Modelling Microstructure and Function of Cortical Layers with MRI



Jiewon Kang

Wolfson College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2024

Acknowledgements

I would like to thank everyone who has supported me up until this day. I remember all the encouraging words of support that helped me endure this process.

Firstly, I would like to thank my supervisors, Prof. Saad Jbabdi and Prof. Mark Woolrich. It was a great opportunity to work with you on a challenging project, and I hope to have learned from your ways of thinking in research. Thank you for your patience with me and for understanding when I was struggling with various issues, from Covid to just me being myself, which is 'thinking too much.' I have learned a lot from both of you over these past few years.

Also, I thank my viva assessors, Dr. Rezvan Farahibozorg and Dr. Eugene Duff, for taking the time to go through the thesis and discuss the content with me.

For some reason, my parents have always been ever so supportive of me, which I can never take for granted. I thank you for the trust you always have in me that exceeds the confidence that I have. I know it is your unconditional love and support that have enabled me to become strong and independent up to this day.

I could go on and thank all my friends and colleagues in Oxford, London, and South Korea, but there are too many of you who probably won't be reading this page or thesis. Thank you to all who have waited for me to finally make up my mind and 'just finish' the thesis. I remember all your kind words of encouragement that have kept me going.

Finally, I thank my best friend and kindhearted Sejin, who has seen me struggle and also be extremely chatty about almost every single bit of my life for the past couple of years. Your words of encouragement and care for me, regardless of where I am in the world, have supported me immensely throughout my DPhil journey.

Abstract

The cerebral cortex is a thin sheet of grey matter surrounding the brain. It is particularly well developed in primates, especially in humans, and it is thought to play a key role in all cognitive functions. Cortical structure and microstructure has been studied for over a century, and modern imaging methods such as MRI have accelerated research on its function. The cortex has a well characterised layered structure, anywhere between 3 to 6 layers, and sometimes with more differentiated layers. Cortico-cortical connections are also organised in a laminar fashion, and it has been theorised that these well established anatomical facts have important consequences for brain functional organisation. Most of what we know about the cortex comes from invasive, ex vivo, histological studies. However, understanding how cortical structure and connectivity relate to cortical function requires in vivo non-invasive methods. Advances in high resolution, and usually high field, fMRI are beginning to address this need, but sub-millimetre MRI is hard to achieve, SNR-poor, and still falls short of resolving layers as the cortex can be as thin as 1 millimetre.

In this thesis, we present new non-invasive MRI-based methods to study cortical laminar structure and connectivity. We introduce a new functional MRI technique, based on inversion recovery, that can modulate the measured signal in a laminar fashion through exploiting the variation of longitudinal relaxation across cortical layers. We also introduce two new modelling approaches to infer laminar connectivity in resting state time series: an extension of structural equation modelling tailored to the new inversion-recovery acquisition method, applied to resting state data, and an extension of dynamic causal modelling to cortical layers, accounting for effects of venous blood draining and applied to task fMRI. Finally, we introduce an approach to combining ultra-high resolution whole brain microscopy with high resolution MRI data for in vivo histology.

Contents

List of Figures viii

List of Tables x

List of Abbreviations xi

1 Introduction 1

2 Background 3

2.1 Study of cytoarchitecture and myeloarchitecture . . . 3

2.1.1 Cytoarchitecture . . . 4

2.1.1.1 Classical methods . . . 4

2.1.1.2 Modern day histology . . . 6

2.1.1.3 Cytoarchitecture of neocortex . . . 8

2.1.2 Myeloarchitecture . . . 10

2.1.3 Combined studies on Cyto-/Myeloarchitecture . . . 11

2.1.3.1 Sensory cortex . . . 11

2.1.3.2 Myelination in MRI . . . 12

2.1.4 Summary . . . 14

2.2 Anatomical connectivity in layers . . . 14

2.2.1 Laminar connectivity of visual areas . . . 14

2.2.2 Limitations of tracer studies . . . 16

2.2.3 Neuronal connections and Feedforward, Feedback activity . . . 17

2.2.3.1 Canonical microcircuit theory . . . 17

2.2.3.2 Hierarchical relationships between cortical areas . . . 19

2.3 Vasculature . . . 20

2.3.1 Intracortical vessels . . . 21

2.3.2 Capillaries . . . 23

2.3.2.1 Vasculature differences and effects . . . 24

2.4 MRI for cortical layers . . . 24

2.4.1 Laminar fMRI . . . 24

2.4.1.1 High-field acquisition . . . 25

2.4.1.2 BOLD fMRI . . . 26

2.4.1.3 Vascular space occupancy (VASO) MRI 27
2.4.1.4 Laminar activity reports from previous studies . . . 28
2.4.1.5 Prospect of laminar fMRI 30
2.5 Diffusion MRI (dMRI) for Grey Matter (GM) 30
2.6 Structural Inversion Recovery-MRI method 32
2.7 Overview and contribution of research chapters 33

3 Modelling Laminar Functional Connectivity using Inversion-Recovery BOLD fMRI 37
3.1 Introduction 38
3.2 Methods 42
3.2.1 IR-BOLD fMRI 43
3.2.2 Structural equation model (SEM) with partial volume (PV) modelling 44
3.2.2.1 SEM for functional connectivity 45
3.2.2.2 SEM for laminar connectivity 46
3.2.3 Experiments: establishing face validity 48
3.2.4 Real data acquisition 49
3.3 Results 50
3.3.1 Simulations for model validation 52
3.3.2 Hierarchy recovery test 52
3.3.2.1 Simulated data 52
3.3.2.2 Hierarchy test for 3-ROI model in early visual regions - real data 53
3.3.3 Pairwise hierarchy test 54
3.3.3.1 Test for comparing 2-ROI models on 9 visual regions 56
3.3.4 Comparison of hierarchy diagrams 58
3.4 Discussion 59
3.4.1 Considerations on model architecture 60
3.4.2 Evaluation of pairwise hierarchy test 61
3.4.3 Limitations in comparing hierarchical systems 62
3.4.4 Extension of current work 63
3.5 Summary 65

4 Cortical layer microarchitecture: combining high-resolution diffusion MRI and histology 66
4.1 Introduction 67
4.2 Materials and Methods 69
4.2.1 Rationale and analysis plan 69
4.2.2 Data 70

4.2.2.1 Histology data from BigBrain 70
4.2.2.2 Diffusion measures from high-resolution dMRI . . . 71
4.2.3 Defining cortical types 72
4.2.4 Design of study 74
4.2.4.1 BigBrain proxy measures of laminar features 74
4.2.5 Analysis 1 77
4.2.5.1 Clustering BB laminar features 77
4.2.5.2 Comparing new cluster labels and cortical types from literature 78
4.2.6 Analysis 2 and 3 81
4.2.6.1 Analysis 2 – High-resolution MRI feature to cortical type prediction 81
4.2.6.2 Analysis 3 82
4.3 Results 82
4.3.1 Analysis 1: BB to Cortical types 82
4.3.1.1 K-means clustering of BB laminar feature metrics . 82
4.3.1.2 New cluster label comparison with published cortical types 85
4.3.2 Analysis 2 – High-resolution MRI to Cortical types 86
4.3.2.1 Supervised learning of MRI features to predict cortical types 86
4.3.3 Analysis 3 – BB to High resolution MRI 90
4.3.3.1 GLM fitting 90
4.4 Discussion 91
4.4.1 Analysis 1 – Clustering BB histology as cortical types 92
4.4.1.1 Observations 93
4.4.1.2 Limitations 94
4.4.1.2.1 Limitations of inspecting staining intensity profiles 95
4.4.1.2.2 Uncertainty on ground truth 96
4.4.2 Analysis 2 and 3 – Diffusion MRI measures to Cortical types 97
4.4.2.1 Observations 97
4.4.2.2 Limitations 98
4.5 Conclusion 99

5 Multilayer Dynamic Causal Modelling with Blood Draining 101
5.1 Introduction . . . 101
5.2 Methods . . . 104
5.2.1 Background on DCM for fMRI . . . 105
5.2.1.1 Neuronal activity state variables . . . 105
5.2.1.2 Comparison of DCM and SEM state equations . . 106
5.2.1.3 Haemodynamic state equations . . . 106
5.2.1.4 Extension to multiple layers . . . 108
5.2.1.5 Resting-state simulation using stochastic DCM . . 110
5.2.2 Experiments for layer DCM study . . . 111
5.2.2.1 Multi-layer DCM for a single region . . . 112
5.2.2.2 Extension of layer DCM for 2 ROIs . . . 114
5.2.2.3 Resting-state simulations using stochastic DCM . . 115
5.3 Results . . . 115
5.3.1 Simulation on multi-layer DCM for a single region . . . 116
5.3.1.1 Impact of blood draining on haemodynamic variables 116
5.3.1.2 Assessing increase of BOLD in mid-cortical depth . 119
5.3.2 Inter-laminar connectivity between 2 regions . . . 121
5.3.3 Resting-state stochastic DCM analyses . . . 124
5.4 Discussion . . . 126
5.4.1 Considerations on model architecture settings . . . 129
5.4.2 Analysis on data covariances . . . 130
5.4.3 Comparison of simulated laminar BOLD against literature . 131
5.4.4 Blood-draining laminar DCM implications for future laminar fMRI . . . 132
5.5 Future plan . . . 133
5.6 Summary . . . 134

6 General Conclusion 135
6.1 Contributions of Thesis . . . 135
6.2 Chapter specific future directions . . . 137
6.2.1 IR-BOLD fMRI and modelling approaches . . . 137
6.2.2 High-resolution diffusion to histology . . . 138
6.3 Concluding remarks . . . 138

Appendices

A Supplementary Figures 141

References 149

List of Figures

2.1 Cytoarchitecutre, myeloarchitecture from Brodmann, Vogt and Vogt 5
2.2 von Economo and Koskinas micrograph cortical types 9
2.3 Cytoarchitecture, myeloarchitecture of V1, V2, 3b, M1 13
2.4 Canonical microcircuit 18
2.5 Hierarchy diagrams 19
2.6 Intracortical vessels 22
2.7 Laminar fMRI 29
2.8 T1 distribution human cortex 33

3.1 3 types of connections - FFW, FB, lateral 39
3.2 T1s and TIs for modelling signal 44
3.3 FFW, FB, FVE model architecture 47
3.4 T1 probability distribution 49
3.5 Model inversion simulations 51
3.6 Hierarchy recovery test 53
3.7 Negative log-likelihood of 6 models fit to simulated data 54
3.8 Real data hierarchy recovery test results 55
3.9 Pairwise hierarchy test 2 candidate models 56
3.10 Pairwise hierarchy test on 9 visual regions results 57
3.11 Hierarchy indices of 9 visual regions 58
3.12 Subject, group-level comparison of FVE labels and Hierarchy indices 59

4.1 Analysis plan 70
4.2 BigBrain histology and staining intensities 71
4.3 Mean MRI metrics over all layers shown on cortical surface 72
4.4 Cortical type parcellation from literature 73
4.5 Computed 6 feature metrics on brain surface 83
4.6 K-means clustering on laminar feature metrics 84
4.7 Mean staining intensities of 7 clusters left, right 85
4.8 Mean feature metric value of clusters 86
4.9 Neighbourhood component analysis on BB staining intensities vs. feature metrics 87

4.10 Cortical type parcellation v. laminar feature based clusters 88
4.11 Supervised learning on high-resolution MRI features to predict cortical types 89
4.12 Mean MRI metric of right hemisphere cortical type labels predicted using Random Forests 90
4.13 Linear regression on diffusion features to predict BB intensity profiles 91

5.1 Single region models tested in chapter 113
5.2 2 ROI model with 3 layers for multi-layer DCM 114
5.3 CBV, dHb for different blood-draining strengths λ_d 116
5.4 BOLD % signal change for models with/without inter-laminar connections 117
5.5 BOLD % change for inputs durations of 1s, 10s 118
5.6 Time delay to reach peak BOLD response in different models 119
5.7 BOLD and CBV changes in 6 layer single ROI model with input applied to middle layer 120
5.8 Covariances of data from SEM, x of DCM, BOLD of DCM 122
5.9 Comparison of covariance matrices of 2 ROI models with input stimulus applied on different layers 123
5.10 Covariances of neuronal activity from SEM vs.Covariances of neuronal activity x from stochastic DCM 125
5.11 Covariance of BOLD signal simulated from stochastic DCM that accounts for blood-draining 126
5.12 Correlation of data covariances between neural activity simulated with stochastic DCM, BOLD simulated with stochastic DCM, and signal simulated with SEM 127

A.1 Pseudocode for simulating data with SEM 142
A.2 Paired t-test on hierarchy test results on 50 simulated data 143
A.3 Paired t-test on hierarchy test with T1 sampling on real data 144
A.4 K-means clustering on 6 laminar features using 4 cluster outputs . . 145
A.5 K-means clustering on 6 laminar features using 10 cluster outputs . 146
A.6 Original diagram for Figure 5.9 147
A.7 Original diagram for Figure 5.10 148

List of Tables

4.1 Laminar features and proxy measures for determining cortical types from staining intensity profiles 78
4.2 Cortical types adapted for current work 80

5.1 Constants for haemodynamic state equations and BOLD signal formula 111
5.2 Haemodynamic parameters related to blood-draining 112

List of Abbreviations

1D, 2D, 3D	One-, two-, three-dimensional, referring in this thesis to spatial dimensions in an image.
SNR	Signal-to-noise ratio
ROI	Region of Interest
MRI	Magnetic Resonance Imaging
fMRI	functional MRI
dMRI	diffusion MRI
BA	Brodmann Area
GM	Grey matter
WM	White matter
IR	Inversion Recovery
TI	Inversion time
T1	Longitudinal relaxation time
T2	Transverse relaxation time
PV	Partial volume
FFW	Feed-forward
FB	Feedback
BOLD	Blood-Oxygen-Level-Dependent
VASO	Vascular space occupancy
CBV	Cerebrovascular blood volume
SE	Spin-echo
GE	Gradient-echo
SEM	Structural Equation Modelling
V1	Primary visual cortex
FEF	Frontal Eye Field

BB . . . . . BigBrain

FA . . . . . Fractional Anisotropy

MD . . . . . Mean Diffusivity

NCA . . . . . Neighbourhood Components Analysis

SL . . . . . Supervised learning

SVM . . . . . Support Vector Machine

DT . . . . . Decision Trees

RF . . . . . Random Forests

DCM . . . . . Dynamic Causal Modelling

dHb . . . . . deoxy-Hemoglobin

CBF . . . . . Cererbral blood flow

FVE . . . . . Felleman and Van Essen

VEK . . . . . von Economo and Koskinas

FSL . . . . . FMRIB Software Library

BEDPOSTX . Bayesian estimation of diffusion parameters obtained using sampling techniques (X: for crossing fibres)

1 Introduction

For centuries, the quest to unravel the human brain structure and how it gives rise to emergent functions has been a major interest in neuroscience research. The cerebral cortex, the thin layer of grey matter surrounding the brain, is of particular interest. It is most developed in the human brain and is thus thought to play a key role in higher brain functions. The cerebral cortex contains dozens of billions of neurons, each connecting to hundreds to thousands of other neurons, which makes it one of the most complex systems in the known universe. Much is known about the structure of the cerebral cortex. Most famously, it is organised into distinct layers containing different types of neurons. Different areas of the cortex have similar laminar organisation but can differ in the prominence of these layers. Despite these well-documented anatomical facts from past observations, the mechanisms by which specific laminar architectures give rise to high-level functions remain elusive. Therefore, emerging approaches for studying brain functions underscore the need for an integrative method that can bridge the gap between microscale neuronal organisation and macroscale functional dynamics.

The internal structural organisation of the brain was traditionally examined by visually inspecting post-mortem tissue with the aid of microscopy and tissue staining. While these methods illuminated much on the anatomy of the brain, they fall short in elucidating how brain functions arise in the living brain. It is only in recent decades

that Magnetic Resonance Imaging (MRI) has provided unprecedented insights into human brain functions and macroscale brain connectomics. High-level areas related to uniquely human brain cognition are inherently difficult to investigate through animal experiments, so non-invasive imaging techniques applicable to humans have a wider scope of capabilities to study. Now, more recent technical advancements in MRI are enabling scientists to examine cortical layers and functional activity on a laminar scale that were conventionally impossible to achieve. There still exists a discernible gap between what can be examined using invasive methods in animals, but various modalities within MRI, such as functional MRI (fMRI) and diffusion MRI (dMRI), are offering more detailed investigations on functional and structural connectomics.

This thesis proposes new methods for studying the human cerebral cortex, focusing on its layered structure, and contemplates how laminar organisation can be related to studying how the brain functions. But before we get to that, the next chapter covers the relevant background on traditional and more modern methods for the study of cortical anatomy and function.

2
Background

Contents

2.1 Study of cytoarchitecture and myeloarchitecture . . . 3
2.1.1 Cytoarchitecture . . . 4
2.1.2 Myeloarchitecture . . . 10
2.1.3 Combined studies on Cyto-/Myeloarchitecture . . . 11
2.1.4 Summary . . . 14
2.2 Anatomical connectivity in layers . . . 14
2.2.1 Laminar connectivity of visual areas . . . 14
2.2.2 Limitations of tracer studies . . . 16
2.2.3 Neuronal connections and Feedforward, Feedback activity 17
2.3 Vasculature . . . 20
2.3.1 Intracortical vessels . . . 21
2.3.2 Capillaries . . . 23
2.4 MRI for cortical layers . . . 24
2.4.1 Laminar fMRI . . . 24
2.5 Diffusion MRI (dMRI) for Grey Matter (GM) . . . 30
2.6 Structural Inversion Recovery-MRI method . . . 32
2.7 Overview and contribution of research chapters . . . 33

2.1 Study of cytoarchitecture and myeloarchitecture

This section aims to provide an overview of classical approaches that form our fundamental understanding on human cortical layers to date. Historically, the

most straightforward method for studying the brain was by examining cellular compositions of a stained post-mortem brain tissue visually through microscopes. As result of those past observations, qualitative descriptions of cellular architecture, from now on referred to as cytoarchitecture (e.g. Brodmann's work), and myeloarchitecture (e.g. Vogt) of the human brain has been studied. Researchers would visually inspect varying cellular organisation and characteristics across the brain and define brain maps (i.e. brain parcellations) still used in the modern era. Although numerous definitions of borders between brain areas have been proposed by many neuroanatomists, these often involve subtle distinctions depending on criteria used to divide areas. In other words, there is no universally fixed nomenclature of brain areas, and researchers continue to come up with alternative parcellations to describe brain areas more effectively.

The study of cytoarchitecture and myeloarchitecture have been instrumental in identifying these 'layers' in the cortex and were previously observable only in post-mortem for humans. Nowadays, with MRI we can detect cortical layer structure and have begun to study functional activation patterns in certain layers that can allow past research on cytoarchitecture or myeloarchitecture to be extended.

2.1.1 Cytoarchitecture

2.1.1.1 Classical methods

Histology requires preparation of specimen that include dissection of tissue samples and cell staining for better visibility. Then follows a meticulous process of visually inspecting and identifying differentiable properties, for instance, composite cell type, size, cell density, and layer thickness. In animal studies, it is possible to conduct invasive electrophysiology type stimulation followed by immediate dissection and cell staining to observe neurons involved in the process [1]. However, such invasive procedures are not feasible in the human brain, therefore observational work reveal how the cortical cytoarchitecture appears in the humans. Cytoarchitecture, or in other words, cellular compositions of the cortex that have a laminar organisation was able to reveal the basic structure of the cortex. There were several

cytoarchitectural studies that resulted in brain maps that segregated brain areas as observed by changing cytoarchitectural properties in different parts of the brain. These observations initially provided anatomical ground that would later allow for functional activity of the human brain to be studied.

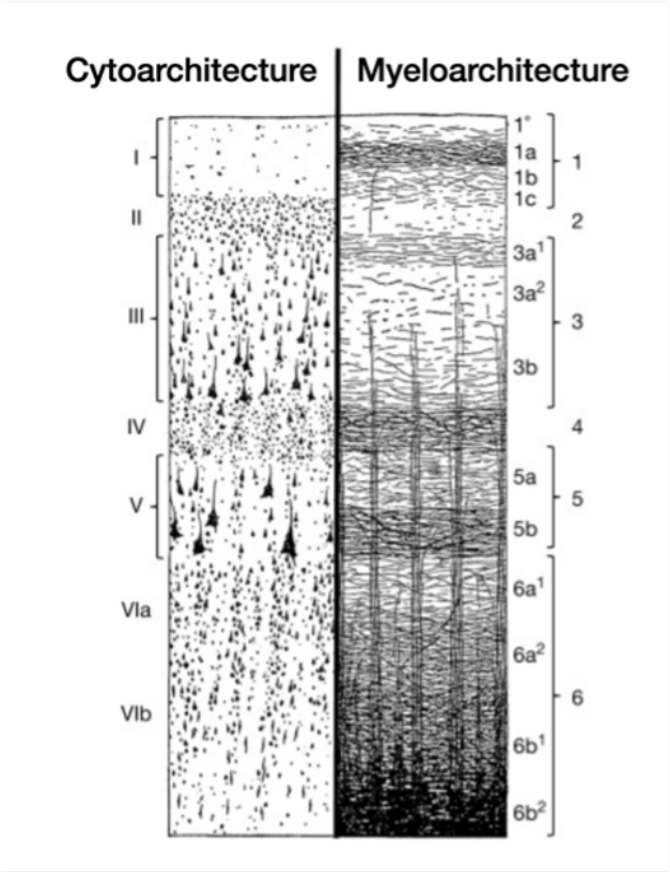


Figure 2.1: Layers in cytoarchitecture described by Brodmann (left) and myeloarchitecture as observed by Vogt and Vogt (right). The figure was modified from [2].

Brodmann's cytoarchitectural map [3] defined borders of cortical areas based on similarity of cell sizes or shapes, varying layered structure by referring to how the human cortex differed from other species [4]. As a result, Brodmann divided the human and non-human primate brain into 44 and 52 cortical areas respectively and areas were generally divided into 6 cortical layers. Brodmann's brain map for humans that was originally a 2D monograph was later adapted into a 3D brain parcellation that is often used as a reference brain map nowadays [5]. Such cytoarchitecture-based brain parcellations gained more integrity later along with

functional and structural MRI studies that attempted to map brain areas and functions. One problem when using either version of Brodmann's map to localise brain areas was that brain areas could be located erroneously as the Brodmann map did not fully distinguish gyral and sulcal surfaces [4, 6].

Another cytoarchitectural map suggested by von Economo and Koskinas Economo and Koskinas [7] includes more detailed labelled descriptions of classified areas. Unlike conventional methods, von Economo and Koskinas sliced cortical tissue into slices perpendicular to gyral surfaces than simply using axial slices. Also, whereas Brodmann simply used numbers to refer to Brodmann Areas (BA) 1 – 52, von Economo and Koskinas took photomicrographs of stained cortical sections (112 micrographic plates) and produced more informative annotations by using initials for the brain lobe (7 lobes), gyrus within the lobe, and cytoarchitectural property. Consequently, the von Economo and Koskinas (VEK) map classified the cortex into 5 types of cortical structures: 1 - agranular, 2 - frontal, 3 - parietal, 4 - polar, 5 - granulous (or koniocortex; 'konio' or 'dust-like' appearance of granular cells).

2.1.1.2 Modern day histology

Earlier on, it has been acknowledged by scientists that due to individual variability of brains, sizes of brain areas and corresponding borders are not identical [8]. As a consequence, probabilistic cytoarchitectural maps were suggested to circumvent the problem of misclassifying brain areas that can differ between individuals [9, 10]. Especially as cytoarchitecture of an individual human brain is accessible only post-mortem subject to availability, probabilistic cytoarchitectural maps [11, 12] were invaluable resources that could be compared against structural or functional MR images and act as a reference for validation.

The Juelich brain atlas was produced as part of a collective effort to reconstruct 2D whole brain histology into 3D and to generate a probabilistic cytoarchitectural map that accounts for individual variability in cytoarchitecture or boundaries of sulci, gyri, and brain regions [11]. It is important that brain areas are demarcated accurately before conducting analyses for examining how brain functions are related

to its underlying microstructural organisation. These probabilistic cytoarchitectural maps can inform us on cytoarchitecture of higher order non-sensory regions such as the less well understood association cortex.

High-resolution histology data sets such as BigBrain provide us with a detailed account of cytoarchitecture. The BigBrain dataset includes 7400 slices of dissected and stained human brain tissue (65 year old) of 20 µm thickness [13] and this is an example of a densely sampled slices from a whole brain as result of a laborious time consuming process. Such meticulously curated anatomical data exceeds resolution commonly offered by MRI therefore researchers can capitalise on BigBrain to compare against their own cytoarchitectural studies conducted on humans through MRI. Along with the data itself, there have also been efforts to develop systematic ways to analyse such compact histology for instance automated segmentation of cortical layers [14, 15] can facilitate quantitative analyses of BigBrain data.

MRI for brains can be performed in both in-vivo and ex-vivo conditions. For ex-vivo MRI, additional agents can be applied to increase image contrast and the brain can subsequently be dissected for a comparative analysis of cytoarchitecture accessed in higher resolution. Ex-vivo MRI scans are readily scanned at strong field strengths for longer durations so that images can be obtained in higher resolution and less distortion caused by fast acquisition methods [16, 17]. In addition, it has a broad application in animal studies (e.g. mice) combined with in-vivo imaging at a preclinical stage [18]. Ex-vivo imaging will also avoid distortions caused from tearing of tissue or damage during mounting into a slide [19].

Then, is in vivo MRI ready to be used as a tool to conduct non-invasive histology? In modern neuroimaging, data quality of histology and MRI has improved immensely and it is now common to analyse these multiple modalities in the same space. There are, however, many challenges in data post-processing steps to follow histological staining of brain slices or acquisition of brain images through MRI. This is even the case for reconstruction of a single brain that requires registration of consecutive 2D brain slices that have gaps in between. Then, images from different modalities must be accurately registered and corrected for so that areas and areal borders would

correspond.

To conduct these processes in a consistent manner, many have endeavoured to automate workflows for curating histology and MRI data [15, 20, 21]. More recent research also used deep learning approaches to detect cytoarchitectonic features within high-resolution histology that enabled them to distinguish areas within the visual cortex [22]. These efforts paved way for reducing the manual burden of visually inspecting the densely populated cells within the cortex thus allowing an observer-independent analysis of cytoarchitecture. As result of such advances, although currently in early stages, it is becoming possible to examine the structure-function relationships in several sensory (visual or motor) areas with respect to similarities in cytoarchitecture [20, 23, 24] and even attempted to relate with brain networks [25].

It is anticipated that access to high-resolution microstructural images, automated workflows for differentiating cytoarchitectural features that lead to better identification of brain areas will provide insights into the less well understood higher-order cortical areas that regulate complex cognitive functions such as working memory or learning. Moreover, access to powerful computational nodes have accelerated the overall process of analysing such compact histology data sets worth many Terabytes. Therefore, there is potential that we can surpass our past understanding on human cytoarchitecture established upon seminal past histological works that could have been prone to subjectivity.

2.1.1.3 Cytoarchitecture of neocortex

Our current understanding of cortical layer composition and structure is based on numerous past observations on histology including those introduced above. Neocortical areas generally have 6 conspicuous layers with exceptions in several parts of the brain such as in the allocortex. In this section, we describe general properties known of human cytoarchitecture for each cortical layer. Cytoarchitecture was explained conjunctly by cell type, size, shape, packing density or layer thickness as was done in many prior studies [26, 27]. Also, we focused on cell bodies of neurons when describing cellular organisation rather than providing a complete account

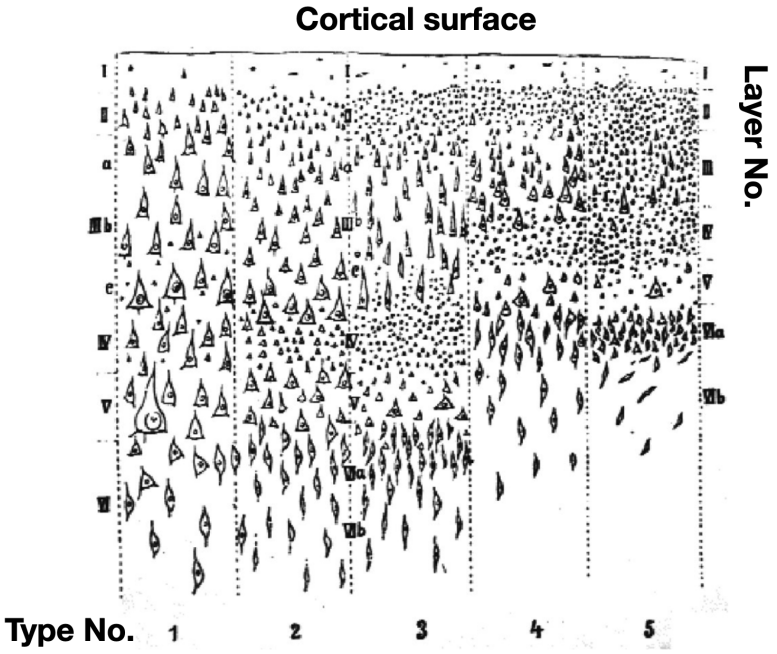


Figure 2.2: Von Economo and Koskinas divided the isocortex into cortical types I – V Economo and Koskinas [7]. Cortical types were differentiated by the overall size or density of cells composing a typical layer.

on the morphology of a neuron including initiation and termination locations. The more precise termination locations of axons, branching of axon collaterals, or input/output synapses involved in communication of neurons will be visited later in Section 2.2 focusing separately on connectivity architectures.

- **Layer I (molecular/plexiform layer)** This uppermost layer beneath the cortical surface lacks cell bodies compared to other layers as it often serves as a termination location for axons or dendrites of neurons conveying signals from the thalamus or other areas. Therefore, cellular organisation appears sparse.

- **Layer II (outer granular layer)** Smaller pyramidal or stellate cells compose layer II, and axons of these neurons are sent into deeper cortical layers. On the other hand, apical dendrites are sent upwards with respect to the cell body and terminate in layer I. Small cells are densely packed in this layer.

- **Layer III (external pyramidal layer)** Small to medium sized pyramidal cell bodies emerge while extending their apical dendrites to layer I. Many

cortico-cortical connections arise from pyramidal cells in sublayers of layer III. Cell density is lower than layer II.

- **Layer IV (inner granular layer)** Layer IV is densely populated with granular or stellate cells that are smaller than layer III. Input from the thalamus synapse with elements in layer IV such as stellate cells. Also, basal dendrites of pyramidal cells in the nearby layer III and apical dendrites of pyramidal cells of layers V-VI reside in layer IV, thus receive thalamic input.

- **Layer V (internal pyramidal layer)** Layer V is generally known for large pyramidal cells and giant pyramidal cells (Betz cells) in specific regions like the motor cortex.

- **Layer VI (spindle cell multiform layer)** Pyramidal cells, stellate cells, and spindle-like cells exist in the innermost layer. Pyramidal cells send axons towards WM to form cortico-thalamic connections, and dendrites are sent to layers above.

2.1.2 Myeloarchitecture

Cortical white matter (WM) is known for its rich myelination, but myelinated axons also exist in grey matter (GM) in lesser amounts. Pyramidal cells in cortical GM send dendrites and axons towards other layers within the same or proximate regions as short-length connections. There are also those with longer axons that pass through the GM and enter the WM to form long-range connections sent to more distant areas. Myelinated axons that are observed in the GM are most likely neurons with shorter axons, hence less myelinated axons. Here we introduce signatures of myelination observed in different parts of GM that occur as a feature of a typical cortical layer.

Tangential (horizontal) stripe patterns often observed in GM by Baillarger even earlier than when staining of tissue was used for histology Baillarger [28]. These structures were called Baillarger stripes and can be divided into a band near cytoarchitectonic layer IV (outer Baillarger stripe) and Vb (inner Baillarger stripe). In other locations such in cytoarchitectonic layer III, there are also less

common tangential Kaes stripes identified only in some sensory areas [2]. The motor cortex has an additional tangential stripe in layer 1. Those past observations revealed that strong myelination patterns were less often observed in upper layers. When referring to signatures of myelination, there are a few notable structures that can for instance be observed even with the naked eye. The Stria of Gennari is one representative example of myelination; a type of outer Baillarger stripe that is easily detectable in GM of the primary visual cortex as a thick white pattern around layer IV. However, such distinct patterns of myelination is not a shared feature of other cortical areas [29].

Varying patterns of myelination cause the GM to appear as if there are laminar borders with different myelination densities in each layer. Some used those borders of myelination to define layers with respect to differing degrees of myelination or myeloarchitecture [27]. Vogt and Vogts' work renown for observing myelination patterns was based on myelin-specific staining using Weigert staining Vogt and Vogt [30]. They focused on the organisation of radial and tangential fibres and divided the brain into 200 areas. Brain regions were defined by intensity or thickness differences of radial fibres and based on visibility of the various tangential myelinated bands [2].

Other studies also observed myeloarchitecture [31] sometimes jointly with cytoarchitecture [32] to differentiate brain areas. Many past studies still form a huge part of our current knowledge on myeloarchitecture thus still referred to in modern-day studies. For instance, several myeloarchitectonic maps from the past were revisited to adapt into a new unified myeloarchitectonic map of the cortex [33, 34].

2.1.3 Combined studies on Cyto-/Myeloarchitecture

2.1.3.1 Sensory cortex

Observations on cytoarchitecture and myeloarchitecture has since shaped our knowledge of how cortical layers in brain regions are organised as will be introduced next. A comparative study dissects properties of cortical areas according to cytoarchitecture, myeloarchitecture and even receptoarchitecture [27] and we summarise some significant characteristics for areas V1, V2, M1 and S1 in the

following segment. Those areas within the sensorimotor cortices were chosen for marked laminar properties such as certain structural elements that were easier to observe due to greater sizes. The primary sensory cortex (V1, BA17) is notable for a differentiated layer IV and can be further divided into sublayers. Sublayers IV a-c can be determined by cell density differences in which layer IVc tends to have greatest cell density and layer IVb is least dense. Unlike other areas, layer IVc can be further subdivided into sublayers α and β. In terms of myeloarchitecture, dense myelinated axons also known as the Stria of Gennari can be observed near cytoarchitectural layer IVb. Otherwise, radial fibre bundles extend up to layer III.

Next, the secondary visual cortex (V2, BA18) has large pyramidal cells in layer IIIc and larger granular cells in layer IV. V2 also has a thicker layer I compared to V1. Unlike V1, V2 lacks a Stria of Gennari and has a more prominent inner Baillarger stripe. Despite being adjacent visual cortices, such features enable V1/V2 borders to be distinguished easily. The primary motor cortex (M1, BA4) is thicker than the V1 and V2 in overall thickness. It is well known for its giant pyramidal cells also known as Betz cells and these largest cells are in layer V as layer IV is an almost lacking structure in the motor cortex. Accordingly, in some cases, the motor cortex is described as having a total of 5 layers rather than 6. The bottom layers of M1 are mostly radial myelinated fibres and horizontal fibres are rather meagre.

Finally, the primary somatosensory cortex (S1, 3b) is thinner than M1 with highest cell density in layers II – IV. A dense outer Baillarger stripe exists near layer IV and vertical fibres also extend from innermost layers up to layer III.

2.1.3.2 Myelination in MRI

The advent of MRI was an impetus for expanding ways to study microstructure of the human cerebral cortex. Patterns of myelination were observable in both histology and as different contrasts in structural MRI [35]. Therefore, some prominent structures such as the Stria of Gennari could be observed and compared between the two modalities [29]. Studies have also demonstrated that myelination can be quantified using indirect measures derived from MRI, hence, it is now possible

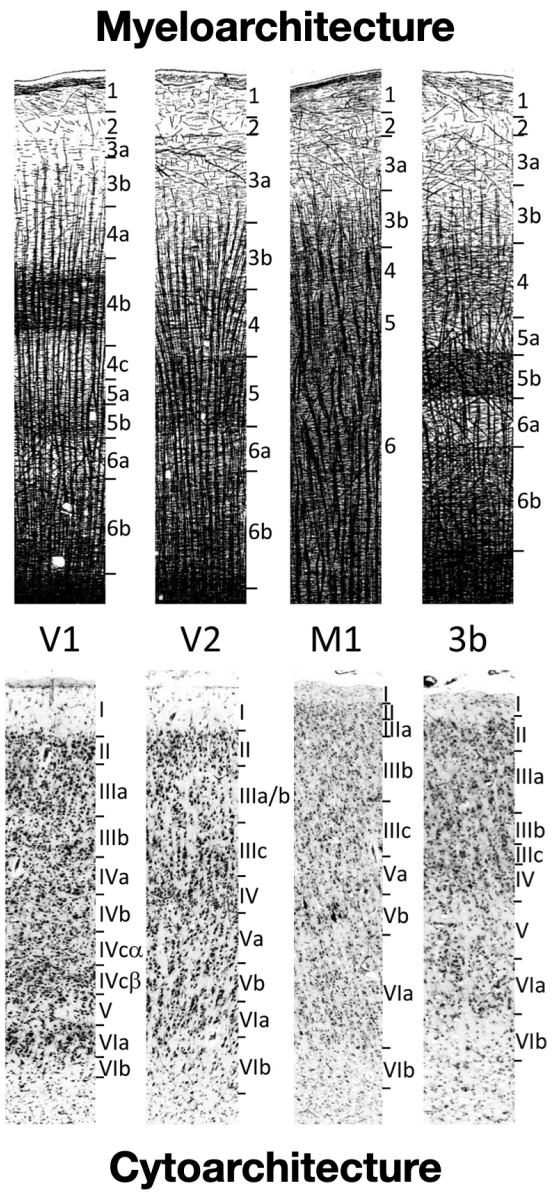


Figure 2.3: Darker horizontal bands can be observed in V1, V2 and 3b, whereas M1 lacks marked horizontal fibre bands. Cytoarchitectural differences can be determined in terms of density and cell size changes between layers. For instance, M1 has largest cells in its layer V. However, M1 has a thin almost indistinguishable layer IV. The above illustrations were originally from [30] and we adapted the version from [27].

to study myelination of brain areas beyond V1. Specifically, contrasts that are correlated with amounts of myelination include quantitative T1 relaxation mapping [36] or semi-quantitative approaches combining T1 and T2 weighted imaging [37]. More recently, diffusion MRI has been used in conjunction with neurite density quantification to reveal correlation with myelination [38]. The search for measures

that can serve as a tool to conduct in-vivo microscopy in MRI such that it can level with anatomical detail offered from classical histology is an ongoing area of research [39].

2.1.4 Summary

To summarise, classical studies focusing on microscopy are not complete nor obsolete to remain as historical methods from past research. Rather, continuous efforts are being put to quantify structural elements of the cortex more precisely for instance with MRI. It is important that reliable and reproducible interpretations on microscopic histological images are produced while avoiding observer-dependent bias. Ongoing research takes advantage of high-resolution MRI and joint analyses on high-resolution post-mortem histology data that are emerging. Yet the question of how cortical microstructure relates to certain complex brain functions remains to be answered.

2.2 Anatomical connectivity in layers

2.2.1 Laminar connectivity of visual areas

We now move on to consider connections that arise between neurons. A single neuron generally spans multiple layers and depending on where dendrites synapse with other neurons, different connections are evoked. Thus, a neuron can be described simply by the location a connection initiates and finally terminates in [40]. Specifically, a connection initiates in a layer that holds the cell body (soma) and terminates in layers that axon terminals finally reach. Such reduced annotations of neuronal connections offer a simplified way of describing complex neuronal wiring.

Such fundamental knowledge on how neurons connect were initially studied through injection of anatomical tracers which indicated initiation and termination locations of neurons. Decades of tracer studies uncovered how neurons were connected to each other, hence formed the basis of our knowledge on laminar connectivity between cortical areas [41]. Yet a complete map of neuronal connections with precise initiation and termination patterns is not available even in the current

era. Tracer studies still face shortcomings as exact termination locations of axons are unclear. This makes it difficult to identify neuronal pathways involved in longer-range connections. In addition, functional roles related to many of these anatomical connections have not been studied exhaustively. Especially neurons that terminate in multiple locations are still poorly understood as we do not know whether there are separate roles involved in each pathway. Furthermore, these tracer-based assessments for investigating laminar connectivity are not possible to conduct in humans for its invasive nature.

Laminar connections between areas in the sensory system such as the visual system, although have been studied more extensively as it is easier to manipulate experimental stimuli that can activate sensory areas. In the following section, we first introduce how areas within the visual system are connected to each other with respect to its composite layers to provide an example of how laminar organisation relates to functions involved in the visual information processing pathway.

Visual input that enters the retina and passes through the optic tract is sent to the Lateral Geniculate Nucleus (LGN) of the thalamus. This input that reached the thalamus is then sent to visual areas in order of ascending hierarchy starting from the primary visual cortex (V1) towards higher order visual areas (V2, V3, V4, V5/MT). From a laminar perspective, input is sent from the LGN to layer IV of V1. Then the signal is internally relayed from layer IV to superficial layers (layers II/III) within V1 before leaving towards another area. The signal from layer II/III is sent to layer IV of a separate visual area and such activity is referred to as feedforward (FFW) activity. FFW activity usually targets approximately middle layers such as deep layer III or layer IV. Meanwhile, layer IV of the visual area receiving FFW input from V1 can coincide with another input sent directly from the LGN to the same layer (layer IV). Input signals that enter layer IV either way can be relayed to Layer VI to be sent back to the thalamus. Input also propagates further up the hierarchy and reaches the middle layer of other higher-order visual areas. Within the same brain area, intrinsic connections relay signals to other layers in the same region while feedback (FB) activity propagates from higher-order visual areas to

lower-order areas down the cortical hierarchy. FB connections from layers V/VI of higher order areas are sent down the visual hierarchy towards supragranular (superficial/upper) layers (layer II/III) or infragranular (inferior/lower) layers (layer V/VI) avoiding middle layers [40, 42]. FB connections are not strictly contingent to this pattern and are more numerous and variable, hence it has been difficult to probe all existing kinds of connections [43]. Adjacent or closer regions are connected by FB pathways that generally conform to a typical FB connection pattern avoiding the middle layer but does not always target a specific layer in a fixed manner. Also, it has been considered that layer VI often sends FB connections towards other regions (cortico-cortical connections) whilst layer V is more relevant to sending FB to the geniculate nucleus (cortico-thalamic connections) [40]. Without specifying excitatory or inhibitory roles of neurons, visual input from lower order to higher order areas drives FFW activity while FB modulates activity in a more diffuse way [44, 45].

2.2.2 Limitations of tracer studies

Not all sources and termination locations of neurons connecting layers of visual regions involved in visual information processing are clearly revealed in terms of its anatomical substrates. Such difficulty in elucidating how neurons connect to each other are attributable to the complex morphology of neurons. For instance, FFW input towards layer IV can be received by basal dendrites of pyramidal cells in deep layer III or apical dendrites of pyramidal cells whose cell bodies are in lower layers [46]. Each of the neurons receiving layer IV input were neurons that send of FFW and FB activity respectively, therefore it could become a difficult task to disentangle the actual pathway induced by a task input. Even for separate cortical areas, differences between FB connections emanating from deeper layers V/VI is still an ongoing area of research [41, 47].

Modern tracers lack the ability to explain those complex connections that could change our understandings such as the role of collateral branches with respect to a FB connection. To address such problems, we would require tracer experiments

that exploit multiple retrograde tracers and this way termination locations for more than two regions can be probed simultaneously.

2.2.3 Neuronal connections and Feedforward, Feedback activity

2.2.3.1 Canonical microcircuit theory

The canonical microcircuit theory demonstrates FFW and FB connections that exist between neurons from a perspective of a cortical column. The right diagram of Figure 2.4 briefly demonstrates a general model of FFW and FB connections. In forward propagating directions (Figure 2.4 right, blue arrows) superficial neurons (layer 2/3) project to layer 4 of higher order areas. In other words, an excitatory process initiates from this extrinsic input sent to layer 4. Continuing the FFW process in the higher order area, layer 4 neurons relay the input to layer 2/3 and then from layer 2/3 to layer 5/6 as intrinsic connections. FB processes can also spur from layer 5/6 neurons to supragranular or infragranular layers of a lower order region. Critically, these connections are not confined within a single area but propagate towards other cortical areas through those FFW and FB connections. There are also intrinsic inhibitory connections such as inhibitory layer 5/6 neurons connecting to layer 4 neurons as part of a FB process and layer 4 inhibitory neurons sending connections to layer 2/3 neurons as an inhibitory FFW process. The extent in which such 'canonical' models can be applied to areas other than the visual cortex or sensory cortices is still an active area of research.

According to this concept, such laminar connection patterns could underly a hierarchy between cortical areas thus can help us understand how information propagates between cortical areas. This theory can be tested on higher-order cortical areas beyond sensory cortices to

investigate whether more complex cognitive processes also adhere to similar laminar connectivity principles. Alternatively, the canonical microcircuit theory can help elucidate the less well understood relationships between distant regions

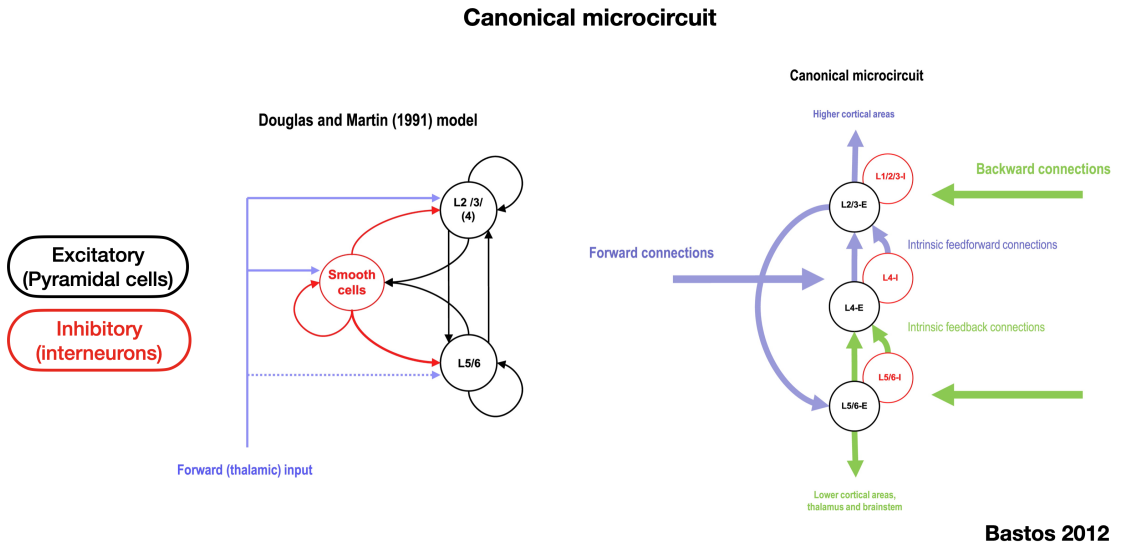


Figure 2.4: On the left is a basic representation of the canonical microcircuit taken from [48]. This diagram is a simplified adaptation of the classical canonical microcircuit diagram from [49] and shows thalamic inputs that target pyramidal cells in superficial layers (layers 2/3/4) or granular layers (layers 5/6). Input to granular layers marked with broken lines imply weaker connections. Here, not only excitatory pyramidal cells, but also inhibitory (smooth) cells were shown as they also receive thalamic input and are interconnected with excitatory neurons. Next, the diagram on the right side demonstrates in detail FFW and FB pathways related to inhibitory and excitatory activity of neurons. A canonical microcircuit diagram often focuses on how neurons within a cortical column connect to each other, but this view can be extended beyond a cortical column within a single area. Thus, FFW and FB connections projecting towards other lower or higher order cortical areas can be added on the canonical microcircuit representation from [49], accordingly.

connected by long-range extrinsic connections. Some claim that these canonical cortical microcircuits are prevalent across the entire brain [50].

It is often the case that the canonical microcircuit theory seems complex due to different roles of neurons (inhibitory or excitatory) or various cell types (pyramidal cell, spiny stellate cell, interneuron and others) but in this thesis we kept laminar connectivity models simple rather than further differentiating neurons by their excitatory or inhibitory role. It is more intuitive to apply the canonical microcircuit theory in understanding laminar connections in animal brains that are capable of cellular resolution, invasive electrophysiology as well as injection of tracers. However, along with higher resolution MRI studies that have started to account for layers, this concept also must be tested in the human brain. With current

MRI systems, it is impossible to measure human brain activity on a neuronal microcircuit level resolution. It has continuously been stressed that there is a need for modelling approaches that account for neuronal micro-circuitry underlying larger-scale connectivity [46, 51].

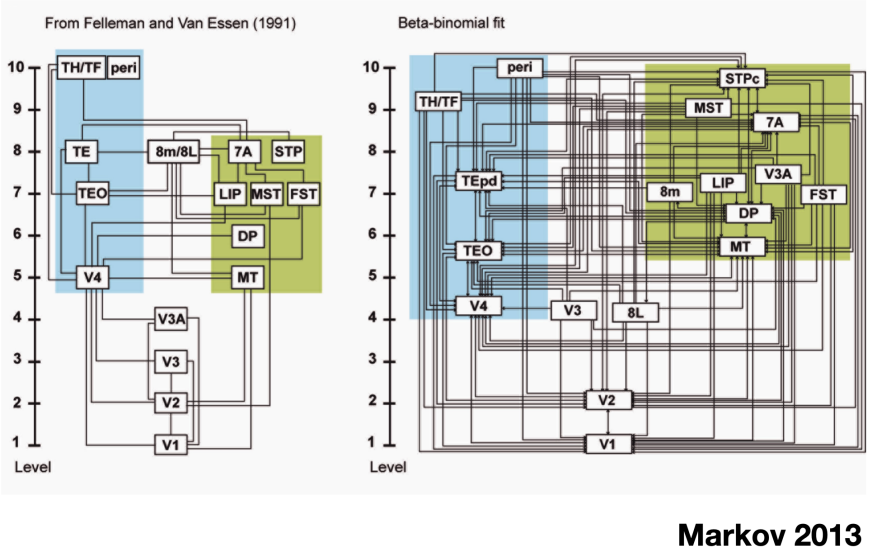


Figure 2.5: The original hierarchy diagram from Felleman and Van Essen [52] (left) is shown alongside an updated hierarchy diagram that showed study results from [43] (right). The earlier hierarchy diagram from [52] was established by conducting a meta-analysis on connections identified in multiple published tracer studies, while the latter hierarchical model was based on counts of superficial layer neurons detecting with modern tracers. Both representations in this figure were adapted from [43] and mark visual areas involved in the ventral stream in blue shades and those in the dorsal stream in green. As such, models of hierarchy are not yet complete and are still being updated depending on connection patterns determined.

2.2.3.2 Hierarchical relationships between cortical areas

When attempting to understand how information is processed across different brain regions, we regularly study connectivity patterns that possibly have a hierarchical relationship between each other. One seminal study from Felleman and Van Essen attempted to establish a functional hierarchy between cortical areas based on their laminar connections Felleman and Van Essen [52]. They explored this concept by a meta-analysis compiling dozens of tracer studies on macaques and introduced how FFW connections arise from upper layers and terminate in Layer IV (approximately

middle layer) of the other region. Conversely, FB connections initiate from cells from an inferior (bottom) layer and terminates in superior or inferior layers of the other area. In other words, FB connections would terminate by avoiding Layer IV. There were also lateral connections that terminated approximately in all layers within a region and these connections are more difficult to quantify due to the numerosity of neurons in the human brain. These patterns of FFW and FB connections underly the concept of hierarchical processing of regions as FFW activity occurs from hierarchically lower to higher regions and FB activity in the opposite direction. Felleman and Van Essen classified observed laminar connections collated from their tracer database and ultimately suggested an account of hierarchy between cortical areas as shown in Figure 2.5. There have been several updated models of hierarchy since and one alternative model is shown as the right diagram from [43] (Figure 2.5). Markov and colleagues' modern tracer study focused on quantifying percentages of supragranular and infragranular projecting neurons and they suggest alternative hierarchical relationships in some regions. For instance, they subdivided the frontal eye fields into 8L and 8m and unlike Felleman and Van Essens' model of hierarchy they were placed in different hierarchical levels.

2.3 Vasculature

So far, our discussion of the laminar organisation of the cortex has focused on neurons and their connections. However, the cortex also contains a complex vasculature that is essential for sustaining all the cells. Vascular organisation is particularly important in this thesis as blood vessels are the major source of contrast in fMRI which is the main technique used here.

The cortical vascular system has a crucial role of transporting blood and providing oxygen to the brain parenchyma (brain tissues composing cerebral cortex). When neural activation occurs, the increased demand for oxygen is supplied by an influx of blood at the cortical site of activation. This hemodynamic reaction was referred to as neurovascular coupling. Blood enters through the cortical surface and passes through arteries to reach capillaries. Then Oxygen held in Hemoglobin (Hb) is released near

capillaries therefore venous blood leaving the capillaries is in a deoxygenated state. Deoxygenated blood drains out of the cortex through ascending veins then larger pial veins in the cortical surface. It is this deoxygenated blood (dHb) that is the main source of fMRI contrast (so called BOLD contrast mechanism).

In this section, we introduce basic morphology of cortical vasculature that spans the cortical depth overlapping with cytoarchitectonic cortical layer boundaries. The hemodynamic response measured in fMRI is a delayed signal following actual neural activation, hence a peak response is usually reached with a delay of 4 to 6 seconds. fMRI often faces scrutiny that temporal resolution is poor compared to the time scale of neural activity so research in MRI often focuses on developing faster acquisition methods as well as improving spatial resolution to resolve cortical layers. Furthermore, recent studies also analyse the different sections of the hemodynamic response function to understand causes and relevant substructures related to the initial dip below baseline or post-stimulus undershoot [53, 54]. With high-resolution MRI, the study of laminar activity has become imminent so it has become ever so important to assess the process in which vascular effects arise as they do not match the actual neuronal activation both spatially and temporally.

The cortical vasculature can be divided into sub-compartments by their location with respect to the cortical depth.

2.3.1 Intracortical vessels

Intracortical blood vessels indicate the vertically oriented arteries (descending arteries) and veins (ascending veins). Blood entering and exiting the cortex are in contact with intracortical vessels. These intracortical vessels can also be further subdivided by their lengths, which is approximately the distance between the cortical surface and depth that a blood-vessel finally terminates to branch into smaller vessels. The following structural details were observed by Duvernoy and colleagues in their study on 25 post-mortem human brains that had undergone intravascular injection [55]. Intracortical vessels typically have a main trunk segment that extends perpendicularly from the cortical surface into the cortical depth. In

Duvernoy's study, intracortical vessels were divided into 6 artery groups and 5 vein groups differentiated by their lengths or by the layer in which they terminate before merging into the capillary bed (Figure 2.6). Group 1 arteries had trunks terminating in upper layers whilst arteries in groups 2 – 4 each terminated in layer III, layers IV-V, and layer VI respectively. The thickest and longest arteries in group 5 or 6 extended from pial vessels down to the WM. Arterial groups 2 - 4 were the most frequently observed arteries and typically had a candelabra-like structure (Figure 2.6 A2, A3, A4) unlike veins in counterpart groups (vein groups 2 – 4). Arteries of group 5 had a similar branching structure with groups 2 - 4 but passed both GM and WM. Finally, group 6 arteries headed straight to the WM without any branches extending into any of the cortical layers in GM. Most intracortical vessels in groups 5 and 6 that reached into deeper parts of the cortex had larger vessel diameters than upper layer vessels such as groups 1 – 2.

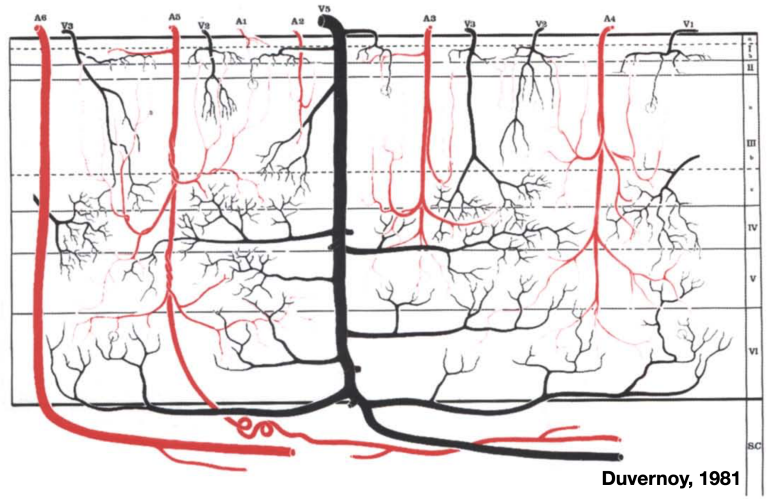


Figure 2.6: Arterial groups 1 – 6 and vein groups 1 – 5 extend through the cortical depth as sketched in cortical tissue from [55]. Vessels in different groups have varying penetration depths and thickness. One example vessel V5 is referred as having a notable conical structure that becomes wider (extends horizontally) in deeper layers.

Likewise with arterial groups, veins were divided into groups originating from layer I, layer II/III, layer IV, layer VI and WM each (vein groups 1 – 5). Typically, veins in group 4 had long horizontal branches extending from the main trunk and branches were longer if they extended from a deeper point in the cortex.

Group 5 veins had thicker trunks and branches stemming from the WM and is especially known for its conical structure (Figure 2.6 V5). On the other hand, veins in groups 1 – 3 had relatively smaller diameters. Duvernoy and colleagues observed that the most frequently identified veins and arteries were those in groups 2 and 3. Auxiliary vascular borders can be formed based on main trunk lengths of intracortical vessels but these do not necessarily coincide with the previous cytoarchitecturally defined layers.

2.3.2 Capillaries

While intracortical vessels have a role of transporting arterial or venous blood in and out of the cortical parenchyma, oxygen and glucose are released in capillaries from oxygen-rich inflowing blood. Descending arteries branch into smaller arterioles and then merge into the capillary bed. Capillaries are spread across the cortex excluding the cortical surface in form of a wide mesh (50 µm width), thus overlap with many vessels in the cortex. Among all vascular compartments, capillaries have the smallest vessel diameter size around 4 µm [56]. It is generally considered that the more proximal a measured signal is to microvasculature, the closer it is to actual neuronal activity [57]. However, microvasculature is an umbrella term that encompasses arterioles, capillaries, and venules that all have smaller diameters than intracortical penetrating vessels.

It is still difficult to delineate capillary specific changes using blood oxygen level dependent (BOLD) fMRI measures that do not offer sufficient spatial resolution. In the macaque visual cortex, layer IV had highest vascular density [56]. Even in a numerical simulation study focusing on humans, the number of arterial branches feeding into capillary beds were highest around layer IV [58]. As consequence, higher density of neurons and blood supply near layer IV potentially contribute to increased demand of oxygen following neuronal activity. In non-BOLD fMRI measures, cerebrovascular blood volume (CBV) weighted acquisition is often preferred for its greater specificity to the middle layer and this is possible due to the method's higher sensitivity to arterioles than venous vessels. However, these issues still need

to be further validated to deduce how vascular or neuronal elements affect metabolic responses in that layer when neural activity is induced.

Venous blood leaves capillaries through thin venules before entering vertically oriented ascending veins. Finally, blood drained through ascending veins is transported back into the pial surface and exits the cortex through pial veins.

2.3.2.1 Vasculature differences and effects

Vasculature can vary depending on a brain region of interest and even between species. For instance, the primate striate cortex is rich in vasculature especially layer IV of primary visual cortices [56]. Also, differences in vascular density between primary and non-primary cortical areas could affect different hemodynamic responses in laminar fMRI although such hypotheses should be further examined [59].
Inherent differences in vascular organisation between species could be one of the reasons why animal and human fMRI studies often show conflicting results. Ascending veins were more numerous than descending arteries in rodents [60] whereas in all primates (humans and monkeys) artery to vein ratio was the opposite from rodents [56].

2.4 MRI for cortical layers

2.4.1 Laminar fMRI

We previously overviewed methods for studying cortical layers such as post-mortem histology that reveal the structural organisation and tracer studies that show layer-specific connections between cortical areas. Most conventional methods excluded the possibility of non-invasive in-vivo studies of the living brain but this is no longer true after a decade of rapid development in MRI hardware, acquisition and analysis methods. We are currently at the stage in which human cortical layers can be relatively resolved by virtue of human MRI being pushed to stronger field strengths. Accordingly, current research aims to understand connectivity relationships between brain regions on a laminar level as would be done invasively using anatomical tracers in animal studies. Although it is not possible to achieve the same neuron-level

resolution, it is becoming possible to assess laminar connectivity by leveraging high field MRI. It has become common to use high-field MRI to probe laminar-specific brain activity in small patches of visual [61] or sensorimotor areas [62]. Yet the challenge remains for assessing of laminar connectivity in higher-order cognitive areas beyond the sensory cortex or between multiple regions and more towards the whole brain [63, 64]. The following section covers some methods that have been used in the recent field of laminar fMRI. To note, there is not yet a consensus on a winning MRI method of widespread use to study human cortical layers as methods have trade-offs in sensitivity and specificity and these will be discussed further.

2.4.1.1 High-field acquisition

Higher magnetic field strengths at 7T now has more widespread use in human brain imaging thus allowing greater detail of cortical layers to be revealed [65]. Even in lower magnetic field strengths at 3T, significant landmarks such as Stria of Gennari was easy to detect near Layer IV of V1 [66]. Such distinct features could be used as reference points that help differentiate cortical layers without the need for all 6 layers to be clearly disambiguated [67]. Now higher SNR can be achieved with high-field MRI [68] so laminar detail and cortical thickness differences in various regions beyond the visual cortex can be detected [62, 69, 70].

However, high-field MRI was made possible by development of advanced hardware, complex acquisition sequences and pipelines. Pipelines for processing MRI data acquired in ultra-high field must cope with artefacts that emerge due to Radio Frequency (RF) inhomogeneity, magnetic field susceptibility and thermal noise amplifies in greater degrees compared to low-field MRI [71]. Not only are there such technical considerations but due to trade-offs in sensitivity and specificity, it is challenging to acquire good quality images in optimal conditions.

With sufficient statistical power it may eventually become possible to exploit high-field MRI to assess the individual brain on a laminar scale in the future. In human neuroimaging, many studies analyse large data sets when searching for clinically viable biomarkers that can be globally applicable. Such big data

approaches have proven powerful but when it comes to observing mesoscale activity with respect to layers in larger populations careful action should be taken. A recent review by [72] discusses avenues that ultra-high resolution MRI studies should overcome to observe laminar activity of individual brains with improved sensitivity. Other researchers have commented on how individual variability can be coped with in laminar-specific fMRI. Current studies have addressed problems related to individual variability through repeated sampling of an individual brain but the argument on how much of individual-specific observations should be accounted for insofar as group-level results hold is very much an open area of discussion [63].

2.4.1.2 BOLD fMRI

Next, we introduce common fMRI methods used to study laminar activity. In laminar fMRI, Gradient Echo (GE)-BOLD acquisition has been favoured over other acquisition methods for its greater sensitivity. A GE-BOLD sequence uses a T2* weighted contrast and is proportional to deoxy-Hemoglobin (dHb) concentration as paramagnetic dHb affects T2* of spins [73]. However, its specificity or ability to detect true laminar activity has often been questioned as GE-BOLD is sensitive to signals in large blood-draining vessels such as ascending veins and pial veins in the cortical surface [74]. An alternative Spin Echo (SE)-BOLD acquisition is relatively unaffected by pial and ascending draining veins, thus, results in laminar signals more specific to microvasculature than in GE-BOLD. SE-BOLD contrasts are achieved when a 180° radiofrequency pulse refocuses spins resulting in a T2 rather than T2* decay. The refocusing pulse reduces extravascular BOLD signals that are greater around larger vessels, hence, SE-BOLD signals result in being more specific to microvasculature. However, the correlation of SE-BOLD signals with the origin of neuronal activity is somewhat scrutinised, as these BOLD signals are an indirect, delayed response of neural activity. Additionally, as haemodynamic responses spread towards nearby areas away from the true site of neural activation the fMRI signal effectively smears out [75]. One sequence design suggested to mitigate such limitations of GE and SE-BOLD was a combined 3D GRASE

sequence that improved sensitivity and specificity in detecting signals pertaining to microvasculature even in high-field scans [76].

When measuring laminar activity with BOLD fMRI, it is often the case that voxels including upper cortical layers have amplified signal affected by blood draining through larger veins. It is difficult to prevent such phenomenon from occurring especially in GE-BOLD that is more susceptible to signal amplification in upper layers. However, initial onset HRFs were similar for SE and GE-BOLD [77] and signal amplification occurring towards the cortical surface could be additionally corrected for. For instance, there were several modelling approaches that eliminated unwanted signal amplification through regression or separate modelling of laminar compartments across the cortical depth [53, 78].

2.4.1.3 Vascular space occupancy (VASO) MRI

VASO MRI is a non-BOLD measure that is sensitive to CBV and has gained popularity in the laminar fMRI community for its good balance in sensitivity and specificity. For this reason, many fMRI studies incorporated VASO imaging to detect laminar activity in regions in the visual, motor, somatosensory cortices even to observe double peaks across the cortical depth (Figure 2.7). VASO MRI exploits the fact that T1 properties differ for blood and brain tissue so tissue contrast can be retained when nulling blood signals using an inversion-recovery acquisition [79]. Following an inversion pulse, CBV increases in arterioles and capillaries causing the MR signal to drop [80]. It should also be noted that VASO is not completely free of signals from intracortical vessels such as pial and descending arteries although is more specific to microvasculature compared to GE-BOLD acquisition.

One shortcoming of VASO is its low sensitivity which in turn requires longer acquisition sessions. However, in more recent studies, there are acquisitions combining VASO imaging with perfusion contrasts to make it possible to detect laminar activity for the whole brain rather than from thin slices of brain areas [81].

As was the case for other BOLD acquisition methods, to understand the vascular nature of VASO signals, different experimental stimuli with various duration lengths

should be tested and validated for further applications in observing laminar activity. For instance, a recent study inspected how signals change in different cortical depths when applying stimuli with varying durations. Here, with a shorter stimulus (up to 4 seconds), VASO signal had larger percentage signal change in superficial layers whereas with a longer stimuli (12, 24 s) signal changes were greater in middle layers. This study suggested that longer stimulus durations could be building up gradual changes in capillaries [82].

2.4.1.4 Laminar activity reports from previous studies

There tends to be a consensus in the pattern of neural activity expected to be observed in cortical layers involved in neuronal pathways especially in the visual system. That is, through previous animal experiments involving anatomical tracers or non-invasive electrophysiology methods for humans that both demonstrated layer-specific neural activity. However, depending on the acquisition sequence used or experimental conditions, inconsistent results on laminar activation were often reported in MRI studies.

In human fMRI, there were mixed reports on layer IV activation that were detected in some studies [67, 85] but also were not identifiable in slightly different experimental conditions [74, 76] (Figure 2.7A). These inconsistent observations on middle layer activation could be related to specific experimental set-up such as the choice of duration or repetition frequency of stimuli. It became increasingly emphasised to carefully curate experiments when it came to laminar fMRI studies aiming to selectively activate layers [86, 87],.

Some studies later successfully delineated separate laminar activity that could potentially differentiate FFW and FB activity using specific tasks. Using a visual illusion experiment, Kok and colleagues were able to observe FB activity towards V1 activating deeper layers Kok et al. [83] (Figure 2.7B). In another study, there was also evidence of internally generated processes being evoked without FFW inputs. FB activity was observed as targeting superficial layers in V1 [84] and reports on the laminar activation in these two studies summarised in Figure 2.7B.

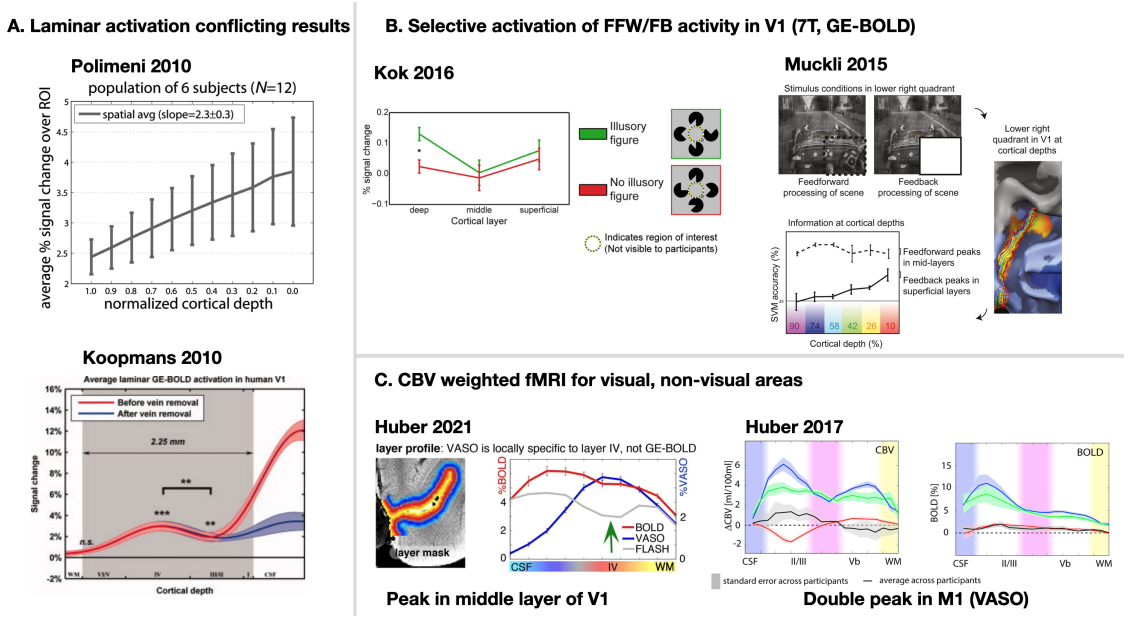


Figure 2.7: Here we collate several reports of laminar activity in published laminar fMRI studies. (A) In an earlier laminar fMRI study that used GE-BOLD acquisition in 7T, [74] observed that signals closer to pial vessels had greater signal readout, but also larger error (top graph). The graph below from another study is shown as an example that demonstrated middle layer activation that was unseen in many other GE-BOLD studies such as [74]. Here, laminar activation was observed also in layer IV of V1 [67] even after removing voxels contaminated by surface veins (bottom graph). (B) Researchers endeavoured to push laminar fMRI paradigms further through experimental designs involving visual illusion tasks or visual occlusion experiments. As consequence, some studies were able to selectively activate different layers of the visual cortex as shown in these reported depth-dependent activation profiles. For instance, [83] identified increased deeper layer activity during visual illusion intending to induce FB activity. Using a visual occlusion experiment, [84] observed peaks in V1 middle layers when a complete image was presented, whereas superficial layers were most strongly activated when participants were shown an image missing a quadrant (this would correspond to inducing FB activity). (C) These final activation profiles are results from non-BOLD fMRI acquisitions that were compared against BOLD activation profiles. In V1 using VASO acquisition, middle layer peaks were detected with improved specificity (left) and in a resting-state CBV-weighted acquisition, signatures of separate activity from input and output of M1 were identified in upper and lower layers, respectively [69, 81].

Another recent integrative fMRI-EEG approach was able to validate layer-specific activity observed in fMRI BOLD studies [88]. In their study, BOLD signal in upper layers correlated with gamma band EEG recordings whereas lower layers correlated negatively with beta band activity therefore suggesting that non-invasive methods could potentially be detecting distinct processes occurring on a laminar level.

2.4.1.5 Prospect of laminar fMRI

One branch of connectivity models that have potential for microcircuit-level models to be tested is by using Dynamic Causal Modelling (DCM) that estimates effective connectivity between nodes (introduced in detail in Chapter 4). With DCM models, we can study relationships between populations of neurons on the scale of inhibitory and excitatory neurons within a cortical column of a same region as between-neuron connections. Data acquired with EEG or magnetoencephalography (MEG) have superior temporal resolution than fMRI so by treating inhibitory and excitatory neurons as separate nodes, DCM models can be used to build models of effective connectivity by combining with fMRI data [48]. Moreover, DCM style modelling can be useful to independently model vascular and neuronal components [53, 89, 90]. Previously, DCM has been used to validate hierarchical relationships between visual regions V1 – V4 [91] and suggested that hierarchical relationships between regions could be assessed with respect to laminar connectivity using a DCM approach. To move towards more accurate meso-scale non-invasive imaging, it may be timely and useful to develop models that assume specific crosstalk between neurons that pertain to excitatory or inhibitory activity with the aid of DCM models.

2.5 Diffusion MRI (dMRI) for Grey Matter (GM)

dMRI is a widely used neuroimaging modality that provides information on tissue microstructure and connectivity. dMRI measures diffusion of water molecules which can be restricted or hindered by cell membranes. Anisotropic diffusion occurs when tissue is structured by the presence of neuronal pathways and this can be used to infer the directions of axons. Fibre orientations can be estimated and reconstructed from diffusion data and inform on structural connectivity of the brain.

Most dMRI studies of the cortex have looked at two aspects: anisotropy and fibre orientation. With improved acquisition sequences allowing higher resolution, anisotropy measures such as fractional anisotropy (FA) and mean diffusivity (MD) changes across the cortical depth became observable in GM [92, 93]. Factors that

usually contribute to diffusion anisotropy were the amount of myelination in axons [38] or structural arrangement of dendrites [94].

WM is formed of more myelin rich fibre bundles therefore are highly anisotropic. Naturally, researchers were more interested in reconstructing and quantifying fibre orientations in the more anisotropic WM. However, dMRI approaches have recently been extended to observe GM structure as well. In GM, radially oriented fibres perpendicular to the cortical surface were the dominant fibre orientations often detected by dMRI [95, 96]. More complex fibre orientations involving crossing fibres can occasionally be identified in GM such as in the Stria of Gennari located in layer IV of V1 [97]. Those crossing fibres occur by radial fibres coinciding with tangential fibres [98].

Recent studies that computed fibre orientations for GM observed differences between some regions and even within a region with different layers. Typically, the primary motor cortex (M1) and somatosensory cortex (S1) with distinct cytoarchitectural features are areas with well-known fibre orientations. One study observed radially oriented fibres in M1 and with a sharp change in its adjacent S1, tangential fibre orientations were observed [93]. Some emerging methods such as neurite orientation dispersion and density imaging (NODDI) have improved estimation of multiple tissue specific fibre orientation distributions so that researchers can reconstruct fibres that reach into the GM beyond WM [99]. Also, with high angular resolution diffusion MRI on post-mortem visual cortex, the less commonly observed tangential fibre elements were visible in particular layers, especially markedly along the cortical surface [97]. dMRI acquisition and modelling efforts are ongoing therefore now it is possible for subtle diffusivity differences to be distinguished and GM fibre orientations to be reconstructed even in areas beyond M1 and S1 [100].

One prolonging problem involved in interpreting fibre directions from tractography is that the exact point of termination of a tract is not exact. In GM, this problem persists for both fibres parallel to the cortical curvature and radial fibres that terminate in a layer in GM [24]. It has been shown that origins of WM fibres

located in the GM could be reached given sufficient SNR [99]. There is possibility that research in dMRI focusing on meso-scale organisation of cortical layers can provide insight into ambiguous neuronal pathways that hold information on certain less well-known connections between cortical and subcortical areas.

2.6 Structural Inversion Recovery-MRI method

Finally, in this section, we introduce a different approach within the breadth of MRI that has been used to study cortical layers. It was suggested by [101] that instead of endlessly increasing spatial resolution in MRI to physically disambiguate cortical layers, inversion recovery (IR) MRI can exploit varying T1 relaxation properties across the cortical depth to separate cortical layers (Figure 2.8). Several former studies have used IR-MRI, even in lower-resolution (e.g. at 3T), to differentiate cortical layers whilst achieving reasonably spread-out T1 distributions across the cortex [102, 103].

In Chapter 3, we develop a method based on this IR acquisition concept, but with two additional elements. First, we combine IR with a BOLD acquisition to achieve an IR-sensitised BOLD measure. Then, we incorporate an additional connectivity modelling element in order to examine layer-specific connectivity relationships between regions.The acquisition requires an IR pulse (180° pulse) that is applied prior to a SE pulse sequence (90° pulse). The time between the inversion pulse and the excitation pulse for SE is the inversion time (TI), and by measuring signals at different TIs a T1-related contrast can be obtained. In a nutshell, our new approach; an IR-BOLD acquisition accompanied by mathematical modelling of layer-specific connections aims to conduct layer-fMRI using the more commonplace lower-resolution MRI. Later, in Chapter 3 and 5 we cover details of mathematical models used to model connectivity and test whether our new methodology is feasible.

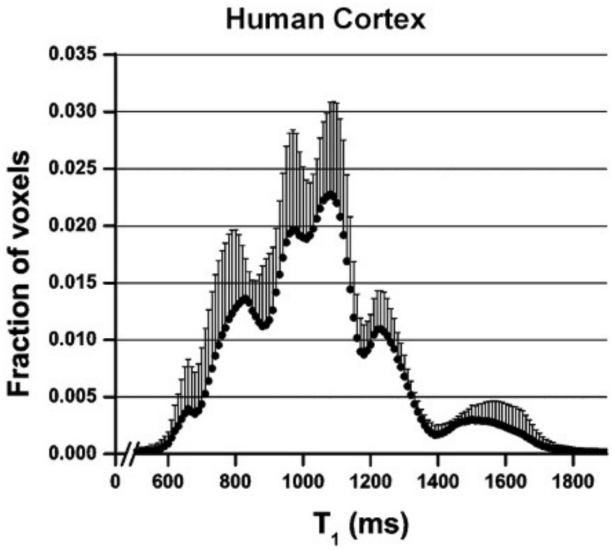


Figure 2.8: In [101], it was proposed that T1 differences across the cortical depth be used rather than resolving cortical layers by improving spatial resolution. The above figure from the aforementioned study shows a distribution of T1 relaxation times across the cortical depth in 6 subject data acquired with an inversion-recovery sequence.

2.7 Overview and contribution of research chapters

The remaining chapters study cortical layers based on the above background literature in 3 perspectives. Firstly, in Chapter 3 we tested a novel IR-BOLD paradigm that is accompanied by fitting models to identify layer-specific connectivity relationships. We test capabilities of our novel laminar fMRI method (IR-BOLD accompanied by SEM modelling of connectivity) in assessing laminar connections underlying hierarchical relationships between brain regions. An outline of goals and research gaps to be tackled is as follows.

- **IR-BOLD** is suggested as the modality to conduct laminar fMRI. High-resolution fMRI still faces limitations in resolving cortical layers, therefore, we seek a way to perform laminar fMRI that is feasible in lower, more standard resolutions widespread in research and clinical facilities.

- Even though higher-resolution MRI is now available to observe cortical layers in spatial detail which was not once possible, assessing laminar functional

connectivity is still challenging. The problem is that measured signals are still unspecific as they smear out across layers. This causes many common BOLD acquisition sequences difficult from accurately observing specific layer activation, hence, making it difficult to assess connectivity on the laminar level. We propose that with an additional mathematical modelling technique along withour IR-BOLD measures, we can model layer-specific connections between brain areas. **SEM** has conventionally been used to estimate effective connectivity, and we adapt this model for the purpose of assessing layer-specific connections.

- After testing models that assume typical connection patterns we can further assess hierarchical relationships between regions. To date, determining hierarchical relationships between brain regions in the humans was not easy to perform as ways to observe laminar-level connections (closer to neuronal-level) was limited. In this work, we examine hierarchical relationships in human visual areas and attempt to replicate the macaque brain hierarchical organisation diagram from Felleman and Van Essen's classical study.

This chapter is currently in preparation for release as a preprint.

Next, in Chapter 4 we move on to explore structural aspects of cortical layers.

- In studying human cortical layers a major challenge is to non-invasively access cortical layers in separable detail (cortical thickness being 2-3 mm). Histology is of invasive nature, but microscopic approaches in much higher resolution offer observation of cellular detail in layers. However, dissection of a brain into thin slices and registration to study the brain as a whole entity entail laborious curation processes. On the other hand, non-invasive imaging is promising, but resolution is compromised in comparison to resolution referred to in histological practises.

- This chapter aims to explore the extent in which 'in-vivo histology' can be done when using two extremely high-resolution (especially, at the time of writing) data sets.

- We used **dMRI** and a high-resolution histology data (**BigBrain**) registered on the same space to examine whether shared features or cortical layer properties can be observed in the two different modalities. We systematically examined what cortical microstructural features can be observed or captured from these high-resolution data sets.

Finally, in Chapter 5 we conduct a simulation study using a layered DCM models to examine the effects of including blood-draining parameters when measuring laminar connectivity.

- Prior to this study, SEMs were used to model connectivity. In this chapter, we attempt to fit DCM models to the same IR-BOLD data. A key distinction of laminar DCMs is that blood-draining effects are modelled, and this was not examined in the previous SEM study. Blood-draining is likely to interfere with BOLD signals measured in a laminar-level, therefore, we examine this alternate modelling technique directly against SEMs.

- Also, with data simulated from multi-layer DCMs we try to understand how haemodynamic parameters, activity, and connectivity change when there are multiple layers and when blood-draining strengths are modulated. Our understanding on how blood-draining affects signals measured on a laminar-level is still lacking. This simulation was a crucial step that was explored in this chapter before fitting multi-layer DCM to real laminar data.

Observations outlined in this chapter were based on preliminary simulation results and we plan to release the toolbox (Toolbox API contributed by Professor Saad Jbabdi). We can simulate laminar BOLD data using the layer DCM tool to further investigate whether haemodynamics interfere in other published or measured layer fMRI data that try to cope with blood-draining effects.

3

Modelling Laminar Functional Connectivity using Inversion-Recovery BOLD fMRI

Contents

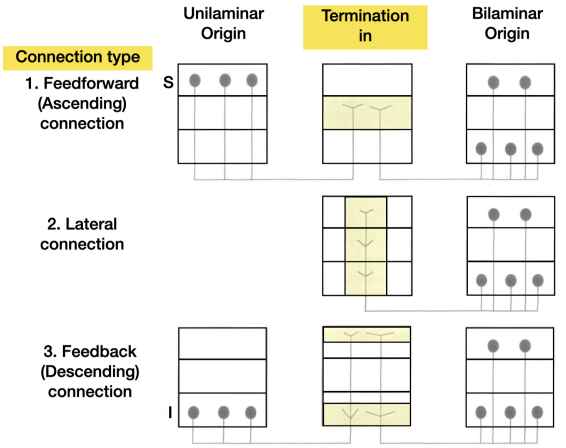
3.1 Introduction 38
3.2 Methods 42
3.2.1 IR-BOLD fMRI 43
3.2.2 Structural equation model (SEM) with partial volume (PV) modelling 44
3.2.3 Experiments: establishing face validity 48
3.2.4 Real data acquisition 49
3.3 Results 50
3.3.1 Simulations for model validation 52
3.3.2 Hierarchy recovery test 52
3.3.3 Pairwise hierarchy test 54
3.3.4 Comparison of hierarchy diagrams 58
3.4 Discussion 59
3.4.1 Considerations on model architecture 60
3.4.2 Evaluation of pairwise hierarchy test 61
3.4.3 Limitations in comparing hierarchical systems 62
3.4.4 Extension of current work 63
3.5 Summary 65

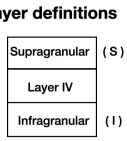
3.1 Introduction

The cerebral cortex, the main subject matter of interest in this thesis, refers to the 2–3 mm thick grey matter (GM) crust surrounding the brain. Cortical GM consists of 5–6 layers, and these layers can be differentiated through their cellular composition. While connections within and between cortical layers, as well as between brain areas, can appear complex, certain regularities in the patterns of connections have been identified. We do not, however, fully understand how these patterns of neuronal connections confer distinct functional roles to cortical areas. This is the main motivation behind the development of techniques for measuring layer-specific connectivity.

Conventionally, tracer studies in animals helped unravel how neurons are connected to each other by highlighting where a tracer terminates with respect to the location of the cell body (soma). In their seminal 1991 paper, Felleman and Van Essen surveyed dozens of tracer studies in macaque monkeys to establish a general model for the connectional organisation of the brain's visual areas. Their survey involved both anterograde and retrograde tracer studies, which means they could analyse both the source (soma) and terminations (synapses) of axonal connections. In consequence, they differentiated two types of connectional patterns: cortical layers involved in feedforward (FFW) connectivity, which mostly emanate from superficial (supragranular) layers and terminate in mid-cortical (granular) layers, and feedback (FB) connections, which originate from inferior (infragranular) layers and terminate in non-granular layers (avoiding the middle layer) [42]. By classifying all the visual connections into one of these two categories (alongside a third category that they termed "lateral connections"), they were able to organise the entire visual cortex into a hierarchy (Figure 3.1).

In their model, visual input from the retina enters the thalamus; it then propagates towards layer IV of the primary visual cortex (V1). After relaying this input in V1 towards superficial layers, it can propagate to layer IV of higher-order visual regions such as V2, V3, V4, and V5 (also known as MT). On the





Connection type	Origin		Termination in
	Unilaminar	Bilaminar	
1. Feedforward (FFW)	S	S, I	Layer IV
2. Lateral		S, I	Across all layers (columnar)
3. Feedback (FB)	I	S, I	Avoiding Layer IV

Figure 3.1: The diagram on the left (adapted from [52], Figure 3) shows 3 basic types of connections: FFW, FB, and lateral connections that are distinguished by layers that neurons initiate from and terminate in. In table on the bottom right, we describe the same aforementioned connection patterns from the left diagram into words.

other hand, FB signals from higher-order visual regions are sent to superficial or lower layers in lower-order visual regions.

While laminar organisation of the visual cortex has been extensively studied, it is still unclear to what extent these models apply to the rest of the brain. It is particularly unclear whether they apply to high-order cortical areas and in cortical functions that are unique to humans. It has been difficult to assess laminar-level connections in the human brain as we are restricted to using non-invasive means of imaging.

Anatomical connections in human brains have been studied, albeit indirectly, using diffusion MRI (dMRI) tractography. Tractography has provided us rich insight into structural connectivity of the brain as it indirectly maps axons by following the diffusion of water. It is still difficult to determine more precise origin and termination locations of tracts, especially for those terminating inside the GM [24] and between distant areas connected through long-range connections [105]. Even though tractography attempts observation of structures in macro- to micro-scale resolution, it lacks directionality information involved in an axonal pathway [104]. On the other hand, with fMRI it is possible to study directionality of pathways when connected areas activate togther. Moreover, laminar activity or inter-areal

laminar connectivity is becoming more feasible to measure with layer-specific fMRI. Laminar fMRI, given the right resolution and cleanness of data, can potentially help us investigate how information flows during neuronal crosstalk.

Laminar connectivity in humans can now be measured using fMRI, although, it requires much higher spatial resolution in the sub-millimetre range. High-resolution MRI became possible with recent advances in high-field MRI (e.g., 7T). As signal-to-noise ratio (SNR) in MRI scales with voxel volume, high field is required for imaging at high spatial resolution with sufficient SNR. Current state-of-the-art techniques enable voxel sizes in the sub-millimetre range (0.5–0.8 mm) [81]. Despite attaining such unprecedented laminar detail, laminar fMRI struggles to clearly resolve cortical layers in all regions of the brain [106].

While high-field scanners are relatively rare and mostly absent in the Global South, their use in research is becoming more widespread. However, when imaging cortical layers, high spatial resolution is not enough on its own. One main issue is that vasculature in the cortical GM interferes with neuronal activity, which is the actual signal we aim to study. fMRI signals are produced by changes in blood oxygenation, mainly in the veins, and blood travelling up and down the cortex effectively smears the fMRI signal across the layers. This, in turn, counteracts the merits of achieving high spatial resolution.

To avoid measuring signals pertaining to vascular effects irrelevant to neural activity, layer fMRI methods tend to avoid gradient echo (GE) BOLD, the dominant fMRI acquisition method, as it is the most sensitive to large draining veins. Other approaches are favoured, such as spin echo (SE) fMRI, which has been claimed to be more sensitive to smaller blood vessels that tend to feed relevant tissue compartments. Another method, VASO (Vascular Space Occupancy), which was developed decades ago to measure blood volume, is now regaining popularity and being adopted in high-resolution laminar fMRI. VASO uses a difference method between signals with and without blood nulling (using inversion recovery) to isolate the contribution of blood from that of tissue.

In summary, laminar fMRI can be used to investigate human cortical layers beyond the visual cortex and potentially into higher-order cortices. In general, this would have to be done using high-resolution MRI, but it is challenging to perform while aiming the entire brain. High-field acquisition is not always feasible due to availability of advanced hardware, and most studies only image a partial section of the brain (e.g., a few slices around the visual cortex). Also, even in high-resolution some inherent problems within fMRI regarding signal specificity to neuronal activation persist or could be more detrimental in inferring layer-specific connectivity.

Here, we introduce a new and entirely different approach to measuring laminar connectivity in fMRI. The idea is as follows: rather than increasing spatial resolution, we instead propose to measure at conventional spatial resolution but aiming at the whole brain. This can be done with a. Inversion-Recovery (IR) BOLD acquisition that is combined with b. modelling of laminar connections using Structural Equation Modelling (SEM). Firstly, using an IR acquisition we can sensitise the MRI signal in a layer-dependent manner. Layer-dependency comes from the fact that longitudinal relaxation (T1) varies across the cortical depth [101]. For a T1 value that roughly corresponds to a layer or certain laminar depth within the cortex, IR acquisition is conducted at multiple inversion times (TIs). Hence, following such acquisition the MRI signal is modulated by weighting different partial volume contributions at each TI. The terminology partial volume contributions is to be distinguished from the commonly encountered partial voluming of different tissue types such as CSF, GM, and WM as we refer to the different fractions of a signal that can be accessed in different layers. In sum, we can build a signal with amplitude variations that are T1-dependent, thus layer-dependent.

Next, to assess laminar connectivity relationships between cortical areas we used SEM extended for multiple layers. Our SEM for multiple layers makes use of IR-BOLD data acquired in multiple TIs, hence, neuronal signal is weighted by a partial volume model that weighs the signal at a T1 and TI. Moreover, along with

assumptions on a noise model we can easily assess fit of model using covariance structure of data and the model. It is convenient to implement SEMs to explore laminar connectivity as we assume a certain connectivity structure. We exploit its versatility for comparing models and test candidate inter-laminar connectivity patterns that underlie hierarchically organised areas across the entire brain.
Using this novel acquisition and connectivity modelling approach, we attempt to perform laminar fMRI in the more general MRI setting. We investigate whether laminar connectivity can be studied without having to push hardware to its limits to obtain extremely high-resolution data. High-field acquisition is often restricted to slice or slab-level coverage of the brain, therefore, whole-brain level fMRI was a major hurdle. Currently, however, laminar connectivity analyses for the whole brain has become available in several VASO studies where signals were controlled to be less affected by blood physiological effects [81]. Using the two innovations associated with our method, we explore the extent that our new method can be used to assess multiple brain areas to deduce whether our method is be feasible for whole-brain connectivity analyses. Since, the method is not confined to high-field acquisition, we investigate capabilities of this technique in this chapter.

Specifically, we first validate the model using simulated data to check for the face validity of the model and inference approach. Then, we apply the approach in real data to estimate FFW and FB connections in the visual cortex, culminating in an attempt to partially reproduce the Felleman and Van Essen hierarchical model of the visual system.

3.2 Methods

Two new ideas are introduced in this chapter. The first idea (Section 3.2.1) is to add an inversion pulse to a standard fMRI SE sequence, and this concept is referred as IR-BOLD fMRI acquisition. By acquiring time series data in this way at multiple TIs, the signal is amplitude-modulated by a factor that depends on the T1 relaxation time. Using the fact that T1 varies across the cortical depth, we use T1 values to represent different cortical layers within the GM. The second idea (Section

3.2.2) is to model this type of data using a modified structural equation model (SEM) that tries to infer effective connectivity between different layers of the cortex.

3.2.1 IR-BOLD fMRI

Figure 3.2a shows the distribution of T1 relaxation values in the human cortex from a previous study that used the IR method to distinguish separate layers [101]. While T1 relaxation is a measure sensitive to differing amounts of myelination, the study focused on using T1 variability across the cortical depth to achieve cortical layer separation within GM in opposition to spatially delineating layers. Some subsequent studies supported the claim that a reasonable spread of T1s in the GM exists in both humans and rats when scanning at both low and high resolutions [102, 103]. In our work, we based the method for data acquisition on this T1 dependent IR scheme, but this time, we move on to extend it to BOLD imaging. Especially, we explore whether we can observe laminar connectivity relationships of interest in this low-resolution MRI setting and for areas across the entire brain.

In Figure 3.2a, T1 ranges from ∼600ms near the white matter border to ∼1600ms at the pial surface. This wide range of about 1 second across the cortex allows us to modulate the signal differently in distinct parts of the cortex using varying inversion times as demonstrated in the example in Figure 3.2b.

We can model the BOLD signal $y(t)$ as a function of time as:

$$y(t) = M_0 \cdot S_{\text{hemo}}(T2) \cdot S_{\text{tissue}}(TI, T1) \tag{3.1}$$

where M_0 is the total magnetisation (a constant), $S_{\text{hemo}}(T2)$ is a modulation that is a function of neuronal activity that affects our signal via haemodynamics (T2-dependent for SE, T2*-dependent for GE BOLD acquisition), and $S_{\text{tissue}}(TI, T1)$ is a modulation due to the additional IR module added prior to the spin-echo module. Assuming that we have reconstructed the magnitude of the signal, the signal can be visualised as Figure 3.2b and this term has the form:

$$S_{\text{tissue}}(TI, T1) = |1 - 2e^{-TI/T1}| \tag{3.2}$$

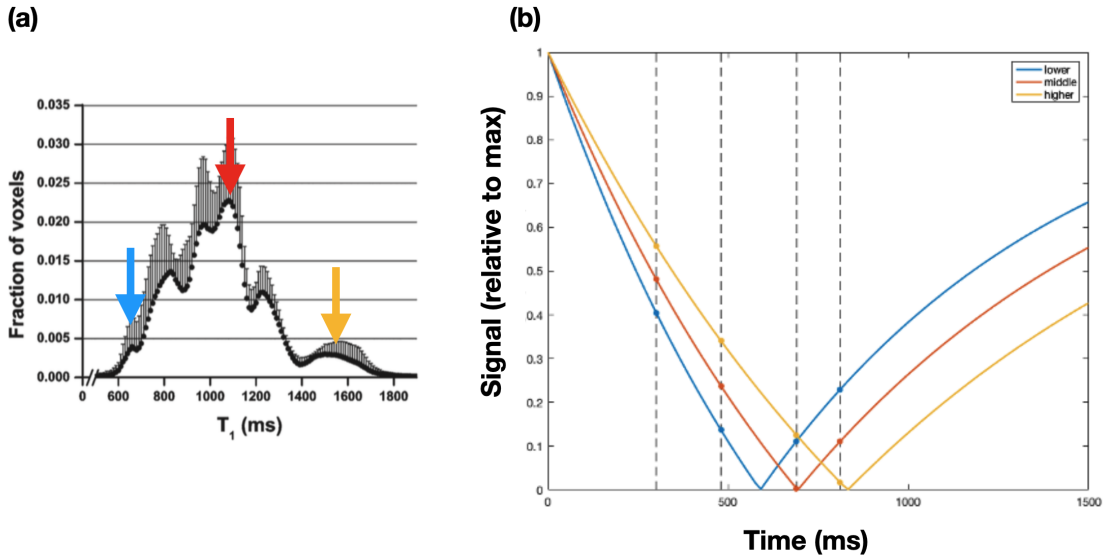


Figure 3.2: (a) Distribution of T1 relaxation values across cortical depth [101]. 3 different T1 values indicated with coloured arrows representative of lower T1 (blue), middle T1 (red), and higher T1 (yellow) ranges. (b) We simulate IR signals that feed into our partial volume model. The vertical dashed lines indicate the inversion times at 300, 480, 690, and 810ms, and the y-axis shows how signals (i.e. computed partial volumes) vary across the acquisitions. For example, upper layers closer to the cortical surface (higher T1, yellow) dominate the signal at low TI but are almost nulled at higher TI, and vice versa for the lower layers closer to WM surface (lower T1, blue).

3.2.2 Structural equation model (SEM) with partial volume (PV) modelling

Structural Equation Models (SEM) are amongst the first types of approaches that have been proposed to study functional (or "effective") connectivity in fMRI [107]. In SEM, a backbone connectivity structure is assumed between a small set of regions, and the connectivity between these regions is estimated from the data. Importantly, the model must explicitly state which connections are absent, i.e., set as null connections. The remaining non-zero elements of the connectivity matrix are estimated to best match the covariance matrix of the time-series data. In following sections, we first the original SEM used to model functional connectivity. Then, we move on to our modified SEM for modelling laminar connectivity. In the latter, we adopt the same formalism but adjust it to account for the fact that our observations are subject to different amounts of partial volume contributions

from distinct cortical layers. That can be done by coupling the signal with partial volume weights as will be formally explained in the following section.

3.2.2.1 SEM for functional connectivity

In the standard SEM formalism, data is modelled through a "noise propagation" equation whereby an uncorrelated white noise variable z is fed through a network with connectivity matrix A leading to observed data vector x. The data at network node i can be written as:

$$x_i = \sum_{j=1}^{N} A_{ij} x_j + z_i \tag{3.3}$$

where each node depends on the other nodes through the matrix A and receives its own noise perturbation (z_i). In matrix form, this can be written as:

$$x = Ax + z \tag{3.4}$$

or

$$x = (I - A)^{-1} z \tag{3.5}$$

which shows more explicitly the idea of noise propagating through the network.

To invert this model, notice that the covariance matrix implied by the above model is

$$\Sigma = \mathrm{cov}(x) = (I - A)^{-T} \mathrm{cov}(z) (I - A)^{-1} \tag{3.6}$$

that is,

$$\mathrm{cov}(x) = \sigma^2 (I - A)^{-T} (I - A)^{-1} \tag{3.7}$$

where σ is the noise standard deviation. By comparing this formula to the observed covariance matrix (data covariance), we can infer the non-zero elements of A alongside the noise standard deviation σ.

3.2.2.2 SEM for laminar connectivity

We extend the above formalism by assuming that x represents not only activity within cortical areas but also within the layers of each area. Thus, at a single time point t, $x(t)$ is no longer a $N \times 1$ vector (where N is the number of areas) but a $(N \times L) \times 1$ vector, L being the number of layers for each area.

We assume, as above, that the activity in x is governed by (what we could call a state equation):

$$x = Ax + z \tag{3.8}$$

The A matrix has a dimension $(N \times L) \times (N \times L)$.

Unlike in the SEM model, we cannot directly observe neural activity in layers x. Instead, we measure another vector $y(t)$ which collapses the signal across the layers of any given region. This gives rise to the observation equation:

$$y = Px \tag{3.9}$$

where P is a $N \times (N \times L)$ matrix encoding the partial volume contributions of each layer to the observed signal. We have already seen the functional form of these contributions in Eq. 3.2, but with a focus on its multi-layer versatility, our partial volume model is:

$$P_{ij} = |1 - 2e^{-TI/T1_{ij}}| \tag{3.10}$$

where TI is the inversion time and $T1_{ij}$ is the longitudinal relaxation time constant in region i and layer j. By acquiring data at multiple inversion times, the model can be inverted to recover the layer-specific non-zero connectivity coefficients in A and the noise variance σ^2 (assuming perfect knowledge of the T1 relaxation times).

As above, model inversion can be achieved by measuring the closeness of the model-predicted covariance ($\sum$) and the empirical data covariance (S) and we use negative log-likelihood (nLL) as the objective function to minimise. The model

covariance ($\sum$) accounting for the addition of the observation equation through the partial volume matrix P can be derived as:

$$\Sigma = \text{cov}(y) = \text{cov}(Px) = P\text{cov}(x)P^T = P(I - A)^{-1}\text{cov}(z)(I - A)^{-T}P^T \quad (3.11)$$

We used Metropolis Hastings sampling to fit the model and the following nLL that assumes a Wishart covariance model as an objective function to:

$$\text{nLL} = -\log(|\Sigma|) - \text{trace}(S\Sigma^{-1}) + \log(|S|) + p \quad (3.12)$$

Here ($\sum$) is the implied covariance of the model, S is the observed covariance or experimental data covariance ($S = \text{cov}(y) = E[y^T y]$), and p is the number of free parameters. In practice, the model is augmented with broad Gaussian priors constraining the free parameters.

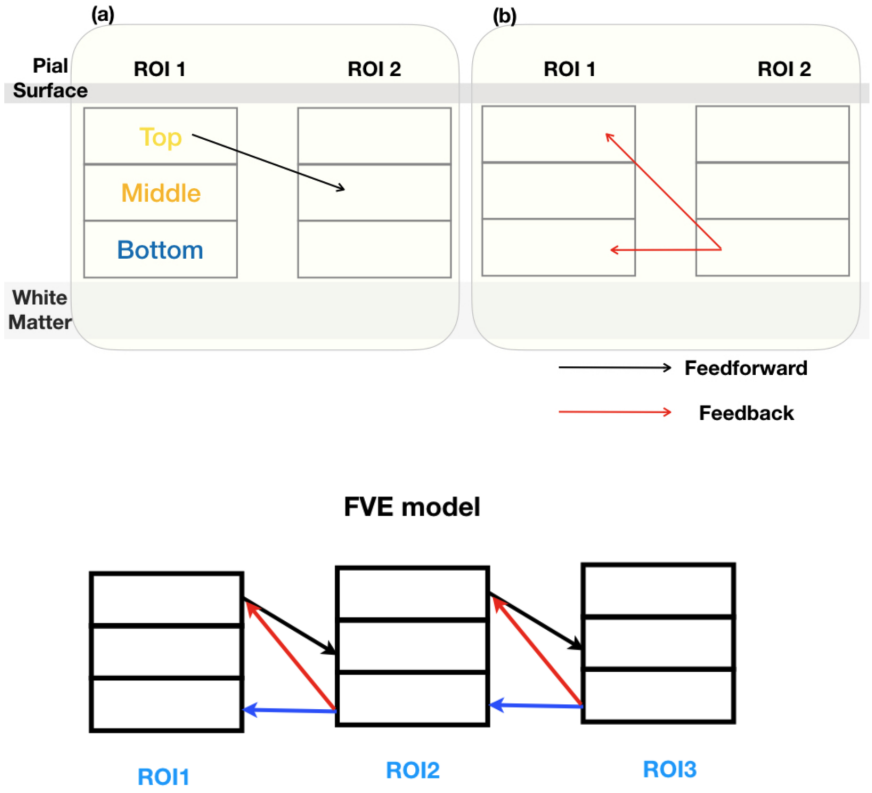


Figure 3.3: (a,b) Schematic of how FFW and FB connections were defined in models of 2 ROIs defined with 3 layers (top, middle, bottom layers). (c) Structure of the Felleman and Van Essen (FVE) model that defines 3 hierarchically organised regions connected with 1 FFW and 2 FB connections. In this study, the architecture of FFW and FB connections were defined in a way that represents typically observed connections from the canonical microcircuit theory and that a simplified version of laminar connections had been observed in past tracer studies [42, 52].

3.2.3 Experiments: establishing face validity

In the section above, we explained how a layer-specific SEM model can be established. Next, we establish the face validity of this model through simulations where we modify the number of regions (N), number of layers (L), and weights in the connectivity matrix A. For simplicity, we assume that all regions consist of 3 layers, resembling superficial, middle, and deep layers. In this work, the connectivity structure of central interest was FFW and FB connections introduced in Felleman and Van Essen's model of laminar connections. We used a simplified version of FFW and FB laminar connections as shown in Figure 3.3, in which FFW connections originate from the supragranular (upper) layer and terminate in the granular (middle) layer in another area (Figure 3.3a). FB connections shown in red arrows initiate from infragranular (bottom) layers and terminate outside of the middle layer and in either the top or bottom layer (Figure 3.3b). From now on, we refer to the model consisting of 1 FFW and the 2 possible types of FB connections as the Felleman and Van Essen (FVE) model, as shown in Figure 3.3c.

Studies preceding our current work first introduced how an acquisition involving multiple inversion-recovery pulses can be used to measure T1s of different layers [101, 102]. As can be seen in the formula for PV (Eq.3.10), the choice of TIs and T1s are related to the PV contribution of a layer. We chose T1s representing top, middle, and bottom layers from a sample structural data with a T1 distribution as in Figure 3.4. Peak values of the T1 distribution were selected and assumed to represent T1s for approximately top, middle, and bottom layers. We also accounted for the possibility of varying T1s in the cortex by randomly sampling a T1 value from each Gaussian curve fitted for each layer in the T1 distribution.

We first checked the invertibility of the layered SEM model using simulated data. Simulations were conducted on two types of models with either 2 or 3 regions of interest. Initially, a 2-ROI model with three sets of laminar connectivity patterns; FFW, FB, and FVE connections (i.e., the combination of FFW and FB connections) were examined. Then we explore a more complex model that assumes FVE connections between adjacent regions in a 3-ROI model. Our simulated data

was generated with random connectivity weights that consisted of both negative and positive real numbers. We show a basic backbone of these simulations as pseudocode in Figure A.1.

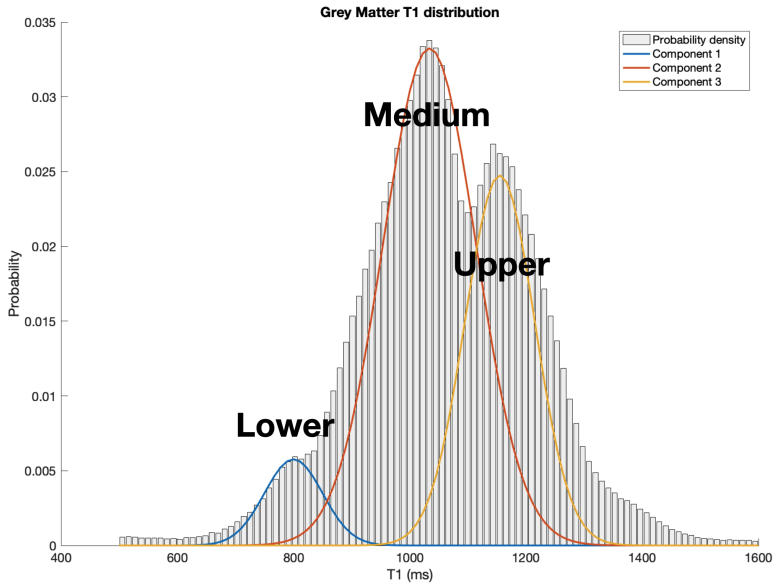


Figure 3.4: An example T1 probability distribution was taken from GM of a sample subject's structural data and Gaussian distributions were fit to peaks. Each peak value of fitted Gaussian curves was chosen as an input T1 required to calculate PV contribution to the signal (Eq.3.10). Lowest to highest T1 values in the graph were assumed to represent T1s of the innermost, middle, and uppermost layers in a cortical depth.

3.2.4 Real data acquisition

We acquired resting-state data for 5 healthy subjects using an SE-BOLD IR acquisition. This acquisition involved an initial SE-BOLD sequence preceded by inversion pulses added with inversion times of 300, 480, 690, and 810ms. Other scanning parameters were as follows: TE=45ms; TR=3s; 160 volumes (8 minutes per TI); voxel size = 3.1 x 3.1 x 4.5 mm. In this set-up, we were able to retrieve fMRI data covering most of the brain. Data was acquired and provided by collaborators in Tel Aviv University, Yaniv Assaf's lab (`https://yanivassaflab.com`). The data were preprocessed using a custom pipeline, developed specifically to address the unique contrast properties of the images, which differed from conventional GE data (contributed by Professor Saad Jbabdi). The preprocessing steps included brain extraction and motion correction, with volumes exhibiting excessive motion

being discarded. Furthermore, a subtle yet noticeable correlation was observed between slices acquired during the same inversion, presenting as streaking artefacts in whole-brain connectivity maps. To mitigate this issue, we regressed out the mean signal from groups of slices acquired at the same TI. To confirm that this adjustment did not adversely impact connectivity estimation, we conducted single-subject ICA on the processed data and verified the presence of all canonical resting-state networks. Additionally, our 5 subjects did not have T1 data, but only IR functional data. Therefore, we continued using a sample T1 distribution from a single, different individual (Figure 3.4).

3.3 Results

In this section, we report outcomes of applying our novel IR-BOLD acquisition with laminar connectivity modelling using SEM. Broadly, we validate models through model inversion using:

- Simulated data
- IR-BOLD data acquired in 5 subjects

In common, we used MCMC sampling for model fitting, and the best model of fit is achieved through minimisation of nLL, essentially, through comparison of data and model covariances. Next, we additionally introduce two types of tests that were used to assess various hierarchical model architectures:

- Hierarchy recovery test suitable for 3 ROI models
- Pairwise hierarchy test suitable for models consisting of numerous ROIs

Finally, we attempt a comparison of our results against a diagram from a seminal work that proposes hierarchical relationships between visual areas in macaque monkey brains [52].

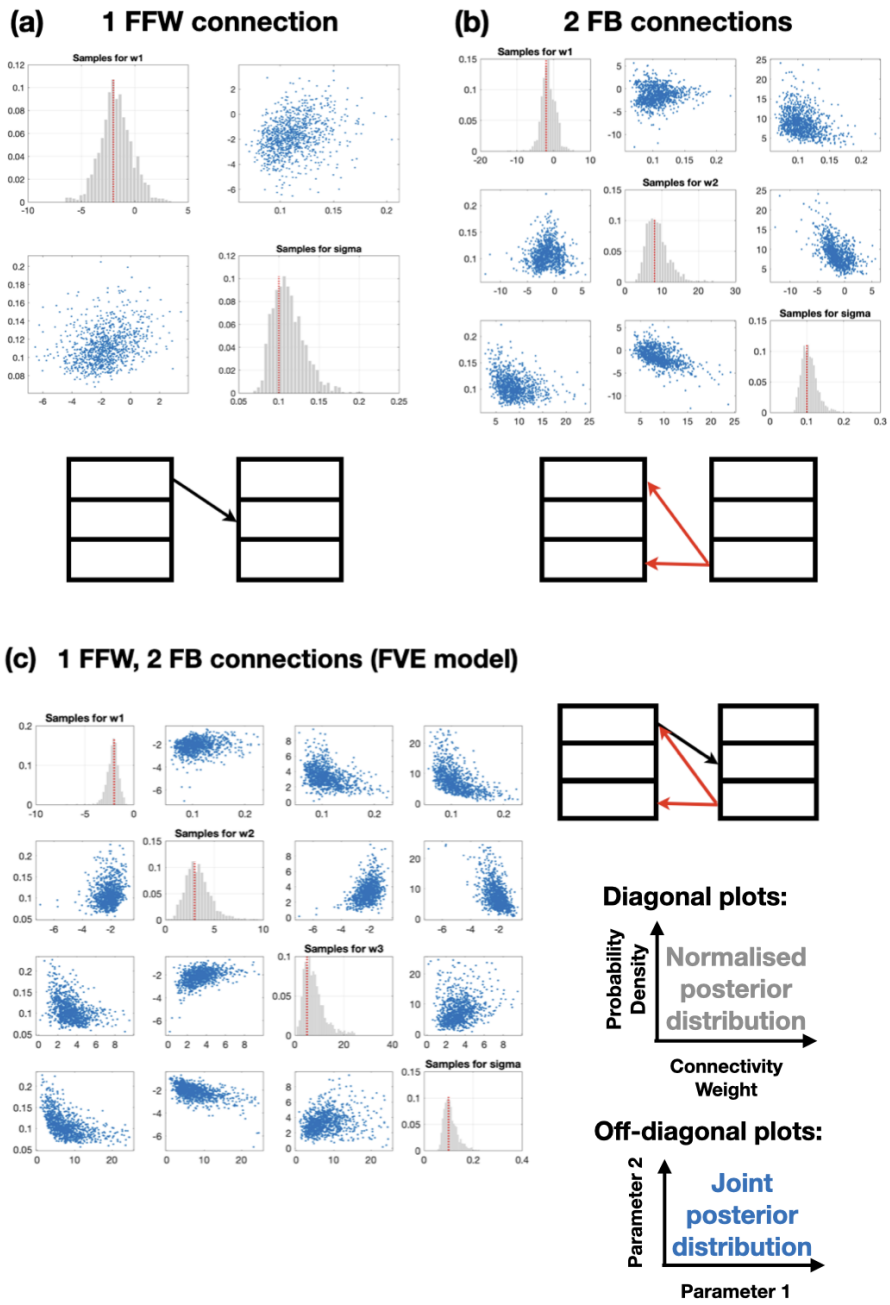


Figure 3.5: We used MCMC-MH sampling to fit model parameters to simulated data. Panels (a)-(c) are each based on models (Figure 3.3) with 1 FFW connection, 2 FB connections, and FVE connections (1 FFW, 2 FB connections). In each panel, diagonal plots show posterior distribution of free parameters and off-diagonal plots show joint posterior distributions between pairs of parameters. We observe sample convergence in these examples and ground truth connection weight values are marked with broken red lines.

3.3.1 Simulations for model validation

We first examined whether our models of interest were invertible using simulated data. In each panel of Figure 3.5, diagonal plots indicate sample convergence towards preinitialised connectivity weight coefficients marked with red vertical broken lines. We simulated data for each of the FFW model, FB model, and combined FVE model.

3.3.2 Hierarchy recovery test

Here, we sought to test whether our modelling approach can recover the visual hierarchy directly from the data based on assumptions on underlying laminar structural connectivity patterns. A key organisational principle of the visual cortex is that brain regions have a hierarchy in processing information. We do not fully know whether laminar connections in the FVE model underlie such relationships between brain regions. Thus, to test this concept using our IR-BOLD acquisition with SEM connectivity modelling approach, we simulate data from 3 regions with connectivity patterns in the FVE model, and conduct model comparisons to test whether we can recover the true hierarchical structure assumed for simulated data (Figure 3.6).

3.3.2.1 Simulated data

We simulated laminar data for 3 regions assumed to have a hierarchical relationship. Six possible models of alternate hierarchies can be defined for a model consisting of 3 ROIs (Region X < Region Y < Region Z, where 3 regions such as BA 17, 18, and 19 can be inserted into X, Y, Z without replacement) (Figure 3.6). The consecutive ROIs within a model were connected through laminar connections previously introduced as FVE connections (Figure 3.3). We simulated 50 random data sets and repeated the procedure of fitting the 6 candidate models to the data. As a result, we can compare mean nLL of the 6 different models and identify the best fitting model as shown in Figure 3.7. The best fitting model with minimum mean nLL value was Model 1. Since Model 1 was the model used to simulate data, our simulation was able to justify how the model with the correct hierarchical order fit best to simulated

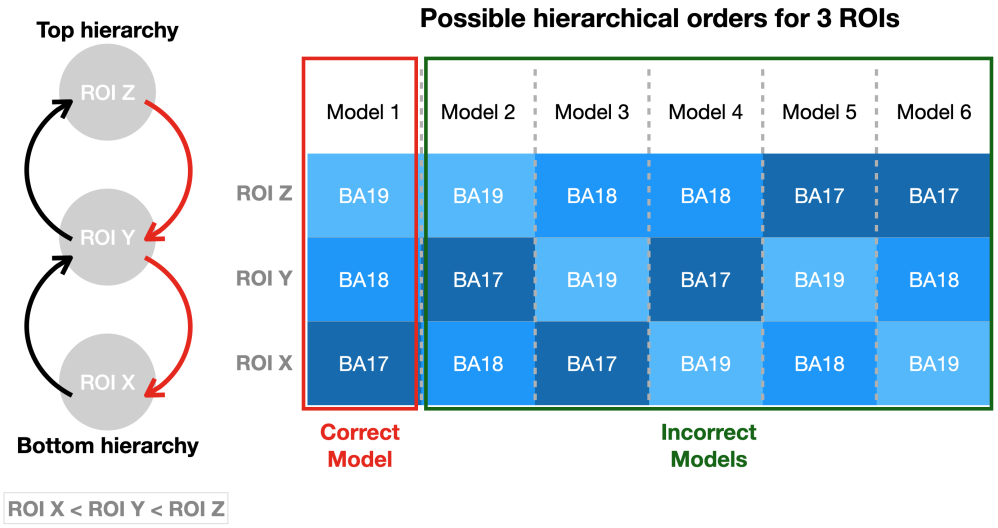


Figure 3.6: With a 3 ROI model we conducted the hierarchy recovery test, which was required to assess which configuration of the 3 ROIs fit a given data (both simulated and real) best. For instance, for real data we used an example of 3 early visual areas (BAs 17 – 19) with a known hierarchical relationship of BA 17 < 18 < 19. We fit 6 types of models by shuffling the hierarchical order assumed between the 3 areas. If the 'correct model' (Model 1) that represents the ground truth hierarchy fits the data best, the hierarchical relationship implied in that model can be considered as the winning model compared to the other 5 candidate models.

data compared to other models with alternate hierarchical orders between regions. In Appendix Figure A.2 we show statistical comparisons between the 6 models.

3.3.2.2 Hierarchy test for 3-ROI model in early visual regions - real data

Following simulations, we examine whether a correct model can be identified when all possible hierarchies in a 3-ROI model is fit to real fMRI data extracted for 3 visual areas: BA 17, 18, and 19. BAs 17 - 19 each correspond to visual regions V1 (primary visual cortex), V2 (secondary visual cortex), and the extrastriate cortex, which approximately indicates higher-order visual areas such as V3, V4, and V5. These regions were chosen for their known ascending hierarchical order, starting from BA 17, BA 18, and to BA 19 (BA 17 < BA 18 < BA 19). We fit the 6 possible model architectures (Figure 3.6) and observed lowest nLL for Model 1 (BA 17 < BA 18 < BA 19) in all 5 subjects. In Figure 3.8a, we ranked models from 1st to 6th in order of lowest to highest nLL. All 5 subjects identified Model 1 as the best fitting model, therefore, all show a rank of 1st place. Next, we repeated this

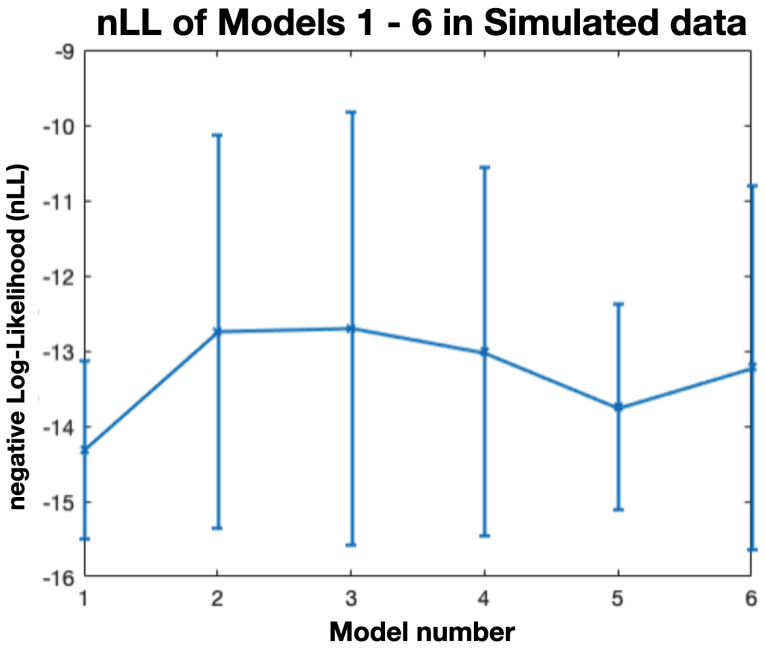


Figure 3.7: Negative log-likelihood (nLL) calculated for 6 models were fit to 50 instances of simulated data. Model 1 had a hierarchy of ROI 1 < ROI 2 < ROI 3 and was the ground truth model that simulated data was augmented on. Six possible models of alternate hierarchies were fit to 50 distinct simulated data sets. The plot shows how mean nLL was lowest for Model 1, hence, was the best fitting model among other models.

hierarchy test for 100 random sets of T1s sampled from 3 different cortical depth levels corresponding to layers in a lower, medium, upper part of the cortex. T1s for each depth range were randomly pooled from Gaussians fit to the sample T1 distribution in Figure 3.4. Figure 3.8b then depicts our results when accounting for variability of T1s in the cortex [102]. Over 5 subjects, an average of 66% (from 100 sample T1 sets) successfully identified the model with correct hierarchy of visual regions as best fitting (Figure 3.8b, red box).

3.3.3 Pairwise hierarchy test

In the previous section, we showed how our model can recover a hierarchical organisation among 3 regions, assuming an FVE connections. We move on to address whether this approach can be used to recover a hierarchy of the whole visual system, which would involve more ROIs. One possible model for more numerous regions would be to define a large connectivity matrix involving many more regions,

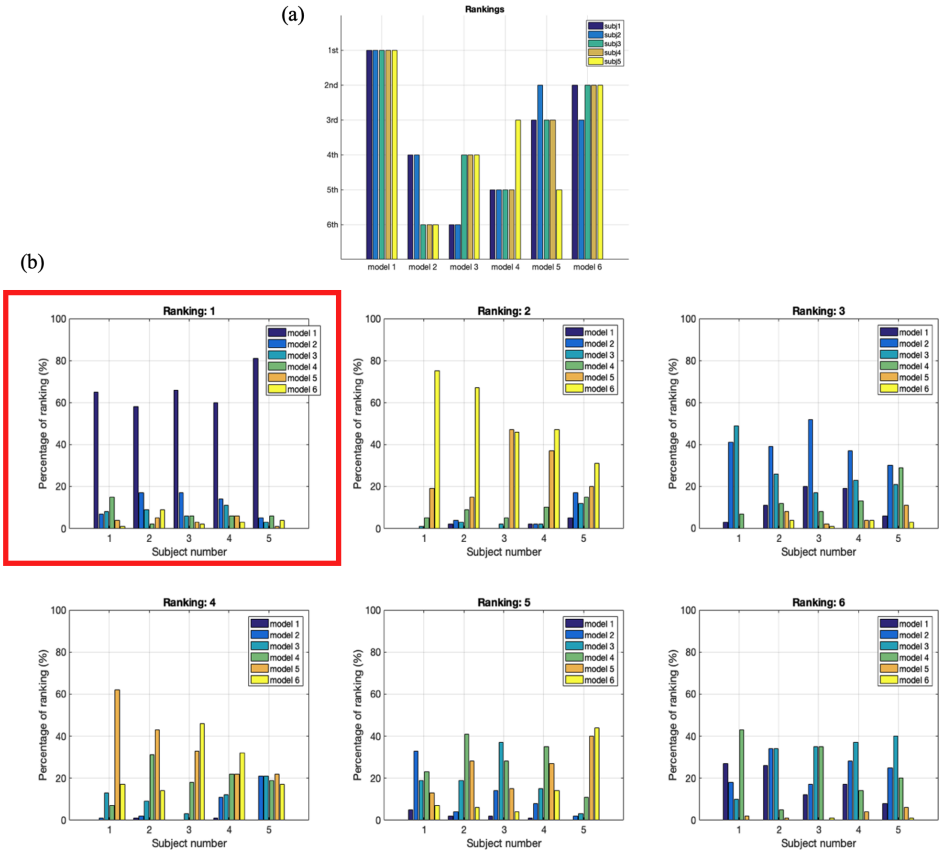


Figure 3.8: (a) For all 5 subjects, Model 1, which had the correct definition of hierarchies between BAs 17, 18, and 19, fit the fMRI data best. Ranking of 1st place corresponded to the model with the lowest nLL, thus was the best fitting model among candidate models. Model 1 was ranked 1st for all subjects, and other models were also ordered according to ascending order of nLL value (i.e. ranks 1 - 6 correspond to low to high nLL values). (b) Next, the schematic in plot (a) was repeated for 100 random sets of T1s, while each set consisted of 3 T1s representing upper, middle, and bottom layers. Here, 6 subplots each display percentages that a model was identified as having a rank between 1 - 6. We focus on the highlighted red box, in which we observed Model 1 ranking 1st (i.e., was best fitting model among candidate models) for all 5 subjects (average of 66%). In the adjacent subplot showing frequency of models ranked 2nd, one can observe that Model 6 was most frequently ranked as having 2nd to lowest nLL in 4 subjects (subjects 1, 2, 4, 5). Details of statistical significance has been addressed in Figure A.3.

but this would have two drawbacks. First, the number of weight parameters to infer becomes too large for efficient inference, therefore, model convergence would be difficult to reach. Second, determining the best model among candidate models will no longer be feasible due to the exponential growth in model architecture combinations. In this section, we introduce an approach to recover the hierarchical organisation amongst an arbitrarily large number of ROIs. To do so, we fit 2-ROI

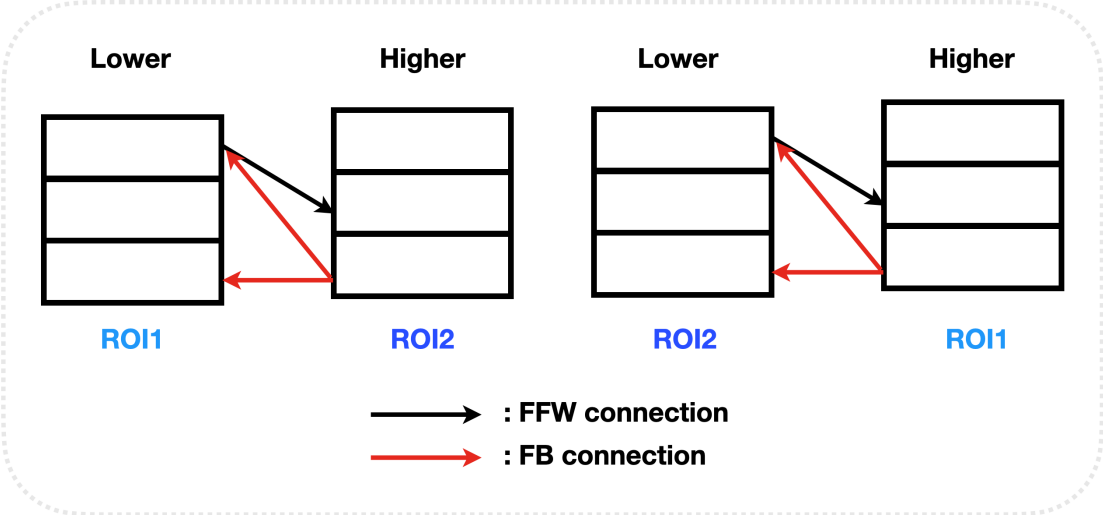


Figure 3.9: In the pairwise hierarchy test, we fit 2 types of candidate models shown above for each pair of ROIs. The model defines FVE connections (1 FFW, 2 FB connections), hence assuming that a hierarchy exists between the two ROIs. The difference between the 2 candidate models is that the hierarchical relationship between ROI 1 and ROI 2 are flipped. We used the same metric nLL to identify the better fitting model between the 2 candidate models.

models to each pair of regions and only compare two models with alternating hierarchies (Figure 3.9). After doing so for all existing pairs of regions, we then simply count the number of times a given region was selected as being lower (or higher). This type of pairwise test results in a scoring of all regions that can serve as a surrogate measure of hierarchy.

3.3.3.1 Test for comparing 2-ROI models on 9 visual regions

We extracted fMRI timeseries from a fine-grained parcellation of visual regions from Glasser's parcellation of the human cortex Glasser et al. [108]. We chose a subset of regions and examined 9 visual areas (V1, V6, V2, V3, V4, FEF, V3A, MT, V4t) that existed in both human and macaque brains at similar locations, so that we can directly compare against regions in the FVE diagram.

Accordingly, fMRI BOLD time series were extracted for those 9 regions and we fit our models to the data. In this experiment, instead of increasing the number of regions in our SEM model to 9 ROIs, we used the approach of comparing pairs of 2-ROI models. Each pair of models would have 2 different hierarchies, that is, when

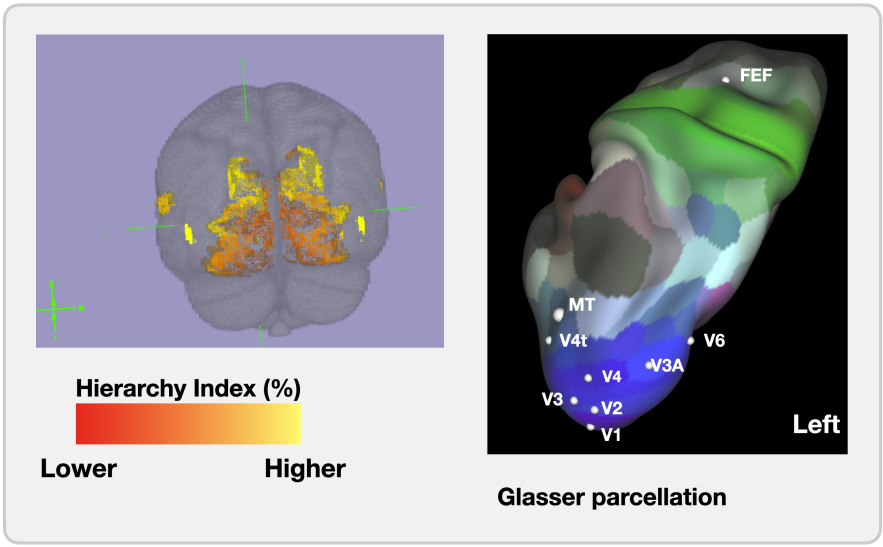


Figure 3.10: As a result of pairwise hierarchy tests, the model with the minimum nLL value indicated a 'better fitting' model. Thus, the hierarchical relationship between 2 ROIs can be revealed. This, however, then had to be evaluated along with all other visual areas to achieve a full account of hierarchical relationships among 9 regions. We used a metric, namely a hierarchy index, that was the percentage that another region was identified having an inferior hierarchical level than oneself. The left figure displays mean hierarchy indices in the 9 visual regions for 5 subjects and demonstrates a clear anterior-posterior gradient of distinct hierarchies. The right figure displays selected 9 visual regions from the Glasser parcellation that were used as masks to extract BOLD time-series data.

hierarchies between two regions are either ROI X $<$ ROI Y or ROI X $>$ ROI Y. We fit these two models of opposite hierarchical relationships (Figure 3.9) and compared nLLs to decide on the better fitting model between the two. In such way, we could avoid increasing the number of parameters to be inferred in a model that would in turn increase the complexity of models and make the model difficult to invert.

After repeating this assessment for all region pairs (36 pairs), we counted the number of regions that were hierarchically below oneself; this value was called the hierarchy index. Finally, hierarchy indices were converted to a percentage value out of the total number of unique ROI pairs.

Average hierarchy indices of 5 subjects were computed for each region and displayed on the brain surface as in Figure 3.10. Here, a clear hierarchical gradient from lower to higher index values appeared in the posterior-anterior direction in the brain. In other words, regions with higher hierarchy indices were hierarchically

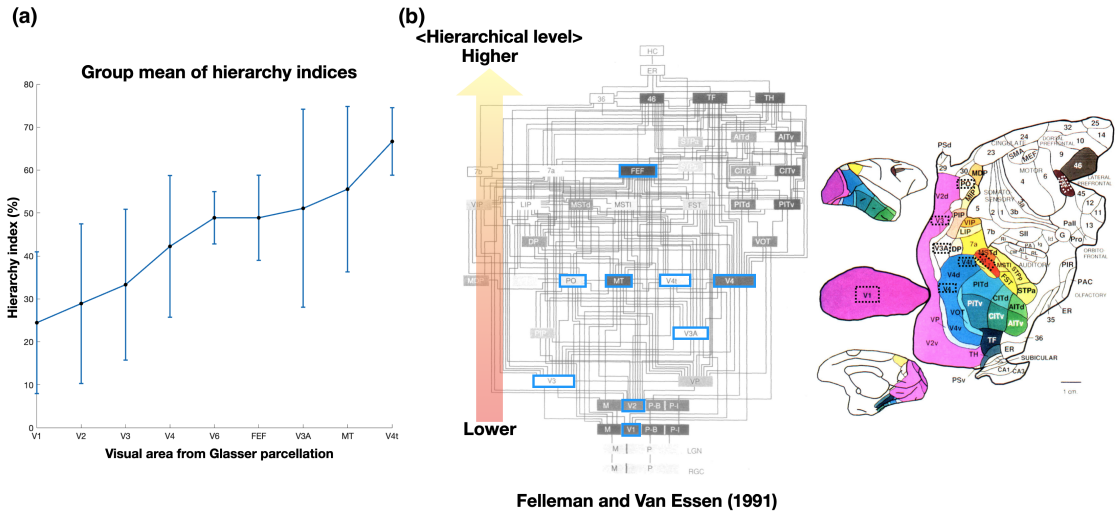


Figure 3.11: (a) We show average hierarchy indices over all subjects while accounting for the potentially variable T1s of individual brains. Visual areas were rearranged in order of increasing hierarchy indices. (b) Hierarchy diagram from Felleman and Van Essen [52]. Rearranging of hierarchy indices in (a) was to generate a similar layout as overlapping visual areas in (b). Felleman and Van Essens' work was on macaques, so we highlighted regions that overlap with our study with blue boxes in the hierarchy diagram and with black broken line boxes on the flattened macaque brain diagram.

superior than those with lower hierarchy indices.

3.3.4 Comparison of hierarchy diagrams

In an alternate configuration, we displayed average hierarchy indices of the 9 visual regions in ascending order as in Figure 3.11a. Similarly to our practise in accounting for variable T1s, we sampled 100 random sets of T1s (i.e. 3 T1s per set, representing each layer). Felleman and Van Essen distinguished laminar connection patterns revealed from their tracer study results and analysed hierarchical relationships between visual areas Felleman and Van Essen [52].Their final diagram informs how visual regions were segregated into distinct hierarchical levels (Figure 3.11b), and we produced a similar representation of hierarchies. To do so, we plotted hierarchical levels from the past study against our hierarchy indices separately for 5 subjects and ona group-level using average index values (Figure 3.12). In most subjects, excluding subject 5, a strong correlation was observed between the two labelling systems (FVE hierarchy levels vs. hierarchy indices in this work), indicating an agreement in the hierarchical trend between the two approaches.

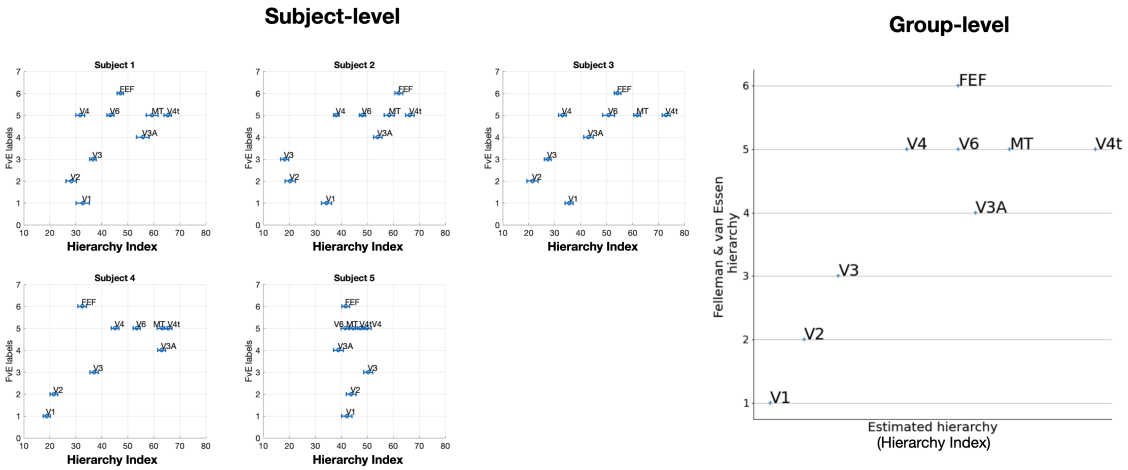


Figure 3.12: On the left side, we show our hierarchy indices computed for each subject against the hierarchical levels from Felleman and Van Essen's study referred to as FVE labels for each individual. The final plot on the right plots the mean hierarchy indices computed in Figure 3.11 against FVE hierarchical levels. This representation intends to show how the trend of hierarchies change in the two different labelling systems.

3.4 Discussion

One of the main goals of this chapter was to use a SE IR-BOLD acquisition and connectivity modelling based on laminar-specific connections to recreate hierarchical relationships between visual areas. Our method combines IR-BOLD acquisition at multiple TIs and modelling with SEM to consider PV effects local to each layer at each TI. We aimed to validate this novel approach in which we can assess layer-dependent connectivity of the human cortex through a different approach from the more prevalent laminar fMRI conducted in high-field MRI. In essence, rather than leveraging high-resolution for disambiguating cortical layers spatially, such approach can be conducted in the more widely available low-field MRI in 3T. Even in high-field acquisition, which is the more favoured method for resolving layers spatially, blood draining effects continue to affect the BOLD signal measured (discussed further in Chapter 5). By implementing SEMs, we can account for PV contributions in each layer when modulating the fMRI signal at different TIs. Following simulations that justified that our models could be inverted, we defined and fit SEMs that assumed connections that relate to hierarchical relationships.

3.4.1 Considerations on model architecture

We first simulated data and validated the invertibility of SEMs with ROIs connected through FFW, FB, and a combination of those connections referred to as FVE connections. Then, with real data in three early visual regions, BAs 17–19, we fit the 3-ROI model as these regions were known to process inputs sequentially following a hierarchy in processing order. The advantage of fitting the 3-ROI model is that we can assess whether laminar connectivity patterns in the FVE model that assume a certain hierarchy (BA 17 $<$ BA 18 $<$ BA 19) can be identified correctly in these early visual regions.

In this study, we used 3 layers to perform multi-layer SEM, and each layer broadly represented the top, middle, and bottom layers. When layering structures were considered in MRI studies, it was often a matter of discussion on the total number of layers required to optimally delineate layers. From past histology studies on post-mortem human brain sections, we already know that human cortical layers are mostly around 5–6 layers, with exceptions in certain areas such as the motor cortex that lacks a layer IV. Therefore, if our goal in this study was to spatially distinguish all layers in different cortical areas, there would be greater emphasis on the accurate distinction of separate layers. However, the aim of this study was to make use of variable PVs in voxels obtained at relatively larger sizes through low-field (3T) MRI. Therefore, we adhered to a less stringent rule and divided the cortical thickness into 3 equidistant layers whereby the cortical depth was divided into bands of equal thicknesses.

After initially using fixed T1 values extracted from structural data, we also sampled random T1s from Gaussians fit to the sample T1 probability map. This way, we tested for situations in which T1s would vary in different brain areas or individuals. Even after sampling T1s, we could still recover the model of correct hierarchical order most frequently among 6 permutations of hierarchies for the 3-ROI model.

We additionally tested a more complex version of this 3-ROI model that also connects the bottom and top region with FVE connections (9 connectivity weights being free parameters) so as to emulate a fully connected model. However, when

using MCMC sampling to predict unknown parameters of complex models, it becomes computationally heavy; therefore, it was extremely time-consuming for the Markov Chain to reach convergence. Moreover, parameters were no longer able to be accurately estimated with too many parameters to estimate. As a resort, in this work, we focused on examining hierarchical relationships between a small number of regions (2–3 ROIs per model) to keep models simple and parameters recoverable. Similarly, experiments on real data were circumscribed to visual regions that were richer in evidence from past animal studies for being hierarchically organised [43, 52, 109] and with well-received studies on its laminar connectivity patterns [42].

3.4.2 Evaluation of pairwise hierarchy test

Neurons in V1 are known for being orientation selective from a renowned past study by Hubel & Wiesel Hubel and Wiesel [110]. They suggested a significant concept related to hierarchical processing of the brain in modern neuroscience. Specifically, they claimed that different regions along the visual hierarchy have different roles and respond to different features within a visual representation, such as contours, orientation, colour, and others. However, we do not have a complete understanding regarding the order or mechanism underlying the processing of those distinct features. Those visual regions are often chosen as potential regions that have a hierarchy for being processed, and likewise, we chose those regions to assess hierarchy through fitting of models that imply certain laminar connections.

We used a pairwise hierarchy test to examine hierarchical relationships between 2 ROIs. This was repeated for different pairs of regions to ultimately establish a hierarchical diagram comparing all the region pairs that was achieved by computing a hierarchy index by counting the number of regions that have a lower hierarchy than oneself. The pairwise hierarchy test for multiple visual regions was an initial attempt to study hierarchical relationships between cortical areas, which has been a difficult problem to tackle in human neuroimaging.

Connectivity relationships between cortical areas are usually complex; for instance, they often simultaneously arise in multiple directions, and some areas are

connected through indirect paths. Therefore, it has been notoriously difficult to clarify what the role of longer-range connections is. However, now there are emerging results in modern tracer work that support the existence of long-range connections relevant to the prefrontal cortex that can be characterised as having weaker connection strengths [111]. There are also new perspectives on the role of FFW and FB connections in constraining local circuits occurring in shorter ranges [44].

We still heavily rely on Felleman and Van Essen's macaque study as invasive approaches are not feasible in the human brain, but those prevailing results on other species still need to be validated in humans. On one hand, the current study focused on model validation and applying the scheme of sensitising BOLD signal with PVs that depend on the choice of T1s and TIs on visual areas with more well-known hierarchical relationships. Next, we can extend the current methodology and examine some less well-studied long-range connections so that we can better understand their roles with respect to connections with shorter distances but stronger connection strengths.

3.4.3 Limitations in comparing hierarchical systems

Due to limited access to studies that investigate hierarchical relationships in cortical areas beyond early visual areas, it was not straightforward to compare nLL-derived hierarchies of our 9 visual regions. We still lack a complete knowledge on hierarchical relationships between cortical areas where research is ongoing. For instance, recent tracer studies have claimed hierarchies between visual regions that deviate from Felleman and Van Essen's hierarchy diagram based on advanced tracers that revealed more subdivided connections than simply being one of FFW, FB, or lateral connections [43, 44].

We chose the 9 visual regions in this study as they anatomically and functionally overlapped with regions in Felleman and Van Essen's tracer study on the macaque brain Felleman and Van Essen [52]. This was for the purpose of directly comparing areas with identical nomenclature. However, this did not mean that these macaque and human brain regions were exactly the same regions as some areas differed

in sizes or in location within the brain. We focused on the relative similarity of functional roles that visual regions have within the visual processing stream, but this comparison between species may be subject to variability. For instance, the frontal eye field (FEF) was segregated into smaller subregions in latter studies [44] compared to FEF from Felleman and Van Essen's study. As such, a direct comparison of hierarchy labels based on overlapping 'region label' was not virtually possible even in our comparison between nLL-based hierarchy indices against area labels in Felleman and Van Essen's study.

Additionally, Markov and colleagues also found that the FEF formed a strong loop with area V4, hence FEF was assigned a hierarchical level lower than what had been previously claimed. An overarching aim of studies investigating hierarchies is to develop a common rule that can help us understand the mechanism in which information is processed in different cortical areas. It is crucial that studies that focus on the human brain continue to seek more reliable repeated observations that can help understand how cortical areas, such as those involved in long-range connections or higher-order areas, communicate with other areas. Moreover, as such hierarchical relationships are defined upon relatively preserved anatomical patterns in FFW and FB activity between layers, studies elucidating such laminar connections can provide a foundation for narrowing the knowledge gap between the structural and functional organisation of human brains.

From a practical point of view, one limitation of our approach is that we assume fixed values for T1 relaxation within each layer, and we further assume that these values remain constant across brain areas. One possible extension of this work could involve fitting T1 values, but this would require denser sampling over inversion times and possibly the use of task data instead of resting state to further constrain the modelling.

3.4.4 Extension of current work

One of the endpoints of this work was to examine the plausibility of an IR-modulated fMRI approach combined with SEM to other cortical areas, especially beyond the

visual system. Initially, we can apply model fitting to other dorsal regions in the visual system to check for the existence of a hierarchical gradient in the system. We can also look at other cortico-thalamic systems or the motor cortex, which lacks layer IV but conducts hierarchical processing through its own underlying laminar connections, to examine whether a hierarchical relationship still holds in such regions with variable laminar structures.

The next phase in multilayer SEM modelling requires modelling of blood haemodynamics using an extended dynamic causal modelling (DCM) [78] and is explored in Chapter 5. A modified version of DCM can allow us to move away from fitting covariances and towards fitting dynamics. In order to accurately assess laminar-specific signals, the haemodynamic model has to be able to model effects of blood flowing perpendicular to the cortical surface. Blood draining effects near the pial surface are not accounted for in the current SEM, so we can test an extended DCM on our data.

Once these models that account for blood haemodynamics are validated, we can also begin to apply certain task-based visual experiments that intend to activate layers involved in, for instance, top-down attention. A recent study using visual illusion was able to disambiguate FFW and FB activity [83], and such tasks can be used to reveal directed pathways involved in human cognition using fMRI. The present work focused on the resting brain to first understand the model to be implemented, and this can potentially be followed up by studies regarding activities pertaining to higher-order cognitive functions. Those cognitive functions and how they arise with respect to the cortical laminar organisation are poorly understood, and we do not yet know of precise physical substrates that give rise to complex functions. Furthermore, with a deeper understanding of laminar connections that evoke FFW and FB activity in cortical areas, there is scope to analyse mechanisms related to predictive coding that is known to be conveyed via FFW and FB connections.

3.5 Summary

Finally, we summarise some key results and address limitations that were encountered.

- With simulated data, we confirmed we can reach MCMC sample convergence when fitting layered SEMs for weight estimation.
- Models that assume a certain hierarchical order based on organisation of FFW and FB connections can be compared and negative log-likelihood can be used to determine a best fitting model.
- For real IR SE-BOLD subject data the hiearchy test was able to determine a model with correct hierarchical order in 3 visual areas.
- For a more fine-grained parcellation of 9 visual areas, we determined a better fitting model between pairs of regions. Based on those pairwise results, hierarchies were deduced for 9 visual regions. They appeared similar to a macaque hierarchicy diagram from Felleman and Van Essens' classical study.
- For real data, we were not able to consider variable T1s of different cortical areas within a subject and also between subjects. This is one possible reason that we were not able compare other pairs of regions beyond visual areas. For future work to explore the whole brain we may require better quantification of subject-specific T1s.
- Moreover, for further work we would benefit from additional acquisition of subject data as we were using a small set of 5 subject data here. Difficulties were encountered when comparing our results against hierarchical relationships of macaques from the past. Thus, we would benefit from acquiring more subject data if we would avoid relying on the limited ground truth data of human cortical areas.

4

Cortical layer microarchitecture: combining high-resolution diffusion MRI and histology

Contents

4.1 Introduction 67
4.2 Materials and Methods 69
4.2.1 Rationale and analysis plan 69
4.2.2 Data 70
4.2.3 Defining cortical types 72
4.2.4 Design of study 74
4.2.5 Analysis 1 77
4.2.6 Analysis 2 and 3 81
4.3 Results 82
4.3.1 Analysis 1: BB to Cortical types 82
4.3.2 Analysis 2 – High-resolution MRI to Cortical types . . . 86
4.3.3 Analysis 3 – BB to High resolution MRI 90
4.4 Discussion 91
4.4.1 Analysis 1 – Clustering BB histology as cortical types . 92
4.4.2 Analysis 2 and 3 – Diffusion MRI measures to Cortical types 97
4.5 Conclusion 99

4.1 Introduction

The study of the cortex is a complex, multifaceted problem that requires understanding of microstructural organisation to how brain activity evolves from those microscale units to allow us to execute cognitive functions. In this chapter, we focus on methods and structural aspects pertinent to investigating cortical layers using histology and diffusion MRI (dMRI) as opposed to focusing on laminar connections with respect to functional connectivity. In the past, human brains were examined only through microscopic observations of stained histological slices post-mortem. Brain histology has previously revealed that that the cortex consists of 5 - 6 layers organised parallel to the cortical surface. Although histology offers immense detail of cortical cytoarchitecture, histology is a limited resource that cannot inform us on how component cells communicate with each other in the living brain.

The majority of dMRI research has focused on anisotropic WM, but grey matter (GM) can also benefit from in vivo studies to go beyond what 2D histology has offered to date. dMRI has mostly been used to examine connections between cortical or subcortical structures that are relayed through white matter (WM). From past histology studies, we know that GM consists mostly of cell bodies (soma) and dendrites of a neuron and span layers of the cortical thickness, whilst its axons are sent below to WM. Many short or long-range connections initiate from and pass the GM before axonal dendrites reach its target areas. Therefore, dMRI at improved resolution has the potential to reveal more on microstructural and laminar organisation of neurons in GM.

Many microstructural aspects observed in GM are area specific and are not too far from cyto- or myeloarchitecture that neuroanatomists had observed in the past [3, 7]. Some brain areas are well known for more characteristic layer-specific cellular composition from histology. For instance, the motor cortex has a prominent layer V with large pyramidal cells known as Betz cells and a less developed layer IV [27]. Also in V1, a thick myelinated segment parallel to the cortical surface (Stria of Gennari) can also be observed. Histology is still integral when it comes

to comparing against other imaging methods as it provides ground-truth anatomy of cortical architecture whilst brain parcellations vary according to a criteria used to divide the brain into sub-compartments [39, 112].

Despite the fact that brain histology has been an indispensable resource for studying cytoarchitecture in earlier years, there are challenges when handling biological tissue. Histological slice preparation is time-consuming and laborious as brain tissue can be damaged from tearing or folding. Even during analysis of the curated data, interpretation of histological observations on cytoarchitecture is of subjective nature and quantitative ways that accurately express cytoarchitectural properties are still lacking [15]. Moreover, histology samples are only retrievable post-mortem so that study of individual brain microstructure during the living time course is not feasible for humans.

Many researchers endeavour to improve spatial resolution and reconstruction of fibres so that we can leverage this non-invasive method and push towards conducting in vivo histology through dMRI. Some recent research on human brains in high-resolution ex-vivo [98], in-vivo high-resolution 3T [113], and 7T [92] have detected varying radial fibre orientations across the cortical GM that were unable to be seen before. There is also evidence that dendrites contribute to diffusion anisotropy of GM [96, 114]. As such, there are anisotropic elements being observed within GM which would have been difficult to access in standard resolutions.

In this current chapter, we conduct a thorough examination of a high-resolution histology (post-mortem) and diffusion data (in vivo) of two different healthy individuals. The aim of the study was to observe whether diffusion measures derived from dMRI acquired in submillimetre resolution of 760 μm reflect any cortical laminar properties observable in high-resolution histology (BigBrain). By taking advantage of the high-resolution permitted in each modality, we attempt to examine microstructural organisation of cortical GM on a laminar level. Firstly, we design proxy measures of laminar properties from BB staining intensities that vary across the cortical depth. We then perform K-means clustering on laminar feature metrics to produce a brain map of cortical areas with unique laminar patterns.

Then, we test whether it is possible to derive histology-based cytoarchitectural parcellations from diffusion measures using a supervised machine learning approach. This way, we can assess whether dMRI features can be used to classify a histological brain (staining intensities) based on cytoarchitectonic properties.

4.2 Materials and Methods

4.2.1 Rationale and analysis plan

The overarching aim of this chapter is to investigate to what extent high resolution in vivo MRI can inform us about cortical microarchitecture. This research chapter is a first step in this direction. To this aim, we conduct a three-way analysis involving three sources of information: (i) a historical subdivision of the cortex into so-called "cortical types" defined on the basis of detailed cytoarchitecture; (ii) the BB dataset which provides a semi-quantitative measure relating to histology; and (iii) in vivo MRI including diffusion and a "myelin-related" MRI-derived map.

In analysis 1, we devise a new method for relating BB to cortical types. While these cortical types are defined through detailed cytoarchitectural descriptions (see Table 4.1), the BB data only provides a coarse metric derived from staining intensity from cell body staining using a Merker stain [115]. We develop 5 quantitative metrics aimed at characterising the BB profiles in a way that they can be related to cortical types and we validate these by using unsupervised clustering of the whole cortex based on these features.

In analysis 2, we directly map cortical features from MRI onto cortical types using supervised learning. The aim is to ask to what extent MRI data contain sufficient information to inform on cytoarchitecture at this level of granularity.

Finally, in analysis 3, we quantitatively relate MRI features to BB profiles via a regression analysis and identify key contributing features. Figure 4.1 summarises the structure of the analyses in this chapter.

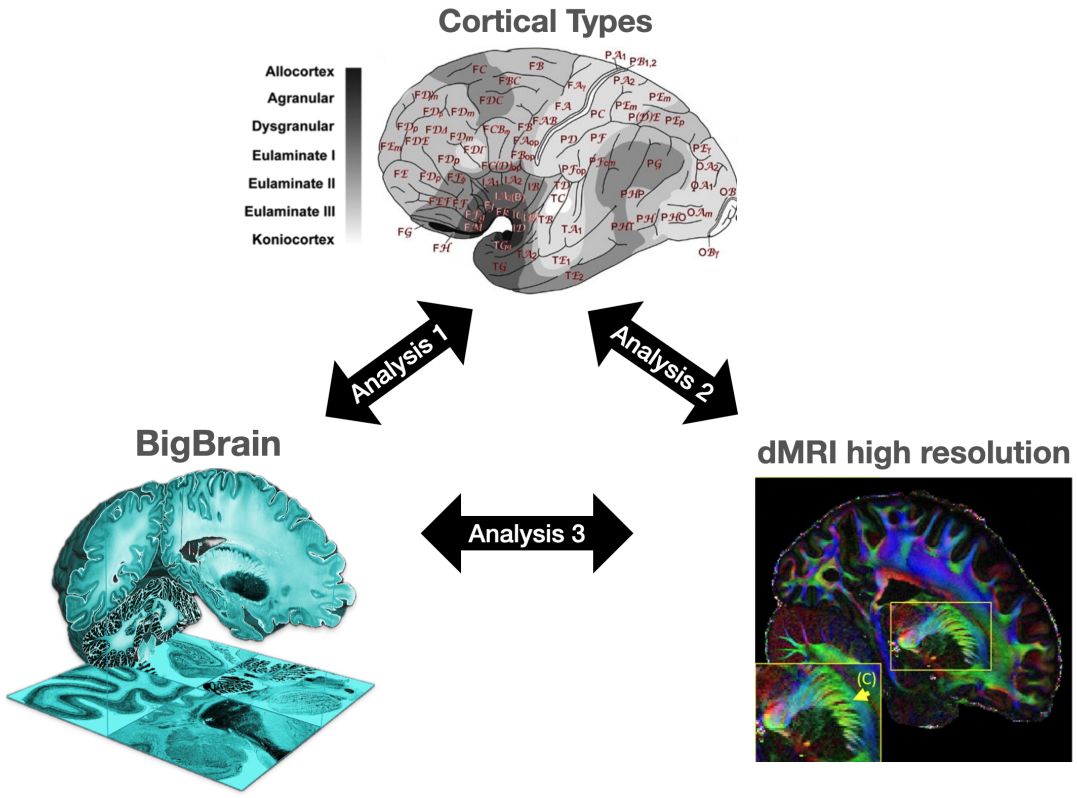


Figure 4.1: Schematic of the analyses conducted in this chapter.

4.2.2 Data

4.2.2.1 Histology data from BigBrain

BB staining intensity profiles were provided as open-source data and is described in detail in the BBWarp Toolbox paper [116]. The data consisted of staining intensities (Merker Staining) BB data acquired at 100 μm resolution [13] (Figure 4.2). Cortical profiles of staining intensities were obtained through construction of 50 equidistant surfaces between the pial and the white matter surfaces. To account for variations in cortical thickness as a function of curvature, an equivolumetric fractional depths scheme was employed [117]. The BBWarp toolbox enables BB histological slices to be transformed and registered onto freesurfer space [116] which is commonly used for MRI data. Following this step, BB staining intensities were transformed to a 3D brain surface formed of a mesh of vertices. Such transformation to an MRI data compatible space was required to conduct the main comparative analysis between BB histology and high-resolution diffusion MRI data.

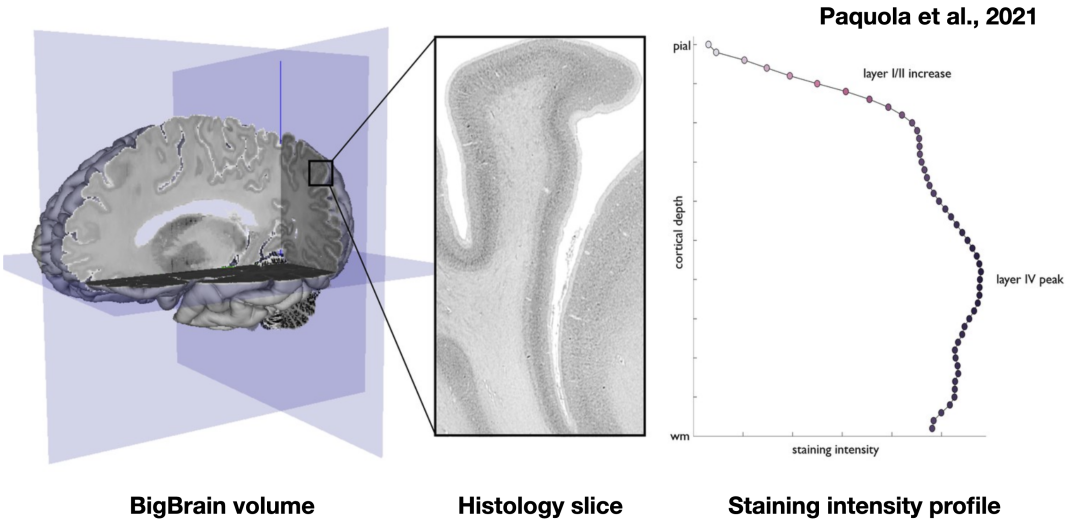


Figure 4.2: Schematic of the format that BB data comes in. Histology slices in the middle were carefully curated in its native histological space is given, and then was registered and reconstructed into a 3D volume. One related resource included were staining intensity profiles of stained cell bodies (Merker staining) that were sampled in points through the cortical depth and corrected for noisy parts, such as by excluding distorted slices. The above figure was taken from [116] and the resolution of BB used to explain the process of BB data curation was 20 μm.

4.2.2.2 Diffusion measures from high-resolution dMRI

We used a high-resolution in-vivo diffusion MRI dataset for a single subject released online from researchers at MGH [118]. At the time of writing, this was the highest resolution in vivo diffusion MRI dataset of a whole human brain ever acquired. This data was selected for our purpose of accessing diffusion related features to examine microstructural organisations in cortex, which requires high-resolution data. Wang and colleagues were able to acquire diffusion MRI of a whole brain on a 3T Connectom scanner capable of achieving gradients up to 300mT/m, which means one can increase the b-value whilst keeping the echo time short. The combination of using the Connectom system and a newly developed multi-slab sequence enabled them to achieve a resolution of 0.76 mm with high SNR. The data was pre-processed by applying distortion and motion correction using FSL's Topup and Eddy prior to modelling the voxel-wise signal.

To relate cortical microstructure to the histology data, we extracted summary metrics from the dMRI data by fitting two models: the diffusion tensor model (DTI

as implemented in FSL's DTIFIT) and a fibre orientation model to examine fibre crossings (ball and stick model as implemented in FSL's bedpostx). From dtifit, our summary metrics were as follows: fractional anisotropy (FA), mode of anisotropy (MO), parallel diffusivity ($L_{\parallel}$), and perpendicular diffusivity $L_{\perp}$. From the ball and sticks model, our features were the fibre volume fractions f_1, f_2, f_3. Finally, we used a proxy measure for myelination computed as the ratio between T1-weighted and the T2-weighted image data (T1w/T2w) [37]. All diffusion measures were initially in volumetric space, hence were resampled onto cortical surfaces to obtain vertex-wise cortical profiles. We used 6 equivolumetric depths between pial surface and WM and Figure 4.3 demonstrates the mean metric value for each vertex on the brain surface.

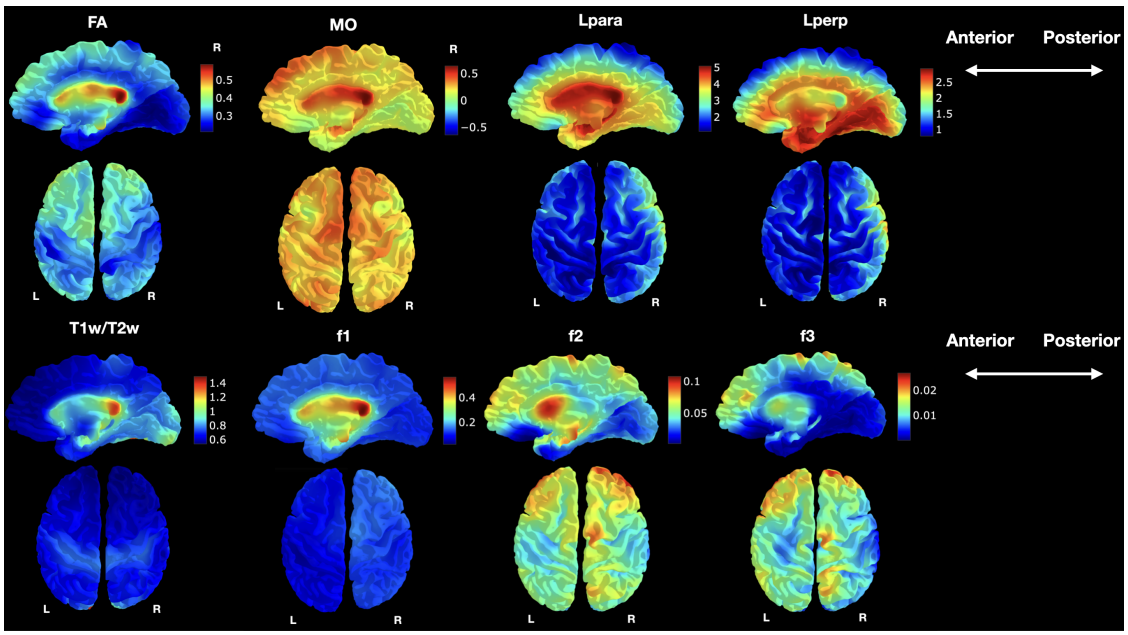


Figure 4.3: Mean MRI metrics over all layers shown on cortical surface.s

4.2.3 Defining cortical types

Cortical areas are often classified by their similarity in laminar structures. Von Economo and Koskinas defined cortical types by granularity of cells, that is, density and sizes of cell bodies. However, the complexity of laminar detail is prone to observer bias as it can also be defined by alternate criteria that an observer selected to differentiate areas. Cortical type definitions depend on several broad but key

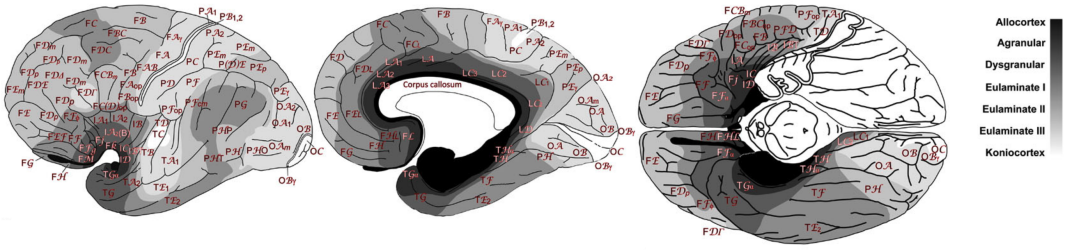


Figure 4.4: Parcellation of the human cortex into cortical types based on cytoarchitecture (from García-Cabezas, Hacker, and Zikopoulos [119]). The cortex is subdivided into 7 cortical types; allocortex, agranular, dysgranular, eulaminate I - III, and Koniocortex through visual inspection of von Economo and Koskinas micrographs and physical specimen for frontal cortex. The figure was modified from the original publication [119]

features of laminar organisation. A more recent study by García-Cabezas, Hacker, and Zikopoulos [119] used example properties of cortical lamination that differentiate areas into cortical types through varying thickness of Layer IV, identification of the more distinct layer group between superficial and infragranular layers, and the degree of differentiation of layers V – VI. As a result, areas with similar laminar properties were categorised into the same cortical type (Figure 4.4). Additionally, in their cortical type definitions along with Von Economo and Koskinas' definition of cortical areas into cortical types, a common topological rule holds for cortical type definitions. The rule imposed was to keep adjacent cortical areas to be ordered by increasing laminar complexity. Cortical types have been defined in variable ways as long as cortical types were defined so that simplest laminar properties are closer to the allocortex in the limbic system, and cortical types with more elaborate laminar patterns are towards primary sensory areas. The benefit of such cortical type definitions is that shared laminar properties of cortical areas can be examined at ease for purposes of searching for a pattern or rule that can explain microstructural organisation in terms of cytoarchitecture. In previous work, such cortical type definitions were used to examine connectional relationships. It has been suggested that cortical types with less layer differentiation tend to connect to more complex cortical types via feedback connections, whilst feedforward connections follow the opposite pattern [120, 121]. This type of repeated pattern can help predict relationships between non-neighbouring areas.

We incorporate a map of cortical types suggested by García-Cabezas, Hacker, and Zikopoulos [119]. Regions from von Economo and Koskinas' micrographs and their own curated histology samples for frontal regions were segregated to cortical types depending on laminar complexity. We used their suggested criteria for describing cortical layers and divided the cortex into allocortex, agranular, dysgranular, eulaminate 1 – 3, and Koniocortex (Figure 4.4). These cortical types were used as a ground-truth parcellation for our study. Outside and adjacent to the allocortex was an agranular cortical type which lacks Layer IV. Next, dysgranular areas have a more developed, although still thin, layer IV. Beyond agranular and dysgranular types are eulaminate cortices, which have around 5 – 6 distinct layers and sublayers, including a reasonably distinct layer IV. The koniocortex was observed as having the thickest layer IV among all other cortical types and primary sensory cortices were usually members of this type.

4.2.4 Design of study

We pre-processed staining intensity profiles assigned for each brain vertex before computing laminar metrics. We first regressed out the linear trend across the cortical depth by fitting a general linear model. This was to enable easier detection of local maxima within a cortical profile. Local maxima in the cortical profile could potentially indicate a change in layer property, but staining intensities across the cortical depth in many regions did not always change abruptly between all existing layers. Usually, changes in staining intensities could be detected between layers with markedly different features.

4.2.4.1 BigBrain proxy measures of laminar features

Von Economo and Koskinas first assigned cortical areas with cortical types based on how 'granular' cell bodies appear, and 5 cortical types (agranular, frontal, parietal, polar, granulous (koniocortex)) were designated so that laminar features of adjacent cortical types changed gradually.

We adapted laminar properties suggested by García-Cabezas, Hacker, and Zikopoulos [119] from their study that endeavoured to define cortical types based on a systematic observation on histology. Those features were converted into proxy metrics that could be applied to BB staining intensity profiles (Table 4.1). BB data used Merker staining, which is a silver impregnation technique for cell body and dendritic shaft staining [115], and is similar to Nissl staining used in those studies for determining cortical types.

The following descriptions (i – iv) explain the specific criteria for distinguishing laminar properties by Garcia-Carbezas and colleagues including i. Thickness of layer IV, ii. Prominent layer group between superficial and deeper layers, iii. Layer with largest pyramidal cells, iv. Sublayers between layers V, VI, and v. How sharp layer I/II boundaries are. Each description of Feature 1 – 5 describe the actual adapted quantity being measured in BB staining intensity profiles.

i. Feature 1 – Layer IV peak intensity

Thickness of layer IV can vary depending on how granular (i.e., how grainy neuron cell bodies appear) or densely populated its constituent cells are.

From BB staining intensity profiles, a peak near the mid-cortical depth of the cortex can be observed (Figure 4.2), which corresponds to an acute change in staining intensity. We attempt to quantify thickness of layer IV without employing more precise detection of layers. To do so, we searched for a maximum peak around mid-cortical depth using a peak detection algorithm in SciPy (`https://docs.scipy.org/doc/scipy/reference/generated/scipy.signal.find_peaks.html`) that searches for local maxima through comparison with neighbouring data points. This was to avoid searching for the exact location of layer IV that could vary in different cortical areas. Following peak detection, a variance around the detected peak was calculated as the first proxy metric.

ii. Feature 2 – Prominent layer group between upper and lower layers

For the next feature, Garcia-Cabezas and colleagues determined which of upper

and lower layers were more differentiated. Upper layers were those above layer IV (i.e. layers II – III as layer I does not include many cell bodies). Layers below layer IV were lower layers (layers V - VI). 'Prominence' of a layer group was defined by how densely packed constituent cells of a layer group were or by overall sizes of cell bodies in their study.

We designated the range for upper layers from the peak near layer I/II border to layer IV peak. The lower layer group ranged from layer IV peak to the end of the profile above the WM surface. Again we used peak detection to identify the layer I/II border and peak within a fixed range near mid-cortical depth. We measured prominence between the two layer groups by calculating the difference of each group mean, which was: (mean intensity of lower layer group) – (mean intensity of upper layer group).

iii. Feature 3 – Location of largest pyramidal cells

Identifying the layer that contains largest pyramidal cells was a way of distinguishing cortical types. However, in our study we could not transfer this feature into a computable metric as we were using only BB staining intensity profiles rather than observing cell body sizes from the original histological sections.

iv. Feature 4 – Layer V/VI differentiation

Between layers V and VI there are often subdivided layers, therefore, number of layer V sublayers can be an indication of a more differentiated layer V.

From the staining intensity profiles, we searched for a range that would include layer V sublayers and derived a metric from 1st order derivatives of the profile. These staining intensities were sampled at 50 equidistant surfaces so we used an approximate range from the 35^{th} and 49^{th} sample points include a layer V sublayer. This was done to include layers V – VI, while excluding middle depth layers (e.g. layers III - IV) when applying the same metric to all areas.

As for detection of layer V sublayers, we assumed that any abrupt change in intensity for significant layer V sublayers would appear as local maxima in the

profile. If a local maximum was found with the peak detection algorithm, an interval was set for a fixed width around that point. If a peak was not detected, the interval was set around the midpoint of the 35^{th} and 49^{th} samples (i.e around the 42^{nd} sampled point).

The final metric was computed by taking a sum of absolute values of 1^{st}order derivatives in the range set accordingly to the above steps. We also used a second metric which was variance of 1^{st}order derivatives from the layer V sublayer interval.

v. Feature 5 – Layer I/II boundary sharpness

The final feature determined cortical types by sharpness of laminar boundary between layers I – II and were divided into three categories: irregular, slightly irregular, and sharp. They observed that layer I/II boundaries for regions near the allocortex were more irregular, while eulaminate cortical types such as sensory cortical areas had sharper boundaries.

The metric for this feature was the sum of absolute values of 1^{st} order derivatives near the layer I/II transition peak.

4.2.5 Analysis 1

4.2.5.1 Clustering BB laminar features

We conducted k-means clustering on laminar features 1, 2, 4, and 5 (excluding feature 3) to test whether similar cortical types determined in the Garcia-Cabezas study could be derived from our laminar features from BB staining intensities. Clustering was performed on concatenated feature data consisting of 6 laminar features **(A)** – **(F)** (in Table 4.1 BB proxy measures) for each hemisphere (data matrix size: 6 features x number of vertices, for each hemisphere). We chose a total number of 7 clusters (N = 7) to generate clusters that were neither too fine-grained nor too broad. Then, clusters were assigned with numerical labels in increasing order (1, 2, ... 7) and were compared against their published cortical types to examine if they corresponded to the allocortex, agranular, dysgranular, eulaminate I - III, or konicortex (Figure 4.4).

Feature No.	Laminar feature	Feature specification	BB proxy measures
1	Layer IV thickness	Is layer IV absent, or if exists, is it prominent or thin?	Find interval around layer IV peak **(A)** sum or **(B)** variance of staining intensities in interval
2	Prominent layer group (superficial vs. deep layers)	Are superficial layers II – III more or less prominent than deeper layers V – VI?	**(C)**(Mean of layer V/VI range) - (Mean of layer II/III range)
3	Largest pyramidal cells	Which layers are largest pyramids located in?	N/A
4	Layer V/VI differentiation	Are there sublayers between layer V and VI?	**(D)** Sum, **(E)** variance of 1st order derivatives in layer V/VI range
5	Layer I/II boundary sharpness	Is layer I/II boundary sharp or irregular?	**(F)** Sum of 1st order derivatives near layer I/II transition peak

Table 4.1: Laminar features were suggested in a past study [119] to define cortical types for human cortical areas. Here, we described features they used and describe how each feature was adapted for the current work so that we could compute proxy laminar measures from BB staining intensity profiles.

We also assessed whether clustering based on our hand-crafted laminar feature proxy measures was more effective than applying K-means clustering directly on detrended BB staining intensities profiles. We applied Neighbourhood Components Analysis (NCA) on the laminar feature matrix (vertices $\times$ number of features) and detrended BB staining intensity profiles (vertices $\times$ number of staining intensity samples). Also, to minimise training time in NCA we used a random subset (10%) of the data which was the size of all brain vertices (14k).

4.2.5.2 Comparing new cluster labels and cortical types from literature

A way to validate new clusters formed from BB laminar features was to compare them against cortical types already published in past literature. We first examined the overlap between the two labelling systems.

We adapted the published cortical types from results shown in ([119], tbls. 4–7) to a 3D surface parcellation that was not given in the publication. All cortical areas from von Economo and Koskinas' micrograph images were assigned a cortical type. Then, we transferred those cortical types to an open source Von Economo and Koskinas (VEK) parcellation (`http://www.dutchconnectomelab.nl/economo/`) that was released as a freesurfer compatible 3D surface file [122]. However, in the process of converting cortical type labels into numerical class labels, some modifications were inevitable due to either (1) multiple cortical types being assigned to an area or (2) a slight difference in cortical area boundaries in the 3D parcellation that became the template in which labels were transferred to. Examples of the first case are marked with asterisks in Table 4.2.

For instance, FJK and FLMN were classified as Agranular and Dysgranular types and PE, PH, TA, and TE each were a mixture of types Eulaminate I – III, hence, those regions were assigned a single numerical label as shown in the last column of Table 4.2. Also, there were discrepancies between region boundaries of same areas in the 3D VEK parcellation adapted by Scholtens et al. [122] and VEK regions demonstrated by Garcia-Cabezas and colleagues in 2D (Figure 4.4).

Finally, to inspect the published cluster label that overlaps most with our new clusters, we calculated the dice similarity index. Then, we reordered our cluster labels so a corresponding cortical type numerical label was assigned. Consequently, we reassigned our new clusters with a cortical type label that overlapped most with a cluster. Ultimately, we aim to compare against the Garcia-Cabezas cortical types converted into numbered labels (last column in Table 4.2).

VEK Parcel No. [122]	Region name [119]	Cortical type [119]	7 class labels
0	Unknown	N/A	0
1	Corpus callosum	N/A	0
2	FA	Eu-III	6
3	FB *	Eu-I, Eu-II	5
4	FC	Eu-I	4
5	FCBm	Eu-II	5
6	FD	Eu-II	5
7	FDΔ	Eu-II	5
8	FDT	Eu-III	6
9	FE	Eu-II	5
10	FF *	Dys, Eu-I	4
11	FG	Eu-I	4
12	FH	Eu-I	4
13	FJK *	Ag-Dys	3
14	FLMN *	Ag-Dys	3
15	HA	Allocortex	1
16	HB	Allocortex	1
17	HC	Allocortex	1
18	IA	Dys	3
19	IB	Eu-I	4
20	LA1	Ag	2
21	LA2	Ag	2
22	LC1	Eu-I	4
23	LC2	Dys	3
24	LC3	Ag	2
25	LD	Ag	2
26	LE *	Ag-Dys	3
27	OA	Eu-III	6
28	OB	Eu-III	6
29	OC	Koniocortex	7
30	PA	Eu-II	5
31	PB	Koniocortex	7
32	PC	Eu-III	6
33	PD	Eu-II	5
34	PE *	Eu-II, Eu-III	5
35	PF *	Eu-I, Eu-II	5
36	PG	Eu-I	4
37	PH *	Eu-II, Eu-I, Eu-III	5
38	TA *	Eu-II, Eu-III, Eu-I	6
39	TB	Eu-III	6
40	TC	Koniocortex	7
41	TD	Koniocortex	7
42	TE *	Eu-III, Eu-I	4
43	TF	Eu-I	4
44	TG	Dys	3

Table 4.2: In García-Cabezas, Hacker, and Zikopoulos [119] a single region was often assigned multiple cortical types such as the highlighted cortical areas marked in this table. Therefore, it was necessary to adapt those cortical type labels into single numerical labels between 1-7 as shown in the final column of the table. (Abbreviations in table; Agranular: Ag, Dysgranular: Dys, Eulaminate: Eu)

4.2.6 Analysis 2 and 3

4.2.6.1 Analysis 2 – High-resolution MRI feature to cortical type prediction

The data used in analysis 2 was from a high-resolution study conducted on a single human brain scanned using the MGH-USC 3T Connectom scanner [118]. The relevant imaging parameters were: voxel resolution 0.76 x 0.76 x 0.76 (mm), 2 shells with b-values = 1000, 2500s/mm^2 (number of directions=420, 840 respectively). Additional T1w and T2w images were also acquired at a resolution of 0.7mm isotropic from which we calculated a ratio map as a proxy for myelination [37]. At the time of writing, this was the highest resolution in vivo human diffusion MRI scan ever acquired.

In analysis 2, we aim to compare the high-resolution MRI data with cortical types. First, diffusion measures (FA, MO, $L_\perp$, $L_\parallel$, f_1, f_2, f_3) and a measure of myelination (T1w/T2w). These features were first smoothed with a kernel (sigma = 8mm) and sampled for 6 equidistant layers (layers I – VI). Now that there were 8 features for 6 layers for all vertices (∼164k), we computed summary metrics instead of using features of all layers. The 3 summary measures calculated for 8 MRI features were: 1. mean value, 2. variance between layers III-VI (i.e. middle-deep layers), and 3. layer III (middle layer) feature value. These summary metrics were concatenated into a feature data matrix (vertices × (8 × 3 features)).

Finally, we used 5-fold cross-validation to train the above 24 MRI laminar features for the left hemisphere to predict cluster labels of the right hemisphere. Several supervised learning methods including linear support vector machine (SVM), Decision Trees (DT), and Random Forests regression (RF) were explored in this analysis. Then the predicted right hemisphere labels were compared against cluster labels (7 classes) to inspect prediction accuracy. Also, feature importance was computed to inspect which MRI derived features were involved in prediction of cortical type labels.

4.2.6.2 Analysis 3

The purpose of analysis 3 was to link the MRI metrics directly to the BB cortical profiles. We conducted a simple linear regression on all mean profiles of our diffusion features extracted and observed how well the mean BB profile would be predicted. By comparing absolute values of coefficients from GLM fitting when including and excluding the measure of myelination content (T1w/T2w) as regressors, we assessed which features contributed to prediction of BB mean profiles.

4.3 Results

Now, we show results according to the separate analysis approaches that were referred earlier as Analysis 1 - 3.

4.3.1 Analysis 1: BB to Cortical types

4.3.1.1 K-means clustering of BB laminar feature metrics

In Figure 4.5, we first show the 6 laminar feature metrics computed for all vertices (A-F Table 4.1). From each metric displayed on the cortical surface, we also demonstrate cortical profiles representative of the 25^{th} and 75^{th} percentiles to qualitatively appreciate the difference in the profiles.

K-means clustering was conducted on the above summary features extracted from BB staining intensities. Then, clusters were labelled in increasing order so that there would be a gradual increase in laminar complexity similar to the how cortical type definitions change across the cortex [121]. Figure 4.6 demonstrates our clusters with labels. The innermost cluster was designated cluster 1, so that it roughly corresponds with the cortical type 'allocortex' which was also included in this figure for easier comparison. From inner towards outer areas, cluster labels numbers were assigned so that they increase smoothly between adjacent clusters until reaching a final cluster that includes the primary visual cortex. Cortical types including the visual cortex were those with more differentiated laminar properties,

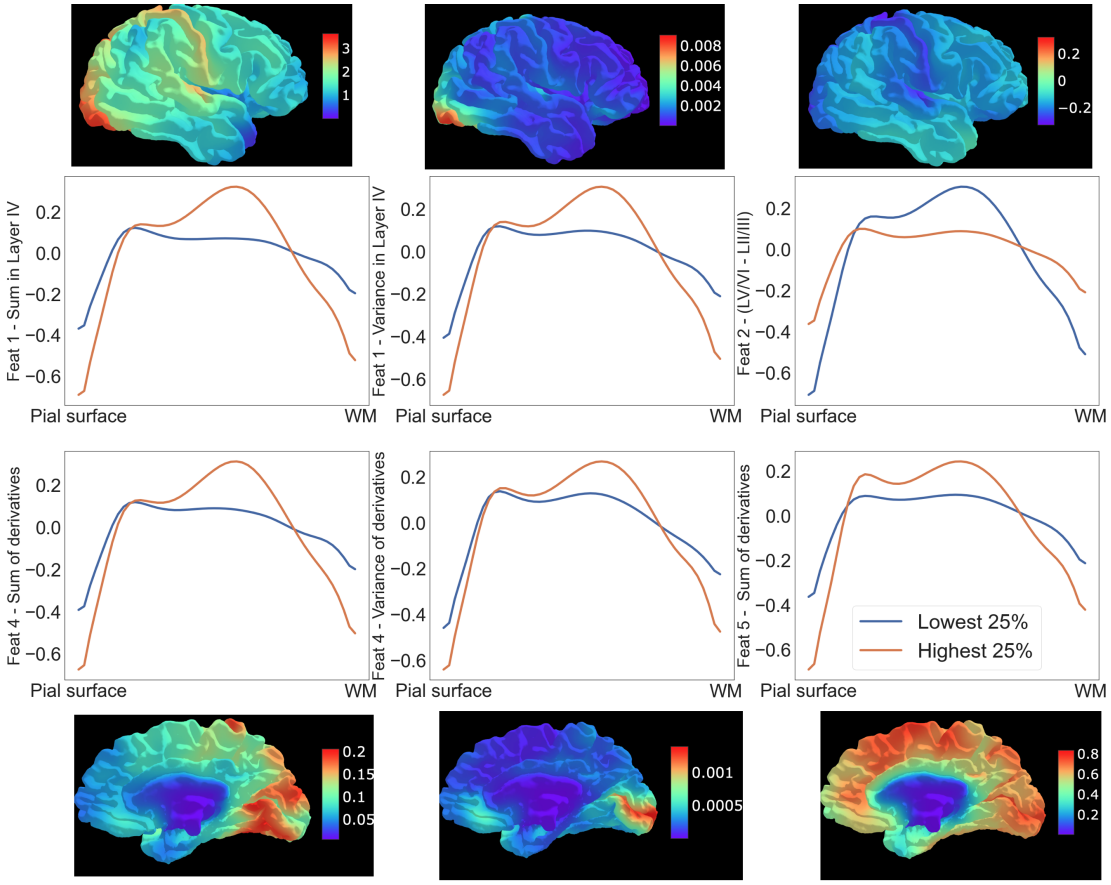


Figure 4.5: Each subplot shows the mean of detrended BB profiles in the lowest and highest quartile for 6 feature metrics. Also, distributions of feature measures are shown on lateral (3 features in first row subplots) or medial brain surfaces (3 features in second row subplots).

and clusters 6, 7 roughly corresponded to those granular cortical types in terms of their spatial locations in the cortex.

We then reassess mean BB staining intensity profiles for our newly generated clusters in Figure 4.6. Here, we observed that cluster 7 mean BB staining intensity profile changes more dynamically in deeper layers compared to, e.g., cluster 1 that had less variation across the entire cortical depth.

As a separate sanity check on the validity of clusters, we also inspected whether previously known laminar properties of the motor cortex (M1 (primary motor cortex) / area FA (from von Economo and Koskinas)) could be identified in the cluster (cluster 3) that included M1. We inspected whether cluster 3 had staining intensities representative of cytoarchitectural properties known of M1 such as layer

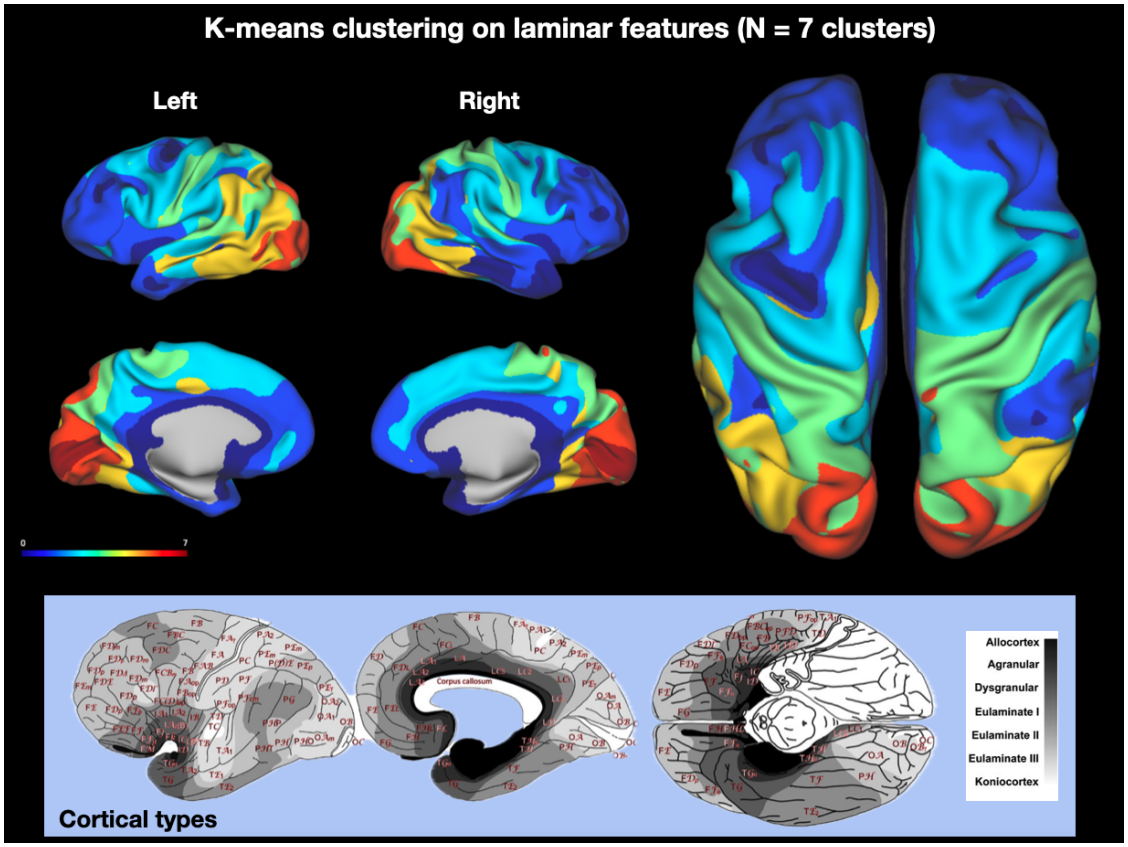


Figure 4.6: We conduct K-means clustering on laminar features concatenated in a data matrix with dimensions 6 features × number of vertices. The 7 output clusters were given numerical labels 1 - 7 ordered in a similar fashion to the spectrum of cortical types shown in the panel below [121]. Cluster 1 corresponds approximately to the allocortex (innermost cortical type) and clusters 6, 7 include the visual cortex which was classified as roughly Eulaminate III or Koniocortex.

V consisting of giant pyramidal cells or lacking a clear layer IV. In Figure 4.7, we can observe that cluster 3 mean staining intensity profile around mid-cortical depth was broad but with a less acute peak intensity than clusters 4 – 6. Whether this observation is a manifestation of the difficulty of comparing staining intensities that could vary according to two different laminar properties, cell size and cell density, would require further investigation.

Also, it was not clear whether a lacking layer IV of M1 could be identified when comparing these mean staining intensity profiles (Figure 4.7) and mean feature value of clusters (Figure 4.8) alone without referring back to the original histology.

Finally, we tested our approach of clustering BB derived laminar measures against directly clustering raw staining intensity profiles. We used NCA to visualise

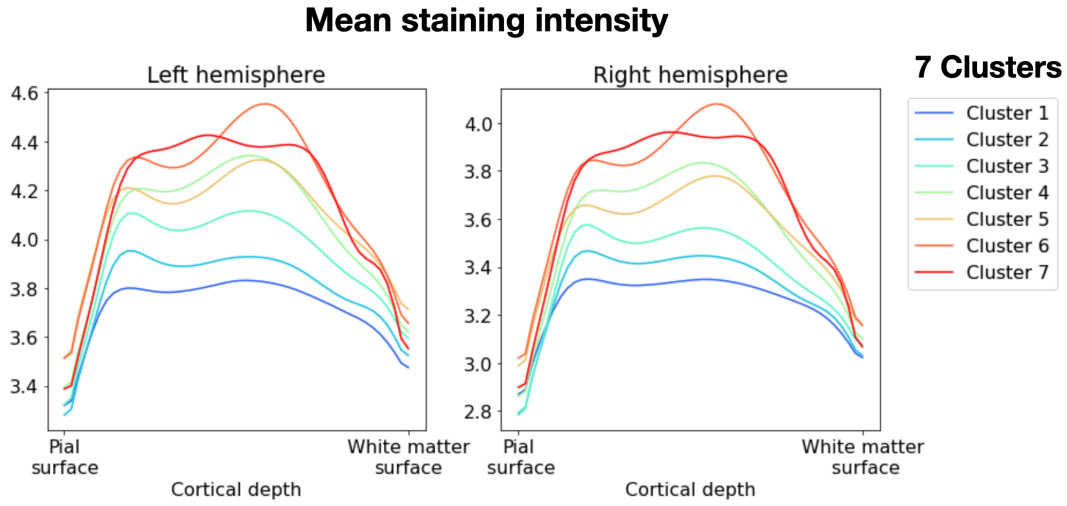


Figure 4.7: Mean detrended BB staining intensity profiles computed for clusters 1-7 were symmetrical in both hemispheres. Cluster 7 located near the primary visual cortex is more differentiated in deeper layers near the WM surface.

on a 2D embedding our summary laminar features and detrended BB staining intensity profiles. Figure 4.9 illustrates how neither method produced perfectly differentiable clusters. In contrast to clustering directly on BB staining intensity profiles (sampled at 50 points across the cortical depth), our proxy measures helped classify, for instance, cluster 7 that had more discernible laminar features. Overall, proxy measures were not sensitive enough to cleanly delineate areas other than the occipital cortex. In Appendix Figures A.4and A.5 we add two alternate scenarios of K-means clustering. Figure A.4 was when the number of laminar features for the input data was greater than the number of clusters (n=4), and Figure A.5 shows increasing cluster numbers (n=10) did not show notable clusters that differed from our original scheme of using 7 clusters.

4.3.1.2 New cluster label comparison with published cortical types

We analytically compared our new clusters against previously published cortical types to check how similar the two schemes were. We searched for a cortical type that maximally overlapped with our clusters according to dice similarity indices, while avoiding repeated assignment of a cortical type to different clusters. In Figure 4.10b, we show our clusters reassigned with cortical type labels defined in Figure

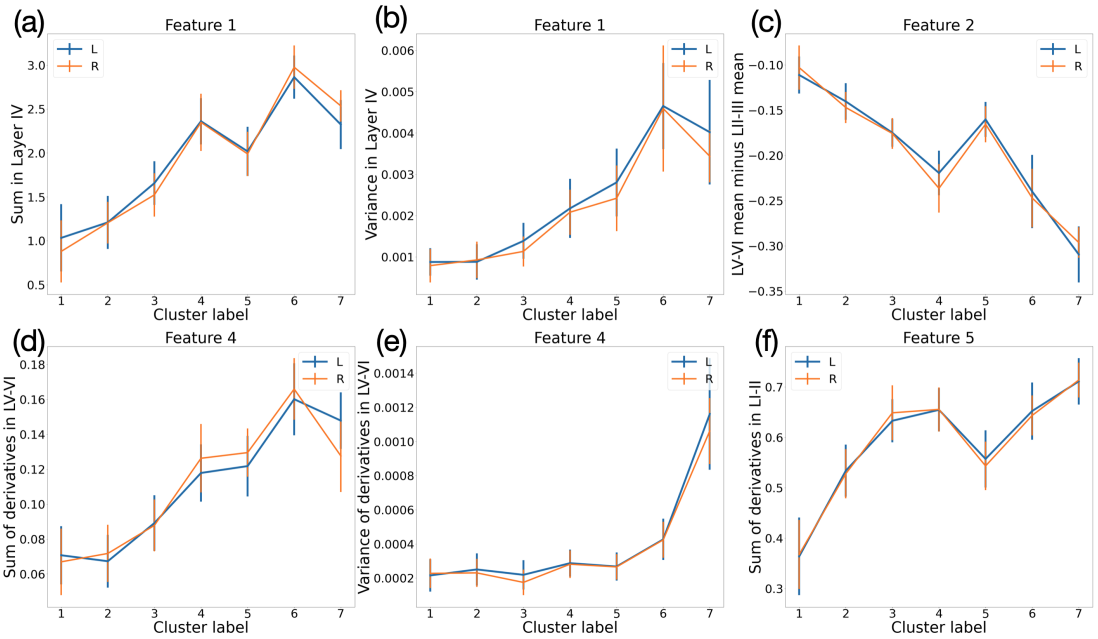


Figure 4.8: Mean of each of the 6 selected features for clusters 1-7 are shown in figures (a)-(f). Features 1 (a,b), 2 (c), 4 (d,e), and 5 (f) are described in Table 4.1 and were computed from BB staining intensity profiles. Feature 1 in (a), (b) express sum and variance of values within the layer IV range, respectively. Feature 2 in (c) were calculated as (mean values in layer V-VI range) – (mean values in layer II-II range). Feature 4 in (d)-(e) were each sum and variance of 1st order derivatives taken from detrended intensity profiles ranging from Layer V to the innermost layer VI. Feature 5 in (f) was sum of 1st order derivatives from a range around the Layer I/II border.

4.10a. Overall, dice similarity indices between our 7 clusters and cortical types were not significant, with a mean overlapping percentage of 28% and 32% for left and right hemispheres, respectively. Thus, our new cluster shapes and borders did not precisely output an identical cortical type parcellation as the previous study.

4.3.2 Analysis 2 – High-resolution MRI to Cortical types

4.3.2.1 Supervised learning of MRI features to predict cortical types

In this analysis, we examined whether dMRI features included useful laminar information that allows the cortex to be distinguished into cortical types.

We applied supervised learning techniques (Linear SVM, DT, RF) to train an estimator that uses MRI summary metrics as inputs to classify cluster (i.e. our new cluster labels) or cortical type labels (i.e. VEK parcellation cortical type labels) of the left hemisphere. We then tested the estimator on unseen right hemisphere

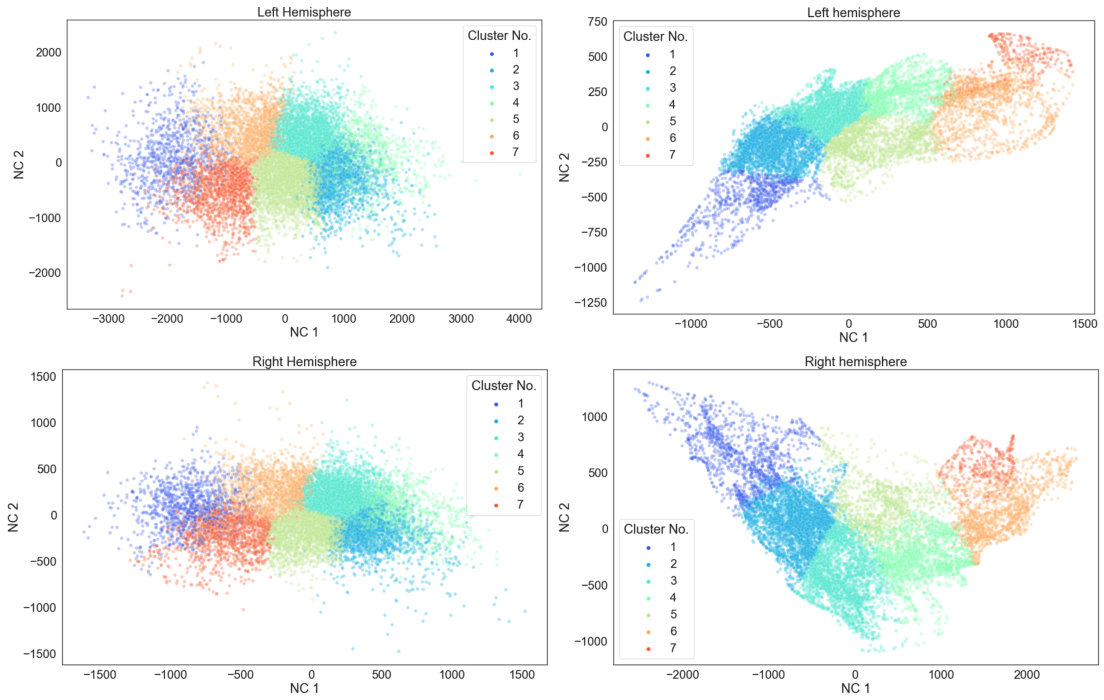


Figure 4.9: Plots on the first column are NCA conducted on left and right hemisphere feature data matrices (vertices × 6 features). To understand the effect of computing specific laminar features, these plotswere compared against NCA performed on detrended BB profile data matrices (vertices × 50 equidistant depths) shown in the second column.

data. Right hemisphere label prediction results showed that cortical type labels transferred to a VEK parcellation (Figure 4.10a) were estimated with a higher accuracy of 64.4% than our cluster labels derived from BB laminar metrics (Figure 4.10b), which had 37.5% accuracy.

Next, we inspected which of the 24 MRI feature summary measures contributed to reasonable classification of cortical type labels with accuracy of 64.4% by computing relevant feature importance methods. Classification accuracy was 55.3% with linear SVM and 64.4% with RF. We determined whether accuracy differences between methods were statistically significant through paired t-tests between pairs of methods RF vs. linear SVM (p=0.00013), RF vs. DT (p=0.0004), linear SVM vs. DT (p=0.002). In Figure 4.11a, we specifically focused on top 8 features with highest feature importance in each method as highlighted in red boxes. RF prediction accuracy was greatest (64.4%) and middle layer myelination (T1w/T2w) had greatest feature importance followed by $L_{\|}$ variance in middle to deep layers. Focusing

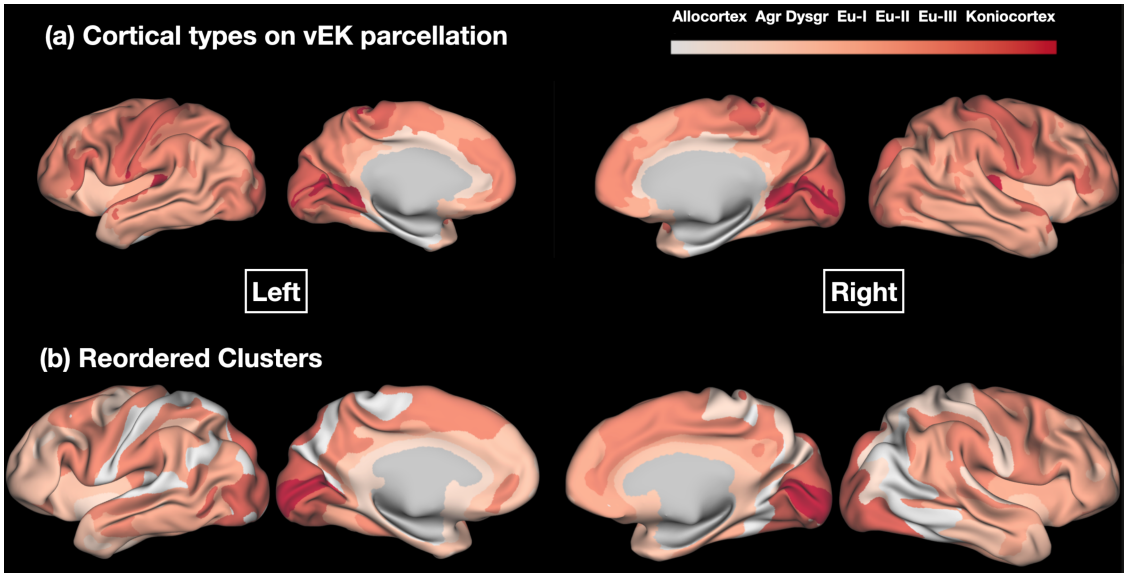
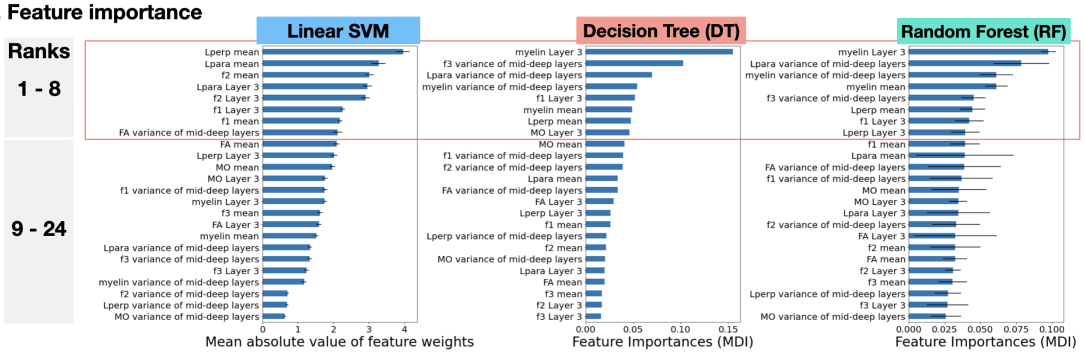


Figure 4.10: On a brain surface, we show 2 parcellations with 7 parcels. This was done in a similar layout to match the original parcellation shown in Figure 4.4 that demonstrates a gradient of cortical type changes across the cortex. As a result, we can compare two labelling systems through a similar perspective. (a) Here, we firstly show 7 cortical type labels from a histology study [119] that was simplified and overwritten on a VEK parcellation [122]. (b) Then, we show our cluster labels that were reordered based on dice similarity indices with cortical type labels in (a).

only on which MRI specific measure a feature defines and without considering the summary metric used to define a feature metric, higher to lower contribution of the features were in the order of myelination, $L_\perp$, and $L_\parallel/f_1/f_2$. Although with an overall lower accuracy of 50% for DT (Figure 4.11c), middle layer variance of middle to deep layers and mean myelination were all identified as one of top 8 feature importances. When including myelination related features as in DT and RF, the accuracy of cortical type labels 5 and 6 was greatest (>60%) as shown in confusion matrices in Figure 4.11b. On the other hand, linear SVM with a moderate prediction accuracy of 55.3% feature importances were greatest for mean $L_\perp$ followed by mean $L_\parallel$ and mean f_2. In a likewise manner with the analysis for RF disregarding which of the 3 summary metrics was computed for features $L_\parallel/f1/f_2$ then $L_\perp$/FA were identified as most frequently appearing features within top 8 feature importances. In linear SVM, cortical type labels 1 and 2 were the more accurately predicted labels ($>$ 60%) (Figure 4.11b).

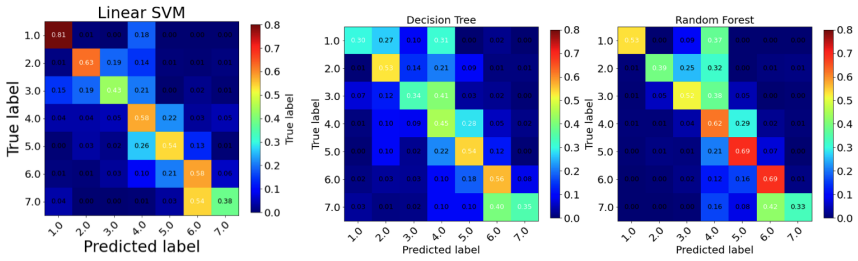
a. Feature importance



Ranks 1 - 8

	Mean			Variance of layers III - VI			Layer III			SVM	DT	RF	Total
FA				8						1	0	0	1
MO								8		0	1	0	1
Lpara	2				3	2	4			2	1	1	4
Lperp	1	7	6						8	1	1	2	4
Myelin		6	4		4	3		1	1	0	3	3	6
f1	7						6	5	7	2	1	1	4
f2	3						5			2	0	0	2
f3					2	5				0	1	1	2

b. Confusion Matrix for each cortical type parcel



c. Accuracy of each SL method

	Linear SVM	Decision Tree	Random Forest
Accuracy (%)	55.3	50.0	64.4

Figure 4.11: Supervised learning on high-resolution MRI features to predict VEK parcellation (published) cortical types. (a) We computed quantities representing feature importance to identify which features contribute to the different prediction results of cortical type labels. To clarify, these labels were not K-means cluster type labels derived in our main analysis. (b) In the confusion matrices for each prediction method, we compute mean prediction accuracy of each cortical type label, separately. When we used linear SVM, cortical types 1 (allocortex) and 2 (agranular) had accuracies above 60%. Similarly, but for different cortical types when using RFs cortical types 4 (Eu-I), 5 (Eu-II), and 6 (Eu-III) had accuracies above 60%. (c) Linear SVM, DT, and RF each showed mean accuracy of 55.3%, 50.0%, and 64.4% in predicting right hemisphere cortical type labels.

Finally, we assessed predicted cluster label properties through observation of diffusion MRI measure mean profiles (Figure 4.12). Most features excluding f_2 and f_3 showed distinct profiles between clusters.

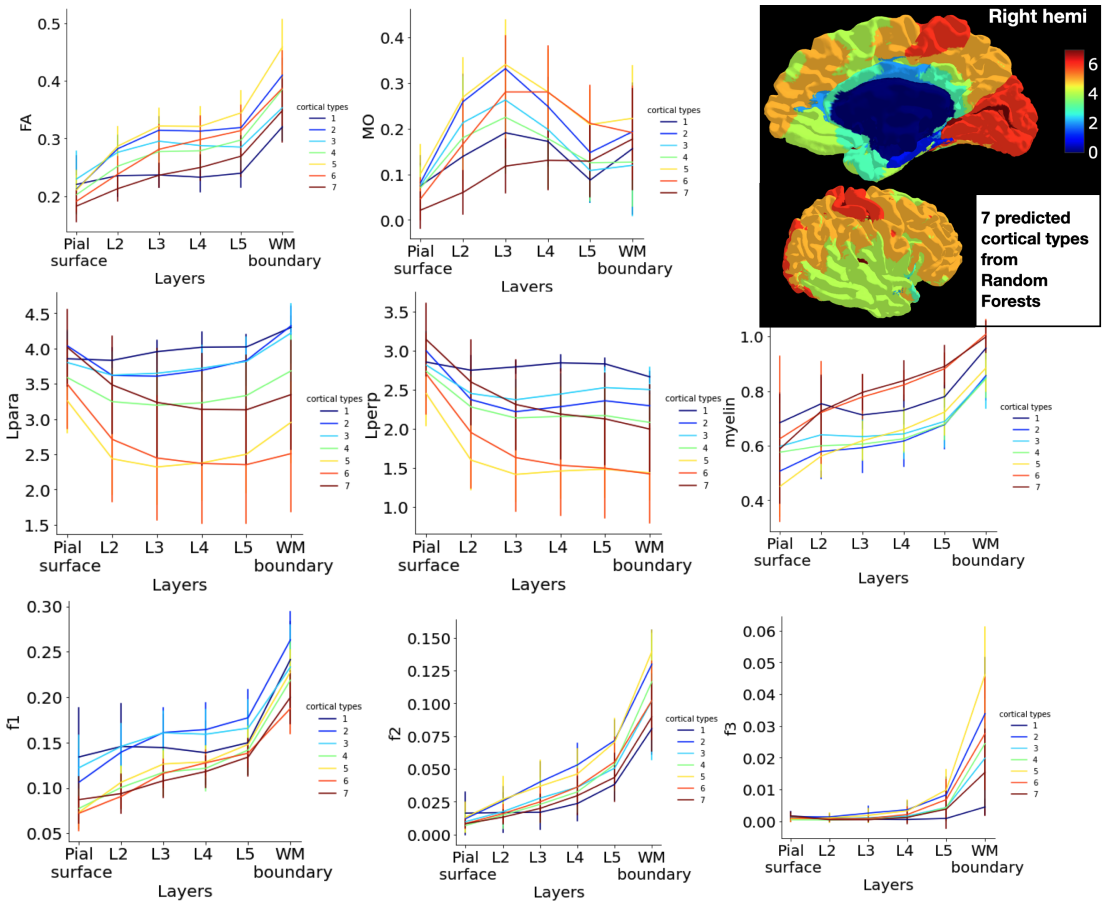


Figure 4.12: Right hemisphere cortical type labels predicted from left hemisphere MRI feature summary metrics using supervised learning with RFs. Mean MRI measures for each predicted cortical type label was computed for all 8 diffusion metrics FA, MO, $L_{\parallel}$, $L_{\perp}$, myelination (T1w/T2w), f_1, f_2, f_3.

4.3.3 Analysis 3 – BB to High resolution MRI

4.3.3.1 GLM fitting

In this final analysis, we asked whether BigBrain cortical profiles could be predicted directly from the MRI metrics sampled across cortical layers. We used a simple linear regression for this analysis (Figure 4.13). We found a good fit between the model prediction and the data and we found that T1w/T2w (i.e., "myelination") was the most important predictor followed by the volume fraction of the dominant fibre

orientation ($f1$). Future work will examine whether this model learned on the basis of the average cortical profile can be used to predict vertex-wise profiles across different brain areas and whether the same model can be extended to work on in vivo data.

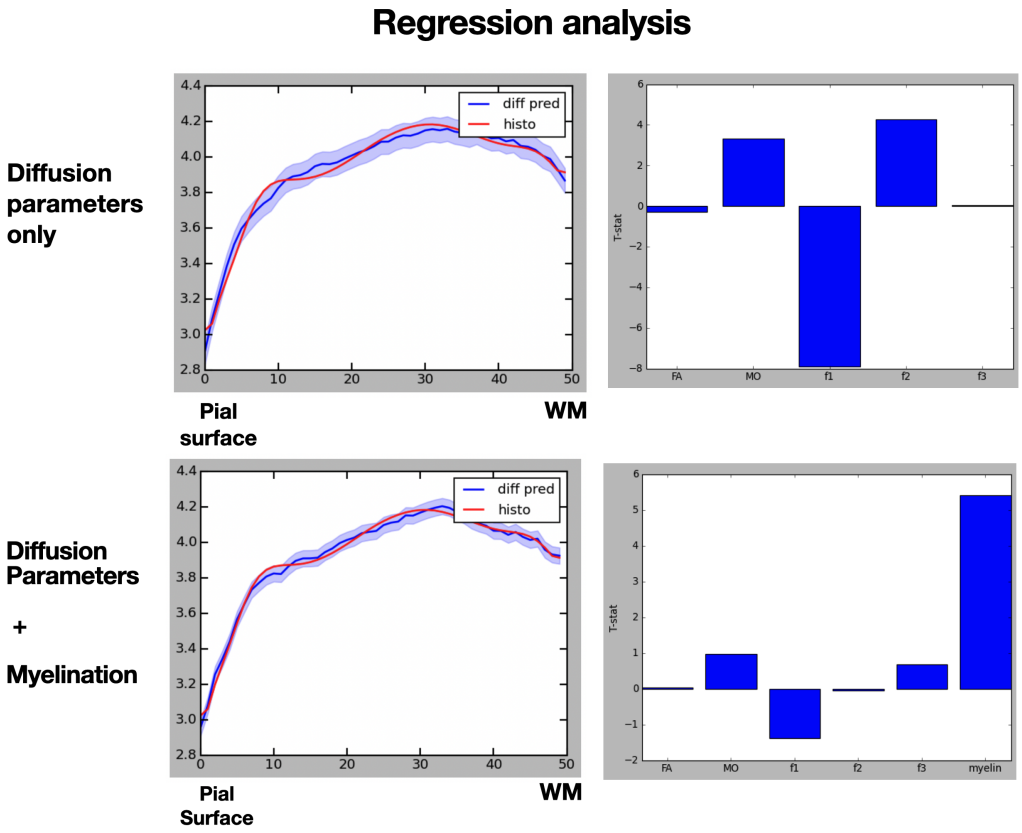


Figure 4.13: We used a GLM to fit diffusion parameters to the BB staining intensities. In analysis 3 we start by asking the question whether we can use mean diffusion profiles to predict mean BB staining intensity profiles. We compared two different scenarios with or without mean myelination profiles. When including diffusion parameters, f_1 contributed most to the fitting. However, with all diffusion and myelination measures, myelination contributed most to the fitting.

4.4 Discussion

In this chapter, we examined whether MRI measures can be used to assess laminar aspects of cortical architecture and to do so, we used cortical types as a medium to relate BB staining intensities with high-resolution MRI measures. Cortical types have been used to distinguish cortical areas based on how complexity of cortical layers changes across regions. There are variable definitions of cortical types in

several other studies [123, 124]. Cortical type definitions can differ in the total number of types defined as long as they adhere to a topological rule in which laminar complexity changes gradually between neighbouring areas from simpler to more elaborate cortical types [119]. Moreover, previous studies have suggested that regions with similar cortical types could be connected to other regions of similar cortical types through the same laminar connection patterns (i.e., same FFW, FB, lateral connection pattern) [123]. This is plausible in theory, as cortical types are determined by laminar properties of regions. However, the potential of applying such concepts in the human cortex was less explored until more recent advances that made possible data sets such as BB histology or high-resolution dMRI at 0.7 mm. Hence in this study, we examined what laminar features could be extracted from the two data sets of interest. We used K-means clustering to examine whether clusters comparable to cortical types could be derived and then tested whether microstructural detail accessed from in vivo dMRI could explain laminar properties observed in the cluster analysis.

4.4.1 Analysis 1 – Clustering BB histology as cortical types

BB histology data can be accessed on the same space with MRI, which allows an insight into microstructural organisation of the brain beyond what MRI studies can provide. Especially, we envision that platforms that enable cross-modality investigations [116] can reveal more on the less well understood areas that regulate higher-order cognition. Other studies have started to show how cytoarchitectural organisations as examined by histology, in fact, underlie cortico-cortical connectivity between brain regions [125].

We first labelled our new clusters so that less characteristic layer properties were closer to the allocortex (part of the limbic system) and the more complex areas were near occipital sensory areas. As a result, clusters had a topological arrangement in which clusters 1 – 2 were adjacent to the allocortex, clusters 3 – 5 were a mixture of frontal and parietal areas, and clusters 6 - 7 were occipital areas (Figure 4.6). We compared these clusters against an adaptation of previously published cortical

types from García-Cabezas, Hacker, and Zikopoulos [119] displayed on a 3D VEK parcellation from [122] to validate our clusters and next address several observations and shortcomings that were encountered.

4.4.1.1 Observations

In this work, we created signatures of laminar differentiation from the staining intensity profiles which was an initial step towards reducing observer-dependent biases that histology studies commonly face. With more automated methods to detect laminar details, we can move towards establishing a framework that can help us observe and understand cytoarchitecture in the entire cortex more systematically. Here, we test this notion on BB staining intensities released in [116]. These staining intensity profiles were sampled from a 100μm BB volume at 50 equivolumetric surfaces [14] while correcting for varying thicknesses at cortical curvatures [117].

On the one hand, we were only able to obtain clusters that had a moderate mean overlap with cortical types in 28% and 32% of vertices in left and right hemispheres. Although we could not replicate those cortical types with our clusters, we made several observations that could help understand our new clusters. We also reiterate that the aim of this study was not to develop laminar metrics that could be applied to derive same cortical types but we rather refer to the previously published cortical types as a reference to compare against our clusters.

Through Figures 4.7 - 4.8, we can assess properties of clusters shown in Figure 4.6 with respect to suggested laminar properties of published cortical types outlined in García-Cabezas, Hacker, and Zikopoulos [119]. In all BB laminar features in Figure 4.8 subplots excluding feature 2 'prominent layer group' (Figure 4.8c), cluster 1 mean feature value was the lowest. Then there is a smoothly increasing trend towards cluster 7 which results in cluster 6 having maximum mean metric value for features 1 and 4 (Figures 4.8a, b, d) and cluster 7 having greatest mean value in features 4 and 5 (Figures 4.8e , f). In other words, layer IV degree of complexity measured by sum (Figure 4.8a) or variance (Figure 4.8b) and differentiation of lower layer group layers V-VI computed as a sum of derivatives (Figure 4.8d) was

greatest in cluster 6. Cluster 7 in appearance resembled a smaller subsection of the koniocortex from Garcia-Cabezas and colleagues' that was the cortical type with greatest laminar differentiation and that included the primary visual cortex. Lower layer group complexity computed as variance of derivatives (Figure 4.8e) and layer I/II sharpness (Figure 4.8f) was greatest in cluster 7.

After producing our new clusters, we returned to Figure 4.7 to re-evaluate whether our feature metrics were representative of the laminar property measured. Specifically, in cluster 7 mean staining intensity profile, we observed that there is no local minimum point but a local maximum that is almost a point with maximum value among all points in that cluster's profile. In clusters 2 – 6, there was a local minimum point between the layer I/II border and lower layer range peak. Clusters 1- 2 mean staining intensities had subtle variation across the entire cortical depth hence had similar local maxima and minima values that were not easy to disambiguate visually. We can postulate that cluster 7 has a mid-layer with greater intensities that are visible compared to other clusters that rarely had a maximum value in a corresponding location.

4.4.1.2 Limitations

We used NCA to visualise in 2D output clusters produced using our proxy laminar feature measures and compared against clustering performed on staining intensity profiles of all brain vertices. We were able to visualise clearer NCA clusters with the extracted laminar metrics. Thus, these laminar feature measures we curated and implemented do seem to divide the cortex into clusters more effectively. However, that does not fully justify whether such handpicked but easily computable laminar feature metrics provide us with clusters that have clear borders in line with past published results on cortical types. Then why were our feature measures not perfect? We address possible sources of error in this study. The first source of errors resides in the approach of adapting hand-written laminar properties to apply to staining intensity profiles. Then another problem relates to challenges that remain in the

field such as differing perspectives on laminar architecture of cortical areas due to observer dependent interpretations of the complex laminar anatomy.

4.4.1.2.1 Limitations of inspecting staining intensity profiles

We computed several laminar feature metrics with respect to the location within a profile layer IV would exist, but such approach could be problematic in certain situations.

In our study, the change of feature 1 values between clusters were essentially different from how layer IV thickness changes between cortical types from the literature. As we were computing different quantities as result of adapting the suggested laminar properties, layer IV in the published work inspects absence or thickness of layer IV if visible [119]. However, for feature 1, we did not precisely identify where layer IV borders initiate and end but approximate a range where layer IV would potentially exist and then compute mean of sum or derivatives. Consequently, the thickest layer in koniocortex was not detectable in cluster 7.

The motor cortex is well known for its absence of layer IV, therefore we inspected area FA that included roughly motor and premotor areas [122] in the VEK parcellation against our clusters that included the motor cortex. Whereas FA was identified as cortical type of Eulaminate III in the reference literature, our results demonstrate that area FA was a mixture of clusters 4 and 5, which had moderate laminar differentiation. The possible discrepancy in properties of the motor cortex could be in virtue of computing feature metrics upon an assumption that layer IV existed rather than detecting the actual location or thickness of layer IV. This would be a problem if the motor cortex did not have a significant layer IV as frequently observed in the past.

There were also other cases in which interpretation of staining intensities was ambiguous as cell body sizes or population densities were not observable.

Feature 2 was calculated as difference between lower group mean and the upper group mean and intended to determine the more prominent layer group between the two. As can be seen in Figure 4.7c, the mean of this metric for all 7 clusters were

negative therefore according to the metric, upper layer groups seemed 'prominent' for all clusters. However, our interpretation on this feature was that in an analysis on staining intensities blinded to original histology, it is not possible to conclude which layer group was more prominent. The original study suggested prominence to be determined depending on how densely organised cells were or how large cell bodies were, but it is difficult to determine which one of these factors contributed to increase of staining intensities in different areas.

One other characteristic of cluster 7 was that although it was expected to have thickest layer IV as it overlaps with the koniocortex, staining intensity mean value around layer IV maximum peak was greater in cluster 6 (Figure 4.7). This could also mean that greater staining intensity value may not be representative of layer thickness and could be signatures of cell density that could be determined by the population or size of cell bodies. Again, the current study did not validate such postulations that would require examination of the original histology slices.

Last but not least, it could be that BB staining intensities did not come in sufficient resolution to reveal differences between layers that can allow clusters to be delineated similarly to existing cortical types.

Upon these observations, we suggest that quantification of laminar properties within staining intensities should be conducted with caution as their variation across the entire cortex is still complex and if one solely relies on staining intensities to conduct a microstructural analysis of the whole brain, the results could compromise the merits having access to such good quality high-resolution histology data.

4.4.1.2.2 Uncertainty on ground truth We also mention that among increasing or decreasing trends observed in each mean feature metric plot in Figure 4.8, cluster 5 deviated from the trend in several instances (Figures 4.8a, c, d, f). It could be that cluster 5 vertices have less meaningful laminar information, but still the mean staining intensity profile of cluster 5 is a unique cortical profile that can be distinguished from other clusters. Therefore, we would require further

studies to validate whether features of these 'new' clusters can help us understand cortical laminar microstructure.

Moreover, only when we have prior knowledge on certain areas like the motor cortex can we assume that it is likely that layer IV did not exist and that observed mid-peaks may be those in layer V rather than layer IV. However, generally we cannot rely on such prior anatomical knowledge as a means of crosschecking.

4.4.2 Analysis 2 and 3 – Diffusion MRI measures to Cortical types

4.4.2.1 Observations

Diffusion metrics such as FA, MO, f_1 and T1w/T2w from structural MRI data in Figure 4.12 demonstrate a sharp increase of mean profile near the GM/WM boundary, which was expected for those quantities. It could also be that increased parallel and perpendicular diffusivities ($L_{\parallel}$, $L_{\perp}$) near the pial surface could be occurring due to partial volume effects that was previously observed to affecting diffusivity near the cortical surface [126].

When observing depth-dependent changes of FA in our clusters, clusters closer to agranular cortical types agreed partially to a finding from past literature that agranular cortices tend to have greater FA than granular cortices [127]. However, clusters 6 – 7, which were the more granular clusters, had greater FA in mid-cortical depths (Figure 4.12 FA plot at cortical depths L4 L5) compared to the more agranular clusters (e.g., clusters 1, 3, 4). It could be that since clusters 6, 7 include the primary visual cortex, L4 or 5 that could include the Stria of Gennari may be related to greater FA compared to other clusters that are closer to agranular types. Clusters 6 and 7 also have higher myelination content in layers ranging from middle layer (L3) to deeper layers near WM surface. FA is often known as a diffusion measure that correlates with the amount of myelination and in literature, it has been suggested that granular cortices can have lower FA due to its complex underlying microstructural organisation that lowers the final value of the measure [127]. We anticipated inspecting other diffusion metrics such as fibre orientations f_1

– f_3 to help explain whether lower FA could be caused by complex fibre orientations in granular cortices. However, it was difficult to observe clear directionality with fibre orientations f_2, f_3 (Figure 4.12) and with the current fitting, we were not able to use f_1 – f_3 to identify significant existence of crossing fibres in any of the clusters.

Based on our observations, we did observe variable diffusion parameters and myelination among different parts of the GM, but notable profiles delineating a typical structure still pertained to visual areas.

4.4.2.2 Limitations

One limitation encountered during data pre-processing steps was that we had to eliminate mean diffusivity (MD) which occurred due to residual bias that were observed even after regressing out spatial coordinates.

Accounting for resolution differences of histology and dMRI, we used a bridging approach through 'cortical types.' We examined whether we could exploit cortical types to compare BB (histology) and diffusion MRI features simultaneously. By definition, cortical types have been used to classify cortical areas based on how differentiated layers are throughout the entire brain. As criteria for dividing cortical types rely on observations on histological properties with respect to laminar organisation such as thickness of layers IV, size of neuron cell bodies, sharpness of laminar borders between layers 1 and 2, existence or complexity of subdivision in layer V [7, 119], it is not straightforward of what information BB staining intensities include or what diffusion measures from high-resolution dMRI data contain. Therefore, we tested to see whether a. If we can perform clustering on BB profiles and retrieve similar cortical types that have been published (analysis 1) and b. If we could use SL machine learning approaches to predict cortical types from our diffusion and structural MRI features (analysis 2).

In analysis 2, we were not able to accurately predict our cluster labels with diffusion (FA, MO, $L_{\parallel}$, $L_{\perp}$, f_1, f_2, f_3) and structural (T1w/T2w) MRI feature summary measures. Our methods for extracting BB laminar features were simple algebraic expressions and these may not have sensitively captured our intended

feature properties thus could have been a source of error producing clusters deviating from published cortical types. Similarly, there could have been an issue with the quality of diffusion measures extracted, therefore other models for estimating fibre orientations in GM [99] are also worth exploring. However, the published cortical type labels were predicted reasonably (64% RF) with MRI metrics often better in cortical types that included parts of the visual cortex (69% Figure 4.11). In this high-resolution MRI dataset (0.76 μm), we were able to observe some signatures of greater laminar differentiation (via BB profiles) and increased myelination and FA (via diffusion MRI features) in parts of the visual cortex or agranular cortex and to a lesser degree and with remaining uncertainty in other frontal or parietal regions.

Finally, analysis 3 currently fitted GLM to mean diffusion and structural features, but on this scale, we could potentially overfit. What makes more sense is to try fitting GLMs that can be applied vertex-wise for the entire cortex. We would like to see if we can derive a single model that can be globally applied to predict BB cortical profiles from MRI features.

4.5 Conclusion

This chapter was an attempt at combining and relating whole brain high resolution histology, whole brain high resolution MRI, and classical anatomy. We generated hand-crafted summary metrics that capture laminar properties of cortex in the BB dataset. This allowed us to parcellate the brain into cortical types that qualitatively matched (although imperfectly) classical descriptions. We also demonstrated that cortical dMRI metrics differentiate cortical areas though there was poor agreement with cortical type definitions. Finally, we developed a preliminary model linking MRI to cell density profiles which could in the future be used to do so-called in vivo MR-based histology.

An ambitious goal in the field studying human cortical structure is to access microstructural detail non-invasively. We examined whether a high-resolution diffusion dataset could reveal insights into cortical GM, but not all diffusion metrics computed for the whole brain were robust enough to distinguish brain areas cleanly.

Similarly, staining intensity profiles from high-resolution microscopy data were not sensitive enough to inform us on laminar structures of the cortex. Through this analysis, we appreciate that even with cross-modality efforts using data of distinct resolutions, we need to extract and carefully curate useful features (such as our laminar features). Those laminar property metrics should be more numerous to benefit from application of data-driven analyses. Otherwise, metrics should be more sensitive to changes across the cortex and this would require better quality images.

5

Multilayer Dynamic Causal Modelling with Blood Draining

Contents

5.1 Introduction 101
5.2 Methods 104
5.2.1 Background on DCM for fMRI 105
5.2.2 Experiments for layer DCM study 111
5.3 Results 115
5.3.1 Simulation on multi-layer DCM for a single region . . . 116
5.3.2 Inter-laminar connectivity between 2 regions 121
5.3.3 Resting-state stochastic DCM analyses 124
5.4 Discussion 126
5.4.1 Considerations on model architecture settings 129
5.4.2 Analysis on data covariances 130
5.4.3 Comparison of simulated laminar BOLD against literature 131
5.4.4 Blood-draining laminar DCM implications for future laminar fMRI 132
5.5 Future plan 133
5.6 Summary 134

5.1 Introduction

BOLD signals measured in fMRI are temporally delayed and spatially blurred haemodynamic responses with respect to its true underlying neural activity. Typi-

cally, a haemodynamic response function (HRF) is used to model this blurring and delay as an impulse response function under the assumption of linearity. Moreover, it was common for traditional fMRI analysis to infer BOLD responses at a level of spatial detail that does not consider the cortical depth.

However, now high-field based laminar-specific fMRI acquisition is becoming more feasible, so there is a pressing need to pay more attention to depth-dependent effects of blood traversing the cortical depth when modelling the BOLD signal. One critical reason to focus on effects of signal blurring in laminar fMRI is that blood draining through vessels with larger diameters affects the final BOLD signal and can potentially appear amplified as if an activation occurred. It is important to distinguish whether a BOLD response is dominated by neuronal activity pertaining to cortical layers or whether it is affected by blood draining up the cortical depth. In this current era, a more comprehensive understanding on how blood physiology influences BOLD measurements in individual layers will enrich the ever so advancing field of laminar fMRI research.

In cerebrovascular physiology, it has already been established that venous blood drains up the cortical depth traversing larger ascending veins and pial veins along the cortical surface. Subsequently, MRI studies have observed that due to blood draining up large vessels, BOLD signals are dominated by macrovascular contributions rather than microvascular effects intended to be measured near sites of neuronal activation [75]. Despite the widespread use of gradient-echo (GE) acquisition to achieve high-resolution laminar fMRI, GE-BOLD signals are biased towards larger veins closer to the cortical surface [128]. In post-processing stages in GE-BOLD frameworks, spatial deconvolution can be used to eliminate unnecessary signal amplification in superficial layers [130]. As consequence of such BOLD signal contamination, other methods have also been investigated to mitigate for unnecessary amplification of BOLD signals in upper cortical layers. For instance, some studies have utilized spin-echo (SE) BOLD acquisitions that have higher specificity, albeit at the cost of lower contrast-to-noise ratio (CNR). Other methods that combine GE and SE-BOLD with vasculature filtering [129] attempt to balance the sensitivity and specificity deficits

involved in each acquisition method. Moreover, alternative non-BOLD modalities such as Cerebral Blood Volume (CBV) based Vascular Space Occupancy (VASO) acquisition respond more specifically to capillaries, therefore, offer an alternative way to study laminar activity and connectivity between distinct regions more precisely.

In this chapter, we explore layer-specific effects when simulating more biologically plausible signals using a dynamic causal modelling approach compared to the SEM modelling approach in Chapter 3. Dynamic Causal Modelling (DCM) has often been used in connectivity studies and DCM models can be modified to fit data acquired in different modalities from faster EEG/MEG data to more sluggish fMRI data [48, 91, 131]. These aforementioned general DCM studies still focused on modelling connectivity between separate brain regions, but more recent studies adapted biophysical models in DCM for fMRI while accounting for physiological effects caused by blood flowing through intra-cortical ascending veins and pial veins. Such approaches started to investigate depth-dependent physiological aspects and their impact on haemodynamic parameters [53, 89].

The method used in this chapter was inspired by work from Heinzle et al. [78] who suggested a haemodynamic model for blood draining across cortical layers. Their models were designed to account for 2 depth levels (two layers) corresponding to the bottom and top cortical layers of a region. Haemodynamic state equations can then, similarly to previous approaches, be incorporated into DCM for fMRI. Hence, BOLD haemodynamic parameters specific to upper and lower depths could be assessed, separately.

Now, in this work we introduce two innovations building on the model introduced in [78]. First, we extend their proposed layer-DCM model to account for an arbitrary number of layers as opposed to their two-layer model. Second, we used this layer-DCM implementation with the intention of combining with our inversion recovery BOLD approach. To begin with, we compare this DCM modelling approach against the SEM work from Chapter 3. As part of investigating those overarching aims, we first used simulations to investigate how blood draining influences laminar activity

and connectivity between cortical areas that model more than two layers. Hence, the more specific research questions we aim to address are the following:

- In Chapter 3 we modelled laminar connectivity, but without considering underlying haemodynamics. We suggest DCM modelling as the next step, so that we can fit models that are more biologically plausible to our IR-BOLD data. In this chapter, we explore this alternative modelling technique.

- We extend DCM to model activity in multiple layers. Through data simulation and model fitting we can examine how blood-draining interferes with layer-specific connectivity relationships. We do not, yet, fully comprehend how blood-draining affects laminar BOLD signals with respect to its underlying activity during rest or task execution. Detecting laminar connectivity signatures is one of the central endpoints in laminar fMRI studies, and we could potentially benefit from this multilayer DCM simulator if it can be used as a benchmark for checking whether measured laminar fMRI signals have been affected by blood-draining.

5.2 Methods

In this section, we introduce basic elements of a DCM model that is generally used in fMRI. Then, we observed how haemodynamic and neural signals change in a laminar-specific manner when using the layered DCM suggested by Heinzle and colleagues [78] extended to multiple layers. Additionally, we introduce a stochastic DCM approach to simulate resting-state conditions.

After introducing models, several experiments were designed to test through simulated data. We conduct a thorough breakdown of haemodynamic and neuronal parameters involved in the DCM model when modelling laminar signals for a single region (Section 5.2.2.1) and also for 2 connected regions (Section 5.2.2.2). Finally, a stochastic DCM was used to simulate a resting-state scenario (Section 5.2.2.3).

5.2.1 Background on DCM for fMRI

A DCM framework allows us to study effective connectivity between different brain areas. DCM models neural dynamics by combining a model of neural activity with a haemodynamic model. Specifically, two main parts of DCM equations for BOLD fMRI are the neuronal state equation (x) and haemodynamic state equations (s, f, v, q) that explain neurovascular coupling mechanisms that relate neuronal activity to neurovascular signals (see Table 5.2.1). In this section we first summarise equations involved in the standard DCM for fMRI.

DCM parameter	Neuro/Hemo	Description
x	Neuronal	"Neuronal" time course
s	Haemodynamic	Neuro-vascular coupling signal
f	Haemodynamic	Blood flow
v	Haemodynamic	Blood volume
q	Haemodynamic	dHb concentration

5.2.1.1 Neuronal activity state variables

In DCM neuronal activity is expressed as the following state equation:

$$\frac{dx}{dt} = Ax + Cu \tag{5.1}$$

Neuronal activity $x(t)$ defined for each node, often a brain region in classical DCM, is multiplied with A that represents a fixed connectivity matrix. Matrix C includes weights for the direct external input $u(t)$ which is the given experimental stimuli. If we were to use an alternative DCM that accounts for interactions in regions due to context-induced inputs u_j that differ from direct external stimuli u_i, bilinear equations can be added to Eq. 5.1,

$$\frac{dx}{dt} = (A + \sum u_j B^j)x + Cu_i \tag{5.2}$$

Here, matrix B^j is the modulated intrinsic (fixed) connectivity induced by context-dependent inputs u_j [132, 133]. In this chapter, we do not make use of this version of DCM.

5.2.1.2 Comparison of DCM and SEM state equations

There is a link between the above Eq. 5.1 and the previously visited structural equation modelling (SEM) in Chapter 3 SEM was formerly used to study causal relationships between regions but for input that was stochastic (random noise) rather than being controlled through exogenous input as in DCM [134]. From Eq. 5.1, input terms can be replaced with an error term ϵ that follows a distribution of Gaussian random noise as shown below:

$$Cu = \epsilon, \quad \epsilon \sim N(0, \sigma^2)$$

Additionally, SEM describes a specific case of DCM where we assume that neuronal activity x dynamics play out on a much faster timescale than the input u. Hence at the point of observation, the neuronal dynamics have reached equilibrium so that the first-order derivative of state variable x is 0 ($dx/dt = 0$). Finally, connectivity matrix A can be described as $\theta - I$ in which θ is an off-diagonal matrix that corresponds to extrinsic connectivity weights and unit inhibitory self-connections are expressed as diagonal matrix I. Using the terms above, Eq. 5.1 can be transformed to:

$$\begin{aligned} \frac{dx}{dt} &= Ax + Cu \\ &= (\theta - I)x + \epsilon = 0 \\ &\text{(where } A = \theta - I, \quad Cu = \epsilon) \\ x &= \theta x + \epsilon \end{aligned} \tag{5.3}$$

Thus, Eq. 5.3 is equivalent to the SEM equation.

5.2.1.3 Haemodynamic state equations

Changes in neural activity are known to induce changes in local cerebral blood flow (CBF). This phenomenon was modelled into neural activity-induced vascular response analogous to a dampened oscillator [135]. The model describes vasodilatory

signal (s) decaying exponentially in response to neural activity x and autoregulatory feedback from incoming CBF (f) as shown in the following equation:

$$\begin{aligned} \frac{ds}{dt} &= x - \kappa s - \gamma(f-1) \\ \frac{df}{dt} &= s \end{aligned} \tag{5.4}$$

In Eq. 5.4, κ and γ are rate constants for vasodilatory signal and autoregulatory feedback for blood flow, respectively. Also, f is normalised CBF at rest.

Next, we require equations that express CBF-induced changes in blood volume (v) and deoxyhemoglobin (dHb) content (q). In the balloon model, changes in v and q can be expressed as differential equations that depend on changes in CBF [136]. First, change in v is calculated as blood flow differences in the venous compartment of interest:

$$\tau \frac{dv}{dt} = f_{\text{inflow}} - f_{\text{outflow}}(v)$$

Here, time delay constant (τ) represents the transit time of blood passing the venous compartment and is equivalent to the ratio of resting volume V_0 over resting blood flow F_0 (V_0/F_0). Also, this equation can be updated so that blood flowing out is expressed as a function of volume:

$$\begin{aligned} f_{\text{outflow}}(v) &= v^{\frac{1}{\alpha}} \\ \tau \frac{dv}{dt} &= f - v^{\frac{1}{\alpha}} \end{aligned} \tag{5.5}$$

Here, Grubb's exponent α represents resistance that allows balloon-like inflation of vessels [137]. Similarly, change in q is calculated as blood flow changes coupled with relevant oxygen extraction fractions so that:

$$\tau \frac{dq}{dt} = f_{\text{inflow}} \frac{E(f_{\text{inflow}}, E_0)}{E_0} - f_{\text{outflow}}(v) \frac{q}{v}$$

while $E(f_{\text{inflow}}, E_0)$ is the fraction of oxygen extracted during blood inflow and $\frac{q}{v}$ is the concentration of dHb. E_0 is the resting oxygen extraction fraction and the above oxygen extraction fraction can be further described as:

$$E(f_{\text{inflow}}, E_0) = 1 - (1 - E_0)^{\frac{1}{f_{\text{inflow}}}}$$

and the final equation of q was updated as:

$$\tau \frac{dq}{dt} = f \frac{1-(1-E_0)^{1/f}}{E_0} - v^{\frac{1}{\alpha}} \frac{q}{v} \tag{5.6}$$

For final Eqs. 5.5,5.6, f_{inflow} was replaced with notation f for simplicity in further extension of the equations.

Finally, the haemodynamic state equations for s, f, v, and q provide a forward model for predicting the BOLD signal change from DCM state equations. The Balloon model [136] was extended [135] so that BOLD signal can be approximated as:

$$\frac{\Delta S}{S} \approx V_0 \left[k_1(1-q) + k_2 \left(1 - \frac{q}{v}\right) + k_3 (1-v) \right] \tag{5.7}$$

$$k_1 = 7E_0$$

$$k_2 = 2$$

$$k_3 = 2E_0 - 0.2$$

Equation 5.7 is a nonlinear equation of v and q in which $k_1 - k_3$ were approximated assuming that field strength was 1.5T and TE was 40 ms in [136]. In other following studies, parameters $k_1 - k_3$ have been modified in other ways to depend on acquisition parameters or magnetic field strength as demonstrated in [132]. For the simulations in this chapter, we used the derivations from [136] to keep things relatively simple rather than adopting the most recent derivations of $k_1 - k_3$.

5.2.1.4 Extension to multiple layers

Heinzle et al. [78] added a blood draining compartment in the original DCM approach so that blood physiology for upper and lower layers could be modelled separately. Their phenomenological model assumed that venous blood drains from lower layers closer to the GM/WM surface towards upper layers near the pial surface. Here, we extended the Heinzle model so that an arbitrary number of layers could be modelled rather than just 2 layers (upper, bottom layers). The adapted model we used undergo several modifications described in the following part. We simply changed annotations from the previous Heinzle model, so that blood draining parts

for model scenarios with more than 2 layers is easier to grasp.

Regardless of the total number of layers defined for a model, blood draining was assumed between two consecutive layers as shown in the following equations where v_k is relative volume and q_k is dHb content of the k^{th} layer:

$$\begin{aligned} \tau_{\text{drain}} \frac{dv_k{}^*}{dt} &= -v_k^* + (1 - v_{k+1}) \\ \tau_{\text{drain}} \frac{dq_k{}^*}{dt} &= -q_k^* + (1 - q_{k+1}) \end{aligned} \tag{5.8}$$

Here, the blood that drained from the k^{th} layer has volume (v_k^*) and dHb content of (q_k^*) and traverses towards the layer above (the $(k-1)^{\text{th}}$ layer). Equation 5.8 expresses the fact that blood in the lower $(k+1)^{\text{th}}$ layer (v_{k+1}, q_{k+1}) drives draining through to upper layer v_k^*, q_k^*.

This numbering was chosen so that it is kept consistent with layer numbers commonly used in histological descriptions ranging from layers I – VI (layer I is the uppermost layer near the pial surface and layer VI is innermost layer above the GM/WM surface). Note that τ_{drain} is the time delay of blood travelling from layer $k+1$ to layer k (k = 1 ... n-1 for integer n). Next, the haemodynamic state equations (Eqs. 5.5,5.6) for the affected superior layer were modified to reflect blood drained in Eq. 5.8 so that:

For layer k (k=1...n-1)

$$\begin{aligned} \tau \frac{dv_k}{dt} &= f_k - v_k^{\frac{1}{\alpha}} + \lambda_{\text{drain}} v_k^* \\ \tau \frac{dq_k}{dt} &= f_k \frac{1 - (1 - E_0)^{1/f_k}}{E_0} - v_k^{\frac{1}{\alpha}} \frac{q_k}{v_k} + \lambda_{\text{drain}} q_k^* \end{aligned} \tag{5.9}$$

The above equations demonstrate state variables v_k, q_k that incorporate draining effects from the inferior $(k+1)^{\text{th}}$ layer and were weighted by a hypothetical blood draining strength λ_{drain}.

Finally, for a multilayer model, the innermost layer on the bottom level follows

Eqs. 5.5, 5.6 giving:

$$\textbf{Bottom (n}^{\textbf{th}}\textbf{) layer}$$
$$\begin{aligned} \tau \frac{dv_n}{dt} &= f_n - v_n^{\frac{1}{\alpha}} \\ \tau \frac{dq_n}{dt} &= f_n \frac{1-(1-E_0)^{1/f_n}}{E_0} - v_n{}^{\frac{1}{\alpha}} \frac{q_n}{v_n} \end{aligned} \tag{5.10}$$

5.2.1.5 Resting-state simulation using stochastic DCM

Laminar activity in resting-state (RS) has been considered less straightforward to observe compared to task-induced activity that could be highlighted through laminar pathways revealed in past invasive animal task-based studies. Many task-based studies have been used to directly investigate activation in or connections between layers of different cortical areas, such as the visual, auditory cortices, or higher-order cognition related areas related to working memory. However, RS-BOLD activity in layers not only has been rarely assessed, but we also do not fully understand whether results from task-based laminar fMRI experiments are generalisable to current understandings on RS activity [138]. In the current chapter, we simulate RS data with stochastic DCM to investigate how neuronal activity evolves when venous draining effects were accounted for at rest.

To simulate resting-state activity using laminar DCM, the deterministic input $u(t)$ must be replaced with a stochastic input. Solving the forward problem requires implementing stochastic differential equations (SDEs). The SDEs assume noise under a standard Weiner process, which here is an integration of Gaussian white noise over time. Here we used a python implementation to solve a stochastic ordinary differentiation equation from (`http://github.com/mattja/sdeint/`). The internally generated noise (dW) is randomly generated and is only affected by noise from a current time point. SDEs were solved through integration with Ito calculus.

There were several past studies that implemented stochastic DCM that assumed more realistic noise models optimal for neural activity [131, 139]. For instance, one recent version of stochastic DCM allowed for temporal correlation of noise modelled along the time course [140]. At this stage in our simulation study, we

did not explore more recent stochastic DCM approaches that assume structured noise following brain network activation [141].

We rather focused on what can be observed when integrating a blood-draining element to a layered DCM model. Hence, we simulated scenarios to observe how RS laminar-specific connectivity is shaped at different blood-draining strengths. Moreover, in this simulation stage, we did not conduct inference on model parameters with real fMRI data nor attempt to search for a best fitting model among candidate models, yet. Noise modelling beyond the stochastic Weiner process in the SDE used here would be critical when it comes to the estimation of unknown parameters.

5.2.2 Experiments for layer DCM study

We implemented our layered DCM (`https://git.fmrib.ox.ac.uk/saad/dcsem`) in python. The following experiments were conducted. For experiments requiring inputs, we used varying lengths of input including a short (1-2s) or long input stimulus (10s).

In this work, we used fixed constants shown in Table 5.1

Constant	Meaning	Value	Units	Source
κ	Rate constant of feedback regulation	1.92	s^{-1}	[78]
γ	Rate constant of vasodilatory signal decay	0.41	s^{-1}	[78]
α	Vessel stiffness	0.32	-	[78]
E_0	Oxygen extraction fraction at rest	0.34	-	[78]
τ	Transit time	2.66	s	[78]
k_1	$k_1 - k_3$ were estimated for field strength B_0=1.5 T, TE=40 ms	$7E_0$	-	[136]
k_2		2	-	[136]
k_3		$2E_0$-0.2	-	[136]
V_0	Venous blood volume fraction at rest	0.02	-	[136]

Table 5.1: Fixed parameters used in haemodynamic state equations (Eqs. 5.4, 5.9, 5.10) and BOLD signal approximation function (Eq. 5.7).

We focused on effects of modulating a subset of model-specific parameters such as layer-dependent parameters (λ_d, τ_d, Table 5.2), number of layers used (n) and neuronal connection weights (A_{ij}). Fixed parameters were set as defined in Table 5.1. Haemodynamic parameters $\kappa, \gamma, \alpha, E_0, \tau$ were taken from Heinzle et al. [78] and $k_1 - k_3$ were as defined in an earlier version of DCM based on a Balloon model [136] extended by [135].

Constant	Meaning		Range for grid search
λ_d	Draining strength	No exact measurable physical value	[0.0, 0.2, 0.4, 0.6, 0.8]
τ_d	Transit time between adjacent layers	1s (but varied between 0.5 – 2.5 s in [78])	Fixed for each instance

Table 5.2: Table of haemodynamic parameters related to blood draining (Eqs. 5.8, 5.9).

5.2.2.1 Multi-layer DCM for a single region

In this section, we simulated data from a single region containing multiple layers. We conducted simulations on various model architecture types that will be elaborated on further. For single region models, we focus on testing various model architectures (e.g. is there blood-draining, or what inter-laminar connections are defined) and input stimulus types (e.g. which layer or layers a stimulus is applied on).

We conducted simulations similar to fig. 1 in [78] to investigate the effect of changing some of the model parameters. For a single region divided into 3 layers across the cortical depth representing upper, middle, lower layers, we observed how blood draining affects haemodynamic parameters v, q, f, s and simulated BOLD signal (y) for each layer. We focused on differences between the 3 layers when varying blood draining strengths λ_d using values: [0.0, 0.2, 0.4, 0.6, 0.8]. We also considered different patterns of neuronal connections. An initial default model was defined so that neuronal connections consisted of only inhibitory self-connections and all other inter-laminar connections were zero (Figure 5.1a, left). All subsequent models to be introduced were variations adding on to this default model (e.g. Figure 5.1a, right).

In models we assumed that inhibitory self-connections that exist within a layer have weights (magnitude) of -2. This negative value was set to such value to enable fast transient neuronal activity within each layer. Then for models involving inter-laminar connections, inputs of 1-2s and 10s duration were equally applied to all layers.

We next observe BOLD signals for each of the 3 layers of a single region when inter-laminar connections between layers are non-zero values (Figure 5.1a, right). Here, we tested a simple model consisting of inter-laminar connections between consecutive layers. In the layer-DCM model, each layer of a region was treated

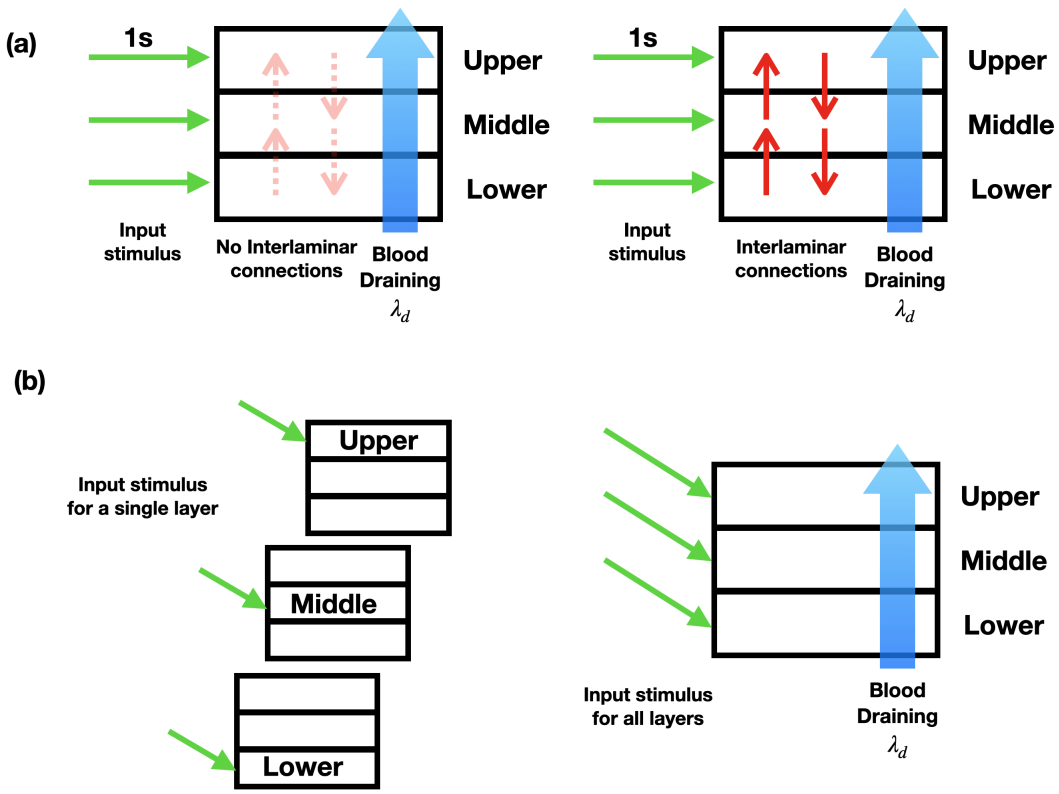


Figure 5.1: We explored various model architectures for single region simulations. Firstly, scenarios in (a) demonstrate when short input stimuli were applied to all layers (upper, middle, lower layers), simultaneously. Inter-laminar connections between neighbouring layers were either suppressed to null (left) or were set as positive weights (right). (b) Otherwise, inputs could also be applied separately to one layer, at a time. On the left we show cases in which an input stimulus targets a single layer (upper, middle, or lower layer). Input stimulus for the model on the right is sent to all 3 layers, and corresponds to the two model cases in (a).

as an individual node, thus, if a connection exists between two layers, a non-zero connection weight would be defined. We also examine two different stimulus conditions; when an input stimulus targets a single layer versus when stimulus is applied to all 3 layers simultaneously.

As a final experiment for a single region model, we increased the number of layers to 6 layers and applied an input stimulus to a middle layer (3^{rd} or 4^{th} layer from bottom). For this 6-layer model, we also included inter-laminar connections between consecutive layers. Especially, we examined how the magnitude of inter-laminar connection weights could influence haemodynamic parameters and the BOLD signal by comparing scenarios with stronger (weight = 1) and weaker

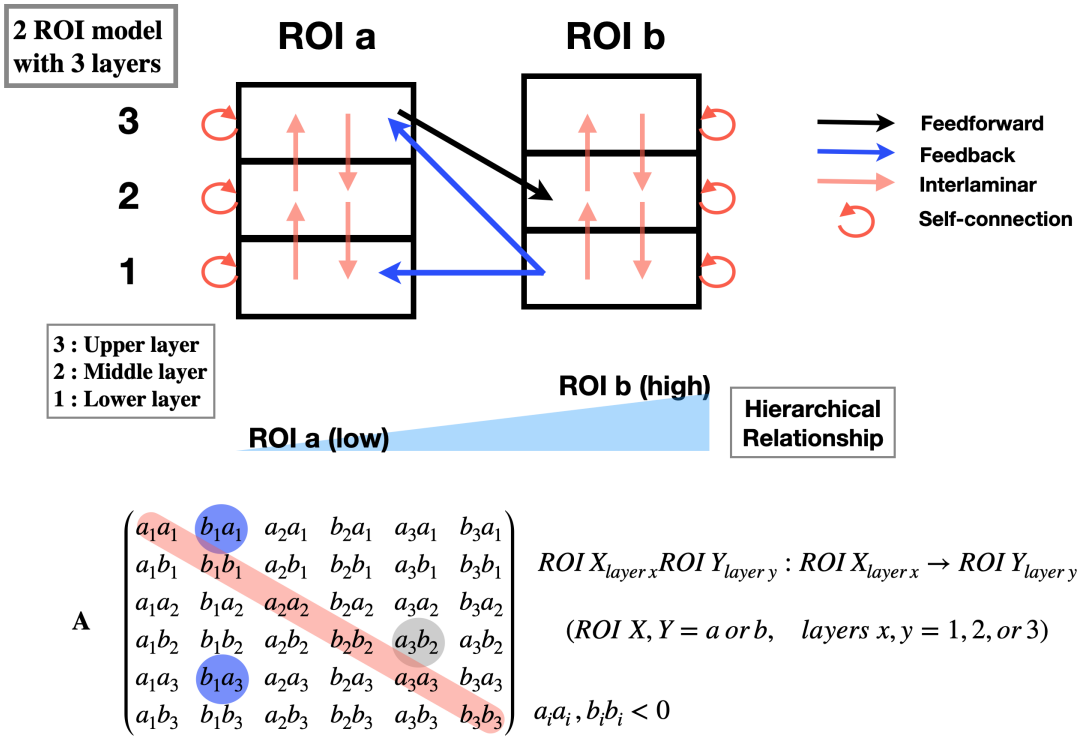


Figure 5.2: We explored multi-layer DCM models defined for 2 ROIS and 3 layers (upper, middle, lower) connected with such FFW and FB connections. In this representation, we explicitly show elements of the defined connectivity matrix A ((2×3)-by-(2×3)) highlighting non-zero connectivity weights (black: FFW, blue: FB, red: intra-laminar). Based on such laminar connections, we can assume that ROI a is hierarchically below ROI b. The main distinction of this DCM model diagram compared to the previously visited SEM models was the definition of self-connections defined for each layer nodes.

connections (weight = 0.5).

5.2.2.2 Extension of layer DCM for 2 ROIs

In the next set of simulations, we defined models with 2 regions and 3 layers (upper, middle, lower) connected via FFW and FB connections (Figure 5.2). We show the defined inter-laminar connections between consecutive layers within each region, inter-areal laminar connections such as FFW and FB connections, and self-inhibitory connections ($x_i x_i (x = a, b)$ in matrix A). Likewise as the previous SEM model, we assume one FFW and two FB connections propagating between regions, so ROI a is hierarchically below ROI b.

Assuming these FFW and FB laminar connections that underly hierarchical differences between 2 regions, we conducted 2 experiments. One was to compare

DCM and SEM covariances when changing the blood draining strength between null and λ_d of 0.5. In the next experiment, we compared 3 different input stimulus scenarios. Again, by comparing covariances of SEM and DCM we examined changes focusing on the condition in which inputs were applied to different layers in each instance and how existence of blood draining affects parameters.

5.2.2.3 Resting-state simulations using stochastic DCM

One of the motivations of this study was to evaluate our layer SEM approach from Chapter 3 in the context of a more realistic dataset that accounts for haemodynamics in a model. We now use a stochastic version of our layer DCM to generate "resting-state" data and examine against the SEM approach. The question we asked was: how different will it be if we fit DCMs that additionally account for haemodynamics and blood draining compared to SEM fitting? Here, we answer this question by directly observing covariance matrices implied by the two models (SEM vs DCM).

5.3 Results

The complexity of a certain model can increase rapidly as we increase the number of regions and layers, therefore we took care to gradually increase complexity in a stepwise manner. Therefore, we first conducted experiments for a single region 3-layer model (Section 5.3.1) and then for a two ROI (3 layer) model (Section 5.3.2) to validate a multi-layered DCM that assumes vascular blood draining. Finally, we inspected a DCM with blood draining for neural activity in a task-free environment to investigate resting-state activity (Section 5.3.3).

5.3.1 Simulation on multi-layer DCM for a single region

5.3.1.1 Impact of blood draining on haemodynamic variables

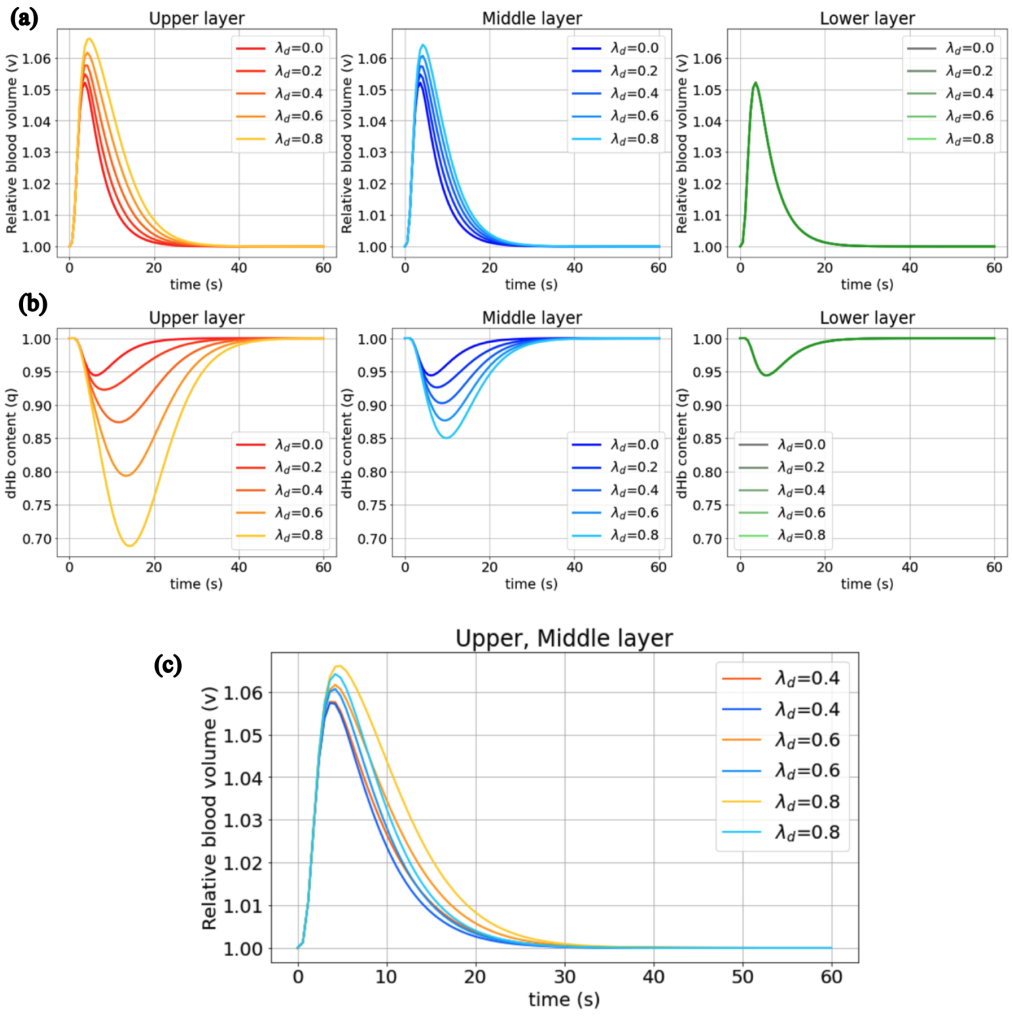


Figure 5.3: CBV and dHb of a default no inter-laminar model (Figure 5.1a, left) are shown. (a) Cerebrovascular Blood Bolume (CBV, v) and (b) deoxy-Hemoglobin (dHb, q) content that are haemodynamic parameters affected by adding a blood draining compartment into state equations. (c) is an enlarged juxtaposed view of both CBV in upper and middle layers in (a) for $\lambda_d = 0.4, 0.6, 0.8$.

The default model we discuss in the current section is a model consisting of 3 layers, a blood-draining compartment with (Figure 5.1a right) or without inter-laminar connections between consecutive layers (Figure 5.1a left). Input stimuli are injected into all upper, middle, and lower layers at the same time in these models. Thenceforth, we refer to the model with inter-laminar connections as

default inter-laminar model and the model without inter-laminar connections as default no inter-laminar model. To begin with, we injected input stimuli in all layers of a default no inter-laminar model and observed that the amplitudes of the haemodynamic variables v, q (Figure 5.3) and the BOLD signal (Figure 5.4a) increased towards upper layers near the pial surface. In particular, we observed amplification of dHb content q (Figure 5.3b) and BOLD signal (Figure 5.4a) in the upper layer for blood draining strengths $\lambda_d > 0.2$ (Figure 5.3b). For different draining strengths (λ_d), relative change in blood volume v was less prominent in the middle and upper layers (Figures 5.3a, c).

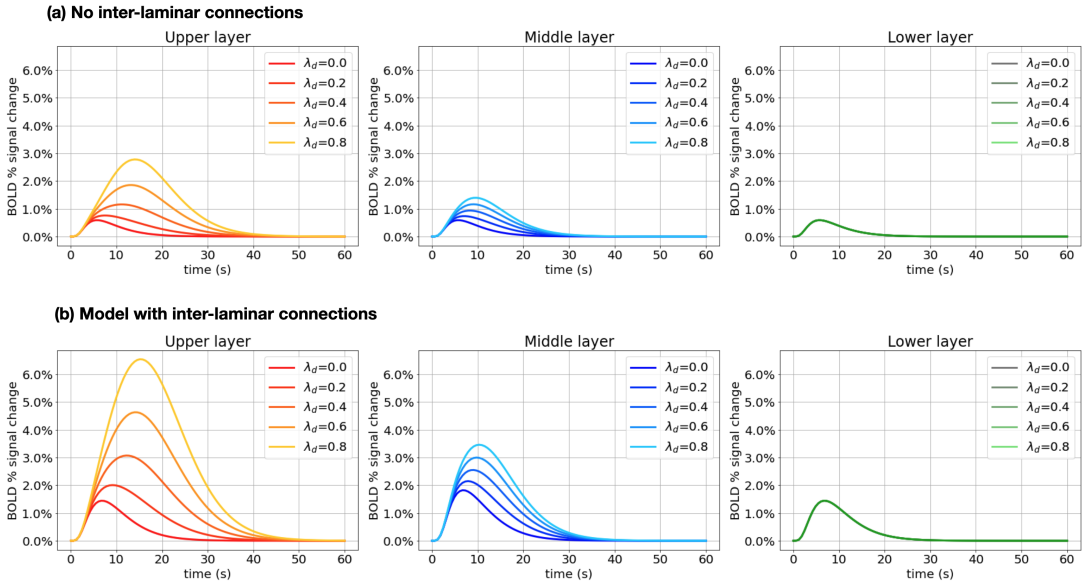


Figure 5.4: We compared BOLD percentage signal change when simulating data with models that lack inter-laminar connections (Figure 5.1a left) and models that assumed inter-laminar connections between neighbouring layers (Figure 5.1a right). For a single region 3 layer-DCM model, defining non-zero inter-laminar connections noticeably increased BOLD signal amplitude, especially, with greater blood-draining strengths.

Next, in Figure 5.4 we compared aforementioned default models with or without inter-laminar connections. The default inter-laminar model has neuronal connections between consecutive layers heading towards either the pial or WM surface. Neuronal connections propagating from the uppermost layer towards the WM surface were parallel to the direction of arterial blood inflow while connections in the opposite direction were parallel to venous blood flowing out through intracortical veins.

Figure 5.4 demonstrates BOLD percentage signal changes. Overall, the model with inter-laminar connections had greater BOLD signal change (%) in all 3 layers (Figure 5.4b) compared to the model without inter-laminar connections (Figure 5.4a). Both models had greater BOLD signal increase at stronger blood-draining ($\lambda_d = 0.6, 0.8$). Similarly, data simulated with longer 10s duration at greater blood-draining strengths was greatest for upper layers (Figure 5.5).

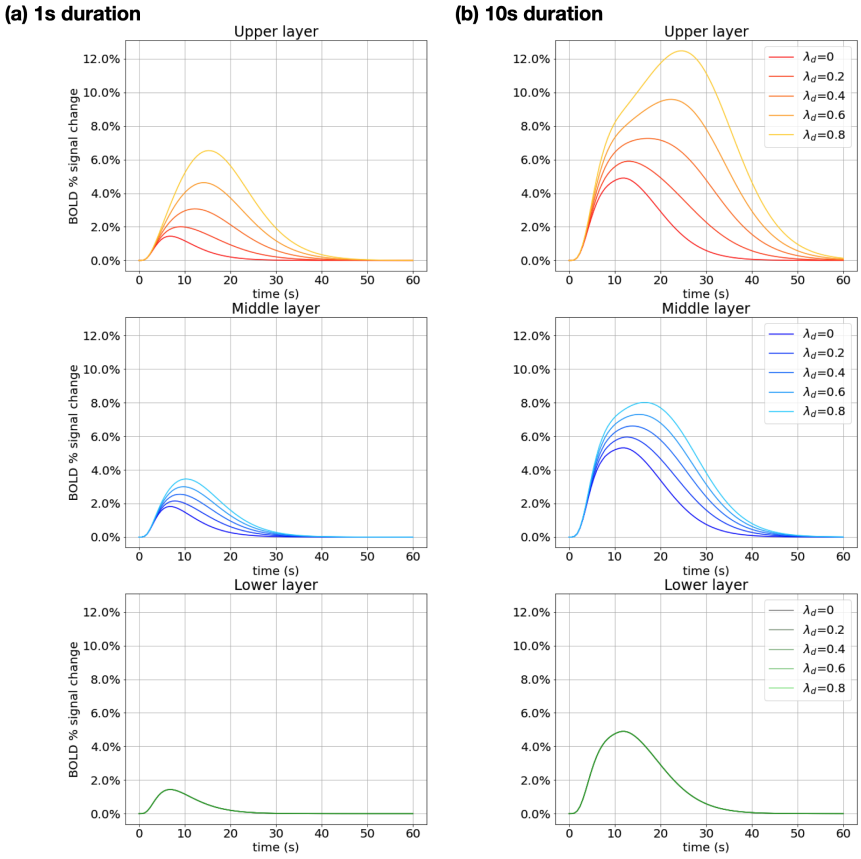


Figure 5.5: Plots of BOLD percentage signal change for variable degrees of draining strengths (0 – 0.8) are displayed for two input durations (1s, 10s). There is an increase in BOLD for all layers when a longer input stimulus of 10s is applied.

Finally, we investigated scenarios in which inputs target a single layer. We compared 3 scenarios that alternate the layer that an input stimulus was applied as in models (i)-(iii) in Figure 5.6a. Time delays to reach peak BOLD response were computed for draining strengths λ_d between 0 – 0.8. We observed that the time delay to reach peak BOLD was shortest in the layer that received the input stimulus when λ_d was nulled ($\lambda_d = 0$) (Figure 5.6a). However, with non-zero draining

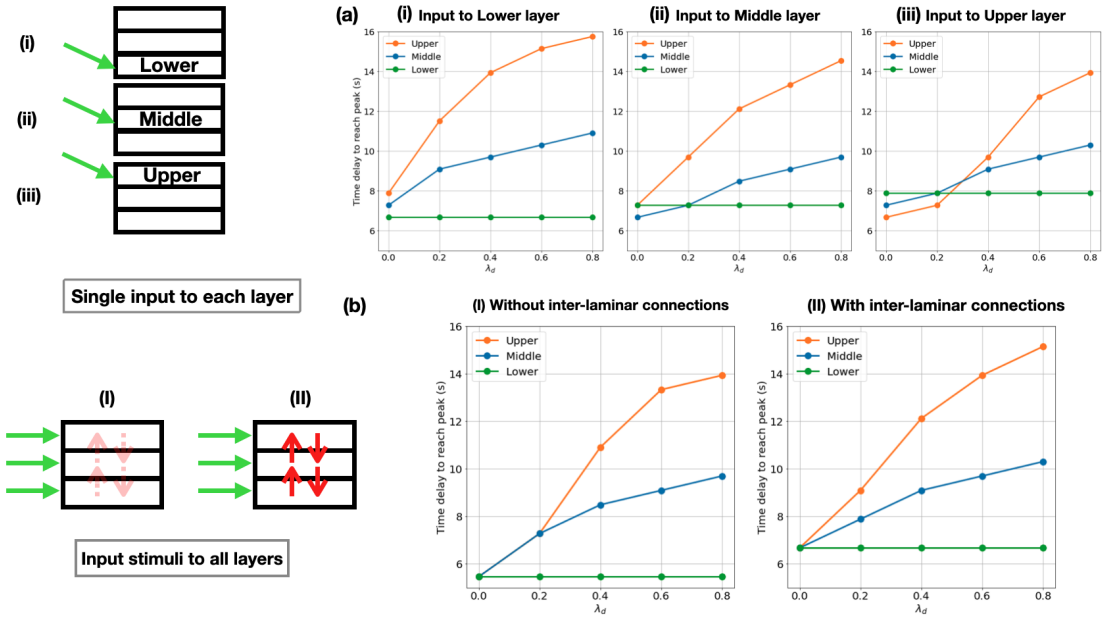


Figure 5.6: Depending on the target layer or layers in which a short input stimulus (2s) was applied it is possible to study time latency differences (time to reach peak BOLD) for each layer. (a) When input stimulus was applied to lower, middle, and upper layers separately as depicted in models (i)-(iii), time delay to reach peak BOLD was shortest in the layer that directly received input, when there was no blood-draining ($\lambda_d = 0$). When there is blood draining at $\lambda_d > 0$ magnitude of peak BOLD increases towards upper layers and time delay to reach peak BOLD also becomes longer. We observed that especially the upper layer is affected by blood-draining compared to the other two layers. (b) Models for the two graphs are shown in I-II and these models are comparable by existence of inter-laminar connections for neighbouring layers. Accordingly, time latency was greater for models with inter-laminar connections (Model II). Also, when input stimuli were simultaneously applied to all 3 layers we observed that time to reach peak BOLD increased for greater draining strengths in upper and middle layers.

strength at $\lambda_d > 0.2$, the time delay increased from lower layer towards upper layers regardless of the layer that the input stimulus was applied to. Furthermore, we observed that time delays to reach the peak BOLD response were longer for models with inter-laminar connections (Figure 5.6b).

5.3.1.2 Assessing increase of BOLD in mid-cortical depth

High resolution fMRI studies often report laminar activity changes as percent signal changes. Many studies have particularly focused on detecting elevated activity in some known layer of interest, such as layer IV in a sensory area. Here we use simulations to assess how peak activity in different cortical depths (as a proxy for

layers) vary and when we account for blood draining effects.

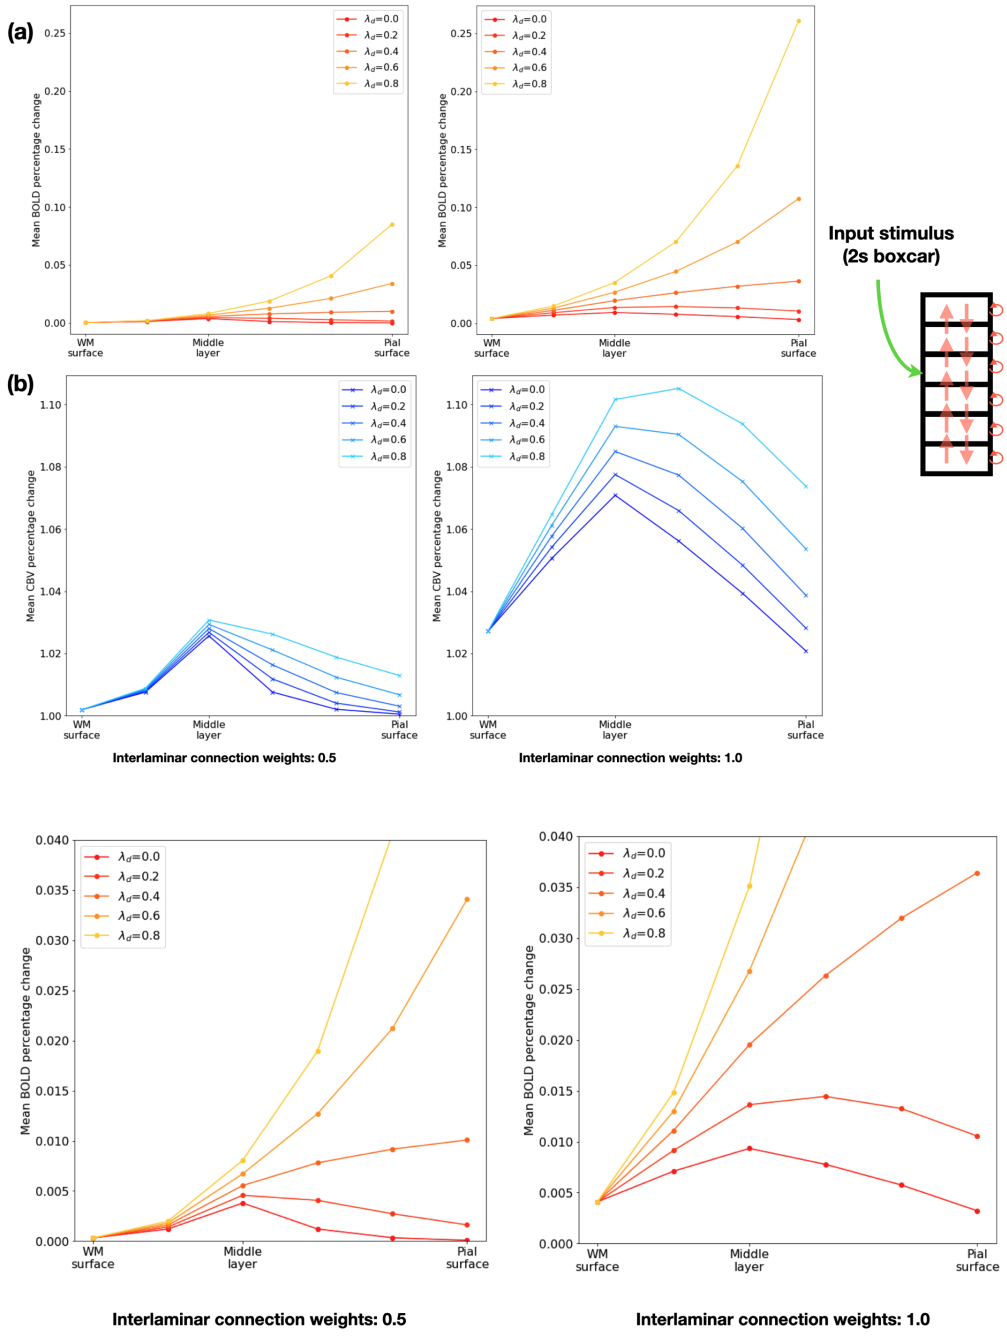


Figure 5.7: In a single ROI model with 6 layers with input stimulus targeting a middle layer (3rd layer from cortical surface in this example) BOLD and CBV increases in the middle layer. However, as blood-draining becomes stronger upper layer signals had substantial increases in mean BOLD percentages, so that the middle layer receiving input was no longer significant.

Using multi-layer DCM, we compared percentage signal changes in BOLD signal

and the relative volume (CBV) when an input stimulus was applied to a middle layer (Figure 5.7). We observed that for weaker draining strengths $\lambda_d < 0.4$, both BOLD and CBV percentage changes were most pronounced in the input receiving middle layer. For stronger blood-draining, BOLD signal amplified towards upper layers closer to the pial surface. As a consequence, at stronger blood-draining strengths, the middle layer no longer had the largest BOLD signal change (%) among other layers, but percentage signal change was much larger in layers above the middle layer (Figure 5.7a (also, enlarged view in bottom row)). In contrast, CBV changes showed maximum activity in the middle layer (albeit a bit broadened).

Additionally, we investigated how BOLD and CBV percentage signals change with different connectivity weights. The left column in Figure 5.7 was simulated from models with inter-laminar connection weights of 0.5 and the right column used models with connectivity weights of 1.0. As a result of increasing interlaminar connectivity weights defined between consecutive layers, we observed that not only BOLD (plot lines with warm colours), but also CBV in upper layers were subject to amplification (plot lines with cold colours).

5.3.2 Inter-laminar connectivity between 2 regions

Next, we examined models with 2 regions connected through laminar connections (Figure 5.8).

In the lower section of Figure 5.8, we show how data covariances of (a) BOLD signal simulated from DCM, (b) neural activity (x) from DCM, and (c) SEM simulated neural activity compared to each other. Here a short input stimulus (1s) was applied to the upper layer of ROI a (of lower hierarchy). At $\lambda_d = 0.5$, we observed that greater covariances for neural activity from DCM shifted towards upper layer nodes in BOLD covariance matrices. In contrast, when blood draining was disregarded ($\lambda_d = 0$), covariance structures between the three data covariances (a)-(c) were relatively similar.

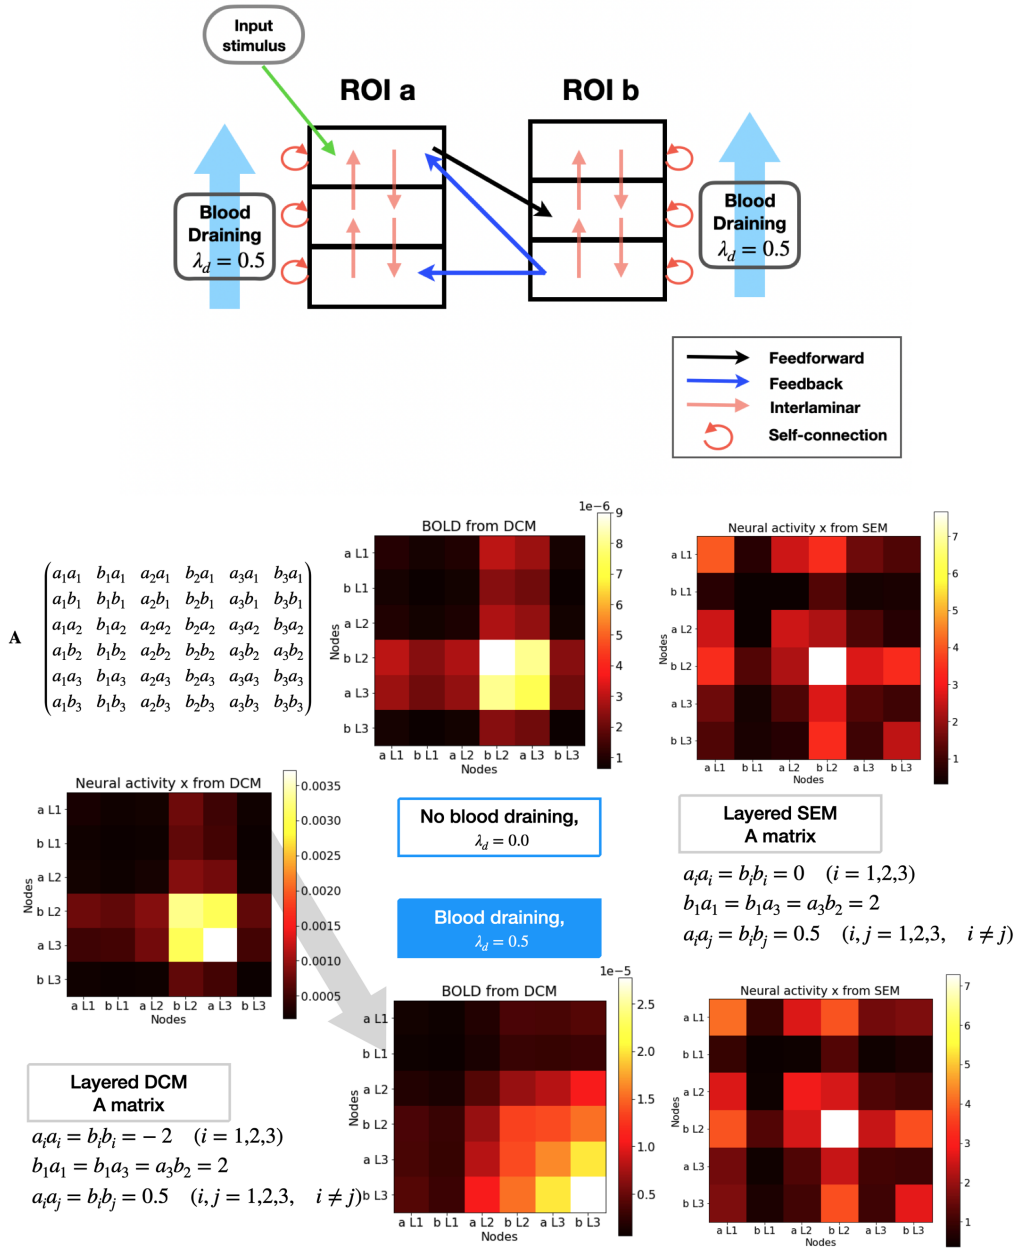


Figure 5.8: The model architecture in the top describes input and connectivity for the DCM model used and the SEM model had equivalent connections but without any input or self-inhibitory connections. To understand the effect of introducing blood-draining into a DCM model, we computed covariances of BOLD data simulated from multi-layer DCM, neural activity (x) simulated from DCM, and SEM simulated data. As shown in the bottom section, gradual changes were observed between each covariance matrix for $\lambda_d = 0.5$. This suggests an example of how integrating blood-draining compartment into a DCM model can change BOLD responses in superficial layers substantially. With blood draining, covariances of final BOLD data deviates from covariances of neural activity from DCM and SEM.

In another model shown in Figure 5.9, we show covariance matrices of data

simulated from 3 different models in which an input stimulus was applied to each layer one at a time. Without blood-draining, the overall covariance was largest in layers that received task input (first row of covariance matrices in Figure 5.9). On the other hand, with blood-draining, variance of ROI a upper layer (L3) was greatest in all 3 cases (A – C) that had different input locations (second row of covariance matrices in Figure 5.9).

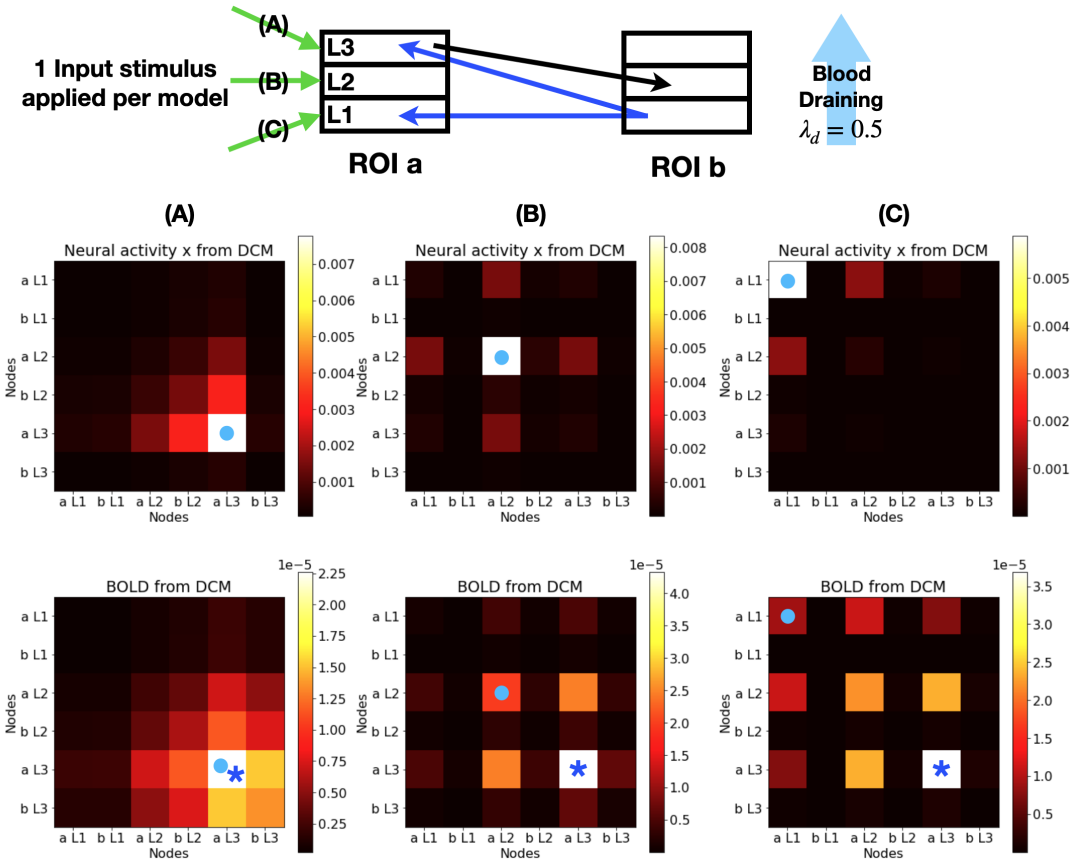


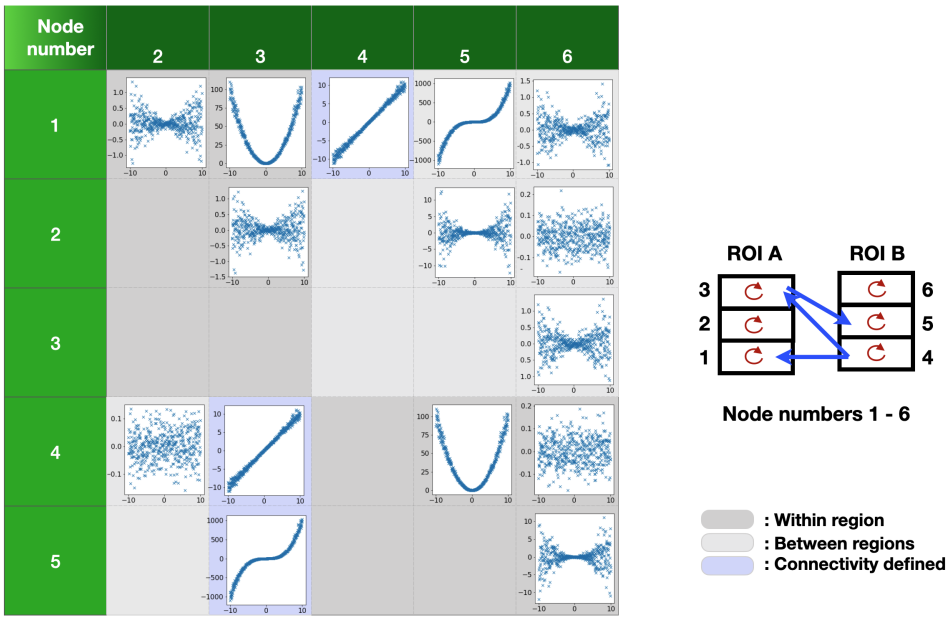
Figure 5.9: We investigated 3 different models with input stimulus (A-C) applied to each layer separately. Inter-laminar connections for consecutive layers within each ROI and self-connections within each layer were defined but omitted in the figure representation above. Computed data covariances of neural activity simulated from DCM each show highest covariance for layers that received input (A: L3, B: L2, C: L1 of ROI a marked with a light blue dot). However, covariances of BOLD data that assume draining at $\lambda_d = 0.5$ all have highest variance at L3 of region **a** (marked with dark blue asterisks). Light blue dots from upper row were simultaneously shown as a reference. Model parameters were set as the following; FVE connections: 1, self-connections: -2, inter-laminar adjacent connections: 0.5, input stimulus duration: 2s, input stimulus amplitude: 2.

5.3.3 Resting-state stochastic DCM analyses

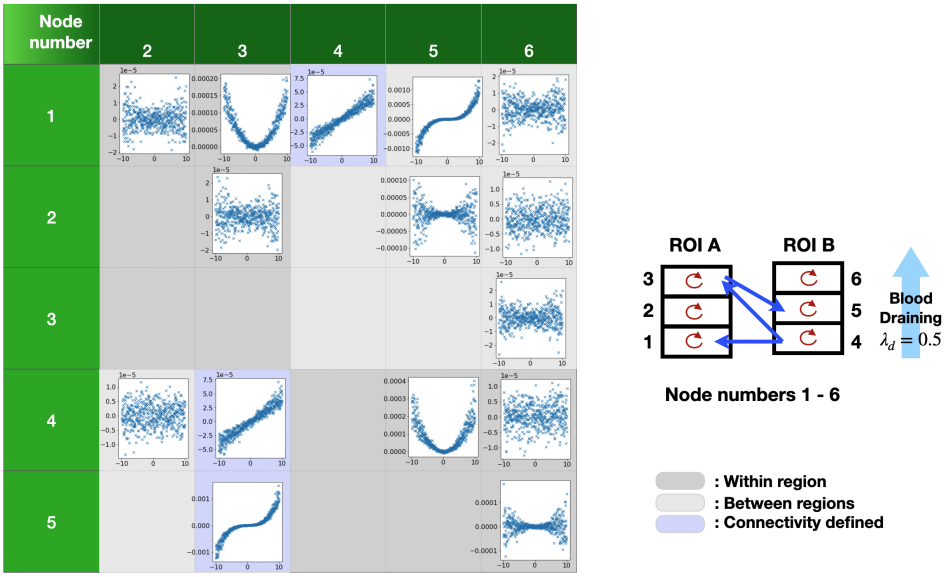
In this final analysis, we simulated RS laminar BOLD using stochastic DCM. As mentioned before, rather than fitting SEM models to DCM-generated data, we simply compare the covariance matrices of SEM-generated vs. DCM-generated data. One advantage of this approach is that we can compare many different scenarios with random datasets with different weight matrices A.

In Figure 5.10, we show how the covariance between ROIs/layers change as a function of the corresponding weights in the connectivity matrix for both SEM (Figure 5.10a) and DCM (Figure 5.10b). Even with a blood draining compartment added in DCM, weight dependent covariance changes are very similar between the SEM and DCM simulations. This was expected, as blood draining does not affect neural activity.

However, in actual fMRI acquistion we do not have access to neuronal activity (x) but only measure BOLD signal (y). We therefore evaluated data covariances when accounting for blood-draining in the overall BOLD response (y) computed from Equation 5.7 (Figure 5.11). In contrast with Figure 5.10, covariances between nodes within the same region (dark grey background) all had a similar parabolic shaped change of covariances for weights [-10, 10]. Covariances between pairs of nodes that were in separate regions also had an increasing pattern that were indistinguishable from each other. Therefore, at blood-draining $\lambda_d = 0.5$, it became difficult to distinguish covariances that arise from connections defined the model. One other noticeable difference observed in Figure 5.11 was that covariances that were randomly scattered around 0 in previous observations (Figure 5.10) gained a clearer structure (either parabolic or increasing pattern).



(a) Covariance between layers in data simulated from SEM.



(b) Covariance between layers in neural activity x simulated from DCM

Figure 5.10: Covariances between layers in our model of 2 regions and 3 layers are shown. We assumed FVE connections between the 2 regions (1 FFW, 2 FB connections, blue arrows on right hand side model). We used stochastic DCM, hence, input stimulus does not need to be defined. However, the model still requires non-zero self-connections (red arrows, weight = -2) in each node. This was to enable signals to be sent out to other nodes (i.e. layers). We observed how covariances changed for different FVE connection weights ranging from -10 to 10 and show plots for each unique node pair (15 pairs). In the table, row and column numbers shown in green were node labels 1-6 labelled on layers of the model (ROI A 1-3, ROI B 4-6). Moreover, covariance plots in the diagram have background colours depending on whether two nodes were connected as within-region connections (dark grey), between-region pairs (light grey). FVE connections defined in the model are shown with a blue background. Overall, figures (a) and (b) have similar weight-dependent changes in covariances. See Appendix A.6 for the original representation of covariance plots before rearranging.

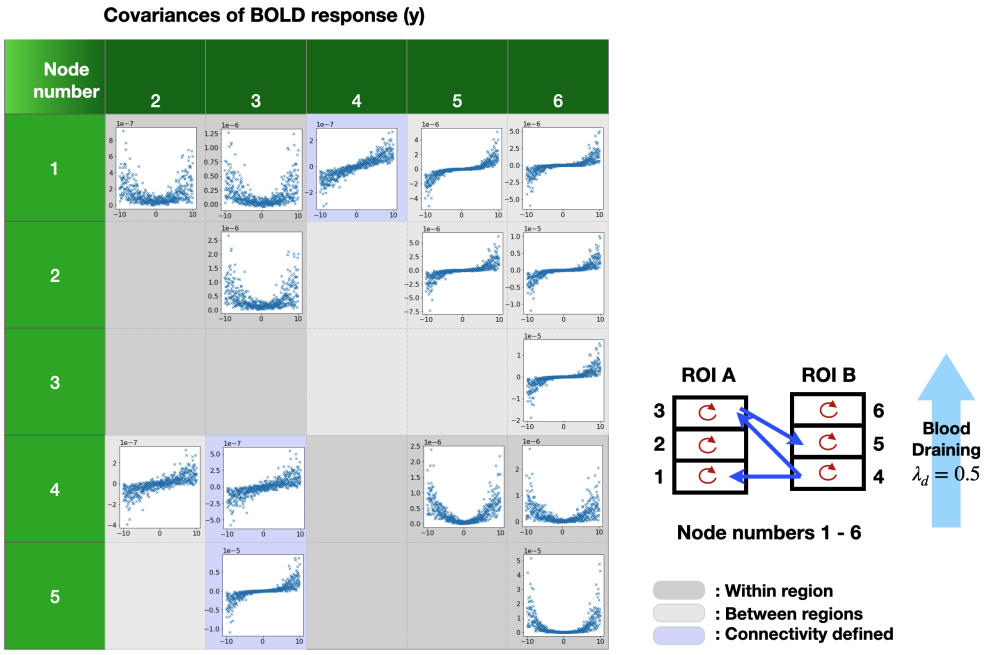


Figure 5.11: In the same format as described in Figure 5.10, covariances between pairs of layers were also shown for the final BOLD signal simulated from the stochastic DCM model that accounts for blood-draining. See Appendix A.7 for the original representation of covariance plots used in this figure.

Finally, we end simulation results with Figure 5.12 where we show how the mean correlation of the above covariance matrices of data simulated from SEM, neural activity from DCM, and BOLD activity from DCM changed with varying blood-draining strength. There was a moderate mean correlation between SEM and x of DCM (orange plot in Figure 5.12) before blood-draining parameters affected BOLD signal approximation in the stochastic DCM. For simulated BOLD at blood-draining strengths greater than 0.2 ($\lambda_d > 0.2$) mean correlation of covariance matrices with SEM simulated data and neural activity (x) simulated from DCM both decreased with greater λ_d (blue, green plots, Figure 5.12).

5.4 Discussion

Interpreting laminar activity in fMRI to infer neuronal connections in neuronal microcircuits [142] has been a long-standing goal that is not easy to achieve using the typical spatial resolutions available in laminar fMRI. We previously suggested

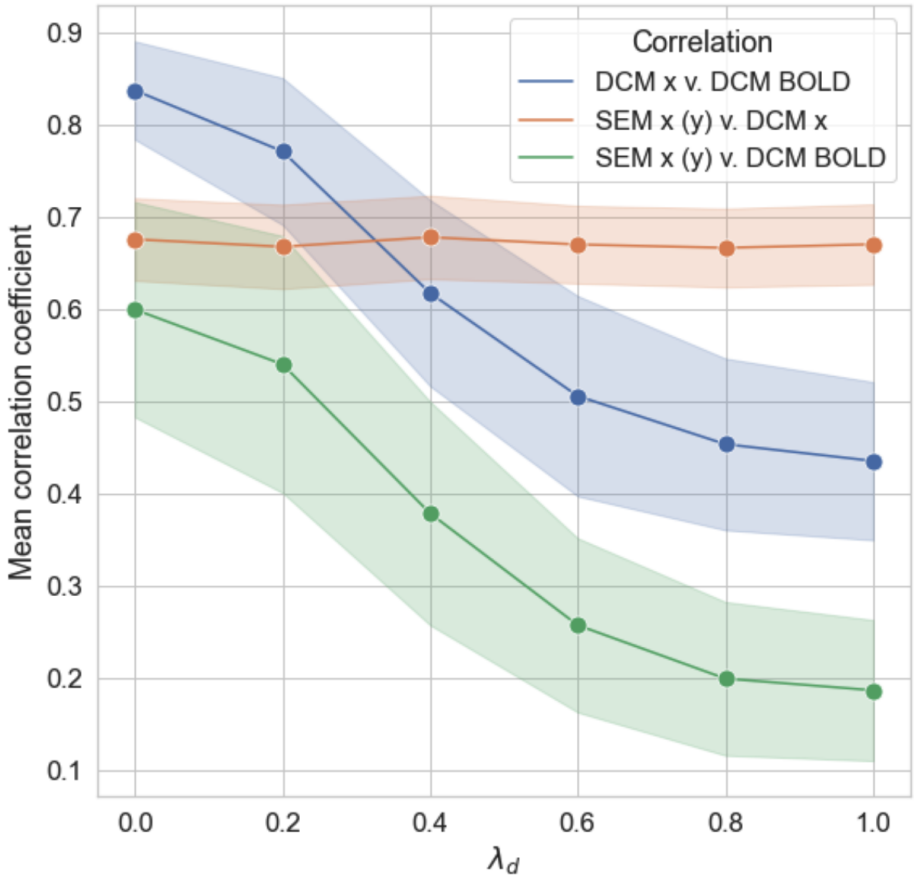


Figure 5.12: We computed correlation of data covariances of our 3 data sets of interest, which were neural activity (x) from a DCM model, BOLD simulated from a DCM model, and signal simulated from an SEM model. Connectivity matrices were equivalent in both DCM and SEM models, except that DCM connectivity matrices required negative value self-connections that connect a layer to itself. In the plot, we show mean and standard deviation of correlation coefficients computed between all pairs of covariance matrices (100 repeats for each λ_d).

an approach that combines an IR-BOLD acquisition with connectivity modelling using SEM that can potentially be used to assess hierarchies of areas that have typical laminar connection patterns (Chapter 3). However, an SEM model does not account for blood haemodynamics that interferes with measured fMRI signals. Closer to the pial surface, fMRI signals are subject to unwanted amplification as venous blood drains out of the cortex. Other studies have started exploring laminar activity and connectivity in light of more careful consideration of blood haemodynamics within the cortex [53, 78, 90, 130, 143]. In this study, we adapt a

laminar DCM model from Heinzle et al. [78] that can be extended beyond 2 layers, and we explore through simulations how activity and connectivity relationships change in multiple layers in context of blood-draining.

Difficulties arising from less specific BOLD contrasts such as from GE-BOLD and less sensitive SE-BOLD acquisition have motivated researchers to develop non-BOLD CBV weighted measures to study laminar activity with fMRI. It is challenging to achieve ample sensitivity and specificity so that haemodynamic changes following neuronal activation can be detected more accurately at the actual site of activation. Although are post-hoc methods to remove unnecessary amplified signals near the cortical surface [130], there often are trade-offs when using a certain image acquisition method in non-invasive fMRI [63]. Furthermore, with MRI we are still limited to low temporal resolution which is far from what invasive electrophysiological recordings regularly offer for probing neuronal activity. Despite advances in ultra-high field MRI that improved spatial resolution, signals are still confounded due to blood haemodynamics. When it comes to inferring laminar-specific connectivity with fMRI, especially, for connectivity based on BOLD measures, it is important that one accounts for blood-draining that amplifies laminar activity in upper layers.

In this work, we conducted simulations with multi-layered DCM that accounts for blood draining across the cortical depth, so that we can study effects of blood flowing through different layers. Based on our observations, we suggest how inferring connectivity relationships from these laminar activity profiles is potentially misleading if blood-draining is not considered in functional connectivity modelling. A cautionary note to add is that we did not perform model comparison and selection based on model fit, but we rather simulated data to understand how laminar activity in certain task-induced or RS conditions can be affected by draining blood. This simulation and respective observations were first required to improve our understanding of laminar BOLD signals.

5.4.1 Considerations on model architecture settings

To conduct a thorough analysis of possible scenarios that could be simulated through a task-evoked DCM and stochastic DCM with multiple layers, we increased complexity of models in a stepwise manner. This way, we could systematically observe when laminar activity of a certain model begins to diverge from another model, while modulating various factors including: connection weights, connectivity architecture, number of layers, number of regions defined, target layers for applying input stimuli, or parameters such as blood-draining strength λ_d.

In our simulations, model parameters were set in a way so that plausible signals would be generated for our intended model architectures. For instance, for any of the models, absolute values non-zero connections had to be smaller than weights of self-inhibitory connections. Conversely, if we assume weak connections close to zero connectivity weight, these connections would be regarded as non-existent as the resulting signals would appear the same as setting weights to 0. In sum, the connectivity matrix A had to follow certain conditions in which absolute values of the negative self-connection weights were the largest among other non-zero connections defined in A. This was so that we could allow for fast transient activity in each layer. Then, weights of other inter-laminar connections or inter-cortical connections such as FFW and FB connections were defined to have smaller absolute values than self-connections.

In 2-region models inter-laminar connections between adjacent layers within a region were required to assess between-region connections. For instance, in an FVE model, in order for a FFW connection sent to the middle layer of an ROI (of higher hierarchy) to continue to the bottom layer that can send FB connections to the other ROI (of lower hierarchy), inter-laminar connections between the middle and bottom layer had to be defined. Therefore, even though we based the FVE model architectures as defined in the classical tracer database study from Felleman and Van Essen [52], it was inevitable to increase the total number of connectivity parameters as inter-laminar connections between adjacent layers were required. Additionally, by defining inter-laminar connections of same magnitude in both directions towards

the pial and WM surface, the BOLD signal was smaller as inter-laminar connections in opposing directions compensate for each other. Likewise, connectivity weight magnitude related interactions should be further examined if we were to impose certain model architectures to ultimately estimate connectivity weights.

In a broader perspective related to models in a canonical microcircuit, our study assumed that connections were excitatory, which is a simplification of the actual neuroanatomy. We referred to the cytoarchitectural organisation of pyramidal cells only, thus, simplified the more intricate account of neuronal connections defined in a canonical microcircuit theory [142]. This allowed us to define inter-cortical connections and other inter-laminar within region connections that did not strictly adhere to how neurons connected with each other between layers. With additional assumptions regarding more detailed inhibitory or excitatory properties of neurons, it could be possible to establish specific DCM models based on precise anatomical organisation of neurons between cortical layers [53, 131, 144].

5.4.2 Analysis on data covariances

One key observation from this simulation study was that after incorporating a haemodynamic model of blood-draining, our suggested connectivity modelling approaches using SEM and DCM had different covariance structures (Figures 5.8, 5.9). We compared covariances of SEM against x of DCM and BOLD of DCM to observe when (at what blood-draining strength) they diverge from each other. One could alternatively investigate diverging covariance matrices by deriving an analytic expression of covariance for the DCM model.

We also simulated RS activity in layers for a 2-ROI model that assumes FVE connections using a simple implementation of stochastic DCM. Before blood-draining parameters came into effect, we could observe very clear covariance structures depending on the connectivity architecture (as defined in our connectivity matrix A). For instance, covariances between two nodes directly connected had different weight dependent covariance structure to those that were indirectly connected through a longer path-length (Figure 5.10). However, after assuming blood-draining towards

the cortical surface, such differences were no longer distinguishable (Figure 5.11). Overall, nodes (layers) within the same region had positive covariances regardless of connectivity weight sign, unlike covariances between layers in different regions. Whereas, nodes only covaried when there were a direct or indirect connection defined 5.10), layers affected by upstream blood-draining (i.e. all within-region nodes) positively covaried (5.11).

Dramatic changes in data covariances were observed after incorporating a blood-draining compartment into the haemodynamic state equations. Therefore, we suggest that when modelling layer-specific connectivity should be approached with caution.

5.4.3 Comparison of simulated laminar BOLD against literature

Our simulations confirmed that venous blood draining contributes to amplification of fMRI BOLD signals in upper layers closer to the cortical surface. It is important that neural activity from laminar fMRI is specific to the true site of activation both for BOLD and CBV-weighted measurements. In previous literature, both SE-BOLD or VASO studies on cortical layer activation reported an observable increase in middle layer activity in, for instance, the primary visual cortex (V1) [67, 145].

In our study, we also observed how CBV and BOLD changes with blood-draining when an input stimulus was sent to a middle layer. For our model consisting of 6 layers of equal thickness, our results suggested that with greater blood-draining strengths, CBV was less affected by signal amplification as the peak signal occurred at approximately the middle layer even with blood draining of $\lambda_d = 0.6$ (Figure 5.7). Next, BOLD signals were more sensitive to blood-draining effects and a middle layer peak was only observed when there was no blood-draining at $\lambda_d = 0$. BOLD signals in layers above the middle layer, which was the layer given an input stimulus, increased towards the cortical surface. There are studies supporting the notion that CBV is simply more sensitive to capillaries or smaller venules which could potentially support our observations on CBV compared to BOLD [146].

There have been other similar other studies that try to understand the source of haemodynamic signals in different cortical layers with regards to onset time delays and signal amplitude biases [147] that we have not, yet, explored.

To answer more specific questions such as how neuronal laminar connections are related to execution of brain functions, a more systematic model comparison may be necessary. Probing laminar communication pathways has long been a challenging problem and many researchers are motivated to devise experiments that aim to evoke a chain of neuronal activities using tasks that target specific neuronal pathways. Signals are often relayed through neurons that connect multiple cortical areas, and this implies that connections targeting a certain layer may not be the final destination of an activation pathway. Hence, we need to use tasks that clearly aim to activate a neuronal pathway, if possible. Also, it may help to further control for the target layer of a task stimulus as preliminarily explored in our simulations. In our current simulations, we assumed a task stimulus enters an area through the middle layer to instigate neural activity. This assumption can be further investigated by searching among competing models the best model that explains how laminar input enters a certain region and how it propagates to its connected regions.

5.4.4 Blood-draining laminar DCM implications for future laminar fMRI

Finally, we suggest some directions that laminar DCM accounting for blood-draining can be used or modified to simulate data that can help elucidate mechanisms underlying BOLD activation signals in laminr fMRI.

We have not yet tried to estimate haemodynamic parameters (v, q) and state variables. Moreover, currently we did not put much emphasis on interpreting the meaning of connectivity weights, and further model comparison would be required to check model invertibility or connectivity weight estimation.

We still do not have a full account of how input stimulus conditions modulate parameters in certain ways. Block-design task studies can be used to control

stimulation timings. As was done in other studies, shorter stimulations can be applied to capture more transient responses in the laminar level.

In theory, it is possible to define more biologically plausible within region (intrinsic) inter-laminar connections based on known neuroanatomical pathways. Neuronal dynamics can be expressed using an alternate neural mass model and laminar connections can be defined as being inhibitory or excitatory [144]. The wiring of supragranular pyramidal cells, granular spiny stellate cells, infragranular pyramidal cells, and inhibitory interneurons imply that there is cell-type specific separation of layers. Thus, communication between those neurons can be defined in a more detailed manner rather than defining layers that indicate locations of cell origin and termination. For instance, FFW connections from pyramidal cells in superficial layers towards layer IV of another region are excitatory connections. On the other hand, FB connections consist of excitatory input from deep pyramidal cells towards superficial pyramidal cells, as well as inhibitory input towards inhibitory interneurons, both towards a hierarchically lower region. If one were to adapt these neuronal populations to fit laminar definitions for layered-DCM in fMRI, model comparison would be required to identify the best model architecture.

5.5 Future plan

We plan to implement a fitting tool for multi-layer DCM and current simulations were required, beforehand. This was to deduce whether laminar DCM can be used to replace SEM that was used in our previous IR-BOLD combined approach for measuring laminar activity. Ideally, we would like to fit models that are tailored to more biophysically plausible signals; thus, we used simulations to check the suitability of laminar DCM. One problem encountered was unintuitive weights predicted when simulating IR data with DCMs and we would require careful inspection to map connectivity weights produced from SEM vs. DCM. This can be done through derivation of matching expressions convertible from SEM to DCM parameters. Otherwise, alternate models of DCMs should be studied to determine whether a more precise model of neuronal connections could explain laminar responses and

activity better. For instance, we can modify DCM to model neural masses of excitatory and inhibitory connections that can be done in bilinear model forms that assume neuronal populations to be a node instead of layers.

5.6 Summary

We state our observations when simulating data with multi-layer DCM and suggest ways that our findings can be used to shape future analyses in laminar BOLD studies.

- Blood-draining included in the model from Heinzle et al. [78] was extended into a multilayer DCM with blood-draining modelling and for greater blood-draining strengths magnitude of BOLD in superficial layers were substantially amplified.
- Blood-draining effects are likely to interfere with upper layers, thus we need additional care when conducting connectivity analyses on layer-specific activity.
- One of our goals for using DCMs was to try to use connectivity models that account for biological phenomenon such as blood-draining towards the cortical surface, later in actual laminar BOLD data. However, we observed that BOLD signals amplify towards upper layers with blood-draining and that data covariances were affected. In this regard, in the future we may not be able to fit SEMs for DCM simulated data or real BOLD data affected by blood-draining in upper layers with the intention of estimating connectivity. We can seek alternative ways such as using SEMs to account for haemodyanmics or can try fitting laminar DCMs to IR-BOLD data as next steps.

6
General Conclusion

Contents

6.1 Contributions of Thesis 135
6.2 Chapter specific future directions 137
6.2.1 IR-BOLD fMRI and modelling approaches 137
6.2.2 High-resolution diffusion to histology 138
6.3 Concluding remarks 138

6.1 Contributions of Thesis

This thesis focused on two methodological perspectives in studying human cortical layers using non-invasive MRI, each intending to study function and structure of cortical layers.

Even though we are currently not at the stage of analysing functional connectivity in terms of neuron-level connections, laminar fMRI introduced the possibility of assessing layer-specific connectivity in the human brain. This has become increasingly possible by virtue of advanced hardware that pushed fMRI to be performed in higher magnetic fields. However, such high-field MRI facilities are not widely accessible worldwide and there are technical impediments that still remain, such as poor SNR and the requirement of complex acquisition sequences. Instead

of attaining higher spatial resolution, we suggest and explore a new method that combines IR-BOLD acquisition combined with SEM modelling. This way, MRI signals were modulated depending on varying T1s across the cortex. We apply this IR-BOLD acquisition with connectivity modelling using SEMs to assess connectivity relationships of visual areas in a standard 3T MRI set-up. We show from our results that this new method allows us to assess hierarchical relationships between 9 visual areas based on assumptions on laminar connection patterns, and this was done without the need to spatially resolve cortical layers. It is, yet, difficult to assess areas beyond the visual system, which we chose in the current study for its relatively well-known hierarchical organisation. As tracer studies are not applicable to the human brain, we still lack a full description of hierarchical organisation of the human brain. Our method has potential to be extended to contribute to the study of other brain regions, but specific conditions such as variability of T1s throughout the entire cortex and assumptions on laminar connections require careful consideration when establishing the laminar SEM model.

Next, we seek to address what diffusion features could be related to known cytoarchitectural or laminar features in BigBrain histology. We used data sets that were some of the most high-resolution histology and dMRI data released at the time of writing. We explored the capability of predicting BigBrain histology from diffusion measures. This was done by seeking reliable diffusion or myelination (T1w/T2w) measures that could predict staining intensities from histological staining. Through our analyses on these data sets, we assessed effects of using a data-driven approach to summarise whole-brain histological staining intensity data. We conclude that we need to derive laminar features derived from BigBrain histology staining intensities (as opposed to using raw staining intensity values) to predict microstructural information as measured by dMRI. Using each modality to derive a cortical parcellation was difficult, so we turned to a secondary cortical type definition to bridge between the two modalities. However, even after computing global summary measures representing cortical layer properties from the raw intensity profiles, we were unable to identify similar cortical types. It could be that due to

resolution gaps between the two modalities laminar features can not be detected robustly for the entire brain (e.g. metrics would not be sensitive in non-sensorimotor areas with less significant laminar properties). Otherwise, it could be that we need more specific, sophisticated features in each or either of the modalities to conduct an improved joint-analysis of BigBrain and high-resolution diffusion data.

6.2 Chapter specific future directions

6.2.1 IR-BOLD fMRI and modelling approaches

There are several ways that work in Chapters 3 and 5 could be extended for new experimental directions. We could use specific task stimuli to evoke specific, existing connections using, for instance, visual illusions [83]. Since there are tasks known for instigating top-down processes targeting deep layers, we can anticipate which connections we would like to observe. We could also try a high-resolution IR approach so that we can try to access layer activity even in thinner slices (1 mm). There are a variety of systems that can also be inspected other than just visual regions, and in this perspective, we could also test connectivity relationships of regions involved in the cortico-thalamic loop.

A possible modification of the current work would be to fit both connectivity models and T1s to more thoroughly account for T1 variations across the cortex. Due to variable sizes of brain regions, brain regions would likely have different SNR values and T1 distributions could vary, so these conditions will require further investigation.

Finally, based on Chapter 5 results, we plan to fit multi-layer DCMs to IR-BOLD task fMRI or resting data. Fitting laminar DCM for IR-BOLD data was our initial aim, but observing that upper layers are largely affected by blood-draining we may also consider extending SEMs to account for haemodynamics.

6.2.2 High-resolution diffusion to histology

There are several other models that can be fitted in GM such as NODDI [148] that we could alternatively use to model microstructure in GM. We could also account for

more quantitative MRI methods that include quantitative T1, T2, magnetic transfer and others [39] to search for other microstructural features other than the more common diffusion measures. In an emerging study with extremely high-resolution data sets cortical area specific differences in microstructure and laminar details have started to be distinguished [149]. The study of cytoarchitecture is on the brink of exciting avenues as there are also other high-resolution multi-modal 3D reconstructed human brain resources that have become available [150]. We suggest that with a plethora of high-quality data, one should try to assess and apply analyses such as dimensionality reduction, clustering, and others to first study the data before applying the more advanced data mining techniques to search for significant features. In the long run, this area of research will greatly from more automated pipelines or harmonisation between data from different scanning environment or sizes.

6.3 Concluding remarks

There are several theoretical motivations underlying research endeavours in this thesis and not many of them are easy to address in the human brain. Longer-range connections between anatomically 'distant' cortical areas are difficult to probe but can in theory be studied if their precise laminar-specific connections can be identified. Moreover, hierarchical relationships between such distant areas may have distinguishable characteristics compared to connections between more local areas as seen in animal tracer studies. The classical canonical microcircuit theory that seems to be behind or driving laminar connections underlying hierarchical relationships may be constrained through the more distant long-range connections. However, even with increasingly higher-resolution MRI we are still struggling to quantify layers or columnar units both structurally and functionally at this era. We are still progressing with a multitude of imaging methods and modalities that are being used in an integrative manner. This thesis concludes by emphasising that there are a myriad of functional connectivity relationships and roles that are not fully understood in sensory and even more beyond sensory areas that regulate more high-order cognitive abilities in the human brain. It is still a difficult realm.

These can be further explored to elucidate what 'hierarchical relationships' actually exist and this is especially an exciting direction that is only starting to be tested on the human brain through non-invasive laminar fMRI. Neuroscience research will continuously require development of safe and accessible methods that can help elucidate how complexly organised neurons are connected and to identify which functional roles are elicited when those wired neurons are in action.

Appendices

A
Supplementary Figures

Figure A.1: Below we show the pseudocode for simulating data and estimating connectivity weights with SEM. All simulations were generated following this format. What is not specified in this code is specific model settings such as types of connection patterns such as FFW, FB, Van Essen and Felleman connections (1 FFW, 2 FB). Results shown in Figure 3.5 are examples of fitting SEMs when examining model invertibilty. Specifically, in our analyses, ground truth weights were usually set as random floating point values between [0, 1] for hierarchy test simulated data, but we also designate weights of wider ranges as in Figure 3.5.

```
// Define ROI pairs in 'pairs'
// Define connection patterns in 'pattern'

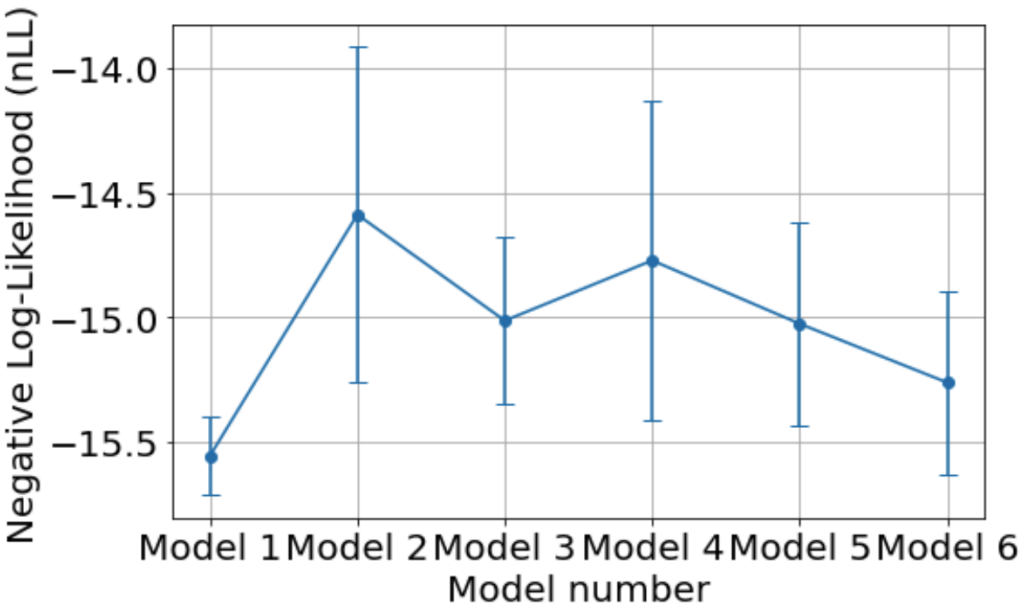
// Set parameters
N = 2; // Number of nodes
T = 160; // Number of samples/time points
L = 3; // Number of cortical layers

// Set TIs and T1s
// Call 'sem_create_Pmatrix(TIs, T1s, N)' to create
   partial volume matrix 'P'

// Set hyperparameters: priorMean and priorVariance
// Create structural matrix 'B' using 'fill_structMat(
   pairs, pattern, N, L)'
// Assign true weights (non-zero elements) to 'B' to get
    weight matrix 'A'

// Simulate ground truth data 'data'
for ti in TIs:
    Initialize matrix 'x' as zeros of size (N * L, T)
    for k in range(T):
        Generate noise vector 'z' as sigma * randn(N * L
           , 1)
        Compute 'x(:, k)' as inv(eye(N * L) - A) * z
    end for
    data{ti} = P{ti} * x;
end for

// Call 'sem_fit_mh(data, B, P, [priorMean,
   priorVariance])' to get 'samples', 'est_A', and 'data
   covariance S'
```




```
Comparison of Model 1 vs. Model 2: Paired t-test p-value = 1.130e-15
Comparison of Model 1 vs. Model 3: Paired t-test p-value = 4.472e-17
Comparison of Model 1 vs. Model 4: Paired t-test p-value = 3.619e-13
Comparison of Model 1 vs. Model 5: Paired t-test p-value = 2.079e-13
Comparison of Model 1 vs. Model 6: Paired t-test p-value = 4.627e-09
```

Figure A.2: We checked statistical significance of the schematic in Figure 3.5 using a paired t-test. When comparing the ground-truth (correct) model negative log-likelihood (nLL) value against those of the other 5 models, p-values were always $p<0.05$ as shown above. For the purpose of running statistical tests, we ran this hierarchy test on 50 data sets simulated using randomly drawn connectivity weights between [0, 1). We also used T1s 800, 1000, 1200 and TIs at 300, 480, 690, 810 and such model settings were slightly different from the random data sets in the main figure (Figure 3.5). Nevertheless, the frequency that the correct model is identified with lowest nLL remains the same. In this plot we again show mean nLL and standard deviations as was done the figure from main text.

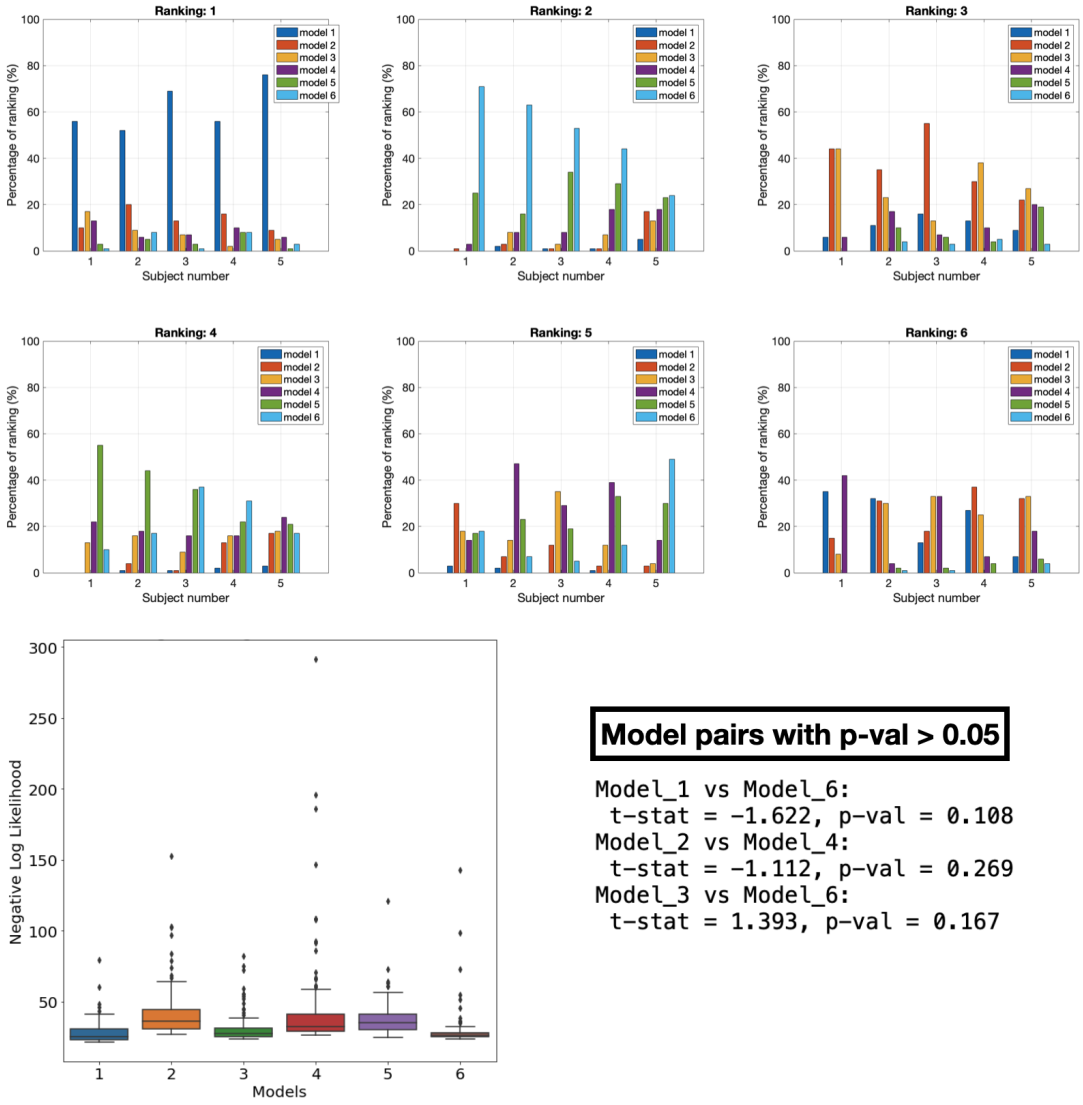


Figure A.3: Here, we run the hierarchy test for 5 subject data whilst accounting for variable T1s. 100 sets of T1s were randomly sampled from Gaussians fit around approximate peaks identified in a sample T1 map. The top figure corresponds to Figure 3.8, but random T1 data were from from a separate run of the hierarchy test. In the main text we focused on demonstrating the frequency that a model is identified as best fitting model, therefore, we count and rank in order of highest to lowest the model with lowest nLL for each subject. In this additional figure, we also checked original nLL values by performing a paired t-test on nLL averaged over subjects. Although models 1 and 6 were each ranked as 1st and 2nd in terms of having lowest nLL, statistical tests on the actual nLL values indicate that differences between those 2 models are not significant ($p>0.05$). On the left bottom we show mean over subjects and spread of nLLs of each model in the 100 random T1 sets and this was used to conduct a paired t-test on.

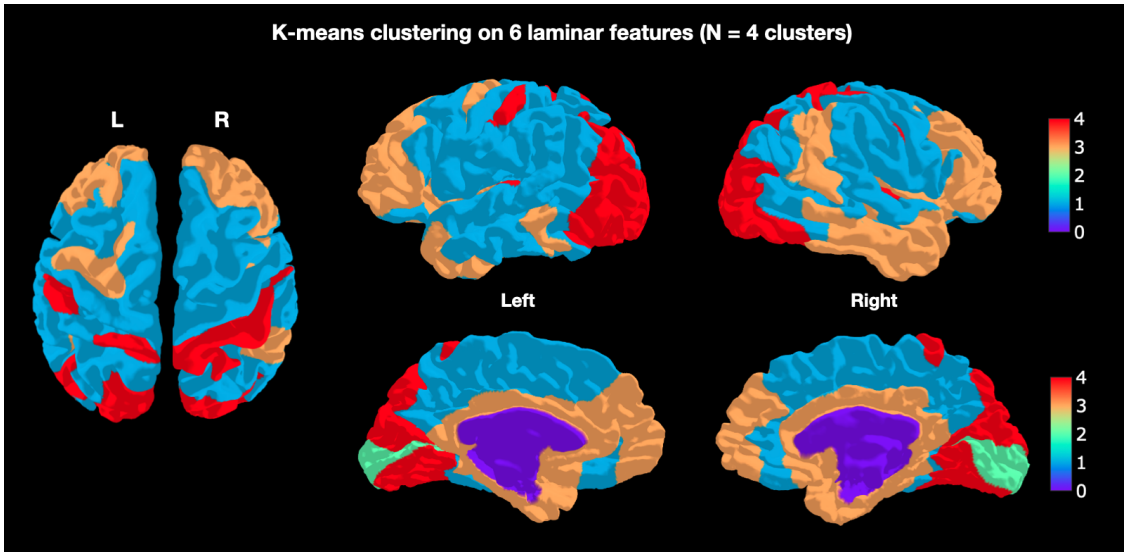
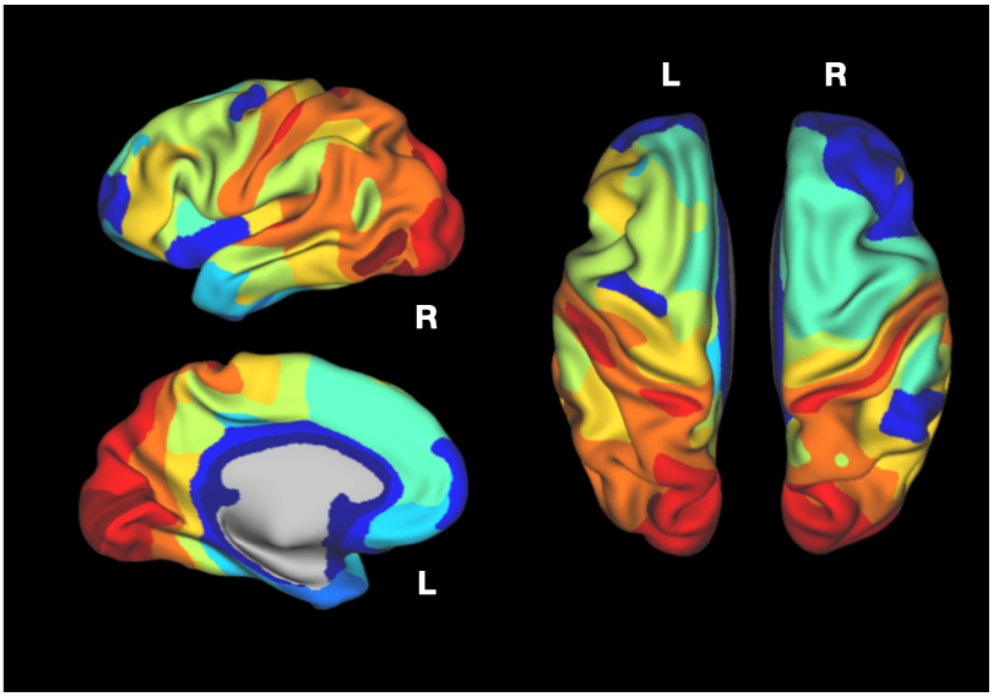
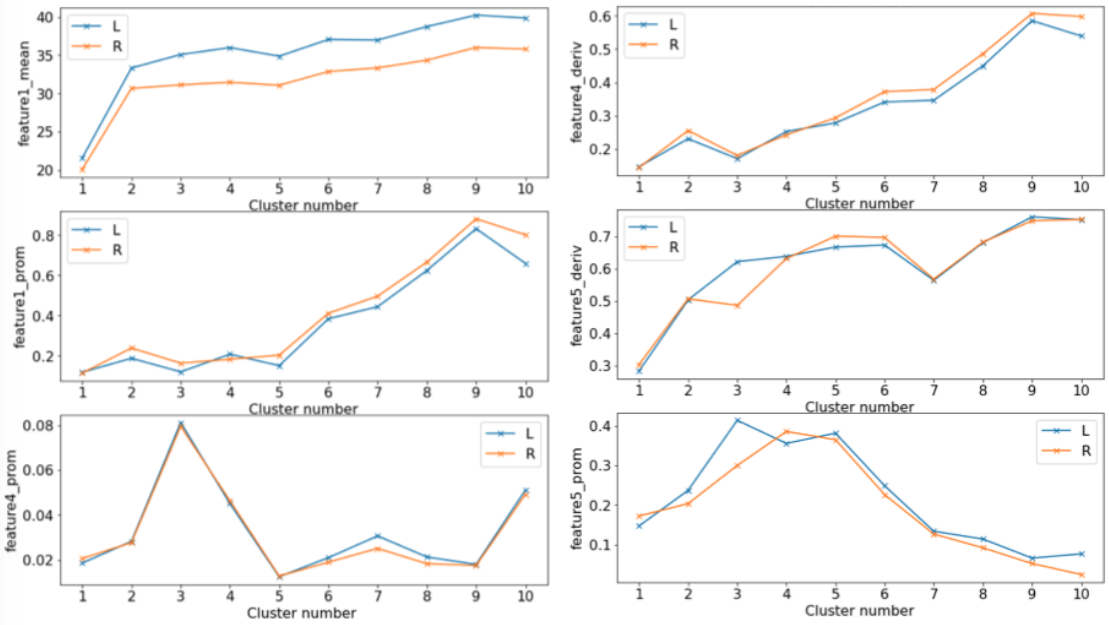


Figure A.4: In the main text, we show K-means clustering results when using 7 clusters. This was intentionally done to approximately match the number of cortical types of the reference literature. Here we show another instance where total cluster numbers were much less (N = 4).



(a) N = 10 clusters in K-means clustering of laminar features



(b) Mean feature values for each cluster

Figure A.5: We also tested granularity of clusters when increasing the number of clusters to 10 clusters. We also show the average feature computations.

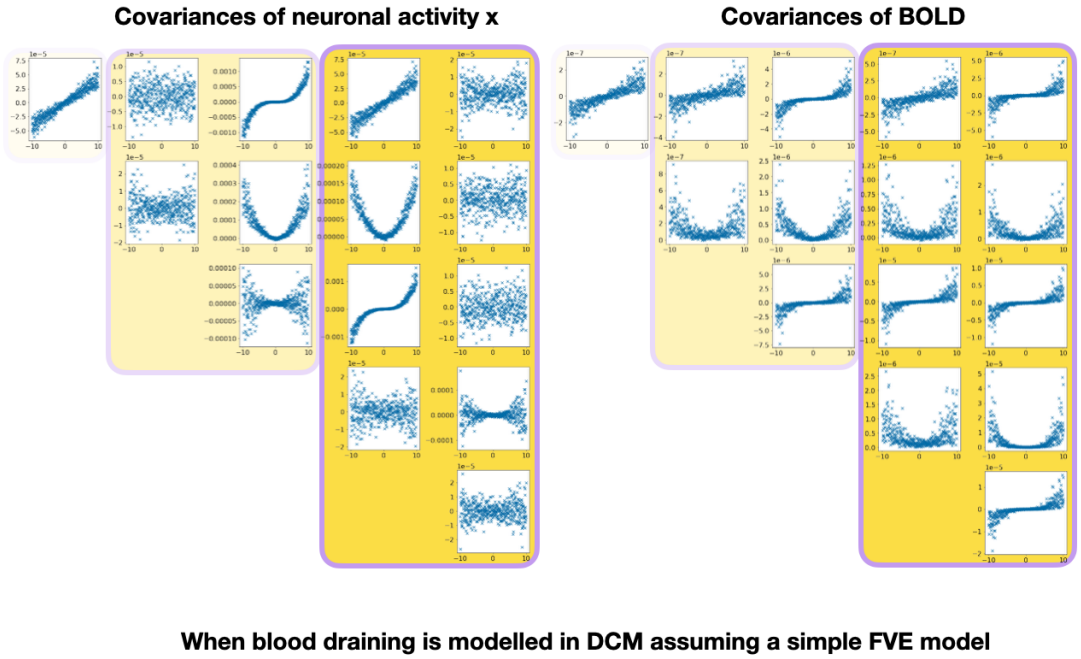


Figure A.6: Here we show the original layout of covariances shown in Figure 5.10a and 5.10b. Layers in a model can be considered separate nodes with individual data, hence, in our model we can observe covariances between 15 unique pairs of nodes. Here, we use an alternative arrangement and colour coding scheme to assist understanding and simply highlight the location of layers (upper, middle, lower) involved in the designated covariance plot.

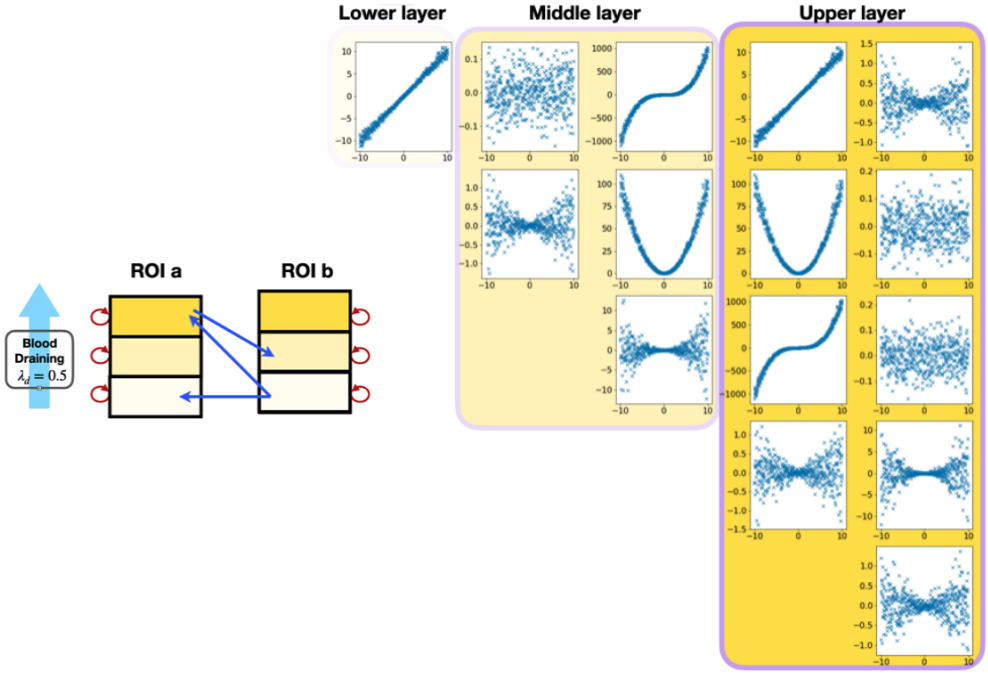


Figure A.7: Same applies to the diagram that shows covariances from multi-layer SEM that was visited in Figure 5.11

References

[1] Stefan Geyer et al. “Microstructural Parcellation of the Human Cerebral Cortex – From Brodmann’s Post-Mortem Map to in vivo Mapping with High-Field Magnetic Resonance Imaging”. English. In: *Frontiers in Human Neuroscience* 5 (Feb. 2011).

[2] K. Zilles, N. Palomero-Gallagher, and K. Amunts. “Myeloarchitecture and Maps of the Cerebral Cortex”. In: *Brain Mapping*. Ed. by Arthur W. Toga. Waltham: Academic Press, Jan. 2015, pp. 137–156.

[3] Korbinian Brodmann. *Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues*. Barth, 1909.

[4] Karl Zilles and Katrin Amunts. “Centenary of Brodmann’s map — conception and fate”. en. In: *Nature Reviews Neuroscience* 11.2 (Feb. 2010). Pp. 139–145.

[5] PT Jean Talairach and P Tournoux. *Co-planar stereotaxic atlas of the human brain*. 1988.

[6] Lazaros C. Triarhou. “The Cytoarchitectonic Map of Constantin von Economo and Georg N. Koskinas”. en. In: *Microstructural Parcellation of the Human Cerebral Cortex: From Brodmann’s Post-Mortem Map to in Vivo Mapping with High-Field Magnetic Resonance Imaging*. Ed. by Stefan Geyer and Robert Turner. Berlin, Heidelberg: Springer, 2013, pp. 33–53.

[7] Constantin Freiherr von Economo and Georg N. Koskinas. *Die cytoarchitektonik der hirnrinde des erwachsenen menschen*. J. Springer, 1925.

[8] K. S. Lashley and George Clark. “The cytoarchitecture of the cerebral cortex of ateles: A critical examination of architectonic studies”. en. In: *Journal of Comparative Neurology* 85.2 (1946). Pp. 223–305.

[9] R. Kötter et al. “A probabilistic atlas and reference system for the human brain: International Consortium for Brain Mapping (ICBM)”. In: *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences* 356.1412 (Aug. 2001). Pp. 1293–1322.

[10] H. B. M. Uylings et al. “Consequences of large interindividual variability for human brain atlases: converging macroscopical imaging and microscopical neuroanatomy”. en. In: *Anatomy and Embryology* 210.5 (Dec. 2005), pp. 423–431.

[11] Katrin Amunts et al. “Julich-Brain: A 3D probabilistic atlas of the human brain’s cytoarchitecture”. In: *Science* 369.6506 (Aug. 2020). Pp. 988–992.

[12] Simon B. Eickhoff et al. “A new SPM toolbox for combining probabilistic cytoarchitectonic maps and functional imaging data”. In: *NeuroImage* 25.4 (May 2005), pp. 1325–1335.

[13] Katrin Amunts et al. “BigBrain: An Ultrahigh-Resolution 3D Human Brain Model”. en. In: *Science* 340.6139 (June 2013), pp. 1472–1475.

[14] Konrad Wagstyl et al. “BigBrain 3D atlas of cortical layers: Cortical and laminar thickness gradients diverge in sensory and motor cortices”. en. In: *PLOS Biology* 18.4 (Apr. 2020). e3000678.

[15] Konrad Wagstyl et al. “Mapping Cortical Laminar Structure in the 3D BigBrain”. In: *Cerebral Cortex* 28.7 (July 2018), pp. 2551–2562.

[16] Bruce Fischl and Martin I. Sereno. “Microstructural parcellation of the human brain”. In: *NeuroImage*. Microstructural Imaging 182 (Nov. 2018), pp. 219–231.

[17] Alard Roebroeck, Karla L. Miller, and Manisha Aggarwal. “Ex vivo diffusion MRI of the human brain: Technical challenges and recent advances”. en. In: *NMR in Biomedicine* 32.4 (2019). e3941.

[18] Kurt G. Schilling et al. *Recommendations and guidelines from the ISMRM Diffusion Study Group for preclinical diffusion MRI: Part 2 – Ex vivo imaging*. Feb. 2023.

[19] Brian L. Edlow et al. “7 Tesla MRI of the ex vivo human brain at 100 micron resolution”. en. In: *Scientific Data* 6.1 (Oct. 2019). P. 244.

[20] Casey Paquola et al. “Microstructural and functional gradients are increasingly dissociated in transmodal cortices”. en. In: *PLOS Biology* 17.5 (May 2019). e3000284.

[21] Christian Schiffer et al. “Convolutional neural networks for cytoarchitectonic brain mapping at large scale”. In: *NeuroImage* 240 (Oct. 2021), p. 118327.

[22] Kai Kiwitz et al. “Deep learning networks reflect cytoarchitectonic features used in brain mapping”. en. In: *Scientific Reports* 10.1 (Dec. 2020). P. 22039.

[23] Linden Parkes et al. “Asymmetric signaling across the hierarchy of cytoarchitecture within the human connectome”. In: *Science Advances* 8.50 (Dec. 2022). eadd2185.

[24] Ittai Shamir and Yaniv Assaf. “Expanding connectomics to the laminar level: A perspective”. In: *Network Neuroscience* 7.2 (June 2023), pp. 377–388.

[25] Casey Paquola et al. “Closing the mechanistic gap: the value of microarchitecture in understanding cognitive networks”. In: *Trends in Cognitive Sciences* 26.10 (Oct. 2022), pp. 873–886.

[26] Elliott M. Marcus, Stanley Jacobson, and Thomas D. Sabin. “325Cerebral Cortex: Cytoarchitecture, Physiology, and Overview of Functional Localization”. In: *Integrated Neuroscience and Neurology: A Clinical Case History Problem Solving Approach*. Ed. by Elliott M. Marcus, Stanley Jacobson, and Thomas D. Sabin. Oxford University Press, Aug. 2014, p. 0.

[27] Nicola Palomero-Gallagher and Karl Zilles. “Cortical layers: Cyto-, myelo-, receptor- and synaptic architecture in human cortical areas”. In: *NeuroImage* 197 (Aug. 2019), pp. 716–741.

[28] Jules Gabriel Fran{\c{c}}ois Baillarger. *Recherches sur la structure de la couche corticale des circonvolutions du cerveau*. chez J.-B. Baillière, 1840.

[29] Holly Bridge and Stuart Clare. “High-resolution MRI: in vivo histology?” In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 361.1465 (Nov. 2005). Pp. 137–146.

[30] Cécile Vogt and Oskar Vogt. “Allgemeinere Ergebnisse unserer Hirnforschung”. In: *J. Psychol. Neurol.* 25 (1919), pp. 292–398.

[31] Adolf Hopf. “Über die Verteilung myeloarchitektonischer Merkmale in der Stirnhirnrinde beim Menschen”. de. In: (1956).

[32] Friedrich Sanides. “Structure and function of the human frontal lobe”. In: *Neuropsychologia* 2.3 (Nov. 1964), pp. 209–219.

[33] Rudolf Nieuwenhuys. “The myeloarchitectonic studies on the human cerebral cortex of the Vogt–Vogt school, and their significance for the interpretation of functional neuroimaging data”. en. In: *Brain Structure and Function* 218.2 (Mar. 2013), pp. 303–352.

[34] Rudolf Nieuwenhuys, Cees A. J. Broere, and Leonardo Cerliani. “A new myeloarchitectonic map of the human neocortex based on data from the Vogt–Vogt school”. en. In: *Brain Structure and Function* 220.5 (Sept. 2015), pp. 2551–2573.

[35] Rosa M. Sánchez-Panchuelo et al. “Correspondence of human visual areas identified using functional and anatomical MRI in vivo at 7 T”. en. In: *Journal of Magnetic Resonance Imaging* 35.2 (2012). Pp. 287–299.

[36] J. Dinse et al. “A cytoarchitecture-driven myelin model reveals area-specific signatures in human primary and secondary areas using ultra-high resolution in-vivo brain MRI”. In: *NeuroImage* 114 (July 2015), pp. 71–87.

[37] Matthew F. Glasser and David C. Van Essen. “Mapping Human Cortical Areas In Vivo Based on Myelin Content as Revealed by T1- and T2-Weighted MRI”. en. In: *Journal of Neuroscience* 31.32 (Aug. 2011). Pp. 11597–11616.

[38] Hikaru Fukutomi et al. “Neurite imaging reveals microstructural variations in human cerebral cortical gray matter”. In: *NeuroImage*. Microstructural Imaging 182 (Nov. 2018), pp. 488–499.

[39] Nikolaus Weiskopf et al. “Quantitative magnetic resonance imaging of brain anatomy and in vivo histology”. en. In: *Nature Reviews Physics* 3.8 (Aug. 2021). Pp. 570–588.

[40] Stewart Shipp. “Structure and function of the cerebral cortex”. en. In: *Current Biology* 17.12 (June 2007), R443–R449.

[41] Kathleen S. Rockland. “What do we know about laminar connectivity?” In: *NeuroImage* 197 (Aug. 2019), pp. 772–784.

[42] Kathleen S. Rockland and Deepak N. Pandya. “Laminar origins and terminations of cortical connections of the occipital lobe in the rhesus monkey”. In: *Brain Research* 179.1 (Dec. 1979), pp. 3–20.

[43] Nikola T. Markov et al. “Anatomy of hierarchy: Feedforward and feedback pathways in macaque visual cortex”. en. In: *Journal of Comparative Neurology* 522.1 (2014). Pp. 225–259.

[44] Nikola T. Markov et al. “Cortical High-Density Counterstream Architectures”. In: *Science* 342.6158 (Nov. 2013). P. 1238406.

[45] Julien Vezoli et al. “Cortical hierarchy, dual counterstream architecture and the importance of top-down generative networks”. In: *NeuroImage* 225 (Jan. 2021), p. 117479.

[46] Kathleen S. Rockland. “Notes on Visual Cortical Feedback and Feedforward Connections”. In: *Frontiers in Systems Neuroscience* 16 (2022).

[47] Farran Briggs. “Organizing principles of cortical layer 6”. English. In: *Frontiers in Neural Circuits* 4 (Feb. 2010).

[48] Andre M. Bastos et al. “Canonical Microcircuits for Predictive Coding”. In: *Neuron* 76.4 (Nov. 2012), pp. 695–711.

[49] R J Douglas and K A Martin. “A functional microcircuit for cat visual cortex.” en. In: *The Journal of Physiology* 440.1 (1991). Pp. 735–769.

[50] David C. Godlove et al. “Microcircuitry of Agranular Frontal Cortex: Testing the Generality of the Canonical Cortical Microcircuit”. en. In: *Journal of Neuroscience* 34.15 (Apr. 2014). Pp. 5355–5369.

[51] Ittai Shamir and Yaniv Assaf. “An MRI-Based, Data-Driven Model of Cortical Laminar Connectivity”. en. In: *Neuroinformatics* 19.2 (Apr. 2021), pp. 205–218.

[52] Daniel J. Felleman and David C. Van Essen. “Distributed Hierarchical Processing in the Primate Cerebral Cortex”. In: *Cerebral Cortex* 1.1 (Jan. 1991), pp. 1–47.

[53] Martin Havlicek et al. “Physiologically informed dynamic causal modeling of fMRI data”. In: *NeuroImage* 122 (Nov. 2015), pp. 355–372.

[54] Jonathan R. Polimeni and Laura D. Lewis. “Imaging faster neural dynamics with fast fMRI: A need for updated models of the hemodynamic response”. In: *Progress in Neurobiology*. How high spatiotemporal resolution fMRI can advance neuroscience 207 (Dec. 2021), p. 102174.

[55] H. M. Duvernoy, S. Delon, and J. L. Vannson. “Cortical blood vessels of the human brain”. In: *Brain Research Bulletin* 7.5 (Nov. 1981), pp. 519–579.

[56] Bruno Weber et al. “The Microvascular System of the Striate and Extrastriate Visual Cortex of the Macaque”. In: *Cerebral Cortex* 18.10 (Oct. 2008), pp. 2318–2330.

[57] Alexander John Poplawsky, Mitsuhiro Fukuda, and Seong-Gi Kim. “Foundations of layer-specific fMRI and investigations of neurophysiological activity in the laminarized neocortex and olfactory bulb of animal models”. In: *NeuroImage* 199 (Oct. 2019), pp. 718–729.

[58] Franca Schmid et al. “Depth-dependent flow and pressure characteristics in cortical microvascular networks”. en. In: *PLOS Computational Biology* 13.2 (Feb. 2017). e1005392.

[59] Franca Schmid et al. “Vascular density and distribution in neocortex”. In: *NeuroImage* 197 (Aug. 2019), pp. 792–805.

[60] Pablo Blinder et al. “The cortical angiome: an interconnected vascular network with noncolumnar patterns of blood flow”. en. In: *Nature Neuroscience* 16.7 (July 2013). Pp. 889–897.

[61] Sriranga Kashyap et al. “Resolving laminar activation in human V1 using ultra-high spatial resolution fMRI at 7T”. en. In: *Scientific Reports* 8.1 (Nov. 2018). P. 17063.

[62] Laurentius Huber et al. “Ultra-high resolution blood volume fMRI and BOLD fMRI in humans at 9.4 T: Capabilities and challenges”. In: *NeuroImage* 178 (Sept. 2018), pp. 769–779.

[63] Peter A Bandettini, Laurentius Huber, and Emily S Finn. “Challenges and opportunities of mesoscopic brain mapping with fMRI”. In: *Current Opinion in Behavioral Sciences*. Deep Imaging - Personalized Neuroscience 40 (Aug. 2021), pp. 189–200.

[64] Samuel J. D. Lawrence et al. “Laminar fMRI: Applications for cognitive neuroscience”. In: *NeuroImage* 197 (Aug. 2019), pp. 785–791.

[65] Robert Trampel et al. “In-vivo magnetic resonance imaging (MRI) of laminae in the human cortex”. In: *NeuroImage* 197 (Aug. 2019), pp. 707–715.

[66] Holly Bridge et al. “Independent anatomical and functional measures of the V1/V2 boundary in human visual cortex”. In: *Journal of Vision* 5.2 (Feb. 2005), p. 1.

[67] Peter J. Koopmans, Markus Barth, and David G. Norris. “Layer-specific BOLD activation in human V1”. en. In: *Human Brain Mapping* 31.9 (2010). Pp. 1297–1304.

[68] Rolf Pohmann, Oliver Speck, and Klaus Scheffler. “Signal-to-noise ratio and MR tissue parameters in human brain imaging at 3, 7, and 9.4 tesla using current receive coil arrays”. en. In: *Magnetic Resonance in Medicine* 75.2 (2016). Pp. 801–809.

[69] Laurentius Huber et al. “High-Resolution CBV-fMRI Allows Mapping of Laminar Activity and Connectivity of Cortical Input and Output in Human M1”. In: *Neuron* 96.6 (Dec. 2017), 1253–1263.e7.

[70] Natalia Zaretskaya et al. “Advantages of cortical surface reconstruction using submillimeter 7 T MEMPRAGE”. In: *NeuroImage* 165 (Jan. 2018), pp. 11–26.

[71] Behroze Vachha and Susie Y. Huang. “MRI with ultrahigh field strength and high-performance gradients: challenges and opportunities for clinical neuroimaging at 7 T and beyond”. In: *European Radiology Experimental* 5.1 (Aug. 2021), p. 35.

[72] Olivia Viessmann and Jonathan R Polimeni. “High-resolution fMRI at 7Tesla: challenges, promises and recent developments for individual-focused fMRI studies”. In: *Current Opinion in Behavioral Sciences*. Deep Imaging - Personalized Neuroscience 40 (Aug. 2021), pp. 96–104.

[73] Hanzhang Lu and Peter C. M. van Zijl. “Experimental measurement of extravascular parenchymal BOLD effects and tissue oxygen extraction fractions using multi-echo VASO fMRI at 1.5 and 3.0 T”. en. In: *Magnetic Resonance in Medicine* 53.4 (2005). Pp. 808–816.

[74] Jonathan R. Polimeni et al. “Laminar analysis of 7T BOLD using an imposed spatial activation pattern in human V1”. In: *NeuroImage* 52.4 (Oct. 2010), pp. 1334–1346.

[75] Robert Turner. “How Much Cortex Can a Vein Drain? Downstream Dilution of Activation-Related Cerebral Blood Oxygenation Changes”. In: *NeuroImage* 16.4 (Aug. 2002), pp. 1062–1067.

[76] Federico De Martino et al. “Cortical Depth Dependent Functional Responses in Humans at 7T: Improved Specificity with 3D GRASE”. en. In: *PLOS ONE* 8.3 (Mar. 2013). e60514.

[77] Jeroen C. W. Siero et al. “BOLD Specificity and Dynamics Evaluated in Humans at 7 T: Comparing Gradient-Echo and Spin-Echo Hemodynamic Responses”. en. In: *PLOS ONE* 8.1 (Jan. 2013). e54560.

[78] Jakob Heinzle et al. “A hemodynamic model for layered BOLD signals”. In: *NeuroImage* 125 (Jan. 2016), pp. 556–570.

[79] Hanzhang Lu et al. “Functional magnetic resonance imaging based on changes in vascular space occupancy”. en. In: *Magnetic Resonance in Medicine* 50.2 (2003). Pp. 263–274.

[80] Hanzhang Lu, Jun Hua, and Peter C. M. van Zijl. “Noninvasive functional imaging of cerebral blood volume with vascular-space-occupancy (VASO) MRI”. en. In: *NMR in Biomedicine* 26.8 (2013). Pp. 932–948.

[81] Laurentius Huber et al. “Layer-dependent functional connectivity methods”. In: *Progress in Neurobiology*. How high spatiotemporal resolution fMRI can advance neuroscience 207 (Dec. 2021), p. 101835.

[82] Sebastian Dresbach et al. *Characterisation of laminar and vascular spatiotemporal dynamics of CBV and BOLD signals using VASO and ME-GRE at 7T in humans*. en. Jan. 2024.

[83] Peter Kok et al. “Selective Activation of the Deep Layers of the Human Primary Visual Cortex by Top-Down Feedback”. In: *Current Biology* 26.3 (Feb. 2016), pp. 371–376.

[84] Lars Muckli et al. “Contextual Feedback to Superficial Layers of V1”. In: *Current Biology* 25.20 (Oct. 2015), pp. 2690–2695.

[85] Peter J. Koopmans et al. “Multi-echo fMRI of the cortical laminae in humans at 7T”. In: *NeuroImage* 56.3 (June 2011), pp. 1276–1285.

[86] Laura D. Lewis et al. “Stimulus-dependent hemodynamic response timing across the human subcortical-cortical visual pathway identified through high spatiotemporal resolution 7T fMRI”. In: *NeuroImage* 181 (Nov. 2018), pp. 279–291.

[87] Natalia Petridou and Jeroen C. W. Siero. “Laminar fMRI: What can the time domain tell us?” In: *NeuroImage* 197 (Aug. 2019), pp. 761–771.

[88] René Scheeringa et al. “The relationship between oscillatory EEG activity and the laminar-specific BOLD signal”. In: *Proceedings of the National Academy of Sciences* 113.24 (June 2016). Pp. 6761–6766.

[89] Martin Havlicek and Kâmil Uludağ. “A dynamical model of the laminar BOLD response”. In: *NeuroImage* 204 (Jan. 2020), p. 116209.

[90] Kamil Uludag and Martin Havlicek. "Determining laminar neuronal activity from BOLD fMRI using a generative model". In: *Progress in Neurobiology*. How high spatiotemporal resolution fMRI can advance neuroscience 207 (Dec. 2021), p. 102055.

[91] A. M. Bastos et al. "A DCM study of spectral asymmetries in feedforward and feedback connections between visual areas V1 and V4 in the monkey". In: *NeuroImage* 108 (Mar. 2015), pp. 460–475.

[92] Robin M. Heidemann et al. "Diffusion imaging in humans at 7T using readout-segmented EPI and GRAPPA". en. In: *Magnetic Resonance in Medicine* 64.1 (2010). Pp. 9–14.

[93] Jennifer A. McNab et al. "Surface based analysis of diffusion orientation for identifying architectonic domains in the in vivo human cortex". In: *NeuroImage* 69 (Apr. 2013), pp. 87–100.

[94] Colin Reveley et al. "Diffusion MRI anisotropy in the cerebral cortex is determined by unmyelinated tissue features". en. In: *Nature Communications* 13.1 (Nov. 2022). P. 6702.

[95] Robin M. Heidemann et al. "k-space and q-space: Combining ultra-high spatial and angular resolution in diffusion imaging using ZOOPPA at 7T". In: *NeuroImage* 60.2 (Apr. 2012), pp. 967–978.

[96] Sune Nørhøj Jespersen et al. "Determination of Axonal and Dendritic Orientation Distributions Within the Developing Cerebral Cortex by Diffusion Tensor Imaging". In: *IEEE Transactions on Medical Imaging* 31.1 (Jan. 2012). Pp. 16–32.

[97] C. W. U. Leuze et al. "Layer-Specific Intracortical Connectivity Revealed with Diffusion MRI". en. In: *Cerebral Cortex* 24.2 (Feb. 2014), pp. 328–339.

[98] Michiel Kleinnijenhuis et al. "Layer-specific diffusion weighted imaging in human primary visual cortex in vitro". en. In: *Cortex* 49.9 (Oct. 2013), pp. 2569–2582.

[99] Alberto De Luca et al. "Spherical deconvolution with tissue-specific response functions and multi-shell diffusion MRI to estimate multiple fiber orientation distributions (mFODs)". In: *NeuroImage* 222 (Nov. 2020), p. 117206.

[100] Tharindu Ganepola et al. "Using diffusion MRI to discriminate areas of cortical grey matter". In: *NeuroImage*. Microstructural Imaging 182 (Nov. 2018), pp. 456–468.

[101] Daniel Barazany and Yaniv Assaf. "Visualization of Cortical Lamination Patterns with Magnetic Resonance Imaging". In: *Cerebral Cortex* 22.9 (Sept. 2012), pp. 2016–2023.

[102] Shlomi Lifshits et al. "Resolution considerations in imaging of the cortical layers". In: *NeuroImage*. Pushing the spatio-temporal limits of MRI and fMRI 164 (Jan. 2018), pp. 112–120.

[103] Ittai Shamir et al. "A framework for cortical laminar composition analysis using low-resolution T1 MRI images". en. In: *Brain Structure and Function* 224.4 (May 2019), pp. 1457–1467.

[104] Saad Jbabdi et al. "Measuring macroscopic brain connections in vivo". en. In: *Nature Neuroscience* 18.11 (Nov. 2015). Pp. 1546–1555.

[105] Mark D. Grier, Jan Zimmermann, and Sarah R. Heilbronner. "Estimating Brain Connectivity With Diffusion-Weighted Magnetic Resonance Imaging: Promise and Peril". In: *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging* 5.9 (2020), pp. 846–854.

[106] Emily S. Finn, Laurentius Huber, and Peter A. Bandettini. "Higher and deeper: Bringing layer fMRI to association cortex". In: *Progress in Neurobiology* 207 (2021). P. 101930.

[107] C Büchel and K J Friston. "Modulation of connectivity in visual pathways by attention: cortical interactions evaluated with structural equation modelling and fMRI." In: *Cerebral Cortex* 7.8 (Dec. 1997), pp. 768–778.

[108] Matthew F. Glasser et al. "A multi-modal parcellation of human cerebral cortex". en. In: *Nature* 536.7615 (Aug. 2016). Pp. 171–178.

[109] Rinaldo D. D'Souza et al. "Hierarchical and nonhierarchical features of the mouse visual cortical network". en. In: *Nature Communications* 13.1 (Jan. 2022). P. 503.

[110] D. H. Hubel and T. N. Wiesel. "Receptive fields, binocular interaction and functional architecture in the cat's visual cortex". en. In: *The Journal of Physiology* 160.1 (1962). Pp. 106–154.

[111] Nikola T. Markov et al. "The role of long-range connections on the specificity of the macaque interareal cortical network". In: *Proceedings of the National Academy of Sciences* 110.13 (Mar. 2013). Pp. 5187–5192.

[112] Katrin Amunts and Karl Zilles. "Architectonic Mapping of the Human Brain beyond Brodmann". en. In: *Neuron* 88.6 (Dec. 2015), pp. 1086–1107.

[113] Trong-Kha Truong, Arnaud Guidon, and Allen W. Song. "Cortical Depth Dependence of the Diffusion Anisotropy in the Human Cortical Gray Matter In Vivo". en. In: *PLOS ONE* 9.3 (Mar. 2014). e91424.

[114] Sune N. Jespersen et al. "Modeling dendrite density from magnetic resonance diffusion measurements". en. In: *NeuroImage* 34.4 (Feb. 2007), pp. 1473–1486.

[115] Björn Merker. "Silver staining of cell bodies by means of physical development". en. In: *Journal of Neuroscience Methods* 9.3 (Nov. 1983), pp. 235–241.

[116] Casey Paquola et al. "The BigBrainWarp toolbox for integration of BigBrain 3D histology with multimodal neuroimaging". In: *eLife* 10 (Aug. 2021). Ed. by Saad Jbabdi et al. e70119.

[117] M. D. Waehnert et al. "Anatomically motivated modeling of cortical laminae". en. In: *NeuroImage*. In-vivo Brodmann Mapping of the Human Brain 93 (June 2014), pp. 210–220.

[118] Fuyixue Wang et al. "In vivo human whole-brain Connectom diffusion MRI dataset at 760 µm isotropic resolution". en. In: *Scientific Data* 8.1 (Apr. 2021). P. 122.

[119] Miguel Ángel García-Cabezas, Julia Liao Hacker, and Basilis Zikopoulos. "A Protocol for Cortical Type Analysis of the Human Neocortex Applied on Histological Samples, the Atlas of Von Economo and Koskinas, and Magnetic Resonance Imaging". In: *Frontiers in Neuroanatomy* 14 (2020).

[120] Helen Barbas. “General Cortical and Special Prefrontal Connections: Principles from Structure to Function”. en. In: *Annual Review of Neuroscience* 38.Volume 38, 2015 (July 2015). Pp. 269–289.

[121] Miguel Ángel García-Cabezas, Basilis Zikopoulos, and Helen Barbas. “The Structural Model: a theory linking connections, plasticity, pathology, development and evolution of the cerebral cortex”. en. In: *Brain Structure and Function* 224.3 (Apr. 2019), pp. 985–1008.

[122] Lianne H. Scholtens et al. “An MRI Von Economo – Koskinas atlas”. In: *NeuroImage*. Segmenting the Brain 170 (Apr. 2018), pp. 249–256.

[123] Claus C. Hilgetag et al. “The primate connectome in context: Principles of connections of the cortical visual system”. en. In: *NeuroImage* 134 (July 2016), pp. 685–702.

[124] Jiahe Zhang et al. “Topography Impacts Topology: Anatomically Central Areas Exhibit a “High-Level Connector” Profile in the Human Cortex”. In: *Cerebral Cortex* 30.3 (Mar. 2020), pp. 1357–1365.

[125] Martijn P. van den Heuvel et al. “Bridging Cytoarchitectonics and Connectomics in Human Cerebral Cortex”. en. In: *Journal of Neuroscience* 35.41 (Oct. 2015). Pp. 13943–13948.

[126] Luke J Edwards et al. “Modelling radial and tangential fibres in the neocortex”. en. In: *ISMRM* (2016).

[127] Yaniv Assaf. “Imaging laminar structures in the gray matter with diffusion MRI”. In: *NeuroImage* 197 (Aug. 2019), pp. 677–688.

[128] Serge O. Dumoulin et al. “Ultra-high field MRI: Advancing systems neuroscience towards mesoscopic human brain function”. In: *NeuroImage*. Neuroimaging with Ultra-high Field MRI: Present and Future 168 (Mar. 2018), pp. 345–357.

[129] SoHyun Han et al. “Improved laminar specificity and sensitivity by combining SE and GE BOLD signals”. In: *NeuroImage* 264 (Dec. 2022), p. 119675.

[130] Irati Markuerkiaga et al. “Estimation of laminar BOLD activation profiles using deconvolution with a physiological point spread function”. In: *Journal of Neuroscience Methods* 353 (Apr. 2021), p. 109095.

[131] Pedro A. Valdes-Sosa et al. “Effective connectivity: Influence, causality and biophysical modeling”. In: *NeuroImage* 58.2 (Sept. 2011), pp. 339–361.

[132] Klaas Enno Stephan et al. “Comparing hemodynamic models with DCM”. In: *NeuroImage* 38.3 (Nov. 2007), pp. 387–401.

[133] Klaas Enno Stephan et al. “Nonlinear dynamic causal models for fMRI”. In: *NeuroImage* 42.2 (Aug. 2008), pp. 649–662.

[134] W. D. Penny et al. “Modelling functional integration: a comparison of structural equation and dynamic causal models”. In: *NeuroImage*. Mathematics in Brain Imaging 23 (Jan. 2004), S264–S274.

[135] K. J. Friston et al. “Nonlinear Responses in fMRI: The Balloon Model, Volterra Kernels, and Other Hemodynamics”. In: *NeuroImage* 12.4 (Oct. 2000), pp. 466–477.

[136] Richard B. Buxton, Eric C. Wong, and Lawrence R. Frank. “Dynamics of blood flow and oxygenation changes during brain activation: The balloon model”. en. In: *Magnetic Resonance in Medicine* 39.6 (1998). Pp. 855–864.

[137] Robert L. Grubb et al. “The Effects of Changes in PaCO2 Cerebral Blood Volume, Blood Flow, and Vascular Mean Transit Time”. In: *Stroke* 5.5 (Sept. 1974). Pp. 630–639.

[138] Gopikrishna Deshpande, Yun Wang, and Jennifer Robinson. “Resting state fMRI connectivity is sensitive to laminar connectional architecture in the human brain”. en. In: *Brain Informatics* 9.1 (Jan. 2022), p. 2.

[139] J. Daunizeau, K. E. Stephan, and K. J. Friston. “Stochastic dynamic causal modelling of fMRI data: Should we care about neural noise?” In: *NeuroImage* 62.1 (Aug. 2012), pp. 464–481.

[140] Baojuan Li et al. “Generalised filtering and stochastic DCM for fMRI”. In: *NeuroImage* 58.2 (Sept. 2011), pp. 442–457.

[141] Karl J. Friston et al. “Network discovery with DCM”. In: *NeuroImage* 56.3 (June 2011), pp. 1202–1221.

[142] Rodney J. Douglas and Kevan A. C. Martin. “NEURONAL CIRCUITS OF THE NEOCORTEX”. en. In: *Annual Review of Neuroscience* 27.Volume 27, 2004 (July 2004). Pp. 419–451.

[143] Irati Markuerkiaga, Markus Barth, and David G. Norris. “A cortical vascular model for examining the specificity of the laminar BOLD signal”. In: *NeuroImage* 132 (May 2016), pp. 491–498.

[144] K. J. Friston et al. “Dynamic causal modelling revisited”. In: *NeuroImage* 199 (Oct. 2019), pp. 730–744.

[145] Atena Akbari et al. “Modelling the depth-dependent VASO and BOLD responses in human primary visual cortex”. en. In: *Human Brain Mapping* 44.2 (2023). Pp. 710–726.

[146] Laurentius Huber et al. “Investigation of the neurovascular coupling in positive and negative BOLD responses in human brain at 7T”. In: *NeuroImage* 97 (Aug. 2014), pp. 349–362.

[147] Cheryl A. Olman et al. “Layer-Specific fMRI Reflects Different Neuronal Computations at Different Depths in Human V1”. en. In: *PLOS ONE* 7.3 (Mar. 2012). e32536.

[148] Hikaru Fukutomi et al. “Diffusion Tensor Model links to Neurite Orientation Dispersion and Density Imaging at high b-value in Cerebral Cortical Gray Matter”. en. In: *Scientific Reports* 9.1 (Aug. 2019). P. 12246.

[149] Alexandru V. Avram et al. “High-resolution cortical MAP-MRI reveals areal borders and laminar substructures observed with histological staining”. In: *NeuroImage* 264 (Dec. 2022), p. 119653.

[150] Anneke Alkemade et al. “A unified 3D map of microscopic architecture and MRI of the human brain”. In: *Science Advances* 8.17 (Apr. 2022). eabj7892.