A Multimodal Investigation into the Relationship Between Ongoing Brain State Dynamics and Visual Response Variability

BRANDON THOMAS INGRAM

A thesis submitted to
The University of Birmingham for the degree of

DOCTOR OF PHILOSOPHY

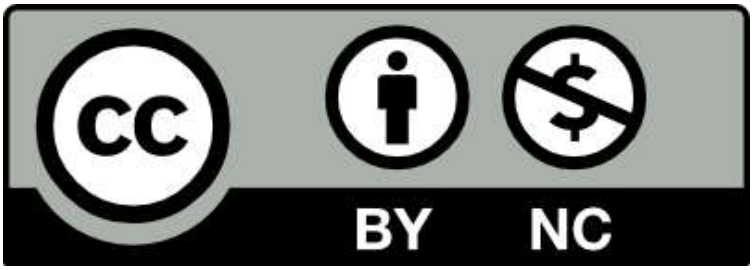


UNIVERSITY OF BIRMINGHAM

School of Psychology
College of Life and Environmental Sciences
University of Birmingham
September 2023

UNIVERSITY OF BIRMINGHAM

University of Birmingham Research Archive
e-theses repository



This unpublished thesis/dissertation is under a Creative Commons Attribution-NonCommercial 4.0 International (CC BY-NC 4.0) licence.

You are free to:

Share — copy and redistribute the material in any medium or format

Adapt — remix, transform, and build upon the material

The licensor cannot revoke these freedoms as long as you follow the license terms.

Under the following terms:



Attribution — You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.



NonCommercial — You may not use the material for commercial purposes.

No additional restrictions — You may not apply legal terms or technological measures that legally restrict others from doing anything the license permits.

Notices:

You do not have to comply with the license for elements of the material in the public domain or where your use is permitted by an applicable exception or limitation.

No warranties are given. The license may not give you all of the permissions necessary for your intended use. For example, other rights such as publicity, privacy, or moral rights may limit how you use the material.

Unless otherwise stated, any material in this thesis/dissertation that is cited to a third-party source is not included in the terms of this licence. Please refer to the original source(s) for licencing conditions of any quotes, images or other material cited to a third party.

Dedication

This thesis is dedicated to my Grandfather, Thomas Edward Jones.

Acknowledgements

First and foremost, I would like to thank my supervisors, Professor Andrew Bagshaw and Professor Stephen Mayhew, for their continual guidance and support throughout my PhD. I would also like to thank the Biotechnology and Biological Sciences Research Council (BBSRC) and the Midlands Integrative Training Programme (MITBP) for funding the project and providing training opportunities. Additionally, I would like to thank all of the members of the Multimodal imaging group for their support, with a special thank you to Sara for her assistance during data collection. Lastly, I would like to give a special thank you to my Mum, Dad, Brother, and my partner, Carl, for their words of encouragement and continual support.

Abstract

The human visual system exhibits a large amount of variability in visual responses, even when presented with identical stimuli, yet the mechanisms behind this variability remain poorly understood. This thesis utilised simultaneous EEG-fMRI and Hidden Markov Modelling to investigate ongoing brain state dynamics and evaluate their impact on visual responses. The first results chapter (Chapter 3) investigates the ability of Hidden Markov modelling to identify differences in BOLD and alpha brain state dynamics between eyes open and eyes closed rest. The results showed that each rest condition shows significantly different BOLD and alpha brain state dynamics, with the open condition being associated with attention-based network states and a reduction in alpha power. Hidden Markov modelling was then applied to a visual dataset, in which participants viewed left visual field checkerboards passively (i.e. no task) and actively (i.e. identifying target stimuli). A total of three separate models were trained: a whole-brain BOLD state model (Chapter 4), an alpha state model (Chapter 5) and a thalamus BOLD state model (Chapter Chapter 6). The results of these models demonstrated that the HMM state active at baseline had a significant impact on the subsequently evoked BOLD, VEP and alpha event-related de/synchronisation responses. Thus, the ongoing brain state dynamics, defined by BOLD and electrophysiological measures, explain a significant proportion of the variability in visual responses, highlighting the importance of the baseline. Overall, this thesis demonstrates that ongoing brain state dynamics interact with visual responses, resulting in the observed variability, highlighting the importance of accounting for the baseline in visual research and neuroimaging.

Acronyms

ANOVA Analysis of Variance.

BCG Ballistocardiogram.

BIDS Brain Imaging Data Structure.

BOLD Blood Oxygen Level Dependent.

CSF Cerebrospinal Fluid.

DMN Default Mode Network.

EC Eyes Closed.

ECG Electrocardiogram.

EEG Electroencephalography.

EMM Estimated Marginal Mean.

EO Eyes Open.

EPI Echo Planar Imaging.

ERDS Event-Related Desyhcronisation / Synchronisation.

ERP Event-Related Potential.

fMRI Functional Magnetic Resonance Imaging.

FO Fractional Occupancy.

FOV Field-of-View.

FPN Frontoparietal Network.

FWHM Full Width at Half Maximum.

GLM General Linear Model.

HMM Hidden Markov Model.

IC Independent Component.

ICA Independent Component Analysis.

ISI Inter-stimulus Interval.

LGN Lateral Geniculate Nucleus.

LMM Linear Mixed Model.

MAR Multivariate Autoregressive.

MEG Magnetoencephalography.

MRI Magnetic Resonance Imaging.

NMR Nuclear Magnetic Resonance.

PSC Percent Signal Change.

PSD Power Spectral Density.

ROI Region of Interest.

SE Standard Error.

TDE Time Delay Embedded.

TE Echo Time.

TMS Transcranial Magnetic Stimulation.

TR Repetition Time.

TRN Thalamic Reticular Nucleus.

VCG Vectorcardiogram.

VEP Visually Evoked Potential.

VFE Variational Free Energy.

WM White Matter.

List of Figures

1.1 Visual System Diagram 18

1.2 MRI Diagram 29

2.1 Example Stimuli 52

2.2 Psychometric Function 53

2.3 Psychometric Function Equation. 53

2.4 VEP and Alpha ERD/S Windows 67

2.5 BOLD Peak and Undershoot Windows 68

2.6 Percent Signal Change Equation 68

3.1 fMRI ICA Components 77

3.2 Alpha Power Analysis 83

3.3 HMM Model Selection 84

3.4 fMRI HMM States 87

3.5 EEG HMM States 88

3.6 Conventional Alpha-BOLD Analysis 90

3.7 FO Sliding Window Analysis 93

3.8 fMRI HMM State Metrics 94

3.9 EEG HMM State Metrics 95

4.1 Group ICA Decomposition 106
4.2 Regions of Interest 109
4.3 GLM Activation Maps 113
4.4 Condition ROI 114
4.5 Stimulus Type ROI 115
4.6 HMM Model Selection 117
4.7 HMM States 120
4.8 Evoked HMM Response 121
4.9 Baseline Distribution and Reliability 122
4.10 HMM State Metrics 123
4.11 HMM Evoked BOLD 124
4.12 Impact of HMM State on BOLD Responses Post Hoc 127
4.13 Impact of HMM State on Alpha ERD/S 129
4.14 Impact of HMM State on VEPs 130

5.1 VEP Condition Comparison 146
5.2 VEP Stimulus Type Comparison 147
5.3 Alpha ERD/S Condition Comparison 148
5.4 Alpha ERD/S Stimulus Type Comparison 149
5.5 HMM Model Selection 150
5.6 HMM Spectra 151
5.7 High Alpha State BOLD Correlates 152
5.8 Evoked HMM Response 153
5.9 Baseline State Distribution 154
5.10 State Metrics 156
5.11 Impact of Alpha State on VEPs 157
5.12 Impact of Alpha State on Alpha ERD/S 159

5.13 Impact of Alpha State on BOLD Responses 161

6.1 LGN ROI 174
6.2 Thalamus ICA Components 178
6.3 Condition Comparison of Right LGN BOLD Responses 179
6.4 Stimulus Type Comparison of Right LGN BOLD Responses 180
6.5 Thalamus HMM Model Selection 181
6.6 Thalamus HMM States 183
6.7 Evoked HMM Response 184
6.8 Thalamus State Baseline Distribution 185
6.9 HMM State Metrics 186
6.10 Thalamus State BOLD Correlates 189
6.11 Thalamus States and Alpha Power 190
6.12 Thalamus States and Alpha Power Post Hoc Tests 191
6.13 Impact of Thalamus States on BOLD Responses 193
6.14 Impact of Thalamus States on BOLD Responses Post Hoc 196
6.15 Impact of Thalamus States on VEP Responses 198
6.16 Impact of Thalamus States on VEP Responses Post Hoc 199
6.17 Impact of HMM State on Alpha ERD/S 201

Contents

Introduction **15**

Abstract 15

The Anatomy of the Human Visual System 16

The Human Eye 16

The Thalamus 17

Primary Visual Cortex 19

Extrastriate Cortex, and the Ventral and Dorsal Pathways 20

Electrophysiology and Neuroimaging Techniques 21

Animal Research 21

Electroencephalography 22

Magnetic Resonance Imaging 27

Silmutaneous EEG-fMRI 32

Variability of Visual Responses 37

Spontaneous Activity 38

Interactions between Spontaneous Activity and Visual Responses 41

Defining Brain State 42

Identifying a Brain State 44

Thesis Aims 47

Methods **49**

Abstract . . . 49

Datasets . . . 49

Eyes Open / Eyes Closed Dataset . . . 49

Visual Response Variability Dataset . . . 50

Analyses . . . 56

High Performance Cluster and Containers . . . 56

BIDS Conversion . . . 56

MRI Preprocessing . . . 57

EEG Preprocessing . . . 61

Hidden Markov Modelling . . . 62

Determining Baseline States . . . 65

VEP and Alhpa ERD/S Analysis . . . 65

BOLD Response Analysis . . . 66

Statistics . . . 67

Brain State Dynamics Differ Between Eyes Open and Eyes Closed Rest **70**

Abstract . . . 70

Introduction . . . 71

Methods . . . 74

Participants . . . 74

EEG-fMRI Data Acquisition . . . 75

fMRI Prepcrocessing . . . 76

fMRI Group ICA . . . 76

EEG Preprocessing . . . 78

GLM Analysis Investigating Alpha-BOLD Relationships . . . 78

Alpha Power Analysis . . . 79

Hidden Markov Modelling . . . 79
Estimating Spectra and Spatial Properties of EEG States . . . 81
State Metric Analysis . . . 81
Results . . . 82
Alpha Power Analysis . . . 82
Model Selection . . . 82
HMM States . . . 85
Conventional Alpha-BOLD GLM . . . 89
Fractional Occupancy Sliding Window Analysis . . . 91
State Metrics Comparison . . . 93
Discussion . . . 96
Summary . . . 96
fMRI HMM States . . . 96
EEG HMM States . . . 98
Limitations and Future Work . . . 100
Conclusions . . . 100

Whole-Brain BOLD States Modulate BOLD Responses and P100 Amplitudes 102

Abstract . . . 102
Introduction . . . 102
Methods . . . 104
Dataset . . . 104
Participants . . . 104
EEG Preprocessing . . . 105
fMRI Preprocessing . . . 105
Group fMRI ICA Decomposition . . . 105
General Linear Modelling . . . 106

Defining ROIs . . . 108
Extracting BOLD Responses . . . 108
Extracting VEPs and Alpha ERD/S . . . 108
Hidden Markov Modelling . . . 108
Evoked Fractional Occupancy . . . 110
Statistics . . . 110
Results . . . 110
General Linear Modelling . . . 110
ROI Analysis . . . 113
HMM Model Selection . . . 116
HMM States . . . 116
Evoked HMM Response . . . 120
Distribution and Reliability of Baseline States . . . 120
State Metric Comparison . . . 121
Impact of Whole-Brain HMM States on BOLD Responses . . . 122
Impact of Whole-Brain HMM States on Alpha ERD/S . . . 126
Impact of Whole-Brain HMM States on VEPs . . . 128
Discussion . . . 131
Summary . . . 131
BOLD Activation Maps Differ Significantly Between Conditions . . . 132
BOLD Responses Differ Significantly Between Conditions . . . 133
Whole-Brain HMM States . . . 134
HMM State Dynamics . . . 134
Impact of Whole-Brain HMM States on Responses . . . 135
Limitations and Future Research . . . 136
Conclusion . . . 137

Alpha State Modulates Stimulus Induced Alpha ERD/S **139**

Abstract . . . 139

Introduction . . . 140

Methods . . . 141

Dataset . . . 141

Participants . . . 141

EEG Preprocessing . . . 142

fMRI Preprocessing . . . 142

Extracting BOLD Responses . . . 142

Extracting VEPs and Alpha ERD/S . . . 142

Alpha HMM . . . 143

Evoked Fractional Occupancy . . . 143

Alpha State GLM . . . 143

Statistics . . . 144

Results . . . 145

Visually Evoked Potentials . . . 145

Event-Related Alpha Synchronisation and Desynchronisation . . . 145

HMM Model Selection . . . 147

HMM States . . . 149

Evoked HMM Response . . . 151

Distribution and Reliability of Baseline States . . . 153

State Metric Comparison . . . 154

Impact of Alpha State on VEPs . . . 155

Impact of Alpha State on Alpha ERD/S . . . 157

Impact of Alpha State on BOLD Responses . . . 159

Discussion . . . 162

Summary . . . 162
Visually Evoked Potentials and Alpha ERD/S Differ Between Conditions . . . 163
HMM States . . . 164
HMM State Metric Comparison . . . 164
Impact of Baseline Alpha State on Stimulus Induced Alpha ERD/S . . . 165
Impact of Baseline Alpha State on BOLD Responses . . . 166
Impact of Baseline Alpha State on VEPs . . . 167
Future Research . . . 167
Conclusion . . . 168

Thalamus BOLD States Modulate BOLD Responses and P250 Amplitudes 169

Abstract . . . 169
Introduction . . . 170
Methods . . . 171
Dataset . . . 171
Participants . . . 171
EEG Preprocessing . . . 172
fMRI Preprocessing . . . 172
Thalamic ICA . . . 172
Thalamic HMM . . . 173
ROI Analysis . . . 173
Extracting BOLD Responses . . . 174
Extracting VEPs and Alpha ERD/S . . . 174
Cortical Correlates of Thalamic States . . . 174
Inferring Spectral Properties of Thalamic States . . . 176
Evoked Fractional Occupancy . . . 176
Statistics . . . 177

Results . . . 177
Thalamic ICA . . . 177
ROI Analysis . . . 178
Model Selection . . . 180
HMM States . . . 181
Evoked HMM Thalamic Response . . . 182
Distribution and Reliability of Baseline States . . . 184
State Metric Comparison . . . 185
Thalamic State Correlates . . . 187
Thalamic States and Alpha Power . . . 189
Impact of Thalamus State on BOLD Responses . . . 192
Impact of Thalamus State on VEPs . . . 197
Impact of Thalamus State on Stimulus Induced Alpha Responses . . . 198
Discussion . . . 201
Summary . . . 201
BOLD Responses in the LGN are Modulated by Attention . . . 202
Thalamic States and Cortical Correlates . . . 203
Alpha Power Correlations . . . 204
Impact of Baseline Thalamus State on BOLD Responses . . . 204
Impact of Baseline Thalamus State on VEPs . . . 205
Impact of Baseline Thalamus State on Stimulus Induced Alpha ERD/S . . . 206
Future Research . . . 207
Conclusions . . . 207

Thesis Summary **209**

Overview . . . 209
Significant Findings . . . 210

Future Directions 212
Conclusion 214

Chapter 1

Introduction

Abstract

The human brain is the most complex organ in the human body. Comprised of billions of neurons, the brain can perform thousands of different computations at any given moment, ranging from controlling homeostatic functions, such as breathing, to orchestrating higher-level cognitive functions, such as attention and memory. Within Neuroscience research, the functionality of the brain has traditionally been investigated from a black-box perspective, in which stimuli are shown to a participant, and the subsequent responses within the brain are measured, revealing when and where different information is processed. Whilst this approach has expanded our knowledge of the functionality of the brain, recent research has moved towards a more internal perspective, in which the neural activity within the brain is examined in the absence of stimuli. This change is partly because the brain exhibits significant sensory response variability when presented with the same stimulus, something the traditional black-box perspective cannot explain. Such research has revealed that the brain performs equally, if not more, complex functions during rest as when presented with stimuli, showing complex network-based activity spanning across both cortical and subcortical regions. This neural activity is often referred to as spontaneous since it occurs in the absence of stimuli. Whilst initially thought to be noise, research has shown that when a participant is performing a task, spontaneous activity is often associated with various elements of cognition and can modulate the responses to stimuli, explaining a part of the variability observed in sensory responses. This has been observed frequently

in visual processing, with pre-stimulus activity modulating the magnitude of the subsequent visual responses and biasing visual perception. Furthermore, spontaneous activity is highly organised and often shows cyclic properties, with specific patterns of activity reoccurring over time, which are now widely referred to as brain states. Therefore, this thesis aims to investigate the influence of these brain states on visual response variability, using simultaneously recorded EEG and fMRI data to investigate the role of and relationships between cross-modal definitions of brain state.

Chapter 1 provides overview of the relevant literature, starting with an overview of the human visual system and the various neuroimaging approaches currently used to investigate visual responses, including single-unit recordings, local field potentials, EEG, fMRI and multi-modal approaches, such as EEG-fMRI. The overview then explores the literature on the variability of visual responses and summarises the behaviour of spontaneous activity at both the cellular level and in terms of whole-brain networks. Lastly, an overview of the concept of brain states is provided, as well as the various methods used to identify them, such as sliding window approaches and, more recently, Hidden Markov Modelling (HMM), which is a key analysis method in this thesis.

The Anatomy of the Human Visual System

The visual system is the most complex sensory system within humans, and as this thesis focuses on variability in visual responses, a brief overview of the visual system is given here, including overviews of the human eye, thalamus, primary visual cortex, and extrastriate cortex.

The Human Eye

The eye is the organ that allows human beings to perceive incoming visual stimuli. As light collides with objects within the environment, some of that light is reflected into the eye via

a small opening known as the pupil. Depending on the object's properties, e.g. its texture or colour, the wavelengths of the reflected light will differ, allowing us to distinguish between different objects. Once the light passes through the pupil, it travels to the back of the eye, colliding with a thin membrane called the retina. The retina is fundamental for vision, as it is where the information transmitted by light is converted into an electrical signal and passed on to the brain. The retina is composed of cells known as photoreceptors, which are dissociated into two groups, rods and cones, based on their structure and function. Rods are the most prevalent photoreceptor within the retina and are sensitive to luminance, making them primarily responsible for low-light vision. On the other hand, cones are less prevalent and are only sensitive to specific wavelengths of light, allowing them to be further divided into long, medium and short-wavelength cones. As cones are selective to different wavelengths of light, they are fundamental for colour perception. Rods and cones are not distributed evenly throughout the retina. Cones are often concentrated within the centre of the retina, known as the fovea, whilst rods are more prevalent within the periphery. Thus, cones are also responsible for foveal vision (e.g. the centre of the visual field), whilst rods are primarily responsible for peripheral vision. It is also key to note that the left and right eye do not solely correspond to the left and right visual fields. Instead, both eyes receive light from both visual fields, with the right eye receiving light from the left nasal visual field and vice versa with the right (see Figure 1.1).

The Thalamus

Once light is converted to electrical signals within the retina, these signals are propagated towards the brain via the optic nerve. Before reaching the brain, signals pass through the optic chiasm, which acts as a crossroads for incoming signals. Here, the signals from the nasal visual fields of each eye switch to the opposite side. This results in visual information from the right visual field being propagated and processed within the left hemisphere of the brain, and vice versa for the left visual field. Once past the optic chiasm, there are three main projections/path-

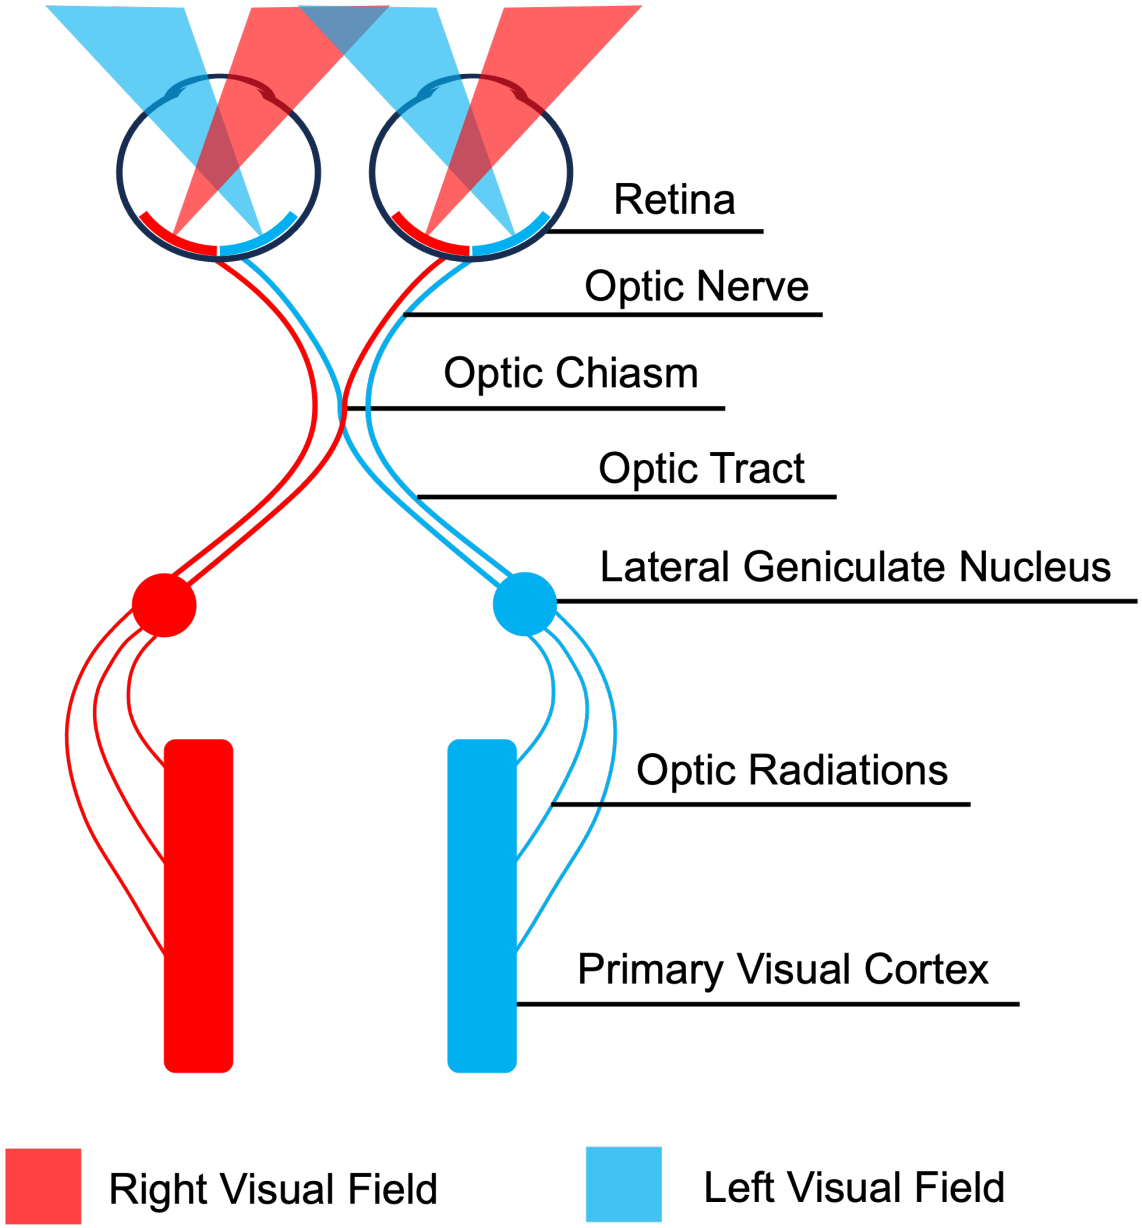


Figure 1.1: Diagram of the human visual system. Blue represents regions responsible for relaying and processing visual information from the right visual field, whilst red represents the left visual field.

ways into the brain: the retinohypothalamic tract, the tectopulvanic and the geniculostriate, of which the latter two pathways either terminate at or pass through the thalamus. The simplest of these pathways is the retinohypothalamic tract, which carries visual information to the suprachiasmatic nucleus of the hypothalamus, where it is used to entrain internal circadian

rhythms. The tectopulvinar pathway carries visual information to the pulvinar via the superior colliculus. This pathway has been associated with the detection of movement and is responsible for rapid saccades to moving stimuli. Lastly, the geniculostriate pathway relays information to the lateral geniculate (LGN) of the thalamus, which, via reciprocal interactions with the cerebral cortex, is responsible for the majority of visual functions, including form, colour perception and object recognition. The LGN is divided into a total of six layers, with cones projecting primarily to the magnocellular layers of the thalamus and rods projecting to the parvocellular layers. These layers then carry the visual information to the rest of the visual system, with most neurons projecting to the primary visual cortex via the optic radiations.

Primary Visual Cortex

Primary visual cortex (PVC) is located along the calcarine sulcus within the brain's occipital lobe and is the first cortical area reached along the geniculostriate pathway. For context, it is important to note that primary visual cortex often goes by multiple names, including but not limited to; primary visual cortex, V1 and striate cortex. A unique feature of PVC is that it is organised retinotopically. This means that each section of PVC responds to a specific portion of the visual field perceived by the retina. Specifically, the ventral regions of PVC respond to the visual field below the midline, whilst the more posterior regions of PVC respond to the more central areas of the visual field (and vice versa, with dorsal PVC responding to regions above the midline and anterior regions responding more to peripheral regions of the visual field). The area's primary function is to perform low-level visual processing, with neurons showing selectivity towards various spatial properties, including but not limited to: orientation (Hubel & Wiesel, 1959), spatial frequency (De Valois & De Valois, 1980) and direction of movement (reviewed in DeAngelis et al. (1995); Carandini et al. (1999)). Though the region primarily receives its input from the LGN, primary visual cortex also receives numerous secondary inputs from other subcortical regions, such as the pulvinar and claustrum. It is believed that these

secondary inputs help to modulate the excitability of neurons within primary visual cortex (Sherman & Guillery, 1998), facilitating processes like visual attention (Harris & Thiele, 2011). Once processed, the information is relayed to multiple regions. A small number of connections carry the visual information back to regions of the thalamus, including the LGN, thalamic reticular nucleus, and pulvinar. However, the majority of the information is relayed upstream to the rest of the visual system, to a regional hierarchy broadly known as extrastriate cortex.

Extrastriate Cortex, and the Ventral and Dorsal Pathways

Extrastriate cortex is primarily composed of four regions; V2, V3, V4 and V5/MT. Within these regions, a clear behavioural dichotomy is observed, unlike in the previous areas of the visual system. Specifically, extrastriate cortex can be split both behaviourally and anatomically into two separate pathways: the ventral and dorsal streams (Goodale & Milner, 1992; Mishkin et al., 1983). The ventral pathway is comprised of V2 and V4, travelling ventrally through the occipital lobe into occipital-temporal cortex. Its primary function is identifying visual information, facilitating colour perception, object recognition and face recognition. Conversely, the dorsal stream consists of V3 and V5, running dorsally through the visual cortex towards the parietal lobe. Unlike the ventral stream, the dorsal stream is primarily concerned with attention and the localisation of visual objects within space, which has led to the ventral stream being referred to as the 'what' pathway and the dorsal stream as the 'where' pathway (Goodale & Milner, 1992; Mishkin et al., 1983). Though there is a large amount of evidence for these streams from neuropsychology (e.g., lesion studies), it is noted that the stream view is likely an oversimplification. Various tractography studies in both primates (Baizer et al., 1993; Webster et al., 1994) and humans (Takemura et al., 2016) have demonstrated that regions within both streams are interconnected, indicating that the visual system is much more complex and is unlikely to operate as a pair of separate, serial pathways. This complexity and interconnectedness results in a large amount of variability in visual responses, even when the stimuli remain constant (Heggelund

& Albus, 1978; Rose, 1979; Schiller et al., 1976; Scobey & Gabor, 1989; Softky & Koch, 1993; Tolhurst et al., 1983; R. Vogels et al., 1989; Whitsel et al., 1977), due to the various top-down and bottom-up interactions both within visual cortex and with other areas of the brain.

Electrophysiology and Neuroimaging Techniques

Neuroimaging and electrophysiology are dedicated to the understanding of both the structure and function of the nervous system. This includes methods such as magnetic resonance imaging (MRI), electroencephalography (EEG), and Magnetoencephalography (MEG). Furthermore, these methods can all be applied to investigate the functionality of the visual system. Thus, this section briefly overviews the electrophysiological methods used in animals, followed by MRI, EEG and EEG-fMRI, as these are the methods used to record human visual responses throughout the thesis.

Animal Research

Though this thesis focuses on in vivo imaging within humans, much of the early literature surrounding neuronal response variability comes from electrophysiological research within animals (e.g. cats, primates, rats and mice). Therefore, a summary of the various in vitro techniques used within the animal literature to record neuronal activity is given here. All methods of in vivo recording are inherently invasive, requiring an electrode (or an array of electrodes) to be directly inserted within the area of the brain that researchers aim to record from. This makes it possible to directly measure electrical activity at a sub-millisecond temporal resolution and single-cell spatial resolution, far surpassing the resolutions of non-invasive methods, e.g., EEG and fMRI.

The two most common types of recordings acquired within animals are single-unit recordings and local field potentials (LFP). Which of these recordings is acquired is determined by the electrode

size and where it is placed relative to the neurons, with local field potentials being recorded from the space surrounding neurons (extracellular space) and single-unit recordings typically being acquired directly from a single neuron (intracellular space). Single-unit recordings directly measure the neuron's membrane potential, allowing action potentials to be recorded directly. Alternatively, LFPs represent the summation of electrical potentials across multiple neurons that surround the electrode. Specifically, they capture the average electrical activity generated by post-synaptic potentials, which are small changes in the neurons' membrane potential caused by the binding of neurotransmitters to the post-synaptic terminal. Typically, these recordings are acquired within behaving animals (typically anesthetised), allowing researchers to observe how neurons respond to sensory stimulation. e.g.whisker twitching in mice. Regarding visual research, many fundamental theories of vision are founded based on animal electrophysiological recordings. For example, Hubel and Wiesel (1959) used single-unit recordings within the visual cortex of cats and discovered that neurons within V1 show specific receptive fields, only firing in response to stimuli that occupied specific areas of the visual field. Lastly, though all these methods are described here in the context of animal literature, these methods can all be applied to humans in a clinical setting and have been successfully used to operate neural prosthetics. However, due to their invasive nature, they are not widely applied in human neuroscience, except in patients already undergoing neurosurgery for medical reasons, e.g. epilepsy.

Electroencephalography

Electroencephalography (EEG) is an electrophysiological method that utilises an array of electrodes to measure the electrical potentials at the scalp generated by neural activity within the brain. These scalp potentials are caused by the summation of post-synaptic potentials, making the recordings comparable to the LFP recording typically acquired in animal research. Unlike action potentials, which occur over approximately 1ms, post-synaptic potentials are much slower, lasting several 10s to 100s of milliseconds. Thus, when they occur in multiple neurons

simultaneously (i.e. in phase), the electrical potentials across those neurons are summated. Therefore, it is hypothesised that EEG signals are primarily generated by post-synaptic potentials within large populations of cortical pyramidal cells (e.g. an area of cortex spanning several centimetres). As the dendritic trees of these neurons align in parallel, when post-synaptic potentials occur in phase, they are summated, generating an electrical potential across the scalp (Murakami & Okada, 2006). To record EEG signals, electrodes are sewn into a cap in a specific spatial configuration known as a montage, which is placed on the participant's head. Caps have varying numbers of electrodes depending on the application, with typical caps consisting of 32 or 64 electrodes, though high-density EEG caps often have 126 or 256 electrodes. Conductive gel is then applied to each electrode, reducing the impedance between the electrode and the scalp. Electrical potentials are then recorded from each electrode relative to a common reference electrode, typically located at the vertex or the mastoids. Unlike fMRI, the temporal resolution is not limited by the sluggish BOLD response (discussed later in the chapter), meaning that EEG can have a much higher sampling rate, capturing changes in electrical potentials at a millisecond resolution, making it well-suited for investigating the temporal dynamics of brain responses.

However, as EEG can only record changes in voltage across the scalp, it is insensitive to deeper sources within the brain, e.g. electrical activity generated by subcortical regions, reducing its overall spatial resolution. This is due to volume conduction, as the electrodes do not directly measure from neurons. Thus, the electrical potentials have to travel through the skull and various other tissues to reach the scalp, distorting the signal and making it difficult to localise the source of the signal. Despite this, sources closer to the scalp (e.g. cortex) can be broadly estimated through source localisation methods, such as beamforming (see review Westner et al. (2022)). This process estimates which brain regions were active according to the observed pattern of electrical potentials across all electrodes. Though this process typically accounts for volume conduction by modelling the head of the participant (e.g. using a boundary estimation

model), the model cannot account for all of the distortions. Thus, localised sources often span several centimetres (though this can be improved using high-density EEG caps or MEG).

Event-Related Potentials

Event-related designs are the most commonly applied design for EEG research. These designs involve participants either being displayed a stimulus or completing a task and then averaging the EEG activity over many repeat trials. These average responses are known as event-related potentials (ERPs), which are primarily used to investigate the spatio-temporal electrical activity associated with processing the stimulus. ERPs are calculated by selecting a window (or epoch) around the onset of each stimulus and averaging the recorded EEG across epochs. ERPs typically consist of multiple components representing different stages of neural processing. For example, in vision research, flickering checkerboard patterns can be displayed to a participant to generate a visually evoked potential (VEP). The classical VEP consists of three main components: P100, N150, and P250 (see Figure 2.4). The letter of each component indicates the polarity, whilst the number represents the time that the component occurred in milliseconds.

Oscillations and Event-Related Synchronisation

As EEG recordings can have a millisecond time resolution, the method can also be used to investigate the role of neuronal oscillations. First discovered by Hans Berger, it was found that specific oscillations in the 7-13 Hz frequency range become more prevalent when participants close their eyes. This oscillation is now widely referred to as alpha oscillations, or the alpha rhythm, and it is one of many frequency bands observed within EEG data. These include: delta (0.5–3.5Hz), theta (4–7Hz) (Walter, 1936), alpha (7-13Hz), beta (13–30 Hz) (Berger, 1929), and gamma oscillations (30–45 Hz) (Jasper, 1938), each of which has been associated with different states and sensory and cognitive processes. Delta oscillations are the slowest oscillations observed in humans and are primarily related to non-rapid eye movement (NREM)

stage three sleep (N3 or slow wave sleep), as well as being involved in various higher-level mental tasks (Harmony, 2013). Like delta oscillations, theta oscillations are also involved with sleep, appearing during NREM stage one and during periods of drowsiness, reflecting sleep pressure (Aeschbach et al., 1997). Additionally, theta oscillations have been associated with various cognitive tasks, including navigation (Araújo et al., 2002; Kahana et al., 1999), memory (Doppelmayr et al., 2008; Jensen & Tesche, 2002; Klimesch et al., 1994; Rizzuto et al., 2006; Sauseng et al., 2010) and attention (Green & McDonald, 2008; Pennekamp et al., 1994; Sauseng et al., 2007).

Alpha oscillations are the most profound oscillations observed during awake rest in humans (Berger, 1929; Niedermeyer & Silva, 2005), and are primarily recorded over posterior electrodes. Though alpha was initially thought to represent cortical idling (Pfurtscheller et al., 1996) in the absence of stimuli, it is now commonly theorised that alpha instead represents functional inhibition (Jensen & Mazaheri, 2010). This process is proposed to facilitate neural processing by inhibiting task-irrelevant regions of the cortex. For example, when covertly attending to visual stimuli within the left visual field, alpha decreases within contralateral (right) visual cortex and increases within ipsilateral visual cortex (left) (Kelly et al., 2009; Rihs et al., 2007; Thut et al., 2006; Worden et al., 2000) suggesting that alpha facilitates attention by inhibiting information coming from the unattended areas of the visual field, allowing more resources to be allocated to task-relevant areas of cortex (Jensen & Mazaheri, 2010). Conversely, beta oscillations are observed across all cortical areas and are most profound over central electrodes. Beta oscillations typically decrease during motor activity but rebound shortly after the cessation of the movement (Jurkiewicz et al., 2006), in what is commonly known as the beta rebound. Lastly, gamma oscillations are most prevalent during tasks such as working memory (Jensen et al., 2007; Kaiser & Lutzenberger, 2005) and visual processing (Bastos et al., 2015; Jensen et al., 2015), and thus an increase in gamma power is typically associated with a decrease in alpha oscillations. However, treating different frequency bands as separate entities is likely

an oversimplification of neural processing, as research has shown that frequency bands often interact with each other. For example, theta and gamma oscillations have been demonstrated to exhibit cross-frequency coupling, in which the phase of the theta oscillations influences the amplitude of gamma oscillations (Lisman & Jensen, 2013), and thus frequency bands appear to work together to facilitate neural processing.

Neural oscillations are typically linked to behavioural and cognitive tasks based on induced changes in the power of oscillatory frequency bands in response to a stimulus. Similarly to ERPs, by averaging larger epochs around a stimulus presentation or task response, it is possible to observe the average changes in oscillatory power in specific frequency bands in response to that event. However, unlike ERPs, the changes in oscillatory power are not phase-locked to the stimulus. These changes are known as event-related synchronisation (ERS), in which the power within a specific frequency band increases in response to the stimuli, and event-related desynchronisation (ERD), in which the power decreases in response to the stimuli. In response to visual stimuli, changes in power are typically observed within the alpha and beta bands, with the ERD lasting for the duration of the stimulus, followed by an ERS in the same frequency bands shortly after the offset of the stimulus (Pfurtscheller & Lopes Da Silva, 1999; Pfurtscheller et al., 1994). Similarly, presenting visual stimuli at a specific frequency makes it possible to extract that frequency component from the evoked response. This is known as a steady state VEP, or SSVEP (Norcia et al., 2015). These frequency-based responses can be extracted even when stimuli are presented beyond the threshold of conscious perception (Seijdel et al., 2023), and this technique is often applied within brain-computer interface (BCI) research (Lesenfants et al., 2014; Siribunyaphat & Punsawad, 2022).

Magnetic Resonance Imaging

MRI is currently one of the most widely used neuroimaging methods for investigating human visual processing. Unlike single-unit recordings, MRI is non-invasive, making it possible to acquire measurements of the brain's structure and function without surgery. First developed in 1973, MRI revolutionised the medical imaging field by providing a more accurate and safer imaging technique than previously available methods, such as computerised tomography (CT), positron emission tomography (PET) and planar X-rays. Unlike previous imaging techniques, MRI can acquire images capable of identifying specific tissue, such as tumours, significantly improving diagnosis accuracy. Furthermore, MRI utilises a combination of magnets and non-ionising radiation to acquire images, allowing an individual to have multiple scans throughout their life without the long-term risks associated with ionising radiation imaging techniques (e.g.CT and PET). Together, these two factors have made MRI and functional MRI (fMRI) the primary methods for investigating human brain structure (MRI) and function (fMRI).

MRI Physics

MRI works on the principle of nuclear magnetic resonance (NMR), a phenomenon in which a combination of magnetic fields and oscillating magnetic fields can be used to obtain information about the composition and structure of a molecule. For NMR to occur, the molecule has to contain nuclei that possess non-zero spin (i.e. the nuclei must have an angular momentum), allowing the nucleus to generate a magnetic field and thus interact with external magnetic fields. Such nuclei with non-zero spin include hydrogen (1H), Carbon 13, and Phosphorus 31, as they all have an odd atomic mass. Whilst all of these nuclei are present within the human body and thus could theoretically be used for MRI, hydrogen is by far the most widely used for two reasons. Firstly, hydrogen is abundant within the human body in the form of water, making up approximately 60% of the overall molecules. Secondly, though other nuclei can be

used for NMR, hydrogen provides the strongest NMR signal due to its high gyromagnetic ratio, meaning that it is more strongly affected by magnetic fields than other elements.

As hydrogen nuclei possess spin, they generate a tiny magnetic field and can therefore be considered bar magnets (dipole) with positive and negative poles. This means that each hydrogen nucleus has a non-zero magnetic dipole moment, which can be considered as an arrow (vector), defining the direction and strength of the nuclei's magnetic field. As measuring the magnetic field of a single hydrogen nucleus is not feasible, it is easier to think about the net magnetisation of an area, as opposed to the magnetic fields of individual nuclei. The net magnetisation is calculated by summing all the nuclei's magnetic dipole moments, providing a single summary measure of the region's magnetic moment. Without the presence of an external magnetic field, the magnetic dipole moments of the hydrogen nuclei are oriented at random. Thus, the net magnetisation is close to zero (see Figure 1.2). However, when the body is placed within the MRI's strong (typically 1.5-7T) homogenous magnetic field (B0), the magnetic dipoles undergo several changes, which can be described from both a classical or quantum physics perspective. From a classical physics perspective, the magnetic dipole moments of a proportion of the nuclei begin to align with the direction of B0, with the majority aligning in parallel and a minority of atoms aligning anti-parallel to the direction of B0, resulting in a small net magnetisation in the direction of B0. These can be thought of as low and high-energy states, respectively, due to the larger amount of energy required to align anti-parallel to B0. Alternatively, from a quantum physics perspective, the magnetic dipole moments of each nucleus begin to process around the axis of B0 at a specific frequency. This frequency is known as the Lamour frequency, and it is dependent on the nucleus itself and the strength of the B0 field. Furthermore, in order for MRI to obtain three-dimensional images, an additional set of gradient magnets (Gx, Gy and Gz) is required. These gradient fields superimpose with the B0 field to introduce a gradient in the net magnetic field along the x, y, and z axes, allowing for the signals from different spatial locations in the tissue to be sampled as different areas of the image, and therefore can be differentiated

from each other.

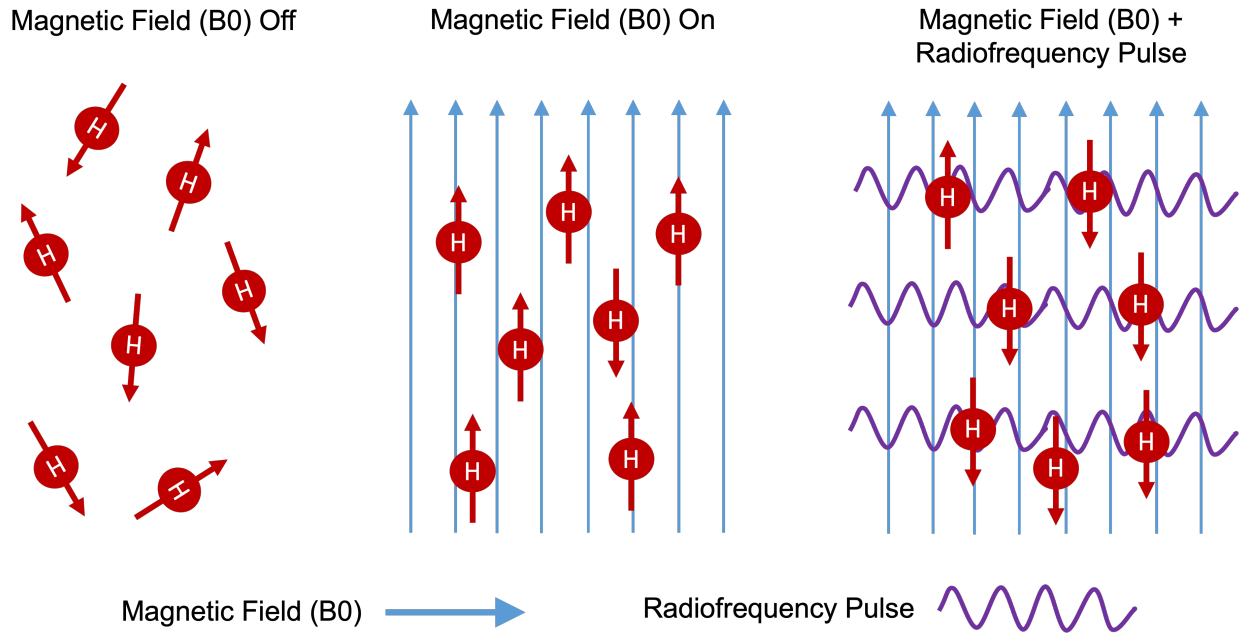


Figure 1.2: Behaviour of hydrogen nuclei in the presence of magnetic field (B0) and radio frequency pulses. The leftmost panel shows hydrogen nuclei without the presence of an external magnetic field. The center panel shows the behaviour of hydrogen nuclei in the presence of B0. The rightmost panel shows the behaviour of hydrogen nuclei in the presence of B0 and radio frequency pulses.

Once the nuclei are aligned with and precessing around B0, an oscillating magnetic field (B1), often known as a radio-frequency (RF) pulse, can be applied perpendicularly to the direction of B0, exciting the hydrogen nuclei. This is achieved by applying an RF pulse at the Lamour frequency of the element, causing the nucleus to resonate and gain energy. This increase in energy perturbs the magnetic dipole moments of the nuclei, forcing them to align away from the direction of B0, subsequently resulting in a net magnetisation away from the direction of B0. The difference between the original orientation of the net magnetisation and the orientation of the net magnetisation after excitation is known as the flip angle and is determined by the strength and duration of the RF pulse. For example, a longer RF pulse will allow more nuclei to become excited, resulting in a larger flip angle of the net magnetisation, e.g.$>$ 90 degrees. Once the RF pulse is stopped, the majority of nuclei return to a low-energy state, aligning parallel

to the magnetic field and releasing energy in the form of electromagnetic waves. This release of energy is known as relaxation and is recorded by receiver coils situated around the body, resulting in the final NMR signal.

There are two main types of relaxation: T1 and T2. T1 relaxation refers to the return of the net magnetisation orientation to be parallel with B0 along the longitudinal plane (e.g. axis of B0), whilst T2 refers to the return along the transverse plane (e.g. plane orthogonal to B0). This loss of net magnetisation within the transverse plane occurs due to the nuclei precessing out of phase with each other. Whilst the radio-frequency pulse is present, all of the hydrogen nuclei precess in phase, resulting in a net magnetisation in the transverse plane as well as the longitudinal plane. However, once the pulse is stopped, the nuclei become out-of-phase exponentially, and the net magnetisation within the transverse plane reduces. Conversely, the T1 NMR signal increases exponentially, taking longer for all the nuclei to realign parallel to B0. Depending on the tissue that is excited, the T1 and T2 relaxation times will vary based on a number of factors. The simplest factor is the number of hydrogen nuclei present in the tissue, but other more complex factors, such as the presence of local magnetic field inhomogeneities (e.g. that occur near air-tissue boundaries), can also impact the relaxation times. It is this difference in relaxation times between tissues that is exploited to gain contrast in images between different tissue types. Thus, in MRI, there are two main categories of images: T1-weighted and T2-weighted images, where each maximises the difference in signal intensity according to its respective relaxation type. These differences are maximised by modulating the timings of the radio-frequency pulse and the subsequent release of energy, also known as repetition time (TR) and time to echo (TE). For T1-weighted images, the TR is short, maximising the amount of T1 signal between tissues, whilst in T2-weighted images, the TE is long, maximising the difference in the amount of T2 signal. Within neuroimaging, fat-rich regions (e.g. meninges) show up brightest within T1s, whilst water-rich regions appear dark (e.g. CSF). Conversely, in T2w images, both water-rich and fat-rich regions appear bright. Furthermore, whilst T2w can be used to generate structural

images, it is also the main contrast type used within functional imaging.

Neurovascular Coupling, the BOLD Signal and the Heamodynamic Response

When neural activity occurs within an area of the brain, for example, in response to visual stimulation, several physiological changes occur. Namely, the increase in neural activity increases the metabolism of oxygen and glucose, inducing vasodilation with the local vasculature, which in turn results in an increase in cerebral blood volume (CBV), cerebral blood flow (CBF), and blood oxygenation levels. This process is known as neurovascular coupling, and it is what enables fMRI research.

These metabolic changes can be recorded due to the specific magnetic properties of deoxyhemoglobin. Specifically, whilst oxyhaemoglobin shows little interaction with magnetic fields, deoxyhemoglobin is strongly paramagnetic, which creates local magnetic inhomogeneities in and around blood vessels. This causes T2 relaxation times to decrease as the concentration of deoxyhemoglobin increases. Consequently, the NMR signal intensity is stronger in the presence of oxyhaemoglobin and weaker in the presence of deoxyhemoglobin. These changes can be recorded by fMRI and are now widely referred to as the Blood Oxygenation Level-Dependent (BOLD) signal. The BOLD signal is most commonly recorded using T2* weighted images, which are even more sensitive to local field inhomogeneities than T2 weighted imaging, making them more sensitive to changes in deoxyhemoglobin. However, this does result in artefacts in areas of the brain where there are naturally larger field inhomogeneities, such as the orbitofrontal cortex, due to the large local bone structures.

When a brain region is active, a specific sequence of BOLD signal changes is observed, known as the Haemodynamic or BOLD response. In response to a short (1s) stimulus impulse, a brief reduction in the BOLD signal is sometimes observed, known as the initial dip, caused by the initial consumption of oxygen, resulting in an increase in deoxyhemoglobin (Hu & Yacoub,

2012; Kim et al., 2000; Tian et al., 2010; Yeşilyurt et al., 2008). This is followed by a large increase in the BOLD signal as fresh blood flows to the capillary bed, increasing the local oxyhemoglobin concentration and causing a peak in the BOLD signal around 5-7 seconds after the initial neural activity. This increase in blood flow is a compensatory mechanism known as functional hyperemia, and it is the main contributing factor towards the BOLD signal within a region, as it induces a large change in the ratio between oxyhemoglobin and deoxyhemoglobin. Lastly, this peak in the BOLD signal is followed by an undershoot at around 10-15 seconds, caused by an increase in deoxyhemoglobin (Van Zijl et al., 2012), before returning to baseline (see Figure 2.5 for example). This canonical BOLD response is taken advantage of within fMRI analysis, allowing researchers to identify which brain regions are involved in processing a particular stimulus based on which regions exhibit a BOLD response. This is typically achieved by a general linear model (GLM) approach, in which an area is deemed active if its evoked BOLD signal matches the pattern of the canonical BOLD response. Due to MRIs high spatial resolution, the results of this analysis provided a map of significant BOLD activity (activation map) across the whole brain, including within smaller subcortical structures, such as the LGN. Furthermore, it is also possible to directly compare the BOLD responses to different stimuli within a specific region, for example, comparing the magnitude of the BOLD response peak in V1 to high-contrast and low-contrast checkerboards. However, though faster sequences can be used to improve the temporal resolution of fMRI data, such as multiband sequences, which acquire multiple slices with each radio-frequency pulse, the temporal resolution remains limited due to the sluggish nature of the haemodynamic response.

Silmutaneous EEG-fMRI

One approach to circumvent the limitations of EEG and fMRI is to record both simultaneously. Combining the two methods makes it possible to measure brain activity with both high spatial resolution (millimetres) and high temporal resolution (milliseconds). This approach has been

applied in alpha research, identifying BOLD correlates of the alpha rhythm (Goldman et al., 2002; Laufs et al., 2006; Mayhew & Bagshaw, 2017; Mayhew, Ostwald, et al., 2013; Moosmann et al., 2003), as well as in clinical epilepsy research, in which epileptic spiking can be localised with millimetre resolution (Bénar et al., 2006; Gotman et al., 2006). As these measures are recorded separately, they need to be integrated during analysis to capitalise on the strength of both methods. There are three main approaches to data integration for EEG-fMRI: asymmetric data integration, neurogenerative modelling and multimodal data fusion (Huster et al., 2012).

Asymmetric Data Integration

Asymmetric data integration is typically biased to one of the two measurements, with one providing information to assist the analysis of the other. This results in two types of analysis: fMRI-informed EEG, which focuses on improving the source localisation of the EEG data, and EEG-informed fMRI, which focuses on enhancing fMRI GLMs by modelling variations in the EEG responses, e.g. ERP amplitude. fMRI-informed EEG works by using the high spatial resolution of fMRI to guide or constrain the source localisation of the EEG. In the case of dipole fitting, fMRI activation maps can be used to determine the number of EEG generators (e.g. the number of regions activated). Thus, individual dipoles can be fitted to each generator, providing time courses of the neural activity within each generator (Huster et al., 2012). Alternatively, when using a distributed source model to localise EEG activity (Hämäläinen & Ilmoniemi, 1994), the fMRI maps can be used to constrain the source space, with sources being more likely to originate from regions of the brain that coincide with the fMRI activation maps (Dale et al., 2000; Wanmei Ou et al., 2009).

On the other hand, EEG-informed fMRI involves utilising the temporal information of the EEG to improve the fMRI analysis (Basha, 2012). An example of this would be adapting the fMRI GLM weights according to single-trial measures of EEG responses, allowing the BOLD responses to vary according to the neural activity instead of assuming a static canonical BOLD

response. Such EEG measures include but are not limited to ERP amplitudes (Bagshaw & Warbrick, 2007; Debener et al., 2005), ERP latencies (Bénar et al., 2007), and frequency band power (Scheeringa et al., 2009). When designing the fMRI GLM design matrix, these measures are used to modulate the expected hemodynamic response, e.g. if a trial has a larger ERP amplitude, the estimated hemodynamic response for that trial is increased. A similar approach can be used for resting state designs, in which EEG measures can be used as events within the fMRI GLM analysis. This approach has been widely applied to the alpha rhythm, in which the amount of alpha power at each fMRI volume is calculated and used as a regressor within the fMRI GLM analysis to identify the BOLD signal correlates of variations in alpha power (Goldman et al., 2002; Laufs et al., 2006; Mayhew & Bagshaw, 2017; Mayhew, Ostwald, et al., 2013; Moosmann et al., 2003).

Neurogenerative Modelling

Unlike asymmetric data integration approaches, neurogenerative models aim to benefit from the information of both the EEG and fMRI data without biasing the analysis towards either modality (Riera & Sumiyoshi, 2010; Rosa et al., 2010; Valdes-Sosa et al., 2009). Neurogenerative modelling approaches aim to model how EEG signals relate to the subsequent BOLD responses and the neurovascular coupling mechanisms that facilitate the interaction. This is known as the generative model, and once fitted to recorded or simulated EEG-fMRI data, it is possible to estimate the inverse model, providing estimations of neural activity, specifically the excitatory and inhibitory postsynaptic potentials. Whilst this approach has been successfully applied to simulated EEG-fMRI data (Bojak et al., 2011), the approach requires multiple assumptions to be made about the mechanisms of neurovascular coupling and can only be applied to a handful of cortical regions due to the computational complexity of a whole-brain neurogenerative model. Therefore, this approach is currently best applied as a framework for comparing neurovascular coupling models and the brain's physiological properties (Huster et al., 2012; Rosa et al., 2010).

Multimodal Data Fusion

An alternative to the model-based approach is to use a data-driven approach to perform a joint decomposition of the data. This has most widely been achieved by using joint independent components analysis (ICA), in which both the EEG and fMRI data are inputted together, resulting in independent components that reflect both the EEG and BOLD signals. (Calhoun et al., 2006; Calhoun, Liu, et al., 2009; Mijović et al., 2012; Moosmann et al., 2008). Specifically, mean fMRI activation maps and ERPs are generated for each subject, and concatenated into a single matrix and inputted into the ICA. The approach assumes that the activity recorded by the two modalities is generated by the same group of neurons (Mijović et al., 2012). Thus, the amplitudes of the ERP components and the hemodynamic responses should exhibit a strong positive covariance. Overall, this approach results in a fusion of the two data modalities, retaining the high spatial and temporal resolution, with independent components reflecting the spatial properties of ERP components (e.g. P100) at a millimetre scale.

Methodological Considerations

Given the plethora of options for integrating EEG-fMRI data, it would initially appear that the method would be routinely applied within cognitive neuroscience research due to the improved temporal and spatial resolution. However, the approach has multiple methodological challenges that make it difficult to be employed within all research institutions and restrict analysis options (Bullock et al., 2021; Warbrick, 2022). When recording EEG and fMRI simultaneously, both methods have an impact on the recording quality of the other; for example, the metallic EEG electrodes cause minor inhomogeneities within B0, causing distortions in the structural MRI images (Bonmassar et al., 2001; Krakow et al., 2000; K. Mullinger et al., 2008). However, the larger issue comes from the effects of the MRI on the quality of the EEG data. The strong static magnetic fields in combination with the radio-frequency pulses generate several artefacts within the EEG data, compromising the quality of the recordings. The largest of these artefacts

is the gradient artefact (P. J. Allen et al., 2000; W. X. Yan, Mullinger, Brookes, et al., 2009). During scanning, the gradient coils rapidly alter the magnetic field to facilitate spatial encoding of the MRI images, subsequently inducing current within the electrodes. This results in large but consistent spiking artefacts that are many orders of magnitude larger than the physiological EEG signals (P. J. Allen et al., 2000; Bullock et al., 2021).

In addition to the gradients, the static magnetic field also generates artefacts within the EEG data due to the induction of currents in the electrodes from participant movement. This occurs even when participants make small, sub-millimetre head movements, and thus, padding is often placed around the head to limit this movement. However, the larger issue comes from non-voluntary movement in the form of the ballistocardiogram artefact, which occurs every cardiac cycle (P. J. Allen et al., 1998; W. X. Yan, Mullinger, Geirsdottir, et al., 2009). Due to the strength of the static magnetic field, even tiny movement can induce a current within the electrodes, and thus, changes in pulsatile blood flow also induce a current. Though this artefact is less consistent than the gradient artefact due to the dynamic nature of heart rate, the amplitude of the artefact is smaller, peaking at approximately 200 μV (Debener et al., 2007). Combined with the gradient artefact, this limits the useable frequency bands to alpha and low beta, with the residual gradient artefacts corrupting the higher frequencies and residual BCG artefacts corrupting the lower frequencies.

Additionally, various environmental factors impact the quality of the EEG recordings, including vibrations from the scanner and electrical noise from the helium pumps. From a safety perspective, there are also heating risks associated with EEG-fMRI. Though the electrodes used for EEG-fMRI are made out of a non-ferromagnetic alloy (typically Ag/AgCl), they remain susceptible to heating from the radio-frequency pulses, and thus images that use higher power RF pulses, such as high multi-band sequences or T2 structural images, cannot be acquired. Fortunately, there are various ways in which the experimental set-up can be optimised to reduce

the risk of heating and the impact of these artefacts, in addition to a wide selection of algorithms that aim to reduce the artefacts during analysis (Bullock et al., 2021; K. J. Mullinger, Castellone, et al., 2013); both of which are discussed within Chapter 5.

Variability of Visual Responses

Neural responses to sensory stimuli are highly variable across different presentations of the same stimuli. This phenomenon was first observed within animal research, in which the spike train recorded from sensory neurons showed large amounts of variability between repeated presentations of the same stimulus (Heggelund & Albus, 1978; Rose, 1979; Schiller et al., 1976; Scobey & Gabor, 1989; Softky & Koch, 1993; Tolhurst et al., 1983; R. Vogels et al., 1989; Whitsel et al., 1977). Initially it was hypothesised that the observed variability in spike trains frequency was a result of cellular mechanisms intrinsic to the neuron being recorded. For example, the opening and closing of ion gates is a stochastic process, contributing towards the observed neuronal noise (O'Donnell & Van Rossum, 2014). However, when a neuron is isolated from its connections to its neighbouring neurons and stimulated with a fluctuating current, it generates spike trains that are essentially deterministic, exhibiting minimal amounts of variability (Mainen & Sejnowski, 1995). This demonstrated that though intrinsic noise from the neuron itself contributes to the variability of the evoked spike trains, most of the variability occurs when the neuron is connected within a network of other neurons. This claim has been further supported through computational modelling, in which simulated networks of deterministic neurons generate variable spike trains to the same stimulation (Brunel, 2000; Mayor & Gerstner, 2005; T. P. Vogels & Abbott, 2005). This further demonstrates that most of the variability in spike trains results from network-related noise as opposed to intrinsic noise.

Though most of the literature on the variability of visual responses comes from animal experiments, similar patterns of variability are observed in human neuroimaging and electrophysiology

studies. For example, both the BOLD responses (Bianciardi et al., 2009; Duann et al., 2002; M. Schölvinck et al., 2012) and the VEPs recorded with EEG (Becker et al., 2008; Mayhew, Ostwald, et al., 2013) show a large amount of variability in response to the same stimuli. However, both measures are inherently more noisy compared to single-cell recordings due to physiological and methodological issues. For example, EEG is subject to muscle and blink artefacts, and fMRI is sensitive to movement, as well as the participants' breathing and heart rate (see Review Chen et al. (2019)). Whilst preprocessing steps are employed to mitigate these artefacts where possible, most research on sensory responses relies on averaging. In a typical study, participants will be presented with the same stimuli numerous times whilst brain responses are recorded (e.g. via EEG or fMRI) and then averaged during the analysis to negate the impact of the noise. This approach assumes that most of the variability observed between trials is noise and, therefore, not of interest. However, given the results from the animal literature, the activity of other neurons within the network impacts subsequent spike trains, and it is unclear whether this represents noise or if the induced changes are functionally informative and serve as a mechanism in neural processing.

Spontaneous Activity

A large amount of activity within neurons occurs in absence of any external stimuli. This is referred to as spontaneous activity, and it occurs inherently within all networks of neurons. As with the variability in evoked responses, the same computational modelling studies have shown that simulated networks of deterministic neurons exhibit similar patterns of spontaneous activity as observed in experimental data, indicating that the spontaneous activity is a result of the network itself (Brunel, 2000; Mayor & Gerstner, 2005; T. P. Vogels & Abbott, 2005). Furthermore, in experimental data, the spontaneous spiking within a single neuron is typically highly correlated with the spontaneous spiking of neurons within the same network (Arieli et

al., 1995; Arieli et al., 1996; Tsodyks et al., 1999). This has been demonstrated in the animal literature, in which optical imaging alongside single-unit recordings in the visual cortex of cats showed that the spiking pattern of a single neuron was significantly correlated with the optical signals and LFP of the surrounding neurons, even in the absence of stimuli (Arieli et al., 1995; Arieli et al., 1996; Nelson et al., 1992; Tsodyks et al., 1999). Additionally, synchronisation of spontaneous spiking is stronger between neurons with similar functions, with single-unit recordings in the visual cortex of cats showing stronger correlations in spike rates between neurons with overlapping receptive fields or similar orientation selectivity (Nelson et al., 1992). Therefore, neurons show strong correlations in their spiking rate to physically connected neurons and neurons with similar functionality that may not share direct synapses.

In addition to synchrony in spiking rate, neurons within the same network also show strong correlations in fluctuations of their sub-threshold membrane potential, indicating synchrony in post-synaptic potentials and spiking. This has been shown experimentally in cats, in which the sub-threshold membrane potentials of pairs of neurons within visual cortex are significantly correlated, thus showing synchronisation in their post-synaptic potentials (Lampl et al., 1999). Interestingly, as this synchronisation in post-synaptic potentials can occur over large areas of the cortex, the fluctuations are large enough to be recorded without electrode insertion via EEG or MEG. At this scale, this spontaneous activity is primarily observed in the form of neural oscillations resulting from the synchronised post-synaptic potentials over a large area of the cortex. For example, as discussed previously, an increase in alpha power is observed upon eye closure despite the absence of external stimuli. Furthermore, with M/EEG it is possible to observe synchronisation in frequency bands over larger distances, e.g. between visual cortex and frontal cortex. This can be observed in awake humans, e.g. increases in alpha as attention lapses, and during sleep, in which the whole brain exhibits significant changes in the power of various frequency bands as it cycles through the different sleep stages (Boostani et al., 2017). Additionally, differences in these spontaneous oscillations have been observed in clinical sub-

populations, including schizophrenia (Koshiyama et al., 2020) and altered states of consciousness (Brenner, 1985).

Unfortunately, it is not possible to record these spontaneous neural oscillations using fMRI, however, various BOLD networks spanning across both cortical and subcortical regions have been identified. During rest, spontaneous fluctuations in the BOLD signal are observed across the brain, with many regions' fluctuations significantly correlating to each other over time. This is commonly referred to as functional connectivity and was first observed in fMRI within motor cortex, where the BOLD signal fluctuations within right motor cortex were found to be significantly correlated with the fluctuations within left motor cortex (Biswal et al., 1995). Since this discovery, many other patterns of functional connectivity have been discovered across the cortex (Fox & Raichle, 2007; Thomas Yeo et al., 2011), as well as within subcortical regions (Barnes, 2010; Thomas Yeo et al., 2011). Regions with correlated activity are typically referred to as being part of resting state / intrinsic networks (even though they are seen both during rest and performance of tasks), each with unique spatial topologies across the brain. Furthermore, these networks have been associated with various elements of cognition; for example, the default mode network (DMN) (Raichle et al., 2001) has been widely associated with self-referential tasks, such as mind wandering (Christoff et al., 2009; Kucyi et al., 2016). Additionally, despite the name, resting state networks are observed during tasks (S. M. Smith et al., 2009), and show unique temporal dynamics, with some networks showing higher levels of BOLD signal during tasks, such as the fronto-parietal networks, whilst others show a decrease in BOLD signal, such as the DMN (Fox et al., 2005). Overall, this highlights that spontaneous activity is non-random, occurs across multiple scales, from small, microscale networks of neurons to long-range macroscale networks across the cortex, and is implicated with various aspects of cognition.

Interactions between Spontaneous Activity and Visual Responses

The brain's spontaneous activity has also been demonstrated to interact with incoming visual stimuli, explaining a part of the variability in visual responses discussed earlier. By recording the neural spiking and local field potentials simultaneously within the visual cortex of cats, researchers found that the amplitude of the membrane potential signal immediately before the stimulus linearly influenced the spiking rate of the subsequent visual response (Arieli et al., 1996; Azouz & Gray, 1999). In other words, when the membrane potential was high prior to the onset of the visual stimuli, the spiking rate response was higher, and vice versa. This effect is known as linear summation, in which the visual responses are essentially added to the baseline activity. Thus, the results of these studies demonstrated that a large amount of the variance in the visual responses could be explained by the amplitude of the baseline activity, in this case, the membrane potential. This indicates that neurons can be thought of as linear dynamic systems, in which the current state of the neuron will affect the magnitude of its outputs to stimuli.

A similar effect can also be observed within humans using M/EEG, with the phase and amplitude of scalp oscillations prior to stimulus onset significantly impacting both the magnitude of the visual and behavioural responses. For example, when alpha power was high prior to stimulus onset, the amplitude of later VEP components was increased (Becker et al., 2008; Reinacher et al., 2009). Furthermore, when observing near-threshold visual stimuli, higher levels of baseline alpha power was associated with reduced chance of the stimuli being detected (Hanslmayr et al., 2007; Kelly et al., 2009; Mathewson et al., 2009; Mazaheri et al., 2009; Schubert et al., 2009; Thut et al., 2006; Van Dijk et al., 2008). Pre-stimulus alpha has also been shown to modulate the magnitude of BOLD responses in EEG-fMRI research, with higher pre-stimulus

alpha power associated with reduced magnitude of visual cortex BOLD responses (Becker et al., 2011; Mayhew, Ostwald, et al., 2013). Additionally, EEG studies have shown that alpha power is reduced pre-stimulus when the stimulus is expected (Bastiaansen et al., 2001; Bastiaansen & Brunia, 2001), indicating that alpha plays a role in visual attention.

Similar effects have been observed for fMRI, with the spontaneous BOLD signal within visual cortex linearly modulating BOLD responses to visual stimuli and influencing how stimuli are perceived. As observed from single-unit recordings (Arieli et al., 1996; Azouz & Gray, 1999), fMRI research has shown that visually evoked BOLD responses linearly summate with the spontaneous activity in visual cortex (Bianciardi et al., 2009; Duann et al., 2002; M. Schölvinck et al., 2012). Despite this, spontaneous activity has shown a close link with behaviour, with spontaneous activity changing with the participant's behavioural state (Bianciardi et al., 2009), as well as influencing perception (**hesselmann_spontaneous_2008.hesselmann_ongoing_2008**; Hesselmann et al., 2010; Pajani et al., 2015). Specifically, fMRI research has shown that spontaneous activity varies significantly between eyes closed and eyes open rest, as well as between rest with fixation and without fixation (Bianciardi et al., 2009). Furthermore, the amplitude of the spontaneous activity within visual cortex has been shown to influence the perception of near-threshold stimuli (Hesselmann et al., 2010; Pajani et al., 2015), visual illusions (Hesselmann, Kell, Eger, et al., 2008) and coherent motion (Hesselmann, Kell, & Kleinschmidt, 2008), suggesting that spontaneous activity plays an important role in visual processing and the anticipation of visual stimuli.

Defining Brain State

Whilst this research highlights the impact of spontaneous activity on visual responses, the definition of spontaneous activity is often oversimplified and not always consistent. In animal literature, spontaneous activity is typically only recorded from a small population of neurons

within visual cortex (Arieli et al., 1996; Azouz & Gray, 1999), and thus does not evaluate the influence of spontaneous activity within other regions of the brain on visual responses. This is primarily due to the methodological constraints of invasive recordings, making it challenging to obtain whole-brain coverage. Whilst EEG provides a more global measure of spontaneous activity that can be recorded from across the majority of the head, its spatial resolution is still poor compared to fMRI. Due to the distance from the scalp, electrical activity from subcortical regions cannot be recorded using EEG. Furthermore, EEG primarily records from the superficial surface of gyri and is less sensitive to electrical activity within the sulci. Though fMRI can circumvent the issue of limited spatial resolution, it is ultimately a proxy measure of neural activity and is not sensitive to changes in the underlying neural oscillations (Swettenham et al., 2013). Thus, whole-brain recordings are required to fully explore how spontaneous activity interacts with visual responses.

One way to go about this is to think of spontaneous activity (and activity in general) as recurring patterns of activity, or brain states, which can be defined according to some characteristic spatio-temporal pattern of activity (e.g. BOLD and/or oscillatory) across the entirety of the brain. However, such a concept can be difficult to define, given the multiple ways spontaneous activity can be recorded (Brown, 2006; A. S. Greene et al., 2023). Furthermore, the term brain state is often used interchangeably to reflect a behavioural state, e.g. sleep and levels of consciousness. To distinguish between these, (Brown, 2006) makes the distinction between the state of the brain and brain state, with the former representing the chemistry of the brain and the latter representing the electrical activity across brain states, with the caveat that brain states occur against the background state of the brain. For example, an 'alpha' brain state may be observed both while awake with eyes closed and during sleep but is dissociated by the chemical state of the brain. Brown (2006) suggests that the ideal definition of brain state would have high spatial resolution and high temporal resolution, defining state by reoccurring patterns in the synchrony in frequencies across the entirety of the brain. In brief, a brain state can be defined

as reoccurring patterns of spontaneous activity seen across the whole brain, with activity being defined based on the modality, e.g. BOLD signal for fMRI, frequency power for EEG, and provides an alternative framework for quantifying the dynamics within the brain.

Identifying a Brain State

These reoccurring patterns of spontaneous activity have been observed experimentally, primarily within fMRI functional connectivity research. In traditional studies, functional connectivity is defined as the temporal correlation in activity (typically the BOLD signal) between a pair of regions throughout a scan, typically lasting 5-10 minutes. However, as this approach inherently assumes that functional connectivity does not change throughout a scan, more recent research has moved towards a dynamic functional connectivity approach, in which the changes in functional connectivity are observed over smaller increments of time. This is typically implemented using a sliding window approach (Ahrends & Vidaurre, 2023; Hindriks et al., 2016; Mokhtari et al., 2019), in which the recording is broken down into windows (e.g. 30s periods), offset from each other (sometimes overlapping), and the functional connectivity is calculated separately for each window. Brain states are then identified by clustering the windows, resulting in a set of reoccurring functional connectivity states (E. A. Allen et al., 2014). This approach has been applied to resting-state fMRI data, identifying unique functional connectivity patterns that are not observed when using conventional functional connectivity analyses (E. A. Allen et al., 2014). Additionally, this analysis has been applied to EEG-fMRI resting state data, in which the spectral properties from the EEG data were used to further characterise the functional connectivity patterns, resulting in states with both high spatial, temporal and spectral resolution (E. A. Allen et al., 2018). Furthermore, by clustering the patterns of activity into brain states, various state metrics can be calculated, summarising the temporal dynamics of each state. For example, the most commonly used state metric is fractional occupancy, which is the percentage of time a

specific state was active during the recording (see Chapter 2 for a complete list of state metrics). These state metrics have been applied in clinical resting-state fMRI data, in which researchers could distinguish between Schizophrenic participants and control participants based on differences in the fractional occupancy measures (Damaraju et al., 2014). Overall, this demonstrates that the spontaneous activity in the brain can be decomposed into reoccurring patterns, leading some researchers to call for a complete characterisation of all possible connectivity brain states in the form of a 'chronnectome' (Calhoun et al., 2014), or 'dynome' for frequency brain states (Kopell et al., 2014).

However, despite the successful application of the sliding window approach for the identification of brain states, its accuracy has been criticised (Hindriks et al., 2016; Laumann et al., 2016; Leonardi & Van De Ville, 2015; Shakil et al., 2016; Zalesky & Breakspear, 2015). The majority of this criticism is targeted at the window size and offset parameters, which are often set arbitrarily, and their value can dramatically affect the results. For example, using a smaller window size will result in a finer temporal resolution; however, the estimations within the window will be subject to more noise due to the small number of timepoints, and vice versa for larger window sizes. This issue is even more problematic for the identification of brain states, with simulations demonstrating that the approach performs poorly when detecting ground truth states, due to the window size often being non-comparable to the duration of the state, resulting in windows often containing a mix of states (Shakil et al., 2016). Furthermore, much of the literature has focused on identifying states according to functional connectivity (E. A. Allen et al., 2018; E. A. Allen et al., 2014; Damaraju et al., 2014), when in practice states can be defined according to the mean activity within regions, as well as spectral information (e.g. high alpha state).

Hidden Markov Modelling

An alternative approach that amends these issues and avoids the limitations of the sliding window-based method is to model the state dynamics using Hidden Markov Models (HMM)

(Baker et al., 2014; Vidaurre et al., 2016). Similar to the sliding window/clustering analysis, the HMM approach also aims to identify recurring brain activity patterns. However, unlike the sliding window-based approach, the HMM approach is entirely data-driven, decomposing the data into reoccurring states without setting a window size or offset. Furthermore, the HMM approach defines states according to an observation model, and this can be changed according to the modality of the data, allowing the approach to identify fMRI network states, as well as oscillatory states in EEG data (see Chapter 2 for details on HMM implementation and observation models). In fMRI research, the approach has been applied to resting-state data (Karapanagiotidis et al., 2020; Vidaurre et al., 2021; Vidaurre, Abeysuriya, et al., 2018; Vidaurre et al., 2017). This approach to modelling the temporal dynamics of resting-state networks has highlighted the hierarchical nature of state transitions, with sensory states being more likely to transition to other sensory states, and higher-order states being more likely to transition to higher-order states (Vidaurre et al., 2017). Furthermore, the states identified within fMRI data have been associated with various psychological processes (Karapanagiotidis et al., 2020; Vidaurre et al., 2021), indicating that they have a functional role.

The approach has also been successfully applied to M/EEG data, identifying states according to their spectral properties (e.g. power and phase). In source localised MEG resting data, HMM was able to identify states based on changes in broadband (4-30Hz) power, which showed spatial properties resembling those of fMRI resting state networks (Baker et al., 2014). A similar approach has been applied to source localised task MEG data, identifying states based on their power and phase via a time delay-embedded (TDE) observation model. The approach identified changes in the state dynamics in response to face stimuli (Quinn et al., 2018), demonstrating that the model can be applied to task data, as well as resting state data. This has further been demonstrated in motor tasks, in which HMM with a TDE observation mode was used to identify beta bursts and investigate their dynamics (Pakenham et al., 2020; Seedat et al., 2020). More recently, the model has also been applied to clinical MEG data from epilepsy patients,

identifying periods of epileptic activity and localising the epileptogenic regions (Seedat et al., 2023).

Of most relevance to the current work, the HMM approach has also been applied to EEG-fMRI data, in which states are inferred from one modality, whilst the other provides more information about the characteristics of the state, similar to traditional EEG-informed fMRI, or fMRI-informed EEG approaches. This approach provides a complete characterisation of each state, allowing for both high temporal and spatial resolution. For example, an EEG-fMRI sleep study recently investigated the links between fMRI resting-state network dynamics and sleep cycles (Stevner et al., 2019). By training the HMM on the fMRI data, they were able to identify states comprised of fMRI resting networks and link them with the various stages of the sleep cycle identified from the EEG. Similarly, the opposite approach has been demonstrated, in which the states were defined based on frequency bands of the EEG data, and the fMRI data was used to identify each state's BOLD correlates (Hunyadi et al., 2019). This study showed that the EEG states showed BOLD correlates resembling resting-state fMRI states, providing states with high spatial, temporal and spectral resolution. Therefore, Hidden Markov modelling offers a promising methodology for defining brain states from E/MEG and fMRI data and investigating their relationship with visual response variability.

Thesis Aims

Based on the literature discussed within this chapter, it is clear that spontaneous activity is much more complex than initially thought and displays reoccurring patterns that can be defined as brain states. Despite this, no current research has explored how these states interact with and potentially shape the brain's responses to visual stimuli. Therefore, the primary goal of this thesis was to identify brain states using Hidden Markov Modelling, characterised using both EEG and fMRI, and evaluate their impact on evoked visual responses, including BOLD

responses, VEPs and alpha ERD/S.

This goal was split into several research questions, each answered within a separate chapter:

1. Can HMMs characterise brain state differences associated with a large and well-characterised behavioural manipulation? (Chapter 3)
2. Does HMM state defined from whole-brain fMRI data explain variability in fMRI and EEG responses to visual stimulation? (Chapter 4)
3. Does HMM state defined from EEG data in the alpha band explain variability in fMRI and EEG responses to visual stimulation? (Chapter 5)
4. Do thalamic HMM states explain variability in fMRI and EEG responses to visual stimulation? (Chapter 6)

Chapter 2

Methods

Abstract

This chapter summarises the various methodologies and analyses used consistently within subsequent chapters. Specifically, details of the participants, tasks, stimuli and EEG-fMRI acquisition are provided, as well as information on the implementation of the EEG and fMRI processing, Hidden Markov Modelling, and statistical analysis. Analyses that are chapter-specific are contained within the Methods section of each chapter.

Datasets

Eyes Open / Eyes Closed Dataset

Overview

The eyes open / eyes closed (EO/EC) dataset was collected within the lab prior to my PhD and consisted of four 10-minute runs of EEG-fMRI resting state, alternating between eyes open and eyes closed with each run, i.e., EO/EC/EO/EC. This data was exclusively used within Chapter 3, where the dynamics between EO and EC rest were compared using a HMM. As Chapter 3 has been submitted for publication, it is presented as submitted, and thus, the relevant methodological details are self-contained within that Chapter's Methods section.

Visual Response Variability Dataset

Overview

The Visual Response Variability dataset is an EEG-fMRI dataset collected during my PhD, in which participants either passively viewed visual stimuli or had to actively identify target stimuli. The dataset was comprised of two sessions: an EEG-fMRI session, where participants completed the active and passive task, and a retinotopy fMRI session, where additional anatomical images were acquired, as well as retinotopy data (retinotopy data not analysed). Additionally, questionnaires were collected after the participant's final session; however, these were not analysed within this thesis. This dataset was used within Chapters 4, 5 and 6 of the thesis, and thus, details of the dataset are provided within this chapter.

Participants

All participants were recruited using opportunity sampling via email advertisement. Participants were only included within the experiment if they were aged between 18 and 35 years old, right-handed, neurologically healthy, and were not taking any medication. A total of thirty participants (Mean age = 22.9y, SD = 3.3y, range: [18, 32]y; Sex: 60.0% females, 40.0% males) completed the experiment. One participant's EEG and fMRI data were removed due to excessive movement during the experiment (mean = >5mm), and another participant's EEG data were removed due to electrode bridging, resulting in 29 participants with fMRI data, and 28 with EEG data.

Ethics

All participants gave written informed consent before completing each session and were compensated for their participation. The experiment was approved by the Research Ethics Committee at the University of Birmingham (ERN_110429AP85).

Stimuli

Visual stimuli were created and presented using the RadialStim module within Psychopy, displayed via a VPixx Projector (120 Hz). The stimuli consisted of reversing black and white checkerboards, presented to the left visual hemi-field (see Figure 2.1). The initial checkerboard was presented for 0.5s before reversing and remaining on the screen for another 0.5s. Each inter-stimulus interval (ISI) was randomly sampled between 16.5s and 21s. A mixture of standard and target stimuli was used within the experiment, with standard stimuli appearing and simply reversing, whilst target stimuli reduced in contrast upon reversal. This change in contrast consisted of a change in the luminance of the stimuli relative to the grey background, which is commonly refereed to as Michelson contrast ($C_{Michelson}$). The contrast used for each subject was determined via a yes / no discrimination task.

Thresholding

Before starting the experiment, participants were presented the stimuli inside the scanner but without the scanner running, with 50% showing a contrast change upon checkerboard reversal (i.e. 50% target stimuli). Participants were instructed to respond via a button box, pressing the left button if they saw the contrast change and the right if they did not perceive the change. A shorter ISI of 1s was used during the task; however, the task would only continue once the participant made a response. After ten trials, the contrast increased, making it more difficult for the subject to detect the target stimuli. Once completed, a psychometric function (see Figure 2.2 and equation 2.3) was fitted to the results, and the contrast at which each participant achieved 80% accuracy was then used for the main task.

Task

The main task consisted of two viewing conditions: active and passive. Both conditions were visually the same, with participants being presented with the aforementioned checkerboard

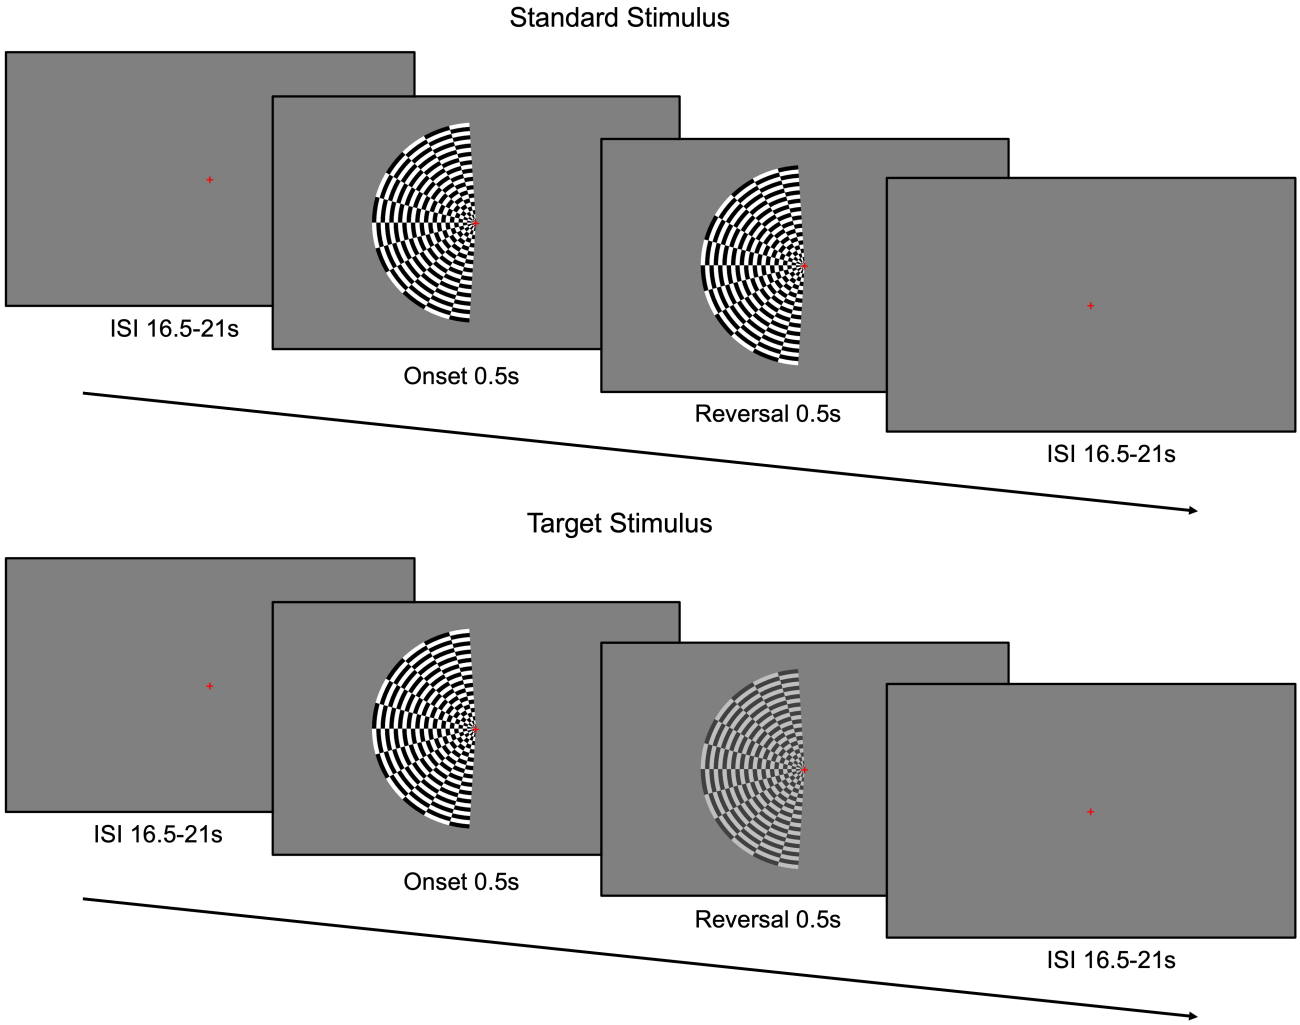


Figure 2.1: Examples of standard and target stimuli with timings.

stimuli. Stimuli were jittered, appearing every 16.5-21 seconds. This duration was chosen to allow single-trial haemodynamic responses to the stimuli to be measured without needing fitting/deconvolution approaches. A total of 36 stimuli were shown, with 27 being standard stimuli without contrast change and the remaining 9 being target stimuli, showing a change in contrast after reversal. The contrast change was subject-specific, as determined by the yes or no thresholding task (see Thresholding). In total, four 10.6-13.4 minute runs were acquired. Conditions were counterbalanced between participants, with half the participants completing the blocks in the order (passive, active, passive, active) and the rest receiving the inverse [active,

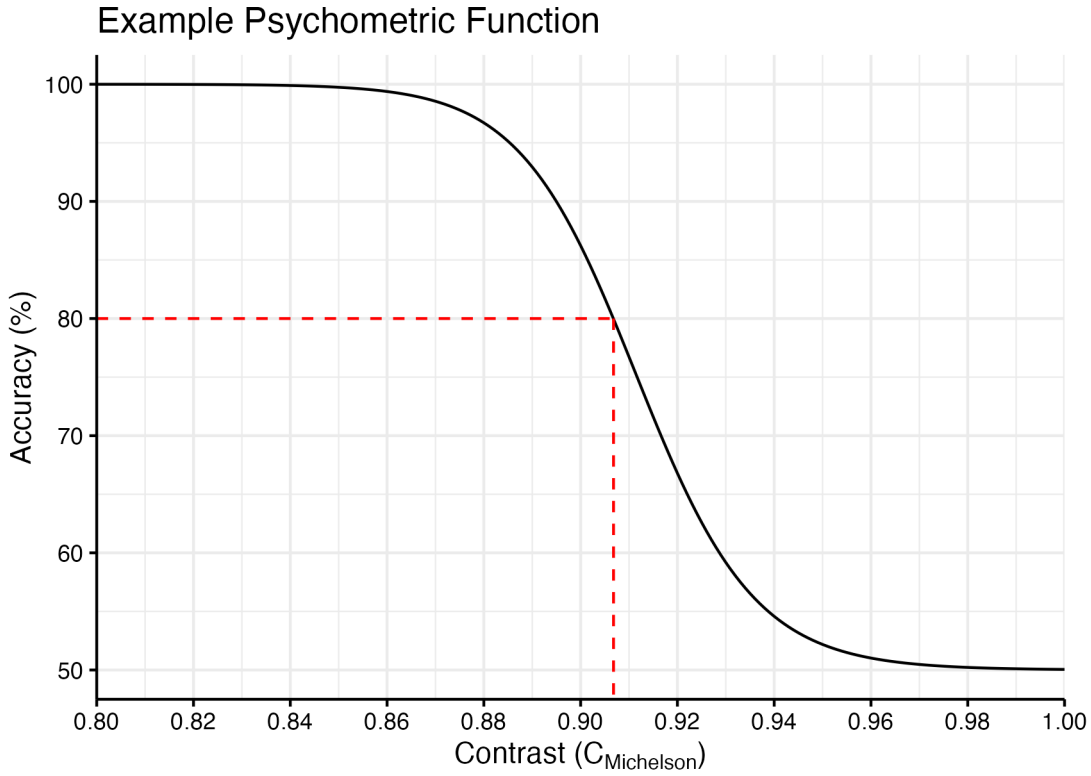


Figure 2.2: Example psychometric function. The y axis represents subjects accuracy during the yes/no task, whilst the x axis represents the contrast of the checkerboard relative to the grey background, known as Michelson contrast ($C_{Michelson}$). When the contrast is low, the change in contrast is easier to detect and thus accuracy starts high. As the contrast increases, the change becomes more difficult to detect, and accuracy decreases, plateauing at chance (50%, as at a contrast of 1, there is no visual difference between the stimuli). Red lines mark the point of 80% accuracy.

$$\left(1 - \frac{0.5}{1 + \left(\frac{x}{\alpha}\right)^{-\beta}}\right) \times 100$$

Figure 2.3: Psychometric function equation.

passive, active, passive]. A red fixation cross was present at the centre of the screen throughout all blocks. During the passive condition, participants were instructed to maintain fixation on the cross throughout the run. Within the active condition, participants were instructed to press the left button of the response box every time they detected a change in contrast. Participants were instructed to respond once the stimuli had left the screen. The trial was recorded as a

miss if the subject did not respond within one second after the removal of the stimulus.

EEG-fMRI Acquisition

Before arriving at the first session, participants underwent MRI screening via email, and again in person on the day of the experiment. Once screened, the EEG cap (64 Channel BrainCap MR with Ag/AgCl ring-type scalp electrodes and carbon wire loops, Brain Products, Germany) was then fitted to the participant's head and aligned centrally to the following landmarks: Nasion, Inion, RPA, and LPA. Electrodes were positioned in the cap according to the 10-20 system. Rubbing alcohol was applied to the scalp below each electrode to remove grease, improving electrode impedance. Abrasive electrolyte gel was then applied using a syringe to each electrode, and a cotton swab was used to rub the gel into the scalp to ensure contact with the electrode. Electrode impedance was kept below 20kΩ. Additionally, an electrode was positioned over the participant's left collarbone to record their ECG. The live EEG recordings were then visually inspected to ensure signal quality. Participants were also shown the impact of blinks and muscle clenching on this signal (e.g. clenching of teeth) and were instructed to avoid excessive movement and blinking to reduce artefacts. Prior to entering the MRI scanner control room, participants were asked to attach disposable electrodes on their chest to record the vectorcardiogram (VCG) (K. J. Mullinger et al., 2008) via the MRI scanner's physiological monitoring system (PMU).

The scanner cold head was turned off and remained off for the entirety of the session. The participant was then brought into the control room while the EEG equipment was set up. A small plastic trolley was used to transport the equipment, consisting of a laptop with Brain Vision Recorder, two 32-channel amplifiers (BrainAmp MR Plus), one BrainAmp EXG MR amplifier, two battery packs, a SyncBox, and a Labjack. The trolley was positioned next to the stimulus PC and connected via USB to send experiment-related triggers, and a parallel cable from the MRI scanner was connected to the SyncBox to record the MRI scanner triggers (i.e. the onset of each volume). Once connected, the amplifiers and batteries were positioned behind the bore of

the MRI scanner on a plastic table weighted down by sandbags to reduce vibrations. Amplifiers were connected to the laptop in the control room via fibre optic cables. The participant was positioned within the MRI scanner, where the VCG electrodes were connected to the PMU, alongside a respiratory belt and pulse oximeter to record the participant's breathing and pulse. The EEG cap connectors were connected to the amplifiers and weighted down with sandbags to further reduce vibrations. EEG electrodes were connected to the two 32-channel amplifiers, and the carbon wire loops were connected to the EXG amplifier. Impedances were checked within the control room, and more gel was applied if any electrodes came loose during positioning. The EEG data were recorded at 5 kHz and synchronised with the MRI scanner clock (K. J. Mullinger et al., 2008), alongside hardware filters of 0.016–250 Hz. EEG was recorded using Brain Vision Recorder (Brain Products, Germany). Once the task was completed, participants were removed from the MRI scanner whilst leaving the EEG cap on. The electrode positions were digitised in a separate room with a Polhemus FastTrak system and Brainstorm (Tadel et al., 2011).

MRI Sequences

All MRI data were acquired with a 3T Seimens MAGNETOM Prisma scanner using a 64-channel head coil. Within the EEG-fMRI session, EPI data (52 slices, 2.5 × 2.5 × 2.5 mm^3, TR = 1010 ms, TE = 36 ms, MB factor = 4, flip angle = 55°) were acquired during each run of the task (four runs total). The number of volumes varied with each run as the ISI was sampled randomly. In addition to the EPI data, a high-resolution anatomical T1 image (256 slices, 1 x 1 x 1mm^3, TR = 2000 ms, TE = 2 ms, flip angle = 8°) was acquired to facilitate registration of the EPI data to standard space, as well as a field map (36 slices, 3 x 3 x 3.75mm^3, TR = 400 ms, TE = 4.92 ms, flip angle = 45°) to correct for B0 field inhomogeneities in the EPI data. This resulted in a total scanning time of approximately 60 minutes.

Within the retinotopy session, EPI data (57 slices, 110 volumes, 2.5 × 2.5 × 2.5 mm^3, TR = 1500 ms, TE = 35 ms, flip angle = 71°) were acquired during each run of retinotopic mapping.

A different EPI sequence was used to ensure that the duration of each retinotopy cycle was a multiple of the TR. Additionally, a second high-resolution T1 image was acquired, as well as another field map and a high-resolution T2 image (320 slices, 0.8 x 0.8 x 0.8 mm^3, TR = 3200ms, TE = 563 ms, flip angle = 120°) to aid segmentation and registration to standard space. For some participants, this T2w sequence exceeded the SAR limit, and thus, a lower resolution T2w image was obtained (256 slices, 1 x 1 x 1 mm^3, TR = 3200ms, TE = 408 ms, flip angle =120°).

Analyses

High Performance Cluster and Containers

All fMRI and EEG preprocessing was performed using the University of Birmingham High Performance Cluster: BlueBEAR (see http://www.birmingham.ac.uk/bear for details). fMRI Preprocessing and analysis were performed using various containers via Singularity/Apptainer, including a custom container (containing FSL, ANTS, and SPM) created using Neurodocker to perform the Nipype (Gorgolewski et al., 2011) and Nilearn analyses.

BIDS Conversion

All data were converted to Brain Imaging Data Structure (BIDS). MRI DICOM images were converted to NIfTI and formatted according to BIDS guidelines using Heudiconv (Halchenko et al., 2023) within a singularity container. Digitised electrode positions were added to the EEG data and formatted according to BIDS using MNE-BIDS (Appelhoff et al., 2019). All behavioural data were outputted in BIDS formatted and organised using a custom BASH script.

MRI Preprocessing

The majority of MRI preprocessing was performed using *fMRIPrep* (Version: 1.0.4) (Esteban et al., 2019; Esteban et al., 2023), a Python-based software package that aims to provide a standardised MRI preprocessing pipeline compatible with BIDS. Specifically, fMRIPrep handles many of the standard preprocessing steps, including but not limited to; field inhomogeneity correction, brain extraction and segmentation, motion correction, standardisation to standard space, and, in this instance, ICA AROMA to identify noise components. As fMRIPrep aims to be a standardised pipeline, it does not include preprocessing steps that will vary between projects, e.g. spatial smoothing and nuisance regression. These final preprocessing steps were performed using a custom Nipype (Gorgolewski et al., 2011) workflow script. One of the main goals of using fMRIPrep is to standardise the preprocessing, thus increasing transparency, the direct outputs from fMRIPrep are given below, followed by an overview of the custom Nipype script.

fMRIPrep Overview

Results included in this manuscript come from preprocessing performed using *fMRIPrep* 21.0.4 (Esteban et al. (2019); Esteban et al. (2023); RRID:SCR_016216), which is based on *Nipype* 1.6.1 (Gorgolewski et al. (2011); Gorgolewski et al. (2016); RRID:SCR_002502).

Preprocessing of B0 inhomogeneity mappings

A *B0* nonuniformity map (or *fieldmap*) was estimated from the phase-drift map(s) measure with two consecutive GRE (gradient-recalled echo) acquisitions. The corresponding phase-map(s) were phase-unwrapped with `prelude` (FSL 6.0.5.1:57b01774).

Anatomical data preprocessing

All T1-weighted (T1w) images were corrected for intensity non-uniformity (INU) with `N4BiasFieldCorrection` (Tustison et al., 2010), distributed with ANTs 2.3.3 (Avants et al., 2008, RRID:SCR_004757). The T1w-reference was then skull-stripped with a *Nipype* implementation of the `antsBrainExtraction.sh` workflow (from ANTs), using OASIS30ANTs as target template. Brain tissue segmentation of cerebrospinal fluid (CSF), white-matter (WM) and gray-matter (GM) was performed on the brain-extracted T1w using `fast` (FSL 6.0.5.1:57b01774, RRID:SCR_002823, Y. Zhang et al., 2001). A T1w-reference map was computed after registration of 2 T1w images (after INU-correction) using `mri_robust_template` (FreeSurfer 6.0.1, Reuter et al., 2010). Brain surfaces were reconstructed using `recon-all` (FreeSurfer 6.0.1, RRID:SCR_001847, Dale et al., 1999), and the brain mask estimated previously was refined with a custom variation of the method to reconcile ANTs-derived and FreeSurfer-derived segmentations of the cortical gray-matter of Mindboggle (RRID:SCR_002438, Klein et al., 2017). Volume-based spatial normalization to two standard spaces (MNI152NLin2009cAsym, MNI152NLin6Asym) was performed through nonlinear registration with `antsRegistration` (ANTs 2.3.3), using brain-extracted versions of both T1w reference and the T1w template. The following templates were selected for spatial normalization: *ICBM 152 Nonlinear Asymmetrical template version 2009c* [Fonov et al. (2009), RRID:SCR_008796; TemplateFlow ID: MNI152NLin2009cAsym], *FSL's MNI ICBM 152 non-linear 6th Generation Asymmetric Average Brain Stereotaxic Registration Model* [Evans et al. (2012), RRID:SCR_002823; TemplateFlow ID: MNI152NLin6Asym].

Functional data preprocessing

For each fMRI run, a reference volume and its skull-stripped version were generated using a custom methodology of *fMRIPrep*. Head-motion parameters with respect to the BOLD reference (transformation matrices, and six corresponding rotation and translation parameters) are

estimated before any spatiotemporal filtering using `mcflirt` (FSL 6.0.5.1:57b01774, Jenkinson et al., 2002). The estimated *fieldmap* was then aligned with rigid-registration to the target EPI (echo-planar imaging) reference run. The field coefficients were mapped on to the reference EPI using the transform. BOLD runs were slice-time corrected to 0.458s (0.5 of slice acquisition range 0s-0.915s) using `3dTshift` from AFNI (Cox & Hyde, 1997, RRID:SCR_005927). The BOLD reference was then co-registered to the T1w reference using `bbregister` (FreeSurfer) which implements boundary-based registration (Greve & Fischl, 2009). Co-registration was configured with six degrees of freedom. Several confounding time-series were calculated based on the *preprocessed BOLD*: framewise displacement (FD), DVARS and three region-wise global signals. FD was computed using two formulations following Power (absolute sum of relative motions, Power et al. (2014)) and Jenkinson (relative root mean square displacement between affines, Jenkinson et al. (2002)). FD and DVARS are calculated for each functional run, both using their implementations in *Nipype* (following the definitions by Power et al. (2014)). The three global signals are extracted within the CSF, the WM, and the whole-brain masks. Additionally, a set of physiological regressors were extracted to allow for component-based noise correction (*CompCor*, Behzadi et al., 2007). Principal components are estimated after high-pass filtering the *preprocessed BOLD* time-series (using a discrete cosine filter with 128s cut-off) for the two *CompCor* variants: temporal (tCompCor) and anatomical (aCompCor). tCompCor components are then calculated from the top 2% variable voxels within the brain mask. For aCompCor, three probabilistic masks (CSF, WM and combined CSF+WM) are generated in anatomical space. The implementation differs from that of Behzadi et al. in that instead of eroding the masks by 2 pixels on BOLD space, the aCompCor masks are subtracted a mask of pixels that likely contain a volume fraction of GM. This mask is obtained by dilating a GM mask extracted from the FreeSurfer's *aseg* segmentation, and it ensures components are not extracted from voxels containing a minimal fraction of GM. Finally, these masks are resampled into BOLD space and binarized by thresholding at 0.99 (as in the original implementation). Components

are also calculated separately within the WM and CSF masks. For each CompCor decomposition, the k components with the largest singular values are retained, such that the retained components' time series are sufficient to explain 50 percent of variance across the nuisance mask (CSF, WM, combined, or temporal). The remaining components are dropped from consideration. The head-motion estimates calculated in the correction step were also placed within the corresponding confounds file. The confound time series derived from head motion estimates and global signals were expanded with the inclusion of temporal derivatives and quadratic terms for each (Satterthwaite et al., 2013). Frames that exceeded a threshold of 0.5 mm FD or 1.5 standardised DVARS were annotated as motion outliers. The BOLD time-series were resampled into standard space, generating a *preprocessed BOLD run in MNI152NLin2009cAsym space*. First, a reference volume and its skull-stripped version were generated using a custom methodology of *fMRIPrep*. Automatic removal of motion artifacts using independent component analysis (ICA-AROMA, Pruim et al., 2015) was performed on the *preprocessed BOLD on MNI space* time-series after removal of non-steady state volumes and spatial smoothing with an isotropic, Gaussian kernel of 6mm FWHM (full-width half-maximum). Corresponding "non-aggresively" denoised runs were produced after such smoothing (i.e. the non-aggresively denoised runs are a copy of the main data that undergo smoothing for the purposes of detecting noise components with ICA AROMA). Additionally, the "aggressive" noise-regressors (i.e. the noise component time courses) were collected and placed in the corresponding confounds file. All resamplings can be performed with *a single interpolation step* by composing all the pertinent transformations (i.e. head-motion transform matrices, susceptibility distortion correction when available, and co-registrations to anatomical and output spaces). Gridded (volumetric) resamplings were performed using `antsApplyTransforms` (ANTs), configured with Lanczos interpolation to minimize the smoothing effects of other kernels (Lanczos, 1964). Non-gridded (surface) resamplings were performed using `mri_vol2surf` (FreeSurfer).

A custom Nipype 1.6.1 (Gorgolewski et al., 2011); Gorgolewski et al. (2016); RRID:SCR_002502)

was was then used to remove the noise components identified by ICA AROMA using the regfilt function in FSL, spatial smoothing (FWHM=4mm) using the SUSAN (S. M. Smith & Brady, 1997) function in FSL, and a apply a high-pass temporal filter (100s). Data were then cropped to 16 volumes after the final stimulus using the fslroi function in FSL. This data was then used in subsequent GLM and ROI analyses. For the ICA analyses, the data were transformed to standard space with the transforms generated by fMRIPrep using the applyWarpTransforms from ANTS.

EEG Preprocessing

EEG Preprocessing was performed using EEGLAB (Delorme & Makeig, 2004) and the FMRIB plugin (Niazy et al., 2005). Gradient artefacts were removed from the data with a template artefact subtraction approach with optimal basis sets (PCA components = 3, 35 artefacts per template) using the FASTR function in the fMRIB toolbox. The data were then downsampled to 250 Hz to reduce the computation time of subsequent analysis. To remove the ballistocardiogram (BCG) artefacts, the physiological data was used to detect the onsets of the R peak in the VCG signal. The VCG recordings underwent bandpass filtering (0.6-18Hz) to remove gradient artefacts and a PCA analysis was conducted across all four leads to remove residual noise. The component with the clearest R peaks was then used for peak detection. A custom Python script was used to identify the R peaks and parse the onsets into the EEG events data. The R peak events were then visually inspected alongside the EEG data to ensure they aligned with the BCG artefacts. Where poor peak detection performance was found, the peak detection parameters (e.g. threshold) were modified to optimise the detection of the R peak for that specific subject. The BCG artefacts were removed using the OBS template approach (PCA components = 3), using the R peaks as markers. In some instances, the participant's VCG data were too poor to identify the R peaks; in these instances, the ECG electrode from the EEG was used (one participant), or no BCG correction was applied (one subject). Runs containing carbon wire

loop (CWL) data underwent CWL regression (Van Der Meer et al., 2016) to remove motion and residual BCG artefacts. This was not possible for one subject as the battery ran out of charge before the end of the session. Bad channels were removed based on their spectral properties using the automated rejchan function in EEGLAB. All rejected channels were interpolated before the dataset was cropped to 16.6 s (16 TR) after the last stimulus to ensure the EEG recordings were the same length as the fMRI data. The data were then re-referenced to the average, and an ICA was conducted on the concatenated data (per subject) to remove any remaining noise artefacts, e.g. blinks. Artefactual components were identified using the automated ICLabel function, with a detection threshold of 75% for ocular, muscular, heart artefacts, and channel noise (Pion-Tonachini et al., 2019). The artefactual components were removed from the data via retro-projection, and two copies of the dataset were created, one with a bandpass filter between 1-25Hz for VEP analysis, and one with a bandpass filter of 7-13 Hz for alpha power-based analyses, i.e. alpha ERD/S analysis.

Hidden Markov Modelling

The Hidden Markov Modelling was performed using the HMM-MAR Toolbox in MATLAB (Vidaurre et al., 2016). As discussed in Chapter 1, the HMM approach aims to decompose inputted time series (BOLD or electrophysiological) into recurring patterns of activity (known as states) in a data-driven manner. These states are defined according to an observation model, and depending on the modality of the data, different observation models can define states on different properties. The simplest observation model is the multivariate Gaussian model, in which states are defined according to the mean values of the inputted time series, and the covariance between the time series (though this can be changed to only model one or the other). When applied to fMRI data, this observation model defines states according to a region's mean activity and functional connectivity with other brain regions. For EEG data, this model is typically used with power time series (e.g. Hilbert transformed data) and thus identifies states

according to instantaneous power changes and correlations between electrodes or regions (with source reconstructed data). EEG data also allows more complex observation models to be used, allowing states to be defined according to other spectral properties. For example, the time delay embedded (TDE) model (Vidaurre, Abeysuriya, et al., 2018) and the multivariate autoregressive (MAR) (Vidaurre et al., 2016) observation models are capable of defining states according to both power and phase relations. Typically, the MAR approach is applied to data with fewer channels/regions to avoid overfitting, whilst more accurately characterising the phase relations of each state, and the TDE is applied to data with a large number of channels/regions but provides a less detailed characterisation of the phase relations. Within this thesis, the multivariate Gaussian model was used for all HMMs. This was due to most of the HMMs being trained on BOLD data and the EEG models being primarily concerned with absolute changes in alpha power, as opposed to phase relations between electrodes/regions.

Model Outputs

Once the HMM has been trained, there are two main outputs: the transition matrix and the state likelihood time series. The former is a matrix containing the probability of transitioning from any one state to another state (including itself). Conversely, the state likelihood time series contains the likelihood of a specific state being active within the data at any given timepoint. However, it is often useful to have a binary representation of which state is most active at any given time series (e.g., for calculating state metrics). This can be estimated via the Viterbi algorithm (Forney, 1973), which predicts the most likely sequence of states given the transition matrix and the initial state estimated by the HMM. This binarised time series is known as the Viterbi path.

State Metrics

From the state time series or Viterbi path, it is possible to calculate various state metrics that describe the temporal dynamics of each state. These metrics consist of fractional occupancy, state lifetimes, state interval and the state switching rate. The fractional occupancy of a state represents the percentage of timepoints that a state is active. The state lifetime represents the duration of each state visit (e.g. how long it consecutively remains active), whilst the interval time represents the duration of time between state visits. Lastly, the switching rate represents how frequently the active state changes over a specified duration of time.

To calculate these state metrics, the state visits need to be binary (i.e. active or not active); however, the state time courses only represent the likelihood of each state being active. One solution to this problem is to threshold the state time courses in order to determine when each state is active (e.g. a state is deemed active if the probability is >=80%). Alternatively, the Viterbi path can be used, which is already binarised. Throughout this thesis, the Viterbi path was used to calculate all state metrics.

Model Selection

Similar to blind source separation techniques, such as ICA, the number of states the data is to be decomposed into must be defined before training the model. This number is determined via a model selection process, in which the most reliable and accurate model is selected. If there are too many states relative to the amount of the data, the model training process is likely to fail, which typically manifests itself in the form of a singular state being active across the entirety of each run, or specific states only being active for particular participants. Furthermore, as the model training process is stochastic (e.g. it will not produce the same solution every time) the reliability of the model needs to be assessed, which typically decreases with the addition of more states.

Therefore, each HMM was trained multiple times (the number of repeats that were used in training is defined in each Chapter) across several potential state numbers (2-15), and three metrics were assessed to determine the state number. These metrics included variational free energy (VFE), repeat reliability, and maximum fractional occupancy. Within the HMM-MAR toolbox, the VFE is the metric that is minimised during the training process, and a lower VFE indicates that the model explains more of the variability within the data. Conversely, the repeat reliability represents the correlation coefficient between two repeats of the model with the same number of states. This metric was calculated across all pairs of repetitions, and the average was taken to give an indication of how reliable the model solutions were for a specific state number. Lastly, maximum fractional occupancy represents a state's highest fractional occupancy value in a particular run. If this metric is on average closer to one, this indicates that a single state dominates the model. Thus, the model is unlikely to truly represent the underlying dynamics, given the dynamic complexity of brain function. Overall, a combination of these metrics was used to determine the state number and ensure that the model reliably and accurately models the temporal dynamics within the data.

Determining Baseline States

In order to determine which HMM state was active during the baseline period of each stimulus, the Viterbi path was split and aligned with the individual runs. According to the Viterbi path, the active state at (t-1) timepoint prior to the stimulus onset was deemed the baseline state. To assess the reliability of using this timepoint, the baseline states were calculated at different timepoints (t0 - t-5) and the calculated baseline state sequences were correlated.

VEP and Alhpa ERD/S Analysis

The VEP and Alpha ERD/s analyses were performed using MNE-Python (Larson et al., 2023). VEPs were extracted from the 1-25 Hz bandpass filtered data from a -200ms to 300ms epoch

for each stimulus. All epochs underwent mean baseline correction (mean of the baseline period subtracted from each timepoint) using data from -200ms to 0ms from each epoch as the baseline period. Any epochs containing a peak-to-peak difference of >=100 µV were marked as bad and rejected from subsequent analysis. As the VEPs showed some spatial variation between subjects, the VEP was selected from the right posterior occipital electrode (O2, PO4, PO8, P2, P4, P6, P8) showing the largest P100 amplitude. Alpha ERD/S were extracted in a similar manner, using the 7-13 Hz bandpass filtered data, with an epoch length of -1s to 5s using the Morlet wavelet transform. Alpha ERD/S epochs were extracted from the same electrode used for the VEP analysis and underwent mean baseline correction (using data from -1s to 0s, defined per epoch). Bad epochs were rejected using the same criteria as for the VEPS. Grand averages of both the VEPs and alpha ERD/S were calculated for each condition, trial type, and baseline HMM state.

The amplitude of the P100, N150, and P250 components were calculated for statistical comparison. The grand average of all epochs was calculated to determine the mean latency of each component's peak. Amplitudes were extracted in a 10ms window surrounding each component, resulting in the following windows for each component; 87-97ms for P100. 159-169ms for N150 and 235-245 for P250 (see Figure 2.4). Similarly, the alpha ERD/S amplitudes were extracted from two separate windows containing the ERD (0.25 s - 1.35 s) and ERS (1.75 s - 4.5 s) (see Figure 2.4).

BOLD Response Analysis

Mean BOLD signal time courses were extracted using the fslmeants function from FSL (ROIs defined within each Chapter). Epochs were extracted around the onset of each stimulus (-2 volumes to 17 volumes) using a custom Python script. Baseline correction was performed by converting the time courses into percent signal change relative to the baseline period. The

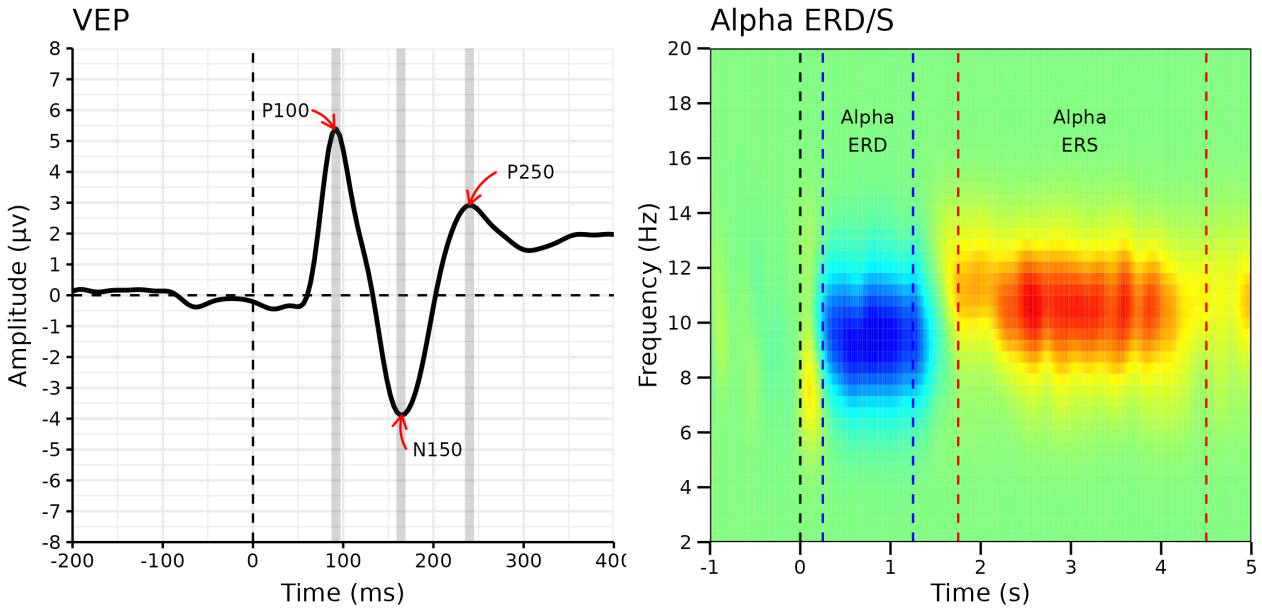


Figure 2.4: The left plot shows the grand average VEP response from the visual response variability dataset. Grey areas mark windows used to calculate the values of each VEP component (10ms window around the peak of each VEP component). The right plot shows the grand average alpha ERD/S response from the visual response variability dataset. Blue dashed lines mark the area used to calculate the alpha ERD values. Red dashed lines mark the window used to calculate the alpha ERS values.

baseline period was defined as two volumes preceding the stimulus onset and the stimulus onset volume (three volumes total). Percent signal change (PSC) conversion was performed individually for each trial. Specifically, the baseline mean was subtracted from each value before being divided by the baseline mean and multiplied by 100 (see Equation 2.6). For statistical comparison, the BOLD peak and undershoot values were also calculated by taking the mean within a set of windows, with the mean peak value being calculated between 2 and 9 TR, whilst the mean undershoot value was calculated between 10 and 17 TR (see Figure 2.5).

Statistics

All statistical testing was performed within R Studio. Plots were created using the ggplot2 package (Wickham, 2016).

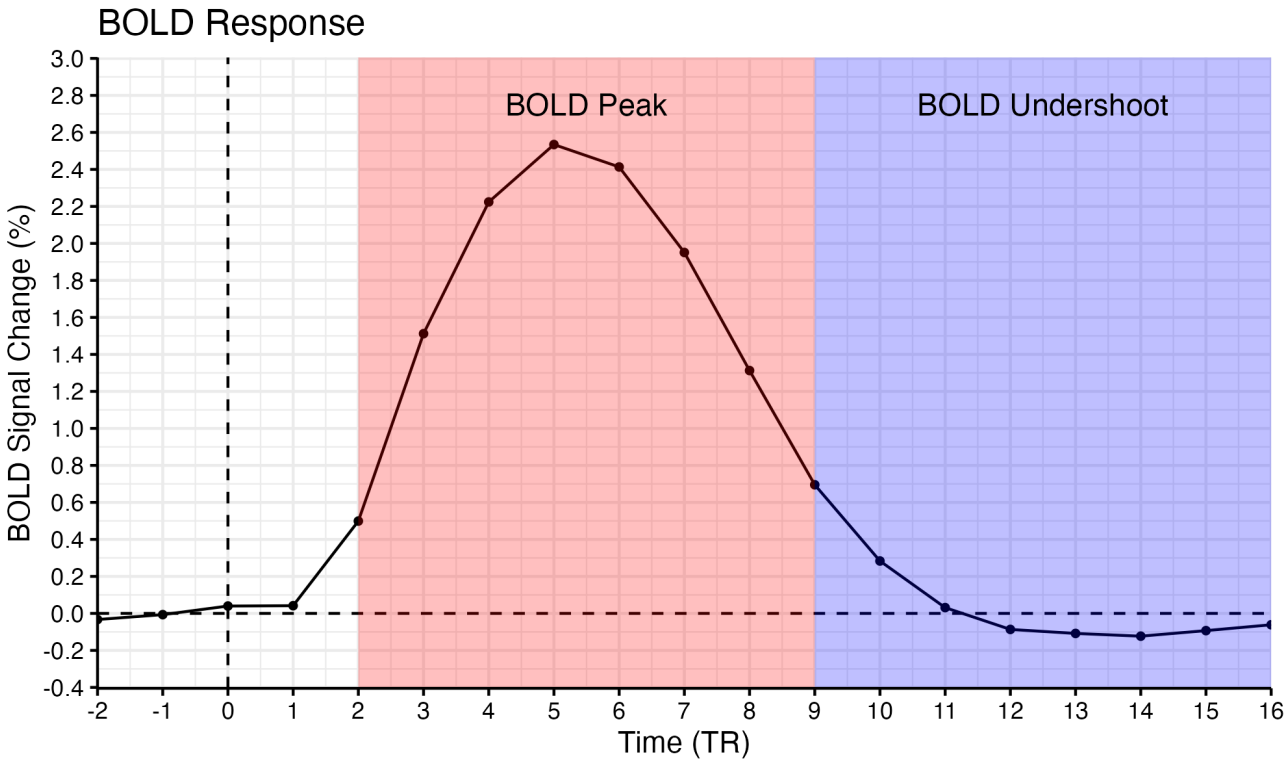


Figure 2.5: Grand average BOLD response in right V1 from the visual response variability dataset. The red area marks the window used to calculate the BOLD peak values. The blue window marks the window used to calculate the BOLD undershoot values.

$$Percent\,Signal\,Change = \frac{(x - Baseline\,Mean)}{Baseline\,Mean} \times 100$$

Figure 2.6: Percent signal change equation. x represents the original BOLD value. The baseline mean was calculated between -2-0TR (-2.02-0s).

Comparing Response Amplitudes Between HMM States

The impact of HMM states on visual responses (e.g. BOLD responses, VEP component amplitudes, and alpha ERD/S amplitudes) was compared statistically via linear mixed modelling (LMM) using the lme4 package within R (Bates et al., 2015). A LMM approach was chosen over a repeated measures ANOVA due to missing data, as which states occurred during the baseline could not be controlled for (i.e. during a run, state x may never be active during the baseline for any target stimuli). Models were trained on the single trial response amplitudes.

All response modalities were compared using the same model specification, with the condition, stimulus type and baseline state being entered as fixed effects and participant being entered as a random effect, resulting in the following model specification: visual response amplitudes $\sim$ condition * baseline state * stimulus type + (1 | subject). P-values were acquired from the LMMs using the LmerTest package (Kuznetsova et al., 2017) using the Satterthwaite method. Post hoc tests compared the estimated marginal means using the emmeans package. All comparisons were corrected for multiple comparisons using Bonferroni correction.

Comparing State Metrics

Fractional occupancy and state switching rate values were compared statistically using a series of paired t-tests. State life and interval times were not compared due to issues with missing data (e.g. if a subject didn't visit a particular state within a run). When comparing the fractional occupancy values, multiple comparisons were corrected for using the Bonferroni method (number of comparisons = number of states). Statistics were formatted using the report package in R.

Chapter 3

Brain State Dynamics Differ Between Eyes Open and Eyes Closed Rest

Abstract

The human brain exhibits spatiotemporally complex activity even in the absence of external stimuli, cycling through recurring patterns of activity known as brain states. However, the optimal method for identifying brain states with both a high spatial and temporal resolution remains unclear. Recent research has applied Hidden Markov modelling (HMM), a data-driven approach to decomposing data into brain states. Here, a HMM is fitted to concurrently collected EEG-fMRI resting state data, to assess the model's ability to distinguish dynamics between eyes open and eyes closed rest, and to examine how this manipulation impacts EEG vs fMRI data. The fMRI HMM identified multiple states with differing dynamics between the two conditions. The eyes closed condition was more associated with sensory states and increased sensory functional connectivity, while the eyes open condition was associated with salience and default mode states. The EEG HMM identified high and low alpha power states, with the higher alpha power state being more active within the eyes closed rest condition, and both states showing BOLD correlates resembling previously identified alpha-BOLD Correlations. Both the fMRI and EEG Hidden Markov models were able to identify spatiotemporal dynamics observed within previous literature, as well as identifying novel brain state properties. A sliding window analysis of each state's fractional occupancy revealed unique temporal dynamics, with multiple

states exhibiting large changes in fractional occupancy upon eye closure, before returning to a baseline after a few minutes. Hidden Markov Modelling can identify novel spatiotemporal dynamics in concurrent EEG-fMRI data and serves as a powerful tool for the research of brain state dynamics within and across modalities.

Introduction

During periods of waking rest, the human brain continues to exhibit spatiotemporally complex activity despite the reduction in external stimuli. Commonly referred to as resting state data, subjects are instructed to relax throughout the duration of a recording (e.g. fMRI or E/MEG), without any task being present. The analysis of such data has led to the discovery of various resting state networks, which have since been shown to persist throughout both rest and task conditions (Biswal et al., 1995; S. M. Smith et al., 2009) and been associated with numerous cognitive processes (Buckner et al., 2008; Corbetta & Shulman, 2002; Vossel et al., 2014), as well as exhibiting differences between clinical and control populations, which could potentially serve as biomarkers for neurological and mental health conditions (Buckner et al., 2008; Calhoun, Eichele, et al., 2009; Greicius et al., 2004; Rosazza & Minati, 2011). Conventionally, these networks have been identified by functional connectivity (FC) based approaches, in which the statistical temporal correlation between the neural activity of multiple brain regions is calculated. Resting-state networks are then defined as a group of regions showing strong FC with each other. These approaches have generally been applied across the entire duration of a recording lasting several minutes, providing a summary measure of the FC between regions over this timescale. Whilst this method has proved useful, it assumes that the FC between regions is static/stationary, or in other words, that the relationship between brain regions does not change over the duration of the recording. Many studies have demonstrated that this is clearly an oversimplification of the complexity of the brain's activity.

Dynamic approaches, such as sliding window analyses (E. A. Allen et al., 2014), leading eigenvector dynamic analysis (LEiDa) (Cabral et al., 2017), and temporal ICA (S. M. Smith et al., 2012), have all demonstrated that FC relationships are not static/stationary; and that the temporal correlations, both between network nodes and between different networks, varies over time. Furthermore, these changes in network FC are non-random, with specific patterns of activity reoccurring throughout the duration of recordings (E. A. Allen et al., 2014; Sakoğlu et al., 2010). This has most commonly been demonstrated using sliding window-based analyses, followed by k-means clustering, resulting in a set of reoccurring FC patterns, commonly referred to as connectivity, or brain "states". Once identified, it is then possible to calculate various temporal metrics for each state, such as the average time a subject spends within each state (state lifetimes), or the frequency at which subjects switch between states (switching rates). Collectively, this has led to the proposal of the chronnectome (Calhoun et al., 2014) or the dynome (Kopell et al., 2014), which aim to provide a comprehensive characterisation of brain FC and dynamics (Calhoun et al., 2014). Such methods have also shown potential clinical application, with people with schizophrenia exhibiting altered connectivity states (Sakoğlu et al., 2010) and state temporal metrics (Damaraju et al., 2014) relative to control participants.

However, sliding window-based analyses have recently been criticised due to the requirement to arbitrarily select parameters, such as the window length and offset (Hindriks et al., 2016; Laumann et al., 2016; Leonardi & Van De Ville, 2015; Zalesky & Breakspear, 2015). Additionally, simulations have shown that the method's ability to detect state transitions and state durations is significantly impacted by the choice of window length (Shakil et al., 2016). An alternative approach is to identify and model brain state dynamics with a Hidden Markov Model (HMM) (Baker et al., 2014; Vidaurre et al., 2016). Unlike sliding windows, HMM aims to decompose the data into a dynamic sequence of distinct states in a data-driven manner, therefore avoiding the need to specify a window length and offset.

However, the current research has exclusively focused on quantifying brain states in one neuroimaging modality, and very little is known about how brain states compare across modalities, or the relationships between neural and haemodynamic brain states, which limits their wider interpretation. Ideally, brain states would be defined with both high spatial and temporal resolution. This would best enable monitoring of dynamic activity patterns and regional interactions for applications such as studying recruitment during task performance and its relation to behaviour or spontaneous alterations during changes in awareness or consciousness. Whilst fMRI can provide millimetre spatial resolution recording of both cortical and subcortical brain activity, it is limited temporally as a result of the sluggish nature of the BOLD signal (Buxton et al., 1998). Conversely, electrophysiological measures such as E/MEG provide millisecond temporal resolution and a direct measure of neural activity compared to the BOLD signal. They also allow brain activity to be analysed from a spectral perspective, e.g. studying specific frequency bands rather than the integration of all underlying local activity represented by the BOLD signal. However, the spatial resolution and sensitivity to subcortical structures of E/MEG are significantly reduced compared to fMRI. One potential method to achieve high spatiotemporal resolution is to employ a multimodal approach, recording EEG and fMRI data concurrently. This also provides a way of simultaneously defining brain states from both EEG and fMRI data, and hence comparing the two directly. For example, Hunyadi et al. (2019) applied HMM to concurrent EEG-fMRI data, defining states by changes in power within the EEG data, and using the state time courses to obtain spatial maps for each state via a GLM analysis of the fMRI data (Hunyadi et al., 2019). To date, however, there has been no explicit attempt to define EEG and fMRI brain states from concurrently recorded data and compare their properties under a behavioural manipulation.

One of the simplest and most fundamental manipulations of brain activity that can be applied is eye closure. Within EEG, there is a significant increase in alpha power upon eye closure (Barry et al., 2007; Boytsova & Danko, 2010; L. Li, 2010). Similarly, fMRI studies have found

that eye closure induces changes in BOLD activity and FC across visual (Patriat et al., 2013), auditory (Marx et al., 2003; Patriat et al., 2013; Zou et al., 2015), somatosensory (Liang et al., 2014; Yang et al., 2007) and insula cortices (Goldman et al., 2002), as well as the default mode (Jao et al., 2013), and dorsal attention networks (DAN) (C. Wang et al., 2016; D. Zhang et al., 2015). Concurrent EEG-fMRI studies have identified correlations between changes in resting-state alpha power and the BOLD signal in these brain regions, both using static (de Munck et al., 2007; Goldman et al., 2002; Laufs et al., 2006) and dynamic (Mayhew & Bagshaw, 2017) analysis methods. However, a deeper understanding of the dynamic and spontaneous alpha-BOLD relationships is required, and how they vary with the manipulation of awareness and attention, which can be provided by analysing fluctuations in EEG and fMRI brain-states.

Therefore, this study used HMM analysis of concurrent EEG-fMRI resting state data and compared brain state dynamics between periods of eyes open (EO) and eyes closed (EC) rest. HMMs were trained separately on each modality to allow comparison with previous EEG and fMRI literature. We quantified the changes in brain state dynamics in the two modalities and used a GLM to identify the BOLD correlates of the EEG states. Our primary aim was to understand how HMM defined states compared across EEG and fMRI data, and how state properties changed with eye closure.

Methods

Participants

A total of 21 healthy participants (11 female, age range = 21-34y, mean age = 25.8y) were recruited via opportunity sampling at the University of Birmingham. Ethical approval for the study was given by the Research Ethics Committee of the University of Birmingham, and all participants gave informed consent before participation. No participants were excluded from

the fMRI analysis; however, three participants were removed from the EEG analysis due to data corruption.

EEG-fMRI Data Acquisition

Concurrent EEG-fMRI data were acquired on a 3T Philips Achieva scanner using a 32-channel SENSE receive head coil and 64-channel EEG system (MR-plus amplifiers, Brain Products, Germany) at the Birmingham University Imaging Centre (BUIC). Four resting state BOLD fMRI runs were acquired using a whole-brain gradient echo EPI sequence (TR = 2000ms, TE = 35ms, flip angle = 80°, voxel size = 3x3x4mm^3, FOV = 240 x 128 x 240mm^3, slices = 32).

Participants were instructed to either keep their eyes open or closed for the duration of one run, alternating after each run (starting with EO). Two runs were acquired for each of the EO and EC conditions. EO and EC were not alternated within a run. Each resting-state run contained 300 whole brain volumes, resulting in 10 minutes per run. In total, 1200 volumes were acquired per participant in a task acquisition time of 40 mins. A high resolution (1 mm isotropic MPRAGE with TR=2000ms, TE=2ms, TI = 880ms, flip angle = 8°, FOV 256x256mm) T1-weighted anatomical image was acquired to facilitate alignment to MNI152 space, resulting in a total run time of 45 minutes.

EEG data were acquired via 63 (including reference and ground) Ag/AgCl MR-compatible scalp electrodes, positioned in accordance with the 10-20 system, with the remaining electrode recording the subject's electrocardiogram (ECG) from the clavicle (EasyCap, Brain Products, Germany). Electrode impedances were kept below 20 kΩ. EEG was recorded using Brain Vision Recorder (Brain Products, Germany) at a sampling rate of 5 kHz, alongside hardware filters of 0.016-250 Hz. Recordings were synchronized with the MR scanner's internal clock (Syncbox, Brain Products, Munich, Germany) (Mandelkow et al., 2006; K. J. Mullinger et al., 2008). The MR scanner's physiological monitoring system was used to record cardiac cycles via

vectorcardiogram (VCG) (K. J. Mullinger et al., 2008).

fMRI Prepcrocessing

All fMRI analysis was performed within FSL (6.0.1) (Woolrich et al., 2009). Automated brain extraction S. M. Smith, 2002, motion correction (MCFLIRT) (Jenkinson et al., 2002), slice timing correction, spatial smoothing (Gaussian Kernel, FWHM = 6mm) and high-pass temporal filtering (>0.01Hz) were applied to all fMRI data. Transforms to MNI152 (2mm) standard space were calculated via non-linear alignment (FLIRT and FNIRT) (Jenkinson & Smith, 2001). Single-subject ICA (Beckmann & Smith, 2004) and FIX (Griffanti et al., 2014; Salimi-Khorshidi et al., 2014) (Training Set = Standard.R, threshold = 20) were applied to the data to identify noise components, which were then removed using `fsl_regfilt`. Cleaned data were transformed to standard space using the applywarp command with the previously calculated transformations.

fMRI Group ICA

As a data reduction step for the fMRI HMM analysis, (Ahrends et al., 2022; Vidaurre, Abeysuriya, et al., 2018; Vidaurre et al., 2017) data were temporally concatenated and decomposed into 25 independent components (ICs) via group spatial ICA (MELODIC). A component dimension of 25 was chosen to acquire ICs containing individual resting state networks. Only ICs that had time courses in the lower frequencies and did not spatially overlap with known artefacts (motion, vascular, physiological, susceptibility) or non-grey matter regions (white matter, CSF) were used in the subsequent analysis, resulting in 13 components being retained (see Figure 3.1). ICs were visually compared with the FIND lab ICA components (Shirer et al., 2012) to assist network labelling. ICs were found within visual (occipital pole, medial and lateral), auditory, dorsal attention, default mode, precuneus, salience, sensorimotor (superior and inferior) and frontoparietal (FPN) (left and right) networks. Dual regression (Beckmann et al., 2009; Nickerson et al., 2017) was applied to estimate BOLD signal time courses for each network IC

in each run of each subject.

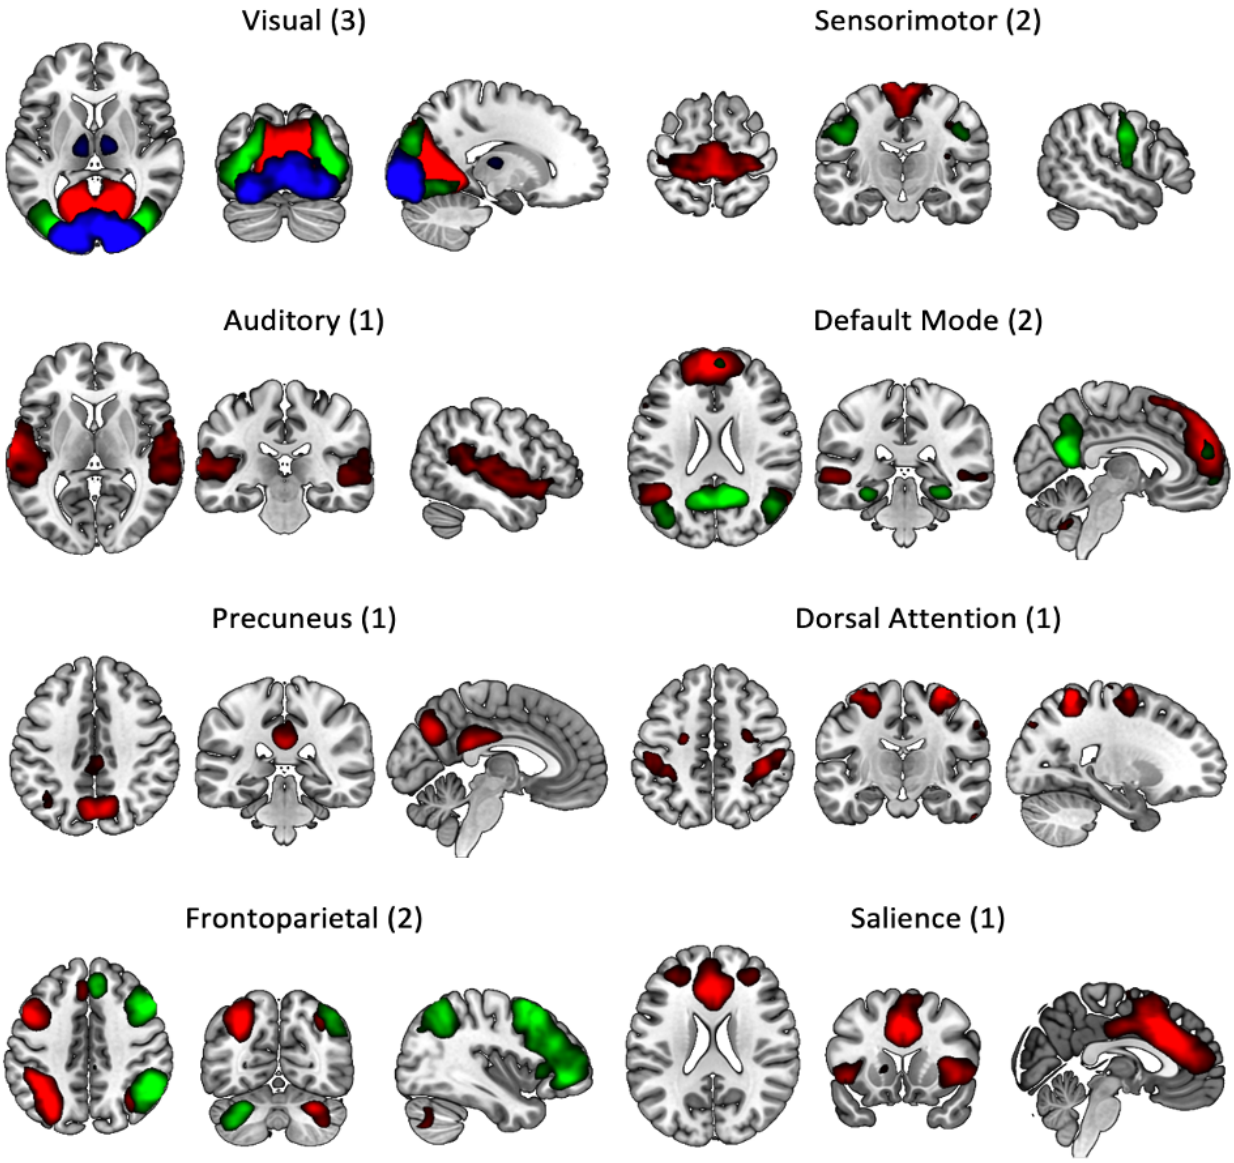


Figure 3.1: Results from the group fMRI decomposition plotted in MNI152 space (left = left). Independent components are grouped into networks. Each colour represents a separate independent component. The number next to each subtitle represents the number of independent components identified for that region/network.

EEG Preprocessing

MR-gradient artefacts were removed from the data using the average template subtraction method with a sliding window (window length = 41) implemented within Brain Vision Analyzer 2.0 (Brain Products, Munich, Germany). Heartbeat events were detected on the VCG and added as markers to the EEG via in-house MATLAB and Perl scripts. Ballistocardiogram (BCG) artefacts were removed via the optimal basis sets method (principal components = 3) using an in-house Python script based on the fMRIB EEGLab plugin (Niazy et al., 2005). Subsequent preprocessing focussed on the alpha band as this was the frequency band expected to be most influenced by the eye closure task. This restricted spectral band also optimises the data quality since both the lower (W. X. Yan, Mullinger, Geirsdottir, et al., 2009) and higher (P. J. Allen et al., 2000) frequencies are contaminated by residual BCG and gradient artefacts, respectively. To isolate alpha band activity and exclude residual MR gradient and BCG artefacts, the data were bandpass filtered between 7 Hz and 13 Hz. Data were then average referenced, epoched by the TR (2 seconds) and downsampled to 40Hz. FastICA (Hyvarinen, 1999) was conducted to further isolate posterior alpha components. Only components with clear alpha-power peaks (7-13 Hz) and posterior scalp topographies were retained. Alpha power components were then retro-projected to channel space.

GLM Analysis Investigating Alpha-BOLD Relationships

A conventional alpha-BOLD GLM approach was conducted to compare with the spatial maps generated by the EEG HMM GLM. The alpha power of each TR epoch was calculated via Welch's method for electrodes Oz, O1/2, PO3/4 and POz. The average alpha power was calculated across these channels for each TR epoch, resulting in a continuous alpha time course at the same sampling rate as the fMRI data (Mayhew & Bagshaw, 2017; Mayhew, Hylands-White, et al., 2013; Mayhew, Ostwald, et al., 2013). The continuous alpha power time courses from each

run per participant were demeaned and normalised to control for differences in maximum alpha power between subjects. The alpha power time courses were then used as explanatory variables in a first-level GLM analysis using FSL FEAT (Woolrich et al., 2001), with contrasts for both the mean positive and negative activations. A second-level FEAT analysis was conducted to obtain the average EO and EC alpha activation maps per subject, followed by a group-level FEAT analysis (FLAME1, Z threshold = 3.1, cluster p threshold = 0.05) (Woolrich et al., 2004) to obtain group average maps for EO and EC, as well as EO>EC and EC>EO contrast map.

Alpha Power Analysis

The average change in alpha power between the EO and EC conditions was quantified and tested statistically using the alpha power time courses created for the GLM. The means of all EO and EC timepoints were calculated for each subject. Alpha power values were then compared statistically via a paired t-test.

Hidden Markov Modelling

The HMM-MAR toolbox (Vidaurre et al., 2016) was used to perform the HMM analyses. BOLD signal time courses from each network IC, obtained via dual regression, were concatenated across all runs and participants and used as the input to the fMRI HMM. To model changes in alpha power within the EEG, the Hilbert transform of the retro-projected alpha components from electrodes Oz, O1/2, PO3/4 and POz were concatenated and used as inputs for the EEG HMM. All inputs underwent amplitude standardisation (SD=1, mean=0). Each state was represented as a multivariate Gaussian distribution, defined by a mean vector and a covariance matrix. Therefore, each state within the fMRI HMM was defined by a vector containing the mean activity of each included network and a covariance matrix describing the FC between each network. Within the EEG HMM, the vector corresponded to the mean alpha power of each electrode, whilst the covariance matrix represented the temporal correlation of alpha power

between electrodes.

Model parameters were inferred based on variational Bayes and the minimisation of free energy (Vidaurre et al., 2016). The inference process determines both the likelihood of any given state being active at each timepoint (state probability time course) and the model parameters of each state. Inference was performed at the group level, resulting in common states being defined across all subjects, whilst the state probability time courses were defined at subject level.

To avoid local minima during the inference process and find the optimum number of states, the inference process was repeated five times for each potential number of states (within a range of 2-15), with each iteration running the process multiple times (fMRI = 10, EEG = 5) and retaining the run with the lowest variational free energy (50 repeats per potential state number). To determine the optimum number of states, the model selection process used a combination of repeat reliability (correlating the state time courses of the five repeats), the distribution of the maximum fractional occupancies (highest fractional occupancy across states) from each run (values closer to 1 are indicative that the model does not accurately capture the dynamics of the data) and the variational free energy values. The final HMM was selected based on a balance between high repeat reliability, low variational free energy, and low max fractional occupancy values. This process was undertaken independently for the EEG and fMRI data.

Once a final model was selected, multiple state metrics for each run were calculated to describe each state's temporal properties: fractional occupancies (% of run each state occupies), state lifetimes (mean duration of a state visit), state interval times (duration between visits of the same states) and the state switching rates. All state metrics were calculated from the Viterbi Path (Forney, 1973) of the state probability time course.

Estimating Spectra and Spatial Properties of EEG States

As the EEG states are defined on the Hilbert transform of the EEG data, the states are defined based on changes in power as opposed to other spectral properties (e.g. phase). Thus, when using wideband data, it is not immediately possible to determine whether a change in power was associated with a specific frequency band. In this instance, the data was bandpass filtered, and thus, it is clear that power changes reflect alpha power. However, it still does not provide the specific spectra of each state, e.g., it is not possible to distinguish between changes in low and high-frequency alpha power. Thus, the state spectra of each state are estimated using a state-wise weighted version of the multitaper approach in which periods of the data where a particular state was more likely to be active contribute more to the spectral estimation (Vidaurre et al., 2016). This provides a full spectral definition of each state.

To determine the spatial properties of each EEG state, the state probability time courses from the EEG HMM were used within a fMRI GLM analysis (Hunyadi et al., 2019). The time courses of each state were convolved with a double gamma function, and downsampled to the same sampling rate as the TR. A separate first-level GLM was run for each state per run, with contrasts for both the positive and negative mean activations. Second-level FEAT analyses generated average state activation maps for both EO and EC conditions at the subject-level, by grouping across runs, whilst third-level FEAT (FLAME1, Z threshold = 3.1, cluster p threshold = 0.05) analyses created a group average map for each condition, as well as contrast maps between the two conditions (EO>EC & EC<EO).

State Metric Analysis

State metrics were statistically compared between the EO and EC conditions using paired t-tests within both the fMRI HMM and EEG HMM analysis. Bonferroni correction was used to control for multiple comparisons. State lifetimes and interval times were not statistically

compared due to being dependent on the number of visits/intervals per condition, often resulting in unequal groups. To observe slower changes in state dynamics across a single run, the fractional occupancies were calculated within a 10s sliding window, resulting in a fractional occupancy time course (Baker et al., 2014). These time courses were averaged across subjects to compare temporal trends in state dynamics between the EC and EO conditions.

Results

Alpha Power Analysis

The alpha power analysis found an apparent increase in alpha power within the EC condition relative to the EO condition (see Figure 3.2A). Alpha power showed larger variability between subjects within the EC condition compared to the EO condition. The difference in alpha power was examined statistically using a paired t-test, which found the difference to be significant with a large effect size (difference = 1.54, 95% CI [0.65, 2.43], $t(17) = 3.66$, $p = 0.002$; Cohen's d = 0.86, Cohen's d 95% CI [0.31,1.40]). The mean alpha power time course further revealed that the largest increase in alpha power was observed during the first half of each run (see Figure 3.2B). Additionally, the initial increase in alpha power appeared larger within the first EC run. Conversely, the alpha power in the second EC remained consistently high throughout the run, unlike the first EC run, which eventually returned to a similar amplitude as observed within the EO runs.

Model Selection

The model selection process aimed to find the highest optimum state number whilst still possessing a relatively low variational free energy value, high inter-correlation between repeats and low maximum fractional occupancy values. For both the fMRI and EEG HMM, the variational

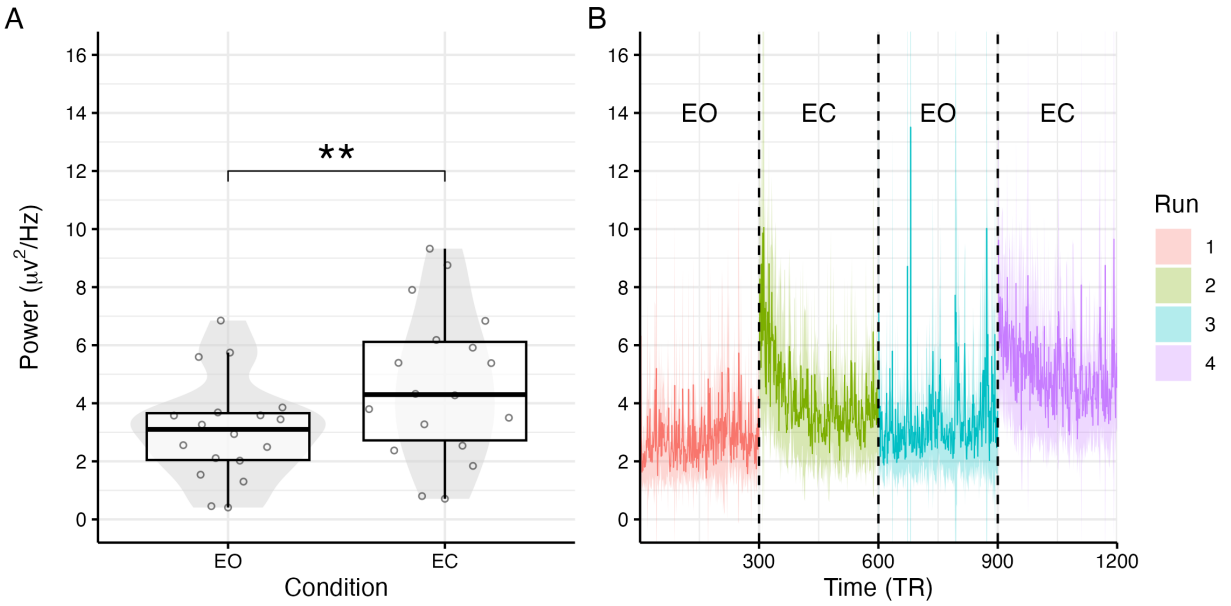


Figure 3.2: Alpha power compared between the EO and EC conditions. A). Comparison of the mean alpha power of each condition. ** = p<0.01. Each point represents a single subject. B) Group average alpha power time course at the sampling rate of the fMRI (TR=2s). Dashed vertical lines mark the boundaries of each run. Shaded regions represent 95% confidence intervals across subjects.

free energy decreased with the addition of each state (see Figures 4.6A and 4.6B). The mean reliabilities remained consistently high between states two to four within both models, with the reliability dropping at approximately six states in the fMRI model (see Figure 4.6C) and five states in the EEG model (see Figure 4.6D). Both state numbers showed high inter-repeat reliability, with all correlations exceeding 0.80 within the fMRI model (see Figure 4.6E) and 0.90 within the EEG model (see Figure 4.6F). Repeat three from the fMRI 6-state solution possessed the lowest maximum fractional occupancies values and was therefore selected for the analysis (see Figure 4.6G). Maximum fractional occupancies within the EEG 5-state solution showed minimal variation, and therefore, the first repeat was selected for analysis (see Figure 4.6H).

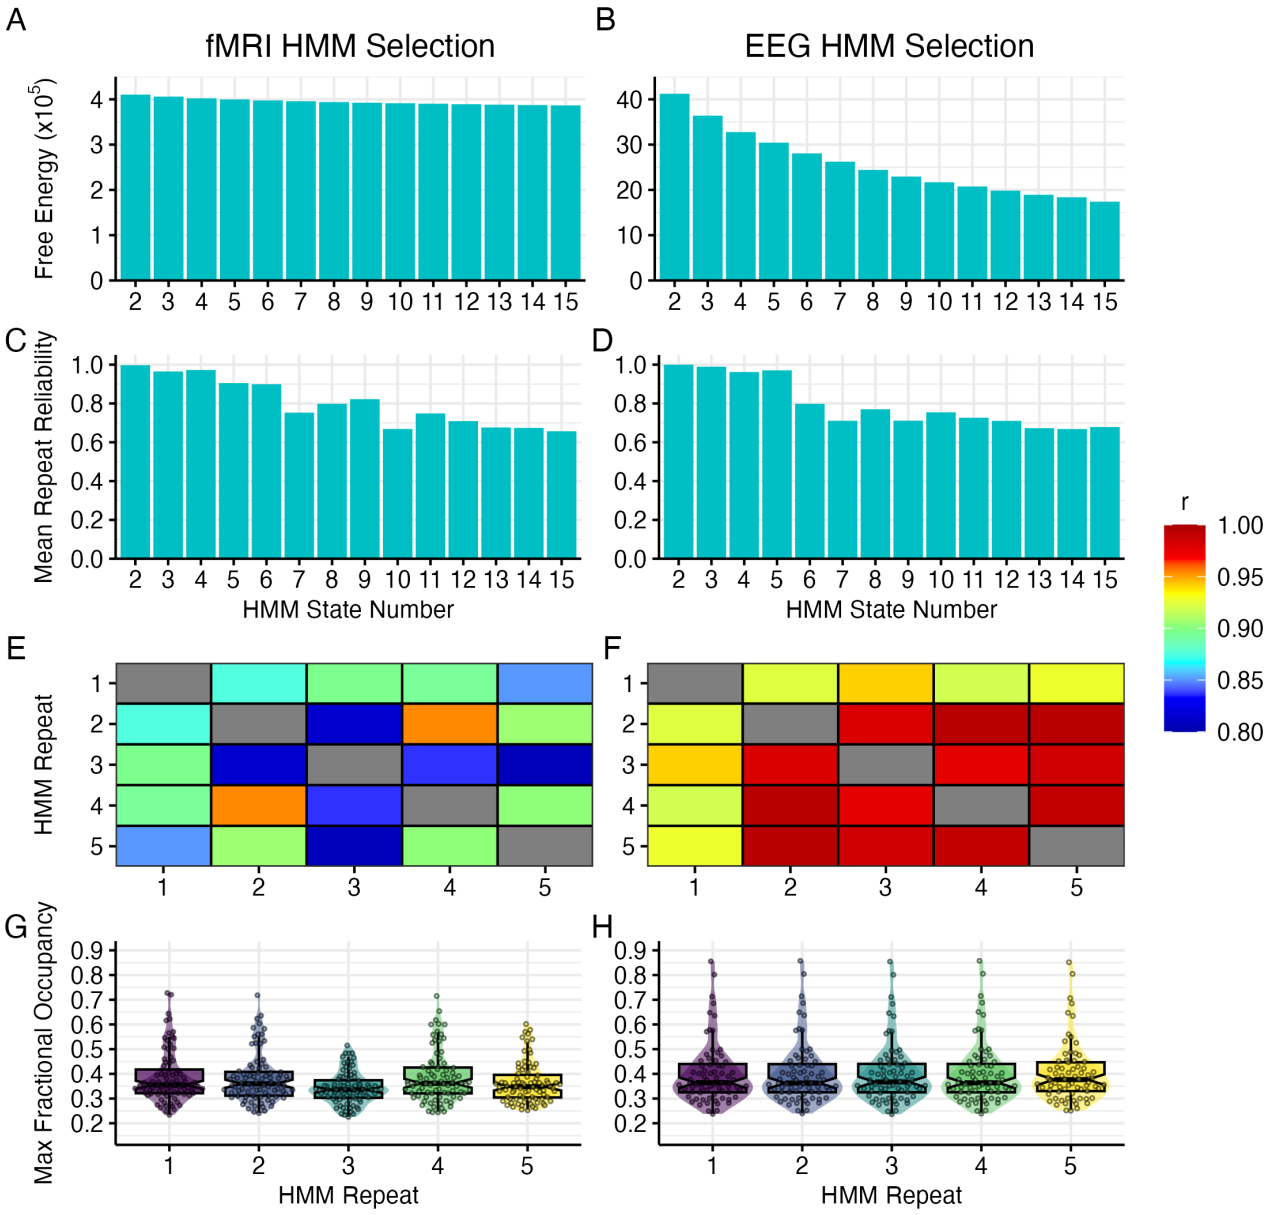


Figure 3.3: Model selection procedure for fMRI (left column) and EEG (right column) HMMs. A & B show the mean free energy for each state number. C & D show the mean correlation between repeats for each state number. E & F show the correlations between each repeat for state numbers 6 (left) and 5 (right). G & H show the max fractional occupancy distributions for each repeat for state numbers 6 (left) and 5 (right). Each point represents the max fractional occupancy from a single run.

HMM States

The fMRI HMM identified a total of six states (labelled F1-F6), defined by the mean activity across the inputted networks and the FC between each network (see Figure 3.4). Observing the states, F1 and F2 each showed large activations/deactivations within multiple sensory regions, including sensorimotor, auditory, and visual cortex. The FC between visual regions was notably higher within these states relative to other states. F3 showed large activations within the salience network and sensory regions, particularly auditory, as well as a deactivation within the DMN. F4's largest activations were predominately within the DMN. F3 showed positive FC between medial visual cortex FC and the DMN and precuneus network, whereas F4's FC pattern on average resembled the other states, with the lowest FC strength. F5 exhibited very low activations across all networks relative to other states, with the DMN, right frontoparietal, and DAN showing the largest de/activations. F6's largest activations were within a combination of the FPN and DAN, alongside deactivations within posterior somatosensory and auditory networks. Both F5 and F6 showed increased FC between the precuneus network and the DMN relative to other states.

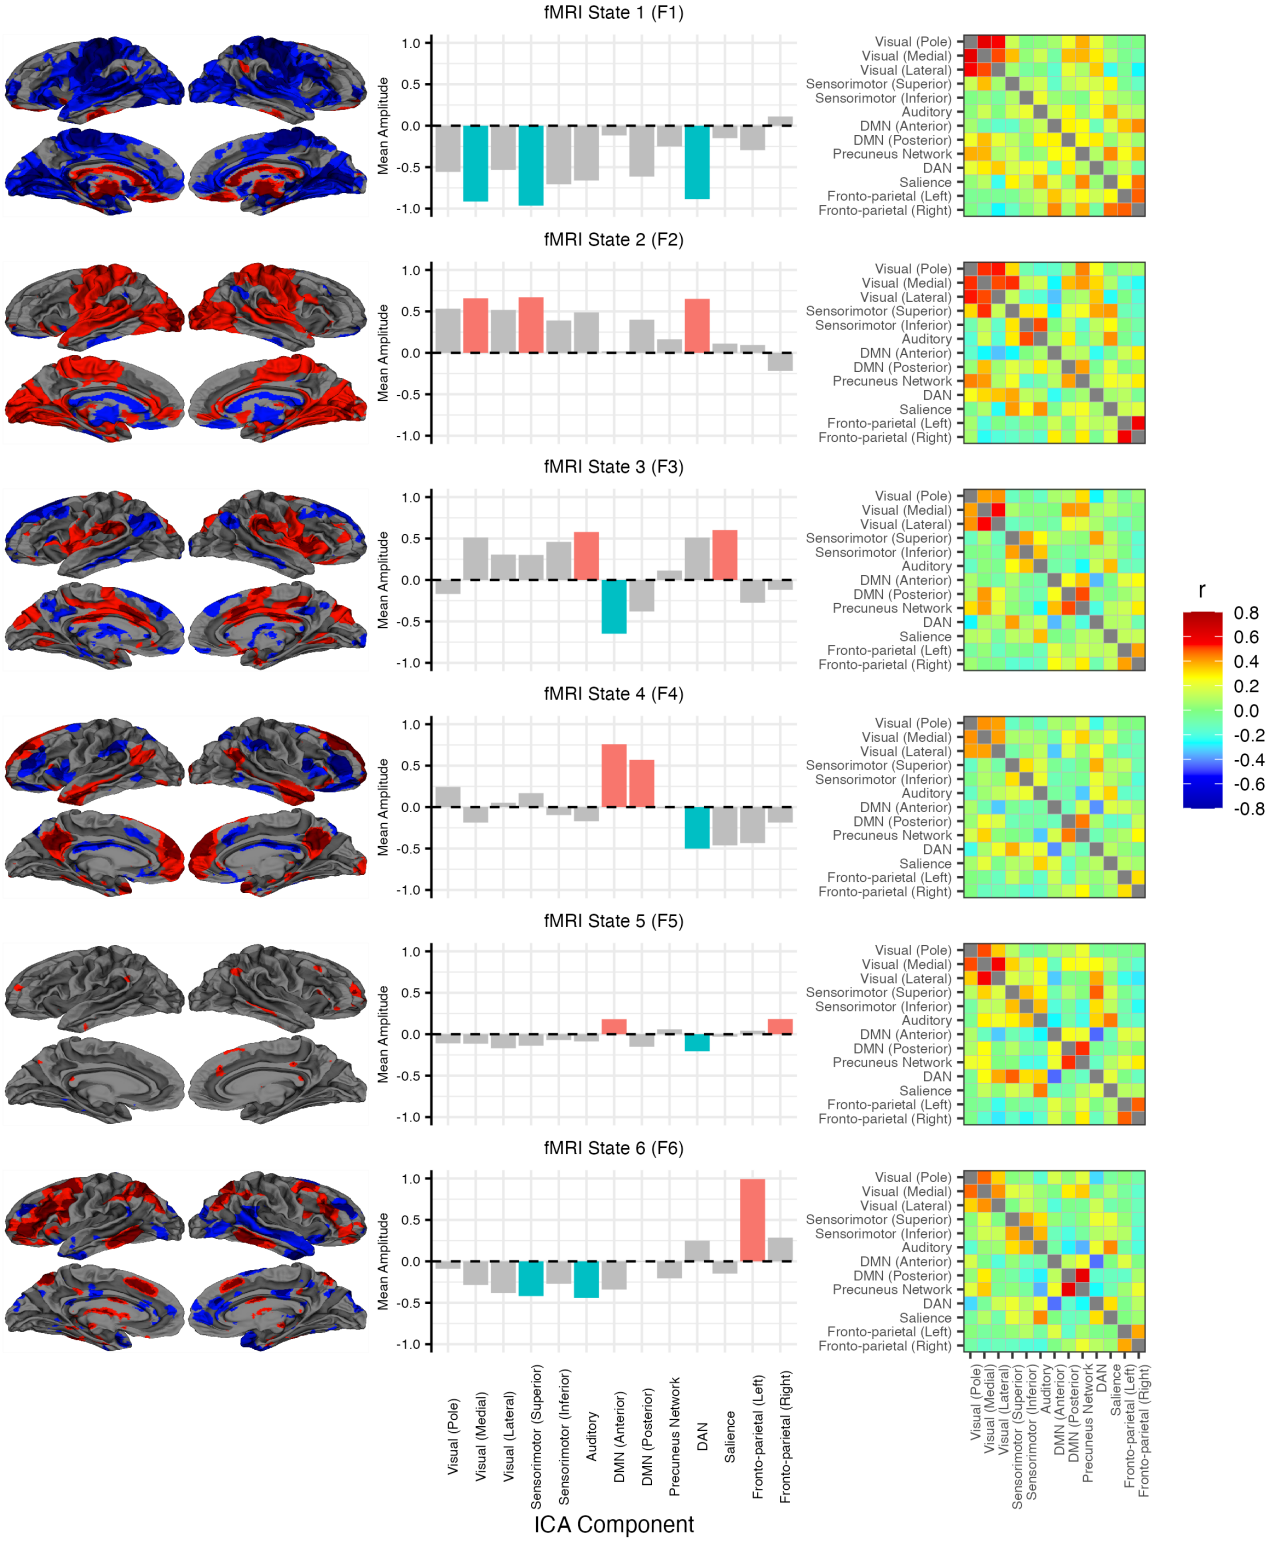

fMRI State 1 (F1)
fMRI State 2 (F2)
fMRI State 3 (F3)
fMRI State 4 (F4)
fMRI State 5 (F5)
fMRI State 6 (F6)
Mean Amplitude
1.0
0.5
0.0
-0.5
-1.0
Visual (Pole)
Visual (Medial)
Visual (Lateral)
Sensorimotor (Superior)
Sensorimotor (Inferior)
Auditory
DMN (Anterior)
DMN (Posterior)
Precuneus Network
DAN
Salience
Fronto-parietal (Left)
Fronto-parietal (Right)
ICA Component
r
0.8
0.6
0.4
0.2
0.0
-0.2
-0.4
-0.6
-0.8

Figure 3.4: fMRI HMM state Plots. Each row represents a different state identified by the HMM (F1 to F6). The left column shows the mean activation maps plotted on brain surfaces. The central column shows the corresponding bar chart for the mean activation maps, displaying the mean contribution of each IC to each HMM state. The coloured bars represent the three largest amplitudes (red: positive, blue: negative). The right column shows the estimated FC for each state.

Five states were inferred with the EEG HMM (E1-E5), defined by changes in posterior alpha (7-13Hz) power amplitude (Hilbert Transform) and alpha power correlations across the posterior electrodes (O1, O2, Oz, POz, PO3, PO4). State spectra were estimated using data from both conditions via a state-wise weighted multi-taper analysis for each channel and averaged. E5 exhibited the highest alpha power, followed sequentially by E4, E2, E3, and E1, which exhibited the lowest alpha power (see examples in Figure 3.5). Only states E1 and E5 showed significant correlations between their state probability time course and the BOLD signal (see Figure 3.5). Within the EC condition, E5 exhibited negative correlations across cortex, including visual cortex, parietal cortex, frontal cortex, and the DAN. In contrast, during the EO condition, E5 showed negative BOLD correlations exclusively within lateral visual cortex. Conversely, E1 displayed positive BOLD correlations with similar activation patterns to E5 during both EO and EC conditions. Unlike E5, the BOLD correlations during the EC condition were notably less widespread across occipital and parietal cortex.

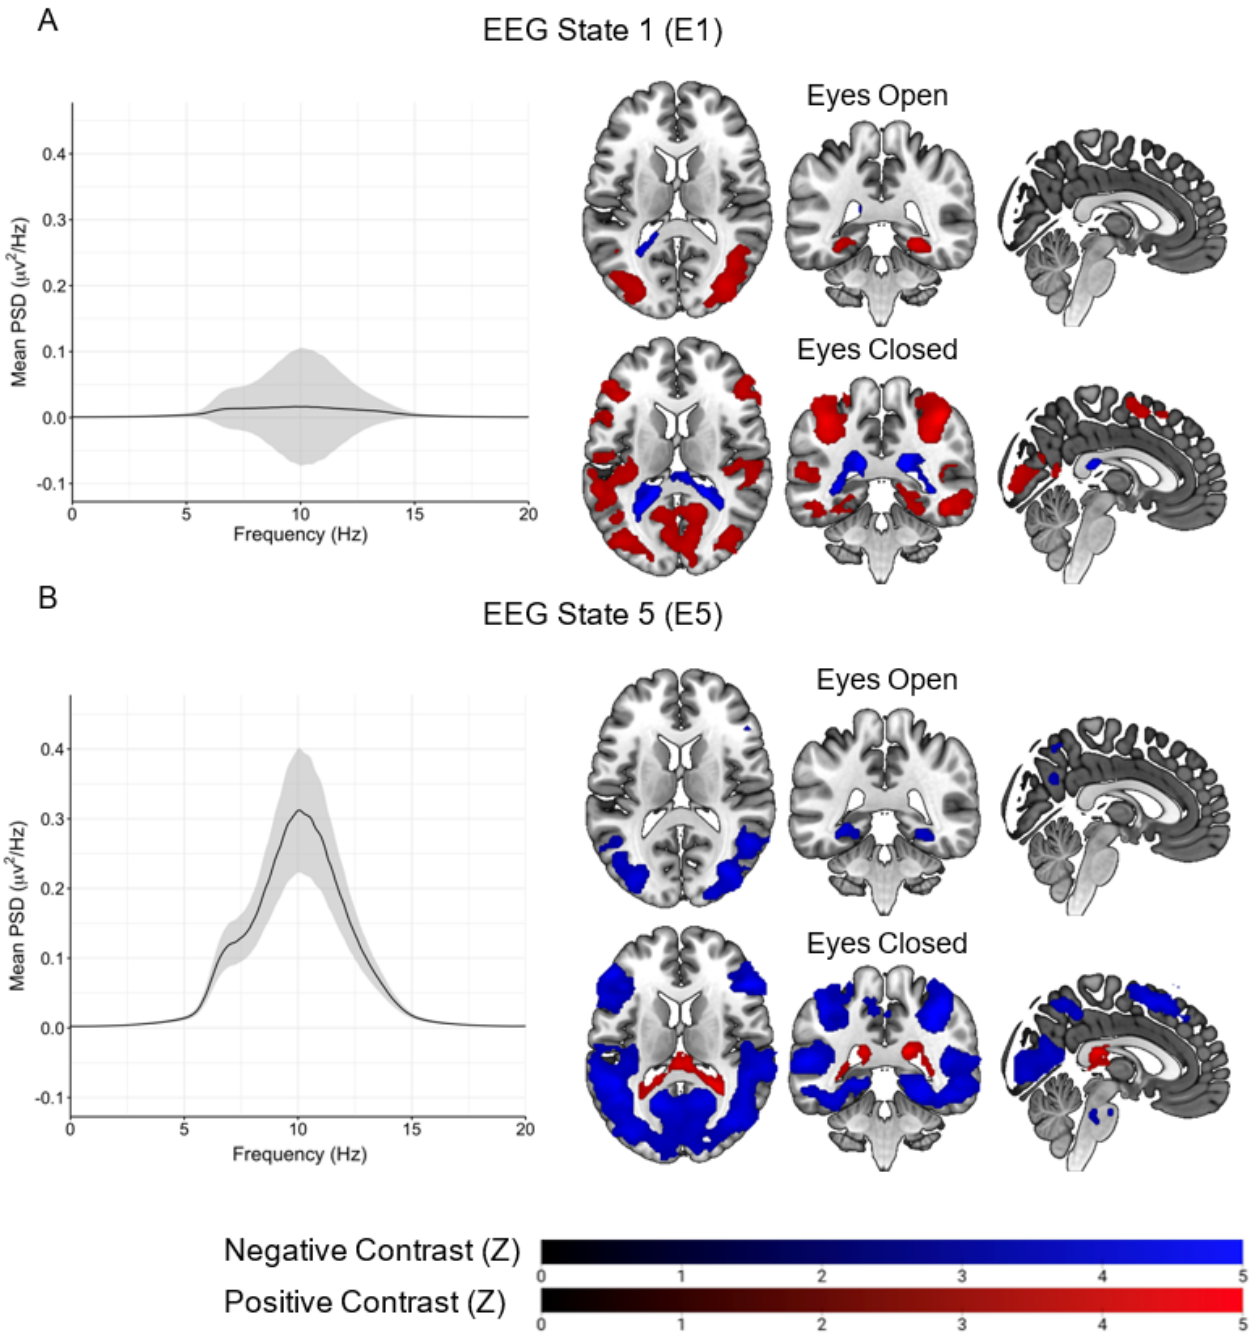


Figure 3.5: EEG HMM state spectral properties and voxel-wise correlations between state probability time courses and the BOLD signal. A) Power spectral density estimate for states E1 (A) and E5 (B). The shaded region represents 95% confidence intervals of the mean alpha power across the posterior channels inputted to the HMM (left). The right column shows the significant BOLD correlates for each of the EEG HMM states during both the eyes open and eyes closed rest conditions. Plotted in MNI152 space (coordinates = 111x96x117). Z threshold = 3.1, cluster p threshold = 0.05. Images are presented in MNI152 space using neurological conventions (left = left). PSD = power spectral density

Conventional Alpha-BOLD GLM

The conventional alpha-BOLD GLM analysis found similar patterns of BOLD activation as the HMM states (see Figure 3.6). During the EC condition, negative alpha-BOLD correlations were observed within primary visual cortex, frontal cortex and DAN. The correlations were notably less widespread compared to the E5 activation map but showed a large amount of overlap. Similarly, negative alpha-BOLD correlations were observed within lateral visual cortex during the EO condition, though notably less than observed for E5. However, unlike E5, positive alpha-BOLD correlations were observed within the insula bilaterally, anterior cingulate cortex, and cerebellum during the EO condition.

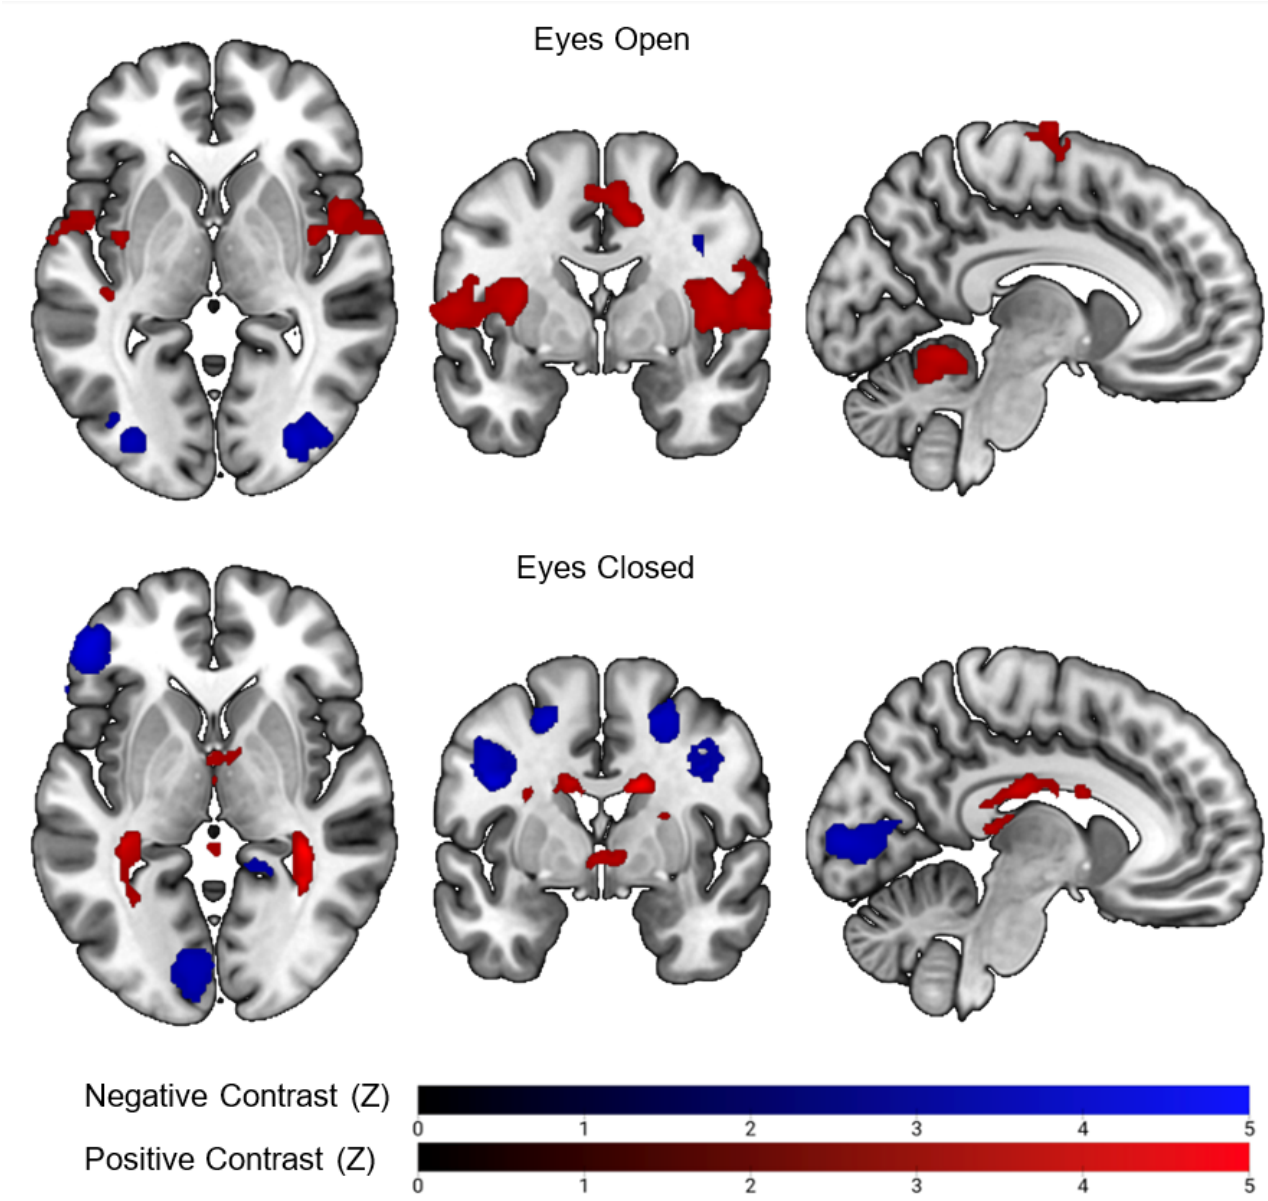


Figure 3.6: Group level statistical maps from the conventional alpha-BOLD correlation approach (coordinates: 92x149x97). Mean positive (red) and negative (blue) correlations between alpha power and BOLD signal are shown for both EO (upper) and EC (lower) rest conditions. Z threshold = 3.1, cluster p threshold = 0.05. Images are presented in MNI152 space using neurological conventions (left = left).

Fractional Occupancy Sliding Window Analysis

The sliding window analysis of the fractional occupancies revealed multiple unique patterns of state dynamics within the fMRI and EEG HMMs. Specifically, the fMRI sliding window analysis found that states F1 and F2 showed gradual increases in fractional occupancy upon eye closure, before decreasing to near zero upon eye opening (see Figure 3.7). States F3 and F4 both increased their occupancy during the EO conditions. F5 showed dramatic increases in fractional occupancy upon eye closure before gradually returning to a baseline level. It was also noted that fractional occupancy of F5 increased upon reopening of the eyes, but to a lesser magnitude than upon eye closure. State F6 exhibited the most static time course of occupancy, with no notable changes between conditions (see Figure 3.7).

Regarding the EEG sliding window analysis, states E1, E2, E4 and E5 all showed changes upon eye closure, with the former two states showing a rapid decrease in fractional occupancy upon eye closure, and the latter two states exhibiting the opposite pattern (see Figure 3.7). All four states then gradually returned to a baseline level. State E3 showed minimal variation throughout the session, except for a minor increase within the final run. State E4 showed a gradual increase across all four runs, with fractional occupancy increasing upon each eye closure.

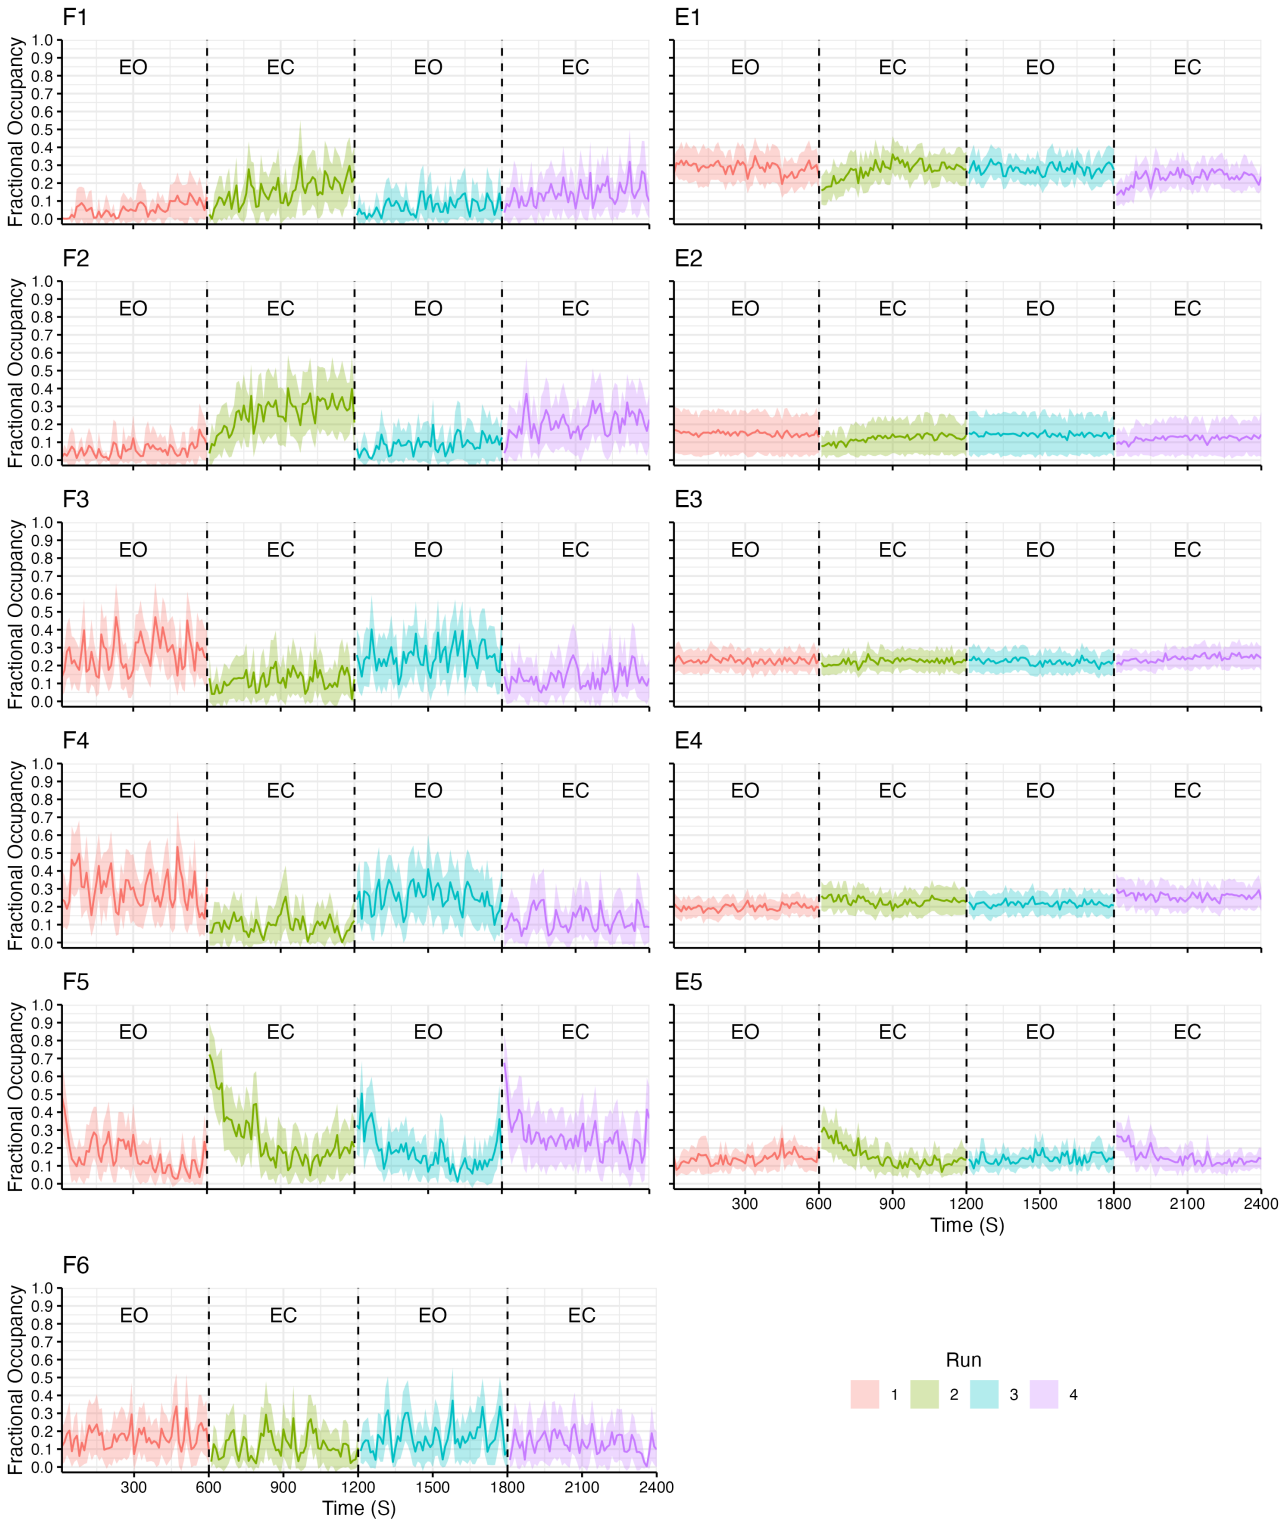

F1
E1
F2
E2
F3
E3
F4
E4
F5
E5
F6
EO
EC
Fractional Occupancy
0.0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1.0
300
600
900
1200
1500
1800
2100
2400
Time (S)
Run
1
2
3
4

Figure 3.7: Fractional occupancy timecourses spanning across all four runs (in order of acquisition EO,EC,EO,EC), calculated using a sliding window analysis (Window size =10s, overlap =0). The left column shows timecourses from the fMRI HMM, and the left column shows timecourses from the EEG HMM.

State Metrics Comparison

Statistical comparison of the fMRI HMM state metrics across EO/EC conditions revealed multiple significant differences (see Figure 3.8. Significant differences in state fractional occupancy between conditions were found for all states. F1, F2 and F5 had a significantly higher fractional occupancy during the EC condition, whilst F3 and F4 had a significantly higher fractional occupancy during the EO condition. State lifetimes also showed differences between conditions, with F1, F2 and F5 having longer state lifetimes during the EC than the EO condition, and F4 showing the opposite (see Figure 3.8C). All states had lower state interval times when in the conditions in which they had higher fractional occupancy (see Figure 8D). The state switching rate was significantly higher for the EO condition relative to the EC condition (difference = 0.01, 95% CI [3.81e-03, 0.02], t(20) = 2.90, p = 0.009; Cohen's d = 0.63, Cohen's d 95% CI [0.16, 1.10]) (see Figure 3.8B).

In comparison, only one significant difference in state metrics was found between the two conditions for the EEG HMM (see Figure 3.9). State E4 had a significantly higher fractional occupancy during the EC condition relative to the EO condition. No significant difference in switching rate was found between the two conditions (difference = -6.63e-04, 95% CI [-2.23e-03, 9.02e-04], t(17) = -0.89, p = 0.384; Cohen's d = -0.21, Cohen's d 95% CI[-0.67, 0.26]). State lifetimes were highest for states E1 and E5, and lowest for E4. Both states E4 and E5 showed a minor increase in state lifetimes when within the EC rest condition relative to the EO rest condition (see Figure 3.9C). Interval times also showed minimal variance between the two conditions, with E1 and E2 showing a minor reduction during the EO rest condition, and E3 showing

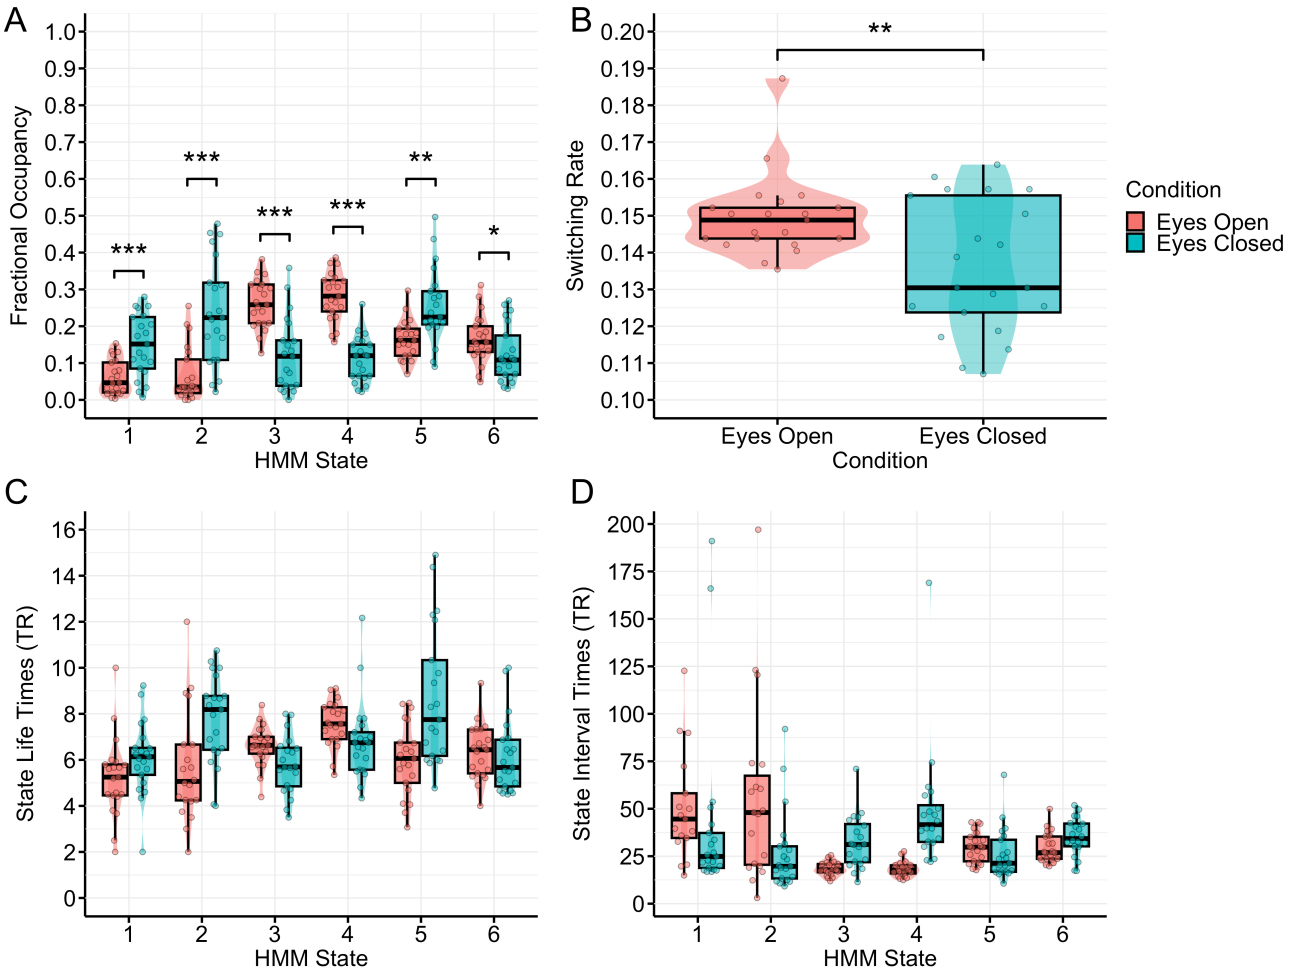


Figure 3.8: Comparison of fMRI HMM states metrics between the EO and EC conditions. Plot A compares max fractional occupancy. Plot B compares the state switching rate. Plot C compares the average state lifetimes. Plot D compares the average state interval times. All p values are Bonferroni corrected (number of comparisons = 6). *** = $p < 0.001$, ** = $p < 0.01$, * = $p < 0.05$. Each point represents a single subject.

a minor increase during the EO rest condition (see Figure 3.9D).

Based upon the results of the sliding window analysis (i.e., time courses in Figure 3.7), the state metrics were additionally compared solely within the first quarter of each run (2.5 mins), since various states demonstrated larger changes within that period. Within the EEG HMM, this revealed significant differences in the fractional occupancies of state E5 (difference = -0.09, 95% CI [-0.15, -0.03], t(17) = -3.32, p = 0.004, adjusted p = 0.020; Cohen's d = -0.78, Cohen's d 95%

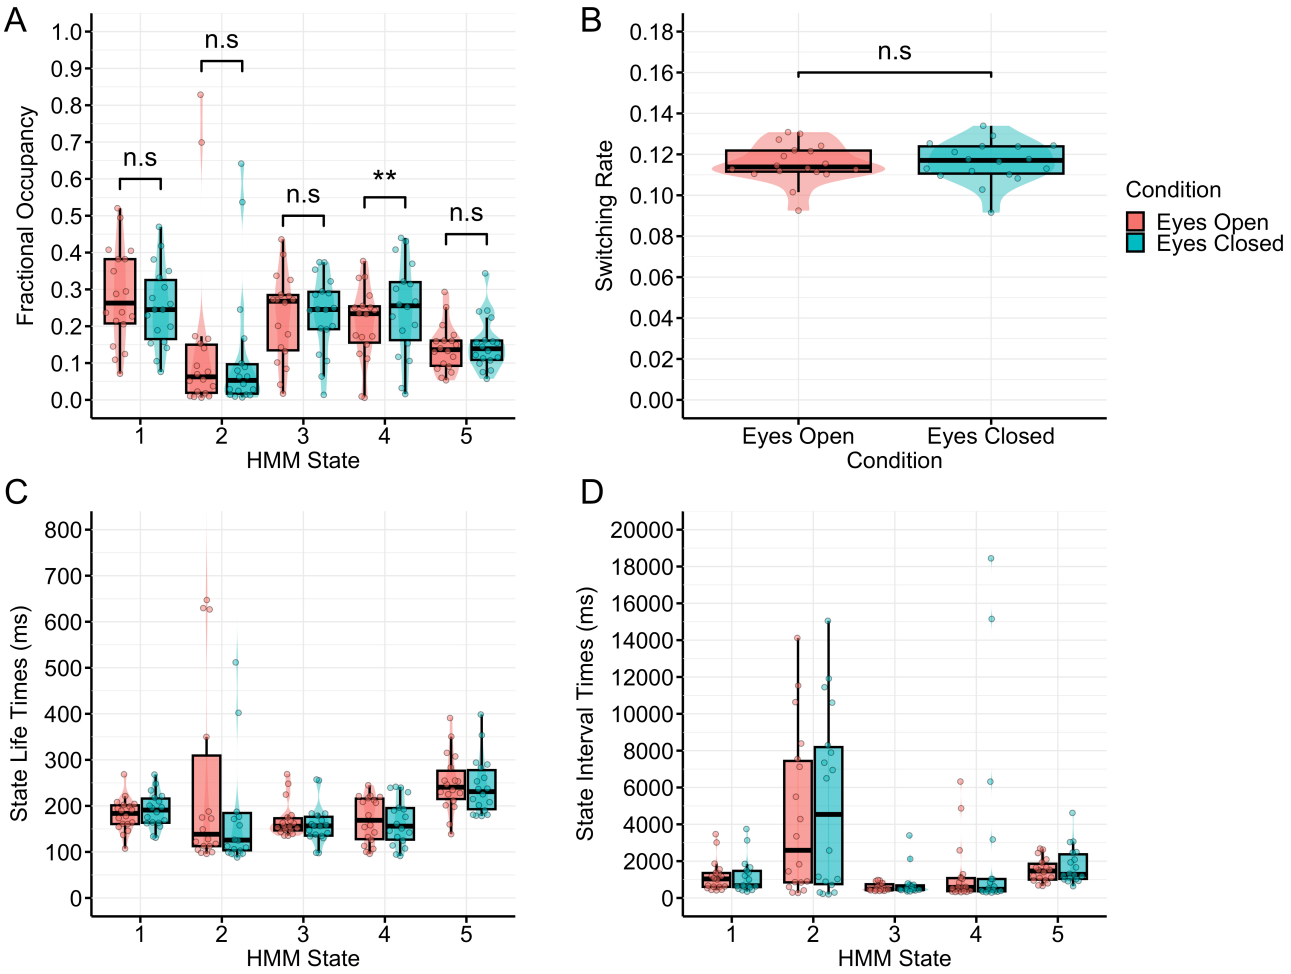


Figure 3.9: Comparison of EEG HMM states metrics between the active and passive conditions. Plot A compares max fractional occupancy. Plot B compares the state switching rate. Plot C compares the average state lifetimes. Plot D compares the average state interval times. All p values are Bonferroni corrected (number of comparisons = 5). ** = $p < 0.01$, n.s = non-significant. Each point represents a single subject.

CI [-1.30, -0.24]) between conditions. The switching rate did not differ significantly between conditions during the first quarter. Conversely, for the fMRI analysis, the significance of state fractional occupancy and switching rates did not change when restricted to the first quarter.

Discussion

Summary

This study applied Hidden Markov modelling to EEG-fMRI data to investigate changes in brain state dynamics during ten-minute runs of eyes open (EO) and eyes closed (EC) resting state. Hidden Markov Models were trained separately on each imaging modality and underwent model selection to determine the optimal number of states (fMRI=6, EEG=5). Comparison of the state metrics between the EO and EC conditions found significant differences in the fractional occupancy of all the fMRI states and one EEG state. Additionally, the fMRI state switching rate was significantly higher during the EO condition, whilst no significant difference was found for the EEG state switching rate. A sliding window analysis further revealed that specific fMRI and EEG states showed significant changes in fractional occupancy upon eye closure before returning to baseline level after approximately two and a half minutes. This pattern closely resembled the dynamics of alpha power at the channel level, in which alpha power increased during the onset of the eye closure condition before returning to a baseline. When state metric comparisons were restricted to the first two and half minutes of each run, EEG state E5 showed a significant increase in fractional occupancy during the EC condition relative to the EO condition. Interestingly, E5 was the EEG state associated with the highest alpha power and showed BOLD signal correlates closely resembling activations maps acquired by conventional alpha-BOLD analysis.

fMRI HMM States

The fMRI HMM identified six states, each with unique activation maps. Of these, F1, F2, and F5 showed significantly increased fractional occupancy during the EC rest, whilst states F3 and F4 showed decreased fractional occupancy during EC. Except for F3, all identified states have been

reported within either previous HMM research (Song et al., 2022) and/or complex PCA resting state research (Bolt et al., 2022), indicating that these states occur reliably. Furthermore, using PCA, the component equivalent to states F1 and F2 within this study was highly correlated with the global BOLD signal (Bolt et al., 2022). Though the origin and significance of the global BOLD signal has been heavily debated, particularly in regards to whether it is a representation of noise or physiologically plausible signal (Liu et al., 2017), animal literature has demonstrated that it correlates strongly with local field potentials, but only during periods of eye closure (M. L. Schölvinck et al., 2010). This suggests that the global BOLD signal exhibits different dynamics within EO and EC rest, a finding that has been reported within previous fMRI literature (Agcaoglu et al., 2019; Weng et al., 2020), with research demonstrating that BOLD signal amplitude shows significantly higher variability within the EC condition (Jao et al., 2013). Given our results, this would seem unsurprising, given that both states were defined by large changes in activity within sensory regions and showed greatly increased FC with eye closure. Furthermore, the reported increase in variability during the EC condition in the literature (Jao et al., 2013) matches the results found within this study, with the state fractional occupancy values exhibiting higher variability during the EC condition relative to the EO condition. As to what these states represent, previous research demonstrating increases in BOLD signal within sensory regions during eye closure has suggested that this represents mental imagery (Marx et al., 2004). Alternatively, given the states' potential association with the global BOLD signal, they could represent a reduction in vigilance, given previous EEG-fMRI research finding correlates between the global BOLD signal and vigilance measures (Wong et al., 2016). It has also been suggested that EC resting state runs contain periods of sleep (Tagliazucchi & Laufs, 2014), and future work will be needed to understand the impact of sleep on HMM states and their properties.

Regarding the remaining states, the interpretation is less complex. State F4 solely consisted of the activity within the DMN and showed a higher fractional occupancy during the EO resting

condition, as previously reported (Jao et al., 2013; Y. Wang et al., 2022; C. Yan et al., 2009). This potentially occurs as a result of the removal of visual stimuli when the eyes are closed, as both posterior cingulate cortex and medial prefrontal cortex have been linked to sensory monitoring (Brandman et al., 2021; Raichle et al., 2001). State F3 on the other hand, primarily consisted of large activations within the salience and auditory networks, with the same pattern of fractional occupancy as state F4. As the salience network is primarily associated with the detection of sensory stimuli (Downar et al., 2000), it can be hypothesised that the state was also more active during EO rest as a result of being able to perceive external visual stimuli. Lastly, state F4 exhibited minimal activation across all networks relative to the other states. However, the sliding window analysis revealed that F5 was the only fMRI state to show an increase in fractional occupancy at eye closure. Additionally, the state also showed a similar increase in functional connectivity between visual regions as in states F1 and F2. Thus, the state appears to represent the initial increase in FC within visual cortex upon eye closure. Together, this suggests that the fMRI HMM was able to accurately identify the changes in dynamics between the two rest conditions.

EEG HMM States

The EEG HMM identified five states with varying alpha power, with E5 exhibiting the highest alpha power, followed by states E4, E2, E3 and E1. Using the state time courses from each state as regressors within a fMRI GLM analysis, significant BOLD correlates were identified for the highest (E5) and lowest (E1) alpha power states in both EO and EC conditions. These activation maps closely resembled the activation maps obtained from the conventional alpha-BOLD approach. Specifically, during the EC condition, state E5 showed widespread negative correlations across occipital and parietal cortex, as well within the DAN. Conversely, state E5 showed bilateral negative correlations within lateral occipital cortex during the EO condition. This suggests that the HMM approach can identify states whose time course of fractional oc-

cupancy has meaningful BOLD correlates, resulting in brain states with both high temporal and spatial properties. Though neither E1 nor E5 showed a significant difference in fractional occupancy across the two conditions, the sliding window analysis revealed that E5 showed a significant spike in fractional occupancy during the first quarter (2.5 minutes) of each EC run before gradually returning to a baseline level. This pattern in fractional occupancy appeared unique to the EEG HMM states, with fMRI state fractional occupancy remaining stationary across the entire duration of each run (except for F4). It is hypothesised that this is due to the fMRI states representing slow global metabolic changes over time, whereas the EEG states solely represent more dynamic changes in alpha power. This is supported by the fact that E5's fractional occupancy time course closely resembled the changes in alpha power observed in the alpha power time courses, suggesting that state E5 specifically represents the initial change in alpha power induced by eye closure. Furthermore, when restricting the fractional occupancy to the first quarter of each run, state E5 showed a significant increase in fractional occupancy during the EC condition, providing further support for this theory. Another interesting feature of E5 is that it possessed a notably lower fractional occupancy relative to the other states, being active 15% of each run, yet still exhibited BOLD correlations that closely resembled the results of conventional approaches. Furthermore, both states E1 and E5 possessed notably lower fractional occupancies relative to the other states, being active 15% of each run, yet still showed BOLD correlates closely resembled the results of conventional approaches. This suggests that BOLD correlates obtained from conventional EEG-fMRI approaches are predominantly driven by short periods of high alpha power, or 'bursts', as opposed to a sustained alpha rhythm that spans across the duration of an entire run (Quinn et al., 2019). A similar conclusion has also been drawn from MEG data, with HMMs identifying beta 'bursts' during a motor task (Seedat et al., 2020).

Limitations and Future Work

This study had several limitations. Firstly, the study did not employ counterbalancing, with all participants completing the runs in the same order. This makes it difficult to discuss any links to vigilance as it is not possible to rule out potential order effects, and therefore, counterbalancing of run order should be employed in future research. Secondly, the EEG analysis was restricted to occipital alpha power. This analysis was chosen due to a priori knowledge of how occipital alpha power changes upon eye closure; however, it only serves as a limited definition of a brain state. In future work, a more complex observation model, such as a Time Delay Embedded model (TDE) (Vidaurre, Hunt, et al., 2018) or Multivariate Autoregressive model (MAR) (Vidaurre et al., 2016) should be employed to directly infer states according to their spectral properties, including both power and phase. However, it still needs to be determined how well these observation models will perform on simultaneous EEG-fMRI data, given the restricted bandwidth available due to the residual harmonics of the gradient and BCG artefacts. It is also acknowledged that the HMM approach itself is not free of limitations. The primary limitation being that it is reliant on the Markov assumption, assuming that the probability of being in any one state at a timepoint is determined solely based on the state active at the timepoint prior. Additionally, the HMM approach requires an extensive model selection process to ensure both the reliability and validity of the model. Despite these limitations, the results here suggest that the HMM approach can detect dynamic changes in both the BOLD signal and alpha power that are in accordance with the literature.

Conclusions

This study identified unique spatiotemporal dynamics during periods of EO and EC rest in concurrent EEG-fMRI recordings. Both the fMRI and EEG Hidden Markov models were able to identify spatiotemporal dynamics observed within previous literature, as well as identifying

novel brain state properties. Specifically, the fMRI HMM found that brain state dynamics during EO rest were primarily associated with vigilance and the monitoring of external stimuli, whilst the EC dynamics were primarily associated with an increase in FC and activity within sensory regions. Conversely, the EEG HMM results demonstrated that upon eye closure, a spike in alpha power lasting approximately two minutes occurred before decreasing to baseline. This suggests that alpha-BOLD correlations are primarily driven by short periods of high alpha power, as opposed to a sustained rhythm. However, it is noted that the observation model used with the HMM likely provides an oversimplification of the dynamics, given that it can only model changes in power within a predefined frequency band, and future research should utilise more advanced observation models, such as the TDE (Vidaurre, Hunt, et al., 2018) or MAR (Vidaurre et al., 2016) model. Overall, Hidden Markov Modelling can identify unique spatiotemporal dynamics in simultaneous EEG-fMRI data, and serves as a powerful tool for researching brain state dynamics within and across modalities.

Chapter 4

Whole-Brain BOLD States Modulate BOLD Responses and P100 Amplitudes

Abstract

Functional magnetic resonance imaging is currently the most commonly applied neuroimaging technique in humans, yet the variability in the BOLD response remains poorly understood. This study investigated the impact of whole-brain baseline BOLD states, defined by Hidden Markov modelling, on the visually evoked BOLD and electrophysiological responses during active and passive conditions. Baseline BOLD HMM state active significantly impacted the amplitude of both the P100 VEP component and evoked BOLD responses, with states associated with higher baseline activity in visual cortex resulting in smaller amplitudes and vice versa. Furthermore, trials preceded by the DMN state showed larger BOLD response amplitudes than other BOLD HMM states. Overall, this study demonstrates that the baseline BOLD activity across the whole brain can modulate the amplitude of visually evoked BOLD responses and VEP components, highlighting the importance of accounting for the baseline state in neuroimaging research.

Introduction

Functional magnetic resonance imaging (fMRI) is one of the most widely applied methods for non-invasively investigating human brain function. Unlike other neuroimaging techniques, fMRI

can measure changes in the BOLD signal across the entirety of the human brain at millimetre resolution. Thus, with fMRI, it is possible to investigate the functions of deeper brain structures (e.g. the thalamus) and to identify large-scale networks (Biswal et al., 1995; S. M. Smith et al., 2009). Despite the method's wide application, research has demonstrated that the BOLD signal exhibits a large amount of variability within individuals, between sessions (Aguirre et al., 1998) and within sessions, with the BOLD response varying to repeated presentations of the same stimulus (Bianciardi et al., 2009; Duann et al., 2002; M. Schölvinck et al., 2012). Research has demonstrated that a proportion of this variability in BOLD responses results from the ongoing BOLD dynamics within the region, with the baseline activity becoming linearly summated with the evoked BOLD activity (Fox & Raichle, 2007; Fox et al., 2006). For example, using a finger-tapping task, Fox and Raichle (2007) demonstrated that the variability in evoked BOLD responses within left somatosensory cortex could be reduced by subtracting the ongoing spontaneous activity from right somatosensory cortex, indicating that the evoked BOLD signal and spontaneous BOLD signals summate linearly (Fox & Raichle, 2007). However, this approach only focussed on the spontaneous BOLD signal within the sensory cortex associated with the stimulus, neglecting the impact of BOLD signals within more distal brain regions. Recent research in rodents has demonstrated that spontaneous BOLD activity across the whole brain (including areas that do not respond to the stimulus) explained approximately 30% of the variability in visual BOLD responses, indicating that the spontaneous activity across the whole brain also explains a large amount of the variance in visually evoked responses (Q. Zhang et al., 2022).

Therefore, given the link between whole-brain BOLD activity and visual response variability, this chapter investigated the impact of whole-brain states, defined by Hidden Markov Modelling, on evoked BOLD, VEP, and alpha ERD/S responses during active and passive conditions. Additionally, this chapter investigated the differences in the whole brain state dynamics between the active and passive conditions and whether such changes in dynamics could explain the

differences in evoked response amplitudes. Firstly, a GLM and ROI analysis were conducted to compare the spatial activation maps and the shape of the BOLD responses between the active and passive conditions. An ICA was then performed to identify resting state / intrinsic BOLD networks, of which the corresponding time series were then inputted into an HMM to identify whole-brain states. Trials were then grouped according to which state was active during the baseline, and the average evoked responses were calculated and compared. The state metrics were also compared across the two conditions to investigate whether the task impacted the ongoing brain state dynamics.

Methods

Dataset

This chapter used the active vs. passive viewing EEG-fMRI dataset collected during my PhD. See Chapter 2 for details on stimuli and tasks.

Participants

A total of 30 participants were collected within the active vs passive EEG-fMRI dataset (see Chapter 2 for complete details on participant demographics and ethics). Two participants were excluded from the EEG data in this chapter: one for excessive movement and one for electrode bridging, resulting in a total of 28 participants being included in all subsequent EEG analyses. Similarly, the participant with excessive movement was also excluded from the fMRI data, resulting in a total of 29 participants being included in all subsequent fMRI analyses.

EEG Preprocessing

Preprocessing consisted of removing gradient artefacts, ballistocardiogram artefacts and movement artefacts via carbon wire loop regression. Bad channels were identified and removed from the data before undergoing average referencing, and ICA to remove any remaining artefacts, such as blinks or muscle-related artefacts (see Chapter 2 for full description and implementation details). Two copies of the EEG data were then created: one bandpass filtered to the alpha band (7Hz - 13 Hz) to be inputted into the HMM and to generate the alpha ERD/S analysis., and one bandpass filtered at 1-25Hz for the VEP analysis (see Chapter 2 for complete details on VEP and alpha ERD/S extraction).

fMRI Preprocessing

fMRI preprocessing consisted of susceptibility correction, motion correction, normalisation to standard space and ICA AROMA (Pruim et al., 2015) via fMRIPrep (Esteban et al., 2017). A custom Nipype (Gorgolewski et al., 2016) script was then employed to remove the noise components identified by ICA AROMA, apply a spatial filter of 4mm using FSL's SUSAN, and apply a high pass filter (100s). See Chapter 2 for complete details on fMRI preprocessing.

Group fMRI ICA Decomposition

A group spatial ICA was performed to decompose the fMRI into resting state networks using FSL's MELODIC function (Beckmann & Smith, 2004) via Nipype (Gorgolewski et al., 2011). Whole-brain decompositions were calculated with 20, 25, and 30 dimensions. These dimensions were selected to isolate networks as individual independent components. The 20-dimension solution provided the clearest networks and was selected for subsequent analysis. ICs were labelled by visual inspection, using the FIND lab ICA parcellation (Shirer et al., 2012) as a reference. Only ICs that had time courses in lower frequencies and did not spatially overlap

with known artefacts (motion, vascular, physiological, susceptibility) or non-grey matter regions (white matter, CSF) were used in the subsequent analysis, resulting in 11 components being retained (see Figure 4.1). The remaining ICs consisted of common networks, including visual, sensorimotor, salience, frontoparietal, dorsal attention and default mode networks (see Figure 4.1). FSL's dualregression function (Beckmann et al., 2009; Nickerson et al., 2017) was then used to obtain the IC time series at the subject level.

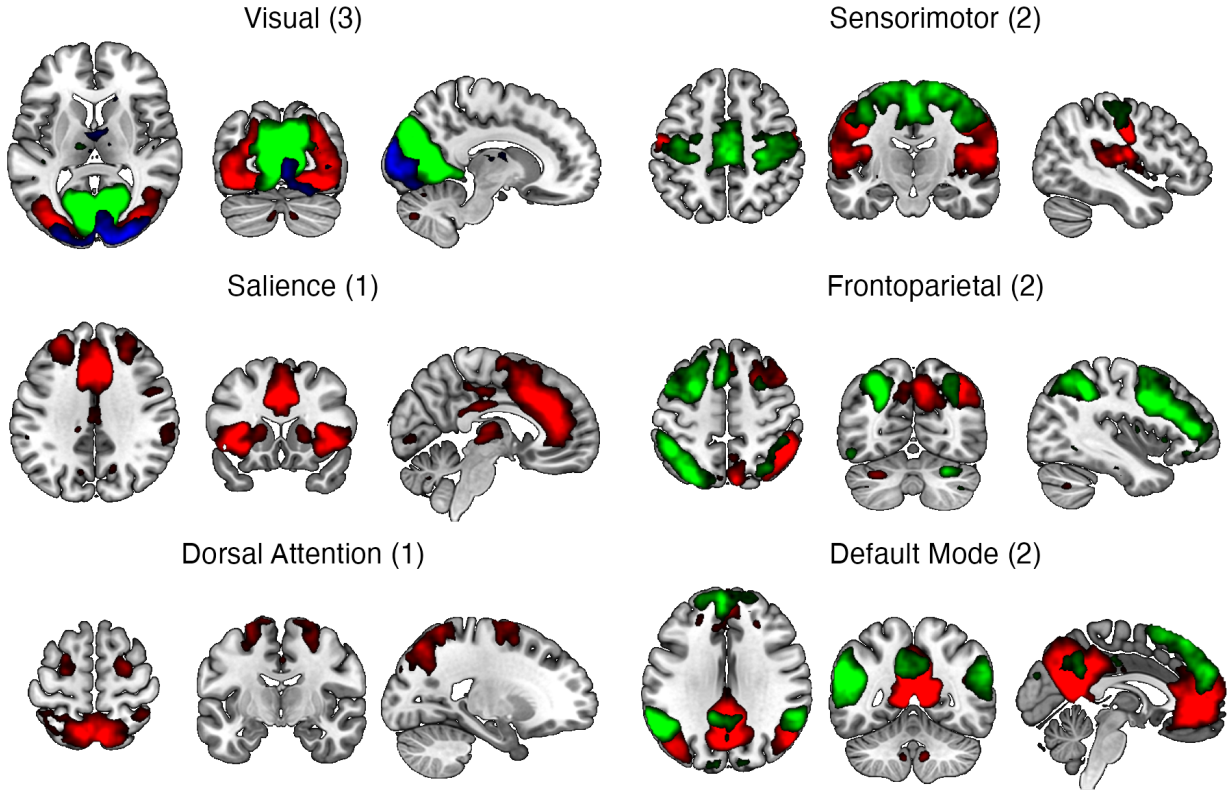


Figure 4.1: Results from the group fMRI decomposition. Independent components are grouped into networks. Each colour represents a separate independent component. The number next to each subtitle represents the number of independent components identified for that region/network. Images are presented in MNI152 space using neurological conventions (left = left).

General Linear Modelling

To investigate the impact of the condition on the stimulus-induced BOLD responses, a GLM analysis was conducted to calculate the spatial activation maps for each condition. The analysis

was performed in Python with the Nilearn package. First-level analyses were conducted at the subject level (T1 space), using the standard and target stimuli onsets as explanatory variables, as well as the temporal derivative of each variable. All explanatory variables were convolved using the SPM haemodynamic response function. In addition to the two explanatory variables, cerebral spinal fluid (CSF) and white matter (WM) regressors generated by fMRIprep were included within the first-level design matrix as regressors of no interest, alongside five polynomial drift regressors to high pass filter the model at 0.01 Hz (100s). Separate first-level GLMs were fitted for each condition, resulting in two first-level models per subject (active and passive). Four contrasts were used within the first-level analysis to calculate both explanatory variables' positive and negative mean effects. The resulting first-level activation maps were transformed to standard space using the transform to MNI152 space calculated in fMRIPrep via the antsApplyWarp function.

The first-level maps were then inputted into second-level group analyses to generate group average maps for each contrast. The second-level design matrices consisted of four regressors for each condition and stimulus type, respectively, and a regressor for each subject to model the within-subject variability, allowing for paired comparisons. Four contrasts were calculated for each model: the mean active activation map, the mean passive activation map, the mean activation map across both conditions (one sample t-tests) and a comparison between the active and passive conditions (active > passive) (paired samples t-tests). Positive and negative maps were calculated for each contrast. Contrast activation maps were then masked using the MNI152_2mm_brain_mask from FSL before being thresholded using Bonferroni correction ($p < 0.05$)

Defining ROIs

Regions of interest (ROIs) were determined from the activation maps generated within the GLM analysis (see Figure 4.2). A total of two ROIs were selected: right v1 and insula. These regions were selected as they showed the strongest activations in the activation maps generated by the GLM analysis. The right V1 ROIs were extracted for each subject individually, using the first-level mean positive activation maps for standard stimuli during the passive condition. The right insula ROI was extracted from the active > passive second-level (group level) activation map. The fslstats function was used to identify the coordinates of the peak response voxel within each activation map. The fslmaths function was then used to generate a 3x3x3 cubic, binarised ROI (27 voxels), centered on these coordinates. The right insula roi was then transformed back to each subject's T1 space using the transforms calculated in fMRIPrep via the antsApplyWarp function.

Extracting BOLD Responses

BOLD responses were extracted from each ROI and converted to percent signal change (see Chapter 2 for details).

Extracting VEPs and Alpha ERD/S

VEP and alpha ERD/S responses were extracted from the right posterior electrode that showed the strongest P100 component and baseline corrected (see Chapter 2 for details).

Hidden Markov Modelling

The Hidden Markov model (HMM) analysis was performed using the HMM-MAR toolbox (Vidaurre et al., 2016). The subject-level IC time-courses were temporally concatenated and used as the input for the model. Inputs were standardised so that the standard deviation was equal

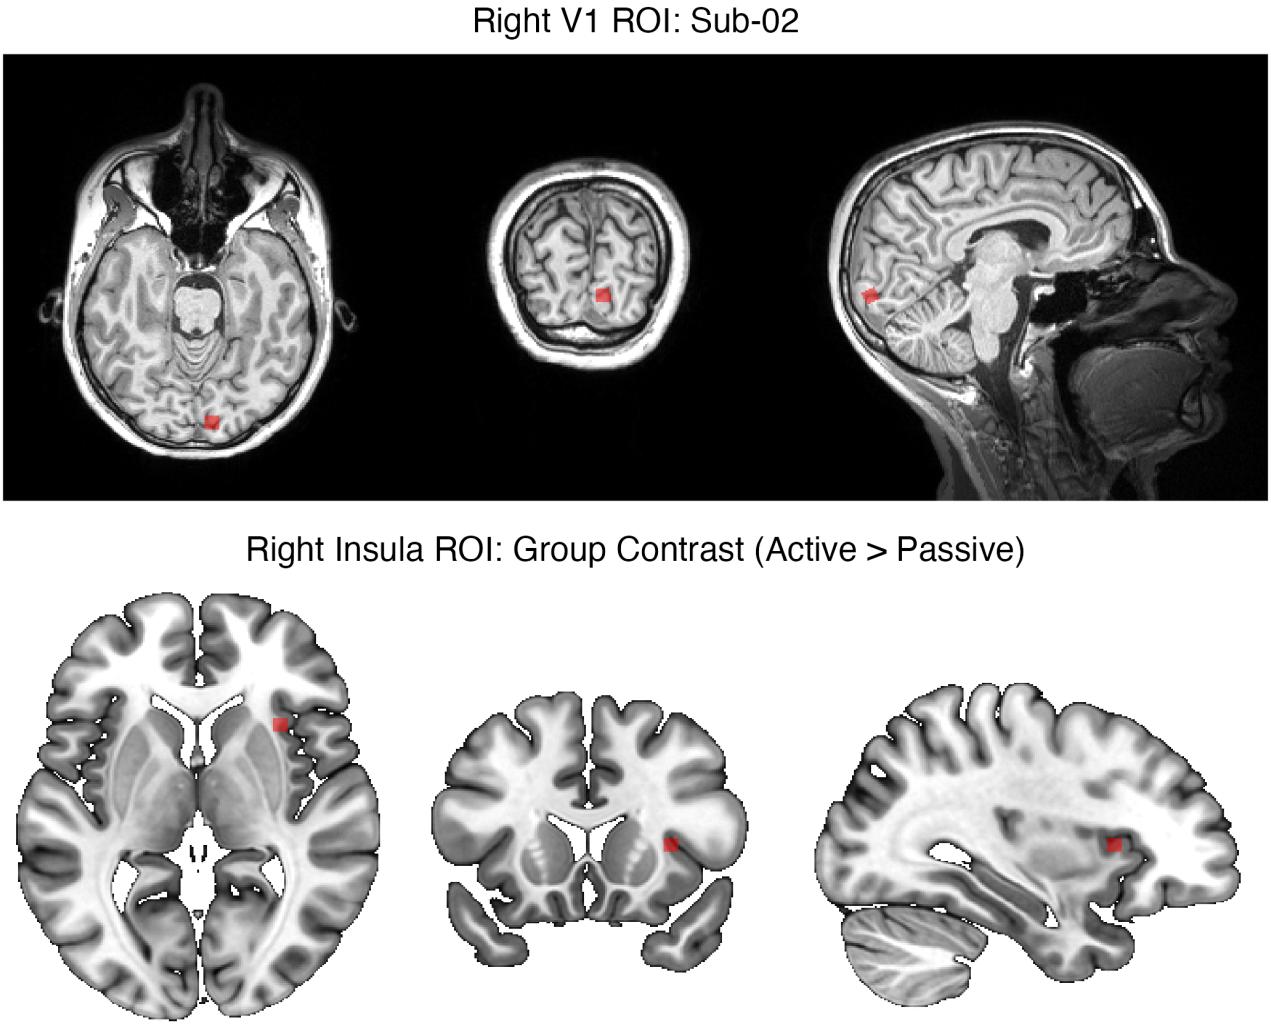


Figure 4.2: The top panel shows an example of a right V1 ROI generated from the sub-02's first-level mean stimulus (positive) contrast. The bottom panel shows the right insula ROI generated from the second level active > passive contrast plotted in MNI152 space. Both images use neurological convention (left = left). Right insula ROI MNI coordinates: 33,22,-3mm.

to one and the mean was equal to zero. A multivariate Gaussian distribution was selected for the observation model, resulting in states being defined by the mean activity of each IC and the covariance between ICs (analogous to functional connectivity).

To avoid local minima during the inference process and find the optimal number of states, the inference process was repeated a total of ten times for each potential number of states. Models

were inferred with state numbers ranging from 2-15. The final HMM was selected based on a balance between high repeat reliability, low variational free energy, and low max fractional occupancy values. See Chapter 2 for an overview). State metrics were calculated using the selected HMM to compare the state dynamics between the active and passive conditions. All state metrics were calculated using the Viterbi path (Forney, 1973).

Evoked Fractional Occupancy

The evoked fractional occupancy responses were calculated to evaluate which states became more or less active in response to the visual stimuli. The responses were calculated by epoching the state likelihood time series around the stimulus onset and baseline corrected using the same window and baseline correction method used for the BOLD response analysis. Average responses were calculated for each subject, which were then averaged to generate a grand average.

Statistics

Statistical testing was conducted to compare the effect of condition on fractional occupancy and switching rate via paired t-tests (see Chapter 2 for details). Additionally, the impact of baseline HMM state, condition, and stimulus type on BOLD, VEP, and alpha ERD/S responses was evaluated via linear mixed modelling (see Chapter 2 for details).

Results

General Linear Modelling

The GLM analysis found significant activations within contralateral (right) visual cortex in response to the stimuli within both the first-level (subject level) and second-level (group level) activation maps (see Figure). Furthermore, significant activity was found within the thalamus

for both the active and passive contrasts, with the activity being more widespread within the active condition, and spatially constrained to the LGN within the passive condition. During the active condition, activations were also observed within the insula, intraparietal sulci, frontal eye fields, putamen, anterior cingulate cortex and cerebellum bilaterally. Activations during the active condition were notably more bilateral within both visual cortex and the thalamus relative to the passive condition, with the passive condition activations being constrained to contralateral visual cortex, and the small areas of ipsilateral visual cortex.

The negative contrasts found significant deactivations during the active condition within the middle frontal gyri, precuneus, middle temporal gyri and angular gyri bilaterally, as well as lateral and middle orbitofrontal cortex. No significant deactivations were observed in response to stimuli during the passive condition. The active>passive contrast found that activations within ipsilateral visual cortex, anterior contralateral visual cortex, the insula, anterior cingulate cortex, and the cerebellum were significantly larger during the active condition relative to the passive condition. Additionally, both the putamen and anterior thalamus showed significantly larger activations bilaterally. Lastly, deactivations within the middle frontal gyri, precuneus, middle temporal gyri and angular gyri bilaterally, lateral and middle orbitofrontal cortex were also significantly larger during the active condition relative to the passive condition.

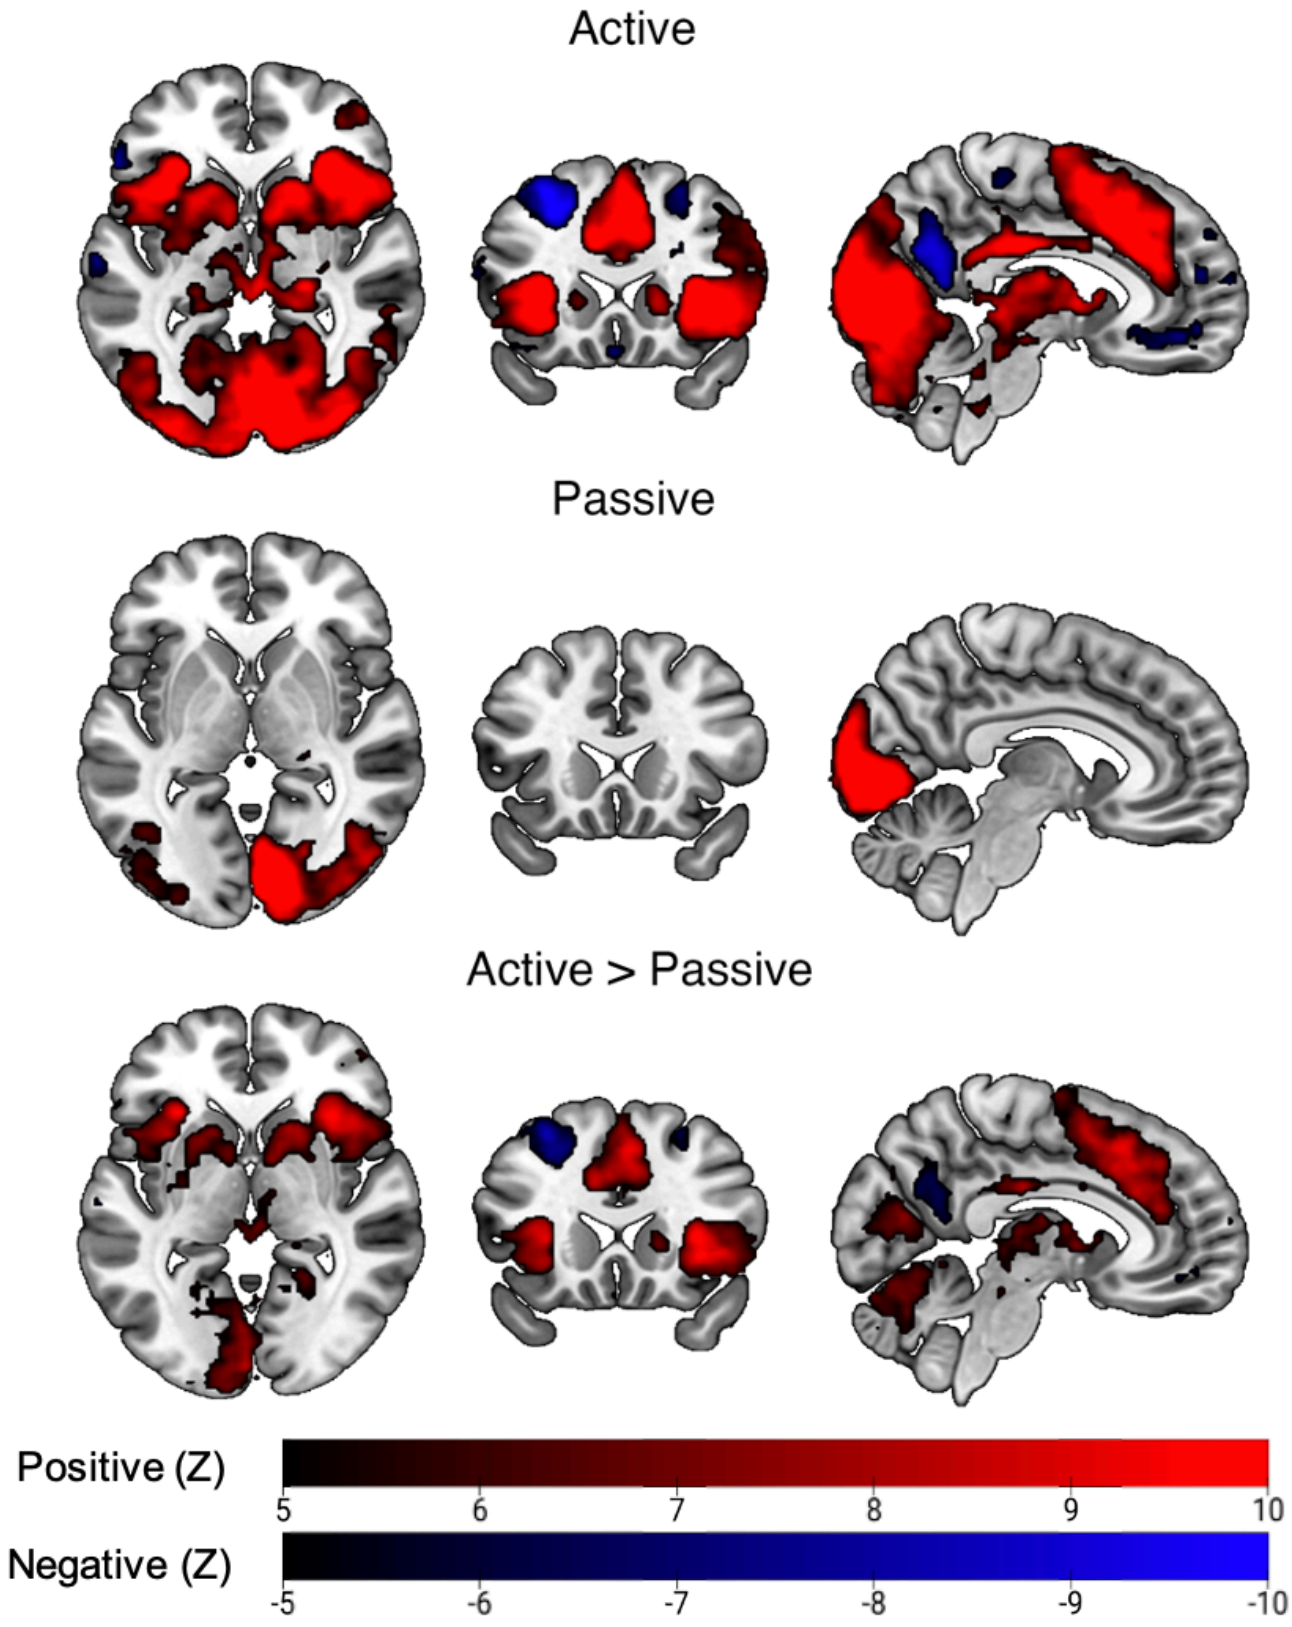

Active
Passive
Active > Passive
Positive (Z)
5
6
7
8
9
10
Negative (Z)
-5
-6
-7
-8
-9
-10

Figure 4.3: Activation maps for each contrast from the GLM analysis. Each row shows the positive and negative mean activation maps for the active, passive and active>passive contrast. All maps are Bonferroni corrected ($p < 0.05$). Images are presented in MNI152 space (x=6, y=20, z=0) using neurological conventions (left = left).

ROI Analysis

To investigate the impact of condition (active vs. passive) and stimulus type (standard vs. target) on the shape/profile of the BOLD response, an ROI analysis was conducted within contralateral V1 and insula (see Figure 4.4). This showed large visually evoked BOLD responses within V1 during both conditions, with the group mean peak signal change occurring at 5 TR (5.05 s), followed by an undershoot between approximately 11-16 TR (11.11s-16.16 s), before returning to baseline. BOLD signal change peaked at 2.6% during the active condition and at 2.3% during the passive condition.

For the insula, the group mean peak signal change also occurred at 5 TR (5.05 s) after stimulus onset, before returning to baseline at 11 TR (11.11 s) without undershoot. Signal change peaked at 0.5% within the active condition and 0.19% within the passive condition.

To determine the impact of stimulus type on the shape of the evoked BOLD response, comparisons were conducted for both ROIs, grouped by both conditions, active condition only and passive condition only (see Figure 4.5). Within right V1, when grouping across both conditions, the standard and target stimuli generated indistinguishable evoked BOLD responses, peaking at 5 TR, with a minor undershoot starting at 11 TR, before returning to baseline (Figure 5). Target stimuli signal change peaked at 2.5%, whilst standard stimuli peaked at 2.4%. Grouping by condition, the responses demonstrated variations in both peak signal change and undershoot. During the active condition, target stimuli showed a larger peak signal change of 2.7% relative

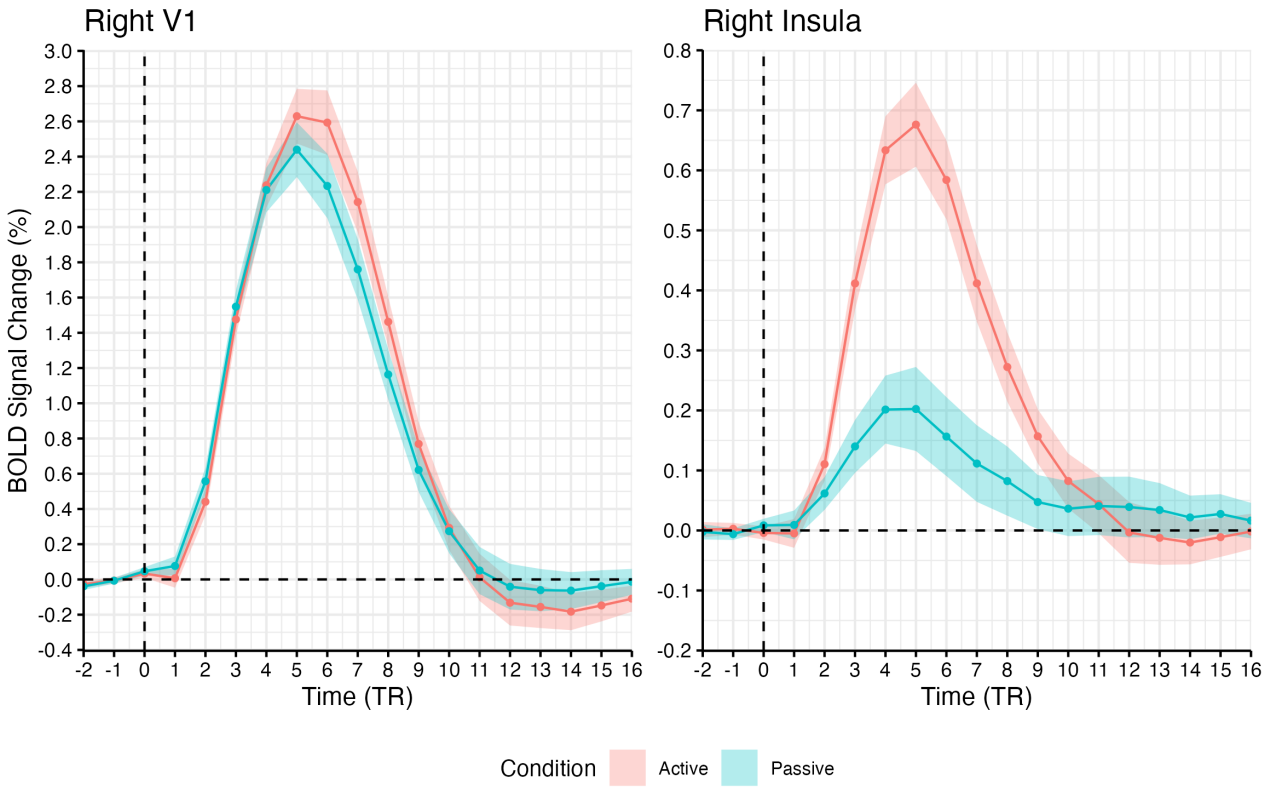


Figure 4.4: Comparison of the group mean BOLD responses within contralateral V1 between conditions for contralateral V1 (left plot) and insula (right plot). Shaded areas represent 95% within-subject confidence intervals.

to the 2.5% in response to standard stimuli. Further, during the active condition, the standard stimuli evoked an undershoot, whilst responses to target stimuli returned to baseline. During the passive condition, no difference in signal change was observed between stimulus types at any point in the response, with both peaking at 2.3% signal change. However, the opposite of the active condition was observed during the undershoot period, with target stimuli generating an undershoot at 11 TR, whilst responses to standard stimuli returned to baseline.

The equivalent analysis was conducted within the right Insula (see Figure 5). Across all three groupings, evoked BOLD responses peaked at 5 TR and returned to baseline without undershoot. Grouping across both conditions, target stimuli generated larger peaks in signal change, with

target responses peaking at 0.4% and standard responses peaking at 0.3%. Similarly, during the active condition, target stimuli generated peaks in signal change of 0.6%, whilst standard responses peaked at 0.5%. During the passive condition, minimal differences in peak signal change were observed between stimulus types, with stimuli generating a peak of 0.2% and standard stimuli generating a peak of 0.15%.

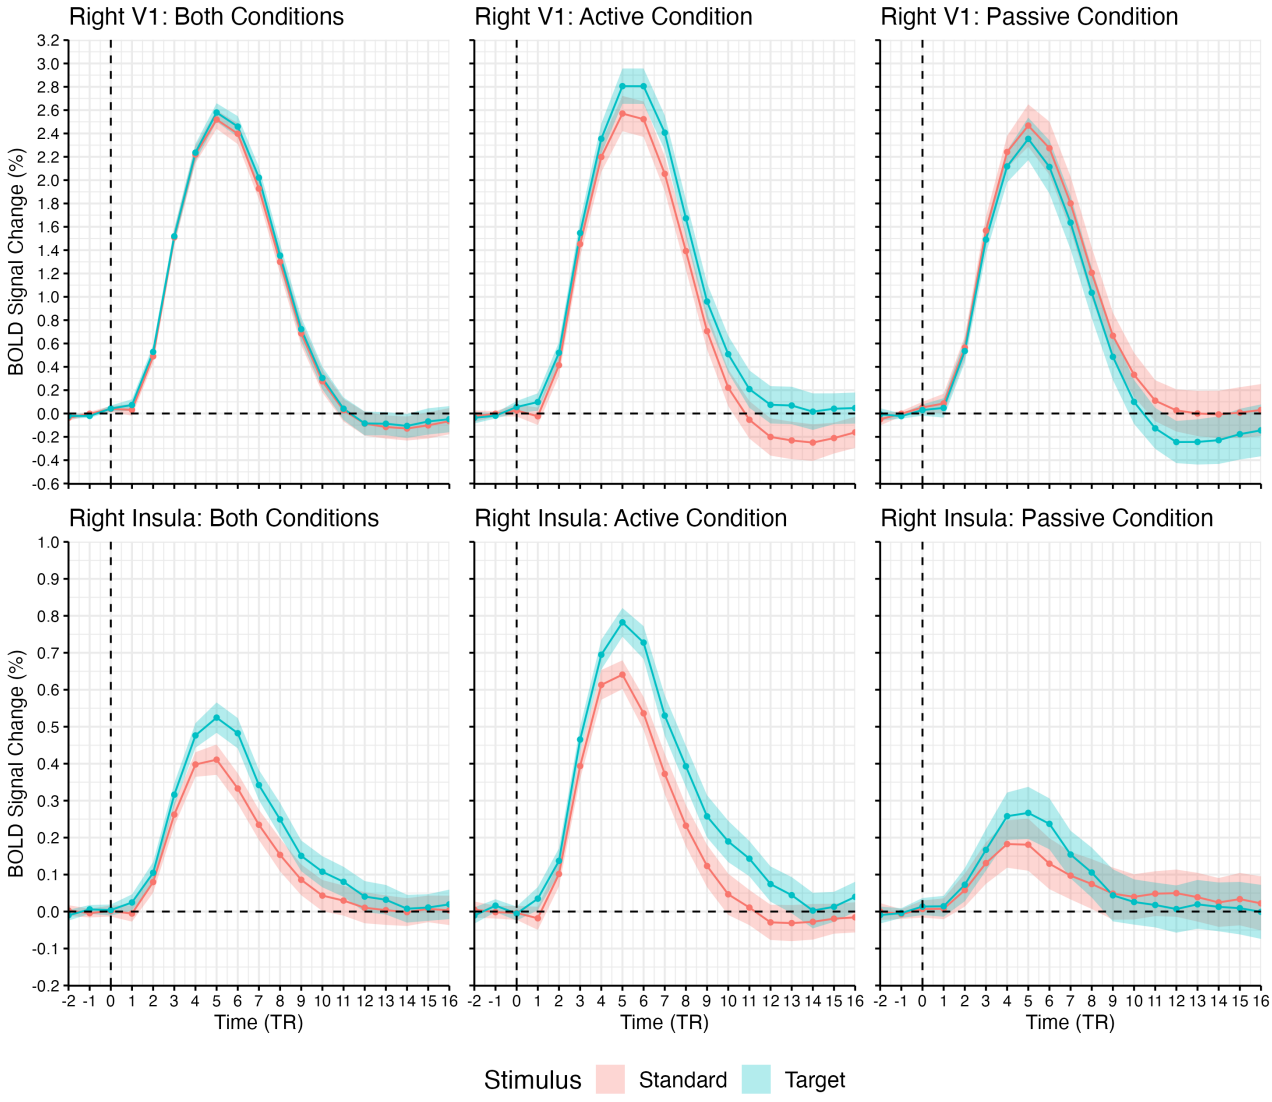


Figure 4.5: The top row of plots compares the V1 haemodynamic responses to the two stimulus types (standard and target) grouped across both conditions, the active, and the passive conditions. The bottom row show of plots shows the same comparisons but for haemodynamic responses within the right Insula. Error represents 95% within-subject confidence intervals.

HMM Model Selection

The model selection process was conducted to find the model with the optimum number of states. A combination of free energy, repeat reliability, and maximum fractional occupancy measures was used to determine the optimal number of states (see Figure 4.6). Free energy decreased with the addition of each new state, making it difficult to determine the optimal state number from this measure alone. Mean repeat reliability showed that the HMM remained reliable up to a 6-state solution ($r > 0.8$). Based on this measure, a five-state solution was chosen for further selection. Though the six-state solution also resulted in a reliable HMM, the five-state solution was selected to reduce the splitting of trials in subsequent analysis. Observing the repeat reliability matrix, all repeats of the model showed strong correlations to each other, with the exception of repeat 9, which had notably weaker correlations with the other repeats. Despite this, all repeats showed an even distribution of max fractional occupancy values, with a median of approximately 0.30. Therefore, repeat one of the five-state solution was chosen for subsequent analysis

HMM States

Five states (W1-W5) were defined by the HMM according to the mean activity and covariance of the inputted dual regression time courses (see Figure). Regarding the mean activity of each state, W1 was characterised by a large increase in activity within both DMN components and a reduction in activity within the dorsal attention network (DAN). W2 was characterised by a large reduction in activity across all three visual networks, inferior sensorimotor network, and the DMN / language network. W3 was characterised by large increases in activity within the salience network and right fronto-parietal network, and a reduction in activity within the DMN. W4 was characterised by a large increase in activity across all visual networks, the inferior sensorimotor network, DMN/ language network and the DAN, as well as a large reduction in

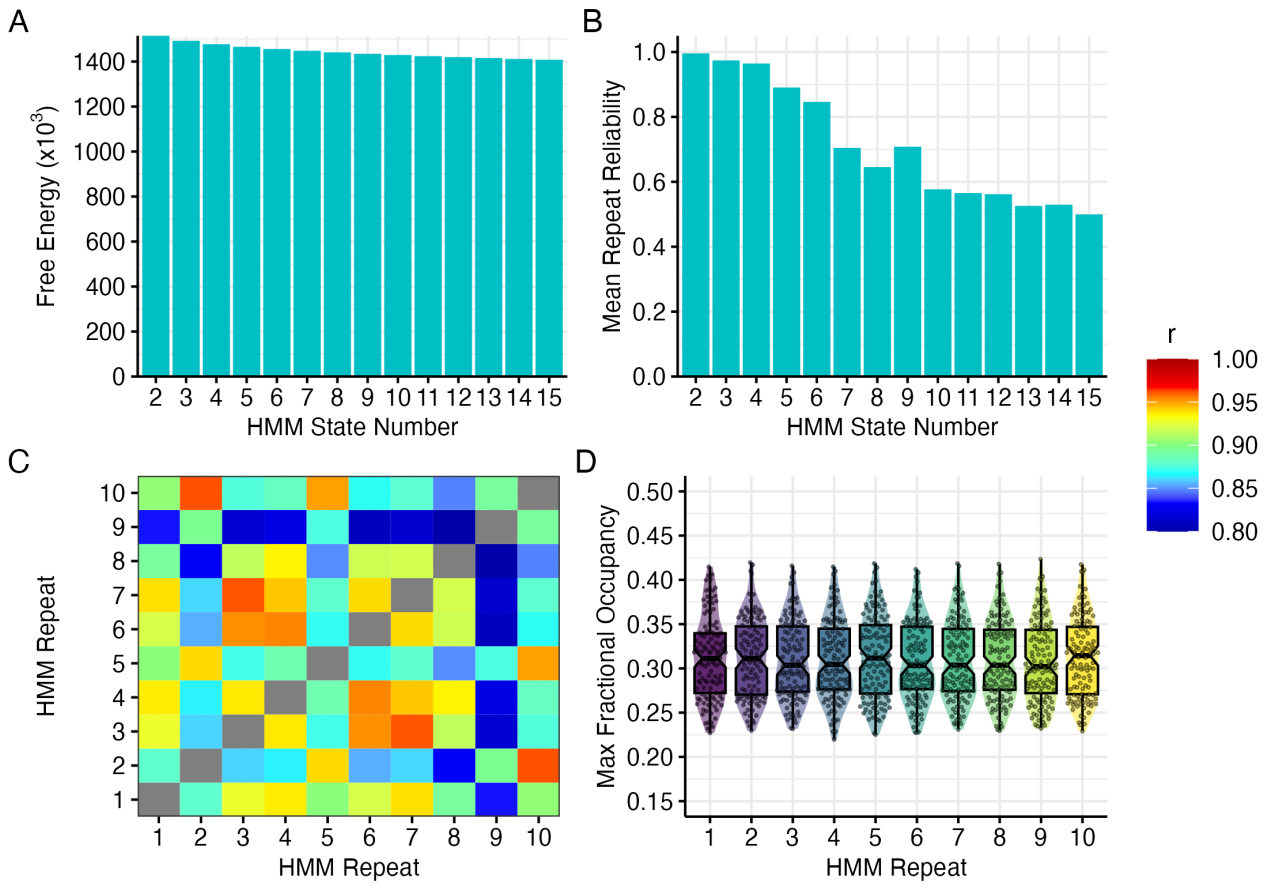


Figure 4.6: The HMM model selection results. A). The mean variational free energy across all 10 repeats for each state number. B). The mean repeat reliability across all combinations of the 10 repeats for each state number. C). Correlation matrix showing the reliability values across all repeat combinations for the five-state HMM model. D). The distribution of the fractional occupancy variables of each block for each repeat of the five-state model. Each point represents the max fractional occupancy from a single run.

activity within the salience network. Lastly, W5 was characterised by a reduction in activity within the lateral visual network and the DMN / language, alongside an increase in activity within the superior sensorimotor network.

With regards to the FC of each state (Figure 7, right panel), all states showed strong intercorrelations between the activity of all the visual networks, as well as correlations between DAN with all visual networks and the inferior sensorimotor network. W3 showed increased intercorrelations within sensorimotor networks and within the DMN. W4 showed a decrease in FC between

visual networks and the superior sensorimotor network, as well as an increase in FC between the DMN / language network and the inferior sensorimotor network. W5 showed a reduction in FC between visual networks and the superior sensorimotor network. Additionally, the state showed an increase in FC between all visual networks and the inferior sensorimotor network, as well as between the DAN and the inferior sensorimotor network.

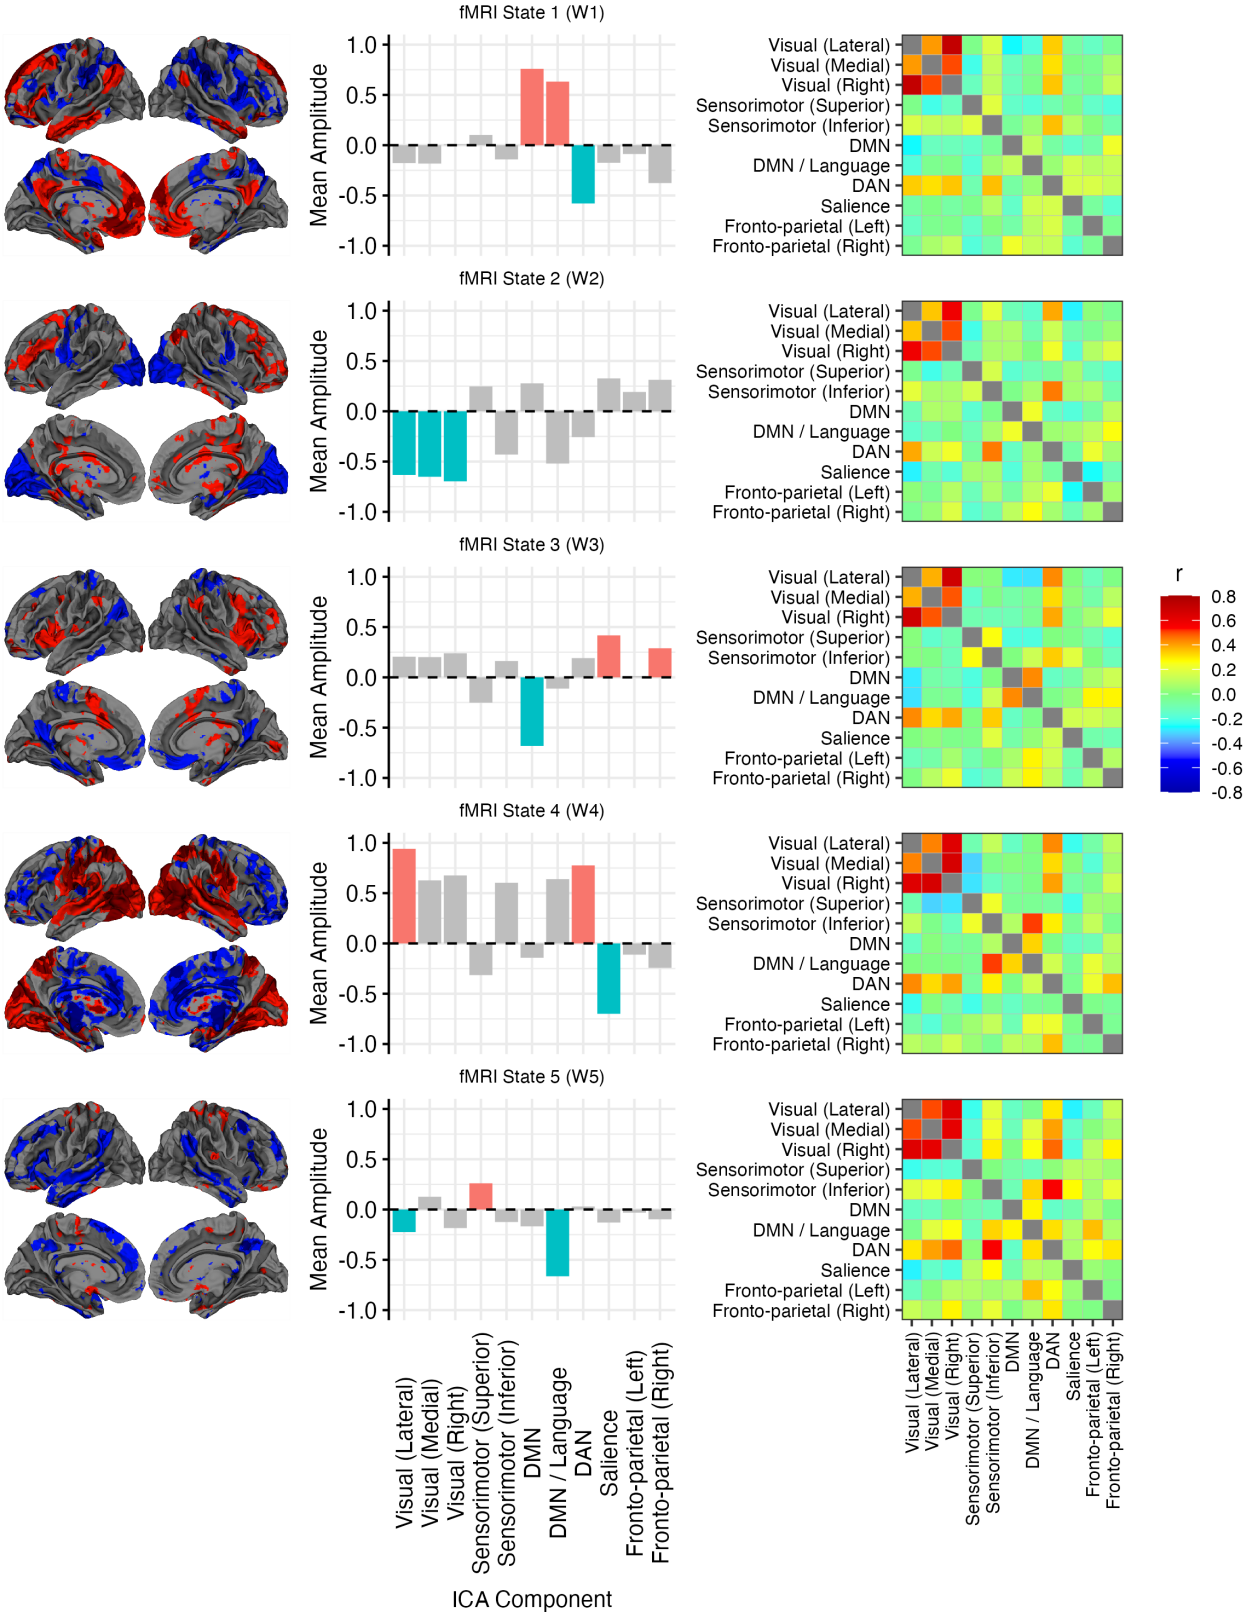

fMRI State 1 (W1)
fMRI State 2 (W2)
fMRI State 3 (W3)
fMRI State 4 (W4)
fMRI State 5 (W5)
Mean Amplitude
1.0
0.5
0.0
-0.5
-1.0
Visual (Lateral)
Visual (Medial)
Visual (Right)
Sensorimotor (Superior)
Sensorimotor (Inferior)
DMN
DMN / Language
DAN
Salience
Fronto-parietal (Left)
Fronto-parietal (Right)
ICA Component
r
0.8
0.6
0.4
0.2
0.0
-0.2
-0.4
-0.6
-0.8

Figure 4.7: States identified by the whole-brain HMM. Each row contains the plots for a single state. The left column shows the mean activation within each ICA component scaled across states. The central column shows the mean activation of each IC component for each state without scaling. Coloured bars represent the three components with the largest amplitudes (red indicates positive amplitude. Blue indicates negative amplitude). The right column shows the functional connectivity between IC components for each state.

Evoked HMM Response

The state time courses were epoched around the stimulus onsets to investigate whether the HMM dynamics were perturbed by stimulus onset (see Figure 5.8). This analysis demonstrated that states W1, W2 and W3 all showed large changes in fractional occupancy in response to the stimuli. State W3 showed the most prominent increase in fractional occupancy and exhibited a similar temporal profile to the BOLD responses from the V1 ROI analysis. Both states W2 and W3 showed a decrease in FO upon stimulus onset. Lastly, state W5 showed a minor increase in likelihood at approximately the seventh TR. This may reflect the button press or the BOLD undershoot given its delay relative to the peak likelihood of states W1, W2 and W3.

Distribution and Reliability of Baseline States

The distribution of baseline states was plotted to evaluate the likelihood of being in any given baseline state at the time of stimulation. Figure 4.9 shows the number of trials where each state was active at the baseline for both conditions. States W1 and W2 were more likely to be the baseline state for both conditions, with an average baseline count of approximately 17 for both states. Likewise, these two states also showed a higher median within the passive condition compared to the active.

To assess whether t-1 was a reliable baseline, a correlation matrix was calculated, with the

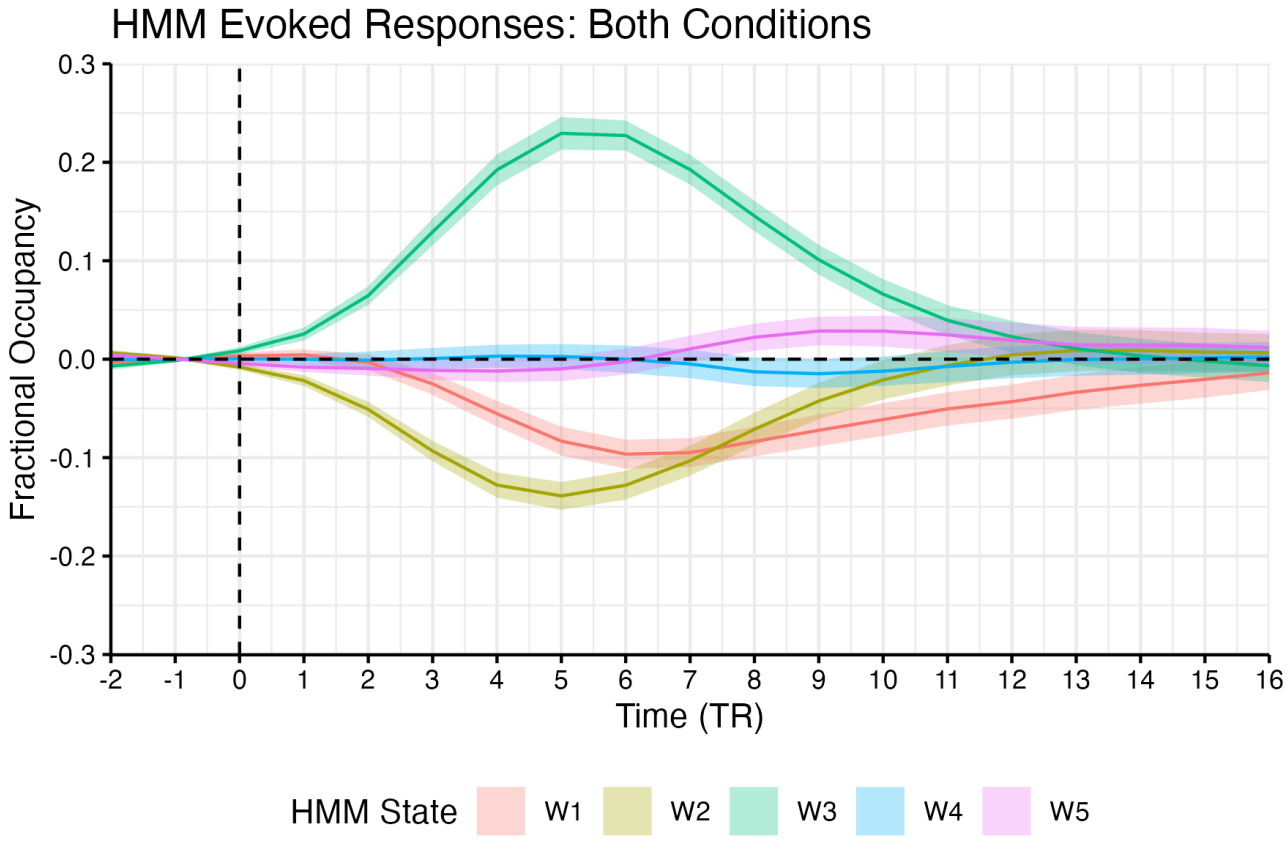


Figure 4.8: Group average changes in state fractional occupancy upon stimulus presentation. Error represents repeated measures 95% confidence intervals.

baseline condition being determined at [t0…t-5] (see Figure 4.9). The order of baseline states was then correlated using Pearson correlations. This demonstrated that baselines determined at t0, t-1, t-2, and t-3 were strongly correlated, suggesting that the baseline is reliable, irrespective of the chosen timepoint, within a 3-4 second bracket.

State Metric Comparison

State metrics were compared between the two conditions to determine whether the condition had a significant impact on the network dynamics (see Figure 5.10). Paired t-tests were run on both the fractional occupancy measures and the switching rate. Statistical tests were not performed on the state lifetimes and interval times due to large differences in group sizes. Bonferroni

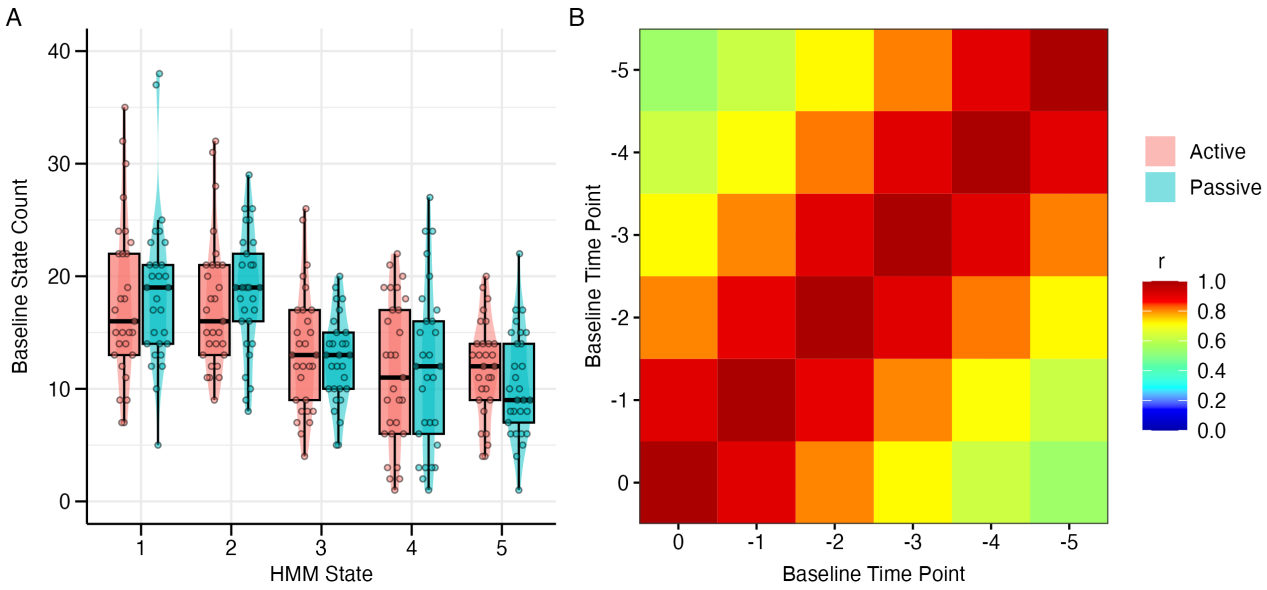


Figure 4.9: Plot A shows the number of times each state was 'active' during the baseline for each condition. Each point represents the number of times a state was active for that condition (combined across runs) and state. Plot B shows the reliability of the baseline states based on whether the baseline state is defined at the 0 – -5 timepoint.

correction was applied to control for the multiple comparisons (N=5). Paired t-tests of the fractional occupancy revealed a significant difference with a large effect size in the fractional occupancy of state W3 between the two conditions (difference = 0.06, 95% CI [0.03, 0.08], t(28) = 4.97, $p < 0.001$, adjusted $p < 0.001$; Cohen's d = 0.92, Cohen's d 95% CI [0.48, 1.35]). No other states showed a significant difference in fractional occupancy between conditions.

Impact of Whole-Brain HMM States on BOLD Responses

Trials were grouped according to baseline state to assess how the baseline state affects the amplitude of BOLD responses (see Figure 4.11). A clear increase in amplitude was observed at the peak of the BOLD response in V1 when the pre-stimulus state was W2 or W4 relative to the other states. Additionally, BOLD responses across all pre-stimulus HMM states exhibited a large amount of variance within the undershoot period (10-16 TR). Responses from W4 remained high after the peak, followed by W1 and W2, whilst W3 and W5 showed a notable

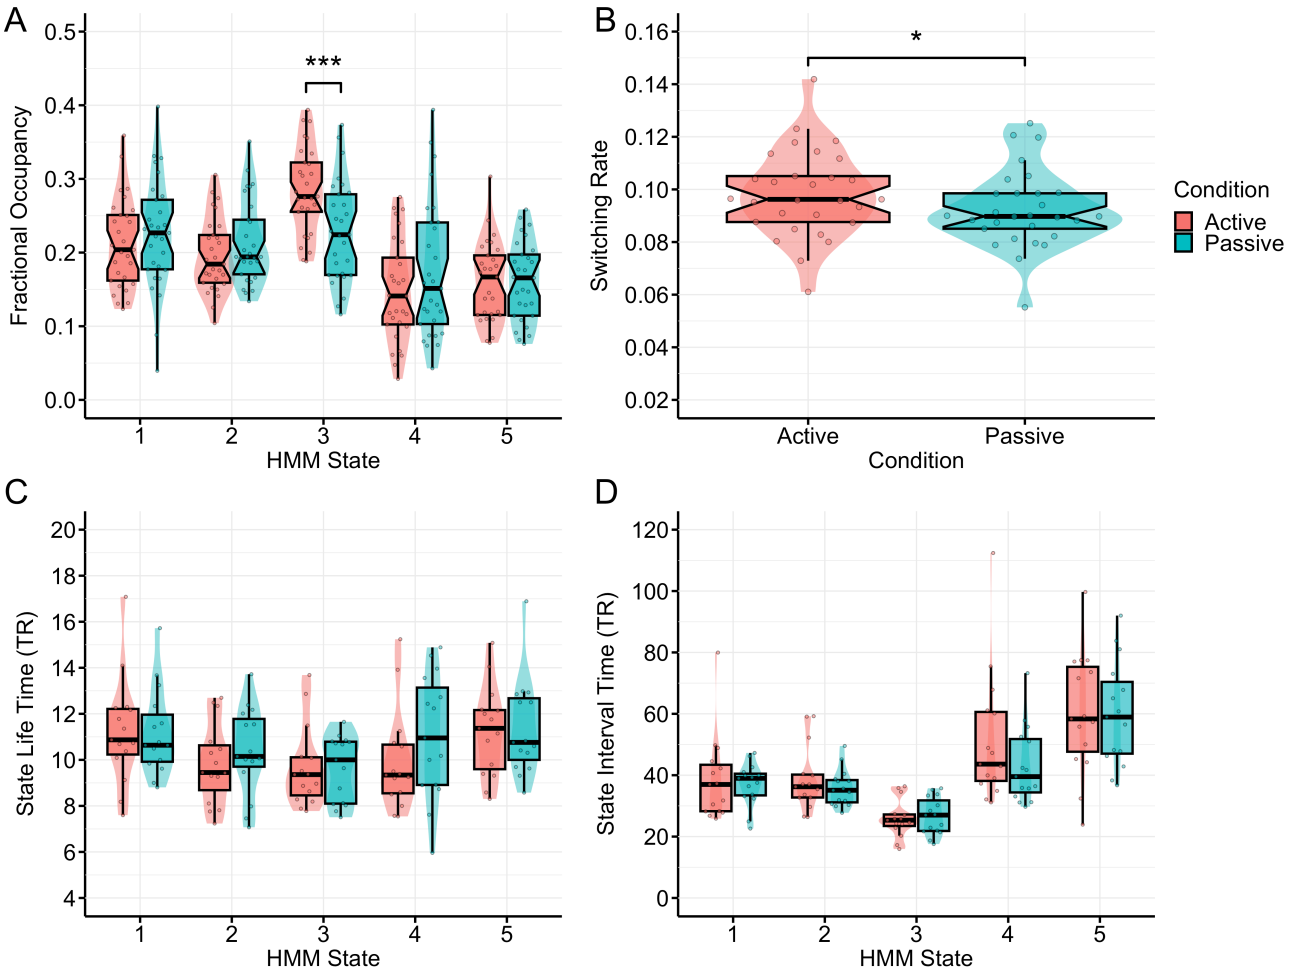


Figure 4.10: Comparison of HMM states metrics between the active and passive conditions. Plot A compares max fractional occupancy. Plot B compares the state switching rate. Plot C compares the average state lifetimes. Plot D compares the average state interval times. Each point represents a single subject.

decrease in amplitude.

The differences in the mean peak (2-9 TR) and undershoot (10-17 TR) BOLD signal change values were compared across condition, stimulus type, and baseline HMM states using linear mixed models (see Chapter 2 for details). The results of the right V1 models found the baseline HMM state main effect to be significant for both the BOLD peak ($F(4, 4136.5) = 44.21, p < 0.001$) and BOLD undershoot ($F(4, 4103.1) = 78.00, p < 0.001$). Additionally, the condition main effect was significant for the BOLD peak ($F(1, 4128.2) = 21.70, p < 0.001$), with the BOLD

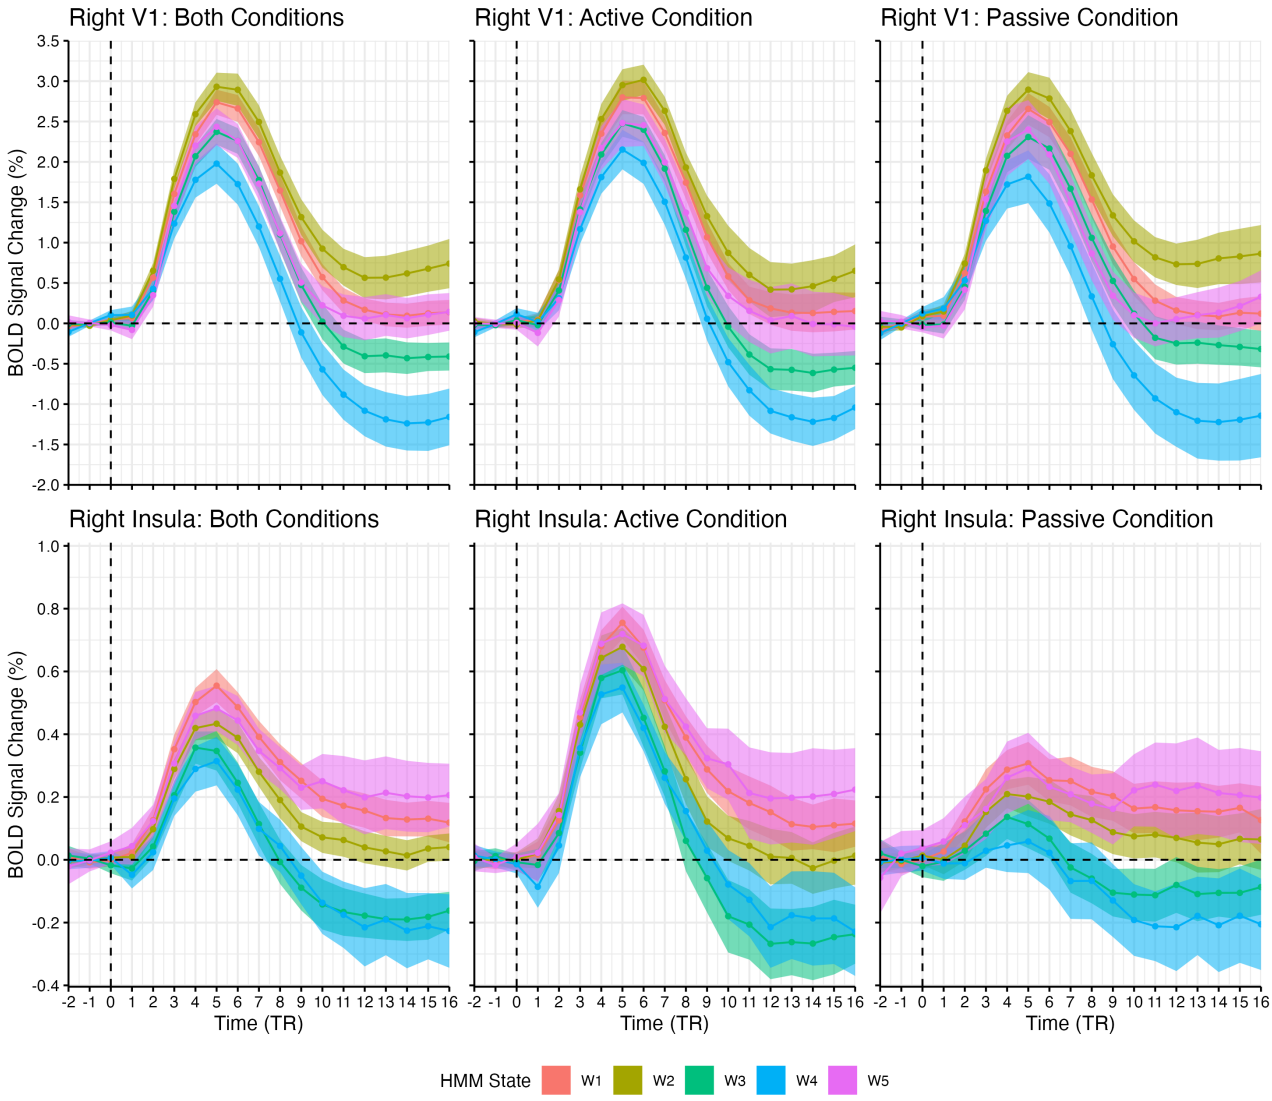


Figure 4.11: Comparison of mean BOLD responses across states. Each row represents an ROI (right V1 and right insula). Each column represents the grouping of trials, with the left column containing trials from both conditions and the centre and the right columns only containing trials from the active and passive conditions, respectively. Error represents 95% confidence intervals.

peak PSC values being significantly larger during the active condition ($EMM = 1.72\%, SE = 0.112\%$) relative to the passive condition ($EMM = 1.46\%, SE = 0.112\%$). Significant interaction effects between condition and stimulus type were found for both the BOLD peak ($F(1, 4128.2) = 10.43, p < 0.01$) and BOLD undershoot ($F(1, 4130.1) = 10.28, p < 0.01$). All

other main and interaction effects were non-significant ($p > 0.05$). The results from the right insula models closely resembled the results in right V1, with the baseline HMM state main effect being significant within both the peak model ($F(4, 4150.0) = 25.03, p < 0.001$) and undershoot model ($F(4, 4087.6) = 41.21, p < 0.001$). Additionally, the condition ($F(1, 4128.8) = 216.91, p < 0.001$) and stimulus type ($F(1, 4128.4) = 21.92, p < 0.001$) main effects were significant for the BOLD peak model, with target stimuli ($EMM = 0.328\%, SE = 0.026\%$) generating significantly larger BOLD peak PSC value than the standard stimuli ($EMM = 0.233\%, SE = 0.022\%$), and stimuli generating significantly larger BOLD peak PSC values during the active condition ($EMM = 0.431\%, SE = 0.024\%$) relative to the passive condition ($EMM = 0.130\%, SE = 0.024\%$). Significant interaction effects between baseline HMM state and stimulus type were found for both the BOLD peak ($F(4, 4135.2) = 2.63, p < 0.05$) and BOLD undershoot ($F(4, 4151.2) = 4.10, p < 0.01$) models. All other main and interaction effects were non-significant ($p > 0.05$).

Post hoc tests for all models revealed significant differences in the BOLD PSC values between the different baseline HMM states (see Figure 4.12 for statistics). State W2 generated the largest changes in BOLD peak signal change within right V1, followed by states W1, W5, W3, and W4. The inverse pattern was observed during the undershoot, with state W4 generating the largest BOLD undershoot PSC values. Within the right insula, state W2 generated the largest changes in BOLD peak PSC values, followed by states W5, W1, W3, and W4. Similar to right V1, the undershoot followed the inverse pattern, with BOLD undershoot PSC values being largest for state W4.

Regarding the significant interaction effects, the post hoc tests (Bonferroni corrected, number of comparisons = 6) on the right V1 peak model found that target stimuli generated larger BOLD PSC values during the active condition ($EMM = 1.83\%, SE = 0.125\%$) compared to target stimuli presented during the passive condition ($EMM = 1.39, SE = 0.125$) ($t(4128) =$

$4.548, p < 0.001$). Conversely, no significant difference was found in the BOLD PSC values for standard stimuli presented in the active condition ($EMM = 1.60\%, SE = 0.112\%$) compared to the passive condition ($EMM = 1.52\%, SE = 0.112\%$) ($t(4128) = 1.434, p = 0.911$). Additionally, target stimuli ($EMM = 1.83\%, SE = 0.125\%$) generated significantly larger BOLD PSC values than standard stimuli ($EMM = 1.60\%, SE = 0.112\%$) during the active condition ($t(4128) = -2.858, p < 0.05$), not during the passive condition ($t(4128) = 1.708, p = 0.526.$). Similarly, post hoc tests (Bonferroni corrected, number of comparisons = 6) for the right V1 undershoot model revealed that standard stimuli ($EMM = -0.213\%, SE = 0.057\%$) generated significantly larger BOLD PSC values relative to target stimuli ($EMM = 0.073\%, SE = 0.093\%$) during the active condition ($t(4129) = -2.745, p < 0.05$), but not during the passive condition ($t(4130) = 1.788, p = 0.443$).

Concerning the right insula, post hoc tests (Bonferroni correct, number of comparisons = 45) on the peak model found that whilst target stimuli always generated larger BOLD PSC values than standard stimuli, target stimuli ($EMM = 0.527\%, SE = 0.049\%$) generated notably larger BOLD PSC values when state W5 was active during the baseline relative to standard stimuli ($EMM = 0.288\%, SE = 0.032\%$) ($t(4134) = -4.618, p < 0.001$). Similarly, for the right insula undershoot model, post hoc tests found that target stimuli ($EMM = 0.408\%, SE = 0.059\%$) generated smaller BOLD PSC values than standard stimuli ($EMM = 0.159\%, SE = 0.033\%$) when state W5 was active during the baseline period ($t(4149) = -3.729, p < 0.001$).

Impact of Whole-Brain HMM States on Alpha ERD/S

The ERD/S analysis showed that all baseline HMM states showed a reliable desynchronisation within the alpha band, followed by a rebound at approximately 1.75 seconds (see Figure 4.13). The results found the condition main effect to be significant ($F(1, 3770.90) = 35.46, p < 0.001$), with the ERS being significantly larger during the active condition ($EMM = 154.4$ $\mu V^2/Hz$,

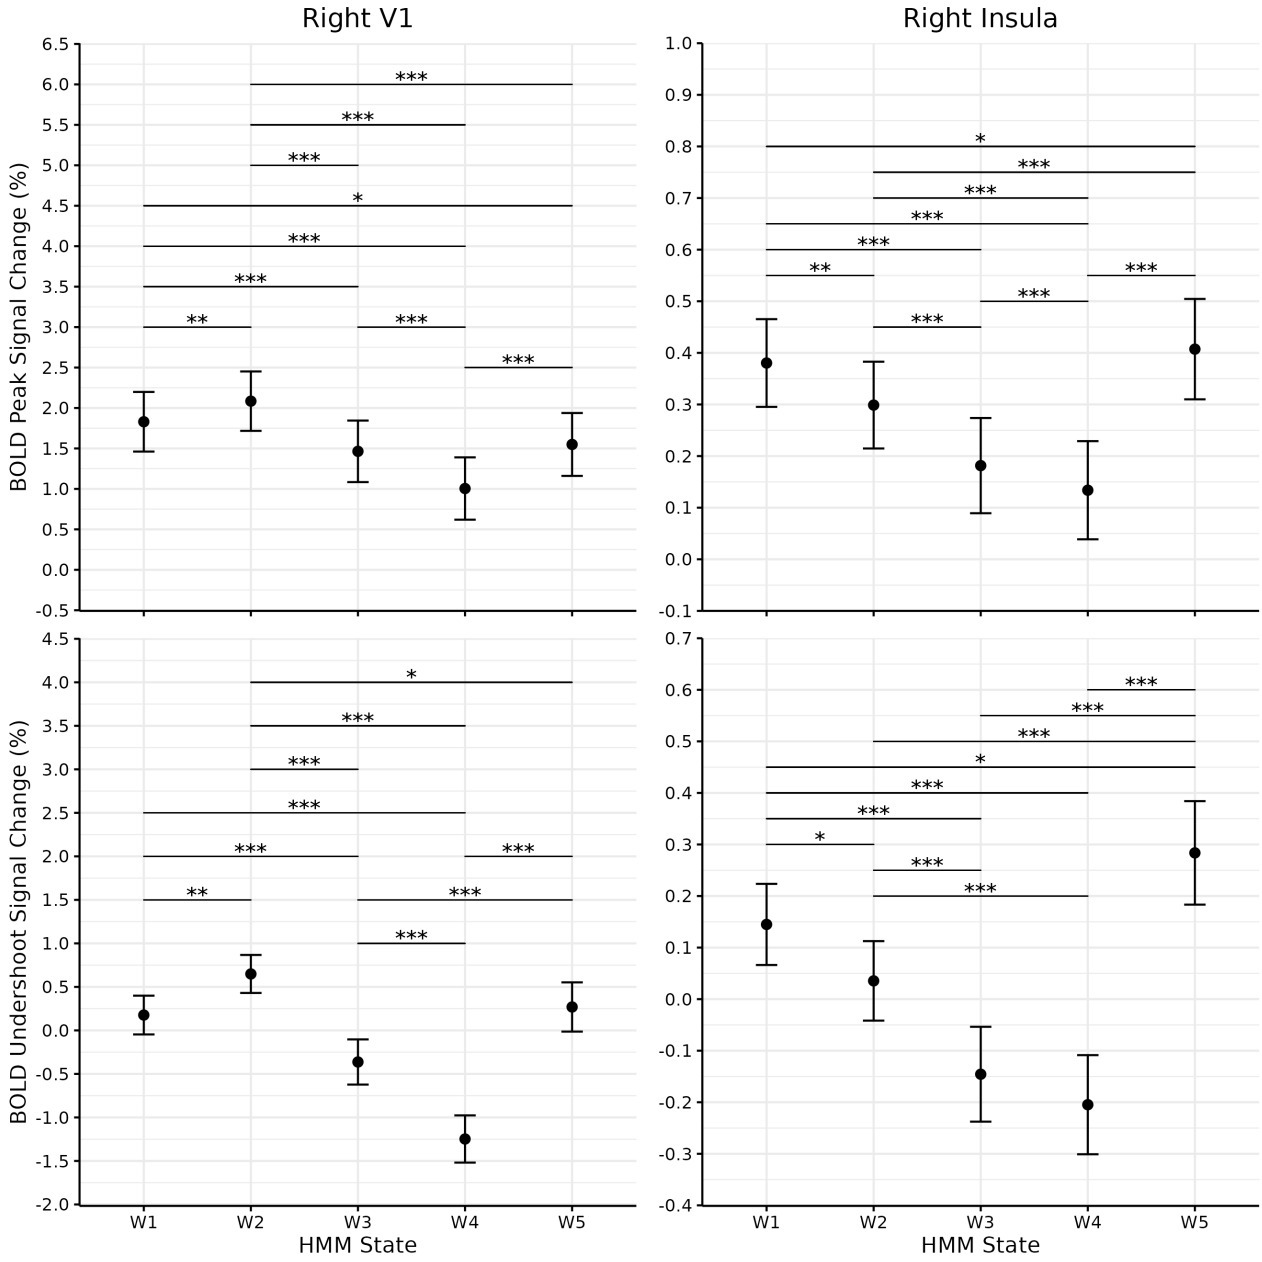


Figure 4.12: Comparison of BOLD peak and undershoot estimated marginal mean amplitudes for each baseline HMM state. Post hoc tests were conducted using an estimated marginal means approach and corrected via Bonferroni correction (number of comparisons = 10). *** = $p < 0.001$, ** = $p < 0.01$, * = $p < 0.05$. Error bars represent 95% confidence intervals.

SE = 20.2 μV^2/Hz) relative to the passive condition (EMM = 54.2 μV^2/Hz, SE = 20.1 μV^2/Hz). Additionally, a significant interaction effect ($F(4, 43779.3) = 2.407, p < 0.05$) was

found between condition and the baseline HMM state. Post hoc tests (Bonferroni corrected, number of comparisons = 45) revealed that trials where state W4 was active during the baseline generated a larger ERS during the active condition (EMM = 227.1 µV^2/Hz, SE = 33.6 µV^2/Hz) compared to the passive condition (EMM = 31.4 µV^2/Hz, SE = 32.3 µV^2/Hz) ($t(3780) = 4.849, p < 0.001$). All other main and interaction effects were non-significant ($p > 0.05$).

Impact of Whole-Brain HMM States on VEPs

Like the alpha ERD/S analysis, all baseline HMM states generated a reliable VEP, peaking at 90ms (see Figure 4.14). Statistical tests found that the condition main effect was significant for the N150 component ($F(1, 3870.9) = 18.126, p < 0.001$), with the amplitude of the deflection being significantly larger during the active condition ($EMM = -4.68$ µV, $SE = 0.636$ µV) relative to the passive condition ($EMM = -2.96$ µV, $SE = 0.635$ µV). Additionally, the HMM baseline state main effect was significant for the P100 component ($F(4, 3887.8) = 2.409, p < 0.05$). Post hoc tests revealed that trials where state W2 ($EMM = 6.03$ µV,$SE = 0.597$ µV) was active during the baseline period generated significantly larger P100 amplitudes than trials preceded by state W4 ($EMM = 4.24$ µV,$SE = 0.668$ µV) ($t(3887) = 2.841, p < 0.05$). All other main and interaction effects were non-significant ($p > 0.05$).

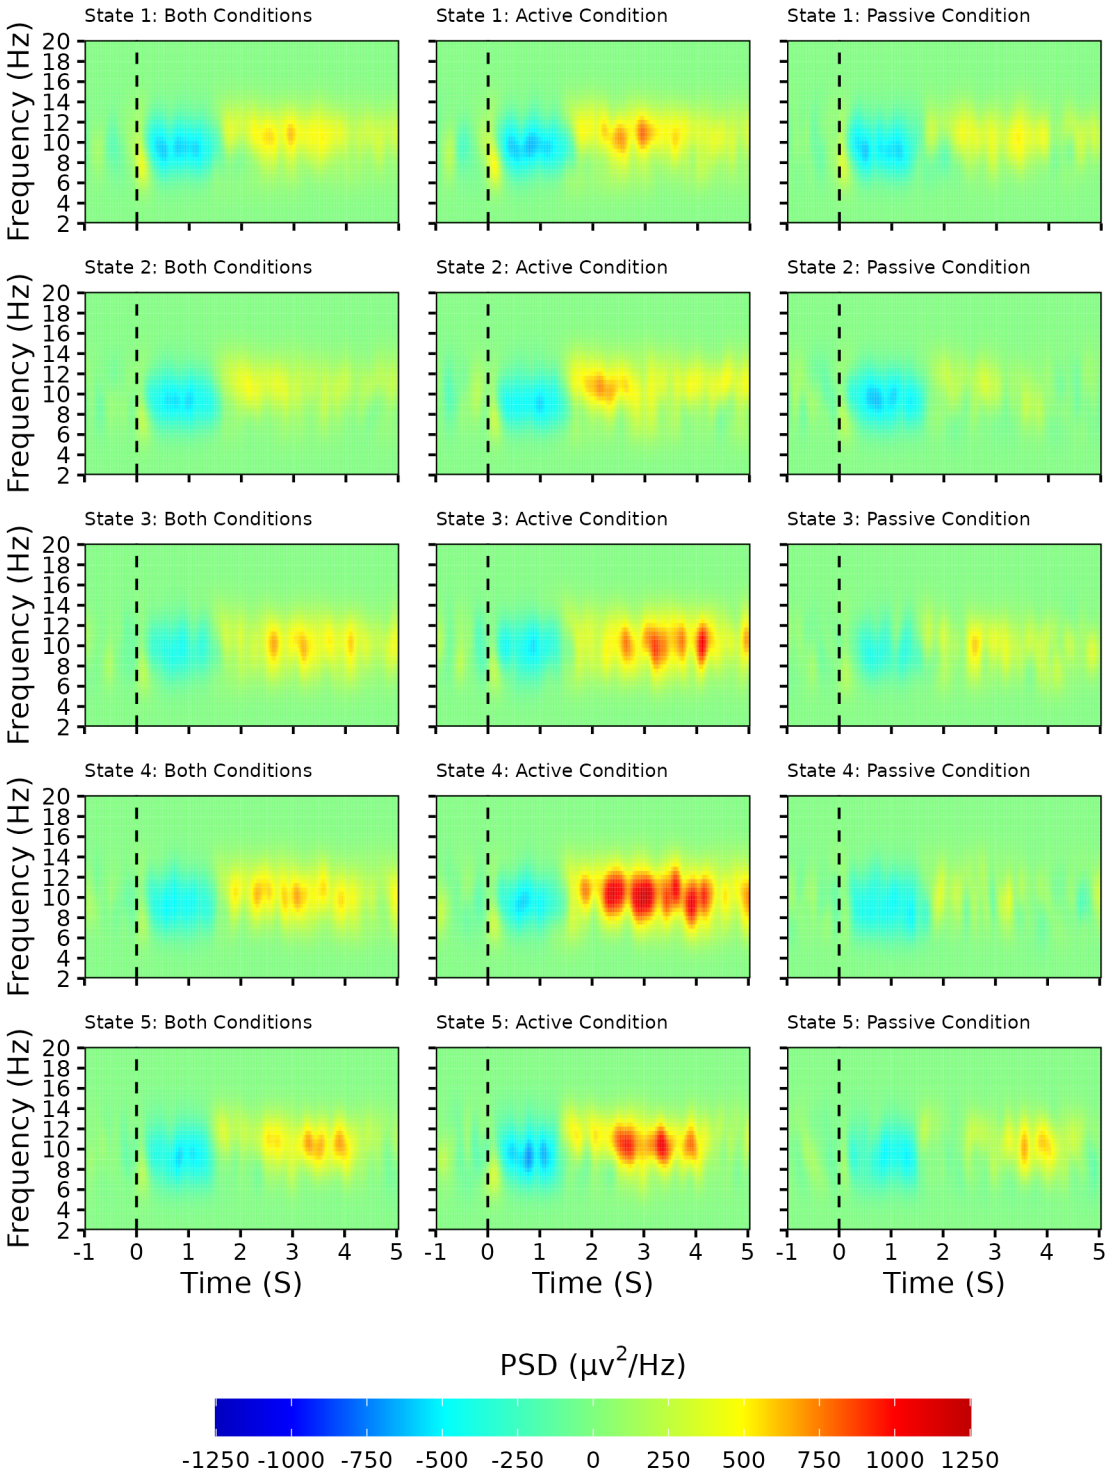


Figure 4.13: Grand average alpha ERD/S for each whole-brain HMM state. The left column shows the averages calculated with all trials, whilst the centre and right columns show the averages calculated using trials from the active and passive conditions, respectively. Error ribbons omitted for visual clarity PSD = power spectral density.

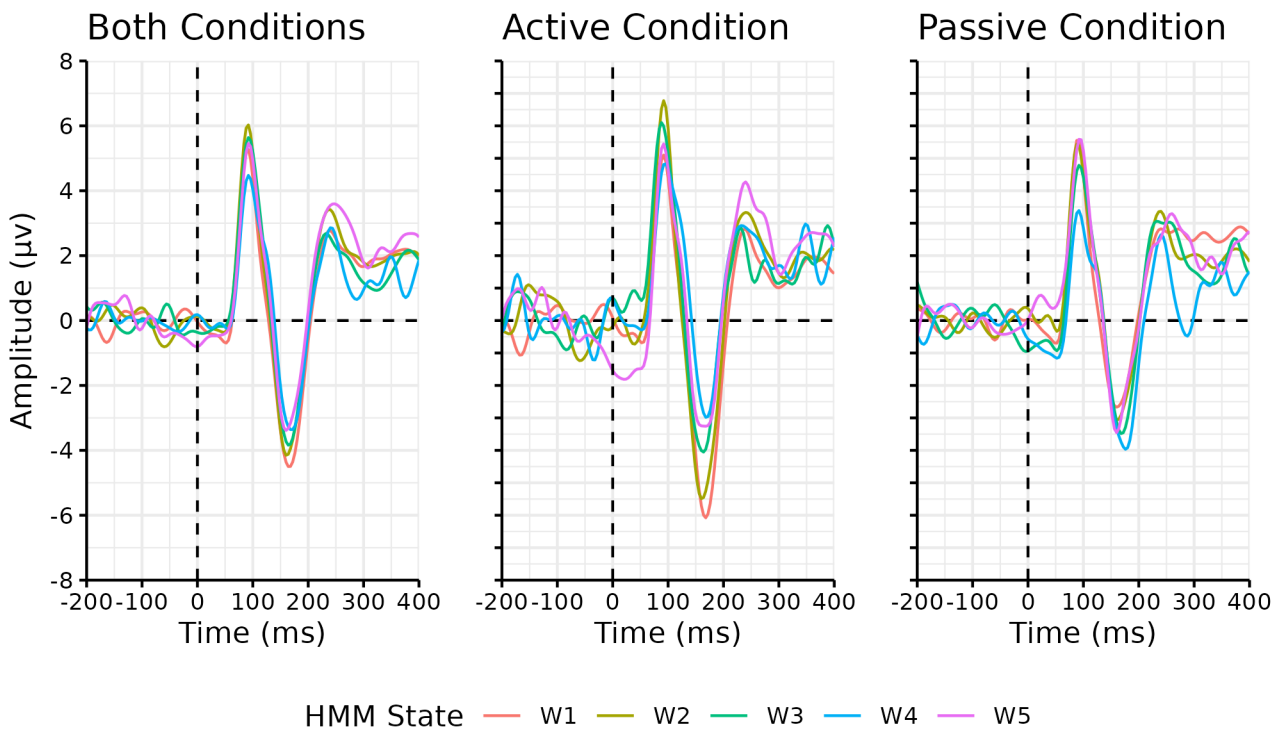


Figure 4.14: Comparison of grand average VEPs across HMM states. The left column shows the averages calculated with all trials, whilst the centre and right columns show the averages calculated using trials from the active and passive conditions, respectively.

Discussion

Summary

This chapter investigated the impact of whole-brain BOLD states on the trial-to-trial variability of evoked BOLD and electrophysiological responses during active and passive viewing of visual stimuli. The GLM analysis revealed that the stimuli during the active condition generated a significantly larger BOLD response within the salience network, particularly within the right insula. Subsequent ROI analysis further confirmed this, with BOLD responses in right V1 and the right insula showing significantly larger BOLD response peaks during the active condition than the passive condition despite the identical stimuli within both conditions. Whole-brain states were then identified via Hidden Markov modelling (Vidaurre et al., 2016) using the BOLD data, inferring a total of five states. Trials were then grouped according to which whole-brian state was active before stimulus onset, and the average BOLD responses, VEPs, and alpha ERD/s responses were calculated. The results found that peak BOLD response and P100 amplitudes varied significantly depending on the baseline state. Specifically, when state W4 (which showed high BOLD activity within visual cortex) was active during the baseline, the peak BOLD responses in right V1 and P100 amplitudes were significantly smaller than when state W2 (which showed low BOLD activity within visual cortex) was active. Furthermore, peak BOLD responses in right V1 and the right insula were significantly larger when state W1 (which showed high BOLD activity within the DMN) was active during the baseline relative to other states. Regarding the state dynamics, no significant difference in fractional occupancy or switching rate was found between conditions.

BOLD Activation Maps Differ Significantly Between Conditions

The results of the GLM analysis found that the stimuli generated significantly different patterns of BOLD activity depending on the condition. Interestingly, significantly more BOLD activity was observed within ipsilateral visual cortex during the active condition relative to the passive condition. As the active conditions required participants to judge the contrast level of the stimulus, this may suggest that participants made saccades towards the stimuli to better detect any change in contrast level. Additionally, the salience network also showed significantly higher BOLD activity during the active condition relative to the passive condition. This finding has been observed in previous research, with nodes within the salience network showing an increase in activity to deviant/salient visual stimuli (Crottaz-Herbette & Menon, 2006; Downar et al., 2001; Linden, 1999), e.g. target stimuli during an oddball task. Thus, the increase in BOLD activity within the salience network was potentially a result of the perceived salience of target stimuli during the active condition relative to the passive condition.

Regarding negative BOLD activity, the DMN exhibited significantly larger deactivations during the active condition relative to the passive condition. The deactivation of the DMN in response to task-related stimuli has been widely reported within the literature (Mayhew, Ostwald, et al., 2013) and has been demonstrated to represent a reduction in neural activity within the network (Miller et al., 2009). Furthermore, the magnitude of the deactivation has been shown to be larger when the task is more difficult, which agrees with the larger deactivation seen within the active condition (McKiernan et al., 2006; Singh & Fawcett, 2008). Despite being reliably observed within the literature (Bressler et al., 2007; Mayhew, Ostwald, et al., 2013; Pasley et al., 2007; A. T. Smith et al., 2004; Wade & Rowland, 2010; Wilson et al., 2020), no significant deactivation was observed within contralateral (left) primary visual cortex.

BOLD Responses Differ Significantly Between Conditions

The ROI analysis showed that all visual stimuli generated strong evoked BOLD responses within right V1 and the right insula, regardless of the condition. However, the peak of the evoked BOLD responses was significantly larger during the active condition relative to the active condition. In line with previous literature, maintaining attention to the stimuli resulted in a larger response despite being identical in both conditions (Ciaramitaro et al., 2007; Gandhi et al., 1999; X. Li et al., 2008; Somers et al., 1999). Additionally, the BOLD responses within the right insula differed significantly depending on the stimulus type, with target stimuli eliciting significantly larger peak BOLD responses compared to standard stimuli. This result matched with the previous literature and the results of the GLM analysis, demonstrating that the insula responds preferentially to deviant/salient stimuli (Crottaz-Herbette & Menon, 2006; Downar et al., 2001; Linden, 1999). Furthermore, an interaction effect was observed between stimulus type and condition within right V1, with target stimuli generating a significantly larger BOLD response peak during the active condition relative to the passive condition. The post hoc tests also revealed that the difference in the peak BOLD response amplitudes between standard stimuli during the active and passive conditions was not significant. This indicates that the increased peak BOLD amplitudes during the active condition are primarily driven by the target stimuli, and that responses to standard stimuli were not enhanced during the active condition. Regarding the undershoot amplitudes, a significant interaction effect was also found between stimulus type and condition, with the post hoc tests revealing that the standard stimuli generated a significantly larger BOLD undershoot relative to target stimuli, but only during the active condition. One explanation for this result is that the larger BOLD undershoot reflects subject's inhibiting responses to the standard stimuli. EEG-fMRI research has demonstrated that BOLD response undershoots were only observed in trials with high post-stimulus alpha ERS (K. J. Mullinger, 2017; K. J. Mullinger, Mayhew, et al., 2013), which has been behaviourally associated

with inhibiting responses and top-down control (Klimesch et al., 2007).

Whole-Brain HMM States

The HMM identified a total of five states, three of which were also identified within the previous chapter (see Chapter 3). Specifically, states W1 (DMN), W3 (salience), and W4 (multi-sensory) have all been previously identified both within Chapter 3 and the literature (Song et al., 2022), providing further evidence that the HMM reliably identifies these patterns of BOLD activity. Interestingly, states W2, and W5 were not identified by the HMM in Chapter 3. As state W5 was primarily defined by an increase in activity within superior sensorimotor cortex, and a decrease within the DMN / Language network, it seems likely that it is associated with the button presses during the active condition. However, the state showed no significant difference in fractional occupancy between active and passive conditions, and as there were no button presses within the passive condition, it remains unclear what W5 represents behaviourally. The results from the evoked analysis showed an increase in the fractional occupancy of state W3 upon stimulus onset, and a decrease in fractional occupancy of states W1 and W2. The changes in fractional occupancy of states W3 and W1 closely resemble the results of the GLM analysis, where the right insula exhibited a significant increase upon stimulus onset, and the DMN showed a negative response. This provides evidence that the HMM approach not only reliably identified states, but also identified evoked state dynamics that closely resemble the results of conventional GLM-based analysis.

HMM State Dynamics

Despite the differences in the evoked BOLD and EEG responses between the two conditions and the significant impact of baseline HMM state on evoked BOLD responses, the HMM state metrics did not differ significantly between the two conditions. The exception was state W3, which showed increased fractional occupancy during the active condition. This likely reflects the

larger salience network responses to stimuli during the active condition, as observed within the GLM and ROI analysis. However, it is still surprising that no other states exhibited significant changes in state metrics, given that the literature has previously identified changes reliably within the DAN, DMN and salience network during sustained attention tasks (Christoff et al., 2009; Fortenbaugh et al., 2018; Leber, 2010; Rosenberg et al., 2015; Rosenberg et al., 2020; Weissman et al., 2006). One potential reason for this discrepancy may be the large inter-stimulus interval (ISI). As each ISI ranged from 16.5-21 seconds, subject's attention likely lapsed during these periods, regardless of the condition, which would explain the lack of variation in state dynamics across the two conditions.

Impact of Whole-Brain HMM States on Responses

The HMM state active during the baseline period had a significant impact on the evoked BOLD responses in both contralateral primary visual and insula cortex. Within primary visual cortex, evoked BOLD response peaks were significantly larger when either state W1 (associated with high DMN activity) or W2 (associated with low visual cortex activity) was active at baseline, and significantly reduced when state W4 (associated with high sensory cortex activity) was active. A similar pattern was observed within the insula, with evoked BOLD response peaks being highest when state W5 (associated with high superior sensorimotor activity) was active at baseline, and lowest when either state W3 or W4 was active. Furthermore, this pattern was also observed within the VEP results, with the P100 amplitude being significantly reduced when state W4 was active during the baseline, and significantly larger when state W2 was active during the baseline. This indicates that states W2 and W4 were especially relevant, as they modulated both the BOLD responses and the P100 amplitudes. These states were primarily characterised by large activations or deactivations within the sensory cortices, including visual cortex. Thus, the modulation may result from the baseline BOLD levels within visual cortex itself, with high BOLD activity reducing the amplitude of BOLD responses and the P100 component, and vice

versa with low BOLD activity.

This inverse relationship between baseline BOLD activity and peak BOLD response amplitudes has been demonstrated in previous literature and was demonstrated to be associated with a reduction in BOLD variability post-stimulus onset (He, 2013). This result contradicts the linear summation hypothesis, as the whole-brain state and BOLD response interact non-linearly, suggesting a more complex relationship between ongoing brain dynamics and visual responses. One explanation of these results is that the differences represent a haemodynamic 'ceiling' effect (Cohen et al., 2002; Mark et al., 2015). Previous calibrated fMRI studies have manipulated the baseline cerebral blood flow (CBF) levels via CO2 inhalation, which revealed an inverse relationship between the baseline levels of CBF and the amplitude of the BOLD response (Blockley et al., 2013; Cohen et al., 2002; Mark et al., 2015; Stefanovic et al., 2006). This effect is so prominent that BOLD responses can be silenced entirely if the baseline CBF levels are sufficiently low enough (Gauthier et al., 2011). Thus, these results would indicate that the BOLD response shows a ceiling effect, in which high baseline CBF levels greatly reduce the level of deoxyhaemoglobin, which results in a lower BOLD response peak due to CBF levels being closer to the 'ceiling' level, resulting in a smaller change in CBF and consequently BOLD signal change (Cohen et al., 2002; Mark et al., 2015). However, given that the baseline state also modulated the amplitude of the P100 VEP component, it is unclear why this haemodynamic mechanism would impact neural responses. One potential explanation is that the baseline state serves a functional role, modulating the amplitude of the VEP intentionally to facilitate visual processing mechanisms, such as visual attention.

Limitations and Future Research

Despite previous literature reporting differences in FC during sustained attention tasks, this effect was not observed within this study. This may result from the long ISI of 16.5-21 seconds,

giving long pauses for the subject's attention to lapse. Though shortening the ISI would amend this, the longer ISI was necessary to observe the entirety of the evoked BOLD responses. An alternative to reducing the duration of the ISI would be to increase the difficulty of the active condition by adding a working memory component or mental task to complete during the intervals. Additionally, during the active condition, the GLM results suggested that participants were not maintaining fixation on the fixation cross for all trials. This could be amended by implementing eye tracking, allowing for the trials where fixation was not maintained to be rejected. In regards to future research, as the results point towards an impact of baseline CBF on evoked BOLD response, simultaneous measurement of BOLD and CBF could be employed in order to further disentangle the exact mechanism behind the observed results (K. Mullinger et al., 2014) (albeit at the cost of TR and signal-to-noise ration). Assuming that the variability in the peak BOLD responses explained by baseline BOLD signal is not relevant function for visual processing, this variability could be regressed out to increase the statistical power of neuroimaging analyses (Krishnamurthy et al., 2020). Lastly, the methodology applied here could be applied to populations known to show differences in BOLD response variability, such as in ageing populations (Garrett et al., 2010; Grady & Garrett, 2014) and in individuals with Autism (Dinstein et al., 2012; Haigh et al., 2015).

Conclusion

This chapter demonstrated that the active whole-brain state during the pre-stimulus baseline period significantly influences BOLD responses in V1 and the insula, as well as P100 amplitudes. This finding shows that baseline BOLD activity defined across the whole brain affects subsequent BOLD responses and electrophysiological responses to visual stimuli, highlighting the importance of accounting for baseline BOLD activity within neuroimaging research. Future studies should aim to replicate the results found within this chapter within other sensory modalities using the Hidden Markov modelling approach. Additionally, the approach should

be applied to data acquired from populations shown to exhibit heightened sensory response variability (e.g. autism) to test whether the baseline BOLD activity explains the increased variability. Lastly, the method should be applied to other modalities to investigate the impact of spectrally defined states on evoked BOLD responses and electrophysiological measures, which is explored in the next chapter.

Chapter 5

Alpha State Modulates Stimulus Induced Alpha ERD/S

Abstract

Pre-stimulus alpha power has a profound influence on both neural and behavioural responses to visual stimuli, impacting the amplitudes of BOLD responses (Becker et al., 2011; Mayhew, Ostwald, et al., 2013), VEPs (Becker et al., 2008; Reinacher et al., 2009) and alpha ERD/S (Mayhew, Ostwald, et al., 2013), as well as determining the likelihood of perceiving near-threshold stimuli (Hanslmayr et al., 2007). Thus, the impact of pre-stimulus alpha power is further investigated here using a novel HMM approach (Baker et al., 2014; Vidaurre et al., 2016) to identify periods of high and low alpha power, and subsequently, group trials according to which state was 'active' immediately before stimulus presentation. The results found a significant impact of baseline HMM state on the magnitude of the alpha ERD and ERS, with the ERD primarily being generated by trials preceded by the high alpha state, and the ERS being generated by trials preceded by the low alpha state. In contradiction to previous literature (Becker et al., 2011; Becker et al., 2008; Mayhew, Ostwald, et al., 2013; Reinacher et al., 2009), the baseline alpha state had no significant impact on the BOLD responses or VEP amplitudes. Overall, the results further support the role of alpha in visual attention, with pre-stimulus alpha power serving as a measure of anticipation, which subsequently modulates the alpha ERD/S amplitudes.

Introduction

The alpha rhythm (7 Hz-13 Hz) is the most prominent neural oscillation visible on the EEG and has been associated with cyclic changes in cortical excitability (Romei et al., 2008). Specifically, periods of high alpha power have been demonstrated to reflect cortical inhibition, whilst periods of low alpha power reflect increased cortical excitability, in both cases indexed by transcranial magnetic stimulation (TMS). The changes in cortical excitability with alpha power have led to the rhythm being implicated in various elements of cognition, with many theorising that alpha represents a form of functional inhibition (Jensen & Mazaheri, 2010), in which unused cortical regions can be suppressed, improving information processing in other cortical regions, facilitating various elements of cognition, such as visual attention (Kelly et al., 2009; Kelly et al., 2006; Rihs et al., 2007; Thut et al., 2006; Worden et al., 2000) and working memory (Bonnefond & Jensen, 2012; Hanslmayr et al., 2009; Klimesch et al., 1994). In the context of visual processing and visual attention, pre-stimulus alpha power can be viewed as an index of attention, with research demonstrating that pre-stimulus alpha power is significantly lower in trials in which the stimulus is anticipated (Bastiaansen et al., 2001; Bastiaansen & Brunia, 2001). This has also been supported behaviourally, with participants being more likely to perceive near-threshold sensory stimuli on trials with lower pre-stimulus alpha power (Hanslmayr et al., 2007; Kelly et al., 2009; Mathewson et al., 2009; Mazaheri et al., 2009; Schubert et al., 2009; Thut et al., 2006; Van Dijk et al., 2008). Furthermore, the levels of pre-stimulus alpha also influence the magnitude of subsequent BOLD responses (Becker et al., 2011; Mayhew, Ostwald, et al., 2013), visually evoked potentials (VEPs) (Becker et al., 2008; Reinacher et al., 2009) and event-related alpha desynchronisation and synchronisation (ERD/S) (Mayhew, Ostwald, et al., 2013), supporting the involvement of the alpha rhythm in visual processing and visual attention.

Given the strong link between pre-stimulus alpha power and visual attention, this chapter investigated the impact of pre-stimulus alpha on BOLD responses, VEPs, and alpha ERD/S during

both active and passive conditions using an HMM approach (Baker et al., 2014; Vidaurre et al., 2016). Firstly, grand average plots of the VEPs and alpha ERD/S were calculated and compared across the conditions to confirm that the condition impacted the electrophysiological responses as expected. Secondly, an HMM was trained on EEG data bandpass filtered to the alpha band (7-13Hz) to identify periods of high and low alpha power. Trials were then sorted according to whether the high or low alpha state was 'active' immediately before stimulus presentation, and grand averages of the BOLD responses, VEPs and alpha ERD/S were calculated and compared. Additionally, the HMM state metrics were compared across the two conditions to determine whether the state temporal dynamics differed across the two conditions and whether either state was more likely to be 'active' at baseline during either of the conditions.

Methods

Dataset

This chapter used the active vs. passive viewing EEG-fMRI dataset collected during my PhD. See Chapter 2 for details on stimuli and tasks.

Participants

A total of 30 participants were collected within the active vs. passive EEG-fMRI dataset (see Chapter 2 for demographics). Two participants were excluded from this chapter; one for excessive movement and one for electrode bridging, resulting in 28 participants being included in all subsequent analyses.

EEG Preprocessing

Preprocessing consisted of removing gradient artefacts, ballistocardiogram artefacts and movement artefacts via carbon wire loop regression. Bad channels were identified and removed from the data before undergoing average referencing, and ICA to remove any remaining artefacts, such as blinks or muscle-related artefacts (see Chapter 2for full description and implementation details). Two copies of the EEG data were then created: one bandpass filtered to the alpha band (7Hz - 13 Hz) to be inputted into the HMM and to generate the alpha ERD/S analysis., and one bandpass filtered at 1-25Hz for the VEP analysis (see Chapter 2 for complete details on VEP and alpha ERD/S extraction).

fMRI Preprocessing

fMRI preprocessing consisted of susceptibility correction, motion correction, normalisation to standard space and ICA AROMA (Pruim et al., 2015) via fMRIPrep (Esteban et al., 2017). A custom Nipype (Gorgolewski et al., 2016) script was then employed to remove the noise components identified by ICA AROMA, apply a spatial filter of 4mm using FSL's SUSAN, and apply a high pass filter (100s). See Chapter 2 for complete details on fMRI preprocessing.

Extracting BOLD Responses

BOLD responses were extracted from each ROI and converted to percent signal change (see Chapter 2 for details).

Extracting VEPs and Alpha ERD/S

VEP and alpha ERD/S responses were extracted from the right posterior electrode that showed the strongest P100 component and baseline corrected (see Chapter 2 for details).

Alpha HMM

The time courses from the right parieto-occipital (O2, P2, P4, P6, P8, PO4, PO8) region were selected from the 7-13Hz bandpass filtered dataset and used as inputs for the HMM analysis (Baker et al., 2014; Vidaurre et al., 2016). Inputted time courses were standardised so the mean was equal to zero and the standard deviation was equal to one. Time courses were then Hilbert transformed and downsampled to 40Hz. A multivariate Gaussian model was selected for the observation model, defining states based on the mean and covariance of alpha power across the inputted electrode sites. (see Chapter 2 for details on model selection).

Evoked Fractional Occupancy

The evoked fractional occupancy responses were calculated to evaluate which states became more or less active in response to the visual stimuli. The responses were calculated by epoching the state likelihood timeseries around the stimulus onset (-2-12s) and baseline corrected using the same baseline correction method used for the BOLD response analysis. Average responses were calculated for each subject, which were then averaged to generate a grand average.

Alpha State GLM

To find the BOLD correlates of the HMM states, a GLM analysis was performed on the fMRI data, using the HMM state likelihood time series as explanatory variables. State likelihood time series were downsampled to match the TR (1.01 s) of the fMRI data and convolved using the SPM haemodynamic response function. The GLM analysis was performed in Python with the Nilearn package. First-level analyses were conducted at the subject level (T1 space), with the downsampled state time series being used as explanatory variables alongside the temporal derivative. Additionally, cerebral spinal fluid (CSF) and white matter (WM) regressors generated by fMRIprep were included within the first-level design matrix as regressors of no interest,

alongside five polynomial drift regressors to high pass filter the model at 0.01 Hz (100s). Separate first-level GLMs were fitted per HMM state and condition, resulting in four first-level models per subject. Two contrasts were used within the first-level analysis to calculate the positive and negative mean effects.

The first-level maps were then inputted into two separate second-level group analyses to generate group average maps for each state, resulting in two second-level models. The second-level design matrices consisted of two regressors for each condition and a regressor for each subject to model the within-subject variability, allowing for paired comparisons. Four contrasts were calculated for each model: the mean active map, the mean passive map, the mean map across both conditions, (one sample t-tests) and a comparison between the active and passive conditions (paired samples t-tests). Contrast activation maps were then masked using the MNI152_2mm_brain_mask from FSL before being thresholded using Bonferroni correction ($p < 0.05$).

Statistics

Statistical testing was conducted to compare the effect of condition on fractional occupancy and switching rate via paired t-tests (see Chapter 2 for details). Additionally, the impact of baseline HMM state, condition, and stimulus type on BOLD, VEP, and alpha ERD/S responses was evaluated via linear mixed modelling (see Chapter 2 for details). Lastly, as a two-state model was chosen, a chi-square test was conducted to test whether either state was more likely to be active during the baseline period within either condition.

Results

Visually Evoked Potentials

Results of the VEP analysis showed that VEPS were reliably generated during both the active and passive conditions (see Figure 5.1). Both conditions generated the classical VEP structure, with a positive peak at approximately 90 ms post-stimulus, followed by a negative deflection at 165 ms (N150) and a positive rebound at 240 ms, lasting to the end of the epoch (P250). The P100 component peaked at 5.5 µV during the active condition, and 5µV during the passive condition. The N150 component showed a large difference in peak amplitudes between conditions, with the deflection being notably larger during the active condition (-4.75 µV) relative to the passive condition (-3 µV). The P250 component peak amplitudes did not differ between conditions, with both peaking at approximately 2.9 µV.

Similarly to the previous comparison, standard and target stimuli generated VEPs irrespective of condition (see Figure 5.2). All components occurred at the same latencies as observed in the VEP comparison across conditions. Peak amplitudes for each component also showed the same patterns as the condition comparison, with both standard and target stimuli showing a larger peak N150 deflection during the active condition relative to the passive condition. No differences were observed in the VEPs for standard and target stimuli during the active condition; however, minor differences were observed during the passive condition. Standard stimuli evoked a larger P250 amplitude, peaking at 3.25 µV and 2 µV for target stimuli.

Event-Related Alpha Synchronisation and Desynchronisation

In addition to the VEPs, alpha ERD/S plots were generated to compare differences in evoked alpha power between conditions and stimulus types. During both conditions, an initial ERD in alpha power was observed shortly after the onset of the stimulus at 0.25 s and lasted until shortly

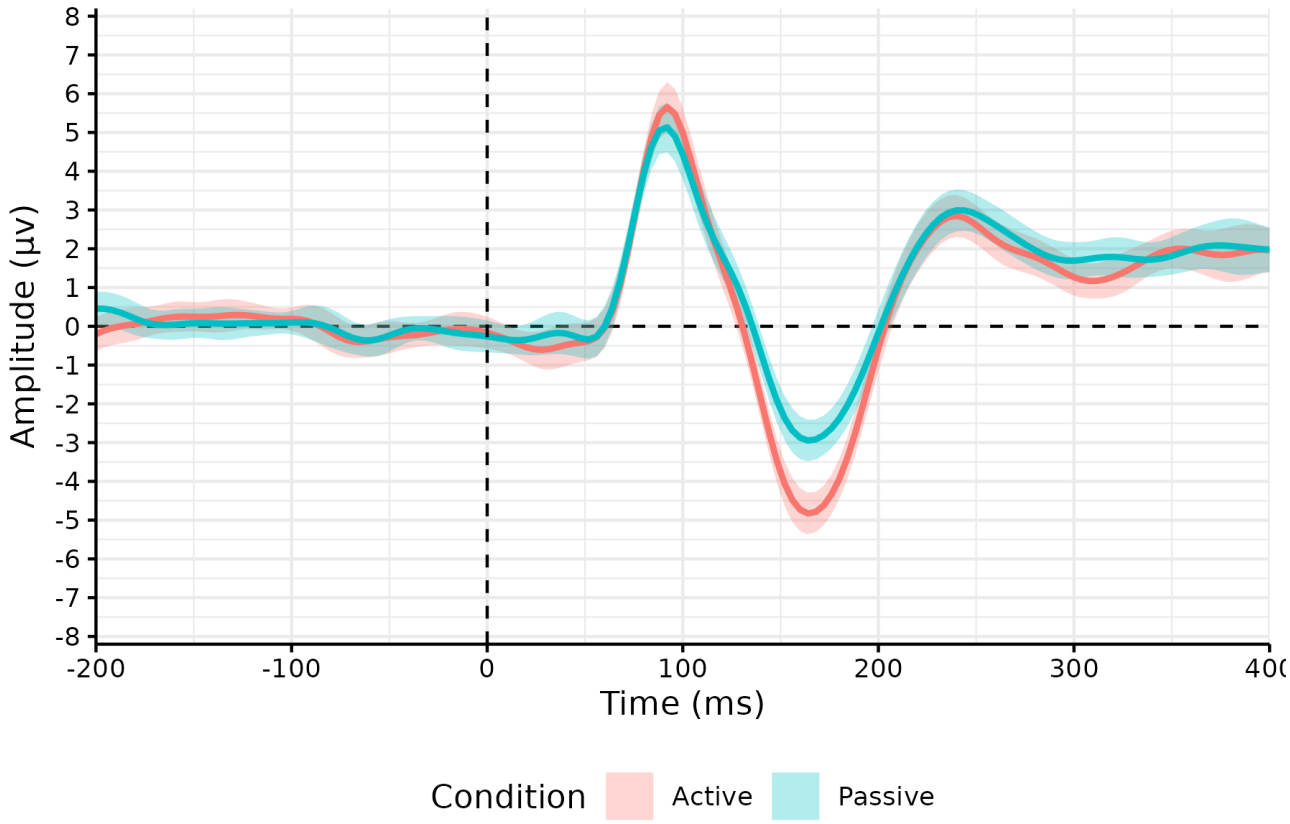


Figure 5.1: Comparison of the grand average VEPs between the active and passive conditions. Error represents repeated measures 95% confidence intervals.

after the offset of the stimulus at 1.2 s. After the initial alpha desynchronisation, a rebound in alpha power was observed, resulting in increased alpha synchronisation between 1.8 s and 4.5 s. Though no difference in alpha power was observed during the initial desynchronisation across conditions, differences were observed during the alpha rebound, with the rebound being notably stronger during the active condition in comparison to the passive condition (see Figure 5.3).

Comparisons across stimulus types showed similar results, with all stimulus types reliably evoking an initial alpha desynchronisation during the presentation of the stimulus, followed by a rebound in alpha synchronisation following the offset of the stimulus (see Figure 5.4). Both standard and target stimuli elicited a larger rebound in alpha synchronisation during the active condition compared to the passive condition. However, no differences in the strength of the

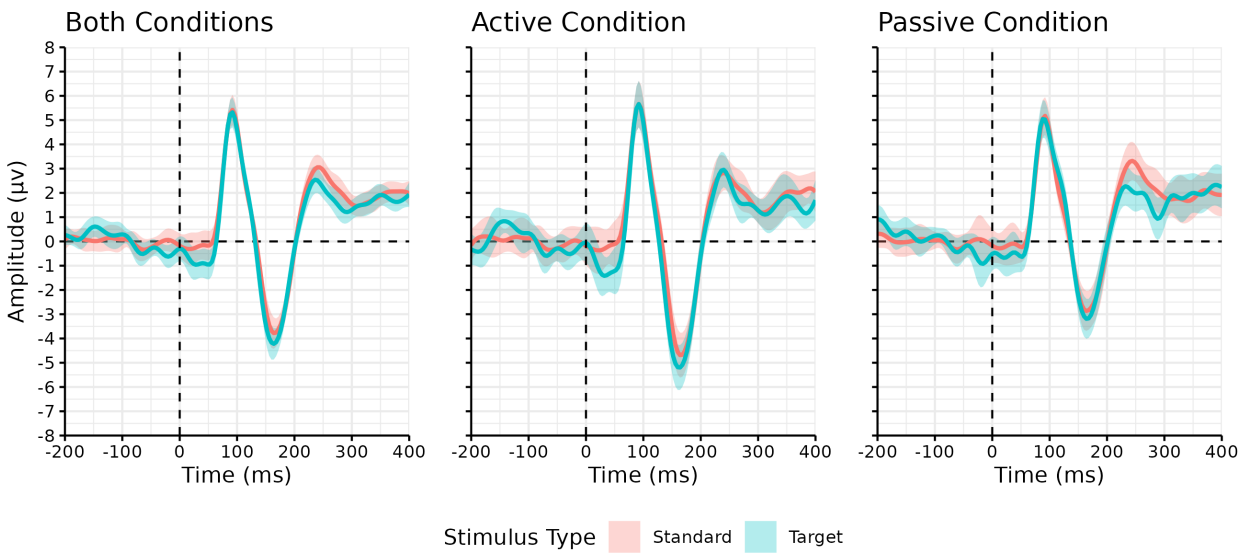


Figure 5.2: Comparison of the grand average VEPs for standard and target stimuli. The left plot shows VEPs with trials included from both conditions, whilst the centre and right plots only include trials from the active and passive conditions, respectively. Error represents repeated measures 95% confidence intervals.

alpha desynchronisation or synchronisation were observed between stimulus types during the active or the passive conditions. The time series comparison of the alpha ERD/S showed that changes in alpha power were notably more variable for target stimuli than for standard stimuli. This increase in variability is potentially due to the limited number of target trials (25%) compared to standard trials (75%).

HMM Model Selection

As we were only concerned with modelling changes in power within the alpha band over right occipital cortex, a two-state model solution was chosen, allowing one state to represent periods of a high alpha power, whilst the other models the absence of alpha, representing a low alpha state. To ensure the reliability of using a two-state model solution, the HMM inference process was repeated a total of 10 times, and the maximum fractional occupancy values and state time course correlation values were calculated (see Figure 5.5. The state time course correlation

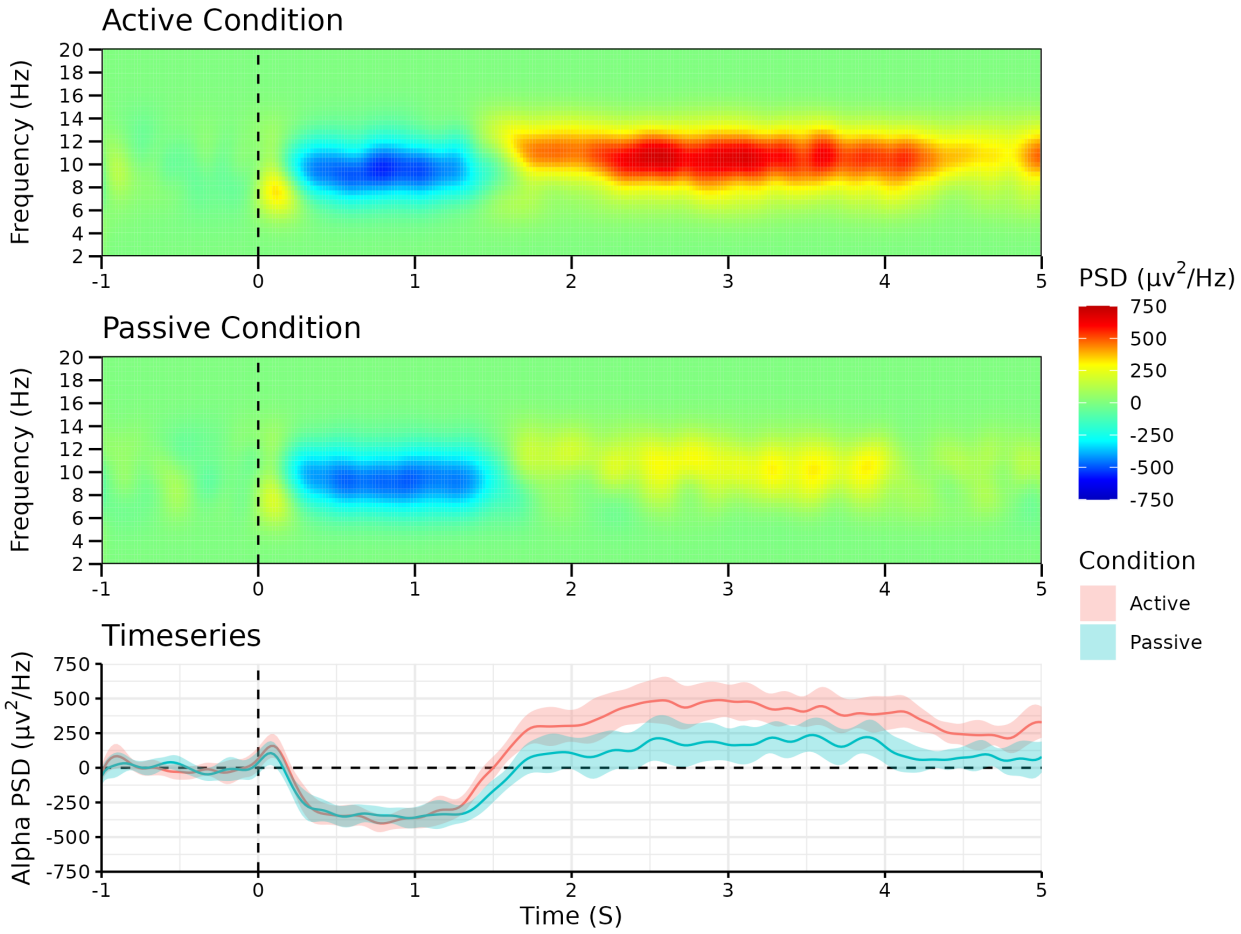


Figure 5.3: Grand average ERD/S plots for the active condition (top plot) and passive condition (middle plot). The bottom plot shows the time series comparison of the two conditions. The time series was created by averaging the values from the above plots across the 7-13Hz band. Error bars represent repeated measures 95% confidence intervals. PSD = power spectral density.

values were indistinguishable, all approaching one, averaging at an r value of 0.99. As the state time course correlations were so high, the runs were essentially identical, and thus, the maximum fractional occupancy distributions for each run were also identical. As the models outputted from the ten runs were identical, the first run model was selected for further analysis.

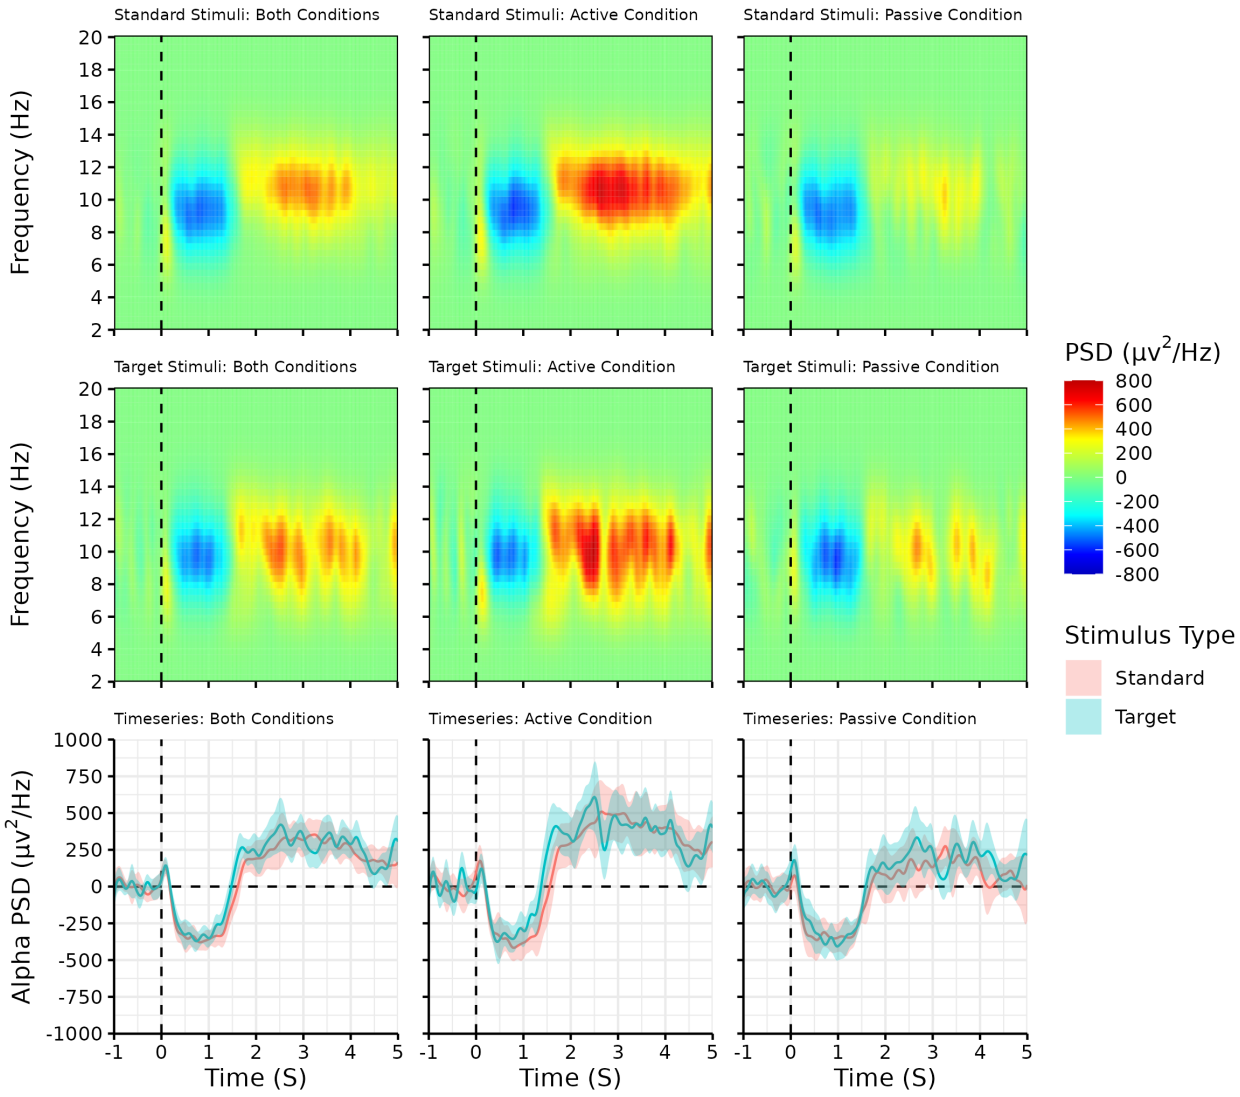


Figure 5.4: Grand average alpha ERD/S plots for standard and target stimuli. Each column represents the trial grouping, with the leftmost column including trials from both conditions, the centre column including trials exclusively from the active condition, and the right column including trials exclusively from the passive condition. The bottom row shows the time series comparisons of the ERD/S plots above them, averaged within the 7-13Hz frequency range. Error represents repeated measures 95% confidence intervals. PSD = power spectral density.

HMM States

Each state's spectral and spatial properties were calculated to determine which state represented the 'high alpha' state. Results from the state-wise multitaper analysis showed a clear separation

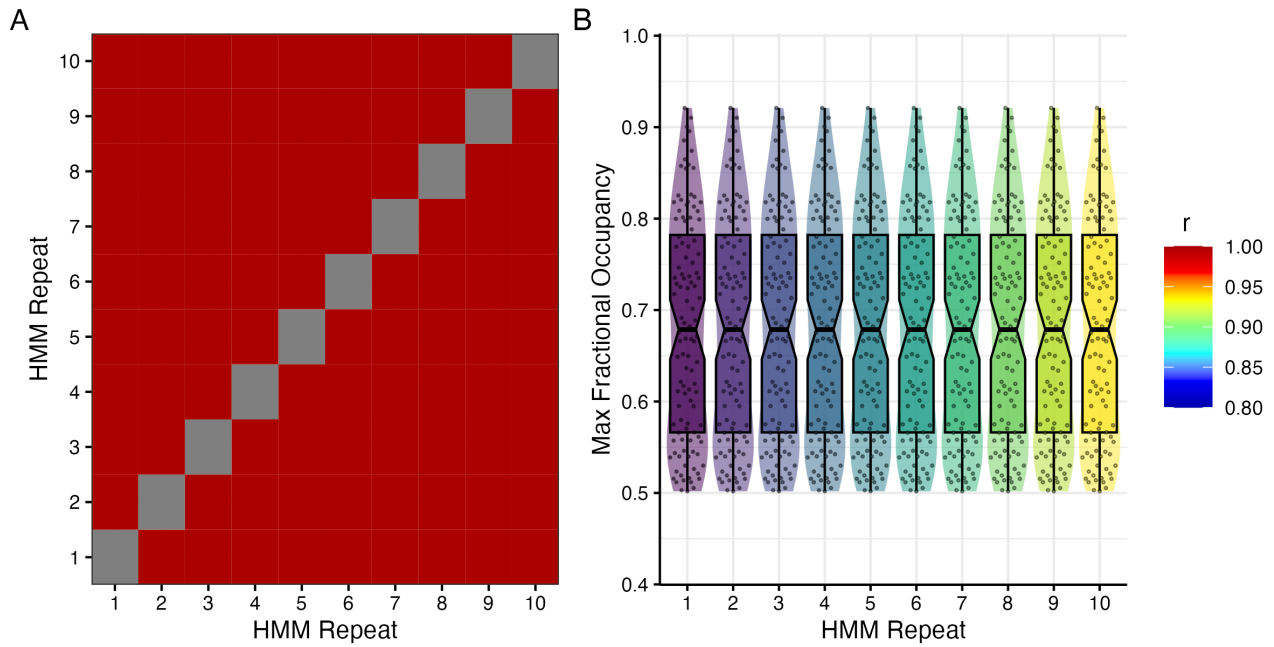


Figure 5.5: Model selection plots for the alpha HMM. Plot A shows the reliability matrix across repeats of the HMM inference process. Plot B shows the maximum fractional occupancy distribution for each HMM repeat. Each point represents the max fractional occupancy from a single run.

between the two states' spectra, with state two showing the highest power peak in the alpha frequency band, and state one plateauing at near zero power (ee Figure 5.6). Thus, throughout the rest of this chapter, state two is referred to as the high alpha state, and state one is referred to as the low alpha state. Though the results showed a clear distinction between the two states, the spectrum of the high alpha state showed a large amount of variance, with artefacts residual gradient artefacts present at 8, 9, 10, 11, and 12 Hz.

The spatial properties of the high alpha state were determined by using its state time course within a fMRI GLM analysis (see Figure 5.7). As there were only two states within the HMM, the low alpha state was not investigated directly, as its BOLD correlates would be the inverse of the high alpha state BOLD correlates. The analysis found significant ($p < 0.05$, Bonferroni corrected) group-level negative correlations across the cortex, including the majority of bilateral occipital cortex, intraparietal sulcus (IPS), insula cortex, and superior frontal cortex, as well

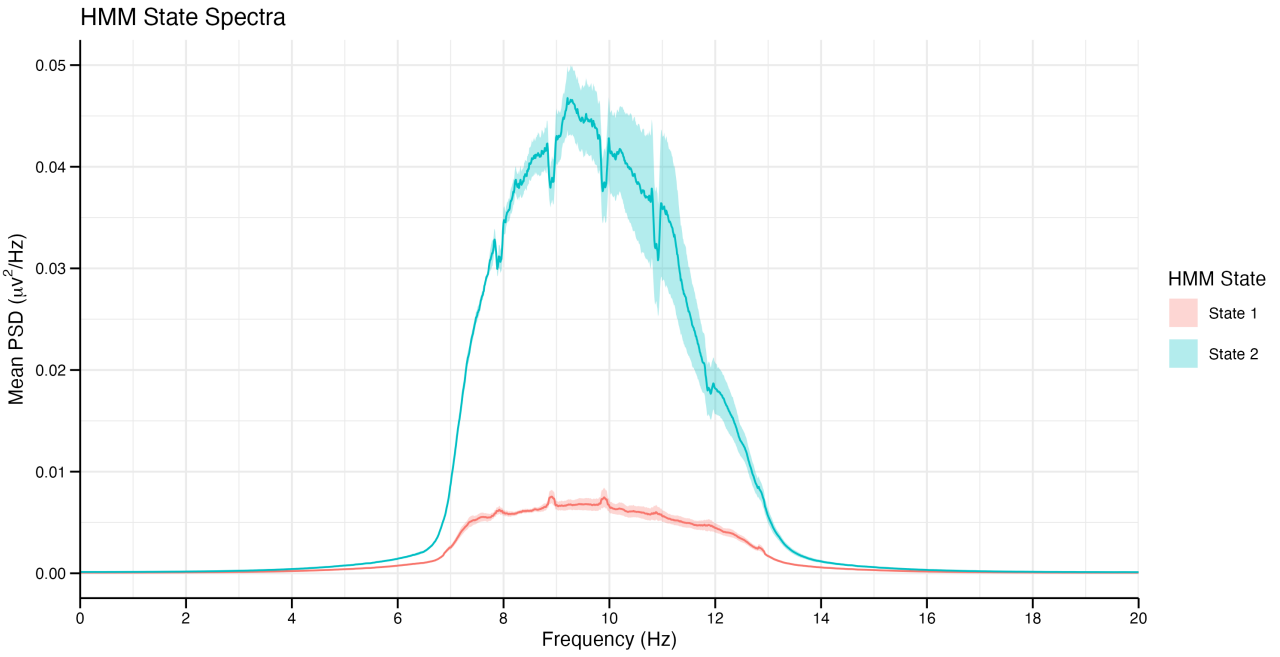


Figure 5.6: Estimated spectra of the two HMM states averaged across all posterior electrodes (O2, P2, P4, P6, P8, PO2 and PO4). Error represents repeated measures 95% confidence intervals. PSD = power spectral density.

as within the anterior cingulate cortex. Significant ($p < 0.05$, Bonferroni corrected) negative correlations were also observed subcortically within the posterior thalamus and the putamen. Significant positive correlations were not observed in the cortex, only within the right midline of the anterior thalamus.

Evoked HMM Response

The mean evoked fractional occupancy responses to the stimuli was calculated across all subjects (see Figure 5.8). As only two states were included in the model, the fractional occupancies are the exact opposite of each other, and thus, the results are discussed in regard to the high alpha power state only. The results showed a clear decrease of 10% in the high alpha states' fractional occupancy following stimulus onset. After stimulus presentation (1s), the fractional occupancy rebounded to a positive increase of 10% compared to baseline, before gradually returning to

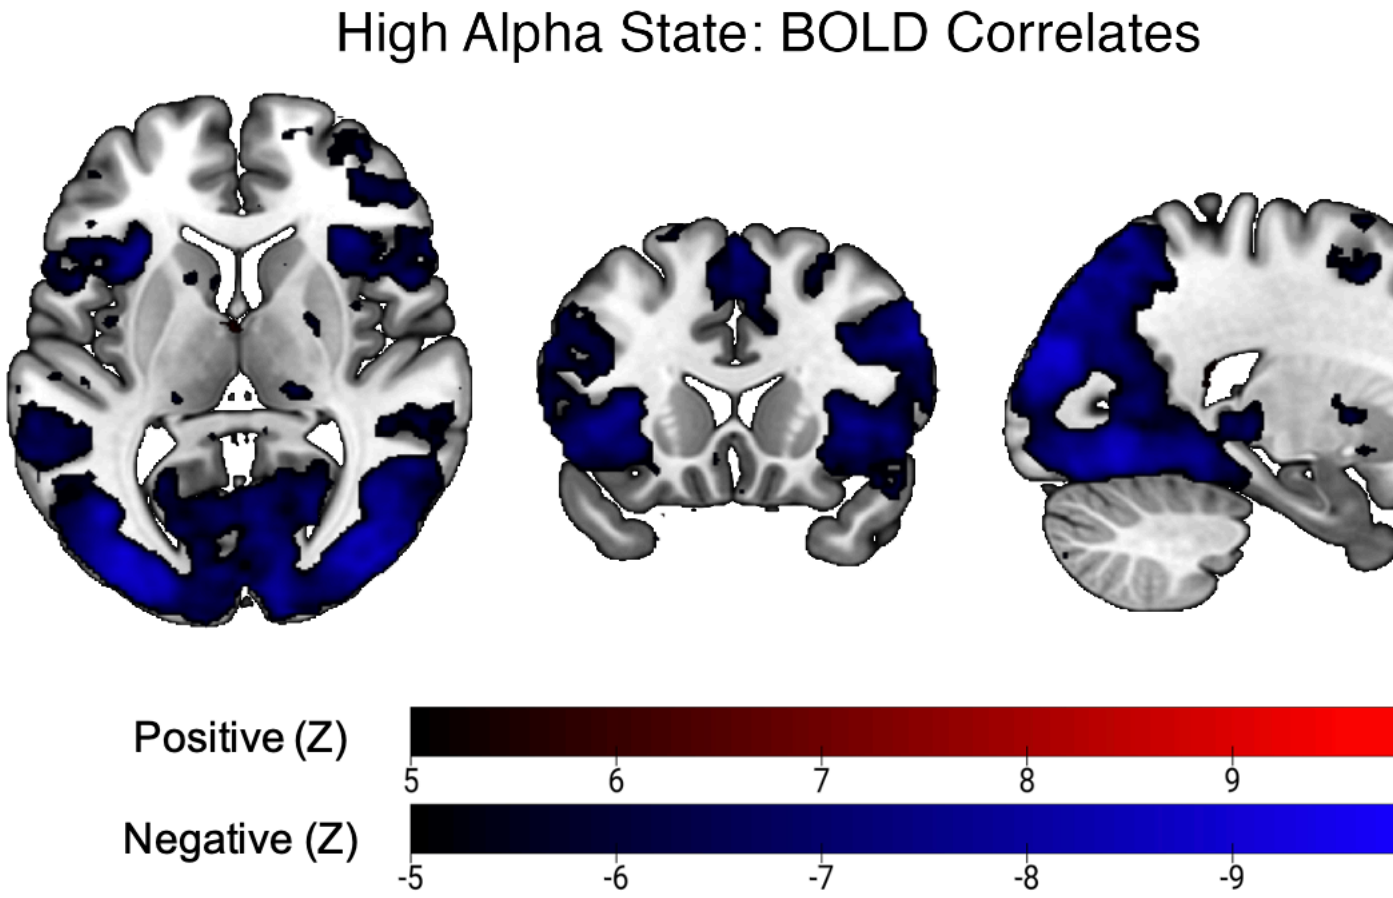


Figure 5.7: Group-level BOLD signal correlates with the high alpha state determined by the GLM analysis. The activation map is Bonferroni corrected ($p<0.05$). Coordinates reprsent MNI152 coordinates. Image presented in MNI152 space using neurological conventions (left = left)

baseline over approximately 12 seconds. These results mirror the pattern of alpha power changes observed within the ERD analysis, with the initial decrease in fractional occupancy matching the latency of the initial alpha ERD, followed by the rebound or alpha ERS. However, unlike the ERD/S plots, the alpha state fractional occupancy returned to baseline at a much slower rate, only returning to baseline approximately 12s after the onset of the stimulus.

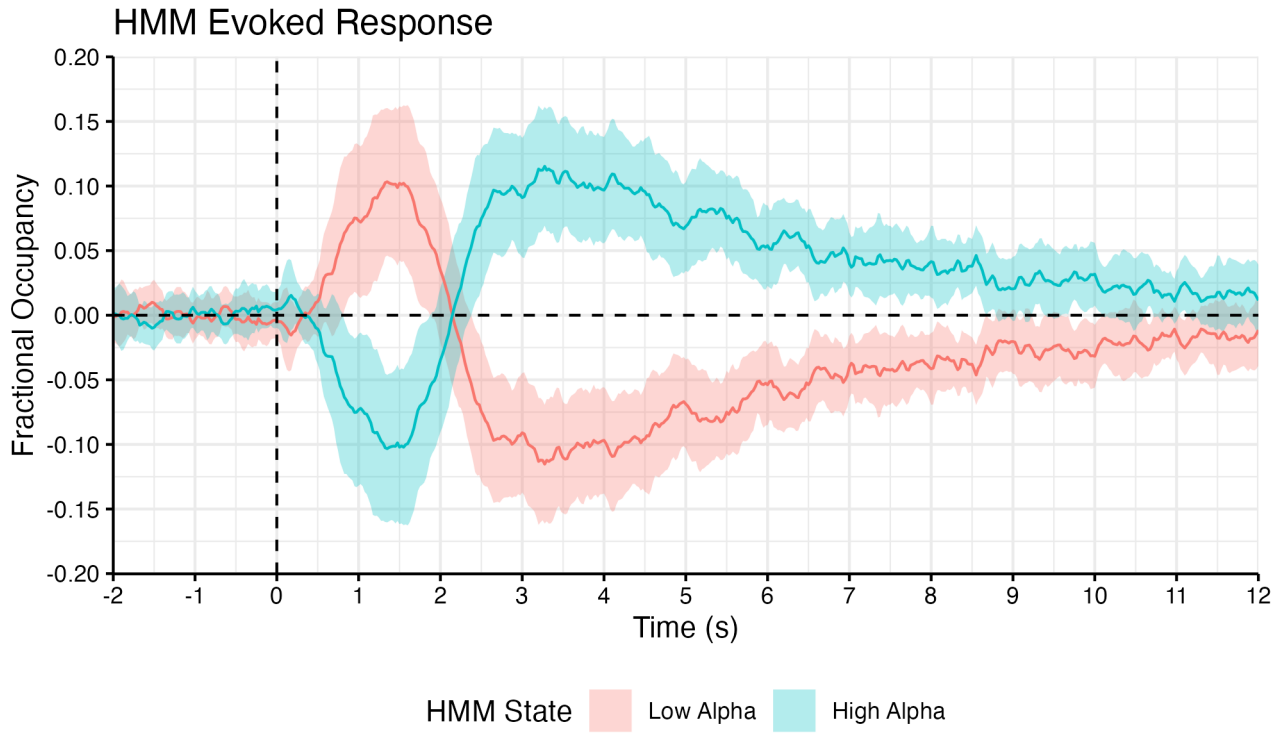


Figure 5.8: Group average changes in state fractional occupancy upon stimulus presentation. Error represents repeated measures 95% confidence intervals.

Distribution and Reliability of Baseline States

To evaluate the likelihood of a specific HMM state being active during the baseline of a trial, the number of baseline state occurrences was counted for each subject (see Figure 5.9). The number of trials where each HMM state was active at baseline was then compared across conditions to observe whether either HMM state was more likely to be active at baseline. The results showed that the low alpha state was active at baseline for most trials, with a median of 53 trials during the active condition, and 47 trials during the passive condition. Alternatively, the high alpha was only active at baseline for a smaller proportion of trials, with a median of 19 trials during the active condition, and a median of 24 during the passive conditions. Both the low and high alpha states showed a bias towards a specific condition. The difference in baseline state frequency was assessed statistically using a chi-square test of independence. The association

between baseline HMM state and condition was significant ($X^2(1) = 20.373, p < 0.01$), showing that the high alpha state was significantly less likely to be active during the baseline of trials during the active condition compared to the passive condition, and vice versa for the low alpha state.

The reliability of using timepoint t-1 to define which state was active prior to stimulus presentation was evaluated by comparing the active baseline states at t-0, to t-5 and correlating the results to generate a correlation matrix (see Figure 5.9). The results showed strong correlations between baseline states determined using timepoints t0, t-1, and t-2.

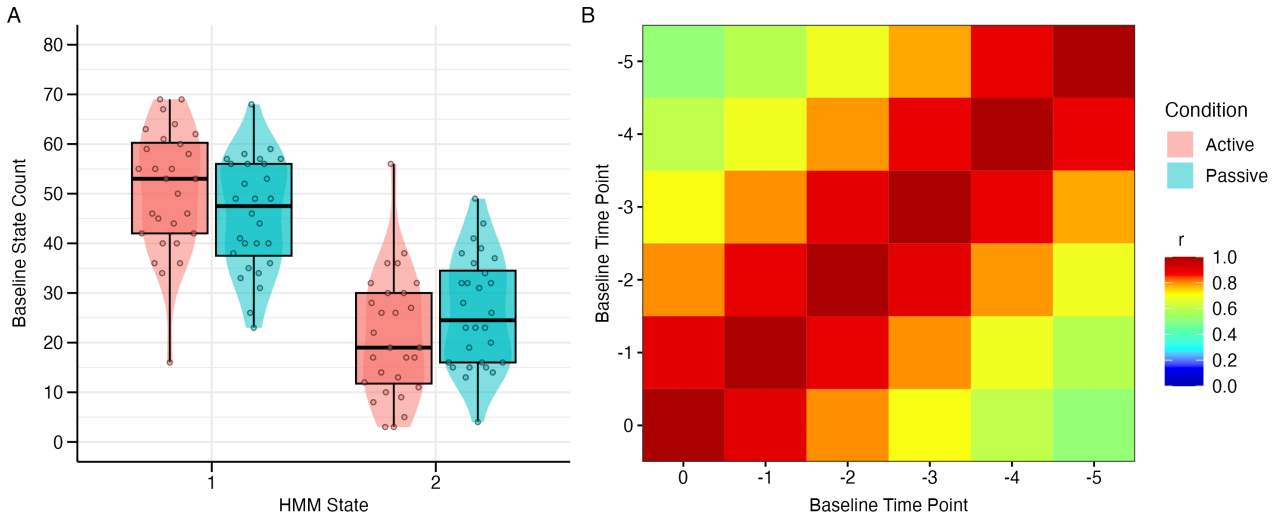


Figure 5.9: Plot A shows the number of times each state was active during the baseline for each condition. Each point represents the number of times a state was active for that condition (combined across runs) and state. Plot B shows the reliability of the baseline states based on whether the baseline state is defined at the 0 – -5 timepoint.

State Metric Comparison

State metrics were calculated over the entirety of each run for each state and compared across conditions to investigate differences in the temporal dynamics of alpha (see Figure 5.10). Fractional occupancy was largest for the low alpha power state, with a median fractional occupancy of approximately 0.67 (67%) during both the active and passive conditions. In comparison, the

high alpha state had a lower median fractional occupancy of approximately 0.33 (33%) across both conditions. Differences in fractional occupancy between conditions were compared statistically with paired t-tests. There were no significant differences in fractional occupancy for the low alpha state (difference = 0.01, 95% CI [-0.02, 0.04], t(27) = 0.99, p = 0.331; Cohen's d = 0.19, Cohen's d 95% CI [-0.19,0.56]) or the high alpha state (difference = -0.01, 95% CI [-0.04, 0.02], t(27) = -0.99, p = 0.331; Cohen's d = -0.19, Cohen's d 95% CI [-0.56, 0.19]). Similarly, switching rates were also unaffected by conditions, with a median switching rate of approximately 0.4 for both conditions. This difference was tested statistically using a paired t-test, which found no significant difference in the switching rate between the two conditions (difference = -7.47e-04, 95% CI [-3.48e-03,1.98e-03], t(27) = -0.56, p = 0.579; Cohen's d = -0.11, Cohen's d 95% CI [-0.48, 0.27]). State lifetimes were, on average, longer for the low alpha state relative to the high alpha state, with the low alpha state having a median state lifetime of approximately 700ms, and the high alpha state having a median state lifetime of approximately 375 ms. No difference in state lifetimes was observed between conditions for either state. As only two states were included within the HMM, the state interval times show the inverse of the state lifetimes i.e. the state lifetimes for the low alpha state are the state interval times for the high alpha state.

Impact of Alpha State on VEPs

In order to evaluate the impact of the HMM state active at baseline, trials were grouped according to baseline HMM state and VEPs were generated for both conditions combined and for each condition separately (see Figure 5.11). We observed no overall impact on the early VEP components but some evidence of an impact on the later components, i.e. P250 onwards. Specifically, trials where the high alpha state was active at baseline showed larger P250 amplitudes that continued to the end of the epoch. The effect was clearest when including trials from both conditions and for trials within the active condition. Though the effect was present

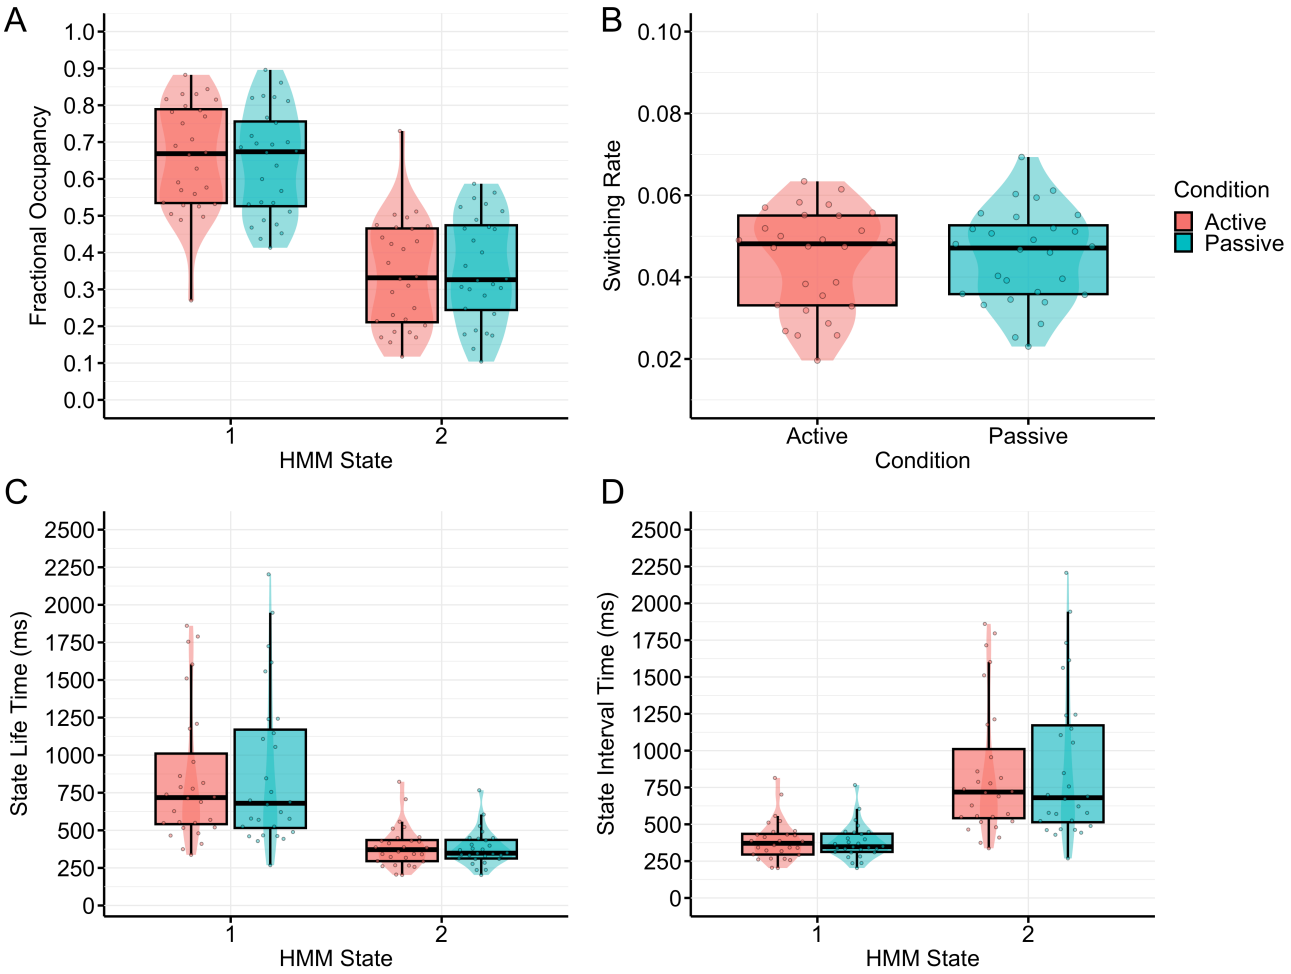


Figure 5.10: Comparison of HMM states metrics between the active and passive conditions. Plot A compares max fractional occupancy. Plot B compares the state switching rate. Plot C compares the average state lifetimes. Plot D compares the average state interval times. Each point represents a single subject.

within the passive condition, the VEPs from this condition showed larger variance than the active condition.

The impact of alpha state, condition and stimulus type on the VEP component (P100, N150, P250) amplitudes were assessed via linear mixed modelling (see Chapter 2) The results of the N150 model found the condition main effect to be significant ($F(1, 3883.5) = 25.322, p < 0.001$), with component amplitudes being significantly larger during the active condition ($EMM = -5.04$ µV, $SE = 0.646$ µV) relative to the passive condition ($EMM = -2.91$ µV, $SE = 0.639$

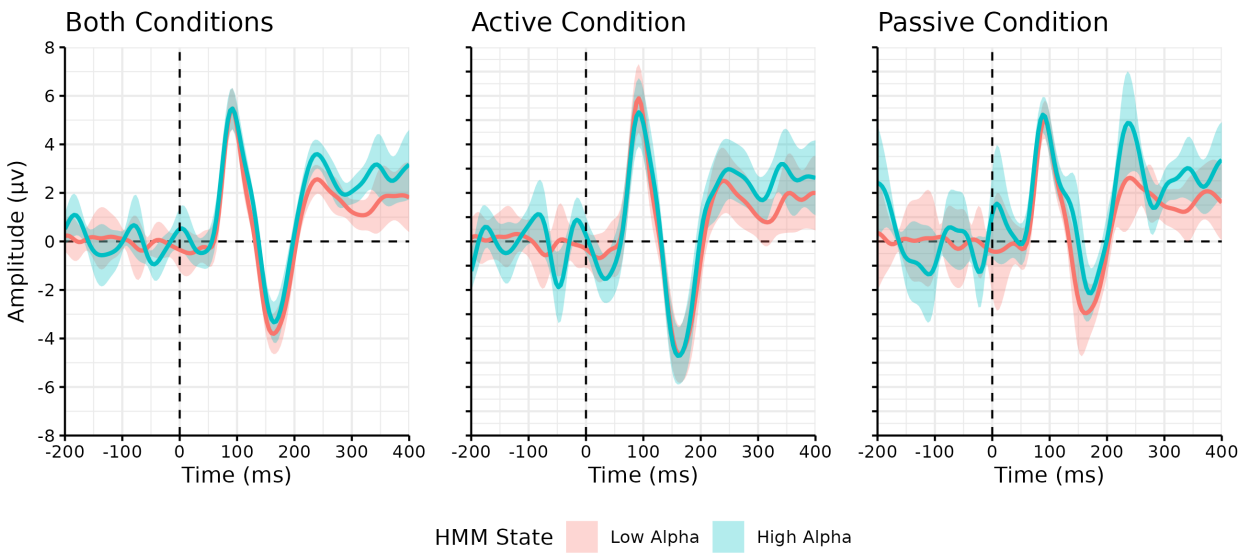


Figure 5.11: Comparison of grand average VEPs for high and low alpha state trials. The left plot shows grand averages across both conditions, while the left and right plots show the grand averages from the active and passive trials, respectively. Error represents repeated measures 95% confidence intervals.

µV). All other main and interaction effects were non-significant ($p > 0.05$).

Impact of Alpha State on Alpha ERD/S

Similarly to the VEP analysis, trials were grouped according to the HMM state active at baseline, and the changes in alpha power were calculated for all conditions together and separately for each condition (see Figure 5.12). This showed that when the low alpha state was active at baseline, only the alpha rebound was observed. This effect was strongest when only averaging trials from the active condition but was also present when averaging trials across both conditions, and weakly present when only averaging trials from the passive condition. Conversely, when the high alpha state was active at baseline, only the initial alpha desynchronisation was observed with minimal to no rebound in alpha synchronisation. Unlike the low alpha state, this effect did not visibly differ across conditions, with the strength of the alpha desynchronisation remaining the same irrespective of how trials were grouped across conditions. Comparing the time series

of the alpha power, there was a clear and significant difference between the two states. Like the VEPs, the ERD/S plots for the passive condition also showed a larger amount of variance relative to the active condition, particularly for the high alpha state trials. This is likely due to the fact there were fewer trials where the high alpha state was active at baseline compared to the low alpha state.

The impact of alpha state, condition and stimulus type on the alpha ERD and ERS were assessed statistically using linear mixed modelling (see Chapter 2 for details). The results of the ERD model found the alpha state main effect to be significant (($F(1, 3806.3) = 184.831, p < 0.001$), with the ERD amplitudes being significantly larger for trials preceded by the high alpha state ($EMM = -245.2$ µV^2/Hz, $SE = 26.8$ µV^2/Hz) relative to the low alpha state ($EMM = -48.6$ µV^2/Hz, $SE = 25.4$ µV^2/Hz). Similarly, the ERS model also found the alpha state main effect to be significant ($F(1, 3772.6) = 111.652, p < 0.001$), with the ERS amplitudes being significantly larger for trials preceded by the low alpha state ($EMM = 162.2$ µV^2/Hz, $SE = 19.2$ µV^2/Hz) relative to the high alpha state ($EMM = -25.7$ µV^2/Hz, $SE = 21.9$ µV^2/Hz). Additionally, the condition main effect was also significant ($F(1, 3783.6) = 16.745, p < 0.001$), with the alpha ERS being significantly larger during the active condition ($EMM = 103.3$ µV^2/Hz, $SE = 20.7$ µV^2/Hz) relative to the passive condition ($EMM = 33.2$ µV^2/Hz, $SE = 20.3$ µV^2/Hz). A significant interaction was also found between condition and alpha state ($F(1, 3793.1) = 3.852, p < 0.05$). All other main and interaction effects were non-significant ($p > 0.05$). Post hoc tests (Bonferroni corrected, number of comparisons = 6) revealed that the effect of alpha state was significant during both the active ($t(3804) = 8.717, p < 0.001$) and passive conditions ($t(3808) = 6.415, p < 0.001$), but was more pronounced during the active condition (EMM Difference= 221.64 µV^2/Hz) relative to the passive condition (EMM Difference= 154.14 µV^2/Hz).

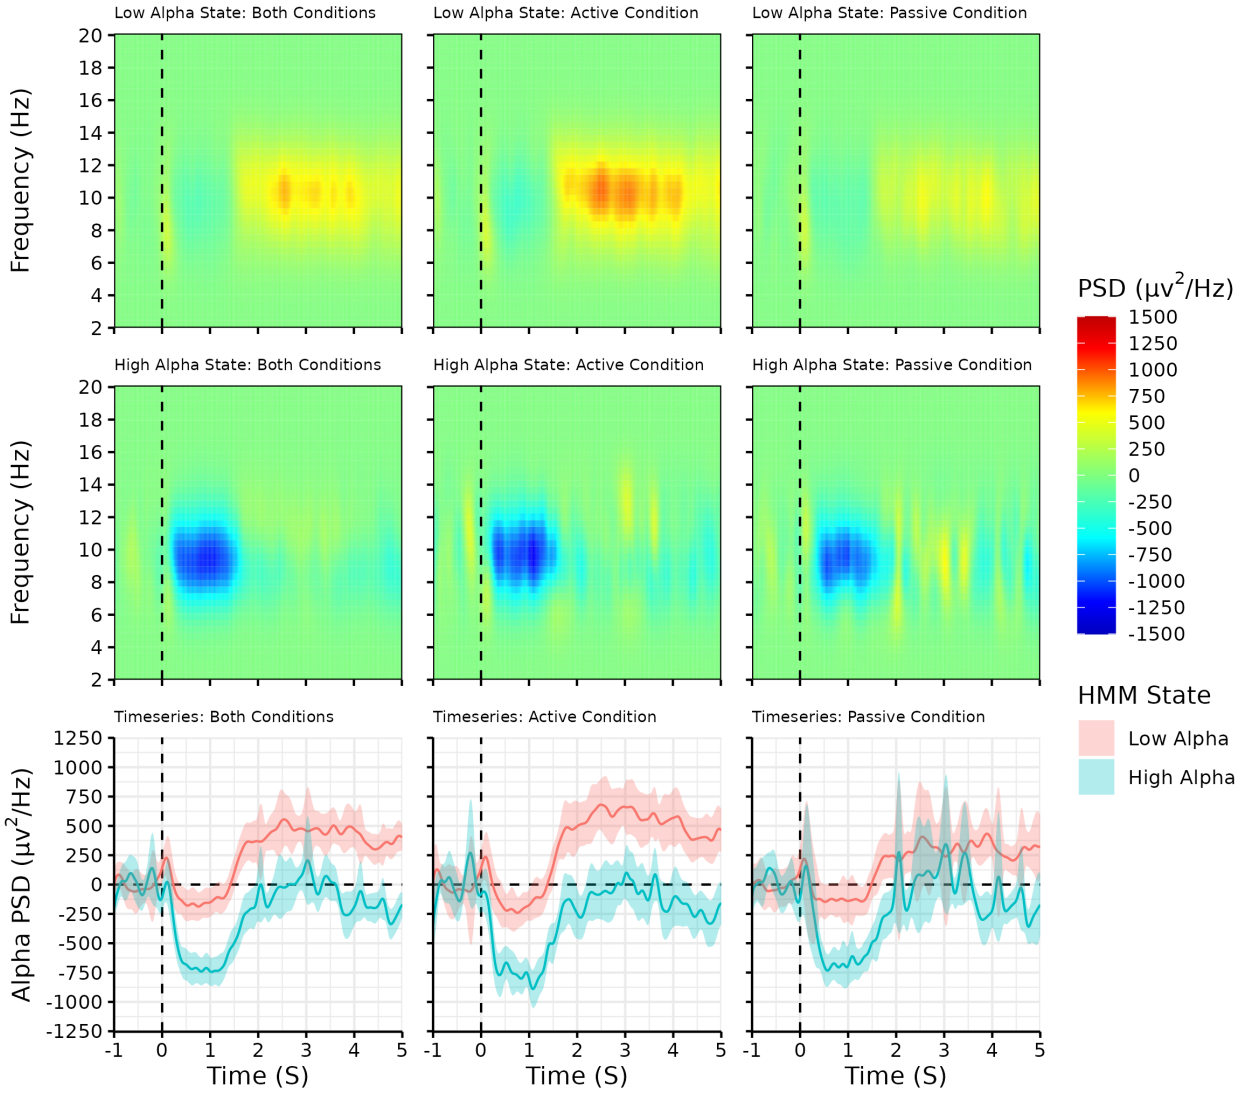


Figure 5.12: Grand average alpha ERD/S for high and low alpha state trials. Each column represents the trial grouping, with the leftmost column including trials from both conditions, the centre column including trials exclusively from the active condition, and the right column including trials exclusively from the passive condition. The bottom row shows the time series comparisons of the ERD/S plots above them averaged within the 7-13Hz window. Error represents repeated measures 95% confidence intervals. PSD = power spectral density.

Impact of Alpha State on BOLD Responses

The impact of alpha state, condition and stimulus type on the BOLD peak and undershoot PSC values in right V1, insula and LGN were assessed statistically using linear mixed modelling (see

Chapter 2 for details) (see Figure 5.13). The results of the right V1 BOLD peak model found the condition main effect to be significant, with BOLD PSC values being significantly larger ($F(1, 3995.9) = 21.428, p < 0.001$) during the active condition ($EMM = 1.80\%, = 0.117\%$) relative to the passive condition ($EMM = 1.52\%, SE = 0.117\%$). Additionally, a significant interaction effect was found between condition and stimulus type for the peak BOLD model ($F(1, 4141.9) = 9.901, p < 0.01$). Post hoc tests (Bonferroni corrected, number of comparisons = 6) revealed that the target stimuli ($EMM = 1.92\%, SE = 0.133\%$) generated significantly larger BOLD PSC values than standard stimuli ($EMM = 1.68\%, SE = 0.117\%$) during the active condition ($t(3995) = -2.804, p < 0.05$), but not during the passive condition ($t(3995) = 1.620, p = 0.632$). A significant interaction effect between condition and stimulus type was also found for the right V1 undershoot model ($F(1, 3999.9) = 10.805, p < 0.01$). The post hoc tests (Bonferroni corrected, number of comparisons = 6) revealed similar results to the peak model, with target stimuli ($EMM = 0.147\%, SE = 0.104\%$) generating significantly larger BOLD PSC values than standard stimuli ($EMM = -0.182\%, SE = 0.061\%$) during the active condition ($t(4003) = -2.748, p < 0.05$), but not the passive condition ($t(4000) = 1.882, p = 0.359$). All other main and interaction effects were non-significant ($p > 0.05$)

Similarly, the right insula peak BOLD model found the condition main effect to be significant (($F(1, 3998.0) = 201.553, p < 0.001$), with the BOLD PSC values being significantly larger in the active condition ($EMM = 0.442\%, SE = 0.026\%$) relative to the passive condition ($EMM = 0.128\%, SE = 0.026\%$). Additionally, the stimulus type main effect was also found to be significant ($F(1, 3996.6) = 15.57, p < 0.001$), with target stimuli ($EMM = 0.328\%, SE = 0.028\%$) generating significantly larger BOLD PSC values relative to standard stimuli ($EMM = 0.241\%, SE = 0.023\%$). Lastly, the right LGN peak BOLD model found the condition main effect to be significant ($F(1, 4004.7) = 9.357, p < 0.01$), with BOLD PSC values being significantly larger during the active condition ($EMM = 0.143\%, SE = 0.012\%$) relative to the passive condition ($EMM = 0.096\%, SE = 0.012\%$). All other main and interaction

effects were non-significant ($p > 0.05$).

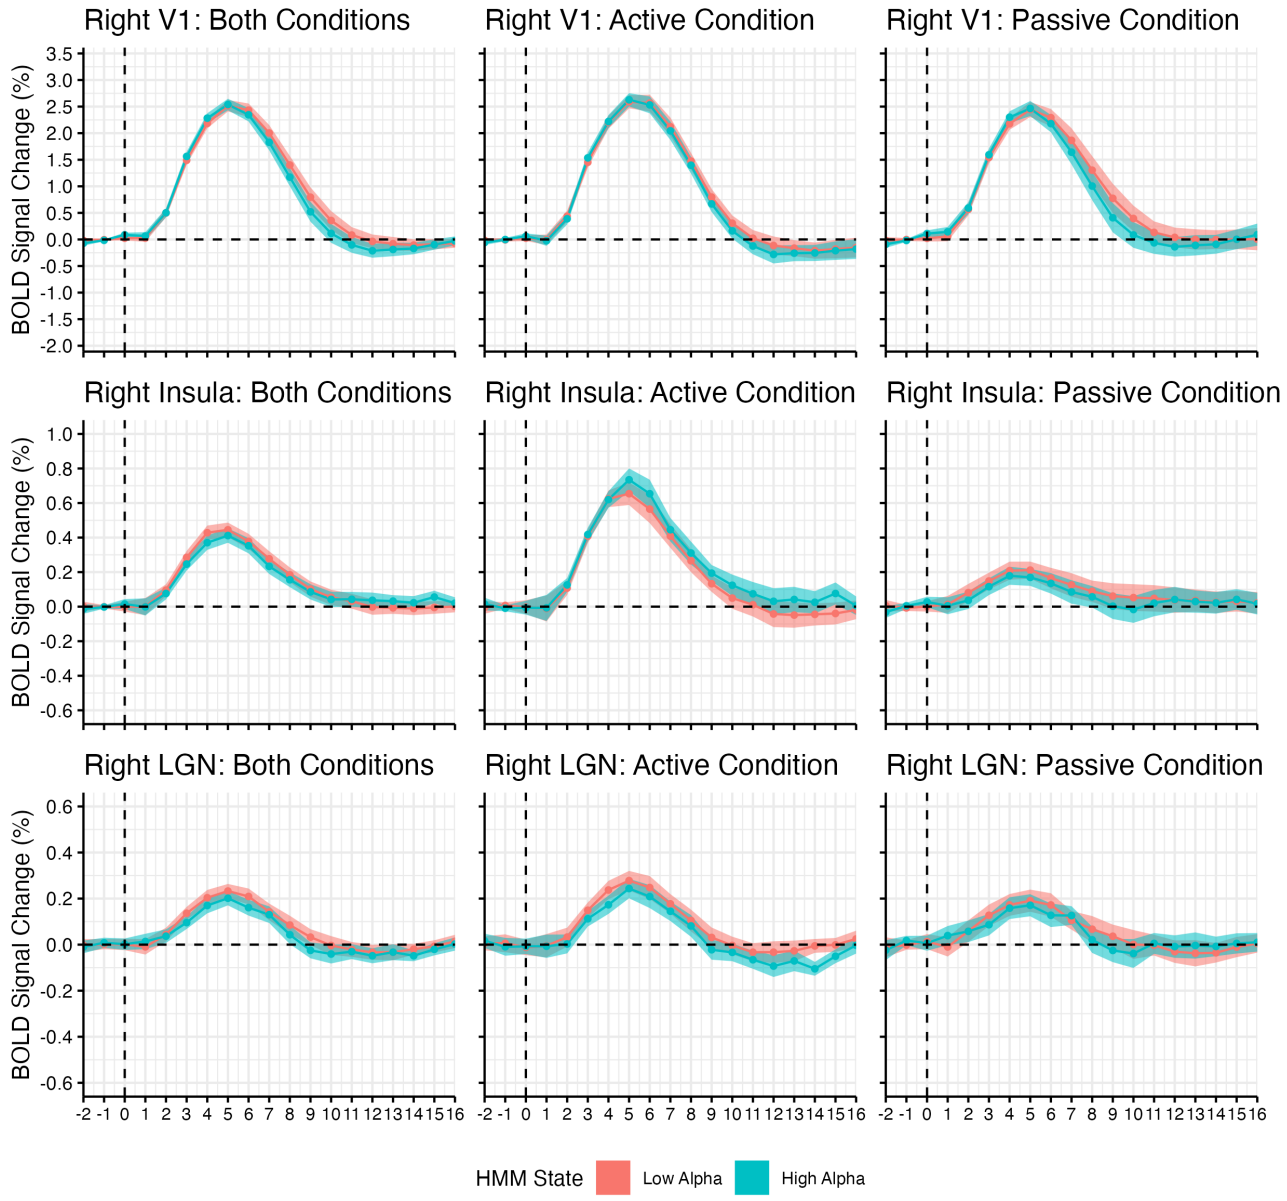


Figure 5.13: Grand average BOLD responses for high and low alpha state trials. Each column represents trials grouping (left column: trials from both conditions, centre column: trials from the active condition, and right column: trials from the passive condition), whilst each row represents the region of interest (top row: right V1, centre row: right insula, and bottom row: right lateral geniculate nucleus). Error represents repeated measures 95% confidence intervals.

Discussion

Summary

This chapter employed an HMM approach to identify periods of high alpha power during both passive viewings of visual stimuli (passive condition) and whilst completing a simple sustained attention task (active condition), with the aim of investigating the impact of alpha power state on evoked electrophysiological and BOLD responses. The average electrophysiological responses (VEP and ERD/S) were calculated and compared between active and passive conditions, revealing significant differences in the N150 component and alpha rebound, respectively. A two-state HMM was trained, underwent model selection, and then underwent a state-wise spectral analysis to determine which of the two states represented high alpha power. In addition, a GLM analysis confirmed the spectral analysis results, showing that the high alpha power state showed negative BOLD correlates within bilateral lateral occipital cortex, IPS, and the salience network. The HMM state metrics were compared across conditions to determine whether the alpha state was more likely to be active during either of the conditions, but no significant differences were found in the fractional occupancy or switching rate across conditions. The Viterbi path of the HMM states was used to group all trials according to whether the low or high alpha state was 'active' prior to stimulus presentation, and the average VEP, alpha ERD/s and BOLD responses were calculated and compared. Alpha ERD primarily occurred in trials preceded by the high alpha state, whilst the alpha ERS primarily occurred in trials preceded by the low alpha state, suggesting that the two phenomena are generated exclusively from different trials. No significant differences were observed between high and low alpha states for the BOLD responses or VEPs. Though no significant difference was observed in the state metrics across conditions, trials were significantly more likely to be preceded by the high alpha state during the passive condition than the active condition. However, the ERS amplitude was still higher

for low alpha state trials during the active condition relative to the low alpha state trials within the passive condition, suggesting that the difference in ERS amplitude across conditions was not a consequence of pre-stimulus alpha power.

Visually Evoked Potentials and Alpha ERD/S Differ Between Conditions

Both the VEP and ERD/S analyses found significant differences in responses between conditions. The VEP analysis found that the N150 component peak amplitude was significantly lower during the active condition compared to the passive condition. This finding is supported by the literature, in which the N150 component has been demonstrated to be modulated by spatial attention (Heinze et al., 1990; Hillyard & Anllo-Vento, 1998; Johannes et al., 1995; Luck et al., 1990; Mangun, 1995), as well as source localisation approaches which have localised the subcomponents of N150 to extrastriate cortex (150ms) and parietal cortex (>155ms) (Di Russo et al., 2002). However, unlike in previous literature, the amplitude of the P100 component showed no significant difference between conditions, although it was slightly higher during the active condition relative to the passive condition.

Regarding the alpha ERD/S results, an alpha ERD was observed shortly after the onset of the stimulus, lasting approximately 1 second, followed by an alpha ERS at 1.7 seconds, lasting approximately three seconds. This was consistent with the results of previous research on visual alpha ERD/S (Brookes et al., 2005; Mayhew, Ostwald, et al., 2013; Woertz et al., 2004). Similarly to the VEP analysis, significant differences in alpha ERD/S were observed between the two conditions, with responses during the active condition having significantly higher alpha rebounds compared to the responses during the passive condition. This finding also fits with the current alpha ERD/S literature, given that an increase in alpha power is hypothesised to represent functional inhibition, facilitating processing within other areas of the cortex (Jensen

& Mazaheri, 2010; Klimesch et al., 2007; Pfurtscheller et al., 1996). Therefore, the larger increased alpha ERS during the active condition may represent increased inhibition, allowing for other areas of the cortex (e.g.prefrontal cortex) to perform higher-order tasks, which in this experiment may represent the decision to press the button. However, this difference in the alpha ERS amplitude is difficult to contrast with the literature, as most studies typically employ much shorter ISIs and do not look at event-related changes over longer periods (i.e. several seconds).

HMM States

The HMM reliably identified two states, with the first representing periods of low alpha power and the second representing periods of high alpha power. The state-wise spectral analysis identified the increase in higher levels of alpha power during the periods where state two was active. There was also evidence of residual MR-related artefacts (see Figure 5.6), such as harmonics of the ballistocardiogram artefact or gradient artefact. However, despite this, the HMM was still able to identify the periods of high alpha power, which, combined with the results of Chapter 3, suggests that HMMs can reliably identify periods of high alpha power from relatively noisy EEG data collected within the MR environment. The high alpha state showed both cortical and subcortical BOLD correlates that closely resembled the analysis within Chapter 3, and within previous literature (Goldman et al., 2002; Laufs et al., 2006; Mayhew & Bagshaw, 2017). The HMM approach provides a novel method for sorting individual trials according to high and low alpha power without setting an arbitrary window or sorting by quartiles as in previous literature (Mayhew, Ostwald, et al., 2013).

HMM State Metric Comparison

Both the fractional occupancy and state switching rate showed no significant differences between the active and passive conditions, suggesting that the alpha power dynamics were similar across the two conditions despite the sustained attention element. At first glance, this result

contradicts the previous findings of the literature, which have consistently found that alpha power reduces during visual attention tasks (Kelly et al., 2009; Rihs et al., 2007; Thut et al., 2006; Worden et al., 2000). However, though the HMM state dynamics did not differ between conditions, the high alpha state was significantly more likely to be active during the baseline of a trial in the passive condition relative to the active condition. This difference indicates that although the overall alpha power dynamics were not modulated by task, the increased attention during the active condition still reduced alpha power at baseline. This likely occurred due to participants becoming entrained with the length of the inter-stimulus interval (ISI), allowing them to preempt the onset of the next stimulus. This would allow the participant to rest during the majority of the intra-trial gaps, resulting in a lapse of attention until just before the next stimulus was anticipated, allowing them to pay attention before the onset of the next stimulus. This explanation is supported by previous literature that has demonstrated that alpha power on average decreases just before the onset of anticipated stimuli (Bastiaansen et al., 2001; Bastiaansen & Brunia, 2001). This also explains why on average the HMM state dynamics were the same between conditions, as across the entire duration of a scan, participants' levels of attention would have been similar across the two conditions.

Impact of Baseline Alpha State on Stimulus Induced Alpha ERD/S

Grouping the ERD/S responses by baseline HMM state found that the alpha desynchronisation was primarily generated by trials in which the high alpha state was active during the baseline. Conversely, the alpha rebound/synchronisation was primarily generated by trials in which the low alpha state was active at baseline. This suggests that the two components (e.g. the alpha desynchronisation and synchronisation) are generated by different trials, with trials exhibiting high alpha power prior to stimulus presentation generating the alpha ERD, and trials with low alpha prior to stimulus onset generating the alpha synchronisation. This shows the importance of accounting for baseline alpha levels, as without sorting the trials by alpha power, it would

appear that, on average, all trials exhibit an alpha desynchronisation and rebound in synchronisation. This further highlights the issues of averaging, where the baseline is assumed to be stochastic and thus irrelevant. In contrast, the results here show that averaging over all trials, irrespective of the baseline, can lead to incorrect interpretations.

Impact of Baseline Alpha State on BOLD Responses

The ROI analysis investigating the impact of the baseline HMM state found no significant differences in the peak amplitude nor shapes of the BOLD response within right V1, right LGN, or right insula. This finding is largely in contradiction with the literature, with previous research demonstrating a significant reduction in the peak amplitudes of the BOLD responses which were prefaced with high alpha power, both within right V1 and right LGN (Becker et al., 2011; Mayhew, Ostwald, et al., 2013). One potential reason for this discrepancy may be the choice of ROI for right V1. Within Mayhew, Ostwald, et al. (2013), only anterior regions of right V1 were modulated by pre-stimulus alpha power, despite showing smaller peak amplitudes relative to the more posterior regions of V1. This distinction between anterior and posterior V1 was determined using a psychophysiological interaction (PPI) GLM approach, allowing for the identification of regions in which visual responses were modulated by alpha power. A similar approach could be adopted for the HMM approach used here, identifying regions in which visual responses are modulated by alpha power to be identified functionally. However, this would not explain why alpha power did not modulate BOLD responses within the LGN. One reason may be the reduced signal-to-noise ratio (SNR) within subcortical regions due to the use of a multi-band sequence (Risk et al., 2018; Risk et al., 2021), though this seems unlikely given that the right LGN both exhibited a BOLD response with similar levels of variability to right V1, and was identified as significant cluster during the GLM analysis within Chapter 4, leaving the cause of this disparity unclear.

Impact of Baseline Alpha State on VEPs

The VEP analysis found no significant impact of the baseline alpha state on any VEP components. However, amplitudes from 200ms onwards were slightly larger for trials preceded by the high alpha state compared to trials preceded by the low alpha state. Though this difference was not significant, the pattern closely matched the results of previous literature, which found that only later VEP components (e.g. >= 200ms) were impacted by pre-stimulus alpha, with trials preceded by higher alpha power generating larger amplitudes (Becker et al., 2008; Reinacher et al., 2009). These results support the findings of Becker et al. (2008), who suggested that alpha power has a modulatory effect on later ERP components as opposed to a causal one, as suggested by the oscillatory theory of event-related potential generation (Makeig et al., 2002; Savers et al., 1974). It is theorised that the effect of the baseline alpha state on the later components was potentially a result of residual MR artefacts, increasing the overall variability of the VEPs. One solution for future research would be to employ functional source separation (Barbati et al., 2008; Barbati et al., 2006; Ji et al., 2019; Porcaro et al., 2010), a semi-blind source technique designed to extract ERPs, which has been demonstrated to improve the overall quality of VEPs (Porcaro et al., 2010) and steady-state VEPs (Ji et al., 2019) recorded within the scanner.

Future Research

Future work should aim to model different frequency bands, as well as using all electrodes, as opposed to solely posterior electrode sites. However, frequency bands outside of the alpha band range often remain corrupted by residual artefacts from the MR environment, with the residual BCG artefact harmonics corrupting the lower frequencies and the residual gradient artefact harmonics corrupting the higher frequency bands. One potential method for including all electrodes and higher frequency bands within the mode would be to perform source localisation

(e.g. beamforming). As well as providing improved spatial information, source localisation acts as a spatial filter, and has been demonstrated to reduce residual gradient artefacts (Brookes et al., 2008), potentially allowing for higher frequency bands (e.g. high beta and low gamma) to be included within the HMM. This could then be modelled using a similar approach as has been applied to MEG data, inputting a series of time series from source localised parcels into an HMM with a time-embedded delay (TDE) observation model (Vidaurre, Abeysuriya, et al., 2018). Alternatively, the study could be replicated using d/EEG, removing the MR-related artefacts, but at the cost of subcortical spatial sensitivity. This would allow the HMM to model changes in power across all frequency bands, and allow for more trials per subject, as the response time would no longer be limited by the sluggish BOLD response, improving statistical power.

Conclusion

This chapter investigated the impact of the active and passive conditions on EEG HMM states and the impact of pre-stimulus alpha power on EEG and fMRI responses. The high alpha state was significantly less likely to precede trials during the active condition, probably reflecting a state of anticipation. Furthermore, the baseline HMM states had a significant impact on the alpha ERD/S responses. The ERD was primarily generated by trials preceded by the high alpha state, while the ERS was primarily generated by trials preceded by the low alpha state. Research has demonstrated that the reduction in alpha power in anticipation of visual stimuli likely reflects a thalamocortical circuit in which the flow of information between the LGN and V1 is moderated (Bastiaansen et al., 2001; Bastiaansen & Brunia, 2001), which is explored in the next chapter. In conclusion, the results of this chapter further support the role of alpha in anticipation and visual attention and demonstrate the importance of accounting for the baseline state during sensory processing. Furthermore, the results show that the HMM approach provides a novel method for sorting individual trials according to alpha power.

Chapter 6

Thalamus BOLD States Modulate BOLD Responses and P250 Amplitudes

Abstract

The thalamus is the largest subcortical structure in humans, and is primarily involved with sensory processing, as well as higher-level cognition, such as attention and state switching. However, it remains understudied within neuroscience research. In this chapter, an exploratory analysis was conducted to investigate the functional dynamics of the thalamus, its relationship with 7-13Hz EEG alpha power, and its impact on visually evoked BOLD and EEG responses. A functional parcellation of the thalamus was obtained via ICA of BOLD data, and the components were then inputted into an HMM to model the state dynamics of the thalamus. The thalamic states showed biologically plausible BOLD correlates, primarily within visual cortex. Additionally, each state showed a unique temporal relationship with posterior EEG alpha power, with several of the states showing an increase or decrease in alpha power approximately five seconds after the thalamic state became active. Furthermore, the thalamic state active during the baseline of each trial significantly impacted the magnitude (% signal change) of the evoked BOLD response peaks and undershoot, as well as the amplitude of the P250 VEP component. Overall, this suggests that the BOLD signal within the thalamus is strongly correlated with the BOLD signal in visual cortex, exhibits a complex temporal relationship with alpha power, and shows a strong association with the modulations observed in visually evoked responses.

Introduction

The thalamus is the largest subcortical structure in the human brain and is widely involved in sensory processing, with visual, auditory, somatosensory, and gustatory signals all passing through their associated thalamic nuclei before being relayed to their respective cortical areas. Due to this, the thalamus was often viewed as a sensory relay station, resulting in the thalamus being widely understudied within neuroscience research (Shine et al., 2023). However, the region has since been associated with more complex functions, including visual attention (Rees, 2009) and the generation of the alpha rhythm (Becker et al., 2015; Suffczynski et al., 2001). These have both been demonstrated to modulate BOLD and EEG responses in the previous two chapters, suggesting that the thalamus may also be involved in orchestrating the observed modulation of visual responses. This has recently been demonstrated in research in cats, with the pulvinar modulating the firing rate of V1 neurons in response to visual stimuli (De Souza et al., 2020). This is facilitated by the dense two-way connections the pulvinar has with both primary visual cortex and extrastriate cortex (Arcaro et al., 2015; Barron et al., 2015; Leh et al., 2008). Furthermore, the inverse relationship has also been observed in mice, with activity in primary visual cortex modulating the firing rate of neurons within the LGN in response to visual stimuli (Reinhold et al., 2023). In addition to the modulation of visual responses, fMRI research in humans has shown that the thalamus, as well as other subcortical regions, such as the putamen, exhibit unique BOLD patterns when different resting-state networks were active (D. J. Greene et al., 2020), suggesting a link with resting-state networks that are generally defined cortically (Hwang et al., 2017).

Given the involvement of the thalamus with resting-state networks (D. J. Greene et al., 2020; Hwang et al., 2017), visual attention (Rees, 2009), and the generation of the posterior alpha rhythm (Becker et al., 2015; Suffczynski et al., 2001), this chapter conducts an exploratory analysis of the BOLD signal dynamics of the thalamus. Firstly, an ROI analysis of the right

LGN was conducted to observe whether responses within the thalamus were modulated by the condition or type of a visual stimulus. A spatial ICA approach was then applied to the BOLD data to obtain a data-driven parcellation of the thalamus (Hale et al., 2015; S. Zhang & Li, 2017), which was subsequently inputted into a Hidden Markov Model (HMM) (Vidaurre et al., 2016; Vidaurre et al., 2017) to model the BOLD dynamics of the thalamus. With the results of the HMM, the study aimed to investigate the cortical correlates of these thalamic states. Additionally, relationships between the BOLD thalamic states and posterior EEG alpha power were investigated to determine whether thalamic states were associated with a significant increase or decrease in alpha power. Lastly, we investigated whether modulations of BOLD thalamic states were associated with significant modulations of BOLD responses, VEPs and alpha ERD/S, as has been observed within the animal literature.

Methods

Dataset

This chapter used the active/passive viewing EEG-fMRI dataset collected during my PhD. See Chapter 2 for details on stimuli and task.

Participants

A total of 30 participants were collected within the active vs passive EEG-fMRI dataset (see Chapter 2 for complete details on participant demographics and ethics). Two participants were excluded from the EEG data in this chapter: one for excessive movement and one for electrode bridging, resulting in a total of 28 participants being included in all subsequent EEG analyses. Similarly, the participant with excessive movement was also excluded from the fMRI data, resulting in 29 participants being included in all subsequent fMRI analyses.

EEG Preprocessing

Preprocessing consisted of removing gradient artefacts, ballistocardiogram artefacts and movement artefacts via carbon wire loop regression. Bad channels were identified and removed from the data before undergoing average referencing, and ICA to remove any remaining artefacts, such as blinks or muscle-related artefacts. Data were then bandpass filtered to the alpha band (7 - 13 Hz) to create an alpha power time course and for the alpha ERD/S analysis (see Chapter 2 for complete details on fMRI preprocessing, VEP, and alpha ERD/S extraction.).

fMRI Preprocessing

fMRI preprocessing consisted of susceptibility correction, motion correction, normalisation to standard space and ICA AROMA (Pruim et al., 2015) via fMRIPrep (Esteban et al., 2017). A custom Nipype (Gorgolewski et al., 2016) script was then employed to remove the noise components identified by ICA AROMA, apply a spatial filter of 4mm using FSL's SUSAN, and apply a high pass filter (100s). See Chapter 2 for complete details on fMRI preprocessing.

Thalamic ICA

To restrict the ICA to the thalamus, the functional images were masked to exclude all regions outside of the thalamus. The thalamus mask was created using the Oxford Thalamic Connectivity Atlas (http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/Atlases) (Behrens et al., 2003; Behrens et al., 2003), in which the maxprob-thr25-2mm atlas was thresholded (>1) and binarised to create a whole thalamus mask using the fsl_maths function from FSL. An ICA was then conducted on the masked ICA data (Hale et al., 2015; S. Zhang & Li, 2017) using the MELODIC (temporal concatenation) function from FSL. A 10 component parcellation was selected based on previous research (Hale et al., 2015) and to minimise the number of free parameters within the HMM. Dual regression was conducted to obtain single-subject time courses for each independent

component (Beckmann et al., 2009; Nickerson et al., 2017).

Thalamic HMM

The dual regression time courses were temporally concatenated and used as inputs for the thalamic HMM analysis (Vidaurre et al., 2016). Inputs were standardised so that the standard deviation was equal to one and the mean was equal to zero. As within Chapters 3 and 4, a multivariate Gaussian distribution model was selected for the observation model, defining states according to the mean activity of each ICA component and the covariance between independent components.

ROI Analysis

An ROI analysis to investigate the impact of condition, state, and baseline HMM states on the shape and amplitude of the BOLD response within the right (contralateral to the stimulus) lateral geniculate nucleus (LGN). The right LGN was chosen as it was the only thalamic region to show significant activation in response to the stimuli (see Figure). Using the positive second-level (group level) passive contrast from Chapter 4 (the right LGN was not reliably observed in first-level contrasts), the thalamus was masked using the mask created for the ICA. The same procedure as used within Chapter 4 was then used to obtain the peak coordinates within the thalamus and create a 3x3x3 voxel cube around the peak voxel (see Figure 6.1). The right LGN roi was then transformed back to each subject's T1 space using the transforms calculated in fMRIPrep via the antsApplyWarp function. Additionally, the right V1 and right Insula ROIS defined in Chapter 4 were also included within the analysis.

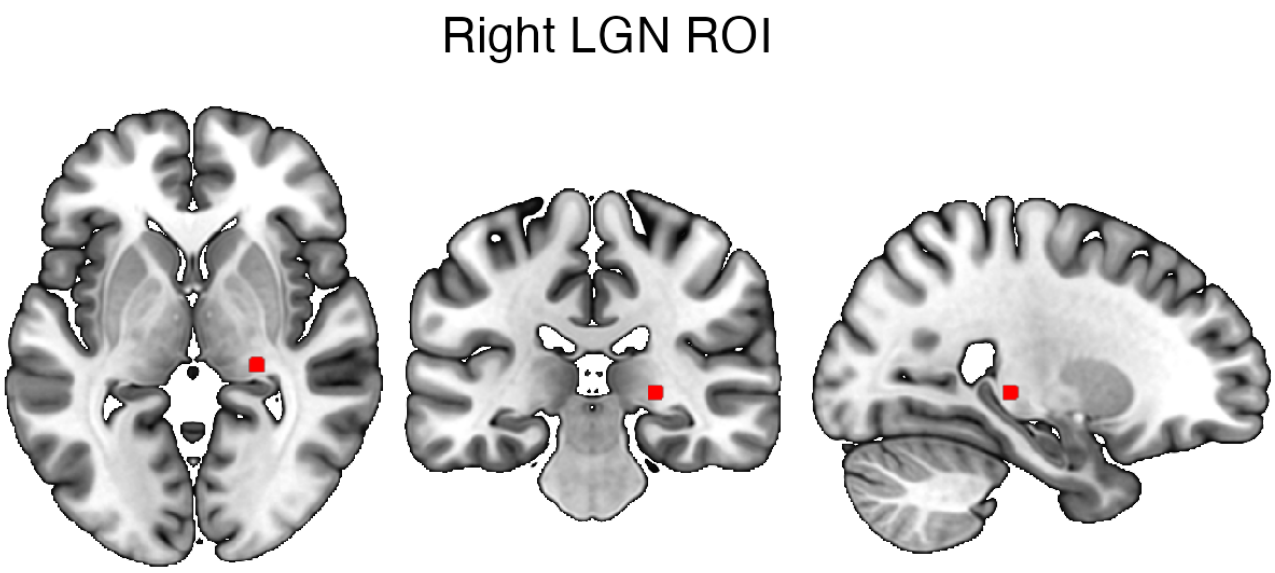


Figure 6.1: Right LGN mask (red). Image presented in MNI152 space using neurological conventions (left = left). ROI MNI coordinates: 24,-27,-1mm.

Extracting BOLD Responses

BOLD responses were extracted from each ROI and converted to percent signal change (see Chapter 2 for details).

Extracting VEPs and Alpha ERD/S

VEP and alpha ERD/S responses were extracted from the right posterior electrode that showed the strongest P100 component and baseline corrected (see Chapter 2 for details).

Cortical Correlates of Thalamic States

To determine the cortical BOLD signal correlates of changes in the activity of each thalamic state, the state probability time courses were extracted and used as explanatory variables within a GLM analysis. The onsets and durations of each state visit were calculated using the getSta-

teOnsets and the getStateLifeTimes (using the Viterbi path) functions within the HMM-MAR toolbox (Vidaurre et al., 2016). The GLM analysis was conducted in Python using the Nilearn package. A first-level analysis was conducted at the subject level (subject T1 space), in which the state visits from each session were entered as explanatory variables. The explanatory variables were not convolved with a haemodynamic response function, allowing for the analysis to identify cortical regions that were active at the same time as the thalamus states (i.e., similar to a functional connectivity analysis). As the regressors were not convolved, temporal derivatives were not included within the model. In addition to the state visits, two regressors of no interest were included within the first-level design matrix: cerebrospinal fluid (CSF) and white matter (WM). These were generated by fMRIPrep during preprocessing, as well as five polynomial drift regressors to high pass filter the model at 0.01 Hz (100s). A separate first-level GLM was performed for each state and each condition, resulting in 12 first-level models (six states x two conditions) per subject. Two contrasts were used within the first-level analysis to calculate state visits' positive and negative mean effects.

A second-level group analysis was then conducted to acquire a group average map for each state and condition. The second level analysis was masked using the MNI152_2mm_brain_mask from FSL. A total of six second-level models were fitted, one for each HMM state, with each including first-level contrast images from both the active and passive conditions. The second level design matrix consisted of two regressors modelling the active and passive condition and a regressor for each subject to model within-subject variability, allowing for paired comparisons. A total of three contrasts were computed: two one-sample t-tests to acquire the mean activation map for both conditions (positive and negative), as well as a paired t-test to contrast the two conditions (active $>$ passive). Contrasts were then thresholded using Bonferroni correction ($p < 0.05$).

Inferring Spectral Properties of Thalamic States

As the alpha rhythm has been previously associated with the thalamus in various computational models (Becker et al., 2015; Suffczynski et al., 2001) and electrophysiological studies (for a review, see Hughes and Crunelli (2005)), the relationship between the occupancy of thalamic states and the power of the posterior alpha rhythm was investigated. The bandpass filtered EEG data were epoched by the TR (1.01s), and alpha power was calculated via Morlet wavelet transform in each epoch, before being averaged across the right parieto-occipital channels (O2, PO4, PO2, P2, P4, P6, P8). This resulted in an alpha power time course at the same sampling rate as the fMRI data. The Viterbi path of the thalamus states was used to determine which state was active for each fMRI volume, and the alpha power time course was split accordingly, allowing for the average alpha power to be calculated for each state. As the temporal relationship between the thalamus states and posterior alpha power was not known, both a lead (e.g. alpha power precedes thalamus state) and lag (e.g. thalamus state precedes alpha power) were introduced on the alpha power time course, allowing for the mean alpha power to be calculated across different time periods. The time period showing the most change in dynamics across all states was then selected for statistical comparison (see Chapter 2 for details on statistical testing).

Evoked Fractional Occupancy

The evoked fractional occupancy responses were calculated to evaluate which states became more or less active in response to the visual stimuli. The responses were calculated by epoching the state likelihood time series around the stimulus onset and baseline corrected using the same window and baseline correction method used for the BOLD response analysis. Average responses were calculated for each subject, which were then averaged to generate a grand average.

Statistics

Statistical testing was conducted to compare the effect of condition on fractional occupancy and switching rate via paired t-tests (see Chapter 2 for details). Additionally, the impact of baseline HMM state, condition, and stimulus type on BOLD, VEP, and alpha ERD/S responses was evaluated via linear mixed modelling (see Chapter 2 for details). Lastly, a repeated measures ANOVA was conducted to compare alpha power values between thalamus states.

Results

Thalamic ICA

An ICA was run on all BOLD signal voxels within the thalamus to obtain a functional parcellation using FSL's MELODIC function (Beckmann & Smith, 2004) via Nipype (Gorgolewski et al., 2011). All identified components showed plausible spatial properties comparable with previous studies (Hale et al., 2015; S. Zhang & Li, 2017), including bilateral and unilateral activations. The THOMAS atlas (Su et al., 2019) was used to assist in the labelling of specific thalamic nuclei present within each component. Components were primarily located within the medial, anterior or posterior regions of the thalamus (see Figure 6.2). ICs one, two, three and five were all located within the medial thalamus, with ICs 1 and 5 overlapping with the midline of the thalamus. These ICs overlapped with medial nuclei, including the medial dorsal nuclei and segments of the ventral lateral posterior nuclei. ICs 6 and 8 were both located within the anterior of the thalamus bilaterally, overlapping with anterior nuclei, such as the ventral anterior nucleus. ICs 4, 7, 9 and 10 were all located within the posterior of the thalamus, with components exhibiting a mix of unilateral and bilateral activity. These components primarily overlapped with the pulvinar, with the exception of IC 9, which overlapped with the LGN bilaterally.

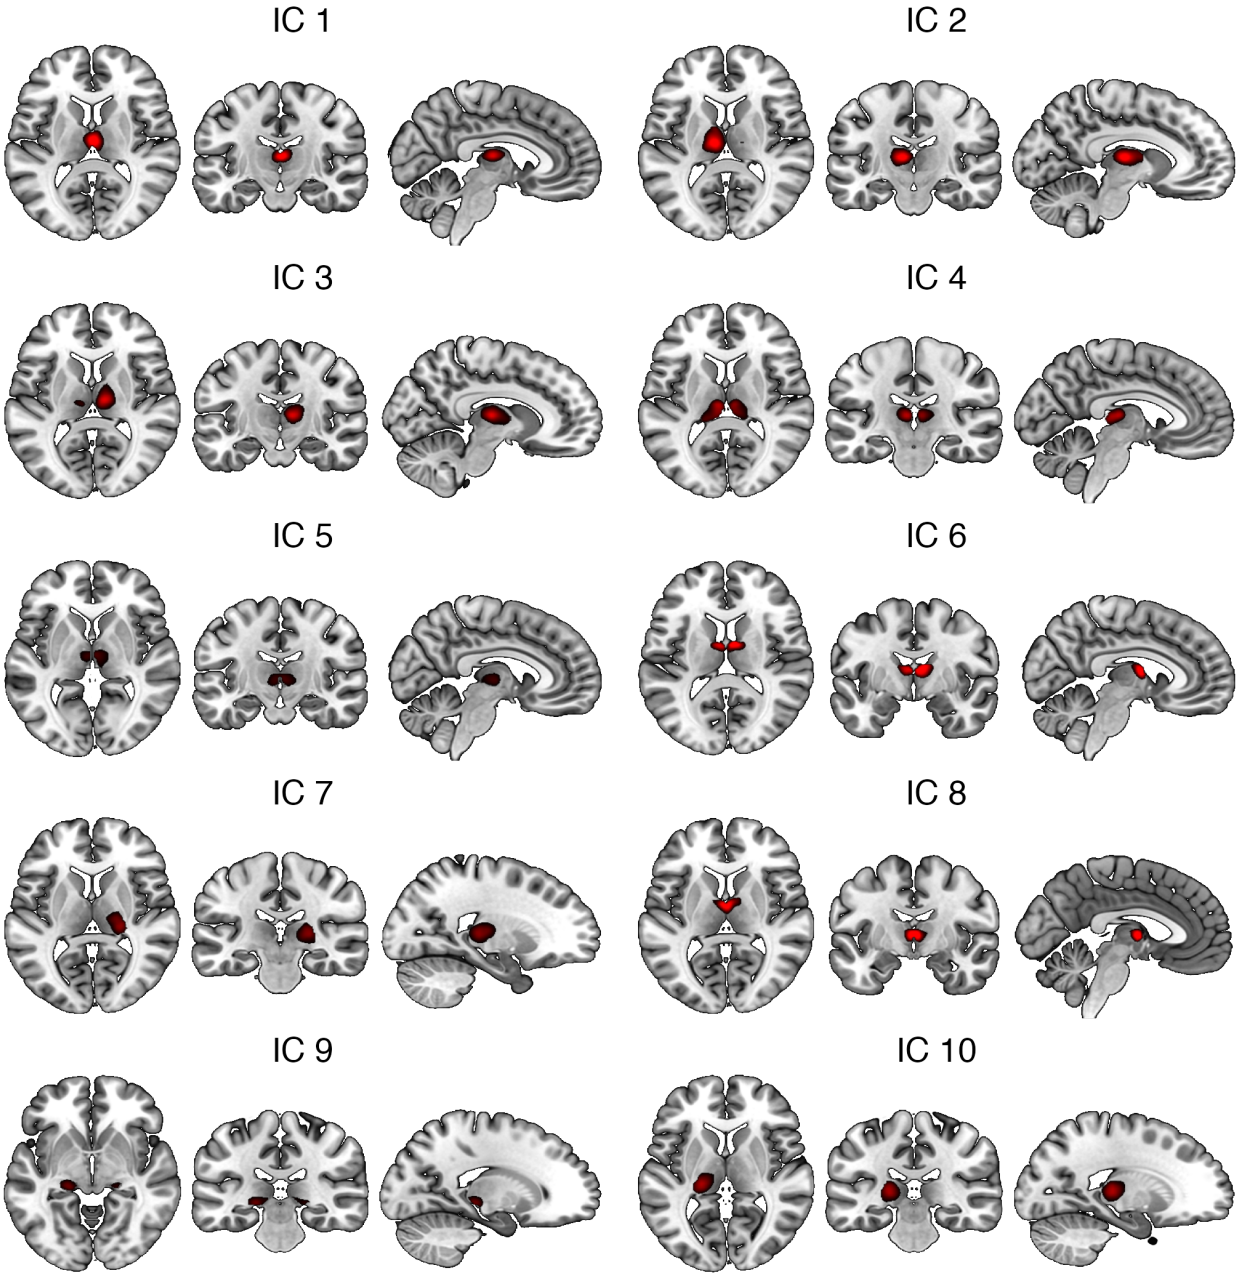


Figure 6.2: Spatial maps of Thalamic independent components (red). Images are presented in MNI152 space using neurological conventions (left = left).

ROI Analysis

The impact of condition and stimulus type on the shape of evoked BOLD responses within the right LGN (see Figure 6.3) were investigated using an ROI analysis. BOLD responses were

reliably generated in the right LGN in both conditions. The BOLD responses peaked at 5 TR (5.05 s), with the active response peaking at 0.26 % signal change, and the passive response peaking at 0.18 % signal change. Additionally, the BOLD response displayed an undershoot at 10 TR (10.1 s) before returning to baseline at 17 TR (17.17 s).

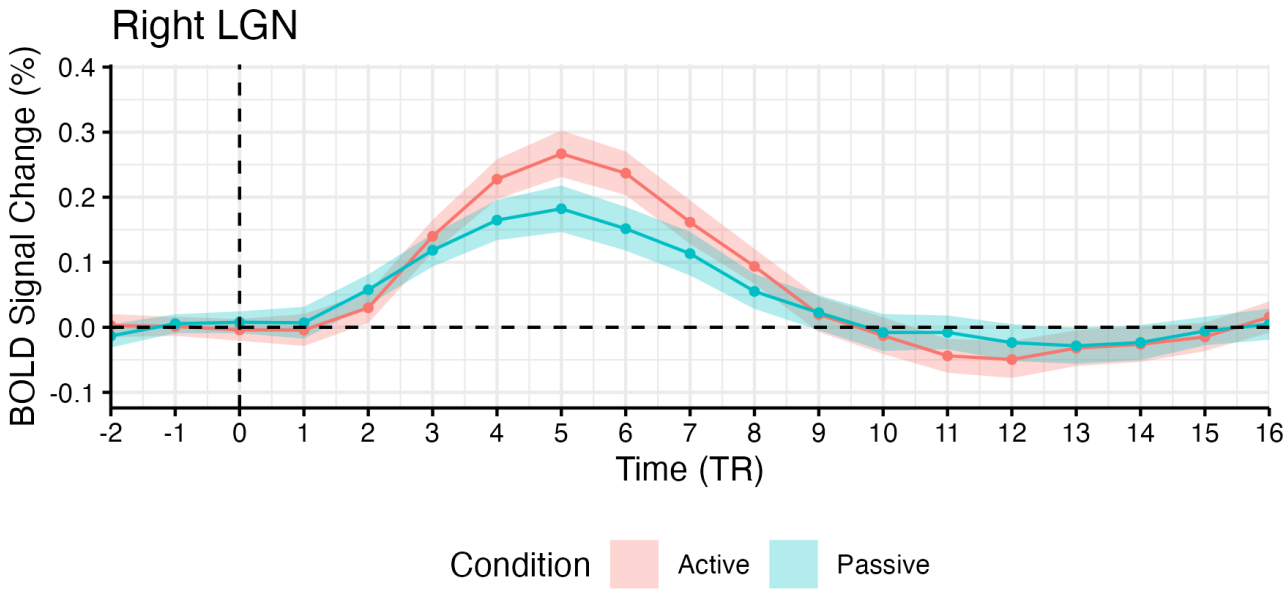


Figure 6.3: Comparison responses within the right LGN across the two conditions. Error represents 95% confidence intervals.

Comparing the evoked BOLD responses to each stimulus type, the standard and target stimuli reliably evoked a BOLD response in the right LGN (see Figure 6.4). Evoked BOLD responses from both stimulus types peaked at 5 TR (5.05 s) irrespective of condition. When grouping trials across both conditions, there was no difference in the peak response amplitudes between target and standard stimuli, with both peaking at 0.23% signal change. When restricted to trials from the active condition, target stimuli peaked at 0.27% signal change, whilst standard stimuli peaked at 0.26% signal change. Lastly, when restricted to trials from the passive condition, target stimuli and standard stimuli peaked at approximately 0.18 % signal change.

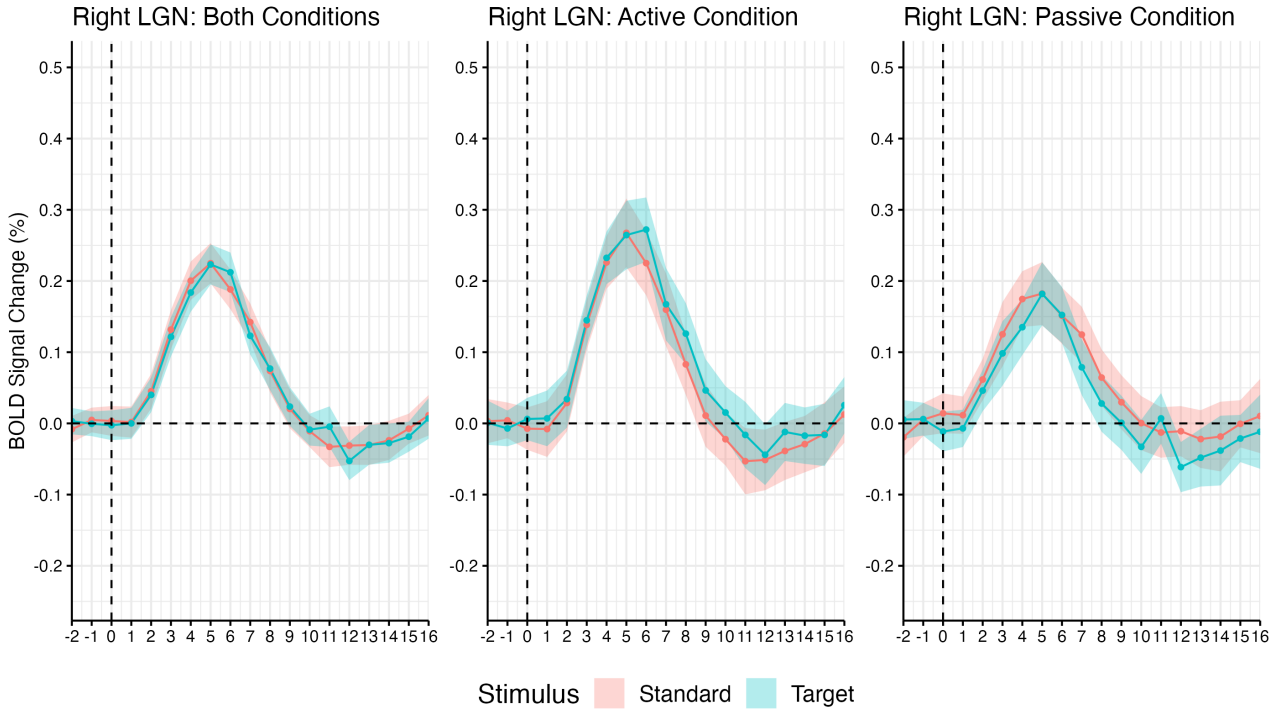


Figure 6.4: Comparison of the mean BOLD responses within the right LGN for each standard and target stimuli. The left plot shows the mean BOLD responses when trials are grouped across both conditions, whilst the centre and right plots show the mean BOLD responses only, including trials from the active and passive conditions, respectively.

Model Selection

A total of 10 model repeats were run per state number, with the state number ranging from 2-15 (see Figure 6.5). The free energy progressively decreased with the addition of each state, making it difficult to use this metric alone to determine the optimal state number. The mean repeat reliability between the 10 runs showed a drop off at six states, with a slight increase in reliability occurring within the nine-state solution. Based on this, a six-state solution was chosen to avoid splitting the trials between nine states in the subsequent response analyses. The repeat reliability matrix was calculated for a six-state solution, which showed that all repeats showed a strong correlation to each other (all correlations above 0.9). The maximum fractional occupancy values were skewed towards 0, with a mean of approximately 0.3, demonstrating that

a single state was not explaining all the variability, and, thus that the state was most likely evenly distributed across the time course.

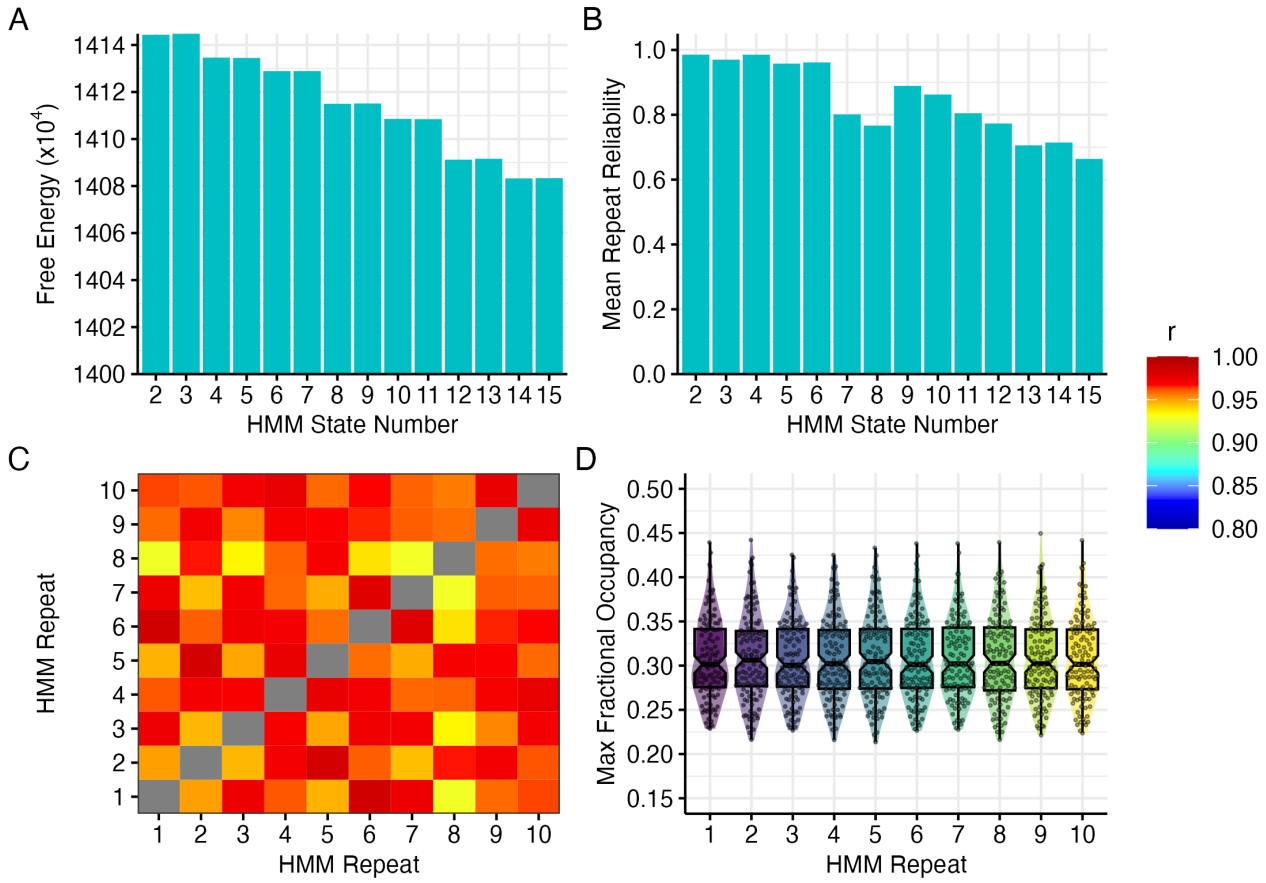


Figure 6.5: Model selection plots for the thalamic HMM. (A) the mean free energy value across runs for each state solution. (B) the mean repeat reliability across runs for each state solution. (C) the repeat reliability matrix across runs for the six-state solution. (D) the distribution of maximum fractional occupancy values for each repeat of the six-state solution. Each point represents the max fractional occupancy from a single run.

HMM States

A total of six thalamic HMM states were inferred (T1-T6) based on the mean activity and covariance of the inputted dual regression BOLD signal time courses (see Figure 6.6). All HMM states showed widespread activations and deactivations, between them spanning the entirety of the thalamus. Specifically, state T1 was defined by positive activations across all IC

components, with the largest responses being within ICs one, two and three. Similarly, state T5 showed the inverse pattern to state T1, with widespread negative mean activity across all components (except IC9), with the largest negative activity also within ICs one, two and three. State T2 exhibited a similar pattern of activity as state T5, with deactivations across all ICs; however, the deactivations were notably smaller relative to state T5 and strongest in ICs three, seven and ten. The remaining states all exhibited lateralised activity, with activations in one hemisphere and deactivations within the other. State T3 showed activations within IC two and IC ten (left lateralised), with deactivations in ICs one, five, six (medial bilateral) and seven (right lateralised). State T4 showed activations across the majority of ICs, with deactivations within ICs seven and nine. Lastly, State T6 was characterised by a large activation in IC seven, alongside deactivations within ICs one, two, four and six. Regarding functional connectivity (FC), states T1, T2, T3 and T4 all showed similar connectivity patterns, with no ICs exhibiting strong correlations to each other. Conversely, states T3 and T5 both showed an increase in FC relative to the other states, with notable intra-correlations between ICs one and six.

Evoked HMM Thalamic Response

To determine whether any of the thalamic states were evoked in response to the visual stimuli, the state time courses were epoched around the stimulus (-2 TR - 17 TR) and then averaged across subjects to generate an average evoked fractional occupancy plot (see Figure 6.7). State T2 showed an increase in fractional occupancy of approximately 7% in response to the visual stimuli. Both states T1, and T4 showed a decrease in fractional occupancy in response to stimuli, with T1 reducing by approximately 3%, and T4 reducing by approximately 4%. Additionally, the peak in fractional occupancy occurred at the fifth post-stimulus TR, irrespective of the direction of change, matching the peak of the BOLD response observed in the LGN, and within V1 and insula within Chapter 4. States T5 and T6 showed minimal change in fractional occupancy in response to stimuli.

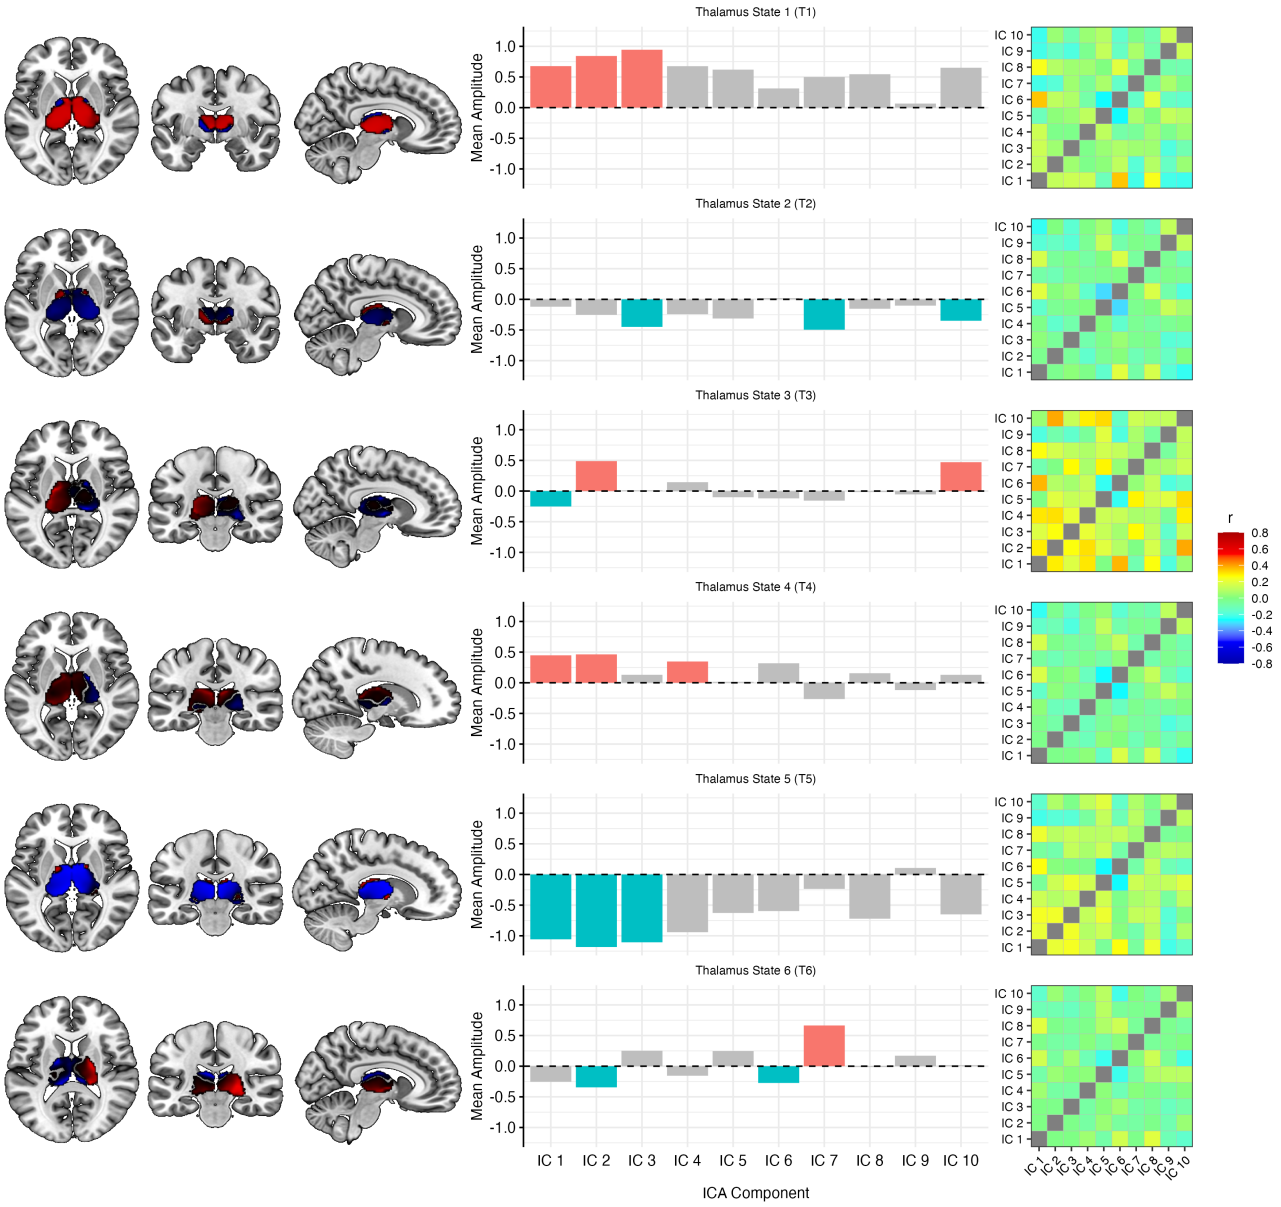


Figure 6.6: States identified by the thalamus HMM. Each row contains the plots for a single state. The left column shows the mean activation within each ICA component scaled across states. The central column shows the mean activation of each IC component for each state without scaling. Coloured bars represent the three components with the largest amplitudes (red indicates positive amplitude. blue indicates negative amplitude). The right column shows the functional connectivity between IC components for each state.

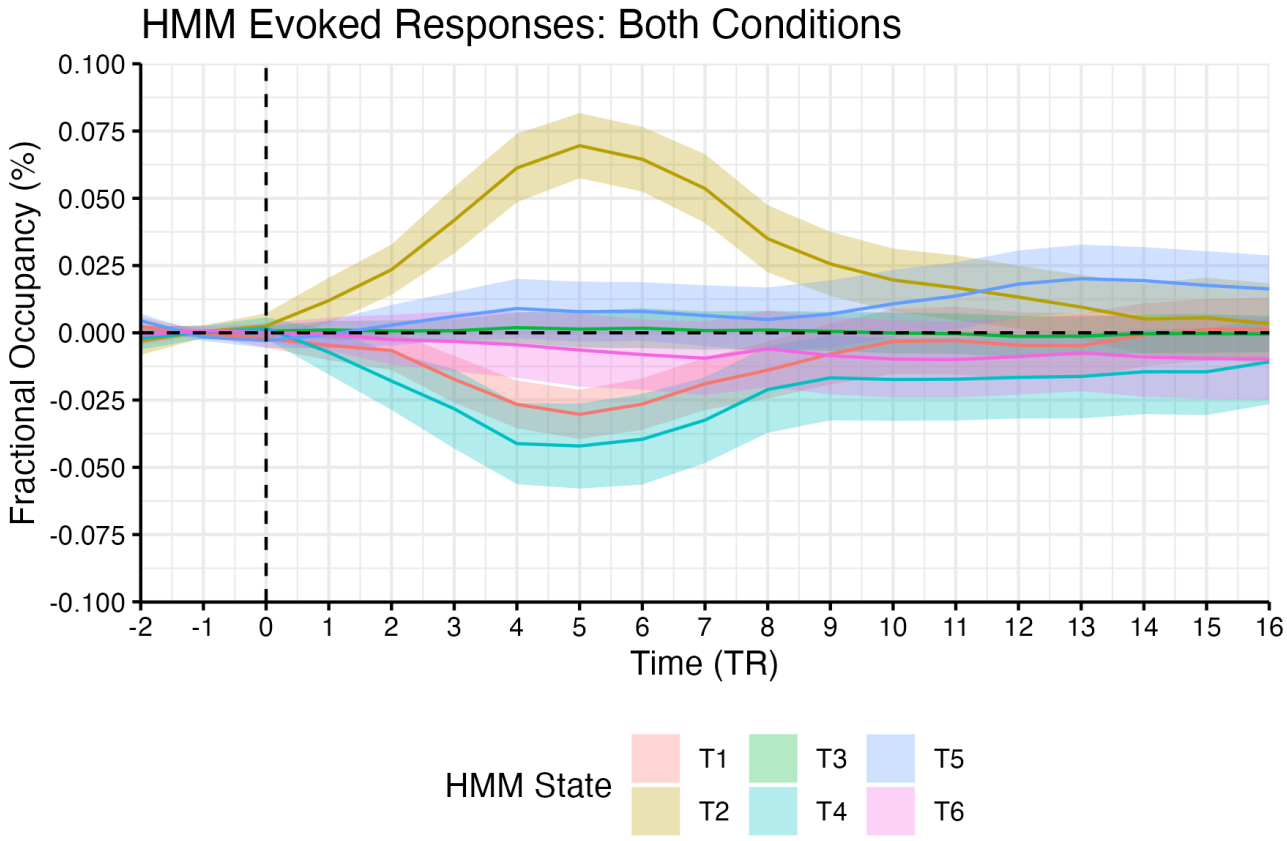


Figure 6.7: Mean changes in state fractional occupancy following stimulus presentation. Error represents 95% confidence intervals.

Distribution and Reliability of Baseline States

To assess the likelihood of a specific state being active before the presentation of each stimulus, the number of occurrences where each state was active at baseline was counted for each subject, compared across conditions and plotted (see Figure 6.8). States T4 and T6 were most likely to be active prior to stimulus presentation. Additionally, states T1, T3, and T5 were active prior to stimulus presentation more frequently during the active condition than the passive condition.

To ensure the reliability of using timepoint t-1 to determine which state was active prior to stimulus presentation, the active baseline states were also determined for timepoints t0 to t-5, and correlated (Pearson's). The resultant correlation matrix showed that the active baseline

states showed strong similarity when determined at t0, t-1 and t-2, demonstrating that the choice of baseline timepoint is reliable within an approximately three-second

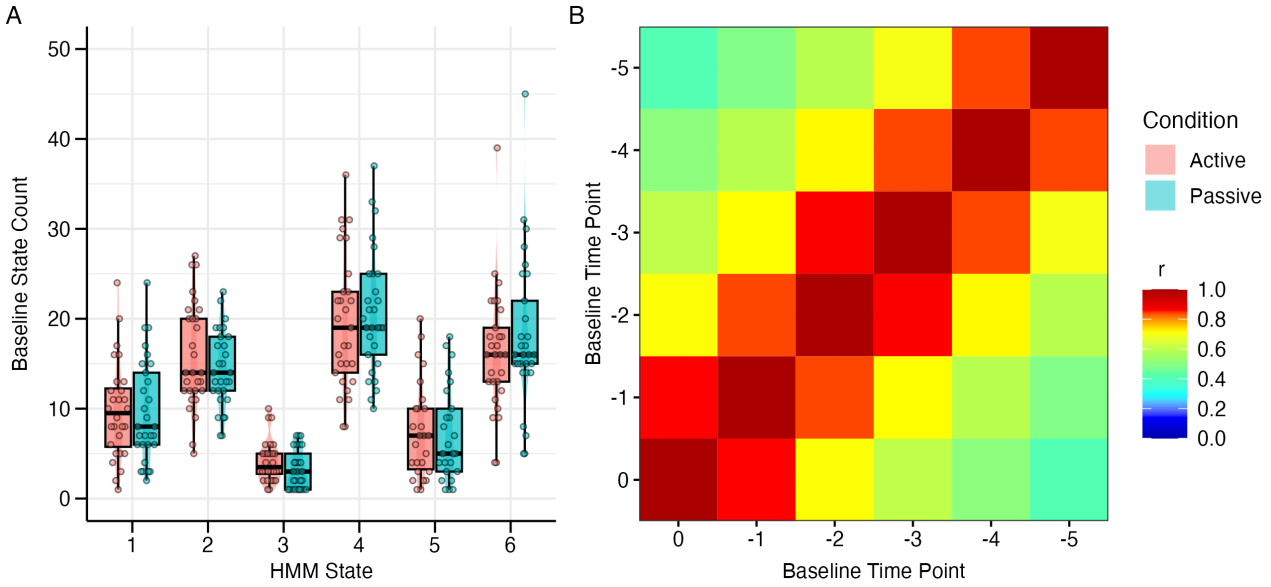


Figure 6.8: Plot A shows the number of times each state was active during the baseline (the TR preceding stimulus presentation) for each condition. Each point represents the number of times a state was active for that condition (combined across runs) and state. Plot B shows the reliability matrix of baseline state timepoints. The matrix is created by calculating the baseline states at each timepoint and correlating the order of baseline states.

State Metric Comparison

State metrics were compared across the two conditions to determine whether thalamic states were impacted by attention. Observing state fractional occupy (see Figure 6.9), states T2, T4, and T6 showed the highest fractional occupancy, averaging above 20%. States T1, T3, and T5 showed lower fractional occupancy, averaging at 10% or less. No states showed any difference in fractional occupancy across conditions. Similarly, no difference was observed in the switching rate condition. Regarding state lifetimes, all states showed similar lifetimes averaging around 8 TR (8.08 s), with the exception of state T6, which averaged at approximately 12 TR (12.12 s). State interval times showed that T3 had a significantly larger average interval within both conditions, averaging approximately 150 TR (151.5 s). All other states showed similar state

interval times, averaging approximately 40 TR (40.4 s). As with fractional occupancy and state switching rate, no obvious differences were observed between the average state lifetimes and state interval times across conditions.

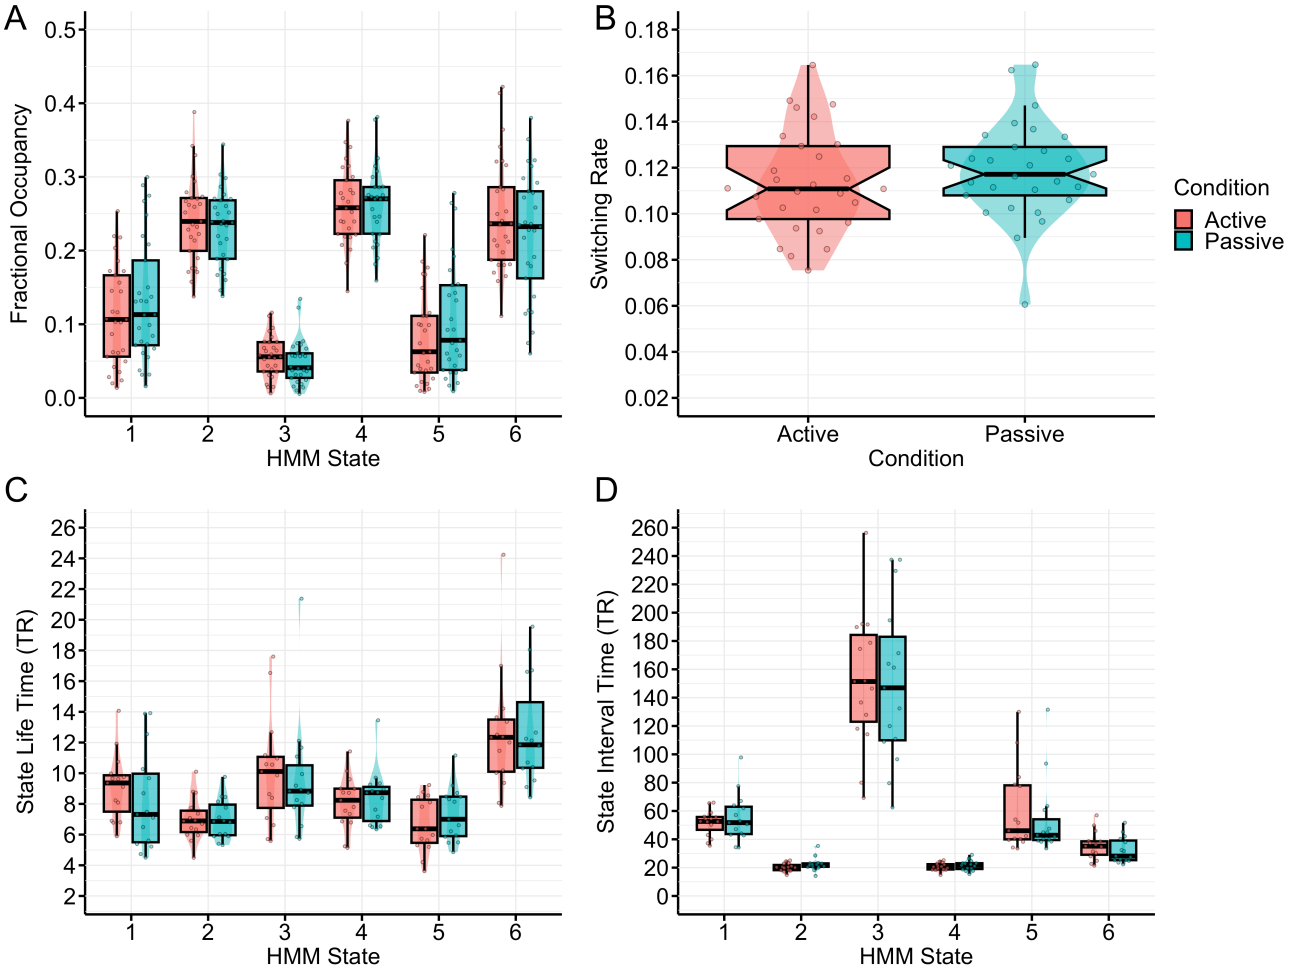


Figure 6.9: Comparison of HMM states metrics between the active and passive conditions. Plot A compares average fractional occupancy. Plot B compares the average state switching rate. Plot C compares the average state lifetimes. Plot D compares the average state interval times. Each point represents a single subject.

To test the differences between conditions statistically, paired t-tests were conducted to compare mean fractional occupancy and switching rate across conditions. Initial results found a significant difference in fractional occupancy between conditions for state five (difference = -0.02, 95% CI [-0.04, -2.79e-03], t(28) = -2.33, p = 0.027; Cohen's d = -0.43, Cohen's d 95%

CI [-0.81, -0.05]); however, the difference was non-significant after Bonferroni correction ($p = 0.162$). All other fractional occupancy tests were non-significant, as was the switching rate test.

Thalamic State Correlates

GLM analysis was conducted to investigate whether changes in the activity of the thalamic states correlated with the BOLD signal across the rest of the brain. Significant correlations (Bonferroni corrected, $p<0.05$) were found across the cortex for each state (see Figure), except for state T3. During the active condition, states T1 and T4 both showed significant deactivations bilaterally within visual cortex, whilst states T2 and T5 showed significant activations within the same regions. Additionally, the activations and deactivations for states T2 and T4 were primarily restricted to visual cortex, whilst states T1 and T5 also showed significant activations/deactivations within the parietal cortex and the intraparietal sulcus. The spatial similarity of their correlations in opposite contrasts is consistent with the inverted representation of T1 and T5 shown in Figure 6.6. State T6 showed unique correlates compared to the other states, with small deactivations solely within the partial cortex. BOLD correlates were observed in thalamic regions for all states except T3. Except for state T6, all states showed the opposite activity within the thalamus to that observed within the visual cortex (e.g. deactivation in visual cortex, activation in the thalamus or vice versa). During the passive condition, the maps closely resemble the activity seen in the active condition; however, activations were typically larger and more widespread across visual cortex, and state T2 specifically showed a stronger lateralisation to the right side of visual cortex. Contrasts between the two conditions were non-significant.

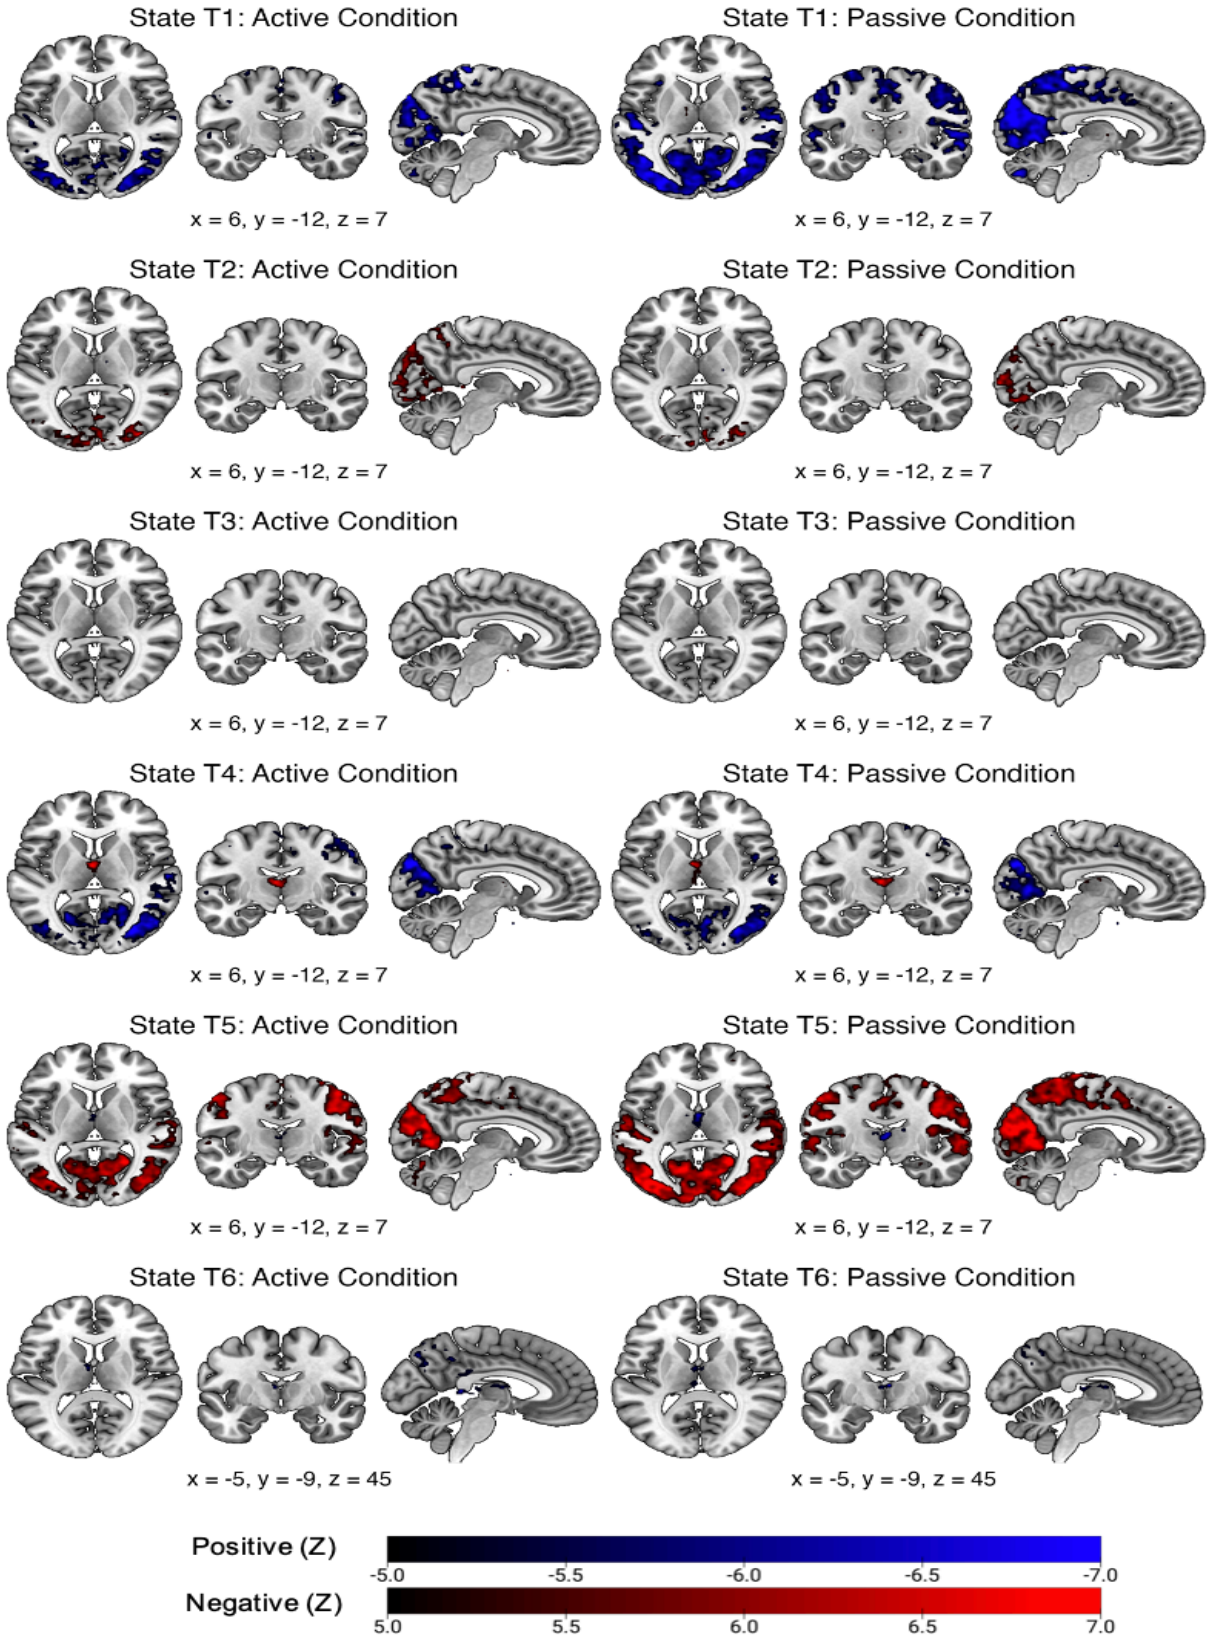

State T1: Active Condition
State T1: Passive Condition
x = 6, y = -12, z = 7
x = 6, y = -12, z = 7
State T2: Active Condition
State T2: Passive Condition
x = 6, y = -12, z = 7
x = 6, y = -12, z = 7
State T3: Active Condition
State T3: Passive Condition
x = 6, y = -12, z = 7
x = 6, y = -12, z = 7
State T4: Active Condition
State T4: Passive Condition
x = 6, y = -12, z = 7
x = 6, y = -12, z = 7
State T5: Active Condition
State T5: Passive Condition
x = 6, y = -12, z = 7
x = 6, y = -12, z = 7
State T6: Active Condition
State T6: Passive Condition
x = -5, y = -9, z = 45
x = -5, y = -9, z = 45
Positive (Z)
-5.0
-5.5
-6.0
-6.5
-7.0
Negative (Z)
5.0
5.5
6.0
6.5
7.0

Figure 6.10: Activation maps for each state across the brain (left = left). Each row contains the activation maps for a single state, with the left column showing the activation maps during the active condition and the right column showing the activation maps during the passive condition. All maps are Bonferroni corrected ($p < 0.05$). Images are presented in MNI152 space using neurological conventions (left = left). Coordinates are in MNI152 space.

Thalamic States and Alpha Power

As BOLD activity within the thalamus has been previously linked with posterior EEG alpha power, we tested whether alpha power differed between states(see Figure 6.11). We found that each state had a unique temporal relationship with alpha power. States T1 and T5 showed a clear change in alpha power with the introduction of a lag, with alpha power increasing with lag for state T1, and decreasing with lag for state T5, both peaking at a lag of 7 TR (7.07 s) and 5 TR (5.05 s) respectively. This is again consistent with the inverse relationship of states T1 and T5 in Figure 2. A similar pattern was observed for state T4, in which alpha power also increased with lag, peaking at 4 TR (4.04s), but to a lesser extent than state T1. States T2 and T6 were also associated with changes in alpha power with the introduction of a lag, with alpha power gradually decreasing with increasing lag for state T2, and increasing for state T6, though much less profoundly than observed for states T1 and T5. State T3 showed a unique temporal relationship with alpha power relative to the other states, with alpha power being highest at the same time as the state was active (lag = 0), and reduced rapidly the introduction of both a lead and lag. Furthermore, the peak alpha power was notably higher for state T3, at more than double the peak power of the other states (note that the y-axis for T3 is larger than for the other states). Generally, introducing a lead to the alpha power had less impact, except for states T2 and T5, in which alpha power increased.

The most notable changes in alpha power occurred approximately 5 TR after the activation of

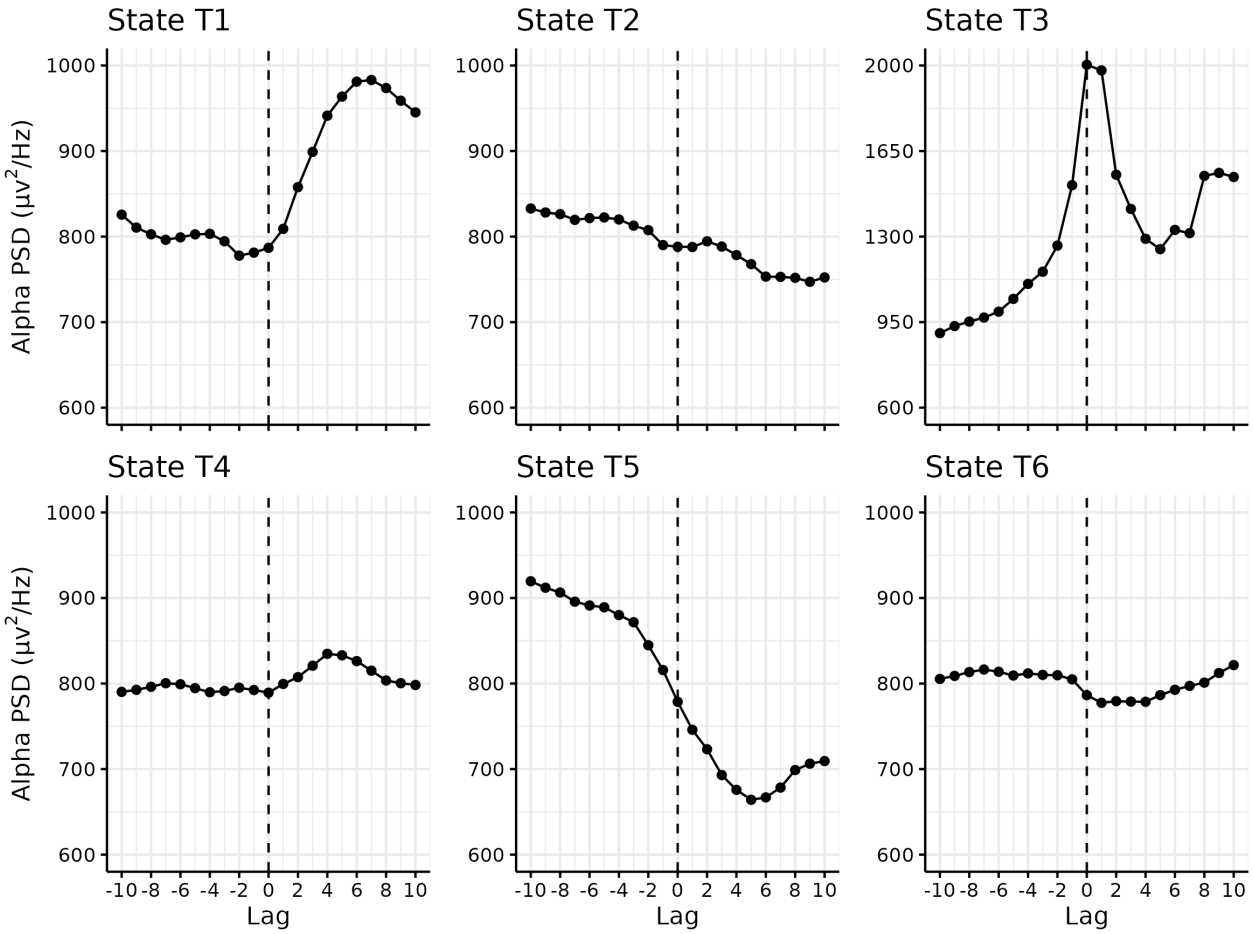


Figure 6.11: Mean alpha power for each state with varying lags. Negative lag values represent a lead, e.g.alpha power changes before the activation of the thalamic state, whilst positive lag values represent a lag (e.g. alpha power changes after the activation of the thalamic state).

the thalamic state. Thus, this time was selected for statistical comparison (see Figure 6.12). A repeated measures ANOVA was conducted to assess whether the thalamic state was significantly associated with alpha power. The alpha power distributions had a strong right skew, so a natural log transform was applied to the data. This demonstrated that the main effect of thalamic state was statistically significant ($F(5, 135) = 24.93$, $p < 0.001$; Eta2 (partial) = 0.48, 95% CI [0.37, 1.00]). Post hoc tests were conducted using paired t-tests with Bonferroni correction to investigate the main effects (see Figure 12). State T3, followed by T1, had the highest lagged alpha power values and were significantly higher than the majority of the other states ($p < 0.05$).

Similarly, state T5 had the lowest lagged alpha power values and was significantly lower than the majority of the other states ($p < 0.05$).

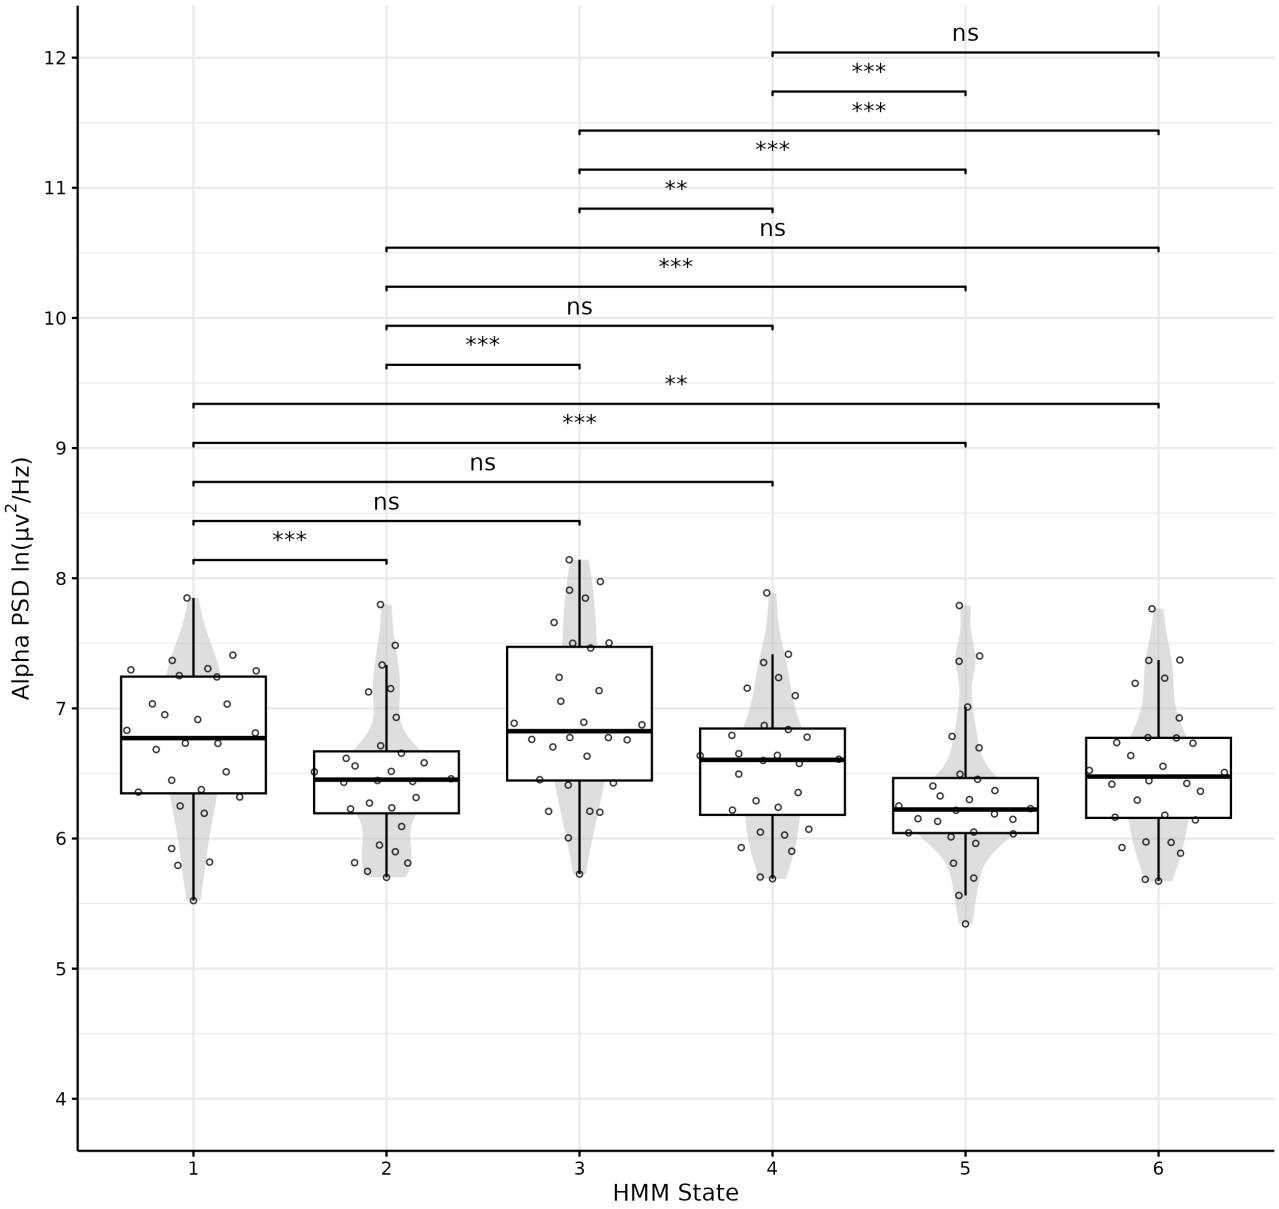


Figure 6.12: Average lagged (5 TR) alpha power values for each HMM state. Post hoc tests represent paired t-tests corrected for multiple comparisons via Bonferroni correction. *** = p < 0.001, ** = p < 0.01, * = p < 0.05, ns = non-significant ($p > 0.05$).

Impact of Thalamus State on BOLD Responses

Next, the impact of the baseline thalamic state on the shape and amplitude of the evoked BOLD responses from right V1, right insula, and the right LGN was evaluated (see Figure 6.13). Qualitatively similar results were observed across both active and passive conditions. Within right V1, the baseline state had a notable impact on the shape of the evoked BOLD response, with the amplitude of the peak and undershoot both being affected. Most notably, the average BOLD peak and undershoot amplitudes were lower when state T5 was active during the baseline. The inverse pattern was observed for state T1, which showed a larger peak and undershoot amplitude relative to the other states. The baseline state had a similar impact on the BOLD response shapes in the right LGN and insula, following a similar pattern to the results in right V1. However, the variability within these regions was larger, making the interpretation within these regions more difficult.

The impact of condition, stimulus type, and thalamus state on BOLD peak and undershoot PSC values was assessed statistically via linear mixed modelling (see Chapter 2). The results of the right V1 models found the thalamus state main effect to be significant for the BOLD peak ($F(5, 4133.2) = 48.397, p < 0.001$) and BOLD undershoot ($F(5, 3922.9) = 71.821.p < 0.001$) PSC values. Additionally, the condition main effect was also significant for the peak model ($F(1, 4124.9) = 26.35, p < 0.001$), with the stimuli generating significantly larger BOLD PSC values during the active condition ($EMM = 1.78\%, SE = 0.114\%$) relative to the passive condition ($EMM = 1.45\%, SE = 0.115\%$).

Significant interaction effects were also found within the right V1 peak model, with condition and stimulus type showing a significant interaction effect ($F(1, 4127.1) = 10.27, p < 0.01$), as well as condition and baseline thalamus state ($F(5, 4124.6) = 2.80, p < 0.05$). For the first interaction effect, post hoc tests revealed (Bonferroni corrected, number of comparisons =6) similar results to the previous chapters, with target stimuli ($EMM = 1.91\%, SE = 0.129\%$) generating signif-

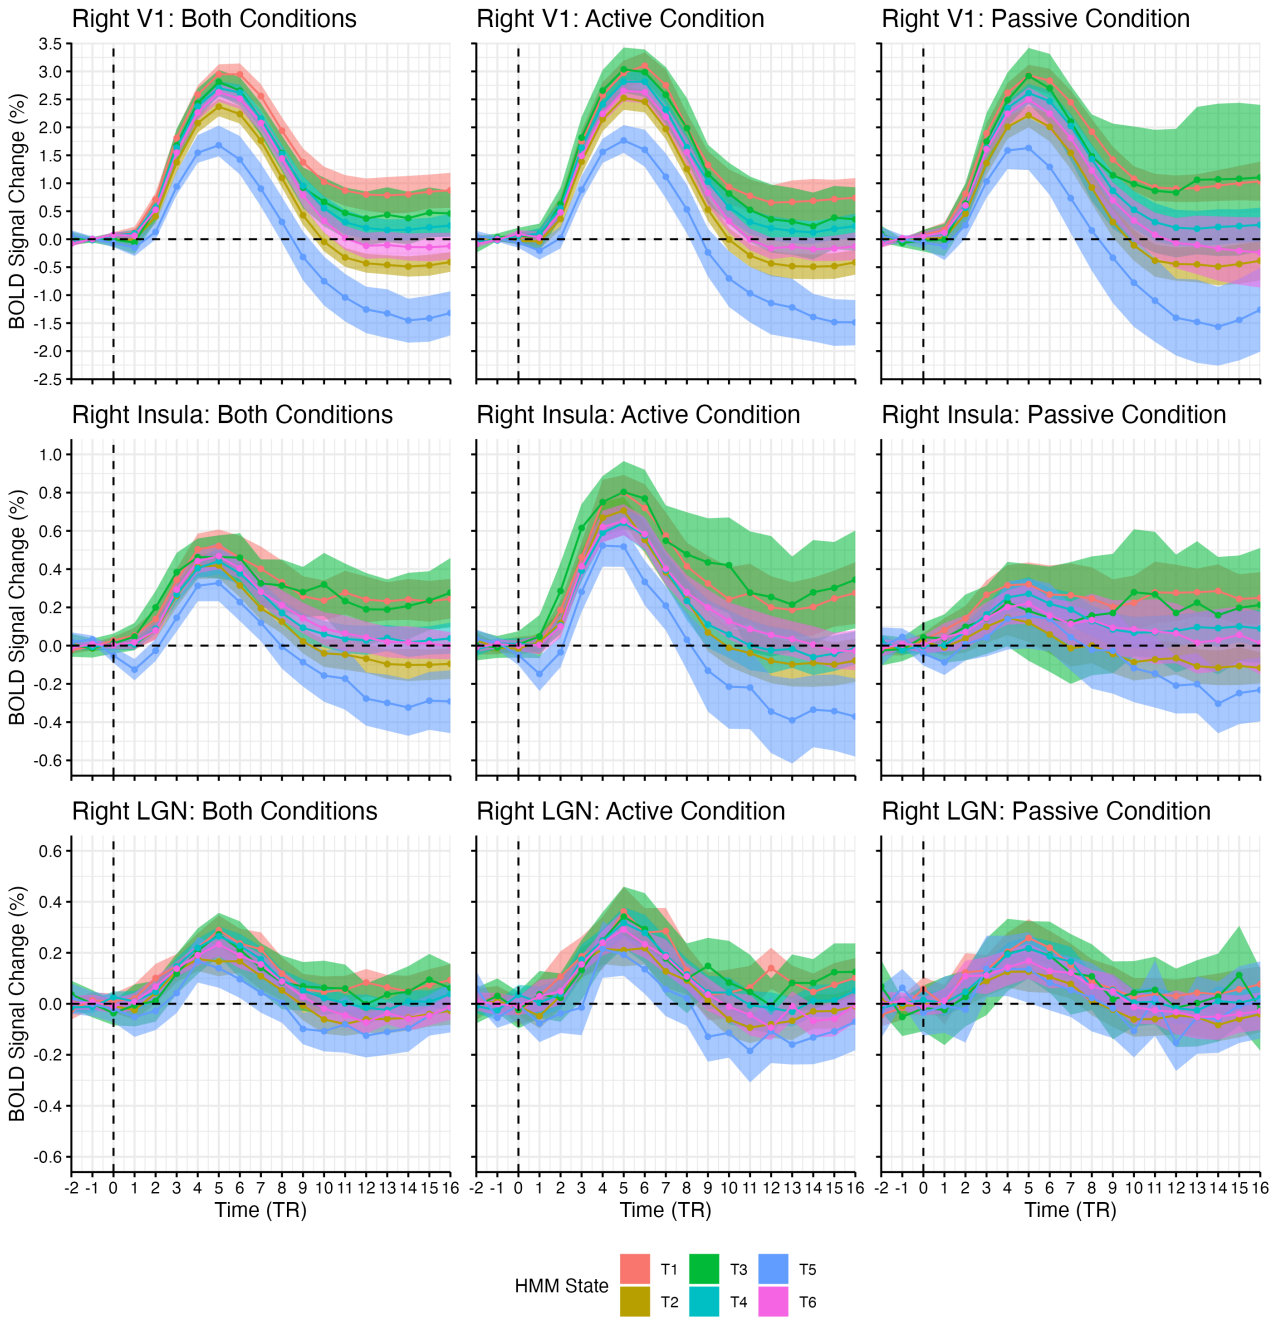


Figure 6.13: Comparison of mean BOLD responses across states. Each row represents an ROI (right V1, right Insula, and right LGN). Each column represents the grouping of trials, with the left column containing trials from both conditions, and the centre and right columns only containing trials from the active and passive conditions, respectively. Error represents 95% confidence intervals.

icant larger BOLD PSC values compared to standard stimuli ($EMM = 1.66\%, SE = 0.114\%$) during the active condition ($t(4125) = -2.863, p < 0.05$), but not during the passive condition ($t(4125) = 1.721, p = 0.512$). For the second interaction effect, post hoc tests (Bonferroni corrected, number of comparisons = 66) revealed that BOLD PSC values were significantly lower for trials preceded by state T5 during the active condition ($EMM = 0.984\%, SE = 0.158\%$) relative to the passive condition ($EMM = 0.185\%, SE = 0.161\%$) ($t(4128) = 4.733, p < 0.001$). Additionally, a significant interaction effect between condition and stimulus type was also observed within the undershoot model ($F(1, 4130.4) = 10.33, p < 0.01$). Post hoc tests (Bonferroni correction, number of comparisons = 6) revealed that target stimuli ($EMM = 0.29\%, SE = 0.114\%$) generated significantly smaller BOLD PSC values than standard stimuli ($EMM = 0.048\%, SE = 0.064\%$) during the passive condition (($t(4134) = 2.677, p < 0.05$)), but not during the active condition ($t(4131) = -1.835, p = 0.399$).

The results from the right LGN models closely resembled the results in right V1, with the baseline thalamus state main effect being significant within both the peak model ($F(5, 3967.0) = 8.11, p < 0.001$) and undershoot model ($F(5, 3965.2) = 11.35, p < 0.001$). Additionally, the condition main effect was significant within the peak model ($F(1) = 7.33, p < 0.01$), with stimuli generating significantly larger PSC values during the active condition ($EMM = 0.148\%, SE = 0.013\%$) relative to the passive condition ($EMM = 0.102\%, SE = 0.014\%$).). Unlike the right V1 models, no interaction effects were significant.

Regarding the right insula, the main effect of thalamus state was significant for both the peak ($F(5, 4136.8) = 12.376, p < 0.001$) and undershoot model ($F(5, 4152) = 20.588, p < 0.001$). The condition main effect was also found to be significant within the peak model ($F(1, 4127.0) = 190.314, p < 0.001$), with stimuli generating larger BOLD PSC values during the active condition ($EMM = 0.446\%, SE = 0.026\%$) relative to the passive condition ($EMM = 0.115\%, SE = 0.027\%$). Additionally, the stimulus type main effect was

also significant within the peak model ($F(1, 4127.3) = 6.255, p < 0.05$), with target stimuli ($EMM = 0.311\%, SE = 0.029\%$) generating significantly larger BOLD PSC values relative to standard stimuli ($EMM = 0.250\%, SE = 0.023\%$). Significant interaction effects were also observed for the peak model, with condition and stimulus type showing a significant interaction effect ($F(1, 4125.9) = 4.031, p < 0.05$), as well as a significant interaction effect between condition and thalamus state ($F(5, 4133.8) = 2.508, p < 0.05$). For the first interaction effect, post hoc tests (Bonferroni corrected, number of comparisons = 6) revealed that target stimuli ($EMM = 0.500\%, SE = 0.034\%$) generated larger BOLD PSC values relative to standard stimuli ($EMM = 0.121\%, SE = 0.037\%$) during the active condition ($t(4126) = -3.336, p < 0.01$), but not during the passive condition ($t(4127) = -0.335, p > 1.00$). Lastly, the post hocs (Bonferroni corrected, number of comparisons = 66) for the second interaction effect revealed that trials preceded by state T3 increased the BOLD PSC values within the active condition but reduced them within the passive condition relative to other states. This effect was most prominent between states T5 and T3, with trials precded by state T3 ($EMM = 0.593\%, SE = 0.062\%$) producing significantly larger BOLD PSC values relative to trials preceded by state T5 ($EMM = 0.218\%, SE = 0.049\%$) in the active condition ($t(4144) = 5.122, p < 0.001$), but not within the passive condition ($t(4150) = 0.509, p = 1.0000$).

As the baseline thalamus state main effect was significant for all three ROIs, impacting both the BOLD peak and undershoot amplitudes, post hoc tests were conducted to investigate the contributions of specific thalamic states (see Figure 6.14 for statistics) (see Chapter 2) for details on post hoc testing LMMs. This showed that the influence of baseline thalamus state remained similar across ROIs. Specifically, the right V1 and right Insula models showed that state T1 typically elicited a larger change in BOLD signal relative to the majority of the other states. Additionally, across all ROIs, state T5 generated smaller changes in the BOLD signal relative to other states. Interestingly, unique to right V1, state T2 generated significantly smaller BOLD signal changes relative to states T4 and T6. Similar results were observed for the undershoot

models, with state T2 and T5 exhibiting larger undershoot relative to the other states and state T1 generating a smaller undershoot relative to the other states.

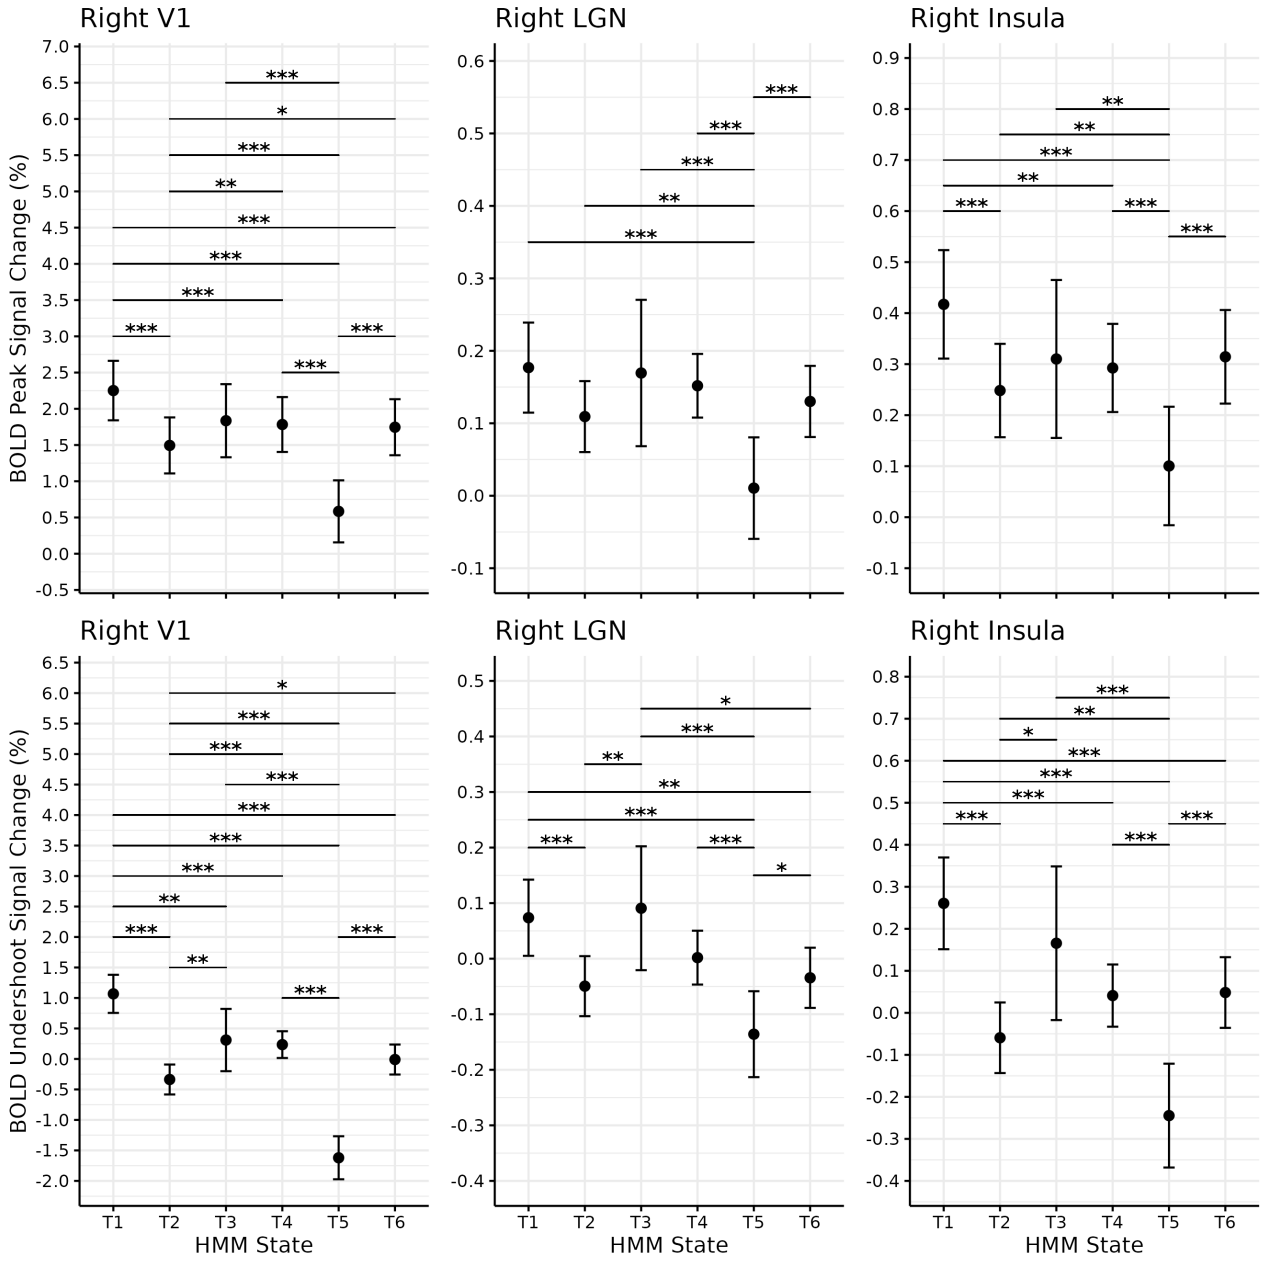


Figure 6.14: Comparison of BOLD peak and undershoot estimated marginal mean amplitudes for each baseline HMM state. Post hoc tests were conducted using an estimated marginal means approach and corrected via Bonferroni correction. *** = $p < 0.001$, ** = $p < 0.01$, * = $p < 0.05$. Error bars represent 95% confidence intervals.

Impact of Thalamus State on VEPs

The grand average VEPs from each state were calculated and compared to evaluate the impact of the baseline state on the subsequent VEP (see Figure 6.15). The results showed that the baseline thalamus states had a notable impact on the amplitudes of many of the VEP components. Specifically, state T4 generated the largest P100 amplitudes when trials were grouped across both conditions for trials, and within the active condition, whilst state T5 showed smaller P100 amplitudes than the other states. Conversely, when averaging across trials from the passive condition, state T3 showed a larger P100 amplitude relative to all other states. The impact of the baseline thalamus state was also observed on the N150 component, with state T3 showing the smallest N150 amplitudes compared with the other states. Lastly, the P250 component showed a large difference in amplitudes between states. State T1 generated the largest P250 amplitudes when grouping across all trials and trials from the active condition, whilst state T5 showed the smallest P250 amplitudes. During the passive condition, state T4 showed the largest P250 amplitude, whilst T5 showed the smallest.

To investigate the impact of condition, stimulus type and thalamus state on the amplitudes of the VEP components (P100, N150, P250), linear mixed models were fited for each component. The main effect of baseline thalamus was significant for the amplitudes of the P100 ($F(5, 3880.4) = 2.55, p < 0.05$) and P250 component ($F(5, 3880.4) = 2.49, p < 0.05$). Additionally, the N150 model found the main effect of condition to be significant ($F(1, 3868.8) = 9.47, p < 0.01$), with the N150 component amplitudes being significantly larger during the active condition ($EMM = -4.36$ µV, $SE = 0.658$ µV) relative to the passive condition ($EMM = -2.85$ µV, $SE = 0.671$ µV). All other main and interaction effects were non-significant ($p > 0.05$).

Post hoc tests were conducted to investigate the thalamus state main effect on the P100 and P250 amplitudes. All comparisons (Bonferroni corrected, number of comparisons = 15) for the P100 component were non-significant ($p > 0.05$). However, a significant comparison was found

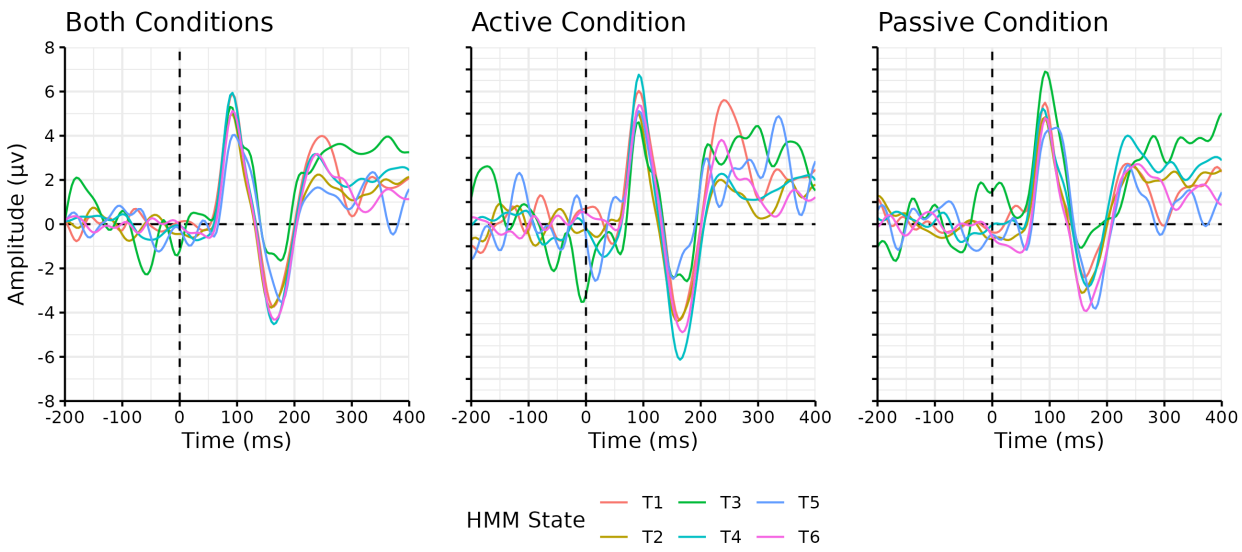


Figure 6.15: Comparison of the grand average VEPs across states. The left plot shows the grand average VEPs across all trials, whilst the centre and right plots show the grand average VEPs from trials within the active and passive conditions respectively. Error ribbons were omitted for visual clarity.

for the P250 component (Bonferroni corrected, number of comparisons =15), with the mean P250 amplitude being significantly larger when preceded by state T1 relative to T5 ($t(3878) = 3.202, p < 0.05$) (see Figure 6.16).

Impact of Thalamus State on Stimulus Induced Alpha Responses

Similar to the VEP comparison, the grand average alpha ERD/S plots were calculated to assess the impact of the baseline state on the subsequent induced alpha desynchronisation and synchronisation (see Figure). Baseline state overall had no effect on the frequency or the magnitude of the alpha desynchronisation, irrespective of the condition. However, a minor impact of state was observed for the alpha synchronisation, with state T5, and, to a lesser extent state T3, both showing an increased magnitude in power relative to other states. This effect was only observed during the active condition, with no obvious changes in the alpha synchronisation for states three or five during the passive condition.

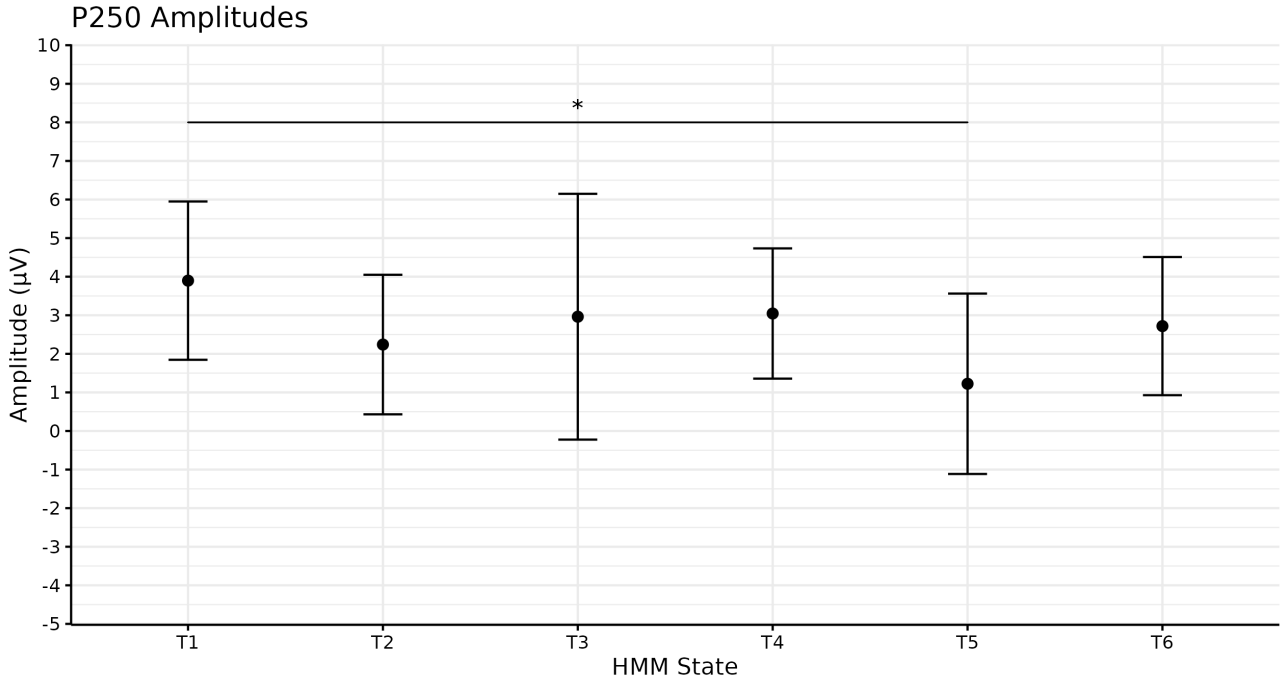


Figure 6.16: Comparison of P250 estimated marginal mean amplitudes for each baseline HMM state. Post hoc tests were conducted using an estimated marginal means approach, corrected using Bonferroni correction. = $p < 0.05$. Error bars represent 95% confidence intervals of the estimated marginal mean.

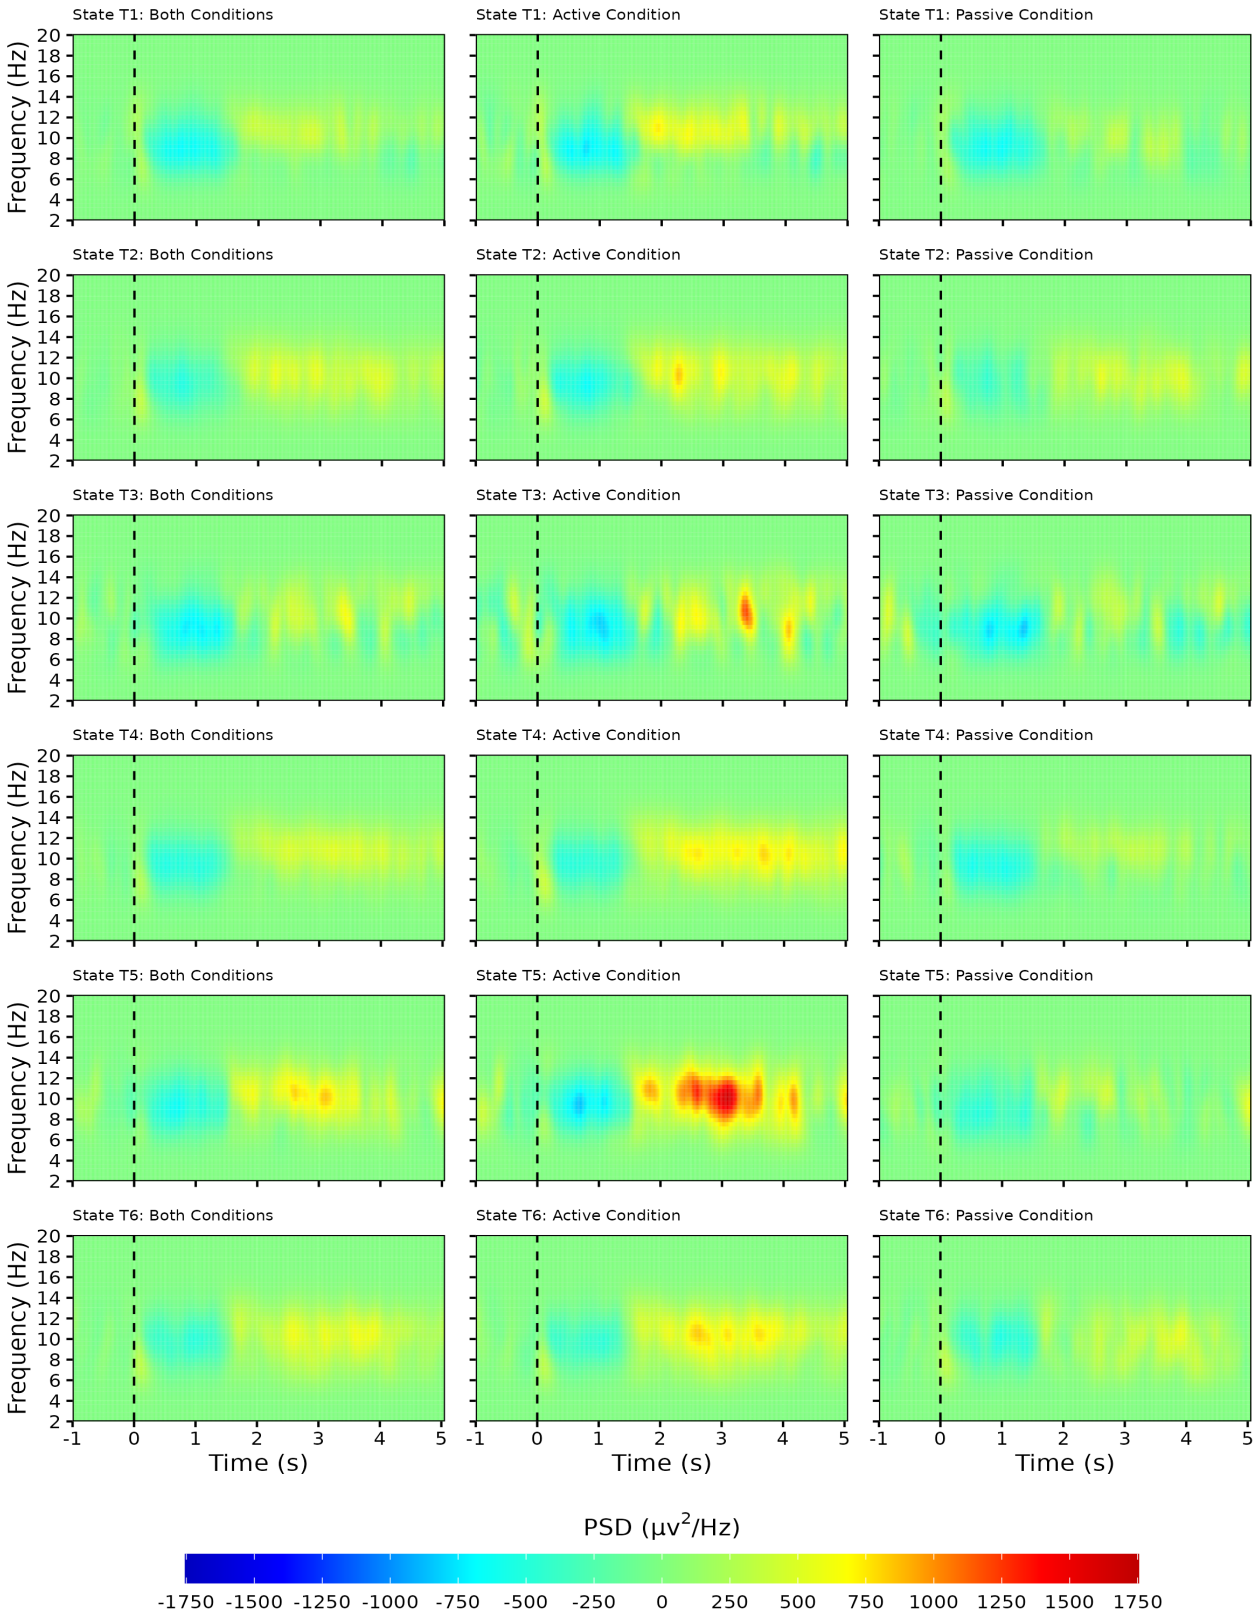

State T1: Both Conditions
State T1: Active Condition
State T1: Passive Condition
State T2: Both Conditions
State T2: Active Condition
State T2: Passive Condition
State T3: Both Conditions
State T3: Active Condition
State T3: Passive Condition
State T4: Both Conditions
State T4: Active Condition
State T4: Passive Condition
State T5: Both Conditions
State T5: Active Condition
State T5: Passive Condition
State T6: Both Conditions
State T6: Active Condition
State T6: Passive Condition
Frequency (Hz)
Time (s)
PSD (μv²/Hz)
-1750 -1500 -1250 -1000 -750 -500 -250 0 250 500 750 1000 1250 1500 1750

Figure 6.17: Grand average alpha ERD/S for each thalamus state. The left column shows the averages calculated with all trials, whilst the center and right columns show the averages calculated using trials from the active and passive conditions, respectively. PSD = power spectral density.

As with the BOLD responses and VEP, LMMs were fitted to the alpha ERD and ERS amplitudes to statistically assess the impact of the condition, stimulus type and baseline thalamus state (see Chapter 2 for details). No significant main or interaction effects were found for the alpha ERD model ($p > 0.05$) Conversely, the results of the ERS model found the main effect of the condition to be significant ($F(1, 3770.5) = 26.57, p < 0.001$), with the ERS being significantly larger within the active condition relative ($EMM = 147.8$ µV^2/Hz, $SE = 21.2$ µV^2/Hz) to the passive condition ($EMM = 43.1$ µV^2/Hz, $SE = 21.8$ µV^2/Hz). All other main and interaction effects were non-significant ($p > 0.05$).

Discussion

Summary

This chapter used HMM of EEG-fMRI data to conduct an exploratory analysis into the BOLD signal dynamics of the thalamus and investigated their relationship with the BOLD signal over the rest of the brain and the power of the alpha rhythm. In addition, we assessed the impact of thalamic states on the shape and amplitude of visually evoked BOLD responses, VEPs, and induced alpha ERD/S. BOLD responses from right V1 and the right LGN were significantly modulated by stimulus condition. A functional parcellation of the thalamus was then acquired using an ICA approach (Hale et al., 2015; S. Zhang & Li, 2017), and the IC time courses were inputted into an HMM (Baker et al., 2014; Vidaurre, Abeysuriya, et al., 2018; Vidaurre et al.,

2016; Vidaurre et al., 2017) to model the BOLD dynamics over time as a discrete set of states. A total of six states were reliably inferred by the HMM, and cortical BOLD signal correlates of each state were discovered by using state visits as an explanatory variable within a GLM analysis. Five of the six states showed significant cortical and thalamic BOLD correlations, primarily within bilateral visual and parietal cortex. Furthermore, each of the states showed a unique temporal relationship with posterior alpha power, with several states showing a change in alpha power approximately five seconds after the onset of the state, whilst one state showed changes in alpha power simultaneously with the state's onset. Regarding the impact of the thalamic states on the evoked responses, BOLD responses within right V1, the LGN and right insula were all significantly modulated by the baseline thalamic state, as was the P250 VEP component.

BOLD Responses in the LGN are Modulated by Attention

Right LGN peak BOLD responses were significantly larger during the active, attend, condition relative to the passive condition despite the stimulus input being identical across conditions. This finding is supported by the literature, with primate research demonstrating that visual attention increases spike frequency with the LGN in response to visual stimuli (McAlonan et al., 2008), and human fMRI research showing an increase in the peak LGN BOLD response to attended visual stimuli (O'Connor et al., 2002). The exact mechanism behind this modulation is not entirely understood. However, research in mice has shown that the connections between primary visual cortex and the dorsolateral geniculate nucleus (Sherman & Guillery, 2002) can modulate responses within the LGN through a form of cortico-thalamic feedback (Reinhold et al., 2023). Additionally, the LGN is connected to other subcortical structures, both within the thalamus and outside of the thalamus, including the thalamic reticular nucleus (TRN) and brainstem, which may also be capable of modulating the responses within the LGN. This is again supported by primate research, which has demonstrated that the initial increase in spike

frequency within the LGN was accompanied by a decrease in the spike frequency within TRN (McAlonan et al., 2008). As the TRN inhibits the other nuclei of the thalamus, this initial inhibition in spike frequency subsequently results in the release of inhibition over the LGN, increasing the spike frequency. Therefore, our findings are supported by the current literature and confirm that the task produced the expected results. Although the exact mechanism is not understood, the increase in BOLD response amplitudes within the LGN is hypothesised to be a result of a combination of cortico-thalamic feedback from primary visual cortex, and thalamo-thalamic connections from the TRN to the LGN.

Thalamic States and Cortical Correlates

The HMM identified a total of six thalamic states, each with unique BOLD activity. States T1, T2 and T5 showed large activations or deactivations spanning the entirety of the thalamus, whilst states T3, T4 and T6 showed lateralised activity, with activations in one hemisphere and deactivations within the other. Despite this, none of the cortical correlations were lateralised, with the majority (T1, T2, T4, and T5) showing bilateral activations or deactivations within the visual cortex, as well as parietal and motor regions for states T1 and T5. No resemblance was found between our results and the patterns of thalamic BOLD correlated with resting-state networks (D. J. Greene et al., 2020), which were more spatially constrained compared to the broad activations and deactivations observed within this study. However, the cortical patterns of states T1 and T4 strongly resembled the BOLD correlations typically observed with posterior alpha power, with bilateral deactivation of visual cortex and activations within the thalamus (Goldman et al., 2002; Laufs et al., 2006; Mayhew & Bagshaw, 2017). Therefore, this suggests that the HMM was able to identify meaningful thalamic states that showed plausible BOLD correlates, closely resembling the BOLD correlates of alpha activity, indicating a potential relationship between the thalamic states and posterior alpha power.

Alpha Power Correlations

We observed a significant difference in posterior alpha power between thalamic states at a lag of 5 TR. This suggests that the occupation of a thalamic state occurs prior to the change in alpha power, potentially inducing the change in posterior alpha power. States T1 and T3 showed significantly higher alpha power than the majority of the other states, whilst T5 showed significantly lower alpha power. Interestingly, state T1 also showed strong negative BOLD correlations within visual cortex and parietal cortex, resembling the BOLD correlates of posterior alpha power (Goldman et al., 2002; Laufs et al., 2006; Mayhew & Bagshaw, 2017), whilst T5 showed the same pattern of activity, but inverted. Thus, it appears that states T1 and T5 both have a close relationship with alpha power, with both states showing significant differences in posterior alpha power and BOLD correlates similar to the alpha BOLD correlations seen within the literature and Chapters 3 and 5. In addition, state T3 showed interesting dynamics with alpha power, with alpha power being highest at the time of state onset and rapidly reducing with the introduction of a lead or lag, suggesting the increase in alpha occurs simultaneously with the onset of the state. However, unlike state T1, the state showed no significant BOLD signal correlations with alpha power, and had the lowest fractional occupancy out of all the states, during both the active and passive conditions, indicating that the state may represent noise. Despite this, the results show that BOLD dynamics within the thalamus exhibit unique temporal relationships with posterior alpha power, indicating a highly complex spatiotemporal relationship between the thalamus and the alpha rhythm.

Impact of Baseline Thalamus State on BOLD Responses

The baseline thalamic states had a significant impact on the amplitudes of the BOLD response peak and undershoot within right V1, LGN and insula. Most notably, when state T5 was active during the baseline period prior to stimulation, the subsequent BOLD responses showed

significantly smaller peak amplitudes and larger undershoot amplitudes (e.g. more negative). The opposite effect was observed for state T1, which showed significantly larger peak amplitudes and smaller undershoot amplitudes. Interestingly, the BOLD correlates of states T1 and T5 were the opposite to each other, with state T1 showing large deactivations in visual and parietal cortex, whilst state T5 showed activations in the regions. These findings closely match the results found in Chapter 4, as well as within the literature (Blockley et al., 2013; Cohen et al., 2002; Gauthier et al., 2011; He, 2013; Mark et al., 2015; Stefanovic et al., 2006), with an increase in the BOLD signal within visual cortex resulting in smaller BOLD responses, and vice versa. As discussed within Chapter 3, these results contradict the linear summation hypothesis, indicating that baseline thalamus state and BOLD response amplitude interact non-linearly. One explanation is that these results represent a ceiling effect: when the BOLD signal is higher at baseline, it is closer to the 'ceiling', and therefore, subsequent BOLD responses are smaller, as they reach the ceiling quicker than when the BOLD signal is lower at the baseline (Mark et al., 2015). However, similar to the results in Chapter 3, states T1 and T5 also modulated the P250 amplitudes in the same way as the BOLD responses, indicating that the baseline thalamus state also impacts the neural activity in visual cortex. One explanation is that the thalamus intentionally modulates BOLD activity within visual cortex to modulate subsequent neural responses, facilitating visual attention.

Impact of Baseline Thalamus State on VEPs

Similar to the results of the BOLD analysis, the baseline thalamus state also had a significant impact on the amplitude of the P250 component. Specifically, P250 amplitudes were significantly higher when trials were preceded by state T1 relative to state T5. Though the temporal relationship between alpha power and thalamic states was at a lag, these results resemble what would be expected for a high and low alpha power state, with both Chapter 5 and previous literature demonstrating an increase in the amplitude of the P250 component when preceded by

high alpha power, both for visual (Becker et al., 2008) and somatosensory responses (Reinacher et al., 2009). Though the mechanism behind this effect is still poorly understood, one possible explanation is that the increase in P250 amplitude represents a form of top-down control from higher-order areas (Reinacher et al., 2009). Regardless of this, the results show a clear impact of thalamic baseline activity on the amplitude of the P250 component, suggesting that the thalamus may also be involved in facilitating top-down control over visual cortex (Griffiths et al., 2022; Kosciessa et al., 2021; Rikhye et al., 2018; Schmitt et al., 2017; Shine, 2021).

Impact of Baseline Thalamus State on Stimulus Induced Alpha ERD/S

Unlike the BOLD response and the VEP components, the baseline thalamus state had no significant impact on the alpha ERS or ERD. Given that many of the thalamic states appeared to show a link with alpha power, this result was unexpected given the results of Chapter 5 and the fact that previous literature (Mayhew, Ostwald, et al., 2013) has shown a link between prestimulus alpha and the magnitude of the alpha ERD/S. One explanation for this discrepancy may be because the alpha power typically decreased/increased 5 TR after the onset of the thalamic state, and therefore, whether a state is active at the time of baseline does not directly reflect a high or low alpha state, as within Chapter 5. Alternatively, it may be a result of splitting the trials over a large number of states. Within Chapter 5, a two-state HMM was used, as opposed to a six-state HMM here, meaning that the trials are being split between more states, potentially increasing the variability. Lastly, the discrepancy could simply result from temporal resolution, in which alpha power may only modulate the visual responses when it is high immediately before the stimulus (i.e. within a few tens of milliseconds), as opposed to within a second-long window.

Future Research

Within this chapter, an ICA approach was used to parcellate the thalamus (Hale et al., 2015; S. Zhang & Li, 2017), an approach that is generally preferred for Hidden Markov Modelling over atlas-based approaches (Ahrends et al., 2022). Though ICA decompositions of the thalamus have been used previously (Hale et al., 2015; S. Zhang & Li, 2017), as with whole-brain ICA, it is often difficult to determine what each IC represents. Therefore, future research could employ an atlas-based ROI approach, e.g. the THOMAS atlas (Su et al., 2019), allowing for more specific claims to be made about the involvement of specific thalamic nuclei. However, it is noted that the thalamic nuclei often show a large amount of anatomical variation between individuals, and thus, using an atlas would come with a level of anatomical uncertainty. Additionally, in this study, the thalamic states were defined from data acquired while participants performed a visual task. This may have biased the states towards ones more involved with visual processing. This potentially explains why four out of the six HMM states showed significant cortical correlations within visual cortex. Thus, a similar analysis should be conducted on resting-state state data to better understand the BOLD dynamics of the thalamus without external stimuli or while participants perform tasks requiring different sensory modalities.

Conclusions

Overall, the results of this chapter demonstrated that the thalamus exhibits complex BOLD dynamics during a visual task, with different states showing unique relationships with the rest of the brain and posterior alpha power. Furthermore, these thalamic states significantly modulated responses to visual stimuli when active at baseline, influencing the amplitude of visual BOLD response peak and amplitude, as well as the amplitude of the P250 VEP component. Additionally, applying a similar analysis technique to resting-state data or under different sensory stimuli would identify whether the states observed here occur exclusively during a visual

task or are equally likely to occur during rest. The results presented here confirm the complex role of the thalamus, considerably beyond what would be expected of a simple sensory relay. What is happening in the thalamus impacts haemodynamic and electrophysiological responses to stimulation, and HMM provides a clear approach to extracting and quantifying the complex dynamics of the thalamus. The thalamus exhibits multiple unique BOLD states, with strong relationships with the rest of the brain and posterior alpha power, as well as modulating visual responses, highlighting the involvement of the thalamus in both the generation of the alpha rhythm and visual attention.

Chapter 7

Thesis Summary

Overview

Variability in the brain's response to identical visual stimuli has been observed across all scales,from single neuron recordings in animals (Arieli et al., 1996; Azouz & Gray, 1999), to fMRI (Bianciardi et al., 2009; Duann et al., 2002; M. Schölvinck et al., 2012) and EEG (Becker et al., 2008; Mayhew, Ostwald, et al., 2013) recordings within humans. This thesis applied EEG-fMRI to investigate how visual responses interact with the ongoing state dynamics within the brain.Throughout all the chapters, Hidden Markov models (Baker et al., 2014; Vidaurre et al., 2016) were used to quantify ongoing state dynamics, with chapters focusing on different definitions of state depending on the modality. Within Chapter 3, the ability of the HMM approach to identify broad changes in brain state in EEG-fMRI data was tested by investigating the differences in BOLD and alpha dynamics during EO and EC rest. Chapter 4 followed a similar model to Chapter 3, defining whole-brain BOLD states according to activity within intrinsic BOLD networks and investigating how they interact with visual responses. Chapter 5 defined states according to alpha power from EEG data, investigating the impact of high pre-stimulus alpha power on visual responses. Lastly, given the involvement of the thalamus with both the generation of the alpha rhythm (Becker et al., 2015; Suffczynski et al., 2001) and the coordination of intrinsic BOLD networks (D. J. Greene et al., 2020; Hwang et al., 2017), Chapter 6 defined states according to the BOLD activity solely within the thalamus. Furthermore, the alpha power correlates and whole-brain BOLD correlates of each thalamus state were

also investigated.

Significant Findings

Throughout this thesis, the impact of baseline state was demonstrated to significantly modulate the amplitudes of subsequent visual responses. In Chapter 4, the baseline state was defined according to the activity within intrinsic BOLD networks and had a significant impact on both the amplitudes of BOLD percent signal change and the amplitude of the P100 component,indicating that the baseline levels of BOLD modulate both the BOLD response and the neuronal activity in that region. Chapters 5 and 6 showed similar results,with alpha states correlating with the magnitude of the alpha ERD/S response and the thalamus states associated with modulations in the BOLD response and P250 VEP component. Overall, these results highlight the importance of accounting for the baseline state within both BOLD and electrophysiological research, with both being significantly associated with the variability in BOLD responses and electrophysiological responses. They demonstrate the impact of state across modalities and the influence specific brain regions like the thalamus can have on cortical visual responses.

Regarding the state dynamics themselves, the results of Chapter 3 found significant difference sin the state dynamics of both BOLD networks and alpha between EO and EC rest. This finding aligns with previous research (Agcaoglu et al., 2019; Jao et al., 2013; Marx et al., 2004; Weng et al., 2020), which has also shown the EO and EC conditions to be associated with similar patterns of BOLD activity as those identified by the HMM. Additionally, the results highlight the importance of choosing between EO and EC rest within resting state studies, given that the dynamics between the two vary substantially. Interestingly, when moving to the visual task data in Chapters 4-6, no significant differences were found in the state dynamics between the active and passive conditions in any of the HMMs. This suggests that the requirement to maintain attention throughout the active condition was insufficient to modulate the ongoing

state dynamics, which again highlights the dramatic changes observed between EO and EC rest. It is hypothesised that this occurred due to the long ISI periods, resulting in a lapse in participant's attention.

Within Chapters 3 and 5, the alpha HMMs were capable of identifying high alpha power states which showed BOLD correlates closely resembling the results of prior alpha-BOLD research (Goldman et al., 2002; Laufs et al., 2006; Mayhew & Bagshaw, 2017). However, unlike previous research, the BOLD correlates could be defined using only a small proportion of each run (e.g. 15% in Chapter 3, 40% in Chapter 5) instead of using the entire length of each run. This suggests that the canonical BOLD correlation with alpha power is primarily driven by short periods of high alpha instead of an ongoing rhythm. This finding agrees with the hypothesis that neural oscillations represent transient bursts as opposed to sustained rhythms (Quinn et al., 2019; Van Ede et al., 2018), which should be considered when investigating neural osscilations and their BOLD correlates.

Lastly, Chapter 6 demonstrated that, as with the cortex, reliable BOLD states can also be observed within the thalamus. Furthermore, these thalamic states show significant BOLD correlations with the rest of the brain and show a complex temporal relationship with alpha power. This result demonstrates that the thalamus exhibits complex spatiotemporal BOLD dynamics, highlighting the complexity of the thalamus and its involvement in other processes besides simply relaying sensory information to the cortex (Shine et al., 2023). Additionally, the results further support the involvement of the thalamus with the generation of the alpha rhythm (Becker et al., 2015; Suffczynski et al., 2001), suggesting that specific spatial configurations of thalamic activity is associated with an increase in alpha power.

Future Directions

Throughout the PhD, various behavioural measures were acquired, such as questionnaire data and active task performance data (e.g. accuracy and reaction times). This data was collected with the aim of investigating the HMM dynamics with the behavioural measures to observe whether the baseline state also influenced behavioural response variability. Given that previous literature has demonstrated that pre-stimulus alpha activity can impact visual perception (Hanslmayr et al., 2007; Kelly et al., 2009; Mathewson et al., 2009; Mazaheri et al., 2009; Schubert et al., 2009; Thut et al., 2006; Van Dijk et al., 2008), it is likely that whole-brain network states and the thalamus states defined in Chapters 4 and 6 may also influence the participant's behavioural responses. However, as the thesis primarily aimed to investigate the impact of baseline state on fMRI and EEG visual responses, there was not enough time to investigate the impact of baseline state on behavioural data. Thus, future research should aim to uncover the impact of baseline state on behavioural responses.

Additionally, various neurodevelopmental populations show an increased level of sensory response variability in EEG (Milne, 2011; Milne et al., 2009) and fMRI recordings (Dinstein et al., 2012; Haigh et al., 2015), such as autism and ADHD. Furthermore, research has shown that BOLD variability during resting state scans reduces with ageing (Garrett et al., 2010), indicating that brain state dynamics may also vary within ageing populations (Grady & Garrett, 2014). Thus, various questionnaires were acquired (e.g. the autism quotient) to investigate the influence of neurodevelopmental traits on their visual and behavioural response variability; however, due to time constraints, these analyses were not conducted. Given that the results of this thesis indicate that the baseline states can explain a significant amount of the variability in visual responses, future research should utilise the questionnaire data to investigate the impact of neurodevelopmental traits on state dynamics and visual response variability.

While the EEG-fMRI data collected could help characterise the spectral and spatial properties

of the HMM states, one method was ultimately used to inform the other instead of combining the spatial and temporal information within a singular HMM. Thus, future research should aim toward finding a method that uses both the high temporal resolution of the EEG data and the high spatial resolution of the fMRI data within one HMM. One way this could be achieved is by including a downsampled alpha and beta power time series along with the fMRI network time series, allowing the HMM states to be defined according to intrinsic network activity and alpha/beta power. Alternatively, a joint ICA approach (Calhoun et al., 2006; Calhoun, Liu, et al., 2009; Mijović et al., 2012; Moosmann et al., 2008) could be used to obtain time series derived from both the EEG and fMRI data, which could then be inputted within a single HMM, again resulting in states defined by the properties of both modalities.

Future research should also explore alternative preprocessing methods to help reduce the impact of residual gradient artefacts on subsequent analyses. Such preprocessing methods include source localisation, which has been demonstrated to attenuate a proportion of the remaining gradient artefact harmonics (Brookes et al., 2005). Additionally, blind source separation techniques, such as functional source separation (Barbati et al., 2008; Barbati et al., 2006; Porcaro et al., 2010) could be used to extract task-related components, which would help reduce the gradient artefact's negative impact and provide more accurate single-trial measures.Furthermore, reducing the impact of the residual gradient artefacts would make it possible to investigate the impact of baseline states on ERD/S within higher frequency bands (e.g. high beta or low gamma). This would also allow for the HMMs to be trained on data containing higher frequencies (I.e. using a more complex observation model, such as a TDE model (Vidaurre, Abeysuriya, et al., 2018)), and thus, it would be possible to investigate the impact of other frequency bands on visual responses.

Furthermore, the definition of baseline could be expanded to further explore the relationship between ongoing state dynamics and visual response variability. Within this thesis, the baseline

was defined as the state active one timepoint before the stimulus onset, which does not take into account any of the state dynamics themselves, e.g. switching rate or state lifetimes. An alternative approach could obtain the state dynamics within a window before stimulus on set and group trials according to these state metrics to evaluate how these dynamics influence visual response variability.

Lastly, future research could expand on the results found within this thesis by investigating the impact of ongoing state dynamics on spatial variability within the fMRI data. The results of this thesis primarily focussed on the temporal element of the visual responses, e.g. finding the peak amplitude of a response. However, it is likely that the visual responses also show a level of variability spatially. For example, trials preceded by state x may generate a response within the DAN, whilst states preceded by state y do not. This could be achieved via a GLM approach and would help reveal if ongoing state dynamics also affect the spatial variability of visual responses as observed in the temporal responses.

Conclusion

This thesis used simultaneous EEG-fMRI data to investigate the relationship between ongoing brain state dynamics and visual response variability. The impact of whole-brain BOLD states, alpha states, and thalamus BOLD states was investigated, with each having a significant association with visual BOLD responses, VEPs, and alpha ERD/s. Together, the results of this thesis highlight the importance of accounting for the ongoing brain state dynamics within visual imaging studies and neuroimaging research as a whole.

Bibliography

Aeschbach, D., Matthews, J. R., Postolache, T. T., Jackson, M. A., Giesen, H. A., & Wehr, T. A. (1997). Dynamics of the human EEG during prolonged wakefulness: Evidence for frequency-specific circadian and homeostatic influences. *Neuroscience Letters*, *239*(2), 121–124.

Agcaoglu, O., Wilson, T. W., Wang, Y.-P., Stephen, J., & Calhoun, V. D. (2019). Resting state connectivity differences in eyes open versus eyes closed conditions. *Human Brain Mapping*, *40*(8), 2488–2498.

Aguirre, G., Zarahn, E., & D'Esposito, M. (1998). The variability of human, BOLD hemodynamic responses. *NeuroImage*, *8*(4), 360–369.

Ahrends, C., Stevner, A., Pervaiz, U., Kringelbach, M., Vuust, P., Woolrich, M., & Vidaurre, D. (2022). Data and model considerations for estimating time-varying functional connectivity in fMRI. *NeuroImage*, *252*, 119026.

Ahrends, C., & Vidaurre, D. (2023). Dynamic functional connectivity [Publisher: arXiv Version Number: 1].

Allen, E. A., Damaraju, E., Eichele, T., Wu, L., & Calhoun, V. D. (2018). EEG signatures of dynamic functional network connectivity states. *Brain Topography*, *31*(1), 101–116.

Allen, E. A., Damaraju, E., Plis, S. M., Erhardt, E. B., Eichele, T., & Calhoun, V. D. (2014). Tracking whole-brain connectivity dynamics in the resting state. *Cerebral Cortex*, *24*(3), 663–676.

Allen, P. J., Josephs, O., & Turner, R. (2000). A method for removing imaging artifact from continuous EEG recorded during functional MRI. *NeuroImage*, *12*(2), 230–239.

Allen, P. J., Polizzi, G., Krakow, K., Fish, D. R., & Lemieux, L. (1998). Identification of EEG events in the MR scanner: The problem of pulse artifact and a method for its subtraction. *NeuroImage*, *8*(3), 229–239.

Appelhoff, S., Sanderson, M., Brooks, T., Van Vliet, M., Quentin, R., Holdgraf, C., Chaumon, M., Mikulan, E., Tavabi, K., Höchenberger, R., Welke, D., Brunner, C., Rockhill, A., Larson, E., Gramfort, A., & Jas, M. (2019). MNE-BIDS: Organizing electrophysiological data into the BIDS format and facilitating their analysis. *Journal of Open Source Software*, *4*(44), 1896.

Araújo, D. B. D., Baffa, O., & Wakai, R. T. (2002). Theta oscillations and human navigation: A magnetoencephalography study. *Journal of Cognitive Neuroscience*, *14*(1), 70–78.

Arcaro, M. J., Pinsk, M. A., & Kastner, S. (2015). The anatomical and functional organization of the human visual pulvinar. *Journal of Neuroscience*, *35*(27), 9848–9871.

Arieli, A., Shoham, D., Hildesheim, R., & Grinvald, A. (1995). Coherent spatiotemporal patterns of ongoing activity revealed by real-time optical imaging coupled with single-unit recording in the cat visual cortex. *Journal of Neurophysiology*, *73*(5), 2072–2093.

Arieli, A., Sterkin, A., Grinvald, A., & Aertsen, A. (1996). Explanation of the large variability in evoked cortical responses. *273*.

Avants, B., Epstein, C., Grossman, M., & Gee, J. (2008). Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. *Medical Image Analysis*, *12*(1), 26–41.

Azouz, R., & Gray, C. M. (1999). Cellular mechanisms contributing to response variability of cortical neurons *In Vivo*. *The Journal of Neuroscience*, *19*(6), 2209–2223.

Bagshaw, A. P., & Warbrick, T. (2007). Single trial variability of EEG and fMRI responses to visual stimuli. *NeuroImage*, *38*(2), 280–292.

Baizer, J. S., Desimone, R., & Ungerleider, L. G. (1993). Comparison of subcortical connections of inferior temporal and posterior parietal cortex in monkeys. *Visual Neuroscience*, *10*(1), 59–72.

Baker, A. P., Brookes, M. J., Rezek, I. A., Smith, S. M., Behrens, T., Probert Smith, P. J., & Woolrich, M. (2014). Fast transient networks in spontaneous human brain activity. *eLife*, *3*, e01867.

Barbati, G., Porcaro, C., Hadjipapas, A., Adjamian, P., Pizzella, V., Romani, G. L., Seri, S., Tecchio, F., & Barnes, G. R. (2008). Functional source separation applied to induced visual gamma activity. *Human Brain Mapping*, *29*(2), 131–141.

Barbati, G., Sigismondi, R., Zappasodi, F., Porcaro, C., Graziadio, S., Valente, G., Balsi, M., Rossini, P. M., & Tecchio, F. (2006). Functional source separation from magnetoencephalographic signals. *Human Brain Mapping*, *27*(12), 925–934.

Barnes. (2010). Identifying basal ganglia divisions in individuals using resting-state functional connectivity MRI. *Frontiers in System Neuroscience*.

Barron, D. S., Eickhoff, S. B., Clos, M., & Fox, P. T. (2015). Human pulvinar functional organization and connectivity. *Human Brain Mapping*, *36*(7), 2417–2431.

Barry, R. J., Clarke, A. R., Johnstone, S. J., Magee, C. A., & Rushby, J. A. (2007). EEG differences between eyes-closed and eyes-open resting conditions. *Clinical Neurophysiology*, *118*(12), 2765–2773.

Basha, M. (2012). Simultaneous EEG and fMRI: Recording, analysis, and application, markus ullsperger, stefan debener (eds.). oxford university press (2010), $125.00, ISBN: 978-0-19-537273-1. *Journal of the Neurological Sciences*, *316*, 191.

Bastiaansen, M. C., Böcker, K. B., Brunia, C. H., De Munck, J. C., & Spekreijse, H. (2001). Event-related desynchronization during anticipatory attention for an upcoming stimulus: A comparative EEG/MEG study. *Clinical Neurophysiology*, *112*(2), 393–403.

Bastiaansen, M. C., & Brunia, C. H. (2001). Anticipatory attention: An event-related desynchronization approach. *International Journal of Psychophysiology*, *43*(1), 91–107.

Bastos, A. M., Vezoli, J., Bosman, C. A., Schoffelen, J.-M., Oostenveld, R., Dowdall, J. R., De Weerd, P., Kennedy, H., & Fries, P. (2015). Visual areas exert feedforward and feedback influences through distinct frequency channels. *Neuron*, *85*(2), 390–401.

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using **lme4**. *Journal of Statistical Software*, *67*(1).

Becker, R., Reinacher, M., Freyer, F., Villringer, A., & Ritter, P. (2011). How ongoing neuronal oscillations account for evoked fMRI variability. *Journal of Neuroscience*, *31*(30), 11016–11027.

Becker, R., Knock, S., Ritter, P., & Jirsa, V. (2015). Relating alpha power and phase to population firing and hemodynamic activity using a thalamo-cortical neural mass model (J. Daunizeau, Ed.). *PLOS Computational Biology*, *11*(9), e1004352.

Becker, R., Ritter, P., & Villringer, A. (2008). Influence of ongoing alpha rhythm on the visual evoked potential. *NeuroImage*, *39*(2), 707–716.

Beckmann, C., Mackay, C., Filippini, N., & Smith, S. (2009). Group comparison of resting-state FMRI data using multi-subject ICA and dual regression. *NeuroImage*, *47*, S148.

Beckmann, C., & Smith, S. (2004). Probabilistic independent component analysis for functional magnetic resonance imaging. *IEEE Transactions on Medical Imaging*, *23*(2), 137–152.

Behrens, T. E. J., Johansen-Berg, H., Woolrich, M. W., Smith, S. M., Wheeler-Kingshott, C. A. M., Boulby, P. A., Barker, G. J., Sillery, E. L., Sheehan, K., Ciccarelli, O., Thompson, A. J., Brady, J. M., & Matthews, P. M. (2003). Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging. *NATURE NEUROSCIENCE*, *6*(7).

Behrens, T., Woolrich, M., Jenkinson, M., Johansen-Berg, H., Nunes, R., Clare, S., Matthews, P., Brady, J., & Smith, S. (2003). Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magnetic Resonance in Medicine*, *50*(5), 1077–1088.

Behzadi, Y., Restom, K., Liau, J., & Liu, T. T. (2007). A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. *NeuroImage*, *37*(1), 90–101.

Bénar, C.-G., Grova, C., Kobayashi, E., Bagshaw, A. P., Aghakhani, Y., Dubeau, F., & Gotman, J. (2006). EEG–fMRI of epileptic spikes: Concordance with EEG source localization and intracranial EEG. *NeuroImage*, *30*(4), 1161–1170.

Bénar, C.-G., Schön, D., Grimault, S., Nazarian, B., Burle, B., Roth, M., Badier, J.-M., Marquis, P., Liegeois-Chauvel, C., & Anton, J.-L. (2007). Single-trial analysis of oddball event-related potentials in simultaneous EEG-fMRI. *Human Brain Mapping*, *28*(7), 602–613.

Berger, H. (1929). Über das Elektrenkephalogramm des Menschen. *Archiv für Psychiatrie und Nervenkrankheiten*, *87*(1), 527–570.

Bianciardi, M., Fukunaga, M., Van Gelderen, P., Horovitz, S. G., De Zwart, J. A., & Duyn, J. H. (2009). Modulation of spontaneous fMRI activity in human visual cortex by behavioral state. *NeuroImage*, *45*(1), 160–168.

Biswal, B., Zerrin Yetkin, F., Haughton, V. M., & Hyde, J. S. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar mri. *Magnetic Resonance in Medicine*, *34*(4), 537–541.

Blockley, N. P., Griffeth, V. E. M., Simon, A. B., & Buxton, R. B. (2013). A review of calibrated blood oxygenation level-dependent (BOLD) methods for the measurement of task-induced changes in brain oxygen metabolism: A REVIEW OF CALIBRATED BOLD METHODS. *NMR in Biomedicine*, *26*(8), 987–1003.

Bojak, I., Oostendorp, T. F., Reid, A. T., & Kötter, R. (2011). Towards a model-based integration of co-registered electroencephalography/functional magnetic resonance imaging

data with realistic neural population meshes. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *369*(1952), 3785–3801.

Bolt, T., Nomi, J. S., Bzdok, D., Salas, J. A., Chang, C., Thomas Yeo, B. T., Uddin, L. Q., & Keilholz, S. D. (2022). A parsimonious description of global functional brain organization in three spatiotemporal patterns. *Nature Neuroscience*, *25*(8), 1093–1103.

Bonmassar, G., Hadjikhani, N., Ives, J. R., Hinton, D., & Belliveau, J. W. (2001). Influence of EEG electrodes on the BOLD fMRI signal. *Human Brain Mapping*, *14*(2), 108–115.

Bonnefond, M., & Jensen, O. (2012). Alpha oscillations serve to protect working memory maintenance against anticipated distracters. *Current Biology*, *22*(20), 1969–1974.

Boostani, R., Karimzadeh, F., & Nami, M. (2017). A comparative review on sleep stage classification methods in patients and healthy individuals. *Computer Methods and Programs in Biomedicine*, *140*, 77–91.

Boytsova, Y. A., & Danko, S. G. (2010). EEG differences between resting states with eyes open and closed in darkness. *Human Physiology*, *36*(3), 367–369.

Brandman, T., Malach, R., & Simony, E. (2021). The surprising role of the default mode network in naturalistic perception. *Communications Biology*, *4*(1), 79.

Brenner, R. P. (1985). The electroencephalogram in altered states of consciousness. *Neurologic Clinics*, *3*(3), 615–631.

Bressler, D., Spotswood, N., & Whitney, D. (2007). Negative BOLD fMRI response in the visual cortex carries precise stimulus-specific information (P. Zak, Ed.). *PLoS ONE*, *2*(5), e410.

Brookes, M. J., Gibson, A. M., Hall, S. D., Furlong, P. L., Barnes, G. R., Hillebrand, A., Singh, K. D., Holliday, I. E., Francis, S. T., & Morris, P. G. (2005). GLM-beamformer method demonstrates stationary field, alpha ERD and gamma ERS co-localisation with fMRI BOLD response in visual cortex. *NeuroImage*, *26*(1), 302–308.

Brookes, M. J., Mullinger, K. J., Stevenson, C. M., Morris, P. G., & Bowtell, R. (2008). Simultaneous EEG source localisation and artifact rejection during concurrent fMRI by means of spatial filtering. *NeuroImage*, *40*(3), 1090–1104.

Brown, R. (2006). What is a brain state? *Philosophical Psychology*, *19*(6), 729–742.

Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience*, *8*(3), 183–208.

Buckner, R. L., Andrews-Hanna, J. R., & Schacter, D. L. (2008). *The Brain's Default Network*: *Anatomy, Function, and Relevance to Disease. Annals of the New York Academy of Sciences*, *1124*(1), 1–38.

Bullock, M., Jackson, G. D., & Abbott, D. F. (2021). Artifact reduction in simultaneous EEG-fMRI: A systematic review of methods and contemporary usage. *Frontiers in Neurology*, *12*, 622719.

Buxton, R. B., Wong, E. C., & Frank, L. R. (1998). Dynamics of blood flow and oxygenation changes during brain activation: The balloon model. *Magnetic Resonance in Medicine*, *39*(6), 855–864.

Cabral, J., Vidaurre, D., Marques, P., Magalhães, R., Silva Moreira, P., Miguel Soares, J., Deco, G., Sousa, N., & Kringelbach, M. L. (2017). Cognitive performance in healthy older adults relates to spontaneous switching between states of functional connectivity during rest. *Scientific Reports*, *7*(1), 5135.

Calhoun, V., Adali, T., Pearlson, G., & Kiehl, K. (2006). Neuronal chronometry of target detection: Fusion of hemodynamic and event-related potential data. *NeuroImage*, *30*(2), 544–553.

Calhoun, V., Eichele, T., & Pearlson, G. (2009). Functional brain networks in schizophrenia: A review. *Frontiers in Human Neuroscience*, *3*.

Calhoun, V., Liu, J., & Adalı, T. (2009). A review of group ICA for fMRI data and ICA for joint inference of imaging, genetic, and ERP data. *NeuroImage*, *45*(1), S163–S172.

Calhoun, V., Miller, R., Pearlson, G., & Adalı, T. (2014). The chronnectome: Time-varying connectivity networks as the next frontier in fMRI data discovery. *Neuron*, *84*(2), 262–274.

Carandini, M., Heeger, D. J., & Anthony Movshon, J. (1999). Linearity and gain control in v1 simple cells [Series Title: Cerebral Cortex]. In P. S. Ulinski, E. G. Jones, & A. Peters (Eds.). E. G. Jones & A. Peters (**typeredactors**), *Cerebral cortex* (pp. 401–443). Springer US.

Chen, X., Xu, X., Liu, A., Lee, S., Chen, X., Zhang, X., McKeown, M. J., & Wang, Z. J. (2019). Removal of muscle artifacts from the EEG: A review and recommendations. *IEEE Sensors Journal*, *19*(14), 5353–5368.

Christoff, K., Gordon, A. M., Smallwood, J., Smith, R., & Schooler, J. W. (2009). Experience sampling during fMRI reveals default network and executive system contributions to mind wandering. *Proceedings of the National Academy of Sciences*, *106*(21), 8719–8724.

Ciaramitaro, V. M., Buračas, G. T., & Boynton, G. M. (2007). Spatial and cross-modal attention alter responses to unattended sensory information in early visual and auditory human cortex. *Journal of Neurophysiology*, *98*(4), 2399–2413.

Cohen, E. R., Ugurbil, K., & Kim, S.-G. (2002). Effect of basal conditions on the magnitude and dynamics of the blood oxygenation level-dependent fMRI response. *Journal of Cerebral Blood Flow & Metabolism*, *22*(9), 1042–1053.

Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, *3*(3), 201–215.

Cox, R. W., & Hyde, J. S. (1997). Software tools for analysis and visualization of fMRI data. *NMR in Biomedicine*, *10*(4), 171–178.

Crottaz-Herbette, S., & Menon, V. (2006). Where and when the anterior cingulate cortex modulates attentional response: Combined fMRI and ERP evidence. *Journal of Cognitive Neuroscience*, *18*(5), 766–780.

Dale, A. M., Fischl, B., & Sereno, M. I. (1999). Cortical surface-based analysis. *NeuroImage*, *9*(2), 179–194.

Dale, A. M., Liu, A. K., Fischl, B. R., Buckner, R. L., Belliveau, J. W., Lewine, J. D., & Halgren, E. (2000). Dynamic statistical parametric mapping. *Neuron*, *26*(1), 55–67.

Damaraju, E., Allen, E., Belger, A., Ford, J., McEwen, S., Mathalon, D., Mueller, B., Pearlson, G., Potkin, S., Preda, A., Turner, J., Vaidya, J., van Erp, T., & Calhoun, V. D. (2014). Dynamic functional connectivity analysis reveals transient states of dysconnectivity in schizophrenia. *NeuroImage: Clinical*, *5*, 298–308.

De Souza, B. O. F., Cortes, N., & Casanova, C. (2020). Pulvinar modulates contrast responses in the visual cortex as a function of cortical hierarchy. *Cerebral Cortex*, *30*(3), 1068–1086.

De Valois, R. L., & De Valois, K. K. (1980). Spatial vision. *Annual Review of Psychology*, *31*(1), 309–341.

DeAngelis, G. C., Ohzawa, I., & Freeman, R. D. (1995). Receptive-field dynamics in the central visual pathways. *Trends in Neurosciences*, *18*(10), 451–458.

Debener, S., Strobel, A., Sorger, B., Peters, J., Kranczioch, C., Engel, A. K., & Goebel, R. (2007). Improved quality of auditory event-related potentials recorded simultaneously with 3-t fMRI: Removal of the ballistocardiogram artefact. *NeuroImage*, *34*(2), 587–597.

Debener, S., Ullsperger, M., Siegel, M., Fiehler, K., Von Cramon, D. Y., & Engel, A. K. (2005). Trial-by-trial coupling of concurrent electroencephalogram and functional magnetic resonance imaging identifies the dynamics of performance monitoring. *The Journal of Neuroscience*, *25*(50), 11730–11737.

Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, *134*(1), 9–21.

de Munck, J., Gonçalves, S., Huijboom, L., Kuijer, J., Pouwels, P., Heethaar, R., & Lopes da Silva, F. (2007). The hemodynamic response of the alpha rhythm: An EEG/fMRI study. *NeuroImage*, *35*(3), 1142–1151.

Di Russo, F., Martínez, A., Sereno, M. I., Pitzalis, S., & Hillyard, S. A. (2002). Cortical sources of the early components of the visual evoked potential: Cortical sources of VEP. *Human Brain Mapping*, *15*(2), 95–111.

Dinstein, I., Heeger, D. J., Lorenzi, L., Minshew, N. J., Malach, R., & Behrmann, M. (2012). Unreliable evoked responses in autism. *Neuron*, *75*(6), 981–991.

Doppelmayr, M., Finkenzeller, T., & Sauseng, P. (2008). Frontal midline theta in the pre-shot phase of rifle shooting: Differences between experts and novices. *Neuropsychologia*, *46*(5), 1463–1467.

Downar, J., Crawley, A. P., Mikulis, D. J., & Davis, K. D. (2000). A multimodal cortical network for the detection of changes in the sensory environment. *Nature Neuroscience*, *3*(3), 277–283.

Downar, J., Crawley, A. P., Mikulis, D. J., & Davis, K. D. (2001). The effect of task relevance on the cortical response to changes in visual and auditory stimuli: An event-related fMRI study. *NeuroImage*, *14*(6), 1256–1267.

Duann, J.-R., Jung, T.-P., Kuo, W.-J., Yeh, T.-C., Makeig, S., Hsieh, J.-C., & Sejnowski, T. J. (2002). Single-trial variability in event-related BOLD signals. *NeuroImage*, *15*(4), 823–835.

Esteban, O., Blair, R., Markiewicz, C. J., Berleant, S. L., Moodie, C., Ma, F., Isik, A. I., Erramuzpe, A., Goncalves, M., Poldrack, R. A., & Gorgolewski, K. J. (2017, September 25). *Poldracklab/fmriprep: 1.0.0-rc5* (Version 1.0.0-rc5). Zenodo.

Esteban, O., Markiewicz, C. J., Blair, R. W., Moodie, C. A., Isik, A. I., Erramuzpe, A., Kent, J. D., Goncalves, M., DuPre, E., Snyder, M., Oya, H., Ghosh, S. S., Wright, J., Durnez,

J., Poldrack, R. A., & Gorgolewski, K. J. (2019). fMRIPrep: A robust preprocessing pipeline for functional MRI. *Nature Methods*, *16*(1), 111–116.

Esteban, O., Markiewicz, C. J., Goncalves, M., Provins, C., Kent, J. D., DuPre, E., Salo, T., Ciric, R., Pinsard, B., Blair, R. W., Poldrack, R. A., & Gorgolewski, K. J. (2023, August 1). *fMRIPrep: A robust preprocessing pipeline for functional MRI* (Version 23.1.4). Zenodo.

Evans, A. C., Janke, A. L., Collins, D. L., & Baillet, S. (2012). Brain templates and atlases. *NeuroImage*, *62*(2), 911–922.

Fonov, V., Evans, A., McKinstry, R., Almli, C., & Collins, D. (2009). Unbiased nonlinear average age-appropriate brain templates from birth to adulthood. *NeuroImage*, *47*, S102.

Forney, G. (1973). The viterbi algorithm. *Proceedings of the IEEE*, *61*(3), 268–278.

Fortenbaugh, F. C., Rothlein, D., McGlinchey, R., DeGutis, J., & Esterman, M. (2018). Tracking behavioral and neural fluctuations during sustained attention: A robust replication and extension. *NeuroImage*, *171*, 148–164.

Fox, M. D., & Raichle, M. E. (2007). Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature Reviews Neuroscience*, *8*(9), 700–711.

Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., & Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences*, *102*(27), 9673–9678.

Fox, M. D., Snyder, A. Z., Zacks, J. M., & Raichle, M. E. (2006). Coherent spontaneous activity accounts for trial-to-trial variability in human evoked brain responses. *Nature Neuroscience*, *9*(1), 23–25.

Gandhi, S. P., Heeger, D. J., & Boynton, G. M. (1999). Spatial attention affects brain activity in human primary visual cortex. *Proceedings of the National Academy of Sciences*, *96*(6), 3314–3319.

Garrett, D. D., Kovacevic, N., McIntosh, A. R., & Grady, C. L. (2010). Blood oxygen level-dependent signal variability is more than just noise. *The Journal of Neuroscience*, *30*(14), 4914–4921.

Gauthier, C., Madjar, C., Tancredi, F., Stefanovic, B., & Hoge, R. (2011). Elimination of visually evoked BOLD responses during carbogen inhalation: Implications for calibrated MRI. *NeuroImage*, *54*(2), 1001–1011.

Goldman, R. I., Stern, J. M., Engel, J., & Cohen, M. S. (2002). Simultaneous EEG and fMRI of the alpha rhythm: *NeuroReport*, *13*(18), 2487–2492.

Goodale, M. A., & Milner, A. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences*, *15*(1), 20–25.

Gorgolewski, K. J., Burns, C. D., Madison, C., Clark, D., Halchenko, Y. O., Waskom, M. L., & Ghosh, S. S. (2011). Nipype: A flexible, lightweight and extensible neuroimaging data processing framework in python. *Frontiers in Neuroinformatics*, *5*.

Gorgolewski, K. J., Esteban, O., Burns, C., Ziegler, E., Pinsard, B., Madison, C., Waskom, M., Ellis, D. G., Clark, D., Dayan, M., Manhães-Savio, A., Notter, M. P., Johnson, H., Dewey, B. E., Halchenko, Y. O., Hamalainen, C., Keshavan, A., Clark, D., Huntenburg, J. M., … Ghosh, S. (2016, April 21). *Nipype: A flexible, lightweight and extensible neuroimaging data processing framework in python. 0.12.0-rc1* (Version 0.12.0-rc1). Zenodo.

Gotman, J., Kobayashi, E., Bagshaw, A. P., Bénar, C.-G., & Dubeau, F. (2006). Combining EEG and fMRI: A multimodal tool for epilepsy research. *Journal of Magnetic Resonance Imaging*, *23*(6), 906–920.

Grady, C. L., & Garrett, D. D. (2014). Understanding variability in the BOLD signal and why it matters for aging. *Brain Imaging and Behavior*, *8*(2), 274–283.

Green, J. J., & McDonald, J. J. (2008). Electrical neuroimaging reveals timing of attentional control activity in human brain (M. Posner, Ed.). *PLoS Biology*, *6*(4), e81.

Greene, A. S., Horien, C., Barson, D., Scheinost, D., & Constable, R. T. (2023). Why is everyone talking about brain state? *Trends in Neurosciences*, *46*(7), 508–524.

Greene, D. J., Marek, S., Gordon, E. M., Siegel, J. S., Gratton, C., Laumann, T. O., Gilmore, A. W., Berg, J. J., Nguyen, A. L., Dierker, D., Van, A. N., Ortega, M., Newbold, D. J., Hampton, J. M., Nielsen, A. N., McDermott, K. B., Roland, J. L., Norris, S. A., Nelson, S. M., … Dosenbach, N. U. (2020). INTEGRATIVE AND NETWORK-SPECIFIC CONNECTIVITY OF THE BASAL GANGLIA AND THALAMUS DEFINED IN INDIVIDUALS. *Neuron*, *105*(4), 742–758.e6.

Greicius, M. D., Srivastava, G., Reiss, A. L., & Menon, V. (2004). Default-mode network activity distinguishes alzheimer's disease from healthy aging: Evidence from functional MRI. *Proceedings of the National Academy of Sciences*, *101*(13), 4637–4642.

Greve, D. N., & Fischl, B. (2009). Accurate and robust brain image alignment using boundary-based registration. *NeuroImage*, *48*(1), 63–72.

Griffanti, L., Salimi-Khorshidi, G., Beckmann, C. F., Auerbach, E. J., Douaud, G., Sexton, C. E., Zsoldos, E., Ebmeier, K. P., Filippini, N., Mackay, C. E., Moeller, S., Xu, J., Yacoub, E., Baselli, G., Ugurbil, K., Miller, K. L., & Smith, S. M. (2014). ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging. *NeuroImage*, *95*, 232–247.

Griffiths, B. J., Zaehle, T., Repplinger, S., Schmitt, F. C., Voges, J., Hanslmayr, S., & Staudigl, T. (2022). Rhythmic interactions between the mediodorsal thalamus and prefrontal cortex precede human visual perception. *Nature Communications*, *13*(1), 3736.

Haigh, S. M., Heeger, D. J., Dinstein, I., Minshew, N., & Behrmann, M. (2015). Cortical variability in the sensory-evoked response in autism. *Journal of Autism and Developmental Disorders*, *45*(5), 1176–1190.

Halchenko, Y., Goncalves, M., Velasco, P., Di Oleggio Castello, M. V., Satrajit Ghosh, Salo, T., Wodder, J. T., Hanke, M., Sadil, P., Christian, H., Dae, Michael, Tilley, S., Kent, J., To,

I., Brett, M., Amlien, I., Gorgolewski, C., Markiewicz, C., … Lurie, D. (2023, May 23). *Nipy/heudiconv: V0.13.1* (Version v0.13.1). Zenodo.

Hale, J. R., Mayhew, S. D., Mullinger, K. J., Wilson, R. S., Arvanitis, T. N., Francis, S. T., & Bagshaw, A. P. (2015). Comparison of functional thalamic segmentation from seed-based analysis and ICA. *NeuroImage*, *114*, 448–465.

Hämäläinen, M. S., & Ilmoniemi, R. J. (1994). Interpreting magnetic fields of the brain: Minimum norm estimates. *Medical & Biological Engineering & Computing*, *32*(1), 35–42.

Hanslmayr, S., Spitzer, B., & Bauml, K.-H. (2009). Brain oscillations dissociate between semantic and nonsemantic encoding of episodic memories. *Cerebral Cortex*, *19*(7), 1631–1640.

Hanslmayr, S., Aslan, A., Staudigl, T., Klimesch, W., Herrmann, C. S., & Bäuml, K.-H. (2007). Prestimulus oscillations predict visual perception performance between and within subjects. *NeuroImage*, *37*(4), 1465–1473.

Harmony, T. (2013). The functional significance of delta oscillations in cognitive processing. *Frontiers in Integrative Neuroscience*, *7*.

Harris, K. D., & Thiele, A. (2011). Cortical state and attention. *Nature Reviews Neuroscience*, *12*(9), 509–523.

He, B. J. (2013). Spontaneous and task-evoked brain activity negatively interact. *The Journal of Neuroscience*, *33*(11), 4672–4682.

Heggelund, P., & Albus, K. (1978). Response variability and orientation discrimination of single cells in striate cortex of cat. *Experimental Brain Research*, *32*(2).

Heinze, H., Luck, S., Mangun, G., & Hillyard, S. (1990). Visual event-related potentials index focused attention within bilateral stimulus arrays. i. evidence for early selection. *Electroencephalography and Clinical Neurophysiology*, *75*(6), 511–527.

Hesselmann, G., Kell, C. A., Eger, E., & Kleinschmidt, A. (2008). Spontaneous local variations in ongoing neural activity bias perceptual decisions. *Proceedings of the National Academy of Sciences*, *105*(31), 10984–10989.

Hesselmann, G., Kell, C. A., & Kleinschmidt, A. (2008). Ongoing activity fluctuations in hMT+ bias the perception of coherent visual motion. *The Journal of Neuroscience*, *28*(53), 14481–14485.

Hesselmann, G., Sadaghiani, S., Friston, K. J., & Kleinschmidt, A. (2010). Predictive coding or evidence accumulation? false inference and neuronal fluctuations (J. Lauwereyns, Ed.). *PLoS ONE*, *5*(3), e9926.

Hillyard, S. A., & Anllo-Vento, L. (1998). Event-related brain potentials in the study of visual selective attention. *Proceedings of the National Academy of Sciences*, *95*(3), 781–787.

Hindriks, R., Adhikari, M., Murayama, Y., Ganzetti, M., Mantini, D., Logothetis, N., & Deco, G. (2016). Can sliding-window correlations reveal dynamic functional connectivity in resting-state fMRI? *NeuroImage*, *127*, 242–256.

Hu, X., & Yacoub, E. (2012). The story of the initial dip in fMRI. *NeuroImage*, *62*(2), 1103–1108.

Hubel, D. H., & Wiesel, T. N. (1959). Receptive fields of single neurones in the cat's striate cortex. *The Journal of Physiology*, *148*(3), 574–591.

Hughes, S. W., & Crunelli, V. (2005). Thalamic mechanisms of EEG alpha rhythms and their pathological implications. *The Neuroscientist*, *11*(4), 357–372.

Hunyadi, B., Woolrich, M., Quinn, A., Vidaurre, D., & De Vos, M. (2019). A dynamic system of brain networks revealed by fast transient EEG fluctuations and their fMRI correlates. *NeuroImage*, *185*, 72–82.

Huster, R. J., Debener, S., Eichele, T., & Herrmann, C. S. (2012). Methods for simultaneous EEG-fMRI: An introductory review. *The Journal of Neuroscience*, *32*(18), 6053–6060.

Hwang, K., Bertolero, M. A., Liu, W. B., & D'Esposito, M. (2017). The human thalamus is an integrative hub for functional brain networks. *The Journal of Neuroscience*, *37*(23), 5594–5607.

Hyvarinen, A. (1999). Fast and robust fixed-point algorithms for independent component analysis. *IEEE Transactions on Neural Networks*, *10*(3), 626–634.

Jao, T., Vértes, P. E., Alexander-Bloch, A. F., Tang, I.-N., Yu, Y.-C., Chen, J.-H., & Bullmore, E. T. (2013). Volitional eyes opening perturbs brain dynamics and functional connectivity regardless of light input. *NeuroImage*, *69*, 21–34.

Jasper, H. H. (1938). ELECTRO-ENCEPHALOGRAPHY: III. NORMAL DIFFERENTIATION OF OCCIPITAL AND PRECENTRAL REGIONS IN MAN. *Archives of Neurology & Psychiatry*, *39*(1), 96.

Jenkinson, M., Bannister, P., Brady, M., & Smith, S. (2002). Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, *17*(2), 825–841.

Jenkinson, M., & Smith, S. (2001). A global optimisation method for robust affine registration of brain images. *Medical Image Analysis*, *5*(2), 143–156.

Jensen, O., Bonnefond, M., Marshall, T. R., & Tiesinga, P. (2015). Oscillatory mechanisms of feedforward and feedback visual processing. *Trends in Neurosciences*, *38*(4), 192–194.

Jensen, O., Kaiser, J., & Lachaux, J.-P. (2007). Human gamma-frequency oscillations associated with attention and memory. *Trends in Neurosciences*, *30*(7), 317–324.

Jensen, O., & Mazaheri, A. (2010). Shaping functional architecture by oscillatory alpha activity: Gating by inhibition. *Frontiers in Human Neuroscience*, *4*.

Jensen, O., & Tesche, C. D. (2002). Frontal theta activity in humans increases with memory load in a working memory task: Frontal theta increases with memory load. *European Journal of Neuroscience*, *15*(8), 1395–1399.

Ji, H., Chen, B., Petro, N. M., Yuan, Z., Zheng, N., & Keil, A. (2019). Functional source separation for EEG-fMRI fusion: Application to steady-state visual evoked potentials. *Frontiers in Neurorobotics*, *13*, 24.

Johannes, S., Münte, T., Heinze, H., & Mangun, G. (1995). Luminance and spatial attention effects on early visual processing. *Cognitive Brain Research*, *2*(3), 189–205.

Jurkiewicz, M. T., Gaetz, W. C., Bostan, A. C., & Cheyne, D. (2006). Post-movement beta rebound is generated in motor cortex: Evidence from neuromagnetic recordings. *NeuroImage*, *32*(3), 1281–1289.

Kahana, M. J., Sekuler, R., Caplan, J. B., Kirschen, M., & Madsen, J. R. (1999). Human theta oscillations exhibit task dependence during virtual maze navigation. *Nature*, *399*(6738), 781–784.

Kaiser, J., & Lutzenberger, W. (2005). Cortical oscillatory activity and the dynamics of auditory memory processing. *Reviews in the Neurosciences*, *16*(3).

Karapanagiotidis, T., Vidaurre, D., Quinn, A. J., Vatansever, D., Poerio, G. L., Turnbull, A., Ho, N. S. P., Leech, R., Bernhardt, B. C., Jefferies, E., Margulies, D. S., Nichols, T. E., Woolrich, M. W., & Smallwood, J. (2020). The psychological correlates of distinct neural states occurring during wakeful rest. *Scientific Reports*, *10*(1), 21121.

Kelly, S. P., Gomez-Ramirez, M., & Foxe, J. J. (2009). The strength of anticipatory spatial biasing predicts target discrimination at attended locations: A high-density EEG study. *European Journal of Neuroscience*, *30*(11), 2224–2234.

Kelly, S. P., Lalor, E. C., Reilly, R. B., & Foxe, J. J. (2006). Increases in alpha oscillatory power reflect an active retinotopic mechanism for distracter suppression during sustained visuospatial attention. *Journal of Neurophysiology*, *95*(6), 3844–3851.

Kim, D.-S., Duong, T. Q., & Kim, S.-G. (2000). High-resolution mapping of iso-orientation columns by fMRI. *Nature Neuroscience*, *3*(2), 164–169.

Klein, A., Ghosh, S. S., Bao, F. S., Giard, J., Häme, Y., Stavsky, E., Lee, N., Rossa, B., Reuter, M., Chaibub Neto, E., & Keshavan, A. (2017). Mindboggling morphometry of human brains (D. Schneidman, Ed.). *PLOS Computational Biology*, *13*(2), e1005350.

Klimesch, W., Schimke, H., & Schwaiger, J. (1994). Episodic and semantic memory: An analysis in the EEG theta and alpha band. *Electroencephalography and Clinical Neurophysiology*, *91*(6), 428–441.

Klimesch, W., Sauseng, P., & Hanslmayr, S. (2007). EEG alpha oscillations: The inhibition–timing hypothesis. *Brain Research Reviews*, *53*(1), 63–88.

Kopell, N. J., Gritton, H. J., Whittington, M. A., & Kramer, M. A. (2014). Beyond the connectome: The dynome. *Neuron*, *83*(6), 1319–1328.

Kosciessa, J. Q., Lindenberger, U., & Garrett, D. D. (2021). Thalamocortical excitability modulation guides human perception under uncertainty. *Nature Communications*, *12*(1), 2430.

Koshiyama, D., Miyakoshi, M., Tanaka-Koshiyama, K., Joshi, Y. B., Molina, J. L., Sprock, J., Braff, D. L., & Light, G. A. (2020). Neurophysiologic characterization of resting state connectivity abnormalities in schizophrenia patients. *Frontiers in Psychiatry*, *11*, 608154.

Krakow, K., Allen, P., Symms, M., Lemieux, L., Josephs, O., & Fish, D. (2000). EEG recording during fMRI experiments: Image quality. *Human Brain Mapping*, *10*(1), 10–15.

Krishnamurthy, V., Krishnamurthy, L. C., Drucker, J. H., Kundu, S., Ji, B., Hortman, K., Roberts, S. R., Mammino, K., Tran, S. M., Gopinath, K., McGregor, K. M., Rodriguez, A. D., Qiu, D., Crosson, B., & Nocera, J. R. (2020). Correcting task fMRI signals for variability in baseline CBF improves BOLD-behavior relationships: A feasibility study in an aging model. *Frontiers in Neuroscience*, *14*, 336.

Kucyi, A., Esterman, M., Riley, C. S., & Valera, E. M. (2016). Spontaneous default network activity reflects behavioral variability independent of mind-wandering. *Proceedings of the National Academy of Sciences*, *113*(48), 13899–13904.

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). **lmerTest** package: Tests in linear mixed effects models. *Journal of Statistical Software*, *82*(13).

Lampl, I., Reichova, I., & Ferster, D. (1999). Synchronous membrane potential fluctuations in neurons of the cat visual cortex. *Neuron*, *22*(2), 361–374.

Lanczos, C. (1964). Evaluation of noisy data. *Journal of the Society for Industrial and Applied Mathematics Series B Numerical Analysis*, *1*(1), 76–85.

Laufs, H., Holt, J. L., Elfont, R., Krams, M., Paul, J. S., Krakow, K., & Kleinschmidt, A. (2006). Where the BOLD signal goes when alpha EEG leaves. *NeuroImage*, *31*(4), 1408–1418.

Laumann, T. O., Snyder, A. Z., Mitra, A., Gordon, E. M., Gratton, C., Adeyemo, B., Gilmore, A. W., Nelson, S. M., Berg, J. J., Greene, D. J., McCarthy, J. E., Tagliazucchi, E., Laufs, H., Schlaggar, B. L., Dosenbach, N. U. F., & Petersen, S. E. (2016). On the stability of BOLD fMRI correlations. *Cerebral Cortex*, cercor, bhw265v1.

Leber, A. B. (2010). Neural predictors of within-subject fluctuations in attentional control. *The Journal of Neuroscience*, *30*(34), 11458–11465.

Leh, S. E., Chakravarty, M. M., & Ptito, A. (2008). The connectivity of the human pulvinar: A diffusion tensor imaging tractography study. *International Journal of Biomedical Imaging*, *2008*, 1–5.

Leonardi, N., & Van De Ville, D. (2015). On spurious and real fluctuations of dynamic functional connectivity during rest. *NeuroImage*, *104*, 430–436.

Lesenfants, D., Habbal, D., Lugo, Z., Lebeau, M., Horki, P., Amico, E., Pokorny, C., Gómez, F., Soddu, A., Müller-Putz, G., Laureys, S., & Noirhomme, Q. (2014). An independent SSVEP-based brain–computer interface in locked-in syndrome. *Journal of Neural Engineering*, *11*(3), 035002.

Li, L. (2010). The differences among eyes-closed, eyes-open and attention states: An EEG study. *2010 International Conference on Computational Intelligence and Software Engineering*, 1–4.

Li, X., Lu, Z.-L., Tjan, B. S., Dosher, B. A., & Chu, W. (2008). Blood oxygenation level-dependent contrast response functions identify mechanisms of covert attention in early visual areas. *Proceedings of the National Academy of Sciences*, *105*(16), 6202–6207.

Liang, B., Zhang, D., Wen, X., Xu, P., Peng, X., Huang, X., Liu, M., & Huang, R. (2014). Brain spontaneous fluctuations in sensorimotor regions were directly related to eyes open and eyes closed: Evidences from a machine learning approach. *Frontiers in Human Neuroscience*, *8*.

Linden, D. E. (1999). The functional neuroanatomy of target detection: An fMRI study of visual and auditory oddball tasks. *Cerebral Cortex*, *9*(8), 815–823.

Lisman, J. E., & Jensen, O. (2013). The theta-gamma neural code. *Neuron*, *77*(6), 1002–1016.

Liu, T. T., Nalci, A., & Falahpour, M. (2017). The global signal in fMRI: Nuisance or information? *NeuroImage*, *150*, 213–229.

Luck, S., Heinze, H., Mangun, G., & Hillyard, S. (1990). Visual event-related potentials index focused attention within bilateral stimulus arrays. II. functional dissociation of p1 and n1 components. *Electroencephalography and Clinical Neurophysiology*, *75*(6), 528–542.

Mainen, Z. F., & Sejnowski, T. J. (1995). Reliability of spike timing in neocortical neurons. *Science*, *268*(5216), 1503–1506.

Makeig, S., Westerfield, M., Jung, T.-P., Enghoff, S., Townsend, J., Courchesne, E., & Sejnowski, T. J. (2002). Dynamic brain sources of visual evoked responses. *Science*, *295*(5555), 690–694.

Mandelkow, H., Halder, P., Boesiger, P., & Brandeis, D. (2006). Synchronization facilitates removal of MRI artefacts from concurrent EEG recordings and increases usable bandwidth. *NeuroImage*, *32*(3), 1120–1126.

Mangun, G. R. (1995). Neural mechanisms of visual selective attention. *Psychophysiology*, *32*(1), 4–18.

Mark, C. I., Mazerolle, E. L., & Chen, J. J. (2015). Metabolic and vascular origins of the BOLD effect: Implications for imaging pathology and resting-state brain function: Understanding BOLD in disease and rest. *Journal of Magnetic Resonance Imaging*, *42*(2), 231–246.

Marx, E., Deutschländer, A., Stephan, T., Dieterich, M., Wiesmann, M., & Brandt, T. (2004). Eyes open and eyes closed as rest conditions: Impact on brain activation patterns. *NeuroImage*, *21*(4), 1818–1824.

Marx, E., Stephan, T., Nolte, A., Deutschländer, A., Seelos, K. C., Dieterich, M., & Brandt, T. (2003). Eye closure in darkness animates sensory systems. *NeuroImage*, *19*(3), 924–934.

Mathewson, K. E., Gratton, G., Fabiani, M., Beck, D. M., & Ro, T. (2009). To see or not to see: Prestimulus α phase predicts visual awareness. *The Journal of Neuroscience*, *29*(9), 2725–2732.

Mayhew, S. D., & Bagshaw, A. (2017). Dynamic spatiotemporal variability of alpha-BOLD relationships during the resting-state and task-evoked responses. *NeuroImage*, *155*, 120–137.

Mayhew, S. D., Hylands-White, N., Porcaro, C., Derbyshire, S. W., & Bagshaw, A. P. (2013). Intrinsic variability in the human response to pain is assembled from multiple, dynamic brain processes. *NeuroImage*, *75*, 68–78.

Mayhew, S. D., Ostwald, D., Porcaro, C., & Bagshaw, A. P. (2013). Spontaneous EEG alpha oscillation interacts with positive and negative BOLD responses in the visual–auditory cortices and default-mode network. *NeuroImage*, *76*, 362–372.

Mayor, J., & Gerstner, W. (2005). Signal buffering in random networks of spiking neurons: Microscopic versus macroscopic phenomena. *Physical Review E*, *72*(5), 051906.

Mazaheri, A., Nieuwenhuis, I. L., Van Dijk, H., & Jensen, O. (2009). Prestimulus alpha and mu activity predicts failure to inhibit motor responses. *Human Brain Mapping*, *30*(6), 1791–1800.

McAlonan, K., Cavanaugh, J., & Wurtz, R. H. (2008). Guarding the gateway to cortex with attention in visual thalamus. *Nature*, *456*(7220), 391–394.

McKiernan, K. A., D'Angelo, B. R., Kaufman, J. N., & Binder, J. R. (2006). Interrupting the "stream of consciousness": An fMRI investigation. *NeuroImage*, *29*(4), 1185–1191.

Mijović, B., Vanderperren, K., Novitskiy, N., Vanrumste, B., Stiers, P., Van Den Bergh, B., Lagae, L., Sunaert, S., Wagemans, J., Van Huffel, S., & De Vos, M. (2012). The "why" and "how" of JointICA: Results from a visual detection task. *NeuroImage*, *60*(2), 1171–1185.

Miller, K. J., Weaver, K. E., & Ojemann, J. G. (2009). Direct electrophysiological measurement of human default network areas. *Proceedings of the National Academy of Sciences*, *106*(29), 12174–12177.

Milne, E. (2011). Increased intra-participant variability in children with autistic spectrum disorders: Evidence from single-trial analysis of evoked EEG. *Frontiers in Psychology*, *2*.

Milne, E., Scope, A., Pascalis, O., Buckley, D., & Makeig, S. (2009). Independent component analysis reveals atypical electroencephalographic activity during visual perception in individuals with autism. *Biological Psychiatry*, *65*(1), 22–30.

Mishkin, M., Ungerleider, L. G., & Macko, K. A. (1983). Object vision and spatial vision: Two cortical pathways. *Trends in Neurosciences*, *6*, 414–417.

Mokhtari, F., Akhlaghi, M. I., Simpson, S. L., Wu, G., & Laurienti, P. J. (2019). Sliding window correlation analysis: Modulating window shape for dynamic brain connectivity in resting state. *NeuroImage*, *189*, 655–666.

Moosmann, M., Eichele, T., Nordby, H., Hugdahl, K., & Calhoun, V. D. (2008). Joint independent component analysis for simultaneous EEG–fMRI: Principle and simulation. *International Journal of Psychophysiology*, *67*(3), 212–221.

Moosmann, M., Ritter, P., Krastel, I., Brink, A., Thees, S., Blankenburg, F., Taskin, B., Obrig, H., & Villringer, A. (2003). Correlates of alpha rhythm in functional magnetic resonance imaging and near infrared spectroscopy. *NeuroImage*, *20*(1), 145–158.

Mullinger, K. J. (2017). Post-stimulus fMRI and EEG responses_ evidence for a neuronal origin hypothesised to be inhibitory.

Mullinger, K., Debener, S., Coxon, R., & Bowtell, R. (2008). Effects of simultaneous EEG recording on MRI data quality at 1.5, 3 and 7 tesla. *International Journal of Psychophysiology*, *67*(3), 178–188.

Mullinger, K., Mayhew, S., Bagshaw, A., Bowtell, R., & Francis, S. (2014). Evidence that the negative BOLD response is neuronal in origin: A simultaneous EEG–BOLD–CBF study in humans. *NeuroImage*, *94*, 263–274.

Mullinger, K. J., Castellone, P., & Bowtell, R. (2013). Best current practice for obtaining high quality EEG data during simultaneous fMRI. *Journal of Visualized Experiments*, (76), 50283.

Mullinger, K. J., Mayhew, S. D., Bagshaw, A. P., Bowtell, R., & Francis, S. T. (2013). Poststimulus undershoots in cerebral blood flow and BOLD fMRI responses are modulated by poststimulus neuronal activity. *Proceedings of the National Academy of Sciences*, *110*(33), 13636–13641.

Mullinger, K. J., Morgan, P. S., & Bowtell, R. W. (2008). Improved artifact correction for combined electroencephalography/functional MRI by means of synchronization and use of vectorcardiogram recordings. *Journal of Magnetic Resonance Imaging*, *27*(3), 607–616.

Murakami, S., & Okada, Y. (2006). Contributions of principal neocortical neurons to magnetoencephalography and electroencephalography signals: MEG/EEG signals of neocortical neurons. *The Journal of Physiology*, *575*(3), 925–936.

Nelson, J. I., Salin, P. A., Munk, M. H.-J., Arzi, M., & Bullier, J. (1992). Spatial and temporal coherence in cortico-cortical connections: A cross-correlation study in areas 17 and 18 in the cat. *Visual Neuroscience*, *9*(1), 21–37.

Niazy, R., Beckmann, C., Iannetti, G., Brady, J., & Smith, S. (2005). Removal of FMRI environment artifacts from EEG data using optimal basis sets. *NeuroImage*, *28*(3), 720–737.

Nickerson, L. D., Smith, S. M., Öngür, D., & Beckmann, C. F. (2017). Using dual regression to investigate network shape and amplitude in functional connectivity analyses. *Frontiers in Neuroscience*, *11*.

Niedermeyer, E., & Silva, F. H. L. d. (2005). *Electroencephalography: Basic principles, clinical applications, and related fields* [Google-Books-ID: tndqYGPHQdEC]. Lippincott Williams & Wilkins.

Norcia, A. M., Appelbaum, L. G., Ales, J. M., Cottereau, B. R., & Rossion, B. (2015). The steady-state visual evoked potential in vision research: A review. *Journal of Vision*, *15*(6), 4.

O'Connor, D. H., Fukui, M. M., Pinsk, M. A., & Kastner, S. (2002). Attention modulates responses in the human lateral geniculate nucleus. *Nature Neuroscience*, *5*(11), 1203–1209.

O'Donnell, C., & Van Rossum, M. C. W. (2014). Systematic analysis of the contributions of stochastic voltage gated channels to neuronal noise. *Frontiers in Computational Neuroscience*, *8*.

Pajani, A., Kok, P., Kouider, S., & De Lange, F. P. (2015). Spontaneous activity patterns in primary visual cortex predispose to visual hallucinations. *The Journal of Neuroscience*, *35*(37), 12947–12953.

Pakenham, D. O., Quinn, A. J., Fry, A., Francis, S. T., Woolrich, M. W., Brookes, M. J., & Mullinger, K. J. (2020). Post-stimulus beta responses are modulated by task duration. *NeuroImage*, *206*, 116288.

Pasley, B. N., Inglis, B. A., & Freeman, R. D. (2007). Analysis of oxygen metabolism implies a neural origin for the negative BOLD response in human visual cortex. *NeuroImage*, *36*(2), 269–276.

Patriat, R., Molloy, E. K., Meier, T. B., Kirk, G. R., Nair, V. A., Meyerand, M. E., Prabhakaran, V., & Birn, R. M. (2013). The effect of resting condition on resting-state fMRI reliability and consistency: A comparison between resting with eyes open, closed, and fixated. *NeuroImage*, *78*, 463–473.

Pennekamp, P., Bösel, R., Mecklinger, A., & Ott, H. (1994). Differences in EEG-theta for responded and omitted targets in a sustained attention task. [Place: Germany Publisher: Hogrefe & Huber Publishers]. *Journal of Psychophysiology*, *8*(2), 131–141.

Pfurtscheller, G., & Lopes Da Silva, F. (1999). Event-related EEG/MEG synchronization and desynchronization: Basic principles. *Clinical Neurophysiology*, *110*(11), 1842–1857.

Pfurtscheller, G., Neuper, C., & Mohl, W. (1994). Event-related desynchronization (ERD) during visual processing. *International Journal of Psychophysiology*, *16*(2), 147–153.

Pfurtscheller, G., Stancák, A., & Neuper, C. (1996). Event-related synchronization (ERS) in the alpha band — an electrophysiological correlate of cortical idling: A review. *International Journal of Psychophysiology*, *24*(1), 39–46.

Pion-Tonachini, L., Kreutz-Delgado, K., & Makeig, S. (2019). ICLabel: An automated electroencephalographic independent component classifier, dataset, and website. *NeuroImage*, *198*, 181–197.

Porcaro, C., Ostwald, D., & Bagshaw, A. P. (2010). Functional source separation improves the quality of single trial visual evoked potentials recorded during concurrent EEG-fMRI. *NeuroImage*, *50*(1), 112–123.

Power, J. D., Mitra, A., Laumann, T. O., Snyder, A. Z., Schlaggar, B. L., & Petersen, S. E. (2014). Methods to detect, characterize, and remove motion artifact in resting state fMRI. *NeuroImage*, *84*, 320–341.

Pruim, R. H., Mennes, M., Van Rooij, D., Llera, A., Buitelaar, J. K., & Beckmann, C. F. (2015). ICA-AROMA: A robust ICA-based strategy for removing motion artifacts from fMRI data. *NeuroImage*, *112*, 267–277.

Quinn, A. J., van Ede, F., Brookes, M. J., Heideman, S. G., Nowak, M., Seedat, Z. A., Vidaurre, D., Zich, C., Nobre, A. C., & Woolrich, M. W. (2019). Unpacking transient event dynamics in electrophysiological power spectra. *Brain Topography*, *32*(6), 1020–1034.

Quinn, A. J., Vidaurre, D., Abeysuriya, R., Becker, R., Nobre, A. C., & Woolrich, M. W. (2018). Task-evoked dynamic network analysis through hidden markov modeling. *Frontiers in Neuroscience*, *12*, 603.

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*, *98*(2), 676–682.

Rees, G. (2009). Visual attention: The thalamus at the centre? *Current Biology*, *19*(5), R213–R214.

Reinacher, M., Becker, R., Villringer, A., & Ritter, P. (2009). Oscillatory brain states interact with late cognitive components of the somatosensory evoked potential. *Journal of Neuroscience Methods*, *183*(1), 49–56.

Reinhold, K., Resulaj, A., & Scanziani, M. (2023). Brain state-dependent modulation of thalamic visual processing by cortico-thalamic feedback. *The Journal of Neuroscience*, *43*(9), 1540–1554.

Reuter, M., Rosas, H. D., & Fischl, B. (2010). Highly accurate inverse consistent registration: A robust approach. *NeuroImage*, *53*(4), 1181–1196.

Riera, J. J., & Sumiyoshi, A. (2010). Brain oscillations: Ideal scenery to understand the neurovascular coupling. *Current Opinion in Neurology*, *23*(4), 374–381.

Rihs, T. A., Michel, C. M., & Thut, G. (2007). Mechanisms of selective inhibition in visual spatial attention are indexed by ?-band EEG synchronization. *European Journal of Neuroscience*, *25*(2), 603–610.

Rikhye, R. V., Gilra, A., & Halassa, M. M. (2018). Thalamic regulation of switching between cortical representations enables cognitive flexibility. *Nature Neuroscience*, *21*(12), 1753–1763.

Risk, B. B., Kociuba, M. C., & Rowe, D. B. (2018). Impacts of simultaneous multislice acquisition on sensitivity and specificity in fMRI. *NeuroImage*, *172*, 538–553.

Risk, B. B., Murden, R. J., Wu, J., Nebel, M. B., Venkataraman, A., Zhang, Z., & Qiu, D. (2021). Which multiband factor should you choose for your resting-state fMRI study? *NeuroImage*, *234*, 117965.

Rizzuto, D. S., Madsen, J. R., Bromfield, E. B., Schulze-Bonhage, A., & Kahana, M. J. (2006). Human neocortical oscillations exhibit theta phase differences between encoding and retrieval. *NeuroImage*, *31*(3), 1352–1358.

Romei, V., Brodbeck, V., Michel, C., Amedi, A., Pascual-Leone, A., & Thut, G. (2008). Spontaneous fluctuations in posterior -band EEG activity reflect variability in excitability of human visual areas. *Cerebral Cortex*, *18*(9), 2010–2018.

Rosa, M., Kilner, J., Blankenburg, F., Josephs, O., & Penny, W. (2010). Estimating the transfer function from neuronal activity to BOLD using simultaneous EEG-fMRI. *NeuroImage*, *49*(2), 1496–1509.

Rosazza, C., & Minati, L. (2011). Resting-state brain networks: Literature review and clinical applications. *Neurological Sciences*, *32*(5), 773–785.

Rose, D. (1979). An analysis of the variability of unit activity in the cat's visual cortex. *Experimental Brain Research*, *37*(3).

Rosenberg, M. D., Finn, E. S., Constable, R. T., & Chun, M. M. (2015). Predicting moment-to-moment attentional state. *NeuroImage*, *114*, 249–256.

Rosenberg, M. D., Scheinost, D., Greene, A. S., Avery, E. W., Kwon, Y. H., Finn, E. S., Ramani, R., Qiu, M., Constable, R. T., & Chun, M. M. (2020). Functional connectivity predicts changes in attention observed across minutes, days, and months. *Proceedings of the National Academy of Sciences*, *117*(7), 3797–3807.

Sakoğlu, Ü., Pearlson, G. D., Kiehl, K. A., Wang, Y. M., Michael, A. M., & Calhoun, V. D. (2010). A method for evaluating dynamic functional network connectivity and task-modulation: Application to schizophrenia. *Magnetic Resonance Materials in Physics, Biology and Medicine*, *23*(5), 351–366.

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.

Satterthwaite, T. D., Elliott, M. A., Gerraty, R. T., Ruparel, K., Loughead, J., Calkins, M. E., Eickhoff, S. B., Hakonarson, H., Gur, R. C., Gur, R. E., & Wolf, D. H. (2013). An improved framework for confound regression and filtering for control of motion artifact in the preprocessing of resting-state functional connectivity data. *NeuroImage*, *64*, 240–256.

Sauseng, P., Hoppe, J., Klimesch, W., Gerloff, C., & Hummel, F. C. (2007). Dissociation of sustained attention from central executive functions: Local activity and interregional connectivity in the theta range. *European Journal of Neuroscience*, *25*(2), 587–593.

Sauseng, P., Griesmayr, B., Freunberger, R., & Klimesch, W. (2010). Control mechanisms in working memory: A possible function of EEG theta oscillations. *Neuroscience & Biobehavioral Reviews*, *34*(7), 1015–1022.

Savers, B. M., Beagley, H. A., & Henshall, W. R. (1974). The mechanism of auditory evoked EEG responses. *Nature*, *247*(5441), 481–483.

Scheeringa, R., Petersson, K. M., Oostenveld, R., Norris, D. G., Hagoort, P., & Bastiaansen, M. C. (2009). Trial-by-trial coupling between EEG and BOLD identifies networks related to alpha and theta EEG power increases during working memory maintenance. *NeuroImage*, *44*(3), 1224–1238.

Schiller, P. H., Finlay, B. L., & Volman, S. F. (1976). Short-term response variability of monkey striate neurons. *Brain Research*, *105*(2), 347–349.

Schmitt, L. I., Wimmer, R. D., Nakajima, M., Happ, M., Mofakham, S., & Halassa, M. M. (2017). Thalamic amplification of cortical connectivity sustains attentional control. *Nature*, *545*(7653), 219–223.

Schölvinck, M. L., Maier, A., Ye, F. Q., Duyn, J. H., & Leopold, D. A. (2010). Neural basis of global resting-state fMRI activity. *Proceedings of the National Academy of Sciences*, *107*(22), 10238–10243.

Schölvinck, M., Friston, K., & Rees, G. (2012). The influence of spontaneous activity on stimulus processing in primary visual cortex. *NeuroImage*, *59*(3), 2700–2708.

Schubert, R., Haufe, S., Blankenburg, F., Villringer, A., & Curio, G. (2009). Now you'll feel it, now you won't: EEG rhythms predict the effectiveness of perceptual masking. *Journal of Cognitive Neuroscience*, *21*(12), 2407–2419.

Scobey, R. P., & Gabor, A. J. (1989). Orientation discrimination sensitivity of single units in cat primary visual cortex. *Experimental Brain Research*, *77*(2), 398–406.

Seedat, Z. A., Quinn, A. J., Vidaurre, D., Liuzzi, L., Gascoyne, L. E., Hunt, B. A., O'Neill, G. C., Pakenham, D. O., Mullinger, K. J., Morris, P. G., Woolrich, M. W., & Brookes, M. J. (2020). The role of transient spectral 'bursts' in functional connectivity: A magnetoencephalography study. *NeuroImage*, *209*, 116537.

Seedat, Z. A., Rier, L., Gascoyne, L. E., Cook, H., Woolrich, M. W., Quinn, A. J., Roberts, T. P. L., Furlong, P. L., Armstrong, C., St. Pier, K., Mullinger, K. J., Marsh, E. D., Brookes, M. J., & Gaetz, W. (2023). Mapping interictal activity in epilepsy using a

hidden markov model: A magnetoencephalography study. *Human Brain Mapping*, *44*(1), 66–81.

Seijdel, N., Marshall, T. R., & Drijvers, L. (2023). Rapid invisible frequency tagging (RIFT): A promising technique to study neural and cognitive processing using naturalistic paradigms. *Cerebral Cortex*, *33*(5), 1626–1629.

Shakil, S., Lee, C.-H., & Keilholz, S. D. (2016). Evaluation of sliding window correlation performance for characterizing dynamic functional connectivity and brain states. *NeuroImage*, *133*, 111–128.

Sherman, S. M., & Guillery, R. W. (1998). On the actions that one nerve cell can have on another: Distinguishing "drivers" from "modulators". *Proceedings of the National Academy of Sciences*, *95*(12), 7121–7126.

Sherman, S. M., & Guillery, R. W. (2002). The role of the thalamus in the flow of information to the cortex (P. Adams, R. W. Guillery, S. M. Sherman, & A. M. Sillito, Eds.). *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, *357*(1428), 1695–1708.

Shine, J. M. (2021). The thalamus integrates the macrosystems of the brain to facilitate complex, adaptive brain network dynamics. *Progress in Neurobiology*, *199*, 101951.

Shine, J. M., Lewis, L. D., Garrett, D. D., & Hwang, K. (2023). The impact of the human thalamus on brain-wide information processing. *Nature Reviews Neuroscience*.

Shirer, W. R., Ryali, S., Rykhlevskaia, E., Menon, V., & Greicius, M. D. (2012). Decoding subject-driven cognitive states with whole-brain connectivity patterns. *Cerebral Cortex*, *22*(1), 158–165.

Singh, K., & Fawcett, I. (2008). Transient and linearly graded deactivation of the human default-mode network by a visual detection task. *NeuroImage*, *41*(1), 100–112.

Siribunyaphat, N., & Punsawad, Y. (2022). Steady-state visual evoked potential-based brain–computer interface using a novel visual stimulus with quick response (QR) code pattern. *Sensors*, *22*(4), 1439.

Smith, A. T., Williams, A. L., & Singh, K. D. (2004). Negative BOLD in the visual cortex: Evidence against blood stealing. *Human Brain Mapping*, *21*(4), 213–220.

Smith, S. M. (2002). Fast robust automated brain extraction. *Human Brain Mapping*, *17*(3), 143–155.

Smith, S. M., & Brady, J. M. (1997). SUSAN—a new approach to low level image processing.

Smith, S. M., Fox, P. T., Miller, K. L., Glahn, D. C., Fox, P. M., Mackay, C. E., Filippini, N., Watkins, K. E., Toro, R., Laird, A. R., & Beckmann, C. F. (2009). Correspondence of the brain's functional architecture during activation and rest. *Proceedings of the National Academy of Sciences*, *106*(31), 13040–13045.

Smith, S. M., Miller, K. L., Moeller, S., Xu, J., Auerbach, E. J., Woolrich, M. W., Beckmann, C. F., Jenkinson, M., Andersson, J., Glasser, M. F., Van Essen, D. C., Feinberg, D. A., Yacoub, E. S., & Ugurbil, K. (2012). Temporally-independent functional modes of spontaneous brain activity. *Proceedings of the National Academy of Sciences*, *109*(8), 3131–3136.

Softky, W., & Koch, C. (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *The Journal of Neuroscience*, *13*(1), 334–350.

Somers, D. C., Dale, A. M., Seiffert, A. E., & Tootell, R. B. H. (1999). Functional MRI reveals spatially specific attentional modulation in human primary visual cortex. *Proceedings of the National Academy of Sciences*, *96*(4), 1663–1668.

Song, H., Shim, W. M., & Rosenberg, M. D. (2022, November 5). *Large-scale neural dynamics in a shared low-dimensional state space reflect cognitive and attentional dynamics* (preprint). Neuroscience.

Stefanovic, B., Warnking, J. M., Rylander, K. M., & Pike, G. B. (2006). The effect of global cerebral vasodilation on focal activation hemodynamics. *NeuroImage*, *30*(3), 726–734.

Stevner, A. B. A., Vidaurre, D., Cabral, J., Rapuano, K., Nielsen, S. F. V., Tagliazucchi, E., Laufs, H., Vuust, P., Deco, G., Woolrich, M. W., Van Someren, E., & Kringelbach, M. L. (2019). Discovery of key whole-brain transitions and dynamics during human wakefulness and non-REM sleep. *Nature Communications*, *10*(1), 1035.

Su, J. H., Thomas, F. T., Kasoff, W. S., Tourdias, T., Choi, E. Y., Rutt, B. K., & Saranathan, M. (2019). Thalamus optimized multi atlas segmentation (THOMAS): Fast, fully automated segmentation of thalamic nuclei from structural MRI. *NeuroImage*, *194*, 272–282.

Suffczynski, P., Kalitzin, S., Pfurtscheller, G., & Lopes Da Silva, F. (2001). Computational model of thalamo-cortical networks: Dynamical control of alpha rhythms in relation to focal attention. *International Journal of Psychophysiology*, *43*(1), 25–40.

Swettenham, J. B., Muthukumaraswamy, S. D., & Singh, K. D. (2013). BOLD responses in human primary visual cortex are insensitive to substantial changes in neural activity. *Frontiers in Human Neuroscience*, *7*.

Tadel, F., Baillet, S., Mosher, J. C., Pantazis, D., & Leahy, R. M. (2011). Brainstorm: A user-friendly application for MEG/EEG analysis. *Computational Intelligence and Neuroscience*, *2011*, 1–13.

Tagliazucchi, E., & Laufs, H. (2014). Decoding wakefulness levels from typical fMRI resting-state data reveals reliable drifts between wakefulness and sleep. *Neuron*, *82*(3), 695–708.

Takemura, H., Rokem, A., Winawer, J., Yeatman, J. D., Wandell, B. A., & Pestilli, F. (2016). A major human white matter pathway between dorsal and ventral visual cortex. *Cerebral Cortex*, *26*(5), 2205–2214.

Thomas Yeo, B. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., Roffman, J. L., Smoller, J. W., Zöllei, 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.

Thut, G., Nietzel, A., Brandt, S. A., & Pascual-Leone, A. (2006). A-band electroencephalographic activity over occipital cortex indexes visuospatial attention bias and predicts visual target detection. *The Journal of Neuroscience*, *26*(37), 9494–9502.

Tian, P., Teng, I. C., May, L. D., Kurz, R., Lu, K., Scadeng, M., Hillman, E. M. C., De Crespigny, A. J., D'Arceuil, H. E., Mandeville, J. B., Marota, J. J. A., Rosen, B. R., Liu, T. T., Boas, D. A., Buxton, R. B., Dale, A. M., & Devor, A. (2010). Cortical depth-specific microvascular dilation underlies laminar differences in blood oxygenation level-dependent functional MRI signal. *Proceedings of the National Academy of Sciences*, *107*(34), 15246–15251.

Tolhurst, D., Movshon, J., & Dean, A. (1983). The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Research*, *23*(8), 775–785.

Tsodyks, M., Kenet, T., Grinvald, A., & Arieli, A. (1999). Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science*, *286*(5446), 1943–1946.

Tustison, N. J., Avants, B. B., Cook, P. A., Yuanjie Zheng, Egan, A., Yushkevich, P. A., & Gee, J. C. (2010). N4itk: Improved n3 bias correction. *IEEE Transactions on Medical Imaging*, *29*(6), 1310–1320.

Valdes-Sosa, P. A., Sanchez-Bornot, J. M., Sotero, R. C., Iturria-Medina, Y., Aleman-Gomez, Y., Bosch-Bayard, J., Carbonell, F., & Ozaki, T. (2009). Model driven EEG/fMRI fusion of brain oscillations. *Human Brain Mapping*, *30*(9), 2701–2721.

Van Der Meer, J., Pampel, A., Van Someren, E., Ramautar, J., Van Der Werf, Y., Gomez-Herrero, G., Lepsien, J., Hellrung, L., Hinrichs, H., Möller, H., & Walter, M. (2016). "eyes open – eyes closed" EEG/fMRI data set including dedicated "carbon wire loop" motion detection channels. *Data in Brief*, *7*(37), 990–994.

Van Dijk, H., Schoffelen, J.-M., Oostenveld, R., & Jensen, O. (2008). Prestimulus oscillatory activity in the alpha band predicts visual discrimination ability. *The Journal of Neuroscience*, *28*(8), 1816–1823.

Van Ede, F., Quinn, A. J., Woolrich, M. W., & Nobre, A. C. (2018). Neural oscillations: Sustained rhythms or transient burst-events? *Trends in Neurosciences*, *41*(7), 415–417.

Van Zijl, P. C., Hua, J., & Lu, H. (2012). The BOLD post-stimulus undershoot, one of the most debated issues in fMRI. *NeuroImage*, *62*(2), 1092–1102.

Vidaurre, D., Llera, A., Smith, S., & Woolrich, M. (2021). Behavioural relevance of spontaneous, transient brain network interactions in fMRI. *NeuroImage*, *229*(37), 117713.

Vidaurre, D., Abeysuriya, R., Becker, R., Quinn, A. J., Alfaro-Almagro, F., Smith, S. M., & Woolrich, M. W. (2018). Discovering dynamic brain networks from big data in rest and task. *NeuroImage*, *180*(37), 646–656.

Vidaurre, D., Hunt, L. T., Quinn, A. J., Hunt, B. A. E., Brookes, M. J., Nobre, A. C., & Woolrich, M. W. (2018). Spontaneous cortical activity transiently organises into frequency specific phase-coupling networks. *Nature Communications*, *9*(1), 2987.

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*(37), 81–95.

Vidaurre, D., Smith, S. M., & Woolrich, M. W. (2017). Brain network dynamics are hierarchically organized in time. *Proceedings of the National Academy of Sciences*, *114*(48), 12827–12832.

Vogels, R., Spileers, W., & Orban, G. A. (1989). The response variability of striate cortical neurons in the behaving monkey. *Experimental Brain Research*, *77*(2), 432–436.

Vogels, T. P., & Abbott, L. F. (2005). Signal propagation and logic gating in networks of integrate-and-fire neurons. *The Journal of Neuroscience*, *25*(46), 10786–10795.

Vossel, S., Geng, J. J., & Fink, G. R. (2014). Dorsal and ventral attention systems: Distinct neural circuits but collaborative roles. *The Neuroscientist*, *20*(2), 150–159.

Wade, A. R., & Rowland, J. (2010). Early suppressive mechanisms and the negative blood oxygenation level-dependent response in human visual cortex. *The Journal of Neuroscience*, *30*(14), 5008–5019.

Walter, W. (1936). THE LOCATION OF CEREBRAL TUMOURS BY ELECTRO-ENCEPHALOGRAPHY. *The Lancet*, *228*(5893), 305–308.

Wang, C., Ong, J. L., Patanaik, A., Zhou, J., & Chee, M. W. L. (2016). Spontaneous eyelid closures link vigilance fluctuation with fMRI dynamic connectivity states. *Proceedings of the National Academy of Sciences*, *113*(34), 9653–9658.

Wang, Y., Li, J., Zeng, L., Wang, H., Yang, T., Shao, Y., & Weng, X. (2022). Open eyes increase neural oscillation and enhance effective brain connectivity of the default mode network: Resting-state electroencephalogram research. *Frontiers in Neuroscience*, *16*(37), 861247.

Wanmei Ou, Nummenmaa, A., Golland, P., & Hamalainen, M. (2009). Multimodal functional imaging using fMRI-informed regional EEG/MEG source estimation. *2009 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, *26*(37), 1926–1929.

Warbrick, T. (2022). Simultaneous EEG-fMRI: What have we learned and what does the future hold? *Sensors*, *22*(6), 2262.

Webster, M. J., Bachevalier, J., & Ungerleider, L. G. (1994). Connections of inferior temporal areas TEO and TE with parietal and frontal cortex in macaque monkeys. *Cerebral Cortex*, *4*(5), 470–483.

Weissman, D. H., Roberts, K. C., Visscher, K. M., & Woldorff, M. G. (2006). The neural bases of momentary lapses in attention. *Nature Neuroscience*, *9*(7), 971–978.

Weng, Y., Liu, X., Hu, H., Huang, H., Zheng, S., Chen, Q., Song, J., Cao, B., Wang, J., Wang, S., & Huang, R. (2020). Open eyes and closed eyes elicit different temporal properties of brain functional networks. *NeuroImage*, *222*(37), 117230.

Westner, B. U., Dalal, S. S., Gramfort, A., Litvak, V., Mosher, J. C., Oostenveld, R., & Schoffelen, J.-M. (2022). A unified view on beamformers for m/EEG source reconstruction. *NeuroImage*, *246*(37), 118789.

Whitsel, B. L., Schreiner, R. C., & Essick, G. K. (1977). An analysis of variability in somatosensory cortical neuron discharge. *Journal of Neurophysiology*, *40*(3), 589–607.

Wickham, H. (2016, September 13). *Ggplot2: Elegant graphics for data analysis* (2nd ed. 2016, Vol. 26). Springer International Publishing : Imprint: Springer.

Wilson, R., Thomas, A., & Mayhew, S. D. (2020). Spatially congruent negative BOLD responses to different stimuli do not summate in visual cortex. *NeuroImage*, *218*(37), 116891.

Woertz, M., Pfurtscheller, G., & Klimesch, W. (2004). Alpha power dependent light stimulation: Dynamics of event-related (de)synchronization in human electroencephalogram. *Cognitive Brain Research*, *20*(2), 256–260.

Wong, C. W., DeYoung, P. N., & Liu, T. T. (2016). Differences in the resting-state fMRI global signal amplitude between the eyes open and eyes closed states are related to changes in EEG vigilance. *NeuroImage*, *124*(37), 24–31.

Woolrich, M. W., Behrens, T. E., Beckmann, C. F., Jenkinson, M., & Smith, S. M. (2004). Multilevel linear modelling for FMRI group analysis using bayesian inference. *NeuroImage*, *21*(4), 1732–1747.

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., Ripley, B. D., Brady, M., & Smith, S. M. (2001). Temporal autocorrelation in univariate linear modeling of FMRI data. *NeuroImage*, *14*(6), 1370–1386.

Worden, M. S., Foxe, J. J., Wang, N., & Simpson, G. V. (2000). Anticipatory biasing of visuospatial attention indexed by retinotopically specific α-bank electroencephalography increases over occipital cortex. *The Journal of Neuroscience*, *20*(6), RC63–RC63.

Yan, C., Liu, D., He, Y., Zou, Q., Zhu, C., Zuo, X., Long, X., & Zang, Y. (2009). Spontaneous brain activity in the default mode network is sensitive to different resting-state conditions with limited cognitive load (A. V. García, Ed.). *PLoS ONE*, *4*(5), e5743.

Yan, W. X., Mullinger, K. J., Brookes, M. J., & Bowtell, R. (2009). Understanding gradient artefacts in simultaneous EEG/fMRI. *NeuroImage*, *46*(2), 459–471.

Yan, W. X., Mullinger, K. J., Geirsdottir, G. B., & Bowtell, R. (2009). Physical modeling of pulse artefact sources in simultaneous EEG/fMRI. *Human Brain Mapping*, *20*(6), NA–NA.

Yang, H., Long, X.-Y., Yang, Y., Yan, H., Zhu, C.-Z., Zhou, X.-P., Zang, Y.-F., & Gong, Q.-Y. (2007). Amplitude of low frequency fluctuation within visual areas revealed by resting-state functional MRI. *NeuroImage*, *36*(1), 144–152.

Yeşilyurt, B., Uğurbil, K., & Uludağ, K. (2008). Dynamics and nonlinearities of the BOLD response at very short stimulus durations. *Magnetic Resonance Imaging*, *26*(7), 853–862.

Zalesky, A., & Breakspear, M. (2015). Towards a statistical test for functional connectivity dynamics. *NeuroImage*, *114*(6), 466–470.

Zhang, D., Liang, B., Wu, X., Wang, Z., Xu, P., Chang, S., Liu, B., Liu, M., & Huang, R. (2015). Directionality of large-scale resting-state brain networks during eyes open and eyes closed conditions. *Frontiers in Human Neuroscience*, *9*(6), RC63–RC63.

Zhang, Q., Cramer, S. R., Ma, Z., Turner, K. L., Gheres, K. W., Liu, Y., Drew, P. J., & Zhang, N. (2022). Brain-wide ongoing activity is responsible for significant cross-trial BOLD variability. *Cerebral Cortex*, *32*(23), 5311–5329.

Zhang, S., & Li, C.-S. R. (2017). Functional connectivity parcellation of the human thalamus by independent component analysis. *Brain Connectivity*, *7*(9), 602–616.

Zhang, Y., Brady, M., & Smith, S. (2001). Segmentation of brain MR images through a hidden markov random field model and the expectation-maximization algorithm. *IEEE Transactions on Medical Imaging*, *20*(1), 45–57.

Zou, Q., Miao, X., Liu, D., Wang, D. J., Zhuo, Y., & Gao, J.-H. (2015). Reliability comparison of spontaneous brain activities between BOLD and CBF contrasts in eyes-open and eyes-closed resting states. *NeuroImage*, *121*(6), 91–105.