a dense network 55
2.8 Group-average resting-state direct functional connections derived from eight subjects, inferred with and without the application of a multivariate spatial leakage correction between ROIs 57
2.9 Alpha-band resting-state network structure surviving the 5% false discovery rate correction for multiple comparisons 58
2.10 Group-average resting-state network matrices 59
2.11 Group-average resting-state correlation matrices 61
2.12 Group-average resting-state unregularised partial correlation matrices 62
2.13 Alpha-band (8-13 Hz) resting-state partial correlations, with a comparison between orthogonalisation methods 63
2.14 Illustration of the effect of the symmetric orthogonalisation process on a time course 64

3.1 Mean alpha-band network matrices at group level 90

3.2 Mean alpha-band network matrices at the group level for all metrics under test 91
3.3 Consistency of network matrix estimation 94

4.1 Contribution of genetic factors and environmental effects shared between twins to functional connectivity 113
4.2 Heritabilities of fMRI connection strengths 116
4.3 Heritability of cortical curvature, and the spatial profile of the heritability of functional connections 116

5.3 Comparison of sparse network modelling methods on simulated datasets 160
5.4 Coverage of posterior intervals on simulated datasets 162
5.5 Scaling of computational cost with changing model dimension 163
5.6 Estimation of network matrices from small samples of data 165
5.8 Displacement of network matrices from their unregularised forms 169
5.9 Performance of the hierarchical model 170
5.10 Group average networks for the MEG beta band 171

List of Tables

3.1 Index of ROI numbers 80
3.2 Categorisation of the tested network metrics 82

4.1 Parameter estimates for the ACE model on functional connectivity 123
4.2 Significance of the heritability of functional connectivity 124
4.3 Correlations between the heritability of cortical curvature and the heritability of functional connectivity 125

5.1 Values of hyperparameters 144
5.2 Description of simulated datasets 148
5.3 Characterisation of methods 150
5.4 Comparative computational cost of sparse Bayesian models 163

List of Abbreviations

AEC	Amplitude Envelope Correlation
APACE	Advanced Permutation inference for ACE models
AR	Auto-Regressive
BOLD	Blood Oxygenation Level Dependent
CCA	Canonical Correlation Analysis
CI	Confidence Interval
CNEURO	Centro de Neurociencias de Cuba
COH	Absolute Coherence
DZ	Dizygotic (twins who share 50% of their DNA)
EEG	Electroencephalography
FDR	False Discovery Rate
FIL	Leopold Muller Functional Imaging Laboratory
fMRI	functional Magnetic Resonance Imaging
FMRIB	Oxford centre for Functional MRI of the Brain
FPR	False-Positive Rate
FSL	FMRIB's Software Library
GGM	Gaussian Graphical Model
GLASSO	Graphical LASSO
HCP	Human Connectome Project
HIPPO	Hierarchical Inference of Posterior Precisions in OSL
HMC	Hamiltonian Monte Carlo
ICA	Independent Component Analysis
i.i.d.	Independent and identically distributed
IMC	Imaginary part of Coherency
IMPC	Imaginary Partial Coherency
LASSO	Least Absolute Shrinkage and Selection Operator
MCMC	Markov Chain Monte Carlo
MEG	Magnetoencephalography
MH	Metropolis-Hastings
MI	Mutual Information

MNI	Montreal Neurological Institute
MRI	Magnetic Resonance Imaging
MWCG	Metropolis-within-collapsed-Gibbs
MZ	Monozygotic (twins who share 100% of their DNA)
NIH	National Institutes of Health (US)
OHBA	Oxford Centre for Human Brain Activity
OSL	OHBA's Software Library
PAEC	Amplitude Envelope Partial Correlation
PCA	Principal Component Analysis
PCOH	Partial Coherence
PDC	Partial Directed Coherence
PET	Positron Emission Tomography
PLI	Phase Lag Index
PLV	Phase Locking Value
PSI	Phase Slope Index
RMS	Root Mean Square
ROI	Region of Interest
RSN	Resting-State Network
SNR	Signal-to-Noise Ratio
SPM	Statistical Parametric Mapping
SSVS	Stochastic Search Variable Selection
UCL	University College London
WPLI	weighted Phase Lag Index

References

Andrews-Hanna, J.R., Snyder, A.Z., Vincent, J.L., Lustig, C., Head, D., Raichle, M.E. & Buckner, R.L. (2007) 'Disruption of large-scale brain systems in advanced aging.' *Neuron*, **56**(5), 924–935.

Astle, D.E., Barnes, J.J., Baker, K., Colclough, G.L. & Woolrich, M.W. (2015) 'Cognitive training enhances intrinsic brain connectivity in childhood.' *Journal of Neuroscience*, **35**(16), 6277–6283.

Astolfi, L., Cincotti, F., Mattia, D., Marciani, M.G., Salinari, S., de Vico Fallani, F., Baccala, L.A., Ursino, M., Zavaglia, M., Ding, L., Edgar, J.C., Miller, G.A., He, B. & Babiloni, F. (2007) 'Comparison of different cortical connectivity estimators for high-resolution eeg recordings.' *Human Brain Mapping*, **28**(2), 143–157.

Atay-Kayis, A. & Massam, H. (2005) 'A Monte Carlo method for computing the marginal likelihood in nondecomposable Gaussian graphical models.' *Biometrika*, **92**(2), 317–335.

Aydore, S., Pantazis, D. & Leahy, R.M. (2013) 'A note on the phase locking value and its properties.' *NeuroImage*, **74**, 231–244.

Baccala, L.A. & Sameshima, K. (2001) 'Partial directed coherence: a new concept in neural structure determination.' *Biological Cybernetics*, **84**(6), 463–474.

Baillet, S., Mosher, J.C. & Leahy, R.M. (2001) 'Electromagnetic brain mapping.' *IEEE Signal Processing Magazine*, **18**(6), 14–30.

Baker, A.P., Brookes, M.J., Rezek, I.A., Smith, S.M., Behrens, T.E., Probert Smith, P.J. & Woolrich, M.W. (2014) 'Fast transient networks in spontaneous human brain activity.' *eLife*, **3**, e01 867.

Baldassano, C., Beck, D.M. & Fei-Fei, L. (2015) 'Parcellating connectivity in spatial maps.' *PeerJ*, **3**, e784.

Barachant, A., Bonnet, S., Congedo, M. & Jutten, C. (2013) 'Classification of covariance matrices using a Riemannian-based kernel for BCI applications.' *NeuroComputing*, **112**, 172–178.

Barbieri, M.M. & Berger, J.O. (2004) 'Optimal predictive model selection.' *Annals of Statistics*, **32**, 81–94.

Barkhof, F., Haller, S. & Rombouts, S.A.R.B. (2014) 'Resting-state functional MR imaging: a new window to the brain.' *Radiology*, **272**(1), 29–49.

Barnes, J.J., Nobre, A.C., Woolrich, M.W., Baker, K. & Astle, D.E. (2016*a*) 'Training working memory in childhood enhances coupling between frontoparietal control network and task-related regions.' *Journal of Neuroscience*, **36**(34), 9001–9011.

Barnes, J.J., Woolrich, M.W., Baker, K., Colclough, G.L. & Astle, D.E. (2016*b*) 'Electrophysiological measures of resting state functional connectivity and their relationship with working memory capacity in childhood.' *Developmental Science*, **19**(1), 19–31.

Bartels, M., Rietveld, M.J.H., van Baal, G.C.M. & Boomsma, D.I. (2002) 'Genetic and environmental influences on the development of intelligence.' *Behaviour Genetics*, **32**(4), 237–249.

Bastos, A., Litvak, V., Moran, R., Bosman, C., Fries, P. & Friston, K. (2015) 'A DCM study of spectral asymmetries in feedforward and feedback connections between visual areas V1 and V4 in the monkey.' *NeuroImage*, **108**, 460–475.

Beal, M.J. (2003) *Variational algorithms for approximate Bayesian inference*. Ph.D. thesis, University College London, London, UK.

Beckmann, C.F., DeLuca, M., Devlin, J.T. & Smith, S.M. (2005) 'Investigations into resting-state connectivity using independent component analysis.' *Philosphical Transactions of the Royal Society B*, **360**, 1001–1013.

Beckmann, C.F., Mackay, C.E., Filippini, N. & Smith, S.M. (2009) 'Group comparison of resting-state fMRI data using multi-subject ICA and dual regres-

sion.' In *The 15th annual meeting of the organization for human brain mapping*. OHBM, San Francisco, CA.

Beckmann, C.F. & Smith, S.M. (2004) 'Probabilistic independent component analysis for functional magnetic resonance imaging.' *IEEE Transactions on Medical Imaging*, **23**(2), 137–152.

Bell, A. & Sejnowski, T. (1995) 'An information-maximisation approach to blind separation and blind de-convolution.' *Neural Computation*, **7**, 1129–1159.

Benjamini, Y. & Hochberg, Y. (1995) 'Controlling the false discovery rate: a practical and powerful approach to multiple testing.' *Journal of the Royal Statistical Society B*, **57**(1), 289–300.

Bhadra, A., Datta, J., Polson, N.G. & Willard, B. (2016) 'The horseshoe+ estimator of ultra-sparse signals.' *Bayesian Analysis*. In Press.

de Bie, H.M., Boersma, M., Adriaanse, S., Veltman, D.J., Wink, A.M., Roosendaal, S.D., Barkhof, F., Stam, C.J., Oostrom, K.J. & Delemarre-van de Waal, H.A. (2012) 'Resting-state networks in awake five- to eight-year old children.' *Human Brain Mapping*, **33**(5), 1189–1201.

Bielczyk, N., Arenas, A., Buitelaar, J., Glennon, J. & Beckmann, C.F. (2016) 'Momentum: a new approach to causality in functional magnetic resonance imaging.' *ArXiv*:1606.08724v1.

Biswal, B., Yetkin, F.Z., Haughton, V.M. & Hyde, J.S. (1995) 'Functional connectivity in the motor cortex of resting human brain using echo-planar MRI.' *Magnetic Resonance in Medicine*, **34**(4), 537–541.

Blackwell, D. (1947) 'Conditional expectation and unbiased sequential estimation.' *The Annals of Mathematical Statistics*, **18**(1), 105–110.

Boomsma, D., Busjahn, A. & Peltonen, L. (2002) 'Classical twin studies and beyond.' *Nature Reviews Genetics*, **3**(11), 872–882.

Brodmann, K. (2006) *Brodmann's Localisation in the Cerebral Cortex*. Springer Science+Business Media, Inc., New York, NY, 3rd edition. Translated by Garey, L. J.

Brookes, M.J., Hale, J.R., Zumer, J.M., Stevenson, C.M., Francis, S.T., Barnes, G., Owen, J.P., Morris, P.G. & Nagarajan, S.S. (2011*a*) 'Measuring functional connectivity using MEG: methodology and comparison with fcMRI.' *NeuroImage*, **56**, 1082–1104.

Brookes, M.J., Liddle, E.B., Hale, J.R., Woolrich, M.W., Luckhoo, H., Liddle, P.F. & Morris, P.G. (2012*a*) 'Task induced modulation of neural oscillations in electrophysiological brain networks.' *NeuroImage*, **63**(4), 1918–1930.

Brookes, M.J., O'Neill, G.C., Hall, E.L., Woolrich, M.W., Baker, A., Palazzo-Corner, S., Robson, S.E. & Barnes, G.R. (2014) 'Measuring temporal, spectral and spatial changes in electrophysiological brain network connectivity.' *NeuroImage*, **91**, 282 – 299.

Brookes, M.J., Stevenson, C.M., Barnes, G.R., Hillebrand, A., Simpson, M.I.G., Francis, S.T. & Morris, P.G. (2007) 'Beamformer reconstruction of correlated sources using a modified source model.' *NeuroImage*, **34**(4), 1454–1465.

Brookes, M.J., Woolrich, M.W. & Barnes, G.R. (2012*b*) 'Measuring functional connectivity in MEG: A multivariate approach insensitive to linear source leakage.' *NeuroImage*, **63**(2), 910 – 920.

Brookes, M.J., Woolrich, M.W., Luckhoo, H., Price, D., Hale, J.R., Stephenson, M.C., Barnes, G.R., Smith, S.M. & Morris, P.G. (2011*b*) 'Investigating the electrophysiological basis of resting state networks using magnetoencephalography.' *Proceeding of the National Academy of Sciences* USA, **108**, 16 783–16 788.

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

Buckner, R.L., Snyder, A.Z., Shannon, B.J., LaRossa, G., Sachs, R., Fotenos, A.F., Sheline, Y.I., Klunk, W.E., Mathis, C.A., Morris, J.C. & Mintun, M.A. (2005) 'Molecular, structural, and functional characterization of Alzheimer's disease: evidence for a relationship between default activity, amyloid, and memory.' *Journal of Neuroscience*, **25**(34), 7709–7717.

Bullmore, E. (2012) 'Functional network endophenotypes of psychotic disorders.' *Biological Psychiatry*, **71**(10), 844–845.

Bullmore, E. & Sporns, O. (2009) 'Complex brain networks: graph theoretical analysis of structural and functional systems.' *Nature Reviews Neuroscience*, **10**(3), 186–198.

——— (2012) 'The economy of brain network organization.' *Nature Reviews Neuroscience*, **13**, 336–349.

Burnham, K.P. & Anderson, D.R. (2002) *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*. Springer-Verlag, New York, ny, 2nd edition.

Cabral, J., Luckhoo, H., Woolrich, M., Joensson, M., Mohseni, H., Baker, A., Kringelbach, M.L. & Deco, G. (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**(0), 423 – 435.

Calhoun, V.D., Adali, T., Pearlson, G.D. & Pekar, J.J. (2001) 'A method for making group inferences from functional mri data using independent component analysis.' *Human Brain Mapping*, **14**, 140–151.

Callard, F. & Margulies, D.S. (2011) 'The subject at rest: novel conceptualizations of self and brain from cognitive neuroscience's study of "resting state".' *Subjectivity*, **4**(3), 227–257.

Canolty, R.T., Edwards, E., Dalal, S.S., Soltani, M., Nagarajan, S.S., Kirsch, H.E., Berger, M.S., Barbaro, N.M. & Knight, R.T. (2006) 'High gamma power is phase-locked to theta oscillations in human neocortex.' *Science*, **313**(5793), 1626–1628.

Carvalho, C.M., Polson, N.G. & Scott, J.G. (2010) 'The horseshoe estimator for sparse signals.' *Biometrika*, **97**(2), 465–480.

Casella, G. & Moreno, E. (2006) 'Objective bayesian variable selection.' *Journal of the American Statistical Association*, **101**(473), 157–167.

Chappell, M.A., Groves, A.R., Whitcher, B. & Woolrich, M.W. (2009) 'Variational bayesian inference for a nonlinear forward model.' *IEEE Transactions on Signal Processing*, **57**, 223–236.

Chen, X. (2014) *Accelerated Estimation and Inference for Heritability of f*mri *Data*. Ph.D. thesis, University of Warwick, Warwick, uk.

Chen, X., Smith, S.M., Viding, E. & Nichols, T.E. (2014) 'apace: Accelerated permutation inference for the ace model.' In *The 20th annual meeting of the organization for human brain mapping*. ohbm, Hamburg, Germany.

Chetty, R., Hendren, N. & Katz, L.F. (2015) 'The effects of exposure to better neighborhoods on children: new evidence from the moving to opportunity experiment.' Working Paper 21156, National Bureau of Economic Research.

Clarke, M.R.B. (1981) 'A Givens algorithm for moving from one linear model to another without going back to the data.' *Journal of the Royal Statistical Society C*, **30**(2), 198–203.

Clyde, M.A., Ghosh, J. & Littman, M.L. (2011) 'Bayesian adaptive sampling for variable selection and model averaging.' *Journal of Computational and Graphical Statistics*, **20**(1), 80–101.

Colclough, G.L., Brookes, M.J., Smith, S.M. & Woolrich, M.W. (2015) 'A symmetric multivariate leakage correction for meg connectomes.' *NeuroImage*, **117**, 439–448.

Colclough, G.L., Smith, S.M., Harrison, S.J., Rojas López, P.A., Valdes-Sosa, P. & Woolrich, M.W. (2016*a*) 'A Bayesian approach for hierarchical modelling of sparse functional networks.' In *The 22nd annual meeting of the organization for human brain mapping*. ohbm, Geneva, Switzerland.

Colclough, G.L., Smith, S.M., Nichols, T.E., Glasser, M.F., Van Essen, D.C. & Woolrich, M.W. (2016*b*) 'Genetic factors influence resting-state connectivity more than common environmental factors.' In *The 22nd annual meeting of the organization for human brain mapping*. ohbm, Geneva, Switzerland.

Colclough, G.L., Smith, S.M., Nichols, T.E., Tewarie, P.K., Brookes, M.J., Winkler, A.M., Glasser, M.F., Van Essen, D.C. & Woolrich, M.W. (2016*c*) 'The heritability of multi-modal functional connectivity.' In *The 29th International Conference on Biomagnetism*. BIOMAG, Seoul, South Korea.

Colclough, G.L., Smith, S.M., Nichols, T.E., Winkler, A.M., Sotiropoulos, S.N., Glasser, M.F., Van Essen, D.C. & Woolrich, M.W. (2016*d*) 'The heritability of multi-modal connectivity in human brain activity.' In revision, *e*life.

Colclough, G.L., Woolrich, M.W., Tewarie, P.K., Brookes, M.J., Quinn, A.J. & Smith, S.M. (2016*e*) 'How reliable are meg resting-state connectivity metrics?' *NeuroImage*, **138**, 284–293.

Cole, S.R., Peterson, E.J., van der Meij, R., de Hemptinne, C., Starr, P.A. & Voytek, B. (2016) 'Nonsinusoidal oscillations underlie pathological phase-amplitude coupling in the motor cortex in parkinson's disease.' BioRxiv:10.1101/049304.

Conley, D., Rauscher, E., Dawes, C., Magnusson, P.K.E. & Siegal, M.L. (2013) 'Heritability and the equal environments assumption: evidence from multiple samples of misclassified twins.' *Behavior Genetics*, **43**(5), 415–426.

Craddock, R.C., Holtzheimer, P.E., Hu, X.P. & Mayberg, H.S. (2009) 'Disease state prediction from resting state functional connectivity.' *Magnetic Resonance in Medicine*, **62**(6), 1619–1628.

Craddock, R.C., James, G., Holtzheimer, P.E., Hu, X.P. & Mayberg, H.S. (2012) 'A whole brain fmri atlas generated via spatially constrained spectral clustering.' *Human Brain Mapping*, **33**(8), 1914–1928.

Craddock, R.C., Jbabdi, S., Yan, C.G., Vogelstein, J.T., Castellanos, F.X., Martino, A.D., Kelly, C., Heberlein, K., Colcombe, S. & Milham, M.P. (2013) 'Imaging human connectomes at the macroscale.' *Nature Methods*, **10**(6), 524–539.

Craiu, R.V., Rosenthal, J. & Yang, C. (2009) 'Learn from thy neighbor: Parallel-chain and regional adaptive mcmc.' *Journal of the American Statistical Association*, **104**, 1454–1466.

Dadi, K., Abraham, A., Rahim, M., Thirion, B. & Varoquaux, G. (2016) 'Comparing functional connectivity based predictive models across datasets.' In *6th International Workshop on Pattern Recognition in Neuroimaging*. PRNI, Trento, Italy.

Dalal, S.S., Sekihara, K. & Nagarajan, S.S. (2006) 'Modified beamformers for coherent source region suppresion.' *IEEE Transactions on Biomedical Engineering*, **53**(7), 1357–1363.

Damoiseaux, J.S., Rombouts, S.A.R.B., Barkhof, F., Scheltens, P., Stam, C.J., Smith, S.M. & Beckmann, C.F. (2006) 'Consistent resting-state networks across healthy subjects.' *Proceedings of the National Academy of Sciences USA*, **103**, 13 848–13 853.

Danaher, P., Wang, P. & Witten, D.M. (2015) 'The joint graphical lasso for inverse covariance estimation across multiple classes.' *Journal of the Royal Statistical Society B*, **76**(2), 373–397.

David, O., Cosmelli, D. & Friston, K.J. (2004) 'Evaluation of different measures of functional connectivity using a neural mass model.' *NeuroImage*, **21**(2), 659–73.

David, O., Kilner, J.M. & Friston, K.J. (2006) 'Mechanisms of evoked and induced responses in MEG/EEG.' *NeuroImage*, **31**, 1580–1591.

De Luca, M., Beckmann, C.F., De Stefano, N., Matthews, P.M. & Smith, S.M. (2005) 'fMRI resting state networks define distinct modes of long-distance interactions in the human brain.' *NeuroImage*, **29**(4), 1359–1367.

Deco, G., Jirsa, V., Robinson, P., Breakspear, M. & Friston, K. (2008) 'The dynamic brain: from spiking neurons to neural masses and cortical fields.' *PLoS Computational Biology*, **4**, e1000 092.

Deco, G. & Kringelbach, M. (2016) 'Metastability and coherence: extending the communication through coherence hypothesis using a whole-brain computational perspective.' *Trends in Neurosciences*, **39**(3), 125–135.

Dellaportas, P. & Forster, J.J. (1999) 'Markov chain Monte Carlo model determination for hierarchical and graphical log-linear models.' *Biometrika*, **86**(3), 615–633.

Denison, D.G.T., Mallick, B.K. & Smith, A.F.M. (1998) 'Automatic Bayesian curve fitting.' *Journal of the Royal Statistical Society B*, **16**, 395–407.

Dennis, E.L. & Thompson, P.M. (2014) 'Functional brain connectivity using fMRI in aging and Alzheimer's disease.' *Neuropsychology Review*, **24**(1), 49–62.

Deuker, L., Bullmore, E.T., Smith, M., Christensen, S., Nathan, P.J., Rockstroh, B. & Bassett, D.S. (2009) 'Reproducability of graph metrics of human brain functional networks.' *NeuroImage*, **47**(4), 1460–1468.

Dhillon, P.S., Wolk, D.A., Das, S.R., Ungar, L.H., Gee, J.C. & Avants, B.B. (2014) 'Subject-specific functional parcellation via prior based eigenanatomy.' *NeuroImage*, **99**, 14–27.

Dosenbach, N.U.F., Nardos, B., Cohen, A.L., Fair, D.A., Power, J.D., Church, J.A., Nelson, S.M., Wig, G.S., Vogel, A.C., Lessov-Schlaggar, C.N., Barnes, K.A., Dubis, J.W., Feczko, E., Coalson, R.S., Pruett, J.R., Barch, D.M., Petersen, S.E. & Schlaggar, B.L. (2010) 'Prediction of individual brain maturity using fMRI.' *Science*, **329**(5997), 1358–1361.

van Drongelen, W., Yuchtman, M., Van Veen, B.D. & Van Huffelen, A.C. (1996) 'A spatial filtering technique to detect and localize multiple sources in the brain.' *Brain Topography*, **9**, 39–49.

Duff, E., Makin, T., Madugula, S., Smith, S.M. & Woolrich, M.W. (2013) 'Utility of partial correlation for characterising brain dynamics: MVPA-based assessment of regularisation and network selection.' In *3rd International Workshop on Pattern Recognition in Neuroimaging*, pp. 58–61. PRNI, Philadelphia, PA.

Duncan, D., Duckrow, R.B., Pincus, S.M., Goncharova, I., Hirsch, L.J., Spencer, D.D., Coifman, R.R. & Zaveri, H.P. (2013) 'Intracranial EEG evaluation of relationship within a resting state network.' *Clinical Neurophysiology*, **124**(10), 1943–1951.

van Dyk, D.A. & Park, T. (2008) 'Partially collapsed Gibbs samplers: theory and methods.' *Journal of the American Statistical Association*, **103**(482), 790–796.

Edelman, S. (2008) *Computing the mind.* Oxford University Press usa, New York, ny.

Edwards, R.G. & Sokal, A.D. (1988) 'Generalisation of the Fortuin-Kasteleyn-Swedensen-Wang representation and Monte Carlo algorithm.' *Physical Review D*, **38**(6), 2009–2012.

Eklund, A., Nichols, T.E. & Knutsson, H. (2016) 'Cluster failure: Why fmri inferences for spatial extent have inflated false-positive rates.' *Proceedings of the National Academy of Sciences USA*, **113**(28), 7900–7905.

Engel, A.K. & Fries, P. (2010) 'Beta-band oscillations—signalling the status quo?' *Current Opinion in Neurobiology*, **20**, 156–165.

Engel, A.K., Gerloff, C., Hilgetag, C.C. & Nolte, G. (2013) 'Intrinsic coupling modes: multiscale interactions in ongoing brain activity.' *Neuron*, **80**(4), 867–886.

van 't Ent, D., van Soelen, I.L.C., Stam, C.J., de Geus, E.J.C. & Boomsma, D.I. (2009) 'Strong resemblance in the amplitude of oscillatory brain activity in monozygotic twins is not caused by "trivial" similarities in the composition of the skull.' *Human Brain Mapping*, **30**, 2142–2145.

Everson, R. (1998) 'Orthogonal, but not orthonormal, procrustes problems.' Technical Report TR99-6, Imperial College, London, uk. Available from empslocal.ex.ac.uk/people/staff/reverson/uploads/Site/procrustes.pdf (accessed 11 August 2014).

Ewald, A., Marzetti, L., Zappasodi, F., Meinecke, F.C. & Nolte, G. (2011) 'Estimating true brain connectivity from eeg/meg data invariant to linear and static transformations in sensor space.' *NeuroImage*, **60**, 476–488.

Fair, D.A., Dosenbach, N.U.F., Church, J.A., Cohen, A.L., Brahmbhatt, S., Miezin, F.M., Barch, D.M., Raichle, M.E., Petersen, S.E. & Schlaggar, B.L. (2007) 'Development of distinct control networks through

segregation and integration.' *Proceedings of the National Academy of Sciences USA*, **104**(33), 13 507–13 512.

Felleman, D.J. & Van Essen, D.C. (1991) 'Distributed hierarchical processing in the primate cerebral cortex.' *Cerebral Cortex*, **1**(1), 1–47.

Felson, J. (2014) 'What can we learn from twin sudies? A comprehensive evaluation of the equal environments assumption.' *Social Science Research*, **43**.

Ferreira, L.K. & Busatto, G.F. (2013) 'Resting-state functional connectivity in normal brain aging.' *Neuroscience and Biobehavioural Reviews*, **37**(3), 384–400.

Ferreira, L.K., Regina, A.C.B., Kovacevic, N., da Graça Morais Martin, M., Santos, P.P., de Godoi Carneiro, C., Kerr, D.S., Amaro, E., McIntosh, A.R. & Busatto, G.F. (2016) 'Aging effects on whole-brain functional connectivity in adults free of cognitive and psychiatric disorders.' *Cerebral Cortex*, **26**(9), 3851–3865.

Filippini, N., J., M.B., Hough, M.G., Goodwin, G.M., Frisoni, G.B., Smith, S.M., Matthews, P.M., Beckmann, C.F. & Mackay, C.E. (2009) 'Distinct patterns of brain activity in young carriers of the APOE-ϵ4 allele.' *Proceedings of the National Academy of Sciences USA*, **106**, 7209–7214.

Finn, E.S., Shen, X., Scheinost, D., Rosenberg, M.D., Huang, J., Chun, M.M., Papademetris, X. & Constable, R.T. (2015) 'Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity.' *Nature Neuroscience*, **18**, 1664–1671.

Fisher, R.A. (1924) 'The distribution of the partial correlation coefficient.' *Metron*, **3**, 329–332.

Fornito, A. & Bullmore, E.T. (2015*a*) 'Connectomics: a new paradigm for understanding brain disease.' *European Neuropsychopharmacology*, **25**(5), 733–748.

——— (2015*b*) 'Reconciling abnormalities of brain network structure and function in schizophrenia.' *Current Opinion in Neurobiology*, **30**, 44–50.

Fornito, A., Zalesky, A., Bassett, D.S., Meunier, D., Ellison-Wright, I., Yücel, M., Wood, S.J., Shaw, K., O'Connor, J., Nertney, D., Mowry, B.J., Pantelis, C. & Bullmore, E.T. (2011) 'Genetic influences on cost-efficient organization of human cortical functional networks.' *Journal of Neuroscience*, **31**(9), 3261–3270.

Fox, M.D., Corbetta, M., Snyder, A.Z., Vincent, J.L. & Raichle, M.E. (2006) 'Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems.' *Proceedings of the National Academy of Sciences USA*, **103**, 10 046–10 051.

Friedman, J., Hastie, T. & Tibshirani, R. (2008) 'Sparse inverse covariance estimation with the graphical lasso.' *Biostatistics*, **9**(3), 432–441.

Friston, K. & Kiebel, S. (2009) 'Predictive coding under the free-energy principle.' *Philosophical Transactions of the Royal Society B: Biological Sciences*, **364**(1521), 1211–1221.

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

Friston, K.J., Harrison, L. & Penny, W. (2003) 'Dynamic causal modelling.' *Neuroimage*, **19**, 1273–302.

Friston, K.J., Kahan, J., Biswal, B.B. & Razi, A. (2014) 'A DCM for resting state fMRI.' *NeuroImage*, **94**, 396–407.

Friston, K.J., Li, B., Daunizeau, J. & Stephan, K.E. (2011) 'Network discovery with DCM.' *Neuroimage*, **56**, 1202–1221.

Friston, K.J., Stephan, K.E., Lund, T.E., Morcom, A. & Kiebel, S. (2005) 'Mixed-effects and fMRI studies.' *NeuroImage*, **24**, 244–252.

Gal, Y. (2015) 'Rapid prototyping of probabilistic models: emerging challenges in variational inference.' Available from `www.approximateinference.org/accepted/Gal2015.pdf` (accessed 21st September 2016).

Gal, Y. & Ghahramani, Z. (2015*a*) 'Dropout as a Bayesian approximation: Representing model uncertainty in deep learning.' *ArXiv*:1506.02142.

——— (2015*b*) 'On modern deep learning and variational inference.' In *Advances in Approximate Bayesian Inference workshop, 29th Annual Conference on Neural Information and Processing Systems*. NIPS, Montreal, Canada.

Garcés, P., Martín-Buro, M.C. & Maestú, F. (2016) 'Quantifying the test-retest reliability of magnetoencephalography resting-state functional connectivity.' *Brain Connectivity*, **6**(6), 448–460.

Gelman, A. (2006) 'Prior distributions for variance parameters in hierarchical models.' *Bayesian Analysis*, **1**(3), 515–533.

Gelman, A., Carlin, J.B., Stern, H.S., Dunson, D.B., Vehtari, A. & Rubin, D.B. (2014) *Bayesian Data Analysis*. CRC Press, Taylor & Francis Group, Boca Raton, FL, 3rd edition.

Gelman, A., van Dyk, D.A., Huang, Z. & Boscardin, W.J. (2008) 'Using redundant parameterizations to fit hierarchical models.' *Journal of Computational and Graphical Statistics*, **17**(1), 95–122.

Gelman, A., Roberts, G.O. & Gilks, W.R. (1996) 'Efficient Metropolis jumping rules.' *Bayesian Statistics*, **5**, 599–607.

Gelman, A. & Rubin, D.B. (1992) 'Inference from iterative simulation using multiple sequences.' *Statistical Science*, **7**(4), 457–511.

George, E.I. & McCulloch, R.E. (1993) 'Variable selection via Gibbs sampling.' *Journal of the American Statistical Association*, **88**(423), 881–889.

——— (1997) 'Approaches for Bayesian variable selection.' *Statistica Sinica*, **7**, 339–373.

Geweke, J. (1996) 'Variable selection and model comparison in regression.' *Bayesian Statistics*, **5**, 609–620.

Ghosh, J. & Clyde, M.A. (2011) 'Rao-Blackwellization for Bayesian variable selection and model averaging in linear and binary regression: a novel data augmentation approach.' *Journal of the American Statistical Association*, **106**(495), 1041–1052.

Glahn, D.C., Winkler, A.M., Kochunov, P., Almasy, L., Duggirala, R., Carless, M.A., Curran, J.C., Olvera, R.L., Laird, A.R., Smith, S.M., Beckmann, C.F., Fox, P.T. & Blangero, J. (2010) 'Genetic control over the resting brain.' *Proceedings of the National Academy of Sciences USA*, **107**(3), 1223–1228.

Glasser, M.F., Coalson, T.S., Robinson, E.C., Hacker, C.D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C.F., Jenkinson, M., Smith, S.M. & Van Essen, D.C. (2016) 'A multi-modal parcellation of human cerebral cortex.' *Nature*, **536**, 171–178.

Glasser, M.F., Sotiropoulos, S., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., Van Essen, D.C., Jenkinson, M. & Wu-Minn hcp consortium (2013) 'The minimal preprocesing pipelines for the Human Connectome Project.' *NeuroImage*, **80**, 105–124.

Gollo, L.L., Mirasso, C., Sporns, O. & Breakspear, M. (2014) 'Mechanisms of zero-lag synchronization in cortical motifs.' *PLoS Computational Biology*, **10**(4), e1003 548.

Gómez-Robles, A., Hopkins, W.D., Schapiro, S.J. & Sherwood, C.C. (2015) 'Relaxed genetic control of cortical organization in human brains compared with chimpanzees.' *Proceedings of the National Academy of Sciences USA*, **112**(48), 14 799–14 804.

Gordon, E.M., Laumann, T.O., Adeyemo, B., Huckins, J.F., Kelley, W.M. & Petersen, S.E. (2016) 'Generation and evaluation of a cortical area parcellation from resting-state correlations.' *Cerebral Cortex*, **26**(1), 288–303.

Granger, C.W.J. (1969) 'Investigating causal relations by econometric models and cross-spectral methods.' *Econometrica*, **37**, 424–438.

Green, P.J. (1995) 'Reversible jump Markov chain Monte Carlo computation and Bayesian model determination.' *Biometrika*, **82**(4), 711–732.

——— (2003) 'Trans-dimensional Markov chain Monte Carlo.' In P.J. Green, N.L. Hjort & S. Richardson, eds., *Highly Structured Stochastic Systems*, Oxford statistical science series, pp. 179–198. Oxford University Press, Oxford, uk.

Green, P.J. & Mira, A. (2001) 'Delayed rejection in reversible jump Metropolis-Hastings.' *Biometrika*, **88**, 1035–1053.

Greicius, M.D. (2008) 'Resting-state functional connectivity in neuropsychiatric disorders.' *Current Opinion in Neurology*, **21**, 424–430.

Greicius, M.D., Krasnow, B., Reiss, A.L. & Menon, V. (2003) 'Functional connectivity in the resting brain: a network analysis of the default mode hypothesis.' *Proceedings of the National Academy of Sciences USA*, **100**(1), 253–258.

Gross, J., Kujala, J., Hamalainen, M., Timmermann, L., Schnitzler, A. & Salmelin, R. (2001) 'Dynamic imaging of coherent sources: studying neural interactions in the human brain.' *Proceedings of the National Academy of Sciences USA*, **98**, 694–699.

Groves, A.R., Beckmann, C.F., Smith, S.M. & Woolrich, M.W. (2011) 'Linked independent component analysis for multimodal data fusion.' *NeuroImage*, **54**(3), 2198–2217.

Guan, Y. & Stephens, M. (2011) 'Bayesian variable selection regression for genome-wide association studies and other large-scale probelms.' *Annals of Applied Statistics*, **5**(3), 1780–1815.

Guo, J., Levina, E., Michailidis, G. & Zhu, J. (2011) 'Joint estimation of multiple graphical models.' *Biometrika*, **98**(1), 1–15.

Gusnard, D.A., Akbudak, E., Shulman, G.L. & Raichle, M.E. (2001) 'Medial prefronal cortex and self-referential mental activity: relation to a default mode of brain function.' *Proceedings of the National Academy of Sciences USA*, **98**(7), 4259–4264.

Gusnard, D.A. & Raichle, M.E. (2001) 'Searching for a baseline: functional imaging and the resting brain.' *Nature Reviews Neuroscience*, **2**(10), 685–694.

Gutchess, A. (2014) 'Plasticity of the aging brain: new directions in cognitive neuroscience.' *Science*, **346**(6209), 579–582.

Hafkemeijer, A., van der Grond, J. & Rombouts, S.A. (2012) 'Imaging the default mode network in aging and dementia.' *Biochimica et Biophysica Acta (bba) - Molecular Basis of Disease*, **1822**(3), 431–441.

Hahn, P.R. & Carvalho, C.M. (2015) 'Decoupling shrinkage and selection in Bayesian linear models: a posterior summary perspective.' *Journal of the American Statistical Association*, **110**(509), 435–448.

Hans, C. (2009) 'Bayesian lasso regression.' *Biometrika*, **96**(4), 835–845.

Hardmeier, M., Hatz, F., Bousleiman, H., Schindler, C., Stam, C.J. & Fuhr, P. (2014) 'Reproducability of functional connectivity and graph measures based on the phase lag index (pli) and weighted phase lag index (wpli) derived from high resolution eeg.' *PLoS ONE*, **9**(10), e108 648.

Harrington, D.L., Rubinov, M., Durgerian, S., Mourany, L., Reece, C., Koenig, K., Bullmore, E., Long, J.D., Paulsen, J.S. & Rao, S.M. (2015) 'Network topology and functional connectivity disturbances precede the onset of huntington's disease.' *Brain*, **138**(8), 2332–2346.

Harrison, S.J., Woolrich, M.W., Robinson, E.C., Glasser, M.F., Beckmann, C.F., Jenkinson, M. & Smith, S.M. (2015) 'Large-scale probabilistic functional modes from resting-state fmri.' *NeuroImage*, **109**, 217–231.

Hastie, T., Tibshirani, R. & Friedman, J. (2009) *The Elements of Statistical Learning: Data Mining, Inference and Prediction*. Springer-Verlag, New York, ny, 2nd edition.

Hauk, O. & Stenroos, M. (2014) 'A framework for the design of flexible cross-talk functions for spatial filtering of eeg/meg data: deflect.' *Human Brain Mapping*, **35**(4), 1642–1653.

Hawrylycz, M., Miller, J.A., Menon, V., Feng, D., Dolbeare, T., Guillozet-Bongaars, A.L., Jegga, A.G., Aronow, B.J., Lee, C.K., Bernard, A., Glasser, M.F., Dierker, D.L., Menche, J., Szafer, A., Collman, F., Grange, P., Berman, K.A., Mihalas, S., Yao, Z., Stewart, L., Barabási, A.L., Schulkin, J., Phillips, J., Ng, L., Dang, C., Haynor, D.R.,

Jones, A., Van Essen, D.C., Koch, C. & Lein, E. (2015) 'Canonical genetic signatures of the adult human brain.' *Nature Neuroscience*, **18**(12), 1832–1844.

de Hemptinne, C., Swann, N.C., Ostrem, J.L., Ryapolova-Webb, E.S., Luciano, M.S., Galifianakis, N.B. & Starr, P.A. (2015) 'Therapeutic deep brain stimulation reduces cortical phase-amplitude coupling in Parkinson's disease.' *Nature Neuroscience*, **18**(5), 779–786.

van den Heuvel M. P., van Soelen, I.L.C., Stam, C.J., Kahn, R.S., Boomsma, D.I. & Pol, H.E.H. (2013) 'Genetic control of functional brain network efficiency in children.' *European Neuropsychopharmacology*, **23**(1), 19–23.

Higham, N.J. (1988) 'Computing a nearest symmtetric positive semidefinite matrix.' *Linear Algebra and its Applications*, **103**(103–118).

Hillebrand, A., Barnes, G.R., Bosboom, J.L., Berendse, H.W. & Stam, C.J. (2012) 'Frequency-dependent functional connectivity within resting-state networks: an atlas-based MEG beamformer solution.' *NeuroImage*, **59**, 3909–3921.

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

Hinne, M., Janssen, R.J., Heskes, T. & van Gerven, M. (2015) 'Bayesian estimation of conditional independence graphs improves functional connectivity estimates.' *PLoS Computational Biology*, p. 10.1371/journal.pcbi.1004534.

Hinne, M., Lenkoski, A., Heskes, T. & van Gerven, M.A.J. (2014) 'Efficient sampling of Gaussian graphical models using conditional Bayes factors.' *Stat*, **3**, 326–336.

Hipp, J.F., Hawellek, D.J., Corbetta, M., Siegel, M. & Engel, A.K. (2012) 'Large-scale cortical correlation structure of spontaneous oscillatory activity.' *Nature Neuroscience*, **15**(6), 884–890.

Hodge, M.R., Horton, W., Brown, T., Herrick, R., Olsen, T., Hileman, M.E., McKay, M., Archie, K.A., Cler, E., Harms, M.P., Burgess, G.C.,

Glasser, M.F., Elam, J.S., Curtiss, S.W., Barch, D.M., Oostenveld, R., Larson-Prior, L.J., Ugurbil, K., Essen, D.C.V. & Marcus, D.S. (2016) 'ConnectomeDB—sharing human brain connectivity data.' *NeuroImage*, **124**, 1102–1107.

Hoerl, A.E. & Kennard, R.W. (1970*a*) 'Ridge regression: applications to nonorthogonal problems.' *Technometrics*, **12**(1), 69–82.

——— (1970*b*) 'Ridge regression: biased estimation for nonorthogonal problems.' *Technometrics*, **12**(1), 55–67.

Huang, M.X., Mosher, J.C. & Leahy, R.M. (1999) 'A sensor-weighted overlapping-sphere head model and exhaustive head model comparison for MEG.' *Physics in Medicine and Biology*, **44**(2), 423–440.

Huang, N.E., Hu, K., Yang, A.C.C., Chang, H.C., Jia, D., Liang, W.K., Yeh, J.R., Kao, C.L., Juan, C.H., Peng, C.K., Meijer, J.H., Wang, Y.H., Long, S.R. & Wu, Z. (2016) 'On holo-Hilbert spectral analysis: a full informational spectral representation for nonlinear and non-stationary data.' *Philosophical Transactions of the Royal Society A*, **374**(2065), 20150 206.

Hui, H, B., Pantazis, D., Bressler, S.L. & Leahy, R.M. (2010) 'Identifying true cortical interactions in MEG using the nulling beamformer.' *NeuroImage*, **49**(4), 3161–3174.

Hyvärinen, A. (1999) 'Fast and robust fixed-point algorithms for independent component analysis.' *IEEE Trans. Neural Networks*, **10**(3), 626–634.

Hyvärinen, A. & Smith, S.M. (2013) 'Pairwise likelihood ratios for estimation of non-Gaussian structural equation models.' *Journal of Machine Learning Research*, **14**, 111–152.

Ishwaran, H. & Rau, J.S. (2005) 'Spike and slab variable selction: frequentist and Bayesian strategies.' *Annals of Statistics*, **33**(2), 730–773.

Jahanshad, N., Hibar, D.P., Ryles, A., Toga, A.W., McMahon, K.L., de Zubicaray, G.I., Hansell, N.K., Montgomery, G.W., Martin, N.G., Wright, M.J. & Thompson, P.M. (2012) 'Discovery of genes that affect

human brain connectivty: a genome-wide analysis of the connectome.' In *Proceedings of the IEEE International Symposium on Biomedical Imaging*, pp. 542–545.

Jamashita, M., Kawato, M. & Imamizu, H. (2015) 'Predicting learning plateau of working memory from whole-brain intrinsic network connectivity patterns.' *Scientific Reports*, **5**, 7622.

Jansen, A.G., Mous, S.E., White, T., Posthuma, D. & Polderman, T.J.C. (2015) 'What twin studies tell us about the heritability of brain development, morphology, and function: a review.' *Neuropsychology Review*, **25**(1), 27–46.

Jenkinson, M., Beckmann, C.F., Behrens, T.E., Woolrich, M.W. & Smith, S.M. (2012) 'FSL.' *Neuroimage*, **62**, 782–90.

Jin, S.H., Seol, J., Kim, J.S. & Chung, C.K. (2011) 'How reliable are the functional connectivity networks of MEG in resting states?' *J. Neurophys.*, **106**(6), 2888–2895.

Jones, S.R. (2016) 'When brain rhythms aren't "rhythmic:" implication for their mechanisms and meaning.' *Current Opinion in Neurobiology*, **40**, 72–80.

Joseph, J. (1998) 'The equal environment assumption of the classical twin method: a critical analysis.' *Journal of Mind and Behaviour*, **19**, 325–258.

Kaminski, M. & Blinowska, K.J. (2014) 'Directed transfer function is not influenced by volume conduction—inexpedient pre-processing should be avoided.' *Frontiers in Computational Neuroscience*, **8**, 61.

Kerman, J. (2011) 'Neutral noninformative and informative conjugate beta and gamma prior distributions.' *Electronic Journal of Statistics*, **5**, 1450–1470.

Kessy, A., Lewin, A. & Strimmer, K. (2015) 'Optimal whitening and decorrelation.' *ArXiv*:1512.00809v1.

Khondker, Z.S., Zhu, H., Chu, H., Lin, W. & Ibrahim, J.G. (2013) 'The Bayesian covariance lasso.' *Statistics and its Interface*, **6**(2), 243–259.

Kiebel, S.J., Garrido, M.I., Moran, R.J. & Friston, K.J. (2008) 'Dynamic causal modelling for EEG and MEG.' *Cognitive Neurodynamics*, **2**, 121–136.

Kingma, D.P. & Welling, M. (2014) 'Auto-encoding variational Bayes.' In *Proceedings of the International Conference on Learning Representations*. iclr, Banff, Canada.

Kochunov, P., Jahanshad, N., Marcus, D., Winkler, A., Sprooten, E., Nichols, T.E., Wright, S.N., Hong, L.E., Patel, B., Behrens, T., Jbabdi, S., Andersson, J., Lenglet, C., Yacoub, E., Moeller, S., Auerbach, E., Ugurbil, K., Sotiropoulos, S.N., Brouwer, R.M., Landman, B., Lemaitre, H., den Braber, A., Zwiers, M.P., Ritchie, S., van Hulzen, K., Almasy, L., Curran, J., deZubicaray, G.I., Duggirala, R., Fox, P., Martin, N.G., McMahon, K.L., Mitchell, B., Olvera, R.L., Peterson, C., Starr, J., Sussmann, J., Wardlaw, J., Wright, M., Boomsma, D.I., Kahn, R., de Geus, E.J., Williamson, D.E., Hariri, A., van 't Ent, D., Bastin, M.E., McIntosh, A., Deary, I.J., pol, H.E.H., Blangero, J., Thompson, P.M., Glahn, D.C. & Essen, D.C.V. (2015) 'Heritability of fractional anisotropy in human white matter: a comparison of Human Connectome Project and enigma-dti data.' *NeuroImage*, **111**, 300–311.

Kohn, R., Smith, M. & Chan, D. (2001) 'Nonparametric regression using linear combinations of basis functions.' *Statistics and Computing*, **11**, 313–322.

Korhonen, O., Palva, S. & Palva, J.M. (2014) 'Sparse weightings for collapsing inverse solutions to cortical parcellations optimize m/eeg source reconstruction accuracy.' *Journal of Neuroscience Methods*, **226**, 147 – 160.

Kucukelbir, A., Ranganath, R., Gelman, A. & Blei, D.M. (2015) 'Automatic variational inference in Stan.' In *The 29th Annual Conference on Neural Information Processing Systems*.

Kunihiro, B., Shibata, R. & Sibuya, M. (2004) 'Partial correlation and conditional correlation as measures of conditional independence.' *Australia and New Zealand Journal of Statistics*, **46**, 657–664.

Kuo, L. & Mallick, B. (1998) 'Variable selection for regression models.' *Indian Journal of Statistics*, **60**, 65–81.

Lachaux, J.P., Rodriguez, E., Martinerie, J. & Varela, F.J. (1999) 'Measuring phase synchrony in brain signals.' *Human Brain Mapping*, **8**, 194–208.

Lachaux, J.P., Rodriguez, E., le van Quyen, M., Lutz, A., Martinerie, J. & Varela, F.J. (2000) 'Studying single-trials of phase synchronous activity in the brain.' *International Journal of Bifurcation and Chaos*, **10**(10), 2429–2439.

Lakatos, P., Barczak, A., Neymotin, S.A., McGinnis, T., Ross, D., Javitt, D.C. & O'Connell, M.N. (2016) 'Global dynamics of selective attention and its lapses in primary auditory cortex.' *Nature Neuroscience*. In Press.

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

Laskey, K.B. & Myers, J.W. (2003) 'Population Markov chain Monte Carlo.' *Machine Learning*, **50**(1), 175–196.

Lee, W. & Liu, Y. (2015) 'Joint estimation of multiple precision matrices with common structures.' *Journal of Machine Learning Research*, **16**, 1035–1062.

Leistritz, L., Pester, B., Doering, A., Schiecke, K., Babiloni, F., Astolfi, L. & Witte, H. (2013) 'Time-variant partial directed coherence for analysing connectivity: a methodological study.' *Philosophical Transactions of the Royal Society A*, **371**, 20110 616.

Lenkoski, A. (2013) 'A direct sampler for G-Wishart variates.' *Stat*, **2**, 119–128.

Lenkoski, A. & Dobra, A. (2011) 'Computational aspects related to inference in Gaussian graphical models with the G-Wishart prior.' *Journal of the American Statistical Association*, **20**(1), 140–157.

Letac, G. & Massam, H. (2007) 'Wishart distributions for decomposable graphs.' *Annals of Statistics*, **35**(3), 1278–1323.

Liang, X., Connelly, A. & Calamante, F. (2016) 'A novel joint sparse partial correlation method for estimating group functional networks.' *Human Brain Mapping*, **37**, 1162–1177.

Liu, C., Rubin, D.B. & Wu, Y.N. (1998) 'Parameter expansion to accelerate em: the px-em algorithm.' *Biometrika*, **85**(4), 755–770.

Liu, Y., Yu, C., Zhang, X., Liu, J., Duan, Y., Alexander-Bloch, A.F., Liu, B., Jiang, T. & Bullmore, E. (2014) 'Impaired long distance functional connectivity and weighted network architecture in Alzheimer's disease.' *Cerebral Cortex*, **24**(6), 1422–1435.

Löwdin, P.O. (1950) 'On the non-orthogonality problem connected with the use of atomic wave functions in the theory of molecules and crystals.' *Journal of Chemical Physics*, **18**, 365–375.

Lozano-Soldevilla, D., ter Huurne, N. & Oostenveld, R. (2016) 'Neuronal oscillations with non-sinusoidal morphology produce spurious phase-to-amplitude coupling and directionality.' *Frontiers in Computational Neuroscience*, **10**.

Luckhoo, H. (2013) *Investigating the role of apoe-ϵ4, a risk gene for Alzheimer's disease, on functional brain networks using magnetoencephalography*. Ph.D. thesis, University of Oxford, Oxford, uk.

Luckhoo, H., Hale, J.R., Stokes, M.G., Nobre, A.G., Morris, P.G., Brookes, M.J. & Woolrich, M.W. (2012) 'Inferring task-related networks using independent component analysis in magnetoencephalography.' *NeuroImage*, **62**, 530–541.

Ludwig, J., Duncan, G.J., Gennetian, L.A., Katz, L.F., Kessler, R.C., Kling, J.R. & Sanbonmatsu, L. (2013) 'Long-term neighborhood effects on low-income families: evidence from Moving to Opportunity.' *American Economic Review*, **103**(3), 226–231.

MacKay, D.J.C. (1995) 'Probable networks and plausible predictions—a review of Bayesian methods for supervised neural networks.' *Network: computation in neural systems*, **6**, 469–505.

——— (2003) *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, Cambridge, uk.

Maldjian, J.A., Davenport, E.M. & Whitlow, C.T. (2014) 'Graph theoretical analysis of resting-state MEG data: identifying interhemispheric connectivity and the default mode.' *NeuroImage*, **96**, 88–94.

Mallick, H. & Yi, N. (2014) 'Bayesian methods for high dimensional linear models.' *Journal of Biometrics and Biostatistics*, **1**(5).

Mantini, D., Penna, S.D., Marzetti, L., de Pasquale, F., Pizella, V., Corbetta, M. & Romani, G.L. (2011) 'A signal-processing pipeline for magnetoencephalography resting-state networks.' *Brain Connectivity*, **1**(1), 49–59.

Mantini, D., Perrucci, M.G., Del Gratta, C., Romani, G.L. & Corbetta, M. (2007) 'Electrophysiological signatures of resting state networks in the human brain.' *Proceedings of the National Academy of Sciences USA*, **104**(32), 13 170–13 175.

Marquand, A.F., Rezek, I., Buitelaar, J. & Beckmann, C.F. (2016) 'Understanding heterogeneity in clinical cohorts using normative models: beyond case-control studies.' *Biological Psychiatry*, **80**(7), 552–561.

Marrelec, G., Krainik, A., Duffau, H., Pelegrini-Issac, M., Lehericy, S., Doyon, J. & Benali, H. (2006) 'Partial correlation for functional brain interactivity investigation in functional MRI.' *NeuroImage*, **32**(1), 228–237.

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

Mazumder, R. & Hastie, T. (2012*a*) 'Exact covariance thresholding into connected components for large-scale graphical lasso.' *Journal of Machine Learning Research*, **13**, 723–726.

——— (2012*b*) 'The graphical lasso: new insights and alternatives.' *Electronic Journal of Statistics*, **6**, 2125–2149.

McIntosh, A. & Gonzales-Lima, F. (1994) 'Structural equation modeling and its application to network analysis in functional brain imaging.' *Human Brain Mapping*, **2**, 2–22.

Meek, C. (1995) 'Causal inference and causal explanation with background knowledge.' In *Proceedings of the Eleventh conference on Uncertainty in artificial intelligence*, pp. 403–410. Morgan Kaufmann, San Francisco, ca.

Mehrkanoon, S., Breakspear, M., Britz, J. & Boonstra, T.W. (2014) 'Intrinsic coupling modes in source-reconstructed electroencephalography.' *Brain Connectivity*, **4**(10), 812–825.

Miller, K.L., Alfaro-Almagro, F., Bangerter, N.K., Thomas, D.L., Yacoub, E., Xu, J., Bartsch, A.J., Jbabdi, S., Sotiropoulos, S.N., Andersson, J.L.R., Griffanti, L., Douaud, G., Okell, T.W., Weale, P., Dragonu, I., Garratt, S., Hudson, S., Collins, R., Jenkinson, M., Matthews, P.M. & Smith, S.M. (2016) 'Multimodal population brain imaging in the uk Biobank prospective epidemiological study.' *Nature Neuroscience*, **19**(11), 1523–1536.

Minka, T.P. (2001) *A family of algorithms for approximate Bayesian inference*. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, ma.

——— (2004) 'Power ep.' Technical report, Microsoft Research, Cambridge, uk.

Mira, A. (2001) 'On Metropolis-Hastings algorithms with delayed rejection.' *Metron*, **59**, 231–241.

Mitchell, T.J. & Beauchamp, J.J. (1988) 'Bayesian variable selection in linear regression.' *Journal of the American Statistical Association*, **83**(404), 1023–1032.

Mohammadi, A. & Wit, E.C. (2015) 'Bayesian structure learning in sparse Gaussian graphical models.' *Bayesian Analysis*, **10**(1), 109–138.

Monk, C.S., Peltier, S.J., Wiggins, J.L., Weng, S.J., Carrasco, M., Risi, S. & Lord, C. (2009) 'Abnormalities of intrinsic functional connectivity in autism spectrun disorders.' *NeuroImage*, **47**(2), 764–772.

Mueller, S.G., Weiner, M.W., Thal, L.J., Petersen, R.C., Jack, C., Jagust, W., Trojanowski, J.Q., Toga, A.W. & Becket, L. (2005) 'The alzhemier's disease neuroimaging initiative.' *Neuroimaging Clinics of North America*, **15**(4), 869–877.

Müller, R.A., Shih, P., Keehn, B., Deyoe, J.R., Leyden, K.M. & Shukla, D.K. (2011) 'Underconnected, but how? A survey of functional connectivity MRI studies in autism spectrum disorders.' *Cerebral Cortex*, **21**(10), 2233–2243.

Murphy, K. (2012) *Machine learning: a probabilistic perspective*. MIT Press, Cambridge, MA.

Murray, I., Ghahramani, Z. & MacKay, D.J.C. (2006) 'MCMC for doubly-intractable distributions.' In *Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence*, pp. 359–366. UAI, Cambridge, MA.

Murray, L. (2010) 'Distributed Markov chain Monte Carlo.' In *Proceedings of 24th Conference on Neural Information Processing Systems: Workshop on Learning on Cores, Clusters and Clouds*, volume 11. NIPS, Vancouver, Canada.

Neisser, U., Boodoo, G., Bouchard, T.J.J., Boykin, A.W., Brody, N., Ceci, S.J., Halpern, D.F., Loehlin, J.C., Perloff, R., Sternberg, R.J. & Urbina, S. (1996) 'Intelligence: knowns and unknowns.' *American Psychologist*, **51**, 77–101.

Ng, B., Dressler, M., Varoquaux, G., Poline, J.B., Greicius, M.D. & Thirion, B. (2014) 'Transport on Riemannian manifold for functional connectivity-based classification.' In *MICCAI – 17th International Conference on Medical Image Computing and Computer Assisted Intervention*. Springer, Boston, United States.

Nichols, T.E. & Hayasaka, S. (2003) 'Controlling the familywise error rate in functional neuroimaging: a comparative review.' *Statistical Methods in Medical Research*, **12**(5), 419–446.

Ninness, B., Wills, A. & Gibson, S. (2005) 'The University of Newcastle Identification Toolbox (UNIT).' In *IFAC World Congress*, pp. 1–6.

Nisbett, R.E., Aronson, J., Blair, C., Dickens, W., Flynn, J., Halpern, F. & Turkheimer, E. (2012) 'Intelligence: new findings and theoretical developments.' *American Psychologist*, **67**(2), 130–159.

Nolte, G., Bai, O., Wheaton, L., Mari, Z., Vorbach, S. & Hallett, M. (2004) 'Identifying true brain interaction from EEG using the imaginary part of coherency.' *Journal of Clinical Neurophysiology*, **115**, 2292–2307.

Nolte, G., Ziehe, A., Nikulin, V.V., Schlögl, A., Krämer, N., Brismar, T. & Müller, K.R. (2008) 'Robustly estimating the flow direction of information in complex physical systems.' *Physical Review Letters*, **100**, 234 101.

Nott, D.J. & Green, P.J. (2004) 'Bayesian variable selection and the Swendsen-Wang algorithm.' *Journal of Computational and Graphical Statistics*, **13**(1), 1–17.

Nott, D.J. & Kohn, R. (2005) 'Adaptive sampling for Bayesian variable selection.' *Biometrika*, **92**(4), 747–763.

O'Hara, R.B. & Sillanpää, M.J. (2009) 'A review of Bayesian variable selection methods: what, how and which.' *Bayesian Analysis*, **4**(1), 85–118.

Omidvarnia, A., Azemi, G., Boashash, B., O'Toole, J.M., Colditz, P. & Vanhatalo, S. (2014) 'Measuring time-varying information flow for scalp EEG signals: orthogonalized partial directed coherence.' *IEEE Transactions on Biomedical Imaging*, **61**(3), 680–693.

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.

O'Neill, G.C., Tewarie, P.K., Colclough, G.L., Gascoyne, L.E., Hunt, B.A., Morris, P.G., Woolrich, M.W. & Brookes, M.J. (2016) 'Measurement of dynamic task related functional networks using MEG.' *NeuroImage*. In Press.

Onoda, K., Ishihara, M. & Yamaguchi, S. (2012) 'Decreased functional connectivity by aging is associated with cognitive decline.' *Journal of Cognitive Neuroscience*, **24**(11), 2186–2198.

Oostenveld, R., Fries, P., Maris, E. & Schoffelen, J.M. (2011) 'FieldTrip: open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data.' *Computational Intelligence and Neuroscience*, **2011**, 156 869.

Pakman, A. & Paninski, L. (2013) 'Auxiliary-variable exact Hamiltonian Monte Carlo samplers for binary distributions.' In C.J.C. Burges, L. Bottou, M. Welling, Z. Ghahramani & K.Q. Weinberger, eds., *Advances in Neural Information Processing Systems 26*, pp. 2490–2498. Curran Associates, Inc.

——— (2014) 'Exact Hamiltonian Monte Carlo for truncated multivariate Gaussians.' *Journal of Computational and Graphical Statistics*, **23**(2), 518–542.

Paluš, M. (1997) 'Detecting phase synchronization in noisy systems.' *Physics Letters A*, **235**, 341–351.

Palva, J.M. & Palva, S. (2011) 'Roles of multiscale brain activity fluctuations in shaping the variability and dynamics of psychophysical performance.' In *Slow Brain Oscillations of Sleep, Resting State and Vigilance*, pp. 335–350. Elsevier bv, Amsterdam, The Netherlands.

Palva, S. & Palva, J.M. (2012) 'Discovering oscillatory interaction networks with m/eeg: challenges and breakthroughs.' *Trends in Cognitive Sciences*, **16**(4), 219–230.

Park, T. & Casella, G. (2008) 'The Bayesian lasso.' *Journal of the American Statistical Association*, **103**(482), 681–686.

Pasarica, C. & Gelman, A. (2010) 'Adaptively scaling the Metropolis algorithm using expected squared jumped distance.' *Statistica Sinica*, **20**.

de Pasquale, F., Della Penna, S., Snyder, A.Z., Lewis, P.S., Mantini, D., Marzetti, L., Belardinelli, P., Ciancetta, L., Pizzella, V., Romani, G.L. & Corbetta, M. (2010) 'Temporal dynamics of spontaneous meg activity in brain networks.' *Proceedings of the National Academy of Sciences USA*, **107**(13), 6040–6045.

de Pasquale, F., Della Penna, S., Snyder, A.Z., Marzetti, L., Pizzella, V., Romani, G.L. & Corbetta, M. (2012) 'A cortical core for dynamic integration of functional networks in the resting human brain.' *Neuron*, **74**(4), 753–764.

de Pasquale, F., Della Penna, S., Sporns, O., Romani, G.L. & Corbetta, M. (2015) 'A dynamic core network and global efficiency in the resting human brain.' *Cereberal Cortex*, **26**(10), 4015–4033.

Patel, R., Bowman, F. & Rilling, J. (2006) 'A Bayesian approach to determining connectivity of the human brain.' *Human Brain Mapping*, **27**, 267–276.

Peltola, T., Havulinna, A.S., Salomaa, V. & Vehtari, A. (2014) 'Hierarchical Bayesian survival analysis and projective covariate selection in cardiovascular event risk prediction.' In K.B. Laskey, J. Jones & R. Almond, eds., *Proceedings of Eleventh UAI Bayesian Modeling Applications Workshop (*bmaw *2014)*, volume 1218 of *CEUR Workshop Proceedings*, pp. 79–88. Quebec, Canada.

Peltola, T., Marttinen, P. & Vehtari, A. (2012) 'Metropolis-Hastings algorithm for variable selection in genome-wide association analysis.' *PLoS ONE*, **7**(11), e49 445.

Peng, H., Long, F. & Ding, C. (2005) 'Feature selection based on mutual information: criteria of max-dependency, max-relevance and min-redundancy.' *IEEE Trans. Pattern Analysis and Machine Intelligence*, **27**(8), 1226–1238.

Penny, W.D., Duzel, E., Miller, K.J. & Ojemann, J.G. (2008) 'Testing for nested oscillation.' *Journal of Neuroscience Methods*, **174**(1), 50–61.

Penny, W.D., Kilner, J.M. & Blankenbur, F. (2007) 'Robust Bayesian general linear models.' *NeuroImage*, **36**, 661–671.

Pereira, F., Mitchell, T. & Botvinick, M. (2009) 'Machine learning classifiers and fmri: a tutorial overview.' *NeuroImage*, **45**(1), S199–S209.

Polderman, T.J.C., Benyamin, B., de Leeuw, C.A., Sullivan, P.F., van Bochoven, A., Visscher, P.M. & Posthuma, D. (2015) 'Meta-analysis of the heritability of human traits based on fifty years of twin studies.' *Nature Genetics*, **47**(7), 702–709.

Polson, N.G. & Scott, J.G. (2010) 'Shrink globally, act locally: sparse Bayesian regularization and prediction.' In *Bayesian Statistics 9: Proceedings of the ninth Valencia international meeting*. Oxford University Press, Oxford, U.K.

——— (2012) 'On the half-Cauchy prior for a global scale parameter.' *Bayesian Analysis*, **7**(2), 1–16.

Posthuma, D., de Geus, E.J.C., Mulder, E.J.C.M., Smit, D.J.A., Boomsma, D.I. & Stam, C.J. (2005) 'Genetic components of functional connectivity in the brain: the heritability of synchronization likelihood.' *Human Brain Mapping*, **26**, 191–198.

Posthuma, D., Neale, M.C., Boomsma, D.I. & de Geus, E.J.C. (2001) 'Are smarter brains running faster? Heritability of alpha peak frequency, IQ, and their interrelation.' *Behaviour Genetics*, **31**(6), 567–579.

Power, J.D., Barnes, K.A., Snyder, A.Z., Schlaggar, B.L. & Petersen, S.E. (2012) 'Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion.' *NeuroImage*, **59**(3), 2142–2154.

Power, J.D., Schlaggar, B.L. & Petersen, S.E. (2015) 'Recent progress and outstanding issues in motion correction in resting state fMRI.' *NeuroImage*, **59**, 2142–2154.

Proudfoot, M., Colclough, G.L., Quinn, A.J., Talbot, K., Woolrich, M.W., Nobre, A.G. & Turner, M.R. (in preparation) 'Functional connectivity changes in ALS confirmed by MEG.'

Qiu, H., Han, F., Liu, H. & Caffo, B. (2015) 'Joint estimation of multiple graphical models from high dimensional time series.' *Journal of the Royal Statistical Society B*, **78**(2), 487–504.

Raftery, A., Madigan, D. & Hoeting, J.A. (1997) 'Bayesian model averaging for linear regression models.' *Journal of the American Statistical Association*, **92**(437), 179–191.

Raichle, M.E. (2009) 'A paradigm shift in functional brain imaging.' *Journal of Neuroscience*, **29**(41), 12 729–12 734.

——— (2015*a*) 'The brain's default mode network.' *Annual Reviews Neuroscience*, **38**, 433–447.

——— (2015*b*) 'The restless brain: how intrinsic activity organizes brain function.' *Philosophical Transactions of the Royal Society B*, **370**(1668), 20140 172.

Raichle, M.E., MacLeod, A.M., Snyder, A.Z., Powers, W.J., Gusnard, D.A. & Shulman, G.L. (2001) 'A default mode of brain function.' *Proceedings of the National Academy of Sciences USA*, **98**, 676–682.

Ramsey, J.D., Hanson, S., Hanson, C., Halchenko, Y., Poldrack, R. & Glymour, C. (2010) 'Six problems for causal inference from fMRI.' *NeuroImage*, **49**, 1545–1558.

Ramsey, J.D., Hanson, S.J. & Glymour, C. (2011) 'Multi-subject search correctly identifies causal connections and most causal directions in the DCM models of the Smith *et al.* simulation study.' *NeuroImage*, **58**, 838–848.

Ramsey, J.D., Sanchez-Romero, R. & Glymour, C. (2014) 'Non-Gaussian methods and high-pass filters in the estimation of effective connections.' *NeuroImage*, **84**, 986–1006.

Rasmussen, C.E. & Williams, C.K.I. (2006) *Gaussian processes for machine learning*. MIT Press, Cambridge, MA.

Rezende, D.J., Mohamed, S. & Wierstra, D. (2014) 'Stochastic back-propagation and variational inference in deep latent gaussian models.' In *Proceeding of the 31st International Conference on Machine Learning*, volume 32. JMLR, Beijing, China.

Richiardi, J., Altmann, A., Milazzo, A.C., Chang, C., Chakravarty, M.M., Banaschewski, T., Barker, G.J., Bokde, A.L., Bromberg, U., Büchel, C., Conrod, P., Fauth-Bühler, M., Flor, H., Frouin, V., Gallinat, J., Garavan, H., Gowland, P., Heinz, A., Lemaître, H., Mann, K.F., Martinot, J.L., Nees, F., Paus, T., Pausova, Z., Rietschel, M., Robbins, T.W., Smolka, M.N., Spanagel, R., Ströhle, A., Schumann, G., Hawrylycz, M., Poline, J.B., Greicius, M.D. & imagen consortium (2015) 'Correlated gene expression supports synchronous activity in brain networks.' *Science*, **348**(6240), 1241–1244.

Roberts, G.O., Gelman, A. & Gilks, W.R. (1997) 'Weak convergence and optimal scaling of random walk Metropolis algorithms.' *Annals of Applied Probability*, **7**(1), 110–120.

Robinson, E.C., Jbabdi, S., Glasser, M.F., Andersson, J., Burgess, G.C., Harms, M.P., Smith, S.M., Van Essen, D.C. & Jenkinson, M. (2014) 'msm: a new flexible framework for Multimodal Surface Matching.' *NeuroImage*, **100**, 414–426.

Robinson, S.E. & Vrba, J. (1999) 'Functional neuroimaging by synthetic aperture magnetometry (sam).' In T. Yoshimoto, M. Kotani, S. Kuriki, H. Karibe & N. Nakasato, eds., *Recent Advances in Biomagnetism*, pp. 302–305. Tohoku University Press, Sendai, Japan.

Roelfsema, P.R., Engel, A.K., König, P. & Singer, W. (1997) 'Visuomotor integration is associated with zero time-lag synchronization among cortical areas.' *Nature*, **385**, 157–161.

Ronan, L. & Fletcher, P.C. (2015) 'From genes to folds: a review of cortical gyrification theory.' *Brain Structure and Function*, **220**(5), 2475–2483.

Rosenberg, M.D., Finn, E.S., Scheinost, D., Papademetris, X., Shen, X., Constable, R.T. & Chun, M.M. (2016) 'A neuromarker of sustained attention from whole-brain functional connectivity.' *Nature Neuroscience*, **19**(1), 165–171.

Roverato, A. (2002) 'Hyper inverse Wishart distribution for non-decomposable graphs and its application to Bayesian inference for Gaussian graphical models.' *Scandinavian Journal of Statistics*, **29**, 391–411.

Rubinov, M. & Sporns, O. (2010) 'Complex network measures of brain connectivity: uses and interpretations.' *NeuroImage*, **52**, 1059–1069.

Ryali, S., Chen, T., Supekar, K. & Menon, V. (2012) 'Estimation of functional connectivity in fmri data using stability selection-based sparse correlation with elastic net penalty.' *NeuroImage*, **59**(4), 3852–3861.

Sala-Llonch, R., Bartrés-Faz, D. & Junqué, C. (2015) 'Reorganization of brain networks in aging: a review of functional connectivity studies.' *Frontiers in Psychology*, **6**, 663.

Sala-Llonch, R., Pena-Gomez, C., Arenaza-Urquijo, E.M., Vidal-Pineiro, D., Bargallo, N., Junque, C. & Bartres-Faz, D. (2012) 'Brain

connectivity during resting state and subsequent working memory task predicts behavioural performance.' *Cortex*, **48**, 1187–1196.

Salimi-Khorshidi, G., Douaud, G., Beckmann, C.F., Glasser, M.F., Griffanti, L. & Smith, S.M. (2014) 'Automatic denoising of functional MRI data: combining independent component analysis and hierarchical fusion of classifiers.' *NeuroImage*, **90**, 449–468.

Sarvas, J. (1987) 'Basic mathematical and electromagnetic concepts of the biomagnetic inverse problem.' *Physics in Medicine and Biology*, **32**(1), 11–22.

Scannell, J.W., Burns, G.A., Hilgetag, C.C., O'Neill, M.A. & Young, M.P. (1999) 'The connectional organization of the cortico-thralamic system of the cat.' *Cereberal Cortex*, **9**(3), 277–299.

Schlee, W., Leirer, V., Kolassa, I.T., Weisz, N. & Elbert, T. (2012) 'Age-related changes in neural functional connectivity and its behavioural relevance.' *BMC Neuroscience*, **13**(1), 16.

Schoffelen, J.M. & Gross, J. (2009) 'Source connectivity analysis with MEG and EEG.' *Human Brain Mapping*, **30**, 1857–1865.

Schutte, N.M., Hansell, N.K., de Geus, E.J.C., Martin, N.G., Wright, M.J. & Smit, D.J.A. (2013) 'Heritability of resting state EEG functional connectivity patterns.' *Twin Research and Human Genetics*, **16**(5), 962–969.

Seeley, W.W., Crawford, R.K., Zhou, J., Miller, B.L. & Greicius, M.D. (2009) 'Neurodegenerative diseases target large-scale human brain networks.' *Neuron*, **62**(1), 42–52.

Seeley, W.W., Menon, V., Schatzberg, A.F., Keller, J., Glover, G.H., Kenna, H., Reiss, A.L. & Greicius, M.D. (2007) 'Dissociable intrinsic connectivity networks for salience processing and executive control.' *Journal of Neuroscience*, **27**(9), 2349–2356.

Shafto, M.A. & Tyler, L.K. (2014) 'Language in the aging brain: The network dynamics of cognitive decline and preservation.' *Science*, **346**(6209), 583–587.

Sheline, Y. & Raichle, M.E. (2013) 'Resting state functional connectivity in preclinical Alzheimer's disease.' *Biological Psychiatry*, **74**(5), 340–347.

Shimizu, S. (2014) 'LiNGAM: non-Gaussian methods for estimating causal structures.' *Behaviormetrika*, **41**(1), 65–98.

Shimizu, S., Hoyer, P.O., Hyvärinen, A. & Kerminen, A. (2006) 'A linear non-Gaussian acyclic model for causal discovery.' *Journal of Machine Learning Research*, **7**, 2003–2030.

Shulman, G.L., Fiez, J.A., Corbetta, M., Buckner, R.L., Miezin, F.M., Raichle, M.E. & Peterson, S.E. (1997) 'Common blood flow changes across visual tasks ii: decreases in cerebral cortex.' *Cognitive Neuroscience*, **9**(5), 648–663.

Siegel, J.S., Mitra, A., Laumann, T.O., Seitzman, B.A., Raichle, M., Corbetta, M. & Snyder, A.Z. (2016) 'Data quality influences observed links between functional connectivity and behaviour.' *Cerebral Cortex*. In Press.

Smit, D.J.A., Stam, C.J., Posthuma, D., Boomsma, D.I. & de Geus, E.J.C. (2008) 'Heritability of "small-world" networks in the brain: a graph theoretical analysis of resting-state eeg functional connectivity.' *Human Brain Mapping*, **29**, 1368–1378.

Smith, D.M. (1989) 'All possible subset regressions using the qr decomposition.' *Journal of the Royal Statistical Society C*, **40**(3), 502–513.

Smith, S.M. (2012) 'The future of fmri connectivity.' *Neuroimage*, **62**, 1257–1266.

Smith, S.M., Beckmann, C.F., Ramnani, N., Woolrich, M.W., Bannister, P.R., Jenkinson, M., Matthews, P.M. & McGonigle, D.J. (2005) 'Variability in fmri: a re-examination of inter-session differences.' *Human Brain Mapping*, **24**(3), 248–257.

Smith, S.M., Miller, K.L., Glahn, D.C., Fox, P.M., Mackay, C.E., Filippini, N., Watkins, K.W., Toro, R., Laird, A.R. & Beckmann, C.F. (2009) 'Correspondence of the brain's functional architecture during activation and rest.' *Proceedings of the National Academy of Sciences USA*, **106**(31), 13 040–13 045.

Smith, S.M., Miller, K.L., Moeller, S., Xu, J., Auerbach, E.J., Woolrich, M.W., Beckmann, C.F., Jenkinson, M., Andersson, J., Glasser, M.F., Van Essen, D.C., Feinberg, D.A., Yacoub, E. & Ugurbil, K. (2012) 'Temporally-independent functional modes of spontaneous brain activity.' *Proceedings of the National Academy of Sciences USA*, **109**(8), 3131–3136.

Smith, S.M., Miller, K.L., Salimi-Khorshidi, G., Webster, M., Beckmann, C.F., Nichols, T.E., Ramsey, J.D. & Woolrich, M.W. (2011) 'Network modelling methods for fMRI.' *Neuroimage*, **54**, 875–891.

Smith, S.M., Nichols, T.E., Vidaurre, D., Winkler, A.M., Behrens, T.E., Glasser, M.F., Ugurbil, K., Barch, D.M., Van Essen, D.C. & Miller, K.L. (2015) 'A positive-negative mode of population covariation links brain connectivity, demographics and behavior.' *Nature Neuroscience*, **18**, 1565–1567.

Smith, S.M., Vidaurre, D., Beckmann, C.F., Glasser, M.F., Jenkinson, M., Miller, K.L., Nichols, T.E., Robinson, E., Salimi-Khorshidi, G., Woolrich, M.W., Barch, D.M., Ugurbil, K. & Van Essen, D.C. (2013) 'Functional connectomics from resting-state fMRI.' *Trends in Cognitive Sciences*, **17**(12), 666–682.

Sotiropoulos, S.N., Chen, X., Smith, S.M., Van Essen, D.C., Behrens, T.E., Nichols, T.E. & Jbabdi, S. (2015) 'Heritability of structural connections from HCP diffusion MRI data.' In *Proceedings of the 23rd Annual Conference of the International Society for Magnetic Resonance in Medicine*, p. 721. Toronto, Canada.

Soto, J.L.P., Lachaux, J.P., Baillet, S. & Jerbi, K. (2016) 'A multivariate method for estimating cross-frequency neuronal interactions and correcting linear mixing in MEG data, using canonical correlations.' *Journal of Neuroscience Methods*, **271**, 169–181.

Stam, C.J. (2014) 'Modern network science of neurological disorders.' *Nature Reviews Neuroscience*, **15**, 683–695.

Stam, C.J., de Haan, W., Daffertshofer, A., Jones, B.F., Manshanden, I., van Cappellen van Walsum, A.M., Montez, T., Verbunt, J.P.A., de Munck, J.C., van Dijk, B.W., Berendse, H.W. & Sheltens, P. (2009)

'Graph theoretical analysis of magnetoencephalographic functional connectivity in Alzheimer's disease.' *Brain*, **132**, 213–224.

Stam, C.J., Nolte, G. & Daffertshofer, A. (2007) 'Phase lag index: assessment of functional connectivity from multi channel EEG and MEG with diminished bias from common sources.' *Human Brain Mapping*, **28**(11), 1178–1193.

Stam, C.J. & van Straaten, E.C.W. (2012) 'The organization of physiological brain networks.' *Clinical Neurophysiology*, **123**, 1067–1087.

Stevner, A. (in preparation) *Whole-brain spatio-temporal characteristics of functional connectivity in transitions between wakefulness and sleep.* Ph.D. thesis, University of Oxford, Oxford, UK.

Storvik, G. (2011) 'On the flexibility of Metropolis-Hastings acceptance probabilities in auxilliary variable proposal generation.' *Scandinavian Journal of Statistics*, **38**(2), 342–358.

van Straaten, E.C.W. & Stam, C.J. (2013) 'Structure out of chaos: functional brain network analysis with EEG, MEG and functional MRI.' *European Neuropsychopharmacology*, **23**, 7–18.

Swendsen, R.H. & Wang, J.S. (1987) 'Nonuniversal critical dynamics in Monte Carlo simulations.' *Physical Review Letters*, **58**(2), 86–88.

Tallinen, T., Chung, J.Y., Rousseau, F., Girard, N., Lefèvre, J. & Mahadevan, L. (2016) 'On the growth and form of cortical convolutions.' *Nature Physics*, **12**, 588–593.

Tashiro, T., Shimizu, S., Hyvärinen, A. & Takashi, W. (2014) 'ParceLiNGAM: a causal ordering method robust against latent confounders.' *Neural Computation*, **26**(1), 57–83.

Tavor, I., Parker Jones, O., Mars, R.B., Smith, S.M., Behrens, T.E. & Jbabdi, S. (2016) 'Task-free MRI predicts individual differences in brain activity during task performance.' *Science*, **352**(6282), 216–220.

Tewarie, P., Hillebrand, A., van Dellen, E., Schoonheim, M.M., Barkhof, F., Polman, C.H., Beauliue, C., Gong, G., van Dijk, B.W. & Stam,

C.J. (2014) 'Structural degrees predicts functional network connectivity: A multimodal resting-state fMRI and MEG study.' *NeuroImage*, **97**, 296–307.

Thompson, P.M., Ge, T., Glahn, D.C., Jahanshad, N. & Nichols, T.E. (2013) 'Genetics of the connectome.' *NeuroImage*, **80**, 475–488.

Tibshirani, R. (1996) 'Regression shrinkage and variable selection via the lasso.' *Journal of the Royal Statistical Society B*, **58**(1), 267–288.

Tomasi, D. & Volkow, N.D. (2012) 'Aging and functional brain networks.' *Molecular Psychiatry*, **17**, 549–558.

Tort, A.B.L., Komorowski, R., Eichenbaum, H. & Kopell, N. (2010) 'Measuring phase-amplitude coupling between neuronal oscillations of different frequencies.' *Journal of Neurophysiology*, **104**(2), 1195–1210.

Trias, M., Vecchio, A. & Veitch, J. (2009) 'Delayed rejection schemes for efficient Markov chain Monte Carlo sampling of multimodal distributions.' *ArXiv*:0904.2207v2.

Uhler, C., Lenkoski, A. & Richards, D. (2016) 'Exact formulas for the normalizing constants of Wishart distributions for graphical models.' *ArXiv*:1406.4901.

Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K. & for the WU-Minn hcp consortium (2013) 'The WU-Minn Human Connetome Project: an overview.' *NeuroImage*, **80**, 62–97.

Van Veen, B.D., van Drongelen, W., Yuchtman, M. & Suzuki, A. (1997) 'Localization of brain electrical activity via linearly constrained minimum variance spatial filtering.' *IEEE Transactions on Biomedical Imaging*, **44**, 867–880.

Varoquaux, G. & Craddock, R.C. (2013) 'Learning and comparing functional connectomes across subjects.' *NeuroImage*, **80**, 405–415.

Varoquaux, G., Gramfort, A., Poline, J.B., Thirion, B., Zemel, R. & Shawe-Taylor, J. (2010) 'Brain covariance selection: better individual functional connectivity models using population prior.' In *Advances in Neural Information Processing Systems*, volume 23. nips, Vancouver, Canada.

Varoquaux, G., Raamana, P.R., Engemann, D., Hoyos-Idrobo, A., Schwartz, Y. & Thirion, B. (2016) 'Assessing and tuning brain decoders: cross-validation, caveats, and guidelines.' *ArXiv*:1606.05201v1.

Vidaurre, D., Quinn, A.J., Baker, A.P., Dupret, D., Tejero-Cantero, A. & Woolrich, M.W. (2016) 'Spectrally resolved fast transient brain states in electrophysiological data.' *NeuroImage*, **126**, 81–95.

Vinck, M., Oostenveld, R., van Wingerden, M., Battaglia, F. & Pennartz, C.M.A. (2011) 'An improved index of phase-synchronization for electrophysiological data in the presence of volume-conduction, noise and sample bias.' *NeuroImage*, **55**, 1548–1565.

Wang, H. (2012*a*) 'Bayesian graphical lasso models and efficient posterior computation.' *Bayesian Analysis*, **7**(2), 771–790.

——— (2012*b*) 'Efficient Gaussian graphical model determination under G-Wishart distributions.' *Electronic Journal of Statistics*, **6**, 168–198.

——— (2015) 'Scaling it up: stochastic search structure learning in graphical models.' *Bayesian Analysis*, **10**(2), 351–377.

Wang, H.E., Benar, C.G., Quilichini, P.P., Friston, K.J., Jirsa, V. & Bernand, C. (2014) 'A systematic framework for functional connectivity measures.' *Frontiers in Neuroscience*, **8**, 405.

Wang, X.J. (2010) 'Neurophysiological and computational principles of cortical rhythms in cognition.' *Physiological Reviews*, **90**(3), 1195–1268.

Wens, V. (2014) 'Investigating complex networks with inverse models: Analytical aspects of spatial leakage and connectivity estimation.' *Phys. Rev. E*, **91**, 012 823.

Wens, V., Marty, B., Mary, A., Op de Beeck, M., Goldman, S., Van Bogaert, P., Peigneux, P. & De Tiege, X. (2015) 'A geometric correction scheme for spatial leakage effects in MEG/EEG seed-based functional connectivity mapping.' *Human Brain Mapping*, **36**(11), 4604–4621.

Wens, V., Mary, A., Bourguignon, M., Goldman, S. & Marty, B. (2014) 'About the electrophysiological basis of resting state networks.' *Clinical Neurophysiology*, **125**, 1711–1713.

Wilmer, A., de Lussanet, M. & Lappe, M. (2012) 'Time-delayed mutual information of the phase as a measure of functional connectivity.' *PLoS ONE*, **7**(9), e044 633.

Winkler, A.M., Ridgway, G.R., Douaud, G., Nichols, T.E. & Smith, S.M. (2016*a*) 'Faster permutation inference in brain imaging.' *NeuroImage*, **141**, 502–516.

Winkler, A.M., Ridgway, G.R., Webster, M.A., Smith, S.M. & Nichols, T.E. (2014) 'Permutation inference for the general linear model.' *NeuroImage*, **92**, 381–397.

Winkler, A.M., Webster, M.A., Brooks, J.C., Tracey, I., Smith, S.M. & Nichols, T.E. (2016*b*) 'Non-parametric combination and related permutation tests for neuroimaging.' *Human Brain Mapping*, **37**, 1486–1511.

Winkler, A.M., Webster, M.A., Vidaurre, D., Nichols, T.E. & Smith, S.M. (2015) 'Multi-level block permutation.' *NeuroImage*, **123**, 253–268.

Womelsdorf, T. & Fries, P. (2007) 'The role of neuronal synchronization in selective attention.' *Current Opinion in Neurobiology*, **17**(2), 154–160.

Woolrich, M.W. (2008) 'Robust group analysis using outlier inference.' *NeuroImage*, **41**, 286–301.

Woolrich, M.W., Behrens, T.E.J., Beckmann, C.F., Jenkinson, M. & Smith, S.M. (2004) 'Multilevel linear modelling for fMRI group analysis using Bayesian inference.' *NeuroImage*, **21**, 1732–1747.

Woolrich, M.W., Hunt, L., Groves, A.R. & Barnes, G.R. (2011) 'MEG beamforming using Bayesian PCA for adaptive data covariance matrix regularization.' *NeuroImage*, **57**, 1466–1479.

Woolrich, M.W., Jbabdi, S., Patenaude, B., Chappell, M., Makni, S., Behrens, T., Beckmann, C., Jenkinson, M. & Smith, S.M. (2009) 'Bayesian analysis of neuroimaging data in FSL.' *NeuroImage*, **45**(1), s173–s186.

Woolrich, M.W. & Stephan, K.E. (2013) 'Biophysical network models and the human connectome.' *NeuroImage*, **80**, 330–338.

Worsley, K.J., Evans, A.C., Marrett, S. & Neelin, P. (1992) 'A three-dimensional statistical analysis for CBF activation studies in human brain.' *Journal of Cerebral Blood Flow and Metabolism*, **12**, 900–918.

Yang, S., Lu, Z., Shen, X., Wonka, P. & Ye, J. (2015) 'Fused multiple graphical lasso.' *Society for Industrial & Applied Mathematics Journal of Optimisation*, **25**(2), 916–943.

Yeo, B.T.T., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., Roffman, J.L., Smoller, J.W., Zollei, L., Polimeni, J.R., Fischl, B., Liu, H. & Buckner, R.L. (2011) 'The organization of the human cerebral cortex estimated by intrinsic functional connectivity.' *Journal of Neurophysiology*, **106**(3), 1125–1165.

Zalesky, A., Cocchi, L., Fornito, A., Murray, M.M. & Bullmore, E.T. (2012) 'Connectivity differences in brain networks.' *NeuroImage*, **60**, 1055–1062.

Zalesky, A., Fornito, A. & Bullmore, E.T. (2010) 'Network-based statistic: identifying differences in brain networks.' *NeuroImage*, **53**, 1197–1207.

Zhu, D., Zhan, L., Faskowitz, J., Daianu, M., Jahanshad, N., de Zubicaray, G.I., McMahon, K.L., Martin, N.G., Wright, M.J. & Thompson, P.M. (2015) 'Genetic analysis of structural brain connectivity using DICCOL models of diffusion MRI in 522 twins.' In *Proceedings of the IEEE International Symposium on Biomedical Imaging*, pp. 1167–1171. ISBI, Brooklyn, NY.

Zou, H. & Hastie, T. (2005) 'Regularization and variable selection via the elastic net.' *Journal of the Royal Statistical Society B*, **67**(2), 301–320.

Zou, Q., Ross, T.J., Gu, H., Geng, X., Zuo, X.N., Hong, L.E., Gao, J.H., Stein, E.A., Zang, Y.F. & Yang, Y. (2013) 'Intrinsic resting-state activity predicts working memory brain activation and behavioural performance.' *Human Brain Mapping*, **34**(12), 3204–3215.