Imperial College London

MEng Individual Project

Imperial College London

Department of Computing

Modelling Neurodegeneration: Investigating Changes in Connectomics of Alzheimer's Patients

Author:
Dyuti Chakraborty

Supervisor:
Dr. Pedro Mediano

Second Marker:
Dr. Islem Rekik

June 17, 2024

Abstract

Alzheimer's Disease (AD) is an irreversible neurodegenerative disease characterized by memory impairment and progressive cognitive decline and is the most common cause of dementia among the elderly. Despite advancements in research, there are no definitive clinical diagnostic methods or effective treatments for AD, highlighting the need to understand the underlying dynamics of neurodegeneration. To address this, we can utilise techniques from computational neuroscience to leverage whole-brain computational models and explore the causal mechanisms behind these changes.

Modern neuroimaging techniques have revealed structural and dynamic brain changes in AD patients that are typically correlated with disease pathogenesis. These include dynamic changes such as Amyloid-β (Aβ) and Tau protein accumulation and structural changes such as disrupted brain network properties. Additionally, medical studies have shown the decline across different clinical measures in AD patients, such as a decline in performance on cognitive tests or anatomical metrics such as ventricular enlargement in the brain. This study integrates connectomics and machine learning to correlate simulated brain connectivity with clinical measures of AD, aiming to understand how the structural and dynamic changes in the former affect the latter.

Our approach adapts the traditional whole-brain modelling pipeline to include the modelling of dynamic changes and the prediction of clinical measures from stimulated functional connectivity (FC). Using the `whobpyt` library, we adapt an RNN-based dynamic mean-field model to simulate functional connectivity from a healthy structural connectivity (SC) matrix, incorporating the effects of Aβ and Tau proteins and fitting the SC matrix in the process. This allows us to model the progression of AD across different patient groups: cognitively normal (CN), mild cognitive impairment (MCI), and AD. By varying structural and dynamic parameters and employing a regressor model, we predict ventricular volume changes for different simulations. This way, we quantify the impact of parameter changes on functional connectivity. This pipeline aims to enhance AD diagnosis and detection by providing a mechanistic understanding of how structural and dynamic alterations influence disease progression.

Overall, this project demonstrates the successful adaptation of the whole-brain modelling pipeline to assess the impact of perturbing structural and dynamic parameters on neural functionality. We have shown that integrating Aβ-Tau degradation into the DMF model can improve the fit between simulated and empirical FC data, demonstrating how combining connectomics and biomarkers can enhance our understanding of AD pathogenesis. Our results suggest that structural changes from a fitted SC matrix have a greater influence on predicted ventricular volume than dynamic parameter changes, indicating the importance of including structural biomarkers in AD modelling in the future.

Acknowledgements

Firstly, I would like to thank my supervisor, Dr. Pedro Mediano, for all the insight, guidance, and encouragement of over the past year. This project would not have been possible without his constant support. Additionally, thanks to Mar Estarellas for providing inspiration and guidance related to ADNI and cognitive clinical measures, and Keenan John Anthony Down for providing the patient fMRI data used throughout the project.

Finally, I would also like to thank my friends and family for their unwavering support over the past four years. To my parents, for their constant optimism and always believing in me. To my friends, you all are the best part of Imperial. I'm glad DoC brought us together, and it's been a privilege sharing the university experience with you. I couldn't have got through this without you guys!

Contents

Chapter 1

Introduction

1.1 Motivation

Neurological disorders account for the leading cause of physical and cognitive disability worldwide, affecting around 15% of the global population [1]. Alzheimer's Disease (AD) is the most common cause of dementia among people over the age of 65 [2]. AD is an irreversible neurodegenerative disease that leads to the impairment of memory, progressive cognitive decline and ultimately severely impacts a person's ability to perform simple daily tasks. With no clinical diagnosis methods or existing treatment, it is crucial to decipher the underlying dynamics of neurodegeneration and its potential link to the onset and progression of AD. Advancements in neuroimaging techniques (such as fMRI and DTI methods) have allowed researchers to identify various structural and dynamic changes within patients that disrupt the brain's functional organisation over the pathogenesis of AD. Examples range from the neuronal accumulation of proteins such as amyloid-beta (Aβ) and Tau [3] to the loss of small-world properties in the structure of brain networks [4]. Additionally, observed symptoms such as cognitive capabilities deteriorate with disease progression. While it can be hypothesised that structural and dynamic changes affect observed patient characteristics, the goal of this project is to find a mechanistic way of illustrating this.

Alongside modern neuroimaging methods, the study of connectomics and whole-brain computational models have provided methodologies to use graph theory and network science to characterise brain systems, understand the link between structural and functional brain connectivity and investigate how this affects the pathogenesis of disease [5]. Whole-brain modelling uses a combination of neurobiological-inspired models and multimodal empirical imaging data to reverse engineer various aspects of brain function [6]. This study primarily focuses on the Dynamic Mean Field (DMF) model - a biologically-informed model that uses coupled stochastic differential equations to simulate mesoscale, regional neural dynamics. The different inputs and biophysical parameters of the model can be systematically perturbed to assess the effects on the overall brain dynamics [7], providing an advantageous tool for investigating the causal mechanisms behind the changes in brain organisation and function in AD patients.

To measure the latter, we investigate methods of correlating connectomics with clinical AD patient measures. These can be either behavioural or biological, but allow us to map the mathematically modelled progression to the tangible impact on real patients. Examples of doing this include predicting the scores of behavioural cognitive tests from functional connectivity [8], or correlating functional connectivity with anatomical values of ventricular volume [9]. By correlating functional connectivity with behavioural or anatomical scores, we can quantify the impact of perturbing parameters on functional connectivity as the difference in the predicted scores of resulting simulations.

In this project, we propose a pipeline for this purpose that applies novel machine-learning techniques to the domain of computational neuroscience. The goal of the pipeline is to quantify how much a structural or dynamical parameter change in a neural model could affect the progression of AD. Although imaging techniques such as DTI and fMRI allow us to form ideas about structural and functional brain activity respectively, we are yet to understand the interrelationship between the two, and quantify how changes in the former can affect the latter.

In this adapted whole-brain modelling pipeline, we use the `whobpyt` library [10] to create an RNN-based, adapted DMF whole-brain neural model to simulate functional connectivity from a

healthy structural connectivity matrix at different stages along the scale of AD progression. To simulate the progression of AD in the whole-brain model, we include the effects of the proteins Amyloid-β (Aβ) and Tau [3] in the DMF model - biomarkers known to be associated with the development of AD [11]. The associated parameters are perturbed, resulting in simulations of functional connectivity across the full range of patient groups: cognitively normal (CN), those with mild cognitive dementia (MCI) and finally those diagnosed with Alzheimer's disease (AD). The respective models give us a learnt parameter set and fitted structural connectivity matrix per patient group. We can experiment with these to vary the extent of structural and dynamic changes in a whole-brain model. Then, using a regressor model, we predict the ventricular volume from the functional simulations of each experiment. The changes in this score between the simulations can quantify the effect of the structural and dynamic changes on functional connectivity. Figure 1.1 illustrates the adapted whole-brain modelling method developed throughout the project, which is explored in more detail throughout this report.

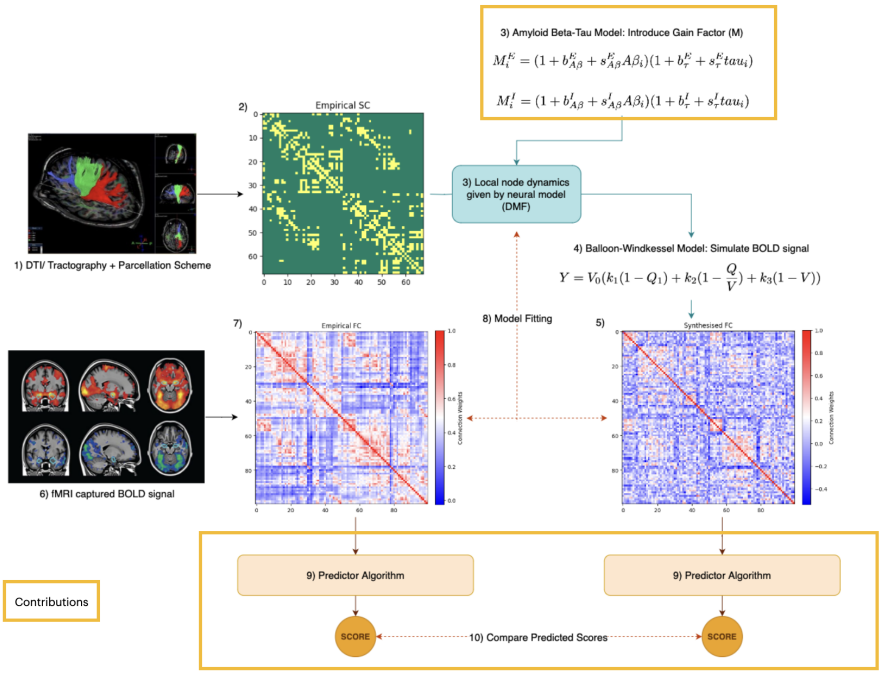

Figure 1.1: Adapted whole-brain modelling pipeline developed through the project, with the main contributions indicated by the yellow boxes

1.2 Objectives

1. **Goal 1:** Using the structural connectivity matrices of cognitively healthy (CN) patients, fit the DMF model using deep learning-based parameter estimation [12] to best reproduce the empirical functional connectivity matrices of these patients.
2. **Goal 2:** Using the parameter values found from fitting the healthy patient group and the same structural connectivity matrices, adapt the DMF model to include Aβ and Tau degradation. Fit this adapted model to the empirical functional connectivity matrices of patients in the groups CN, MCI and AD to learn the different Aβ-Tau model parameter values at each stage.
3. **Goal 3:** Train a regressor model to predict a ventricular volume score from a given functional connectivity matrix. Use this to determine the clinical measure scores associated with the simulated functional connectivity matrices of each patient group.
4. **Goal 4:** Vary different structural and dynamic properties of the Aβ-Tau DMF model and produce functional connectivity simulations for each experiment. Use the regressor model to predict the ventricular volume from each simulation, calculating how much it changes between each experiment.

1.3 Main Contributions and Results

The main contributions and results of this project are outlined below:

- Adapting an existing whole-brain modelling library `whobpyt`[10] to include dynamical changes of biomarkers Aβ and Tau, as shown in Figure 1.1
- Improving the fit between simulated and empirical functional connectivity matrices using DMF models for Schaefer Parcellations [13] for healthy patients by including Aβ-Tau heterogeneities in the DMF model
- Adapting the traditional whole-brain modelling pipeline to include clinical measure prediction, as shown in Figure 1.1: we trained a Gradient Boosting Regressor Model to predict ventricular volume from functional connectivity matrices
- Demonstrated that structural parameter changes have a larger impact on functional connectivity than dynamic changes in CN and AD patient groups, assessed by the difference in predicted ventricular volume score from simulated experiments
- Demonstrated that varying a single dynamic parameter for a CN patient does not have a significant impact on resulting brain functionality, assessed by the difference in predicted ventricular volume score from simulated experiments

1.4 Ethical Considerations

This project involves secondary analysis of pre-collected data, as we will be using the Alzheimer's Disease Neuroimaging Initiative (ADNI) dataset for information on AD patients (further discussed in Section 2.5.1). This dataset uses human data from neuroimaging and clinical studies of a sample of North American study volunteers [2]. However, to protect the personal data, the study authors anonymise the data. The studies have been conducted only after a local ethics review board review.

Chapter 2

Background

2.1 Whole-Brain Modelling

Whole-brain modelling is a tool used to describe the dynamics and interactions between neural populations in different regions of the brain. Underpinned by network science and graph theory, a whole brain model is a mathematical representation of the brain as nodes, with a set of coupled first-order differential equations representing the behaviour of a node (or brain region) over time. The behaviour described is dependent on the specific model used (ranging from the average firing rate to the oscillatory behaviour of neurons in a region) and is affected by interactions both within and between nodes. Neurons are coupled together using empirically derived anatomical connections, such as diffusion tensor image (DTI) analysis for structural connections or functional magnetic resonance imaging (fMRI) for functional connections. Tuning whole-brain models enables the simulation of brain dynamics and can be used to reverse engineer various aspects of brain function to compare against empirical human data from fMRI or EEG results [6]. A whole-brain model is made up of three main concepts: parcellation, structural and functional connectivity, and the dynamic neural model [14].

A parcellation scheme (or atlas) is used to subdivide areas of the brain into distinct subregions or nodes. There are many underlying neurobiological factors affecting the choice of the subregion, but they must be at least internally uniform with respect to a unique functional or biological property. This allows the region to be considered a homogeneous neural population, and so the neurons can be modelled together as one node. There are many atlases such as AAL, Lausanne, and Schaefer parcellations [15, 12] and are chosen based on the purpose of the study.

Once the location of nodes has been decided, connectivity measures are used to define how the nodes in a parcellation are linked together. These define the weighted edges between nodes in a whole-brain model and can either describe structural connections between two nodes or functional connections. Structural connectivity defines which regions in the brain communicate with each other using physical communication pathways, whereas functional connectivity identifies regions of the brain that are co-active at the same time. These can both be determined using empirical measurements, as discussed further in Section 2.2.

Finally, a neuronal model is required to describe the behaviour of a population of neurons within a region. This allows us to reason about how the neural structure can influence neural activity. Neurons in the brain are typically separated into two subclasses: excitatory and inhibitory neurons; determined by the neurotransmitters of the neuron, these increase and decrease the activity of connecting neurons respectively. A neural population can be seen as a dynamical system: the activity of the interacting populations is represented as a set of coupled first-order differential equations, and a non-linear activation function is used to link the input of a population to its output. Although sometimes difficult to solve analytically, geometric methods or numerical estimation techniques can be used to find the solution for the chosen activity variable (e.g. average firing rate) of the system.

Broadly, the spectrum of whole-brain models varies according to the trade-off between conceptual simplicity and neurobiological realism [6]. These range from:

- *Models prioritizing local equations.* These phenomenological models can reproduce empirical data with fewer assumptions or need for tuning, as they typically have a smaller parameter space.

  - *Advantages:* Good for finding general dynamical principles, are faster to simulate and are easier to tune.
  - *Disadvantages:* They lack biological realism, and so conclusions drawn from these models are harder to use to explain empirical behaviour.

- *Models prioritizing biological realism.* These models are more complex as they are interpreted in terms of the biophysics of neurons and therefore have a much larger parameter space.
  - *Advantages:* Realism helps to provide more realistic causal explanations for simulated activity
  - *Disadvantages:* They are harder to fit and train due to the large parameter space.

2.2 Connectivity

Connectivity measures in the brain are crucial in assembling a whole-brain model, as they describe the interactions between regions in the brain. The description of a brain as a network is known as a connectome, and neuroimaging techniques allow us to form both structural and functional connectomes. These form two crucial steps in the whole-brain modelling pipeline: structural connectivity influences functional connectivity, although the extent of this is still unknown [16].

2.2.1 Structural Connectivity

Structural connectivity (SC) describes the anatomical network of fibre tracks in the brain. Techniques such as diffusion tensor imaging (DTI) or diffusion spectrum imaging (DSI) are used to identify white matter tracts of the human brain. White-matter fibres are the long-range myelinated axons that form the brain's communications infrastructure and dictate which regions are interconnected. The pathways are reconstructed from imaging data through tractography. This defines trajectories between nodes with the number of streamlines between regions used to estimate communication strength: a thicker tract (where there are more streamlines) corresponds to an edge with a higher weight. The communication strength is then processed through a threshold or using binarization to filter out tractography noise [17], leading to a sparser SC matrix as shown in Figure 2.1a (normalised and binarised with a threshold of 0.01).

2.2.2 Functional Connectivity

In contrast, functional connectivity (FC) describes which regions of the brain are co-active, and therefore functionally related. This is extracted from time-series measures of neural activity between brain regions, based on functional magnetic resonance imaging (fMRI) and blood-oxygen-level dependent (BOLD) contrasts. Functional connectivity refers to the statistical dependence between the spatially distinct time series of electrophysical activity and de(oxygenated) blood levels [16]. It is most commonly computed using the Pearson correlation coefficient between regional activity, but other statistical measures such as partial correlation, mutual information, and coherence could also be used [18]. From these measures, the matrix undergoes processing (thresholding and binarization) to remove noise and signs (positive or negative) on the correlations, and self-correlations are typically set to 0. This results in an FC matrix such as Figure 2.1b. However, a limitation of using fMRI is that its temporal resolution is restricted by the haemodynamic response time, since it is a measure of metabolism and blood flow on a voxel-by-voxel basis. This means that the resulting time series is on the order of seconds, as opposed to imaging techniques that measure electrical activity (such as EEG and MEG) that give a finer temporal resolution (milliseconds) [16].

2.2.3 Integrating Connectomics into Whole-Brain Modelling

The two connectomes play an integral role in the whole-brain modelling pipeline shown in Figure 2.1c, alongside a neural model describing regional neural dynamics, and a haemodynamic model used to synthesise an FC matrix.

Haemodynamic Model: Balloon-Windkessel Model. A haemodynamic model is required to transform the state variables in the neural dynamic model into a form that can be compared to

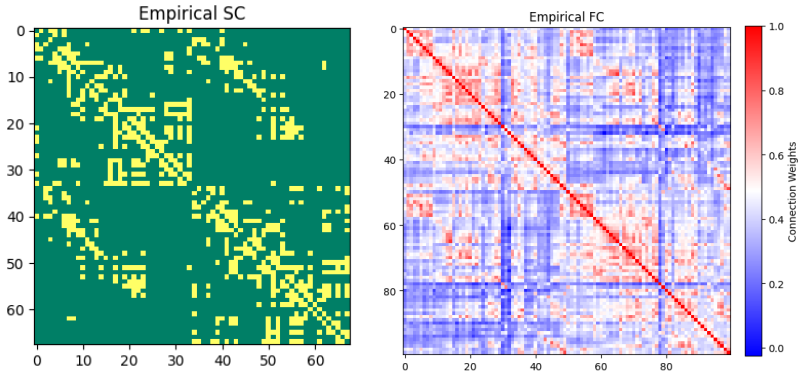

(a) SC Matrix (b) FC Matrix

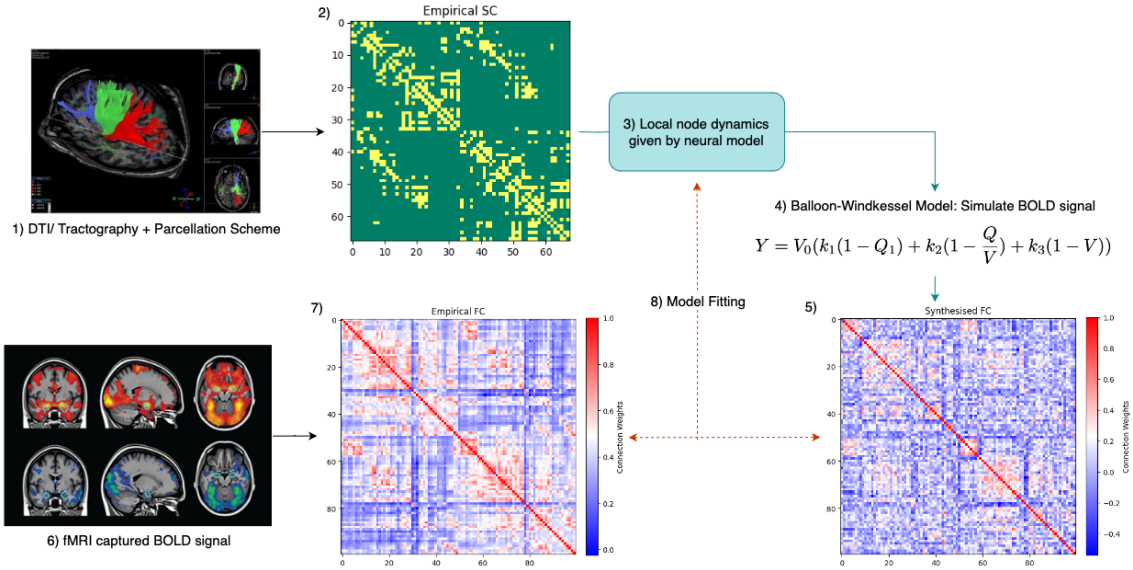

(c) Whole-Brain Modeling Pipeline

Figure 2.1: Example connectivity matrices from a dataset of 70 young healthy adults [19] and how they can be used in the whole-brain modelling pipeline

the empirical BOLD signal from the fMRI measurements [12]. An example of such a model is the Balloon-Windkessel model [20] which describes the biophysics of BOLD signal generation.

In this model, the vascular bed in a small volume of tissue is modelled as an expandable venous compartment in which blood volume changes happen. The main driver of the system is the volume flow rate (ml/s), $F_i(t)$, output of the capillary bed and into tissue. The BOLD signal is a nonlinear function of normalised venous volume (V), normalized total deoxyhemoglobin content (Q), and resting net oxygen extraction fraction by the capillary bed (O_i). The following equations capture the rates of change of each variable as the model expands and contracts:

$$\dot{Q}_i = \frac{1}{\rho\tau_0} F_i(1-(1-\rho)^{\frac{1}{F_i}}) - \frac{1}{\tau_0}(Q_i V_i^{\frac{1}{\alpha}-1}) \tag{2.1}$$

$$\dot{V}_t = \frac{1}{\tau_0}(F_i - V_i^{\frac{1}{\alpha}}) \tag{2.2}$$

$$\dot{F}_i = X_i \tag{2.3}$$

$$\dot{X}_i = O_i - \frac{1}{\tau_s} X_i - \frac{1}{\tau_F}(F_i - 1) \tag{2.4}$$

$\dot{X}_i$ is the derivative of $\dot{F}_i$ and describes the flow between the different steps of the model. The equations make use of the constants $\rho = 0.34$ (rate constant), $\tau_0 = 0.98$ (transit time), $\alpha = 0.32$

(stiffness parameter), $\tau_s = 0.65$ (signal decay) and $\tau_F = 0.41$ (autoregulation). The BOLD signal is then given by Equation 2.5, where V_0, k_1, k_2 and k_3 are constants:

$$Y = V_0(k_1(1 - Q_1) + k_2(1 - \frac{Q}{V}) + k_3(1 - V)) \tag{2.5}$$

Whole-Brain Modelling Pipeline. Putting this all together, we get the pipeline illustrated in Figure 2.1c. Firstly, using a parcellation scheme and neuroimaging techniques such as DTI and tractography (1), an empirical structural connectivity matrix can be found (2). This is then combined with a neural model (3) that describes the mesoscopic population activity at each region. The output of this is then passed through a haemodynamic model (4) to generate a simulated BOLD signal that can be synthesized into a functional connectivity matrix (5). As a result, we can compare the synthesised and empirical (7) FC matrices (the latter captured using fMRI (6)) and calculate a goodness-of-fit metric, such as the Pearson correlation coefficient. This loss function is then used to fit the regional neural model (8) to further tune its hyperparameters.

2.3 Neural Population Models

In this section, we will discuss three different examples of mathematical models for neuronal activity, that have been used for researching AD in recent literature. These range from biologically realistic models - such as the Dynamic Mean Field model - to more abstracted models, such as using the Hopf Bifurcation to model regional dynamics. It is worth noting that there are many other examples of neural population models that span three main groups: rate/ activity-based neural mass models, voltage-based models and oscillatory models [14], that have not been covered in this section.

2.3.1 Dynamic Mean Field Model

The Dynamic Mean Field (DMF) model (also referred to as the Reduced Wong-Wang (RWW) model, used interchangeably in this report) is a two-variable simplification of a biophysically realistic cortical network model [21]. This model approximates the average behaviour of a network of individual integrate-and-fire (IF) spiking neurons by using a 'mean field' approach [22]. Instead, DMF describes how the ensemble activity of different excitatory and inhibitory neural populations (that constitute the IF network) evolves.

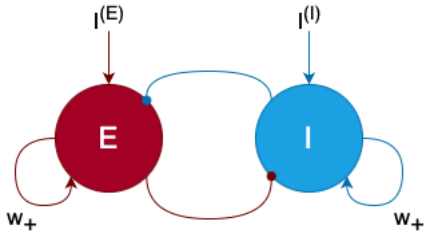

Figure 2.2: Dynamic Mean Field Model [21]

In the 'mean field' approach, the input to a neuron in a large homogeneous population is treated as a Gaussian random process, and therefore the mean activity of the population can be seen as a single unit. As shown in Figure 2.2, the model consists of two homogeneous neural populations (excitatory, E, and inhibitory, I) that are equipped with self-excitation and mutual inhibition. The firing rate for each population is dependent on the input current of that population, but also conversely, the input currents for each population depend on the respective firing rates. This means that the firing rate for each population can be determined self-consistently from a reduced system of coupled nonlinear stochastic differential equations (SDE):

$$I_i^{(E)} = W_E I_0 + w_+ J_{NMDA} S_i^{(E)} + G J_{NMDA} \sum_j Cij S_j^{(E)} - J_i S_i^{(I)} + I_{\text{externals}} \tag{2.6}$$

$$I_i^{(I)} = W_I I_0 + w_+ J_{NMDA} S_i^{(E)} - S_i^{(I)} + \lambda G J_{NMDA} \sum_j Cij S_j^{(E)} \tag{2.7}$$

Variable Name	Description	Value
I_{external}	External simulation for a task-evoked activity	0 for modelling resting-state, 0.02 otherwise
w_+	Local excitatory recurrence	1.4
λ	Indicates whether long-range feedforward inhibition (FFI) is considered	1 if FFI is considered, 0 otherwise
γ	Kinetic parameter	0.641 ms
J_{NMDA}	Excitatory synaptic coupling	0.15nA
J_i	Local feedback inhibitory synaptic coupling	1 for no-FIC case, adjusted independently otherwise
I_0	Overall effective external input	0.382nA
ν_i	Uncorrelated standard Gaussian noise	-
σ	Gaussian noise amplitude for each node	0.01nA

Table 2.1: DMF Model Parameters [22]

$$r_i^{(E)} = H^{(E)}(I_i^{(E)}) = \frac{a_E I_i^{(E)} - b_E}{1 - \exp(-d_E(a_E I_i^{(E)} - b_E))} \tag{2.8}$$

$$r_i^{(I)} = H^{(I)}(I_i^{(I)}) = \frac{a_I I_i^{(I)} - b_I}{1 - \exp(-d_I(a_I I_i^{(I)} - b_I)} \tag{2.9}$$

$$\frac{dS_i^{(E)}(t)}{dt} = -\frac{S_i^{(E)}}{\tau_E} + (1 - S_i^{(E)})\gamma r_i^{(E)} + \sigma\nu_i(t) \tag{2.10}$$

$$\frac{dS_i^{(I)}(t)}{dt} = -\frac{S_i^{(I)}}{\tau_I} + r_i^{(I)} + \sigma\nu_i(t) \tag{2.11}$$

In the equations above, for the brain region i, $r_i^{(E,I)}$ denotes the population firing rate, $S_j^{(E,I)}$ is the average excitatory or inhibitory synaptic gating variable, $I_i^{(E,I)}$ are the input currents and C_{ij} is the SC matrix used to define the anatomical connections between brain regions. The input–output functions of excitatory pools and inhibitory pools are given by $H^{(E)}$ and $H^{(I)}$ respectively. The other constants and values are given in Tables 2.1-2.2.

We consider only the dynamics of the NMDA synaptic receptor that couples neurons since this has the longest decay time constant (100ms) than the other receptors (AMPA and GABA) in the cortical network model. This is because the mean firing rate of a neuron receiving a noisy input is approximated by a sigmoidal function in the first passage time equation [21] - since NMDA has a longer decay time constant, the dynamics of this gating variable will dominate the evolution of the system over the other receptors.

Although this is a biophysically realistic model, certain limitations arise due to its complex nature. Above all, due to the large parameter space associated with the SDEs, the DMF model is computationally expensive to fit and use. Furthermore, to get more realistic dynamics, the Feedback Inhibition Control (FIC) parameters need to be fit to stabilise the excitatory populations, which also inflates the computing power required. This overhead also restricts the spatial resolution of what can be simulated, restricting to a coarse-grained spatial scale for the parcellations used to generate an SC matrix. However, there is a lot of current research regarding parameter optimisation for DMF, for example by leveraging Bayesian Optimisation in the *FastDMF* implementation [6] or using machine learning techniques such as deep neural networks [12] (discussed further in Section 3.1.1).

Excitatory Gating Variables		**Inhibitory Gating Variables**	
Variable	*Value*	*Variable*	*Value*
a_E	$310\text{n}C^{-1}$	a_I	$615\text{n}C^{-1}$
b_E	125Hz	b_I	177Hz
d_E	0.16s	d_I	0.087s
$\tau_E = \tau_{NMDA}$	100ms	$\tau_I = \tau_{GABA}$	10ms
W_E	1	W_I	0.7

Table 2.2: DMF Model Constants [22]

2.3.2 Oscillatory Models: Kuramoto Oscillators

Oscillatory models are a different paradigm of models to the former DMF and can prove to be beneficial when considering synchronisation phenomena in large populations of neural systems. One such example is the Kuramoto Model, created to study collective synchronisation across oscillators [14]. In this model, each population of neurons is characterised by its phase θ and we again use a mean-field coupling of the Kuramoto Order Parameter (KOP), K, given by Equation 2.12 [23].

$$K(t) = \frac{1}{N}\sum_{j=1}^{N}\exp(i\varphi_j(t)) \tag{2.12}$$

This parameter can be used to describe synchronisation phenomena in the brain: the property of coherence is defined as the temporal average of KOP (or mean synchronisation), whereas metastability is defined as the standard deviation of KOP (variation in synchronisation over time).

Following the same assumptions as in DMF, we assume that the neural populations are homogeneous. Additionally, we assume communities of idealised oscillators: the neurons in each population are spiking regularly enough to be considered oscillating, and that there is weak coupling between neurons so that the magnitude of perturbations from the interactions between neurons is lower than the intrinsic frequency of the oscillators. Under these 'mean field' assumptions, the Kuramoto model is a system of differential equations [14, 24] where the phase oscillators are couples through the sine of their phase differences:

$$\dot{\theta}_i = \omega_i + \frac{K}{N}\sum_{j=1}^{N}\sin(\theta_j - \theta_i) \tag{2.13}$$

This model can display a large variety of synchronisation patterns, allowing us to model interacting elements in biological systems with a fewer number of parameters. Although it is flexible enough to be adapted to many contexts (including outside the domain of whole-brain modelling) it is limited to describing synchronous behaviour and therefore tackles a different problem space to that of DMF models.

2.3.3 Phenomenological Models: Hopf Bifurcation

On the other end of the spectrum of models, a whole-brain model can also abstract away from the biophysical properties of the brain. For example, here we discuss a model consisting of a network of nodes (based on a parcellation, such as AAL) that are considered nonlinear oscillators. Each node is modelled using the normal form of a Hopf bifurcation, given by Equation 2.14 [25].

$$\frac{dz_j}{dt} = [a + i\omega_j]z_j - z_j|z_j|^2 \tag{2.14}$$

In the equation above, z_j is the complex variable $z_j = (x_j + iy_j)$ that denotes how the activity of node j changes over time. Additionally, ω_j is the natural oscillation frequency of this node (0.04-0.07Hz), which is determined by finding the average peak frequency of bandpass-filtered fMRI signals for each brain region. Most importantly, a represents the bifurcation parameter that controls the dynamic behaviour of the system. When $a < 0$, the phase space presents a unique stable fixed point at $z_j = 0$ towards which the system asymptotically decays. In contrast, for $a > 0$, the stable fixed point gives rise to a limit cycle and self-sustained oscillations of frequency $f_j = \frac{\omega_j}{2\pi}$ and amplitude proportional to $\sqrt{a}$.

These individual nodes are coupled by an SC matrix, C_{ij}, giving rise to the following set of differential equations:

$$\frac{dx_j}{dt} = \frac{d\mathrm{Re}(z_j)}{dt} = [a - x_j^2 - y^2]x_j - \omega_j y_j + G\sum_{i=1}^{N} C_{ij}(x_i - x_j) + \beta\eta_j(t) \tag{2.15}$$

$$\frac{dy_j}{dt} = \frac{d\mathrm{Im}(z_j)}{dt} = [a - x_j^2 - y^2]y_j - \omega_j x_j + G\sum_{i=1}^{N} C_{ij}(y_i - y_j) + \beta\eta_j(t) \tag{2.16}$$

Here, N represents the total number of nodes, G is the global coupling factor used to scale the anatomical coupling equally between all pairs of nodes, and β_j is the amplitude of the additive Gaussian noise in each node. When such noise is added to the model, dynamics close to the bifurcation point ($a \approx 0$) will switch stochastically between the aforementioned regimes of behaviour, resulting in oscillations with complex amplitude modulations.

While this model illustrates oscillatory behaviour well, its main disadvantage is that - being a phenomenological model - it cannot be directly interpreted in terms of biophysical variables. This makes investigating the causal mechanism behind AD or other neurodegenerative conditions difficult with such an abstracted model. The Hopf Bifurcation can be used to explore mechanisms associated with general aspects of dynamics. However, to make it more meaningful it could be combined with other models, for example using the Kuramoto Order Parameter to give neurophysical meaning to the oscillatory behaviour [23].

2.4 The Effect of AD on Connectomics

While it is not clear what causes the neurodegeneration that leads to Alzheimer's Disease, recent work has been done to identify changes in a patient's connectome that try to explain the underlying mechanisms behind the disease's pathogenesis. Some of these changes are associated with graph theoretic measures or small-world network metrics, whereas others look at the biological factors that influence brain functionality, such as neurotransmitters. This way, network parameters and biomarkers can be used as signs to distinguish healthy ageing from AD, alongside conventional neuroimaging techniques such as fMRI. Here, we outline some structural (anatomical), dynamic (biological factors that affect how neurons fire) and functional (a combination of the aforementioned) changes from recent research.

Changes in AD Patients	**Type**	**Reference**
Presence of anomalous functional connectivity and abnormal white matter in structural connectivity in AD patients compared to healthy controls; white matter pathology could be used as a biomarker for disease progression	*Structural*	[11]
Structural brain networks of AD patients showed reduced local efficiency compared to control patients, with regard to whole-brain white matter connectivity	*Structural, Functional*	[26]
The two most relevant proteins associated with AD development are amyloid-beta (Aβ) and phosphorylated Tau. There is a high correlation between the deposition patterns of these proteins in neuritic plaques in the brain and the onset of AD - patients with AD might show early deposition of these proteins years before the onset of symptoms. Aggregation of Aβ can lead to toxic effects, whereas tau plays a role in neural death and axonal dysfunction, leading to disconnection of the affected areas in the brain. Recent literature shows a clear dominance of Aβ in the early stages of the disease (mild cognitive impairment - MCI - phase), while tau dominates in later stages (AD).	*Dynamic*	[3, 11]

Changes in AD Patients	Type	Reference
Highly connected hub nodes in the brain are most affected by AD, as they have a higher level of amyloid-β deposition from an early stage. Excessive neuronal activity can increase this level of deposition, and as hub regions have the highest levels of activity, this correlation leads to the vulnerability of AD.	*Dynamic*	[11, 27]
Altered modular structure in brain organization patterns: the largest homotopic module in the control group (the insula module) was divided into pieces in the AD patient group. The insula module had lost symmetric functional connection properties and had a lower corresponding gray matter concentration (GMC) in the AD group. Combining structural GMC and functional modular analysis for the insula module, a new biomarker could be developed for a single subject.	*Structural, Functional*	[28]
Disruption due to AD of the functional organization of the brain as a small-world network. The functional networks in AD patients showed the loss of small-world properties such as lower clustering coefficient and characteristic path length and therefore disrupted regional connectivity. These network measures could be useful as imaging-based biomarkers to distinguish AD from healthy ageing.	*Functional*	[4]
Empirical BOLD signal declines with disease progression, correlating with a decline in global interactions (coherence and metastability) and regional functional connectivity strengths. The changes in regional dynamics could lead to a disintegration of activity within the whole-brain model. Furthermore, the functional connectivity differences were correlated to cerebrospinal fluid biomarkers such as Aβ, Tau, and phospho-Tau.	*Functional*	[23]
Decline of neurotransmitter systems that aid synaptic transmissions in the brain, in particular, the cholinergic and glutamatergic systems. Acetylcholine (part of the cholinergic system) is an essential neurotransmitter in the brain, responsible for modulating synaptic signalling, memory consolidation, and memory encoding. Cholinergic modulation is linked to working memory on a functional level - as acetylcholine is crucial for memory consolidation, dysfunction of this system is seen to be related to the development of AD.	*Dynamic*	[11]

Table 2.3: Structural and Dynamic Changes in AD Patients

2.5 Understanding AD Progression using Patient Characteristics

In order to assess the impact of AD on connectomics or build a representative whole-brain model, it is important to be able to measure or track the neurodegeneration that takes place due to AD. Disease progression can be seen in patients either due to behavioural symptoms or changes (measured and monitored through cognitive tests) or biophysical changes in the brain (such as ventricular volume enlargement). There are many global studies that aim to collect such data from patients, ranging from cognitively normal to those diagnosed with AD. The longitudinal data tracked by these studies and discussed further in this section can help us improve the accuracy of our whole brain models, as explained in Chapter 3.

2.5.1 Using ADNI Dataset

For this project, we will be making use of the dataset provided by the Alzheimer's Disease Neuroimaging Initiative (ADNI) [2], based in North America. ADNI is a longitudinal multicentre study aiming to develop clinical, imaging, genetic and biochemical biomarkers for early-stage detection

and tracking of AD. Participants are recruited from across North America and are periodically re-assessed over time to track the progression of the disease. The timeline of AD typically spans normal ageing, early mild cognitive impairment (EMCI), and late mild cognitive impairment (LCMI) to dementia or AD.

The main focus of ADNI is to validate which biomarkers can be used in clinical AD trials. A biomarker is a substance, measurement, or indicator of a biological state that can exist before any clinical symptoms of a disease arise. The five biomarkers measured in ADNI are:

1. Amyloid-beta (Aβ) - measured in cerebrospinal fluid or by amyloid PET imaging
2. Tau proteins - measured in cerebrospinal fluid or by synaptic dysfunction (using FDG-PET)
3. Brain atrophy - measured by structural MRI
4. Memory loss - measured using cognitive tests
5. Clinical function - general cognitive decline measured using cognitive tests

1-3 can be observed before diagnosis, whereas 3-5 are the known symptoms for AD diagnosis. Alongside neuroimaging scans, ADNI data includes cognitive/neuropsychological tests, genetic data sets and biofluids.

However, there are certain limitations to this dataset, most prominently the lack of minority ethnic groups in the sample of participants. To overcome this, ADNI will be expanding internationally as a part of the Worldwide ADNI (WW-ADNI) study [29] with additional sites in India, China, Japan, Korea, Argentina, Australia, and Europe.

2.5.2 AD Indicator: Composite Cognitive Score

Cognitive measures can be more useful than imaging and biochemical biomarkers as they are related to the core systems of disease progression. Since they are sensitive to treatment effects, cognitive measures provide a better mechanism for tracking disease progression. It is often useful to use a combination of such measures to form a composite cognitive function score, as each measure describes proficiency in a different aspect of brain functionality. By using a composite score, researchers can reduce variability and improve the reliability of cognitive assessments.

The ADNI database contains around 20 cognitive scores and composites. Some examples of the latter include:

- ***Pre-Clinical Alzheimer's Cognitive Composite (PACC)*** - combines tests that assess episodic memory, timed executive function and global cognition and is the primary measure used for trials conducted at the asymptomatic or preclinical phase of AD [30]. PACC is a composite of four measures: total recall from the Free and Cued Selective Reminding Test (FCSRT), delayed recall on a subtest of the Wechsler Memory Scale, digit symbol substitution test from the Wechsler Adult Intelligence Scale and the MMSE total score. These measures are known for showing sensitivity to decline in mild dementia and can detect early decline in the preclinical stages of the disease. The composite is determined by normalisation across the measures: the four scores are measured at two points in time, their change is divided by the baseline sample standard deviation of that component to form z scores, and these are summed to form the composite.

- ***Alzheimer's Disease Assessment Scale-Cognitive Subscale (ADAS-Cog)*** - a rating scale developed in the 1980s to assess the level of severity of cognitive dysfunction in AD patients [31]. There are various subscales included in ADAS, both non-cognitive (ADAS-Noncog, not widely used) and cognitive (ADAS-Cog 11 and ADAS-Cog 13). The full ADAS assessment consists of twenty-one tasks and takes around 45 minutes per patient. The number of errors made on each task is summed to give an overall score between 0 and 150, with a higher score indicating worse performance. The most commonly used cognitive subscale is ADAS-Cog 11, consisting of 11 tasks that assess memory, language and praxis (scored from 0 to 70). ADAS-Cog 13 (scored from 0 to 85) adds the tests of delayed word recall and a number cancellation or maze task to this.

- ***ADSP Phenotype Harmonization Consortium (PHC)*** - a co-calibrated composite score generated using psychometric measures, to enhance cross-study comparisons. Composite scores were provided across four domains: the co-calibration of memory, executive function, language and visuospatial abilities [32], all standardised on the same metric to be comparable. Cognitive data for these were collected from different cohorts across four different US-based studies: ADNI, Adult Changes in Thought (ACT), the Religious Orders Study/ Memory Ageing Project (ROS/MAP) and the National Alzheimer's Coordinating Centre (NACC) to increase the heterogeneity of the data. To co-calibrate the composite scores, ADSP required different cognitive tests sharing common items ('anchor' items) that had identical content across the tests carried out in different studies and tested for the same underlying ability. This was identified by performing preliminary confirmatory factor analysis models within each study. The common anchor items were then used to weigh the scores for each domain to a common metric, to allow for comparison.

2.5.3 AD Indicator: Ventricular Volume

A more objective and sensitive measurement of neuropathological change associated with the onset of AD is ventricular enlargement [9]. Using ventricular volume measurements instead of cognitive tests can be beneficial as they are based on biological data (MRI scans), independent of the specific details of any study or test centre and so can assess disease progression for multi-centre studies. It can also be used to examine the neuropathological changes associated with mild cognitive impairment (MCI).

The use of ventricular volume as a measure of AD progression is examined in detail in several studies. For example, it can be seen that hemispheric atrophy rates (indicated by measuring ventricular enlargement) correlate strongly with changes on the aforementioned cognitive tests [33]. Similarly, the rate of ventricular volume change is highly correlated with an increase in the number of senile plaques and neurofibrillary tangles - two pathological features that play a significant role in the diagnosis and understanding of AD.

When considering ventricular volume, it is important to decide which regions in the brain would be best suited to volumetric analysis. MRI techniques that measure neurodegeneration typically focus on the volumetric analysis of the hippocampus. These techniques can be costly and hard to reproduce, as they involve manual or semi-manual tracing techniques. In contrast, cerebral ventricular volume is more amenable to automatic segmentation, as in MRI images there can be a sharp contrast between the signal intensity of cerebral spinal fluid (CSF) in ventricles and the surrounding tissue. Other weaknesses of MRI scans, such as the geometric distortions across the ventricle due to gradient non-linearities, are minimised as the positioning of the ventricles near the centre of the brain means that during scans, the structure is near the magnetic isocentre where the magnetic field is more uniform and stable.

Alongside measuring disease progression, anatomical measurements of ventricular volume can provide meaningful complementary evidence to cognitive testing for disease-modifying therapies. While cognitive tests can illustrate the behavioural effects of AD well, it is possible to use more structural changes to see the biological impact of neurodegenerative disease and assess the efficacy of therapies and treatments on subjects more objectively.

Chapter 3

Related Work

This chapter outlines the recent work in this field that has supported the investigations throughout the project. Section 3.1 explores some recent literature concerning the application of machine learning techniques to whole-brain modelling. Then, in Section 3.2, we present ways in which whole-brain models can be adapted to include the effect of proteins Amyloid-β (Aβ) and Tau over time, as in the onset of AD. Finally, Section 3.3 describes methods of using functional connectivity to predict changes in AD characteristics, such as ventricular volume or cognitive impairment.

3.1 Machine Learning Techniques for Whole-Brain Modelling

There is a wide variety of research regarding the application of machine learning techniques to whole-brain modelling. In this section, we explore some of these methods, namely using deep-learning optimisation to fit the large parameter space of a whole-brain model and using disease-specific atrophy maps as priors to improve the fit of whole-brain models.

3.1.1 Deep-Learning Based Parameter Estimation: `whobpyt` Library

One of the most significant challenges of whole-brain modelling is fitting the many parameters of a neural model to empirical FC data. Griffiths *et al.* proposed using a deep-learning-based approach to parameter estimation [12] by fitting model-generated activity to empirical functional neuroimaging data.

The authors proposed a 'connectome-based neural mass model' (CNMM) approach, implemented in the Python library `whobpyt`. Analogous to the whole-brain modelling pipeline explored in Figure 2.1c, this `whobpyt` pipeline simulates functional neural activity by combining an empirical SC matrix that provides the neural network structure, a two-state RWW model (Equations 2.6-2.11) that describes regional brain dynamics and the Balloon-Windkessel model (Equation 2.5) for the forward haemodynamic model function that maps neural activity as a BOLD time series. The Pearson correlation coefficient is applied between the different regions of this BOLD time series to result in a simulated FC matrix.

In the paper, the authors considered a subset of the within-node parameters from the RWW model (Equations 2.6-2.11) to fit, namely $g^{EI} = J_{NMDA}$ (excitatory-to-inhibitory synaptic gain), $g^{IE} = J_i$ (inhibitory-to-excitatory), $\frac{g^{EE}}{g^{EI}} = w_+$ (excitatory-to-excitatory synaptic gain) and $\frac{g}{g^{EI}} = G$ (long-range global coupling). Modelling whole-brain activity requires modelling large networks of several hundreds of nodes, each with systems of differential equations (often stochastic or non-linear) with sub-millisecond integration time steps. These factors mean that CNMMs can be computationally expensive to run, and as a result, large parameter spaces are difficult to explore. However, taking an empirical SC matrix, C_{ij}, the CNMM model can be seen as a recurrent neural network (RNN) that uses the parameters C_{ij}, g^{EI}, g^{IE}, $\frac{g^{EE}}{g^{EI}}$ and G as its weights. The C-parameters are the connection strengths between network nodes, and the g-parameters concurrently scale connection weights between the nodes in the network. This way, machine learning-based optimisation approaches can be applied to fit the RNN, for example using ADAM and a Variational Bayes objective function.

In the CNMM model, the empirical functional BOLD time series is divided into 30s non-overlapping batches (windows). For each batch, the model is used to simulate a corresponding

BOLD window with the currently estimated model parameter values. A cost is calculated by comparing this to the empirical window, which is then used to update model parameters for the next batch. The ADAM optimiser is used to control this process; ADAM is an adaptive learning rate optimisation algorithm that combines the benefits of RMSprop and Stochastic Gradient Descent with momentum. A key advantage of Adam is that it adjusts the learning rate individually per parameter, beneficial as the different parameter values can be of different scales of magnitude.

With these techniques, the authors used the CNMM model to recover known parameters from synthetic data and estimate parameters from empirical rsfMRI data (from the Human Connectome Project). In both cases, there was a notable performance improvement. In the former case, they achieved a model fit with a Pearson correlation of $R^2 = 0.65$ between synthetic and empirical FC (comparable studies have only achieved R^2 in the range 0.3-0.5). In the latter case, they again achieved an average FC correlation of $R^2 = 0.65$ whereas other studies usually achieved a fit in the range $R^2 = 0.3 - 0.5$. In the latter case, the model additionally fits the elements in the structural connectivity matrix C_{ij}, which previous CNMM studies have chosen to keep fixed. Since this greatly increases the number of parameters to fit (for example, using a 100-region parcellation requires 10000 elements to be fit), this has so far only been done for small-scale neural mass modelling approaches and remains an underdeveloped area of study. Achieving the fit of $R^2 = 0.65$ while fitting numerous connectivity weights shows the power of using deep learning to optimise whole-brain modelling.

As well as the marked performance improvement, the main advantage of the CNMM model is its versatility; the regional neural model discussed here uses the RWW model, but other models such as the Kuramoto model (Equations 2.13) could easily be substituted in its place. Therefore, this deep-learning approach can be used to facilitate a variety of whole-brain studies, independent of the specifics of the investigation.

3.1.2 Using Priors to Improve Model Fit

In experiments conducted by Sanz Perl *et al.* [25], disease-specific atrophy maps were used to improve the fit of whole-brain models to the empirical fMRI data. These maps were used as priors to modulate local model parameters. Anatomical priors considered in this paper include:

- *RSN*: *Resting-State Networks* - functional communities in the brain that have high intra-community functional coupling and low or intermittent inter-community coupling.
- *AD and bvFTD*: Alzheimer's Disease and behavioural variant frontotemporal dementia atrophy maps - combined due to the high spatial correlation between the two anatomical priors
- *PD*: Parkinson's disease atrophy map
- *Random*: random assignment into groups of nodes
- *Equal*: equally sized groups of nodes defined by anatomical proximity

These priors were trialled in 300 independent instances of parameter optimisation using the genetic algorithm from the Matlab Global Optimisation Toolbox. This is a stochastic optimisation process inspired by biological evolution: the most adapted individuals are selected to prevail to the next generation (iteration) of the algorithm. In the case of parameter optimisation, the 'most adapted individuals' are those that minimise a target fitness function (TFF), such as TFF = 1 - GoF (Goodness of Fit) [34]. This algorithm starts by randomly selecting a subset of the whole-brain model parameters (individuals), generating a population of outputs and TFF values, and assigning scores to each individual that is a function of the TFF value. From here, a group of individuals with the best scores are chosen as parents and then the crossover, mutation, and elite selection operations are performed to produce the next generation. The resulting offspring parameters are then used for the next generation of the algorithm, and this continues iteratively until a halting condition has been reached. A halting condition could be either that the limit of iterations has been reached, the population has stayed constant for the past N generations (e.g. $N = 50$) or the average TFF across the last N generations is less than a threshold (such as $1 \times e^{-6}$).

Parameter optimisation using a genetic algorithm (GA) is an alternative to the deep learning approach discussed in the previous section. Since we are dealing with a large and complex parameter space for whole-brain modelling, GAs give us the advantage of being well-suited for global

search and being less prone to getting stuck in a local optimum during optimisation like deep neural networks. Furthermore, genetic algorithms are more computationally efficient and avoid the overhead of large neural network architectures. However, deep learning approaches allow us to exploit data patterns and learn hierarchical representations of input data, useful for complex relationships between parameters.

When considering the original hypothesis of using disease-specific maps to improve model fit [25], the authors found that for AD, using either a combination of AD and bvFTD atrophy maps or RSN obtained the best model fits. This shows that atrophy maps can contain useful spatial heterogeneities to improve the fit of empirical data when used as a prior to the whole-brain model.

3.2 Methods for modelling Amyloid-β and Tau effects on Connectomics

An important subset of the work on whole-brain modelling of AD is that of investigating the biomarkers and dynamic parameter changes that could play a role in the pathogenesis of the disease. This section explores methods of including the role of these factors in the whole-brain modelling pipeline: either by mathematically describing their effect on simulated functional connectivity; or by using VAEs to illustrate the consequences of dynamic parameter changes on FC matrices.

3.2.1 Whole-Brain Modelling of the Influences of Amyloid-Beta and Tau in AD

Amyloid-β (Aβ) and Tau have previously been identified as important biomarkers that impact the onset and progression of AD. The work of Patow *et al.* aims to explore this further by using whole-brain modelling to describe the effects of Aβ and Tau on the excitation-inhibition balance of nodes (regions of interest) in the brain network [3]. In turn, this studies the impact of the proteins on the dynamics of regional node behaviours in the pathogenesis of AD: the authors found that neuronal activity of Aβ dominated over Tau in the early stages of the disease (during mild cognitive impairment, MCI), whereas Tau dominated over Aβ in the later stages (during the AD period). The goal of their work was to try to define whole-brain models of brain function during AD using regional distributions of neuropathology that allow for assessing the relative weights of Aβ and Tau in contributing to abnormal brain activity.

This was done by adapting the DMF whole-brain model presented in Section 2.3.1 and introducing the regional distributions of Aβ and Tau into a regional gain factor $M^{(E,I)}$ for each type of neural population. $M^{(E,I)}$ is a product of linear terms consisting of a bias constant, $b^{(E,I)}_{(A\beta,\tau)}$, and scaling factor, $s^{(E,I)}_{A\beta,\tau}$, for each regional value of Aβ and Tau, as shown in Equations 3.1 and 3.2:

$$M_i^E = (1 + b^E_{A\beta} + s^E_{A\beta} A\beta_i)(1 + b^E_\tau + s^E_\tau tau_i) \tag{3.1}$$

$$M_i^I = (1 + b^I_{A\beta} + s^I_{A\beta} A\beta_i)(1 + b^I_\tau + s^I_\tau tau_i) \tag{3.2}$$

This heterogeneity is introduced into the DMF model by modulating the neuronal response functions $H^{(E,I)}$ from the Equations 2.8-2.9 to those shown below (3.3-3.4):

$$H^{(E)}(I_i^{(E)}) = \frac{M_i^E(a_E I_i^{(E)} - b_E)}{1 - \exp(-d_E M_i^E(a_E I_i^{(E)} - b_E))} \tag{3.3}$$

$$H^{(I)}(I_i^{(I)}) = \frac{M_i^I(a_I I_i^{(I)} - b_I)}{1 - \exp(-d_I M_i^I(a_I I_i^{(I)} - b_I))} \tag{3.4}$$

Given the previously observed physical behaviour of the proteins [11], Patow could introduce bounds on the direction of the effect of Aβ and tau. This affected how many free parameters needed to be fit in the model: for example, as tau appears to only target excitatory neurons, the authors could assume that $b^I_\tau = s^I_\tau = 0$. Similar assumptions could be made for the other scaling parameters [3]: $s^I_{A\beta} < 0$ as Aβ produces inhibitory GABAergic interneuron dysfunction; $s^E_{A\beta} > 0$ as Aβ produces impaired glutamate reuptake and $s^E_\tau < 0$ as Tau reduces the excitatory synaptic neurotransmitter releases by binding to the protein synaptogyrin-3.

Following a similar pipeline to Griffiths *et al.* (Section 3.1.1), the authors used this adapted DMF model in the whole brain modelling pipeline, alongside structural connectomes of healthy control patients (collected via diffusion MRI) and empirical task-free rs-fMRI functional connectivity data to fit to. They first fit a homogenous model where $M^{(E,I)}$ was set to 1 to find the optimal value of the global coupling parameter G (from Equations 2.6-2.7) to move the model to an optimal working point such that the simulated activity best fits the empirical healthy control activity. The Aβ and Tau heterogeneities were then introduced into $M^{(E,I)}$ as described above - adjusting, this would show how local variation in the protein distributions would affect the overall excitability of the population. With the assumption of $b_\tau^I = s_\tau^I = 0$, this left a further 6 parameters (3 bias terms and 3 scaling factors) to be fit in the modified heterogenous model. The fitting process involved Bayesian optimisation and Gaussian processes: using the `g_optimize` function from Python's `scikit-optimize` library that allows the algorithm to choose the next set of parameters to evaluate by selecting the acquisition function over the Gaussian prior. This helps to prevent the optimiser from falling into local minima.

Given the similarity in the research goals of Patow *et al.*, this work has been influential on the methodology followed throughout this project, as discussed in more detail in Chapters 4-6.

3.2.2 Using Variational Autoencoders

Further work by Sanz Perl *et al.* [25] focused on using variational autoencoders (VAEs) to find a low-dimensional representation of brain activity (FC matrices). Since the severity of AD is found on a spectrum from mild cognitive impairment (MCI) to dementia, it is useful to find the trajectory of degree progression by decoding an FC matrix in a two-dimensional latent space.

In this paper, the authors divided the AD patient group into two subgroups: AD- (lower severity) and AD+ (higher severity) based on white-matter hyperintensity scores (a biomarker of cerebrovascular damage associated with AD severity). After encoding four different patient groups - control, AD-, AD+ and FTD (a patient group for a different neurodegenerative disease: frontotemporal dementia) - the authors found that the organisation of the latent space hinted at different underlying pathophysiological mechanisms and was sensitive to the severity of each disease. The two diseases (AD and FTD) were encoded in different directions, proving that latent space encoding is a powerful tool for illustrating the qualitatively different effects on whole-brain connectivity and dynamics of both diseases.

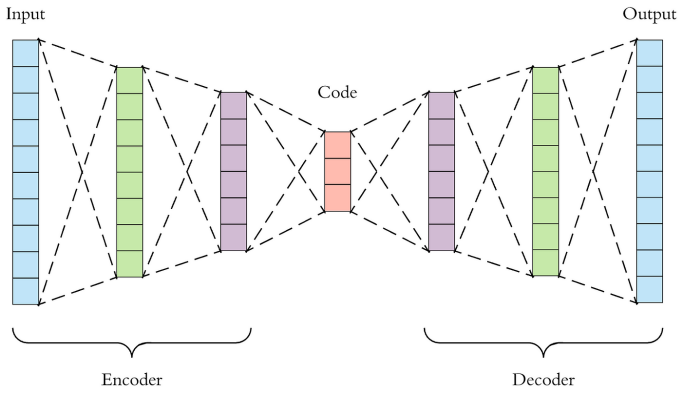

Figure 3.1: Variational Autoencoder Architecture [35]

VAEs are a type of deep neural network that follows an autoencoder (AE) architecture, as illustrated in Figure 3.1, and are trained to map inputs to probability distributions in the latent space by minimising the error between the input and out. They consist of three parts: the encoder, latent space and a decoder. The encoder takes an input (such as the FC matrix) and transforms it to a point sampled from a Gaussian probability distribution in latent space. Sampling introduces a level of randomness that allows the model to generate diverse outputs - ideally, different inputs should map to different outputs in latent space [35]. The decoder network mirrors this architecture to reconstruct the matrices from samples of the distributions. The encoder and decoder are deep neural networks with two layers, using rectified linear units (ReLu) as the activation functions. The networks are trained by minimising the combinations of two losses: a reconstruction error term (to ensure similarity between the encoder inputs and decoder outputs) and a regularisation term for

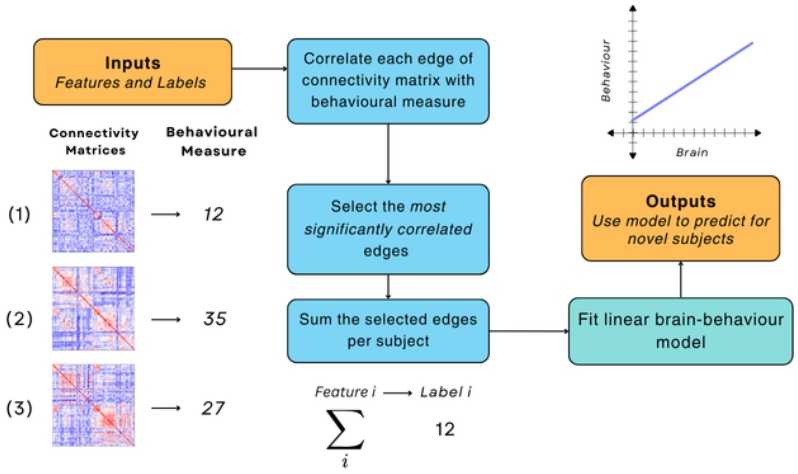

Figure 3.2: Connectome-based Predictive Model (CPM) Outline

the points in latent space (the Kullback-Leibler divergence between the latent space distribution and a standard Gaussian distribution) [25].

The main focus for this application of VAEs is latent space encoding. The latent space in a VAE is a high-dimensional space containing the input points' compressed and abstract representations. Relationships between input data can be captured more meaningfully and facilitate the visualisation of complex perturbations applied to a whole-brain model. This methodology can be used to investigate the dynamical consequences of other parameter changes due to AD on the whole-brain model.

3.3 Methods for quantifying the effect of Functional Connectivity on Clinical Measures

To add to the work in the previous section, the goal of this project is to quantify the degree to which certain biomarkers affect functional connectivity, and in turn, how that affects the AD progression of a patient. To this end, it was useful to investigate the work surrounding using functional connectivity to predict clinical patient measures, such as cognitive score and ventricular volume.

3.3.1 Using CPMs to Predict Cognitive Impairment

The first example of such a technique is connectome-based predictive modelling, developed by Shen *et al.* [36]. This is a method of summarising the most relevant features from brain connectivity data and using it for predictive modelling. As illustrated in Figure 3.2, CPM uses connectivity matrices (or correlation matrices) as input features, each with a target label of a behavioural measure. Firstly, the correlation between each edge in the connectivity matrix and the labelled behavioural measure is found. This is then statistically thresholded to find the most significantly correlated edges, which are then summed together to yield a single summary statistic (network strength) per connectivity matrix. A linear regressor model can then be fit between the summary statistics and labelled behavioural measures, and is consequently used to predict the measures for unseen subjects.

Shen *et al.* addressed the challenge of generalising these results by using cross-validation within the CPM pipeline. In current literature, similar brain-behaviour studies do not perform cross-validation and as a result the models tend to overfit to the data and fail when generalising to novel data. Instead, the authors here employed Leave-One-Out-Cross-Validation (LOOCV). For an input of n subjects, LOOCV trains a CPM model by iteratively training on $n-1$ data points and validating on the left-out data point. The performance metrics obtained from all the iterations are then aggregated to evaluate the model's overall performance.

The CPM pipeline has many benefits, with its main strength being that the linear operations involved at each stage allow for very fast computations in comparison to alternative brain-behaviour

model methods, such as support vector regression (SVR). Additionally, by focusing on the most relevant features or edges in the brain, the amount of noise in the training data is reduced - particularly useful when dealing with large brain parcellations with lots of edge relationships. Finally, due to the use of edges in creating summary statistics, its predictive power is easier for neuroscientists to interpret, as the underlying brain connections that lead to a certain prediction can be easily identified. The latter is useful when comparing results with existing literature: the results can also be interpreted physiologically, helping to generate new hypotheses about network structure and function. However, it also has its limitations, for example, the linear model may not be optimal for capturing all complex or non-linear relationships between brain function and behaviour. Also, the authors reported that the predictive models tend to produce predictions within a smaller range than that of the true values: the model overestimates the behaviour of individuals with low functional measurements, and vice versa. The authors noted that other methods such as SVR may be more beneficial if the goal is to maximise model accuracy.

This method has been successfully used by researchers Lin *et al.* in predicting the ADAS-11 Cognitive Score composite from rs-fMRI data [8]. Here, they used Spearman's rank correlation (ρ) when statistically thresholding edges in the CPM model as the ADAS-11 scores in their chosen sample were skewed. The resulting edges were separated into a positive tail (associated with higher ADAS-11 scores) and a negative tail (associated with a lower ADAS-11 score). The ADAS-11 Cognitive score (as described in Section 2.5.2) is used to monitor the cognitive decline of AD patients - as such, this work shows that rs-fMRI can be used to reveal cognitive impairment in an ageing population. Using data from ADNI-GO and ADNI-2, the authors found that the CPM based on the positive edges significantly predicted the ADAS-11 scores of unseen subjects with $\rho = 0.49$. In contrast, the negative edge CPM did not result in a significant prediction ($\rho = 0.27$).

Chapter 4

Implementation Overview

Over the course of my project, I have amended the traditional whole-brain modelling pipeline described in Figure 2.1c into the pipeline shown below (Figure 4.1):

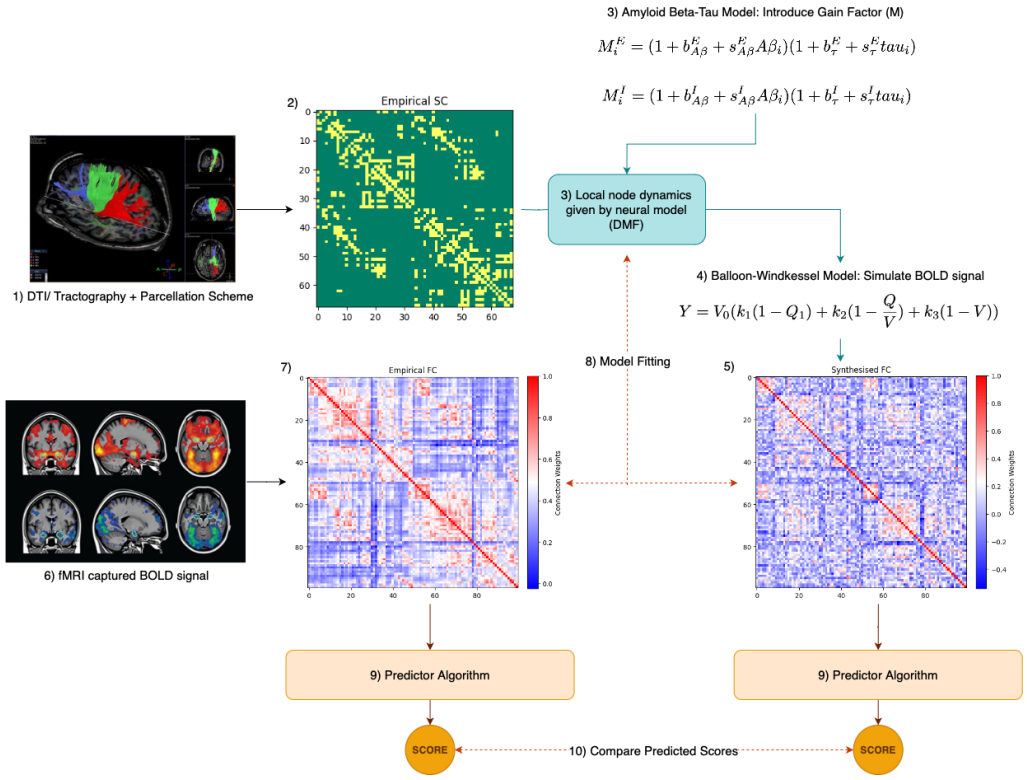

Figure 4.1: Adapted Whole-brain Modelling Pipeline

There are two main changes introduced here: the addition of modelling dynamic changes in Step 3 and the addition of the predictor algorithm in Steps 9-10.

Throughout the course of this project, we primarily focus on modelling the dynamic changes that arise from Amyloid-Beta (Aβ) and Tau depositions in the brain. The influence of these proteins on the DMF whole-brain model is added in Step 3, using inspiration from Patow's model [3] (Equations 3.1-3.4, Section 3.2.1). With this addition, we hope to improve the ability of our whole-brain model in simulating FC matrices of AD patients. In Step 9, the predictor algorithm finds a ventricular volume score for each FC matrix, and in Step 10 we compare these scores in an attempt to quantify the amount of degradation. This way, the pipeline quantifies how much the dynamic changes in Aβ and Tau deposits affected the functional connectivity.

We will be fitting this pipeline at group level for each of the main ADNI patient groups: the healthy controls defined as cognitively normal (CN), those with mild cognitive impairment (MCI) and those diagnosed with Alzheimer's Disease (AD). The process of adding these changes and fitting the Aβ-Tau model is explained in more detail in the following chapters.

Chapter 5

Fitting a Whole-Brain Model on Healthy Subjects

This section explores the initial steps taken to fit a whole-brain model on a control group of healthy patients (group name: Cognitively Normal - 'CN', from ADNI). The objective of this study is to determine constant parameters of the Dynamic Mean Field (DMF) model for cognitively normal subjects prior to introducing the degenerative effects of Amyloid-*beta* ad Tau in the subsequent stage of the pipeline (Chapter 6) In Section 5.2, we explore the steps taken to optimise the DMF model using deep-learning methods in the `Whobpyt` library. We then go on to describe the training methods to fit the DMF to ADNI data in Section 5.3, before concluding with the results and limitations of this fit in Section 5.4.

5.1 Dataset

When using the CNNM modelling pipeline from `whobpyt`, it is necessary to provide both a structural connectivity (SC) matrix and functional connectivity (FC) data as input. As mentioned previously, we will be doing group-level fitting of the DMF model, and as such required average empirical SC and FC matrices for subjects in the CN group.

For structural connectivity, we used healthy participant data from the Human Connectome Project (HCP) database [37] as described in Herzog *et al.*'s 'FastDMF' paper [6]. Here, the authors estimated structural connectivity by using the HCP dMRI dataset. The diffusion data was pre-processed and sampled per participant, transformed into MNI space and standardised to extract the structural connectomes from the Schaefer100 parcellation. Then, the structural connectomes of 32 healthy participants were averaged to obtain the group-level SC matrix.

Functional connectivity data was obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database [2], specifically from the ADNI-2 and ADNI-3 phases of the study that contained fMRI data. We used a sample of pre-cleaned and pre-processed BOLD resting state fMRI (rs-fMRI) data from 367 CN subjects (*mean age:* 70.2, *range of ages:* 56.5-89.1, *gender ratio:* male:female = 151:216) that were also parcellated according to the Schaefer100 parcellation. These rs-fMRI scans were 197 echoes long; while the usual ADNI protocol is to scan for about 10 minutes (200 echoes), we found that some 200-echo length scans were corrupted and opted for the alternative set of 197-echo length scans. To find the group-average FC matrix, we normalised each subject scan by taking the z-score, found the correlation matrix between regions in the BOLD time series using the Pearson Correlation Coefficient and once again averaged the FC matrices across the sample set.

5.2 Model Implementation

We have chosen to fit a DMF whole-brain model to AD patient data due to the biologically realistic nature of the model. Since we hope to model AD biomarkers in the pipeline, it is useful to use a model with a biologically inspired parameter set that can easily incorporate these markers, as seen in Patow's work [3] in Section 3.2.1. We chose to use the DMF model (`whobpyt.models.RWW`) provided in Griffith *et al.*'s `whobpyt` library (discussed earlier in Section 3.1.1), where the training

process is implemented using deep learning methods to optimise parameter fitting. Our training method has deviated from that presented in Griffiths' paper, as we fit a smaller subset of model parameters and adapt the codebase to better suit the project goal and the available patient data.

5.2.1 DMF Model Parameters to Fit

While the notation used by Griffiths closely follows the DMF model outlined by Deco in Equations 2.6-2.11, the authors introduced some alternative notation for terms J_{NMDA}, J_i, w_+ (given in Table 2.1) and G (long-range excitatory couplings) [12]. Instead, Griffiths used the within-node synaptic gains g^{EI} (excitatory-to-inhibitory), g^{IE} (inhibitory-to-excitatory) and g^{EE} (excitatory-to-excitatory), and the long-range global coupling g in the paper and library `whobpyt`, where $g^{EI} = J_{NMDA}$, $g^{IE} = J_i$, $\frac{g^{EE}}{g^{EI}} = w_+$ and $\frac{g}{g^{EI}} = G$.

When considering which subset of these parameters to fit during the training process, we consulted the literature inspiring the next stages of the pipeline to see what model parameters would be the most crucial to fit for a cognitively normal brain. For example, consider the pipeline followed by Patow (Section 3.2) when modelling Amyloid-β and Tau: they chose to find an optimal value for G from the healthy DMF model, which they re-used the Aβ-Tau model. As G scales the input received by each region from the rest of the brain network, tuning it moves the overall model closer to an optimal working point where a desired statistical feature of the empirical data can be reproduced, such as functional connectivity [6]. However, this is under the assumption that all diffusion MRI-reconstructed streamline fibres in the connectome had a similar conductivity, and therefore the coupling between different regions of the brain could be assumed to be scaled by this single global parameter [3].

Following this assumption and methodology, we also chose to find the optimal G, which in the case of `whobpyt` meant that the parameters to tune were g and g_{EI}. The methods used for hyperparameter tuning are discussed later in this chapter.

5.2.2 Modifications to the `whobpyt` Library

To make the `whobpyt` more suitable to the input data and the goals of our whole-brain modelling pipeline, we had to implement two main changes:

1. **Modifying the type of functional connectivity data given as input to the model:** The `whobpyt` library was originally set up to do subject-level fitting by taking as input a single patient's empirical SC matrix and BOLD time series (rs-fMRI data). In contrast, we aim to do group-level fitting for each patient group given an average healthy SC matrix and the average group functional data. However, it is difficult to provide an average BOLD signal for a patient group, as this tends to vary significantly and lack correlation between patients. Instead, it is better to compute an FC matrix from the BOLD signals of each subject in a patient group, and find an average of these. Since the FC matrices provide a common framework for comparing brain activity across subjects, averaging these can help normalise individual differences and help identify consistent patterns of connectivity across a patient group. Therefore, we wanted to provide an empirical group-level average FC matrix as input to `whobpyt`'s training function.

   This required some modifications to the training loop. Originally, it was designed to simulate a short period (window) of a BOLD time series, compare it to the respective window in the empirical BOLD data, and calculate a training loss (as described previously in Section 3.2). This cost function involves calculating an FC matrix from both the simulated and empirical windows, finding the Pearson correlation coefficient between both matrices, and returning the negative log-likelihood of this.

   Instead, we removed the dependency on the empirical BOLD time series and only provided the empirical FC matrix as input to the cost function, reducing the redundancy of computing it again at each window. Overall, by addressing the use of the empirical data when computing the loss, we were able to change the training function to take the average FC matrix as input. A similar modification was done to the evaluation function, as it also followed the previously described steps of converting a time series to a correlation matrix before finding the similarity. This improved the overall fit of the model, since it was optimising to fit the same empirical FC matrix as opposed to varying it per window.

2. **Splitting the prediction and evaluation steps:** Additionally, we modified the evaluation function to split up the prediction step from the evaluation step. The subsequent modification was more of a bug fix than a performance enhancement. We introduce the function `simulate()` in the model to simulate a BOLD time series and corresponding FC matrix. This is then passed as input to the evaluation function, or used independently to simulate FC data from the final fitted model. The scope of the evaluation function is now limited to only calculating the Pearson correlation coefficient (R) between the simulated FC and empirical FC data.

5.3 Model Training

5.3.1 Training and Evaluation Methods

When fitting the DMF model on CN functional patient data, we split the aforementioned ADNI sample into a training and test set (80%-20% respectively). We then computed the average FC matrix for each subset: the average training FC matrix was used as the empirical functional input to the DMF model during training; then, the simulated output of the trained model was compared to the average test FC matrix (Figure 5.2a, left) during model evaluation.

Although, we are interested in finding the value of parameters g and g_{EI}, initial experiments showed that varying the parameter g affected the overall performance of the model more than g_{EI}. For this reason, we focused on fitting g as the training goal and used the same values as Griffiths *et al.* for the other parameters (g_{EI}, g_{EE}, g_{IE}), which are not learnt by the model.

`Whobpyt` makes use of a Variational-Bayes based objective function [12] during training, and calculates the Pearson Correlation Coefficient, R, between a simulated FC matrix (FC_{sim}) and empirical (test) FC Matrix (FC_{test}) during model evaluation. This is given by Equation 5.1, where $\texttt{Cov}(FC_{sim}, FC_{test})$ is the covariance of the simulated and empirical matrices, σ_{sim} is the standard deviation of matrix FC_{sim} and σ_{test} is the standard deviation of matrix FC_{test}:

$$R_{sim,test} = \frac{\text{Cov}(FC_{sim}, FC_{test})}{\sigma_{sim}\sigma_{test}} \tag{5.1}$$

The coefficient R is bounded with $-1 \leq R \leq 1$, with $R = -1$ describing strongly negative correlation, $R = 1$ strong positive correlation and $R = 0$ no correlation. The goal of the model is to maximise the positive correlation between the simulated and empirical FC matrices, thus maximising the value of R.

5.3.2 Hyperparameter Tuning using Optuna

To find the best hyperparameters for the DMF model, we conducted a hyperparameter tuning study of 150 trials using the `Optuna` framework in Python. The `Optuna` framework can be beneficial as it uses Bayesian and stochastic optimisation techniques to focus on the most promising region of a large parameter space. This leads to faster convergence and a more efficient tuning process than the traditional RandomSearch or GridSearch methods from Python's `scipy` library. As previously discussed, we used the training set of rs-fMRI scans to train the DMF model on for each permutation of hyperparameters and used the test set of rs-fMRI scans as the validation set to evaluate their performance. The hyperparameters chosen to tune are as follows:

- *Learning rate* - this determines the size of steps taken by the ADAM optimiser and can affect the convergence, stability, and performance of the network. We wanted to test for learning rates of orders of magnitudes smaller and larger than the value used by Griffiths *et al.* (0.05) [12].

- *Number of Epochs* - this affects how long a model trains for, which is a trade-off between running for too long and overfitting, and not training long enough which results in a high training loss and poor performance. In the case of the DMF model, more epochs also made the training more computationally expensive. We needed to find the balance between this and model performance.

- *Step size* - controls the rate at which DMF model parameters are updated during each forward pass. This is analogous to the learning rate of the `whobpyt` CNNM system, as the parameters

Hyperparameter	Optimal Value
Learning Rate	0.1
Number of epochs	50
Step Size	0.1
Initial value g	1000
Time points per window (`TPperWindow`)	20

Table 5.1: Optimal values for DMF hyperparameters found from `Optuna` tuning study

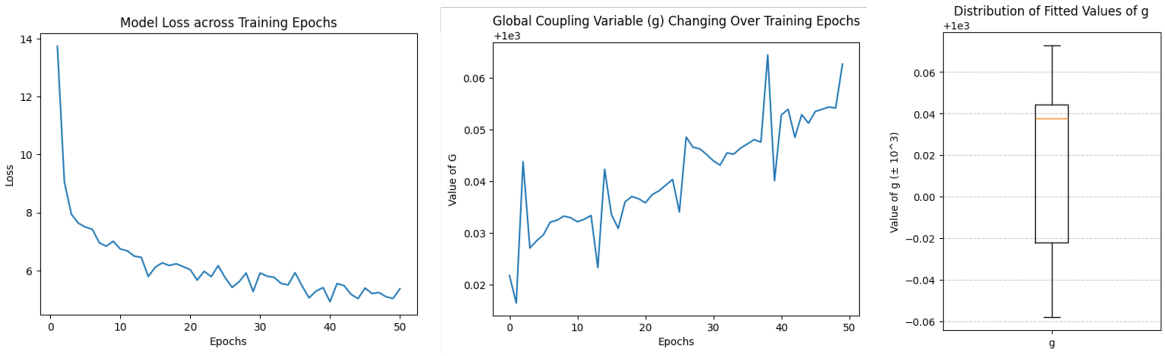

(a) (b) (c)

Figure 5.1: **Training curves of the final model**: (a) Loss curve across training epochs, (b) the evolution of fitting free parameter g for the best set of hyperparameter and (c) the distribution of the final fitted values for g over 20 training runs.

the optimised through ADAM are DMF model parameters too. We investigated a similar range of values as the learning rate for the step size.

- *Initial value for g* - since this is the parameter that the model is attempting to fit, giving a suitable initial value for this can greatly optimise the training process
- *Time points per window* (`TPperWindow`) - this is the number of data points simulated in one window of BOLD time series. Varying this parameter affects the granularity of the model: again a trade-off between computational efficiency and performance. Additionally, having too few data points per window increases the variation between windows, making it harder for the model to generalise.

The optimal set of values from the `Optuna` study can be seen in Table 5.1. As expected, the tuned learning rate of the system has the same value as the DMF step size (as in Griffiths' paper [12]), since they both control the update rate of the DMF model parameters and so should be equal. Figure 5.1 shows the behaviour of the DMF model during training when using these optimal hyperparameter values. Specifically, from Figure 5.1b, we can see that the value of g stays mostly stable during the training process as a result of tuning the initial value. In fact, across 20 training runs using the optimal model hyperparameters, the final fitted values of g have remained quite consistent, with the distribution shown in Figure 5.1c only having an interquartile range of 0.0667 (to 4d.p.). Across the 20 training runs, the mean fitted value for g was found to be 1000.0176 with a standard deviation of ± 0.0433. We will be using this value of g going forward for other models in the pipeline.

5.4 Results and Evaluation of Tuned DMF Model

The final DMF model was evaluated against the average FC matrix from the test set of subjects, yielding a mean Pearson correlation of $R = 0.4309$ (standard deviation of ± 0.0473) across 20 simulations. The results from best fit of this model are shown in Figure 5.2, with a Pearson correlation of $R = 0.5228$. The FC matrix for CN patients simulated by the model is shown in Figure 5.2a (right) and the correlation between the simulated and empirical edges in these matrices

is shown in Figure 5.2b. The latter illustrates the strong positive correlation of 0.5228 and shows a narrow confidence interval around the line of best fit, indicating a low variability between the predictions of each edge-pair.

When comparing these results to current literature, we can see that they align with the fit obtained by Griffiths *et al.* in their original work, where the simulated-empirical FC correlation was in the range $R = 0.4 - 0.9$ [12]. However, the average FC correlation found by those authors was $R = 0.65$, which is better than our average of $R = 0.43$. There are several reasons why our average score could be less, ranging from the parcellation we chose for our neural imaging data to the parameters we chose to fit.

Firstly, a lower correlation score could be attributed to using a higher resolution parcellation. While Griffiths opted for an 83-region Lausanne parcellation, we used a higher resolution Schaefer parcellation with 100 regions. While our fit is still an improvement from other modelling studies that used similar parcellation, there are several limitations associated with using a higher resolution parcellation. For example, authors Harita *et al.* used a 200-region Schaefer-Yeo parcellation when modelling rs-fMRI with the DMF model [13] and obtained an average fit of $R = 0.24(\pm 0.06)$ with a maximum of $R = 0.41$, markedly lower than our reported results. The authors noted that higher resolution brain parcellations increase the number of regions and connections, reducing the accuracy of simulated fits. Increasing the number of regions when parcellating empirical rs-fMRI data means averaging over fewer voxels, leading to a noisier result. Similarly, in the simulated rs-fMRI data, the additional nodes introduce individual noise terms in the DMF model, hence increasing the overall level of stochastic variability in the simulated DMF network. Choosing the resolution of a parcellation means considering the tradeoff between spatial granularity and computational feasibility; typical whole brain studies tend to opt for brain parcellations with around 66 regions, whereas those that use more biophysical models report fits to empirical data with an FC correlation of $R = 0.1 - 0.5$ [13].

Another pitfall of our model could be a result of only choosing to fit the long-range coupling parameter g in our model. As explored in the previous section, our final value found for this parameter was $g = 1000.0176$, which is considerably larger than Griffiths' reported value of $g = 82$ [12]. While this suggests a very strong coupling and coherence between brain regions, it could also be a result of the model over compensating for not learning the other synaptic gain parameters in the model, leading to issues later on in the model (discussed further in Sections 6.2.2 and 8.1).

Although our reported results show a sufficient fit of simulated-empircal FC data (and a marked improvement on other studies that use similar parcellations), it may be beneficial to consider alternative parcellations and more DMF model parameters as ways to improve performance.

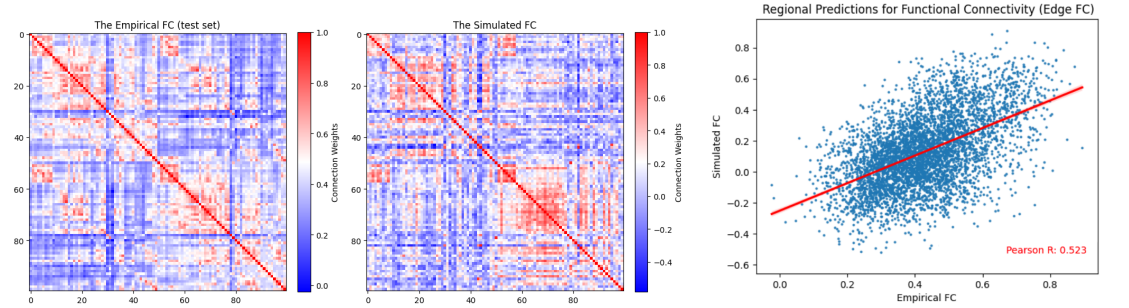

(a)

(b)

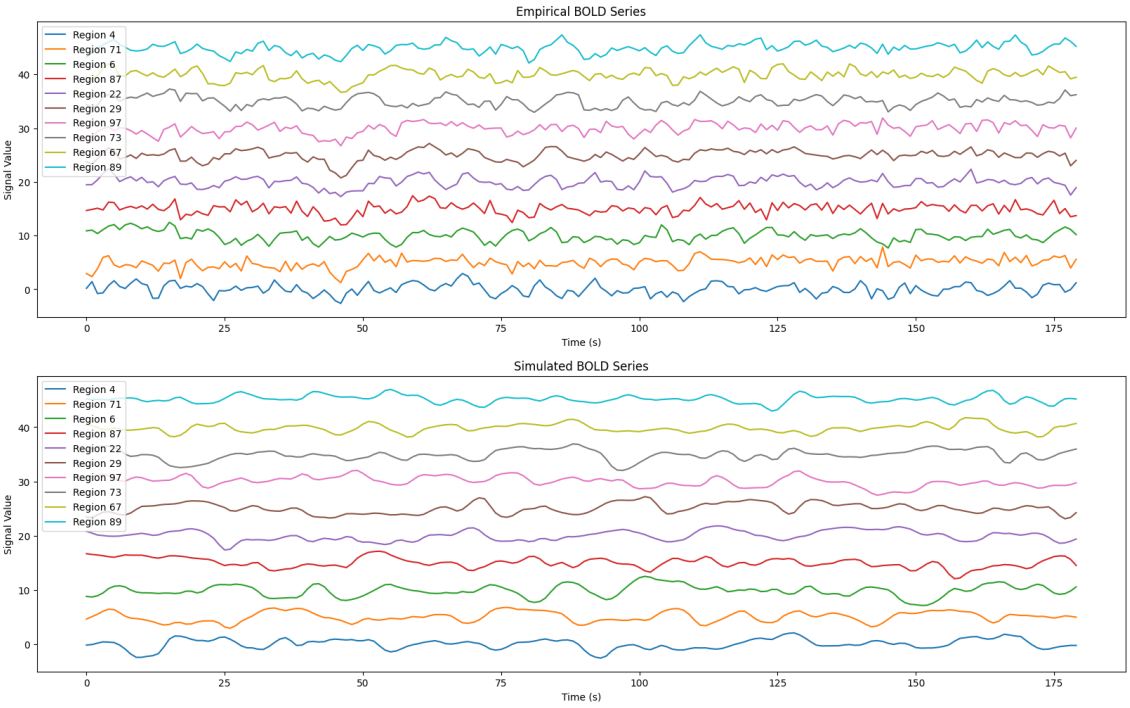

(c)

Figure 5.2: **Final fitted model for average CN patient data (best fit of $R = 0.523$)**: (a) the FC matrices - empirical test set (left) and simulated (right); (b) edgeFC graph: scatter plot of the empirical test set FC matrix values against the simulated estimates of those quantities (simulated FC); (c) comparison of the BOLD signals of a random subset of regions in the Schaefer100 parcellation between the empirical test set BOLD time series (top) and simulated BOLD time series (bottom)

Chapter 6

Whole-Brain Modelling of Aβ and Tau Degeneration

Once we have found the parameter values that fit the original DMF model to healthy patient data, the next step in the pipeline is to introduce the changes due to the proteins Amyloid-β (Aβ) and Tau into the model. This section explores the steps involved in developing the Aβ-Tau (DMF) Model. Section 6.2 covers the addition of `whobpyt.models.RWWABT` to the `whobpyt` library, which closely reflects the work of Patow *et al.* in modelling Aβ and Tau in the DMF model. To fit this new model, we follow a similar methodology as the previous chapter - the details of the hyperparameter tuning study are detailed in Section 6.3. Finally, we fit a unique version of the Aβ-Tau model to the three patient groups (CN, MCI and AD groups) as shown in Section 6.4.

6.1 Data Pre-processing: Aβ and Tau Scans

As before, the `whobpyt` modelling pipeline requires a structural connectivity matrix and functional connectivity data as input. However, with the added Aβ and Tau modelling changes, it is necessary to also provide data describing the Aβ and Tau loads per region of our considered parcellation.

For the SC input, we again use the averaged SC matrix of the healthy participants from the HCP database. As the goal of the overall pipeline is to evaluate Aβ-Tau changes on a healthy DMF model, we use this healthy SC matrix as a base for all three patient groups. Functional connectivity data was also obtained and pre-processed as before, however, this process was extended to a sample of 203 MCI subjects (*mean age:* 69.6, *range of ages:* 55.0-89.5, *gender ratio:* male:female = 105:98) and 35 AD subjects (*mean age:* 70.2, *range of ages:* 55.3-87.1, *gender ratio:* male:female = 20:15). Again, a group-wise average FC matrix was computed per patient group.

Data for Aβ was taken from the AV-45 pre-processed PET scans of 236 subjects (*mean age:* 75.6, *range of ages:* 59.0-95.0, *gender ratio:* male:female = 108:128, *patient group ratio:* CN:MCI:AD = 132:72:32) found in the ADNI database. These scans use a standard image and voxel size and were smoothed to a resolution of 6 mm. The pre-processing steps include applying coregistration, averaging the image across frames and intensity normalisation to cerebellar grey matter. In order to make the AV-45 scans applicable for our use, we undertook further processing steps. Firstly, we registered each image to the Montreal Neurological Institute (MNI) space. Then, we parcellated the scans according to the Schaefer100 parcellation (as shown in Figure 6.1) and

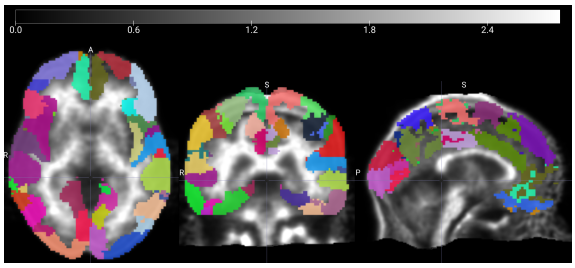

Figure 6.1: Schaefer100 parcellation overlaid on a AV45 PET scan of a cognitively healthy subject

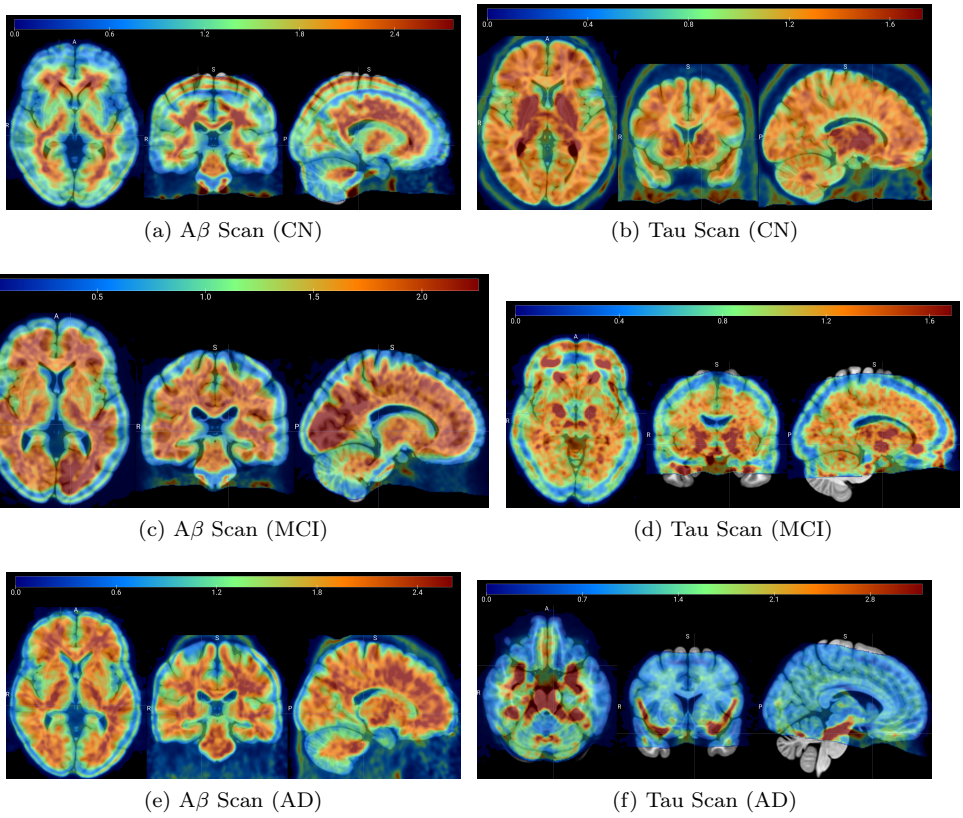

(a) Aβ Scan (CN)

(b) Tau Scan (CN)

(c) Aβ Scan (MCI)

(d) Tau Scan (MCI)

(e) Aβ Scan (AD)

(f) Tau Scan (AD)

Figure 6.2: **PET scans of Aβ (AV45) and Tau (AV1451) deposits in different patient groups**: (a, c, e) visualise the progression of Aβ deposits across patients in CN, MCI and AD groups respectively whereas (b, d, f) visualise the progression of Tau deposits across patients in the same groups

normalised them so that the regions sum to 1. Again, to find a group-level Aβ parcellation, we averaged across the samples in each patient group.

Similarly, Tau data was taken from the AV-1451 pre-processed PET scans of 296 subjects (*mean age:* 74.1, *range of ages:* 57.0-95.0, *gender ratio:* male:female = 128:168, *patient group ratio:* CN:MCI:AD = 168:92:36) from ADNI. The same acquisition and pre-processing steps were taken to get the group-wise average parcellation for each of the three patient groups. Figure 6.2 shows the acquired Aβ (left) and Tau (right) PET scans from a sample of singular subjects across the different patient groups (after being registered in MNI space).

6.2 Modifications of the `Whobpyt` Library

For this stage of the pipeline, we added a new model `whobpyt.models.RWWABT` to the `whobpyt` library. This builds on the previous DMF model (`whobpyt.models.RWW`), but introduces the modelling of Aβ and Tau degradation. Introducing these changes also required optimising other training functions in the library, as detailed below.

6.2.1 Including Aβ and Tau in the DMF Model

We took inspiration from the works of Patow *et al.* (see Section 3.2.1) to model the effect of proteins Aβ and Tau in a whole-brain model. As shown in Equations 3.1-3.4, this works by introducing a gain factor $M^{(E,I)}$ for each region in the DMF model when calculating the firing rates of the excitatory (E) and inhibitory (I) populations. $M^{(E,I)}$ is calculated using bias and scaling terms for the Aβ and Tau quantities, for both the excitatory and inhibitory neuron populations as we've seen in Equations 3.1 and 3.2 (repeated below for clarity):

$$M_i^E = (1 + b_{A\beta}^E + s_{A\beta}^E A\beta_i)(1 + b_\tau^E + s_\tau^E tau_i) \tag{3.1}$$

$$M_i^I = (1 + b_{A\beta}^I + s_{A\beta}^I A\beta_i)(1 + b_\tau^I + s_\tau^I tau_i) \tag{3.2}$$

To implement this in the `whobpyt.models.RWWABT` model, we add the bias ($b_{(A\beta,\tau)}^{(E,I)}$) and scaling ($s_{A\beta,\tau}^{(E,I)}$) terms as the free parameters of the model that are learnt through model training. The model takes the Aβ and Tau parcellations as constant vector inputs, alongside the empirical structural connectivity data. These are all used in the calculation of the gain factor, $M^{(E,I)}$, which is now included in the neuronal response function, as shown in Listing 6.1. The synaptic gain parameters considered in the previous chapter are now all initialised to a constant value: we now use a constant $g = 1000.02$ in the model.

6.2.2 Improving Numerical Stability: Introducing `softplus`

After including the modelling changes in `whobpyt.models.RWWABT` (including the adaptations from `whobpyt.models.RWW` such as supplying the empirical FC matrix rather than BOLD time series as input), we encountered issues when training the model. The modifications resulted in the gradient values becoming undefined (NaN) during the optimization step of the training loop. Although we are not yet certain of the cause of this, one hypothesis is that the introduction of $M^{(E,I)}$ into the neuronal response function was a source of numerical instability. The original implementation of this function was written using PyTorch's mathematical operations defined over tensors. As seen in Equations 3.1-3.4, the division could lead to numerical instability if the numerator becomes substantially larger than the denominator.

To overcome this, we rewrote this operation using PyTorch's `softplus` function, defined as $\texttt{softplus}(x, k) = \ln(1 + \exp^{kx})$ [38] where x is the input and k is a sharpness parameter. This is a smooth approximation of the ReLu function that is often used as an activation function in neural networks. For numerical stability, this function reverts to linear behaviour when $x \times k$ is larger than a threshold value. Using an approximation shown in Equation 6.1, we show below how we can use `softplus` in the neuronal response function (for example, Equation 3.1). The approximation is given by:

$$\ln(1 + \exp(x)) \approx \begin{cases} \ln 2 & \text{if } xd_E \ln 2 > \text{ threshold} \\ \frac{x}{1-\exp(\frac{-x}{ln2})} & \text{otherwise} \end{cases} \tag{6.1}$$

With this approximation, we can recover Equation 3.1. Let $x = M_i^E(a_E I_i^{(E)} - b_E)$. We use the softplus function with the smoothness parameter $k = d_E \ln 2$:

$$\begin{aligned}
\texttt{softplus}(x, d_E \ln 2) &= \frac{1}{d_E \ln 2} \times \ln(1 + \exp(xd_E \ln 2)) && \text{(from the definition of } \texttt{softplus} \text{ [38])} \\
&= \frac{1}{d_E \ln 2} \times \frac{(xd_E \ln 2)}{1 - \exp(\frac{-xd_E \ln 2}{\ln 2}))} && \text{(using the approximation 6.1)} \\
&= \frac{x}{1 - \exp(-xd_E)} \\
&= \frac{M_i^E(a_E I_i^{(E)} - b_E)}{1 - \exp(-d_E M_i^E(a_E I_i^{(E)} - b_E))} && \text{(substituting } x = M_i^E(a_E I_i^{(E)} - b_E))
\end{aligned}$$

This recovers Equation 3.1, as required. The code to implement this is shown below in Listing 6.1.

```
from torch.nn import Softplus

# Neuronal input-output functions of excitatory pools and inhibitory pools.
def h_tf(a, b, d, z, M, softplus_threshold):

    M = M.unsqueeze(1).expand_as(z)

    # input to softplus, x
    x = M * (a * z - b)
```

```
10
11      # sharpness parameter
12      ln2 = torch.log(torch.tensor(2.0))
13      k = d * ln2
14
15      softplus = Softplus(beta=k, threshold=softplus_threshold)
16      return softplus(x)
17
18  # Calculates the gain factor M for the neuronal response function
19  def calc_gain_factor(b_AB, s_AB, AB, b_t, s_t, tau):
20      return (1 + b_AB + s_AB * AB) * (1 + b_t + s_t * tau)
```

Listing 6.1: Example of using `softplus` in the neuronal response function of Aβ-Tau Model

Figure 6.3 shows the result of simulating the neuronal response function for the excitatory population with the original method (left) and the `softplus` implementation (right). For this simulation, we used the values for scaling and bias terms as given in Patow's paper for CN patients, and the default values of arguments a b, z and d provided in Griffith's code. We linearly increased the values of $b_E^{A\beta}$, $s_E^{A\beta}$, b_E^t and s_E^t between the ranges observed in a typical run of the model, and plotted the neuronal response produced by both variations of the function. The results show that for the range of values for bias and scaling terms considered here, both variations produce the same outcome, meaning that `softplus` can be used as a suitable replacement.

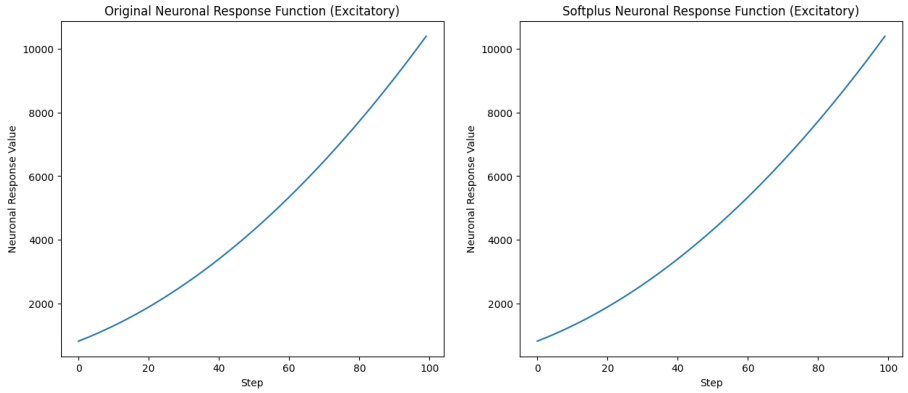

Figure 6.3: Comparison of simulating the neuronal response function for the excitatory population using the original method (left) and the `softplus` approximation (right)

6.2.3 Introducing Early Stopping

However, using `softplus` in the Aβ-Tau DMF Model causes it to deviate from the original DMF model's typical behaviour, reverting to linear behaviour at the points where the threshold limit is reached. This can be seen in the loss curve in Figure 6.4a (top), when at around 30 epochs, the neuronal response function changes to linear behaviour rather than calculating the firing rates based on Equations 3.1-3.4, hence negatively impacting the overall training of the model. This behaviour prevents the model from learning any further about the bias or scaling terms, as shown in the bottom plot in Figure 6.4a, where the values of these parameters remain mostly constant after the adjustment to linear behaviour

There are two ways in which we overcame this: firstly, by hyperparameter tuning the value of this threshold (discussed in Section 6.3); secondly, by introducing early stopping into the training loop.

Early stopping was implemented with a patience period of 5 epochs: if the model training loss did not improve after 5 epochs, the model stopped training. Given the variability of the model loss - particularly in the early stages of training - this patience period was chosen to balance premature stopping and training for too long (falling back into linear behaviour). Figure 6.4b illustrates the benefits of early stopping: as we can see from the top graph, the training curve satisfies the expected shape and the bottom plot shows that the parameter values change (are learnt) through training. By utilising early stopping, we can capture the minimum training loss, and the parameter values that cause this.

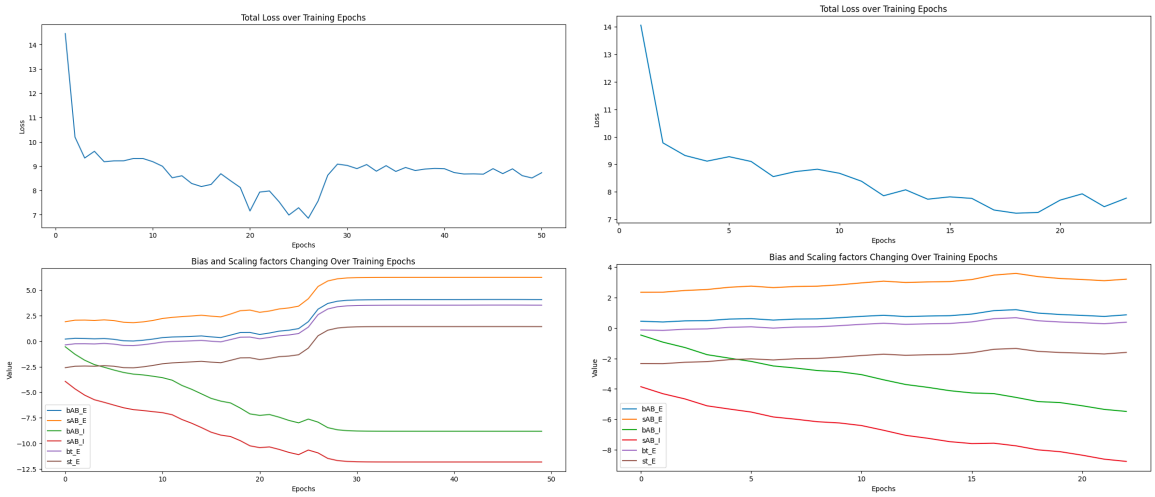

(a) Without early stopping, threshold of 20

(b) With early stopping, threshold of 20

Figure 6.4: Comparing the model behaviour before (a) and after (b) implementing early stopping in the training function

6.2.4 Evaluation of Implementation Changes

It is worth noting that while using `softplus` and early stopping avoids the underlying issue of undefined gradients, it does not address what causes it. One hypothesis is that the firing rates grow too much due to the addition of the gain factor $M^{(E,I)}$ within their calculations, which could push the calculations of the response function to the limit of its numerical stability. This may be a fault in our initial assumption when fitting parameters in the original DMF model, where we assumed that it was sufficient to only fit the single global parameter G as all the streamline fibres have the same conductivity. Instead, it is worth considering that the model parameter g_{IE} should be scaled per node with respect to its coupling strength. Since we do not fit this parameter in the original DMF model, it could be the case that its value is not sufficient to stabilise the firing rates, causing it to reach the limits of numerical stability. This is something to investigate further in the future.

6.3 Hyperparameter Tuning

When finding the optimal values to train the Aβ-Tau model with, we tuned mostly the same parameters as in Section 5.3.2 except for now tuning the bias ($b^{(E,I)}_{(A\beta,\tau)}$) and scaling ($s^{(E,I)}_{A\beta,\tau}$) terms instead of the global coupling parameter g, and tuning the threshold value for the `softplus()` function.

As illustrated in Figure 6.5, the `softplus()` threshold value can greatly impact the model performance and training behaviour. We can see that there is no clear trend in the impact of increasing the threshold: while increasing the threshold from 10 to 20 delays the return to linear behaviour by around 10 epochs, increasing from 20 to 50 does not have the same effect. When the threshold is 50, we see that the model peaks earlier and twice within the same number of epochs. Additionally, when the model first peaks at around 20 epochs, it does not return to linear behaviour but continues to fit the model parameters (as we can see on the right-hand side plots). This is similar to the behaviour when the threshold is 10, as although the second peak in the training curve is not evident, in the top right plot we can see the fit of the parameters significantly change again at around epoch 36.

We once again conducted a hyperparameter tuning study using the `Optuna` framework in Python, this time fitting a model for each patient group separately. Again, we separated the empirical fMRI data for each patient group into a training and testing set and used the training set for hyperparameter tuning and the test set for validation.

Using the healthy DMF model for guidance, the hyperparameters we chose to tune were the learning rate, the threshold for `softplus` and the initial values for the six free parameters that

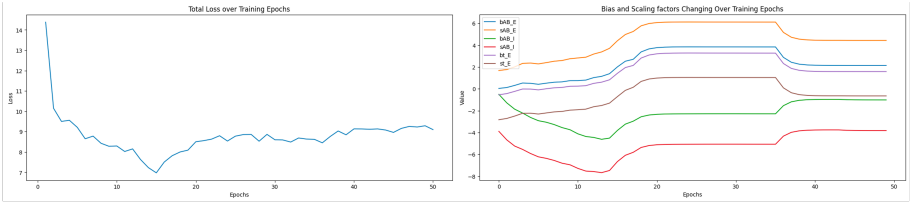

(a) Threshold = 10

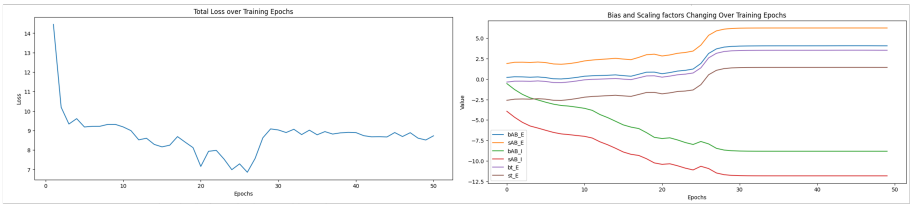

(b) Threshold = 20

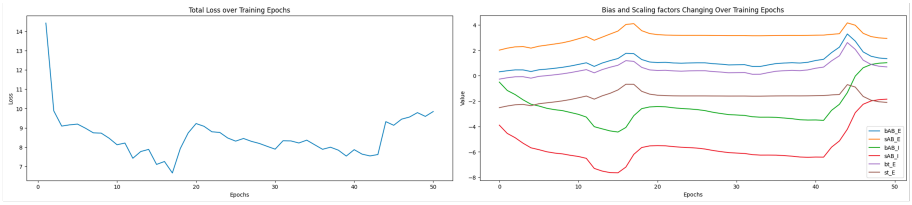

(c) Threshold = 50

Figure 6.5: **Comparing the model behaviour using different threshold values in the `softplus` function**: for all (a, b, c) above, the graph on the left illustrates the training loss across 50 epochs of training, while the graph on the right shows the evolution of the free parameters of the Aβ-Tau Model across 50 epochs of training

arise from the bias and scaling terms of the model: $b_E^{A\beta}$, $s_E^{A\beta}$, $b_I^{A\beta}$, $s_I^{A\beta}$, b_E^t and s_E^t. Note that it is not necessary to fit b_I^t or s_I^t, as Tau only targets excitatory neurons [3] (as discussed in Section 3.2.1). Furthermore, as described in the work of Patow *et al.*, we can constrain the ranges of some free parameters given their biophysical behaviours. Therefore, we have that $s_E^{A\beta} > 0$, $s_E^t < 0$ and $s_I^{A\beta} < 0$ [3]. The optimal set of parameters from the `Optuna` study for each patient group model can be found in Table 6.1.

6.4 Results and Evaluation

The optimal hyperparameters presented in Table 6.1 were used to train three different versions of the Aβ-Tau model for each patient group, ensuring that the matching Aβ-Tau parcellations and empirical fMRI BOLD data for that patient group were used as input. The models were evaluated against the average empirical FC matrix from each patient group test set, resulting in a mean Pearson correlation of $R = 0.6031$ (standard deviation of ± 0.0343) for the CN patient

Hyperparameter	CN Model	MCI Model	AD Model
Learning Rate	0.05	0.05	0.05
Softplus Threshold	31	88	73
Initial $b_E^{A\beta}$	0.4321	-1.0267	-2.0309
Initial $s_E^{A\beta}$	0.3604	3.4428	1.8477
Initial $b_I^{A\beta}$	-0.500	-3.2613	-4.8600
Initial $s_I^{A\beta}$	-3.000	-2.6306	-7.6439
Initial b_E^t	-0.6482	-0.1555	-1.6266
Initial s_E^t	-2.7856	-1.2906	-3.7093

Table 6.1: Optimal hyperparameters for the Aβ-Tau Model for each patient group

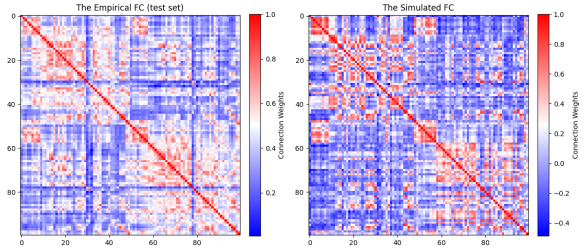

(a) CN: Empirical FC (left), Simulated FC (right), $R = 0.647$

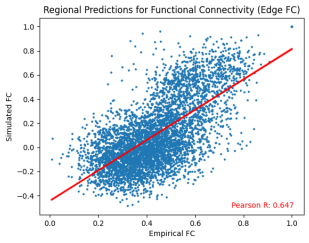

(b) CN EdgeFC graph

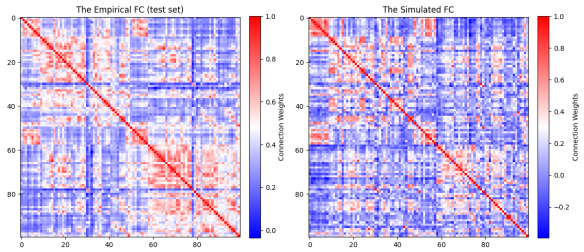

(c) MCI: Empirical FC (left), Simulated FC (right), $R = 0.582$

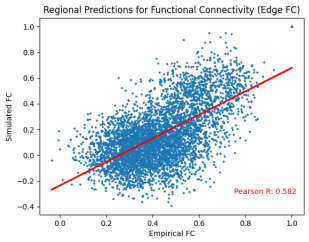

(d) MCI EdgeFC graph

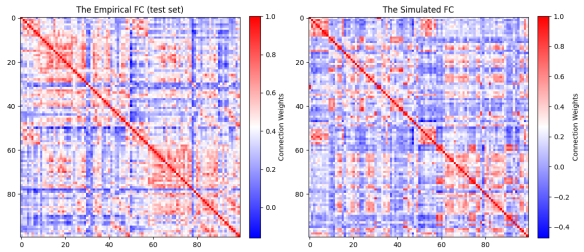

(e) Empirical FC (left), Simulated FC (right), $R = 0.505$

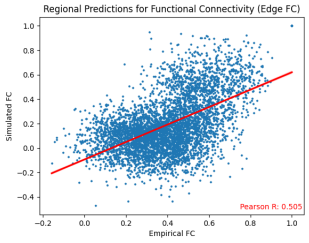

(f) AD EdgeFC graph

Figure 6.6: **Simulations from the best fitted Aβ-Tau DMF model for each patient group**: (a), (c), and (e) show the empirical FC (left) per patient group, and the simulated FC (right), whereas (b), (d) and (f) illustrate the correlation of the values of the empirical test set FC matrix against the simulated estimates per edge

group, $R = 0.5721$ (standard deviation of ± 0.0460) for the MCI patient group and $R = 0.5490$ (standard deviation of ± 0.0313) for the AD patient group, across 20 simulations. We can already see that adding the Aβ-Tau heterogeneity into the DMF model has improved the model fit for the CN group by 39.96% on average (from an average of $R = 0.4309$ to $R = 0.6031$). Figure 6.6 shows examples of the best simulation per model: $R = 0.6474$ for the CN group, $R = 0.5816$ for the MCI group and $R = 0.5054$ for the AD group. The training curves and simulated BOLD time series for each model can be found in Appendix A.1.

Table 6.2 shows the average fitted values of the bias and scaling terms across 20 training runs per model. The distributions of these are illustrated in Figure 6.7. We can see that there is no clear trend for how parameter values differ between patient groups: although the mean and median values of $b_E^{A\beta}$ and $s_I^{A\beta}$ decrease through the patient groups, this is not the case for the other parameters. For all parameters, the mean and median values for the AD patient group are much lower than the CN patient group, however, the behaviour of the MCI group is more variable. This could be because the MCI phase is a transitional phase of neurodegeneration, and therefore harder to distinctly classify. In fact, within the ADNI merge database, the patient data for the MCI phase is split into early-stage MCI (EMCI) and late-stage MCI (LMCI). As we did not have this granularity with the Aβ-Tau PET data, we opted to combine these stages into an overall MCI group. However, it could be that the EMCI and LMCI stages need different parameter values to model sufficiently. For example, both inhibitory Aβ parameters ($b_I^{A\beta}$ and $s_I^{A\beta}$) have comparatively

Patient Group	$b_E^{A\beta}$	$s_E^{A\beta}$	$b_I^{A\beta}$	$s_I^{A\beta}$	b_E^t	s_E^t
CN	3.1147 (0.8022)	3.6591 (0.8637)	-4.2504 (1.0139)	-6.3978 (1.0423)	2.1607 (0.8038)	0.6484 (1.0344)
MCI	1.8500 (0.9603)	7.1747 (1.3572)	-9.5482 (2.1452)	-8.6484 (2.1723)	2.6701 (0.9800)	1.8357 (1.3134)
AD	-5.8092 (2.5754)	-3.8387 (2.9995)	-9.0833 (1.2748)	-11.6546 (1.3000)	-5.3601 (2.5670)	-8.5466 (2.5041)

Table 6.2: **Final learnt parameter values for Aβ-Tau Model:** Mean final value of each parameter across 20 training runs, with the standard deviation given in the corresponding brackets.

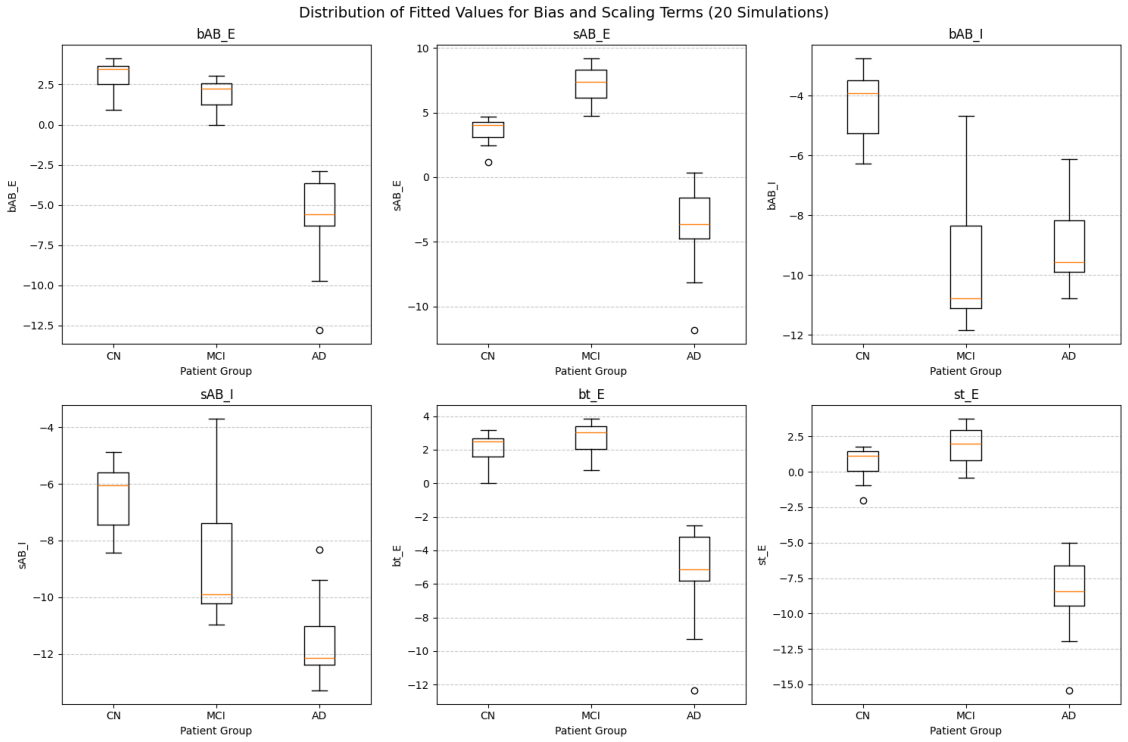

Figure 6.7: Parameter values found across 20 training runs of the final Aβ-Tau Model across the CN, MCI and AD patient groups

large ranges and interquartile ranges in their box plots, indicating significant variability in the parameter values.

Additionally, we conducted two-sample t-tests for each pair of patient groups for each parameter shown in Figure 6.2, correcting for multiple comparisons by using the Benjamin-Hochberg procedure to control the False Discover Rate (FDR). The FDR-corrected p-values for each t-test can be found in Table A.1. Apart from the value of b_E^t between the CN and MCI patient groups, all other learnt parameters are significantly different between the different patient groups. This provides evidence that our different models are capturing the difference in the input features: although all models are given the same SC matrix, the differences in scaling and bias parameters show the contribution of the different Aβ and Tau parcellations for each model, highlighting the heterogeneity it provides. Overall, the p-values between the different parameters across the patient groups are more significantly different between the CN and AD groups, although for some parameters such as $s_E^{A\beta}$ and s_E^t their difference is more significant between the MCI and AD groups. However, since our AD sample size is substantially smaller than the CN and MCI groups, we get more variation in parameter values and a higher frequency of outliers across AD parameters compared to the other patient groups.

It is worth noting that the `whobpyt` CNNM modelling pipeline also works to fit the structural connectivity weights of the provided SC matrix. This means that the final models obtained for each

patient group contain both structural (from the fitted SC matrix) and dynamic (from the fitted Aβ-Tau bias and scaling terms) changes, and combined they result in the reported model performance. The range of our results for Aβ-Tau bias and scaling terms is larger than those reported in the work of Patow *et al.*, and this could be due to the impact of the underlying structural connectivity changes too. The relationship between the structural and dynamic changes is investigated further in Section 7.5.

Chapter 7

Developing a Clinical Measure Predictor

The final stage in the adapted whole-brain modelling pipeline is developing a model to predict observable clinical measures from a subject's functional connectivity data. In this chapter, we outline the different modelling techniques that were investigated for this purpose. The problem can be seen as a regression task: we have a dataset of FC matrices (features) with a corresponding clinical measure (target), and the goal is to predict the clinical measure for an unseen FC matrix. It is important to note that the term 'predict' here is used loosely and is better defined as 'correlate' [36]: we see how functional connectivity can be correlated with patient measures, and how regression techniques and other statistical methods can be employed to use this for forecasting onto novel data.

In this chapter, we discuss the different types of clinical measures considered as the target groups (Section 7.1) and the different modelling techniques trialled per clinical measure (Section 7.3). We discuss the challenges of using data from clinical patient measures and some methods to overcome this in Section 7.2. Finally, we present the performance of our best model that predicts ventricular volume scores (Section 7.4), before seeing the overall effect of combining this with the previous Aβ-Tau Model (Section 7.5).

7.1 Assessing Clinical Target Groups

While developing the predictor model, we considered three clinical measures (available from ADNI) that could be used as the target group: the ADSP Phenotype Harmonization Consortium (PHC) Composite Score, the Alzheimer's Disease Assessment Scale-Cognitive Subscale (ADAS-Cog) 11 Score, or the measure of ventricular volume. The first two are composite cognitive scores, whereas the latter is an objective biological measurement taken from patient scans, as explained in Sections 2.5.2 and 2.5.2.

Composite cognitive scores provide a broader assessment of cognitive abilities and are more stable and reliable than individual studies. These are a better choice of target label than individual tests, as they can be more generalisable. Furthermore, individual tests such as the mini mental state examination (MMSE) are often used for diagnostic purposes, so predicting these from simulated MCI or AD FC matrices could become circular. Previous work [39] has warned of limitations of trying to forecast cognitive scores such as ADAS-Cog 11 from patient data, so we also chose to investigate the neuroimaging biomarker of ventricular volume. Given its biological basis, this is a more objective measure than cognitive studies: there is minimal influence from external factors (from a subject's environment to mood, or examiner bias) on the results.

We explore the challenges of using each clinical measure as a target category in the following sections.

7.2 Data Pre-Processing and Sampling

7.2.1 Pre-processing Features (FC Matrices)

We use the same functional connectivity data as the feature set for the predictor model that we used as empirical input for training the Aβ-Tau Model. This time, however, we combine the samples from each patient group and computed the FC matrix per subject: we normalised each fMRI scan by taking the z-score and found the correlation matrix between regions in the BOLD time series using the Pearson correlation coefficient. Since the FC matrices are symmetrical along the leading diagonal (as functional connectivity between edges is bidirectional), we chose to only use the lower triangle of each FC matrix. This reduces the dimension of each feature in the feature set and improves the overall model performance.

7.2.2 Datasets for Clinical Measures

We obtained two datasets from ADNI for the three different clinical measures considered as target groups, corresponding to the subjects we had functional connectivity data for. The ADAS PHC Composite Dataset gave us details on the four PHC composite cognitive scores, whereas the larger ADNIMERGE table incorporates all the main information per subject from the ADNI database, including both ADAS-Cog 11 scores and ventricular volume measurements. The details for each target category is given below:

1. ***ADAS PHC Composite Dataset*** - the ADAS PHC Dataset (1342 subjects) contains information on four different composite scores: memory (PHC_MEM), executive function (PHC_EXF), language (PHC_LAN) and visuospatial abilities (PHC_VSP). The distributions for these can be seen in Figure 7.1. The distributions are positively skewed for all the composite scores, although more severely in the case of PHC_VSP. This is due to the underlying trend in the dataset that patients with a higher score for any of the composites tend to be diagnosed as more cognitively normal, and our fMRI sample contains a higher proportion of cognitively normal subjects. The ADAS-PHC Composite Dataset contains a derived clinical diagnostic rating for each patient group: '1' for the CN patient group, '2' for MCI and '3' for AD. We can see a negative correlation between the diagnostic score and cognitive score per composite; the Pearson correlation, R and respective p-value, p, (thresholded at $p < 0.05$) per composite score is as follows:

   - PHC_MEM: $R = -0.7571$, $p =$ 6.5404e-249
   - PHC_EXF: $R = -0.6020$, $p =$ 1.2285e-132
   - PHC_LAN: $R = -0.5712$, $p =$ 1.2349e-116
   - PHC_VSP: $R = -0.3550$, $p =$ 5.5023e-41

   As these are independent composite scores and all follow quite different distributions, when using ADAS PHC as the target category we developed a separate regressor model per composite score. This way, we could observe how dynamic or structural changes affected each composite and in what proportion.

2. ***ADAS-Cog 11 Dataset*** - this collection from the ADNIMERGE data set contained 1339 subjects (*mean age:* 71.2, *range of ages:* 55.0-89.6, *gender ratio:* male:female = 660:679) and is distributed as shown in Figure 7.2. Here we see a significant negative skew for the same reason as before: lower ADAS-Cog 11 scores are associated with more cognitively healthy patients (a larger proportion of which we have in our sample). The plot also shows that the dataset is severely imbalanced, which we address using sampling techniques discussed below.

3. ***Ventricular Volume Dataset*** - this collection from the ADNIMERGE dataset contained 1262 subjects (*mean age:* 71.0, *range of ages:* 55.0-89.6, *gender ratio:* male:female = 625:637), again with a negatively skewed distribution (Figure 7.3) due to enlarged ventricular volume being associated with the progression of AD and our dataset containing a higher proportion of healthy patients. All ventricular volume scores are normalised by dividing by the total intracranial volume (referred to as 'Ventricles_ICV' or 'Ventricular_ICV' scores).

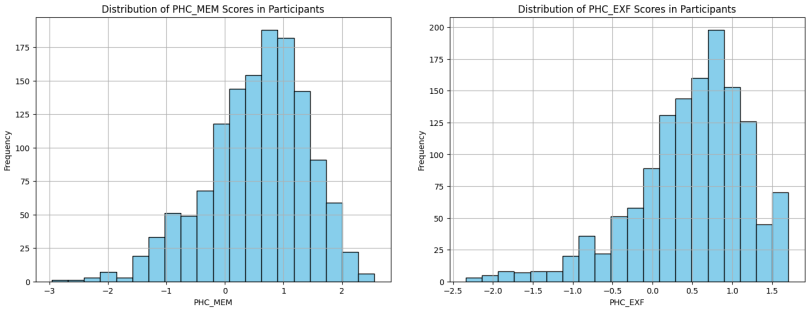

(a) PHC Memory Composite (b) PHC Executive Function Composite

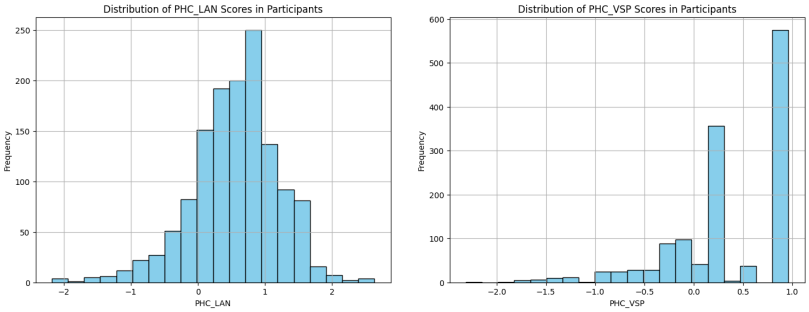

(c) PHC Language Composite (d) PHC Visuospatial Composite

Figure 7.1: Distributions of PHC Composite Scores for each composite type across the used sample

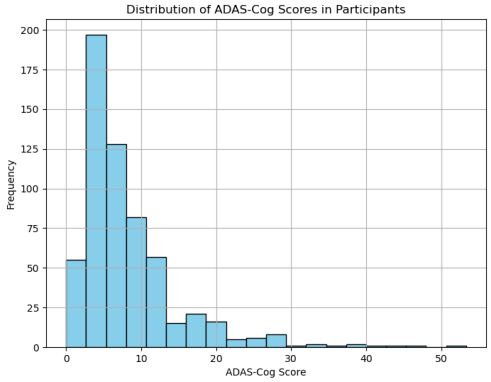

Figure 7.2: Distribution of ADAS-Cog 11 Score across the used sample

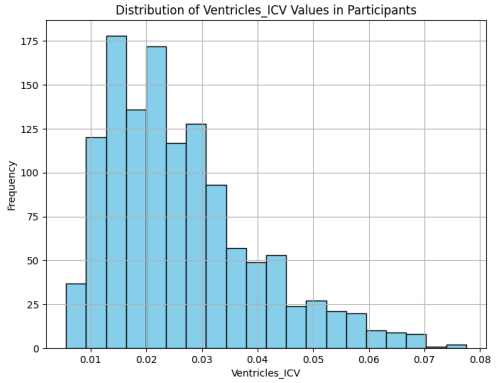

Figure 7.3: Distribution of Ventricular Volume Measure across the used sample

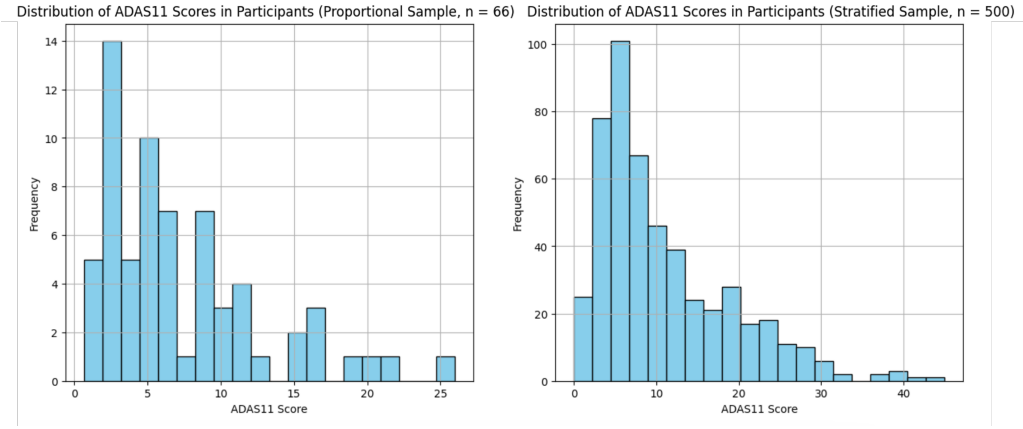

Figure 7.4: Changes in ADAS-Cog 11 Score distribution when employing different sampling techniques: proportional sampling (left) and stratified sampling (right)

7.2.3 Issues within the Dataset

Imbalanced Datasets

A common issue amongst all the target categories was the imbalance is the datasets used. For all scores, this was caused by having more data on the cognitively normal end of the spectrum and fewer samples from AD patients. One way in which we overcame this was to use sampling techniques such as proportional or stratified sampling to see whether this would improve model performance. Figure 7.4 illustrates the results of two types of sampling techniques for the ADAS-Cog 11 target category: proportional (left) and stratified sampling (right). For both sampling methods, we categorised the ADAS-Cog 11 dataset by the baseline diagnosis reported in the ADNIMERGE table. This contains the categories spanning across patient groups: CN, SMC (significant memory concern), EMCI (early stage MCI), LMCI (late stage MCI) and AD. It is worth noting that the AD patient group spans a larger range of ADAS-Cog 11 scores than CN or SMC, resulting in the still seemingly negatively skewed plots of Figure 7.4.

The first type of sampling considered was proportional sampling. Here, we wanted to ensure that each patient group was represented in our final sample, but in proportion to how much they are represented in the original dataset. We chose to randomly select 5% of each patient group in our final sample, leading to an overall smaller sample of 66 subjects. In contrast, in stratified sampling, we chose to randomly select the same number of subjects per patient group. We chose to use the size of the smallest patient group (AD: 100) as our sampling number, resulting in an overall sample of 500 subjects.

Although proportional sampling results in a sample that's more representative of the dataset, stratified sampling ensures that each patient group is adequately represented in the final sample, addressing the dataset imbalance more. For our case, stratified sampling is more suitable, as there are significant differences between the proportions of each patient subgroup. Overall, both sampling methods work to reduce selection bias and variability in the data.

High Feature Dimensionality

As we are using the Scahefer100 parcellation for the fMRI BOLD time series, each feature (FC matrix) consists of 10000 elements. The high feature dimensionality here increases the risk of regressor models overfitting and the risk of redundant features, negatively impacting model performance and increasing computational complexity. While taking the lower triangle from each FC matrix reduces each feature to only contain 4950 elements, we wanted to experiment with further dimensionality reduction to see if this improves model inference.

One method of doing this was using principal component analysis (PCA) to choose, for example, the first 250 principal components from each feature. The number of principal components used was a hyperparameter that needed tuning per model. This was used when considering the CNN, Support Vector Regressor or Gradient Boosting regressor models.

Alternatively, other methods [8] overcame this issue by finding a summary statistic per FC matrix, and training a regressor model using these 1-dimensional summary statistics as the feature set. This is used in methods such as the Connectome-based Predictive Model (CPM), where the summary statistic is the sum of edges in the FC matrix that are significantly correlated with the

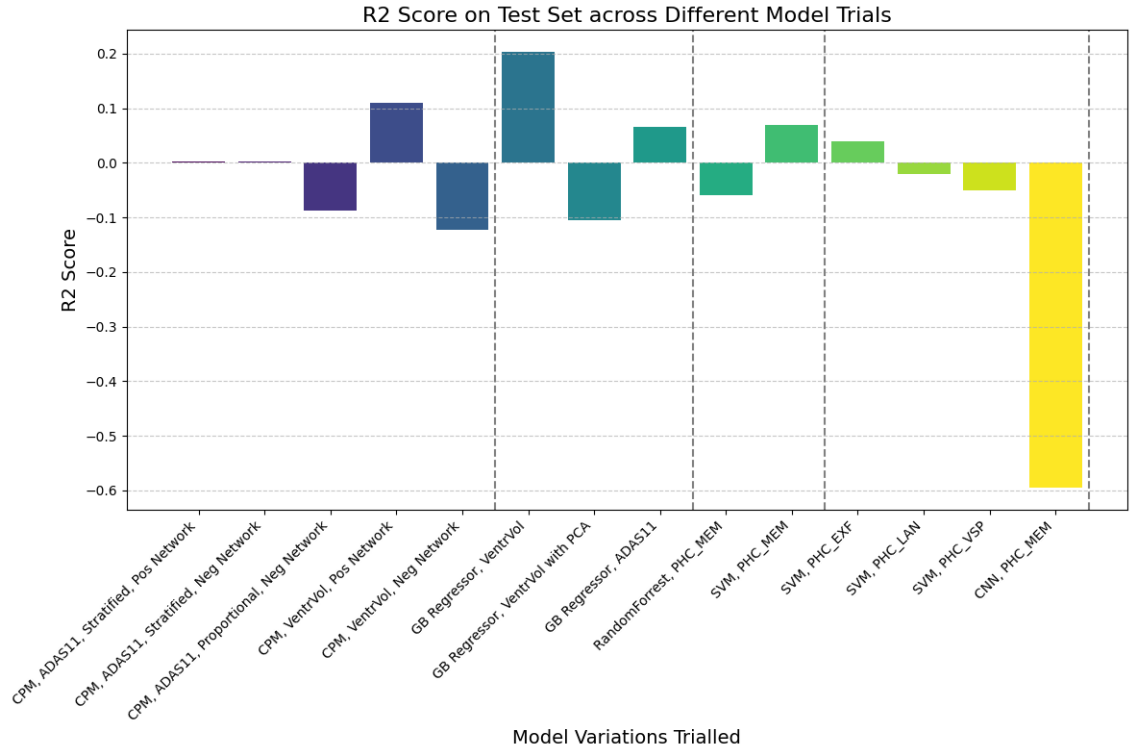

Figure 7.5: **Comparing the R^2 scores across different variations of predicted models trialled over the course of the project, grouped by model type**. *Key*: 'Stratified' - Stratified Sampling Method, 'Proportional' - Proportional Sampling Method, 'Pos Network' - Positive Tail Network Model from the CPM, 'Neg Network' - Negative Tail Network Model from the CPM, 'VentrVol' - Ventricular Volume used as the predictive AD Characteristic

target label.

We experiment with both of the above dimensionality reduction techniques to see which improves model performance the most.

7.3 Predictive Modelling Methods

Over the course of developing a suitable predictor algorithm, we trialled many regression and modelling techniques with different target categories. For each case, we evaluated performance using the R^2 metric. Figure 7.5 illustrates how well the different models performed (using different target categories, as shown) on their respective test sets. We can see that the R^2 score remains relatively low across all models, with the GB Regressor performing the best and CNN model the worst. In this section, we will outline the details of each modelling method trialled.

7.3.1 Model Evaluation Method

To evaluate how well a regression model fits the data, we used the R^2 score. This is a measure of how well a model predicts an outcome of a dependant variable or target. Values for R^2 range from 0 to 1: $R^2 = 0$ means that the model predicts none of the relationships between features and targets, whereas $R^2 = 1$ means that the model predicts all the relationships. It is calculated using the formula below:

$$R^2 = \frac{1 - RSS}{TSS}$$

Here, RSS is the sum of the squares of residuals (the difference between the observed target value and the predicted value) and TSS is the total sum of squares.

It should be noted that while $0 \leq R^2 \leq 1$, the calculations of R^2 for the following models have been done using `scikit-learn`'s `r2_score()` function from `sklearn.metrics` where R^2 can have

a negative value [40]. In the general case (for a non-constant empirical target value) a constant model that always predicts the average value for the target value while disregarding the input features will receive a R^2 score of 0. When this function returns a negative R^2 , it means that the model can be arbitrarily worse than this. In essence, this means that the models in Figure 7.5 that result in negative R^2 scores can effectively be disregarded.

7.3.2 Random Forest (RF) Regressor

Target Categories Trialled: *PHC Composite, ADAS-Cog 11*
When training the RF Regressor model, we split the used dataset into a training set (80%) and a test set (20%) and used `scikit-learn`'s Randomised Search method with 5-fold cross validation to find the optimal hyperparameters that achieves better results. We then tuned this further by conducting a grid search (`scikit-learn`'s `GridSearchCV`) across the best hyperparameters, again using 5-fold cross validation. The hyperparameters we considered were the number of features to consider when looking for the best split, the number of trees in the forest, the minimum number of samples required to be at a leaf node and the minimum number of samples required to split an internal node. 5-fold cross validation was employed during tuning to reduce the variance and bias in performance metrics for each trialled hyperparameter set, since we have a relative large training set. For validation, we found the R^2 score for the optimal hyperparameters found during tuning on the test set.

A Random Forest is a beneficial model to trial as it is comparatively less sensitive to over fitting and can be used to capture the relatively complex relationships between the FC features. However, when used with the PHC Composite or ADAS-Cog 11 as target categories, the RF Regressor performed poorly. The best result came from using the PHC_MEM composite as a target and 464 trees in the random forest, resulting in a score of $R^2 = -0.06$ on the test set. An interesting result was that for all target composites from the ADAS PHC Composite, the RF Regressor returned a positive R^2 score during training, and negative score on the test set, suggesting it was overfitting to the training data and was not able to generalise well to unseen data.

7.3.3 Support Vector Machine (SVM) Regressor

Target Categories Trialled: *PHC Composite*
Next, we tried fitting an SVM Regressor to the PHC Composite scores, and found an improved performance from the RF Regressor model. We followed a similar hyperparameter tuning method to the previous model, by first conducting a Randomised Search using 5-fold cross validation, and then verifying the results with a Grid Search.

With this model, the most interesting parameter to tune was the kernel type. SVMs using linear kernels directly correspond to linear regression models, whereas polynomial or Gaussian kernels perform better at handling complex or non-linear relationships in the data. In our case, we found that the Radial Basis Function (RBF) kernel (also known as the Gaussian kernel), performed the best due to its robustness against the curse of dimensionality. This was advantageous since our feature set contains high-dimensional data, although it could have led to overfitting (accounted for in the poor performances when PHC_LAN or PHC_VSP was used as the target value). As seen in Figure 7.6, the SVM model that uses the PHC_MEM composite target as its target performs the best. This could indicate that out of the different cognitive composites, functional connectivity correlates the most with memory abilities.

7.3.4 Convolutional Neural Networks (CNNs)

Target Categories Trialled: *PHC Composite*
To take advantage of the matrix structure of the features, we tried utilising a deep neural network for regression. We defined a convolutional neural network (CNN) with an architecture consisting of four hidden layers and one dense output layer for calculating the regressor result. The CNN was trained using mean-squared error (MSE) as its loss function, and evaluated using R^2. The difficulty here was to try to find a suitable architecture that could overcome overfitting. As in the RF Regressor, the CNN gave promising results on the training set, but a poor performance on the test set, suggesting overfitting and its inability to generalise to novel data. Although we tried different architectures and adding dropout layers to reduce overfitting, our CNN models performed

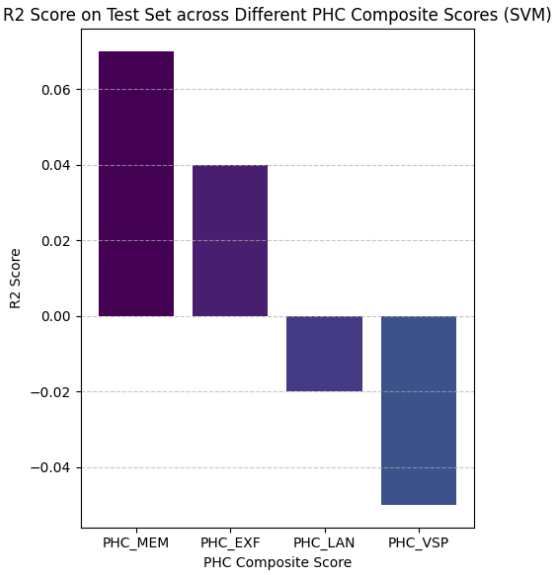

Figure 7.6: Different R^2 scores for the different PHC Composites, using a SVM Regressor

significantly worse than the other regressor models. The best case is shown in Figure 7.5, where the score on the test set was only $R^2 = -0.5954$.

7.3.5 Connectome-based Predictive Model (CPM)

Target Categories Trialled: *ADAS-Cog 11, Ventricular Volume*

We next took inspiration from recent work by Lin *et al.* in predicting ADAS-Cog 11 Scores from functional connectivity data, as described in Section 3.3.1. To overcome the curse of dimensionality, the goal here is to calculate a single summary statistic per FC matrix in the feature set, and use this in place of the features.

We closely followed the method outlined in Section 3.3.1, but instead used the Spearman's rank coefficient (ρ) to calculate the correlation between edges in the FC matrix and the target label, as both ADAS-Cog 11 and Ventricular Volume distributions are negatively skewed rather than normally distributed. Spearman's rank coefficient is a non-parametric measure that does not assume a specific distribution for the data, whereas the Pearson correlation coefficient would assume that the data is normally distributed. After thresholding below a defined p-value, the significant edges could be separated into a positive tail (edges associated with higher ADAS11 scores) and a negative tail (associated with lower ADAS11 scores), resulting in both a positive and negative network model (linear regression models). To test each model, we employed leave-one-out-cross-validation (LOOCV). The choice of significance threshold (p-value), was a hyperparameter to be tuned, although it is considered best to keep it to $p = 0.05$ to avoid the inclusion of weakly correlated edges.

An issue encountered for ADAS-Cog 11 Scores was the lack of edges in the FC matrices that significantly correlated with either target label, even when the significance threshold was increased from $p < 0.05$ to $p < 0.1$. This could be an indication of a lack of correlation between functional connectivity and this particular composite cognitive score. Although using similar data as the literature, we struggled to achieve good performance when using the ADAS-Cog 11 Scores as target labels, as shown in Figure 7.5. We tried utilising the different sampling methods previously mentioned, and although using a stratified sample over a proportional sample increased the R^2 score by 0.0901, the result of stratified sampling ($R^2 = 0.0020$) still shows a poor fit to the data. The only difference in methods between this project and the original authors is that they used a more granular parcellation (Shen 268 functional brain atlas). The added granularity in functional edges could increase the chance of edges in an FC matrix being significantly correlated to the ADAS-Cog 11 Scores, given us a wider variety of summary statistics to train a model on. Furthermore, the choice of fitting a linear model (following Shen *et al.*) to the positive or negative set of statistics could be a limitation, as the relationship between the summary statistics and the ADAS-Cog 11 Scores in our case may not be linear.

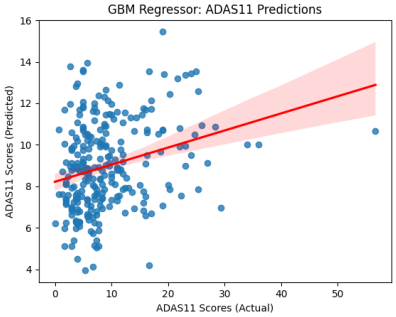

(a) ADAS-Cog 11 GB Regressor

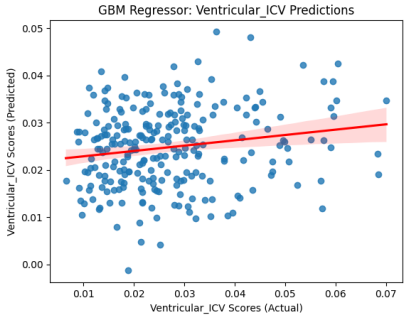

(b) Ventricular Volume GB Regressor, with PCA

Figure 7.7: **Performance of GB Regressor for different target labels and feature sets:** (a) GB Regressor with ADAS-Cog 11 as the target label, (b) GB Regressor for predicting Ventricular Volume, with the dimensions of FC data (features) reduced through PCA

Parameter	Value	Parameter	Value
`n_estimators`	500	`subsample`	0.81
`learning_rate`	0.05	`max_features`	`sqrt`
`max_depth`	2	`min_samples_split`	5

Table 7.1: Tuned parameter values for best performing GB Regressor model ($R^2 = 0.2035$). Note, `max_features`: `sqrt` means that `max_features` = `sqrt`(total number of features)

7.3.6 Gradient Boosting (GB) Regressor

Target Categories Trialled: *ADAS-Cog 11, Ventricular Volume*

Finally, inspired by the methods of the winning team from the TADPOLE challenge for forecasting features of ADNI patient data [39], we attempted using a Gradient Boosting Regressor model to predict the same target categories as the original challenge. The performance of the model was better when using ventricular volume as the target category than when using ADAS-Cog 11, and the GB Regressor is the best performing regressor model on the latter target category (as seen in Figure 7.5).

Training a GB Regressor on ventricular volume gave us the best results seen so far across all models trialled. GB Regressors combine a series of learners (such as decision trees) to create an overall predictive model, but boost accuracy by sequentially correcting errors made by previous models in training. Additionally, the GB Regressor incorporates regularisation techniques such as subsampling that can help prevent over fitting, which was the main issue seen with the other models. We improve the robustness of the model by using Huber loss to train the regressor model, which is a smoother loss that combines MAE (mean absolute error) and MSE, and is less sensitive to outliers in the data. Overall, this improves the ability of the model to generalise to unseen data, as we see by its R^2 performance in Figure 7.5.

We also experimented with dimensionality reduction through PCA, reducing the dimensions of the features from 4950 to 250, although this negatively impacted the performance of the model. The best GB Regressor model was found through hyperparameter tuning the following model parameters: number of boosting stages to perform (`n_estimators`); learning rate to shrink the contribution of each learner (`learning_rate`); maximum depth of each learner or tree (`max_depth`); the minimum samples required to split an internal node (`min_samples_split`); the proportion of samples used to fit an individual learner (`subsample`) and the number of features to consider when looking for the best split (`max_features`). Once again, hyperparameter tuning involved using a Randomised Search and 5-fold cross validation, and then performing a Grid Search (again with 5-fold cross validation) in a smaller range to validate these results. The best hyperparameters across 150 trials can be seen in Table 7.1.

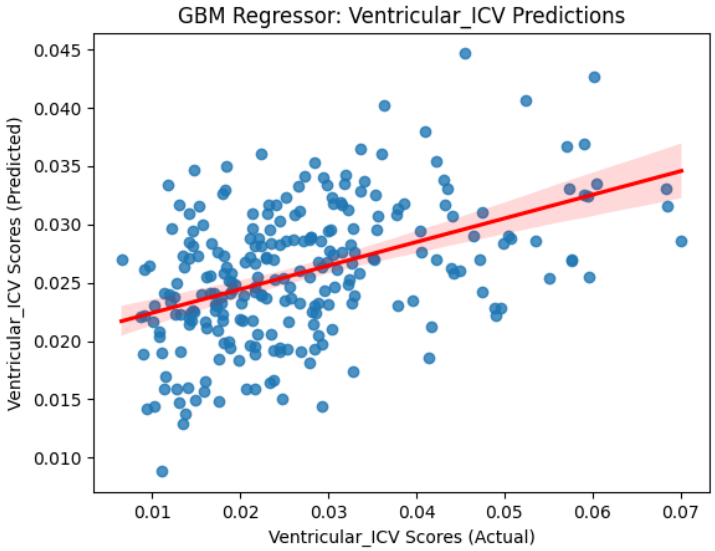

Figure 7.8: **Performance of Final AD Characteristic Predictor Model** ($R^2 = 0.2035$): Predicted ventricular volume scores for each corresponding empirical target value

7.4 Results and Evaluation: Final Predictor Model

After investigating various modelling methods, the most optimal predictor model was found to be the GB Regressor model for predicting ventricular volume. When trained using the hyperparameter values in Table 7.1, we get the predictive performance illustrated in Figure 7.8. This model has a score of $R^2 = 0.2035$, and a Spearman correlation of $\rho = 0.4553$ between the empirical ventricular volume score and the predicted values. While this is not a high score for R^2, we can see in Figure 7.5 that it is significantly better than the other models trialled.

Across all investigations, the greatest challenge has been overfitting to the training data. This issue is noted in current literature on predicting cognitive scores from functional connectivity, for example Shen *et al.* comment in their paper on CPM methods that regression models have a tendency to overfit to input data and as a result struggles to generalise to unseen data [36]. The main advantage of the GB Regressor is using subsampling in the sequential aggregation of individual learners: if each learner only learns from a subset of the training feature set, we can reduce the chance of the combined model overfitting to the overall dataset.

The difficulty of predicting cognitive composites has also been commented on in current literature [39]. Other contemporary studies report a similar low predictive performance: for example, Shen *et al.*'s CPM model achieved a similar Spearman's rank coefficient of $\rho = 0.49$ in their best performing model (the positive network model). As mentioned, our 'predictor' model primarily indicates the correlation between two variables, rather than assuming causation between the independent (FC feature set) and dependant (clinical measure targets) variables. It is not a surprise, then, that the correlation between functional connectivity is found to be stronger with ventricular volume (an objective biomarker) than composite cognitive scores (more subjective and influenced by external factors). As ventricular volume is associated with the neuropathology of the brain and inherently more biological, there are more likely to be underlying similarities between the progression of ventricular volume and functional connectivity (also a biological phenomenon), as opposed to the behavioural traits measured by cognitive studies.

7.5 Using the Predictor Model in the Whole-Brain Modelling Pipeline

In this final section, we look at the results from combining our Aβ-Tau model from previous chapters and the clinical measure predictor algorithm developed in this chapter. We connect the different models in our pipeline: using the Aβ-Tau model to simulate FC matrices, which

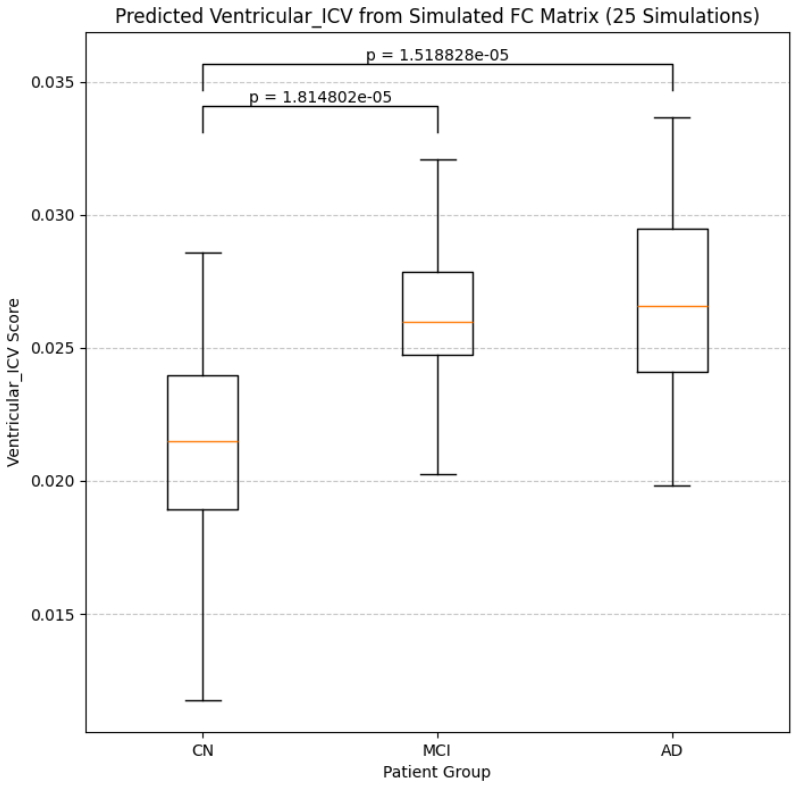

Figure 7.9: Distribution of predicted ventricular volume scores across 25 simulations, using the best fit model per patient group. The FDR-corrected p-values for the significantly different patient groups are shown on the plot. It is important to note that the predicted ventricular volume scores for the CN and AD groups are statistically different

are used as the input feature set to the predictor algorithm. We explore the three key results of this pipeline: predicting ventricular volume using model simulations from each patient group; performing structural and dynamic parameter changes in the Aβ-Tau model to assess their impact on predicted ventricular volume scores and finally varying a singular dynamic parameter in the Aβ-Tau model and assessing the impact on predicted ventricular volume.

7.5.1 Predicting Ventricular Volume From Patient Group Simulations

In our first experiment, we simulate 25 functional connectivity matrices using the best Aβ-Tau model for each patient group (CN, MCI and AD). For each simulated FC, we predict the associated ventricular volume (`Ventricular_ICV` Score) using the GB Regressor model. The distribution of results for each patient group from this experiment can be found in Figure 7.9.

We can see here that the median predicted ventricular volume score increases between patient groups, showing that even with its relatively low R^2 score, the regressor model can retain the underlying trend from the patient groups and empirical ventricular volume scores. We also conducted two-sample t-tests between each pair of patient groups, and the FDR-corrected p-values for these can be found in Table A.2. Again, the p-values were corrected for multiple comparisons by using the Benjamin-Hochberg procedure to control the False Discovery Rate (FDR).

We can see from these p-values that the CN predicted ventricular volume scores were significantly different from both MCI and AD groups, but the MCI and AD predicted scores are not significantly different from each other. This is an expected result, as we expect the healthy controls to be significantly different from either patient group that contains evidence of cognitive decline. In contrast, the MCI patient group is a spectrum of patients from early-stage MCI to late-stage

MCI - there is the possibility of overlap in the condition of patients from the latter category and the AD patient group. These patients could likely share similar ventricular volume scores.

7.5.2 Impact of Structural and Dynamic Changes on Predicted Ventricular Volume

The remaining experiments focus on assessing the impact of varying structural and dynamic parameters of the Aβ-Tau model on the predicted ventricular volume score.

Firstly, we take the fitted SC matrices and the learnt parameter values from both the CN patient model and AD patient model, and train a new model using the permutations of these. Examples of the resulting FC matrix simulations are illustrated in Figures **??-??**. This introduces four models where both the structural and dynamic parameters were varied between CN and AD patient groups. We again simulate 25 FC matrices per model, and use the predictor algorithm to get a ventricular volume measure per simulation. Again, we conducted two-sample t-tests between each pair of models, and the FDR-corrected p-values for these can be found in Table A.3. The distribution of these results alongside the significant p-values are shown in Figure 7.10e.

From these results, it is evident that structural parameter changes have a larger impact on ventricular volume predictions than changing dynamic parameters. For example, when using the CN parameter set, the box plots show that changing the fitted SC matrix from that of the CN model to that of the AD model boosts the median ventricular volume score by around 13%. In contrast, it is interesting to note that the only non-significant p-values arise when the two models being compared use the same fitted SC matrix, but the parameter set used is varied. For example, when using the fitted CN SC matrix, changing the dynamic parameters from that of the CN model to that of the AD model does not cause a significant change to the predicted ventricular volume scores. Similar results are seen between both models using the fitted AD SC matrix.

The final experiment corroborates this result. Here, we use the best CN model and parameter set but independently vary the value of $s_E^{A\beta}$ to assess its impact on ventricular volume. We choose 200 random values for $s_E^{A\beta}$ between the ranges observed across patient groups (-8.0 to 9.0) in Figure 6.7; for each value of $s_E^{A\beta}$, we simulate 5 FC matrices, predict a ventricular volume score per simulation, and average across predicted scores. In each case, all other parameters of the CN model are kept the same. The same experiment is repeated separately for the parameter $s_I^{A\beta}$ (in the range -13.3 to -3.6). Additionally, both experiments were repeated using the fitted SC matrix from the AD model instead. The results of these can be seen in Figure 7.11.

Although some plots (such as Figures 7.11b and 7.11c) show evidence of very weak positive correlation, the p-values for the Pearson correlation coefficient between the dynamic parameter varied and predicted ventricular volume in each of these plots indicate that none of the results are significant. These values are reported in Table A.4. This result is consistent with the previous experiment: solely changing dynamic parameters is not enough to make an impact on the predicted ventricular volume.

However, it could be the case that there is a more complex or non-linear relationship between the dynamic parameter varied and the predicted ventricular that is not captured by Pearson's correlation coefficient. Alternatively, it could be the case that some dynamical parameters are tightly coupled, and both of these need to be varied to see a visible impact on the predicted ventricular volume score. Further experiments are needed in the future to explore this relationship in more detail.

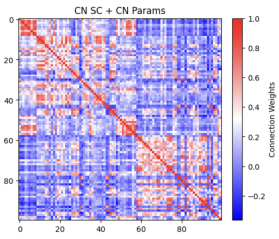

(a) CN SC matrix and CN paramaters

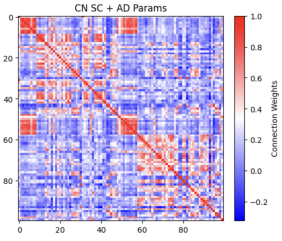

(b) CN SC matrix and AD paramaters

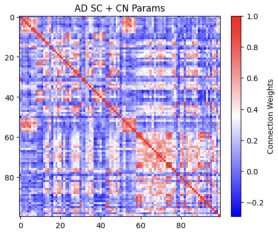

(c) AD SC matrix and CN paramaters

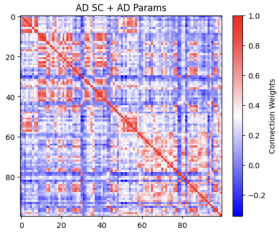

(d) AD SC matrix and AD paramaters

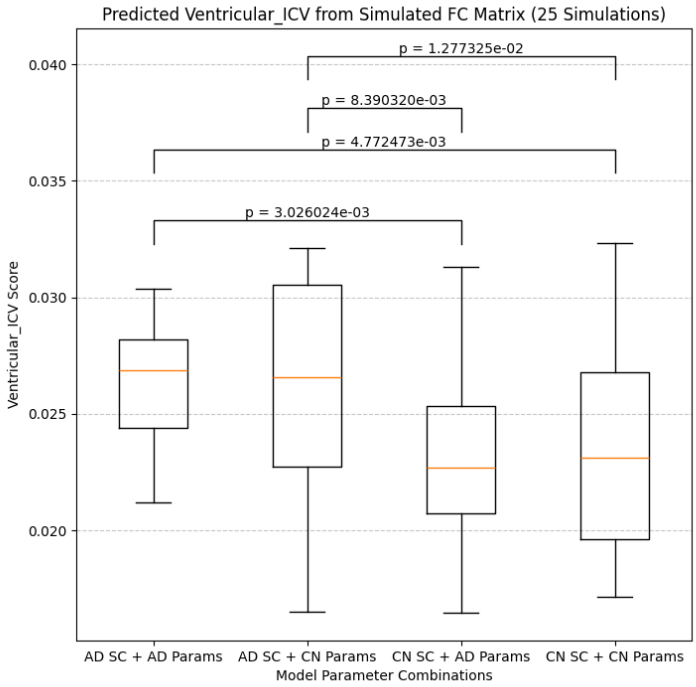

(e) Distribution of Predicted Ventricular Volume Scores across each FC Permutation

Figure 7.10: Assessing the impact of varying structural and dynamic changes on the Aβ-Tau DMF Model: plots (a)-(d) illustrate the simulated FC matrices found from initialising an Aβ-Tau DMF Model with a permutation of AD/CN patient fitted SC matrices and the learnt free parameter values from the AD/CN model; in (e) we show the distribution of predicted ventricular volume scores for each model parameter combination, with the FDR-corrected p-values shown for all pairs of combinations that yield significantly different results.

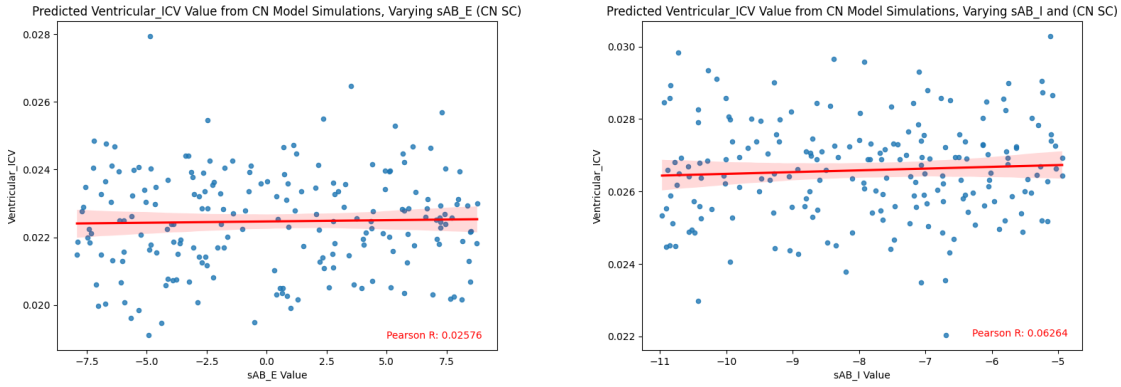

(a) (b)

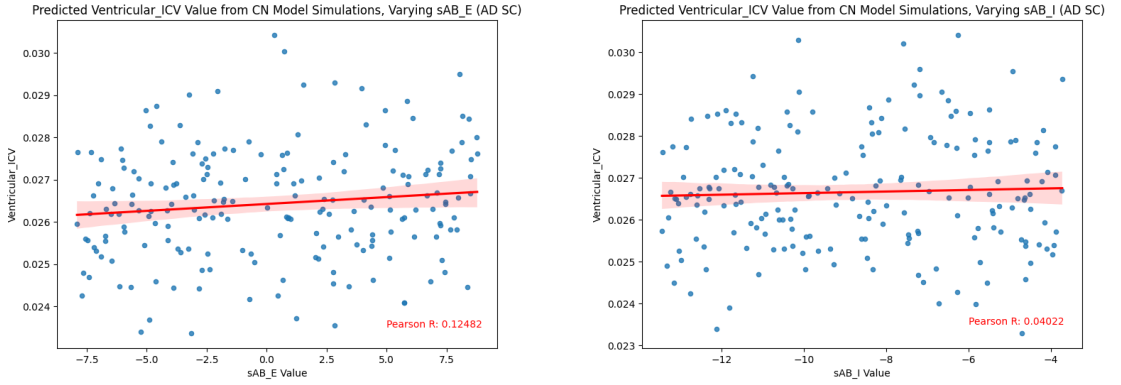

(c) (d)

Figure 7.11: **Varying singular dynamic parameter experiment:** Results from experiments used to assess the changes of varying one dynamic parameter. Here, we used the CN model with Aβ-Tau values and learnt parameters, varying one parameter and the fitted SC matrix to assess their impact on the predicted ventricular volume score of the simulated FC matrices. (a) and (c) show the effect of only varying parameter $s_E^{A\beta}$, while (b) and (d) show the effect of only varying parameter $s_I^{A\beta}$; the CN fitted SC matrix is used in (a) and (b), whereas the AD fitted SC matrix is used in (c) and (d).

Chapter 8

Discussion

In this project, we have successfully adapted the whole-brain modelling pipeline to include structural and dynamic changes and have shown how we can use a predictor model to quantify the effect of changing these parameters on observable clinical measures.

Specifically, we have successfully modelled Aβ-Tau degradation in the DMF whole-brain model and further shown that adding the Aβ-Tau heterogeneity to the DMF model improves the fit between simulated and empirically healthy patient functional data by 39.96% on average. Our results are comparable to other studies employing connectome-based modelling to simulate FC data [12] that do not model dynamic parameter changes. This shows that the amendments we have made to the existing whole-brain modelling pipeline have not negatively impacted performance, but can be instead used to gain a deeper understanding of the underlying mechanisms behind AD. While there is no clear trend in how the bias and scaling terms associated with Aβ and Tau deposits change between patient groups, our results show that for most of the parameters, their value significantly differs between patient groups. This can be useful in illustrating the evolution of the Aβ and Tau neuritic plaques and identifying the distinct stages of the diseases for other parameters. Furthermore, an understanding of this evolution can be beneficial for early detection of cognitive decline: parameters such as $b_I^{A\beta}$ and $s_I^{A\beta}$ change significantly between the clinically normal and MCI patient groups and could serve as biomarkers for detecting the initial stages of disease pathogenesis.

Adding the gradient-boosting regressor model to the pipeline allows us to gain a holistic understanding of the effect of AD on patients by correlating internal neural changes with external clinical measures. Even with its relatively low R^2 score, the regressor model can reproduce the correlation between patient groups and empirical ventricular volume scores within the FC simulations from each patient group and the predicted ventricular volume scores. Further experiments involving varying the fitted SC matrix and learnt parameter set used in an Aβ-Tau model (between AD and CN patient groups) have also been shown to affect the predicted ventricular volume score. Notably, the SC matrix changes caused a greater change in predicted ventricular volume than dynamical parameter set changes. For example, when using the learnt parameters from the CN model, changing from the CN model's fitted SC matrix to the AD model's fitted SC matrix boosted the median predicted ventricular score by 13%. In contrast, varying a single dynamic parameter (bias or scaling term) from the learnt CN model has no significant impact on the predicted ventricular volume score of the model's simulations. It could be the case that the structural changes in the fitted SC matrix data are more highly correlated with the ventricular volume, and so the predictor model performs better on nuanced changes in structural parameters than dynamical parameters.

This project contributes to the suite of novel machine-learning (ML) techniques used to facilitate neuroscience research. One of the biggest challenges of whole-brain modelling remains that of fitting the large parameter space of biophysical models, which current ML research is hoping to address. For example, in this project, we have seen the benefits of using deep learning-based parameter optimisation in `whobpyt`'s connectome-based neural mass model to improve the model fit on empirical data by additionally fitting the many connectivity weights of the input SC matrix. Additionally, ML techniques have been used to gain a more comprehensive understanding of the whole-brain model, for example using VAEs for latent space encoding [25] to visualise complex relationships between neural mechanisms, or using graph-based neural networks to investigate the topological properties in functional data [41]. We hope that the whole-brain modelling pipeline

developed through this project is a meaningful addition to this corpus of research, and can be improved in the future by utilising other ML techniques - for example using the latter graph-based approach to improve the clinical measure predictor algorithm.

From the results of the overall working pipeline, we can see that combining functional connectivity data with influences of other biomarkers could lead to a more meaningful model of Alzheimer's Disease. We believe that this pipeline could assist with the treatment and diagnosis of Alzheimer's Disease, from helping with early detection to tailored treatment strategies. Above all, the forecasting power of being able to switch parameters between the different models means that ventricular volume changes could be predicted before they are visible on traditional imaging scans, leading to an earlier AD diagnosis. These models can act as prognostic indicators to inform patients about expected disease trajectory, helping in treatment planning and management. For treatment methods, understanding the relationships between Aβ and Tau neuritic plaques and brain structure can lead to tailored intervention and treatment strategies based on a patient's specific biomarkers and disease progression patterns. This project introduces a link between dynamic changes (biomarkers), functional connectivity and clinical measures, meaning the pipeline can serve as an objective measure in clinical trials to assess the efficacy of new treatments on the former biomarkers by evaluating the effect on the latter clinical measures.

8.1 Limitations and Future Work

There are still many areas for potential future work to improve the results of this project. Within the overall whole-brain modelling pipeline, two main weaknesses could be addressed: firstly, the performance of the clinical measure predictor, and secondly, the underlying stability issues of the Aβ-Tau model.

Firstly, a notable limitation of the performance of the pipeline is the performance of the ventricular volume predictor algorithm ($R^2 = 0.204$). Although this is the best model out of the clinical measure predictors trialled in this project, we can see that the poor $R2$ score can percolate through the pipeline, affecting the visibility of trends and results in later experiments, such as investigating the changes in ventricular volume when varying one dynamic parameter. In our experiments, we found no significant correlation in this latter experiment, but perhaps a better-performing predictor model could have distinguished more between smaller variations in functional connectivity. Predicting clinical measures from functional connectivity is a challenge in contemporary computational neuroscience, as even the best-performing models in current literature had poor correlation scores between the two. This is an area where ML can be applied in the future, such as using graph-based learning approaches to improve predictive performance.

Then, further work is needed to explore the stability of the firing rates in the developed Aβ-Tau model that leads to undefined gradients during training. Although we have avoided this problem in this project by using early stopping and a more numerically stable computation method, we have not addressed the underlying cause of the undefined gradients. A plausible explanation is that the neuronal response functions (firing rates) of each neural population are not stabilised enough during training. This could be a result of incorrect values for certain synaptic gain parameters in the DMF model that were not learnt during training. A side effect of this could also be the high value learnt for the long-range coupling parameter, g, which we have found to be much higher than other studies. In future, it would be good to utilise the deep-learning capabilities of the `whobpyt` library and fit more of the synaptic gain parameters in the DMF Model.

Finally, in the results of the overall pipeline, we have seen that structural changes have much more of an impact on functional connectivity than dynamic changes (evaluated by their impact on predicted clinical measures). This project has mainly investigated the modelling of dynamic changes, but in future, it would be beneficial to investigate the effect of structural changes further. Structural connectivity weight estimation for whole-brain models remains an under-developed area of study, but by including this in the whole-brain modelling pipeline, in the future, it would be possible to investigate degenerative mechanisms such as white matter pathology or loss of small-world properties in the brain network. Additionally, adapting the pipeline to fit at a subject level as opposed to a group level could help to improve the fit of structural connectivity in the overall model.

8.2 Conclusions

Overall, this project demonstrates the successful adaptation of the whole-brain modelling pipeline in incorporating structural and dynamic changes, and the development of a predictor model to quantify the impact of these changes on observable clinical measures. We have shown that integrating Aβ-Tau degradation into the DMF model can improve the fit between simulated and empirical functional data, showing how we can combine connectomics and biomarkers to gain a better understanding of AD pathogenesis. Our results suggest that structural changes in the fitted SC matrix have a greater influence on predicted ventricular volume than dynamic parameter changes, indicating the importance of including structural biomarkers in AD modelling in the future. Through the development of this new pipeline, we have managed to achieve the four goals outlined in the introduction of this project.

The project contributes to the advancement of applying ML techniques in computational neuroscience to enhance model accuracy while leaving scope for further improvements to use ML to understand relationships between connectomics and clinical measures. Ultimately, we hope our pipeline offers a promising tool in the future for early detection and evaluating treatments for the disease, fostering a deeper understanding of AD progression and management.

Appendix A

Appendix

A.1 Aβ-Tau DMF Model: Additional Training Curves

In Figure A.1, we present the additional training curves for the best Aβ-Tau DMF models, per patient group (CN, MCI and AD). Plots A.1a, A.1c and A.1e show the evolution of the model training loss (above) and the bias and scaling parameters being learnt (below) across the best training run. Plots A.1b, A.1d and A.1f display the simulated BOLD time series (below) for the empirical BOLD time series (above) for 10 randomly chosen regions in the empirical BOLD parcellation.

A.2 Aβ-Tau DMF Model: Parameter T-tests

FDR-correct p-values from the two-sample t-tests conducted between patient groups CN, MCI and AD for each bias and scaling parameter in Aβ-Tau model:

Parameter	CN vs MCI	CN vs AD	MCI vs ADI
$b_E^{A\beta}$	8.328088e-05	1.623497e-16	1.769871e-14
$s_E^{A\beta}$	1.294243e-11	1.404395e-12	1.133556e-16
$b_I^{A\beta}$	1.082182e-11	5.129410e-15	5.129410e-15
$s_I^{A\beta}$	2.282244e-04	7.430379e-16	1.146872e-05
b_E^t	**8.786570e-02**	1.598887e-14	8.208451e-15
s_E^t	3.678277e-03	3.554393e-17	5.348076e-18

Table A.1: P-values from FDR-corrected t-test between each pair of patient groups per parameter. The only non-significant p-value is shown in bold.

A.3 T-tests for the Overall Working Pipeline

A.3.1 Predicting Ventricular Volume From Patient Group Simulations

FDR-corrected p-values from the two-sample t-tests conducted between patient groups CN, MCI and AD:

	CN	MCI	AD
CN	-	0.00002	0.00002
MCI	0.00002	-	0.54462
AD	0.00002	0.54462	-

Table A.2: P-values from FDR-corrected t-test between each pair of patient groups. Significance threshold: $p < 0.05$

A.3.2 Impact of Structural and Dynamic Changes on Predicted Ventricular Volume

We perform two-sample t-tests between the four models resulting from permutations of the fitted SC matrices and the learnt parameter values from the CN and AD patient models. The FDR-corrected p-values are given below:

	AD SC + AD Params	**AD SC + CN Params**	**CN SC + AD Params**	**CN SC + CN Params**
AD SC + AD Params	-	0.974824	0.003026	0.004772
AD SC + CN Params	0.974824	-	0.008390	0.012773
CN SC + AD Params	0.003026	0.008390	-	0.974824
CN SC + CN Params	0.004772	0.012773	0.974824	-

Table A.3: P-values from FDR-corrected t-test between each pair of model combinations. Significance threshold: $p < 0.05$)

Additionally, we evaluate the correlation between the dynamic variable ($s_E^{A\beta}$ or $s_I^{A\beta}$) varied in a CN model and the predicted ventricular volume score. The Pearson correlation coefficient and corresponding p-value are reported below:

Parameter	**Fitted SC Matrix**	**Pearson R**	**p-value**	**Significant Result**
$s_E^{A\beta}$	AD	0.1248	0.0782	**No**
$s_E^{A\beta}$	CN	0.0258	0.7173	**No**
$s_I^{A\beta}$	AD	0.0402	0.5718	**No**
$s_I^{A\beta}$	CN	0.0626	0.3782	**No**

Table A.4: Correlation between dynamic parameter changes and the predicted ventricular volume, with the associated p-value. Significance threshold: $p < 0.05$

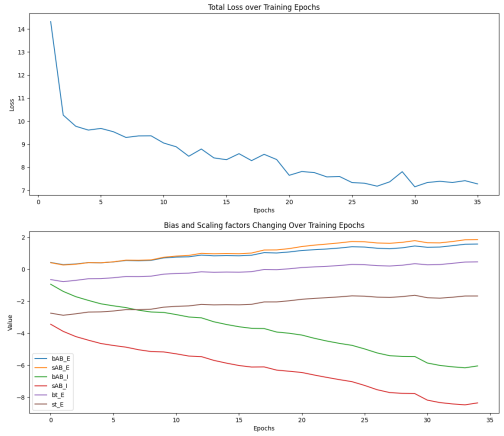

(a) CN: Loss and Model Fitting Curves

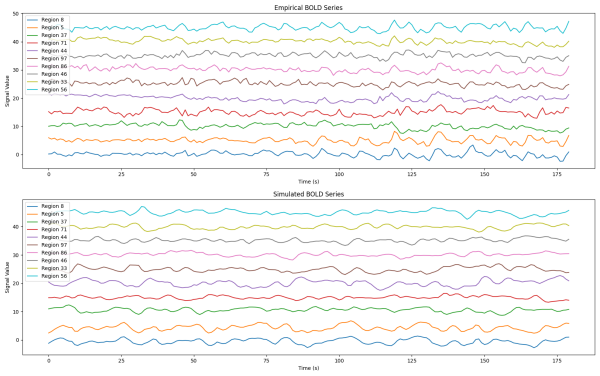

(b) CN: Empirical vs Simulated BOLD

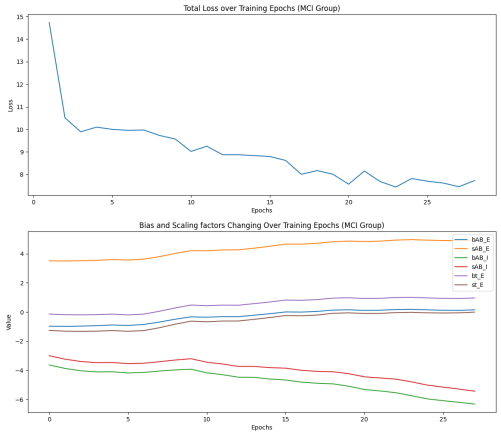

(c) MCI: Loss and Model Fitting Curves

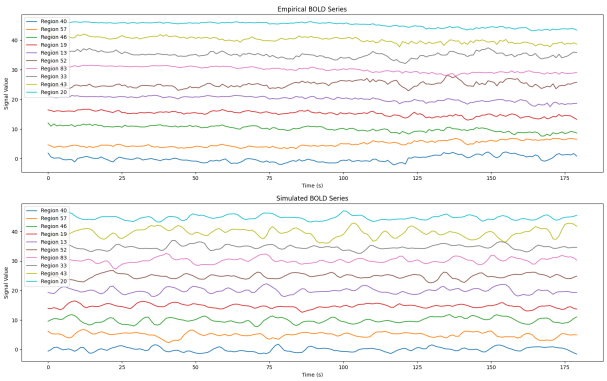

(d) MCI: Empirical vs Simulated BOLD

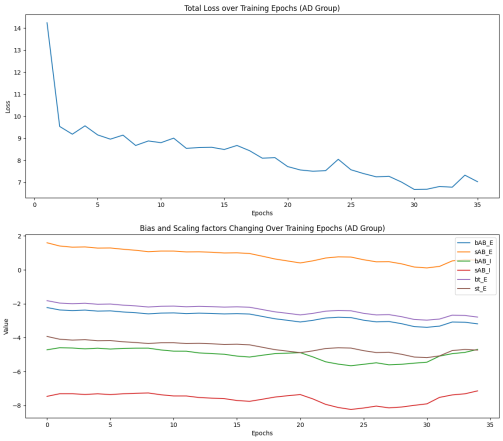

(e) AD: Loss and Model Fitting Curves

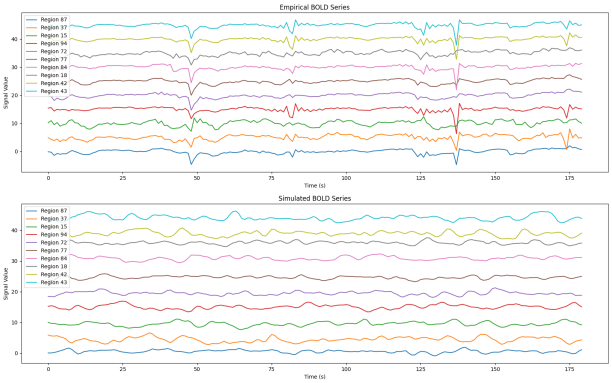

(f) MCI: Empirical vs Simulated BOLD

Figure A.1: **Additional results from the best performing Aβ-Tau model per patient group:** (a), (c) and (d) show the training loss (above) and fitting of bias and scaling parameters (below) across a training run; (b), (d) and (f) illustrated the simulated BOLD time series from each model against the empirical BOLD data, for 10 randomly chosen regions in the empirical BOLD parcellation.

Bibliography

[1] Feigin VL, Vos T, Nichols E, Owolabi MO, Carroll WM, Dichgans M, et al. The global burden of neurological disorders: translating evidence into policy. The Lancet Neurology. 2020 Mar;19(3):255-65. Available from: https://linkinghub.elsevier.com/retrieve/pii/S1474442219304119.

[2] Weiner MW, Aisen PS, Jack CR, Jagust WJ, Trojanowski JQ, Shaw L, et al. The Alzheimer's Disease Neuroimaging Initiative: Progress report and future plans. Alzheimer's & Dementia. 2010 May;6(3):202. Available from: https://alz-journals.onlinelibrary.wiley.com/doi/10.1016/j.jalz.2010.03.007.

[3] Patow G, Stefanovski L, Ritter P, Deco G, Kobeleva X. Whole-brain modeling of the differential influences of Amyloid-Beta and Tau in Alzheimer's Disease. bioRxiv. 2023.

[4] Supekar K, Menon V, Rubin D, Musen M, Greicius MD. Network Analysis of Intrinsic Functional Brain Connectivity in Alzheimer's Disease. PLoS Computational Biology. 2008 Jun;4(6):e1000100. Available from: https://dx.plos.org/10.1371/journal.pcbi.1000100.

[5] Deco G, Kringelbach M. Great Expectations: Using Whole-Brain Computational Connectomics for Understanding Neuropsychiatric Disorders. Neuron. 2014 Dec;84(5):892-905. Available from: https://linkinghub.elsevier.com/retrieve/pii/S0896627314007351.

[6] Herzog R, Mediano PAM, Rosas FE, Luppi AI, Perl YS, Tagliazucchi E, et al. Neural mass modelling for the masses: Democratising access to whole-brain biophysical modelling with FastDMF. Neuroscience; 2022. Available from: http://biorxiv.org/lookup/doi/10.1101/2022.04.11.487903.

[7] Gatica M, Rosas FE, Mediano PAM, Diez I, Swinnen SP, Orio P, et al. High-order functional interactions in ageing explained via alterations in the connectome in a whole-brain model. Neuroscience; 2021. Available from: http://biorxiv.org/lookup/doi/10.1101/2021.09.15.460435.

[8] Lin Q, Rosenberg MD, Yoo K, Hsu TW, O'Connell TP, Chun MM. Resting-State Functional Connectivity Predicts Cognitive Impairment Related to Alzheimer's Disease. Frontiers in Aging Neuroscience. 2018 Apr;10:94. Available from: http://journal.frontiersin.org/article/10.3389/fnagi.2018.00094/full.

[9] Nestor SM, Rupsingh R, Borrie M, Smith M, Accomazzi V, Wells JL, et al. Ventricular enlargement as a possible measure of Alzheimer's disease progression validated using the Alzheimer's disease neuroimaging initiative database. Brain. 2008 Aug;131(9):2443-54. Available from: https://academic.oup.com/brain/article-lookup/doi/10.1093/brain/awn146.

[10] Overview — WhoBPyT 0.0.0 documentation;. Available from: https://griffithslab.github.io/whobpyt/about_whobpyt/overview.html.

[11] Stefanovski L, Meier JM, Pai RK, Triebkorn P, Lett T, Martin L, et al. Bridging Scales in Alzheimer's Disease: Biological Framework for Brain Simulation With The Virtual Brain. Frontiers in Neuroinformatics. 2021 Apr;15:630172. Available from: https://www.frontiersin.org/articles/10.3389/fninf.2021.630172/full.

[12] Griffiths JD, Wang Z, Ather SH, Momi D, Rich S, Diaconescu A, et al. Deep Learning-Based Parameter Estimation for Neurophysiological Models of Neuroimaging Data. Neuroscience; 2022. Available from: http://biorxiv.org/lookup/doi/10.1101/2022.05.19.492664.

[13] Harita S, Momi D, Wang Z, Bastiaens SP, Griffiths JD. The role of inhibition in resting-state fMRI negative correlations; 2024. Available from: http://biorxiv.org/lookup/doi/10.1101/2024.03.01.583030.

[14] Ali H. Computational Whole-Brain Modelling: A Review and Model Comparison. Imperial College London; 2023.

[15] Lawrence RM, Bridgeford EW, Myers PE, Arvapalli GC, Ramachandran SC, Pisner DA, et al. Standardizing human brain parcellations. Scientific Data. 2021 Mar;8(1):78. Available from: https://www.nature.com/articles/s41597-021-00849-3.

[16] Babaeeghazvini P, Rueda-Delgado LM, Gooijers J, Swinnen SP, Daffertshofer A. Brain Structural and Functional Connectivity: A Review of Combined Works of Diffusion Magnetic Resonance Imaging and Electro-Encephalography. Frontiers in Human Neuroscience. 2021 Oct;15:721206. Available from: https://www.frontiersin.org/articles/10.3389/fnhum.2021.721206/full.

[17] Fornito A, Zalesky A, Bullmore ET. Connectivity Matrices and Brain Graphs. In: Fundamentals of Brain Network Analysis. Elsevier; 2016. p. 89-113. Available from: https://linkinghub.elsevier.com/retrieve/pii/B9780124079083000030.

[18] Fornito A, Zalesky A, Breakspear M. Graph analysis of the human connectome: Promise, progress, and pitfalls. NeuroImage. 2013;80:426-44. Mapping the Connectome. Available from: https://www.sciencedirect.com/science/article/pii/S1053811913004345.

[19] Griffa A, Alemán-Gómez Y, Hagmann P. Structural and functional connectome from 70 young healthy adults. Zenodo; 2019. Available from: https://zenodo.org/records/2872624.

[20] Friston KJ, Mechelli A, Turner R, Price CJ. Nonlinear Responses in fMRI: The Balloon Model, Volterra Kernels, and Other Hemodynamics. NeuroImage. 2000;12(4):466-77. Available from: https://www.sciencedirect.com/science/article/pii/S105381190090630X.

[21] Wong KF, Wang XJ. A Recurrent Network Mechanism of Time Integration in Perceptual Decisions. The Journal of Neuroscience. 2006 Jan;26(4):1314-28. Available from: https://www.jneurosci.org/lookup/doi/10.1523/JNEUROSCI.3733-05.2006.

[22] Deco G, Ponce-Alvarez A, Hagmann P, Romani GL, Mantini D, Corbetta M. How Local Excitation-Inhibition Ratio Impacts the Whole Brain Dynamics. Journal of Neuroscience. 2014 Jun;34(23):7886-98. Available from: https://www.jneurosci.org/lookup/doi/10.1523/JNEUROSCI.5068-13.2014.

[23] Demirtaş M, Falcon C, Tucholka A, Gispert JD, Molinuevo JL, Deco G. A whole-brain computational modeling approach to explain the alterations in resting-state functional connectivity during progression of Alzheimer's disease. NeuroImage: Clinical. 2017;16:343-54. Available from: https://linkinghub.elsevier.com/retrieve/pii/S2213158217301961.

[24] Acebron J, Bonilla L, Vicente C, Ritort F, Spigler R. The Kuramoto model: A simple paradigm for synchronization phenomena. Reviews of Modern Physics. 2005 01;77:137.

[25] Sanz Perl Y, Fittipaldi S, Gonzalez Campo C, Moguilner S, Cruzat J, Fraile-Vazquez ME, et al. Model-based whole-brain perturbational landscape of neurodegenerative diseases. eLife. 2023 Mar;12:e83970. Publisher: eLife Sciences Publications, Ltd. Available from: https://doi.org/10.7554/eLife.83970.

[26] Reijmer YD, Leemans A, Caeyenberghs K, Heringa SM, Koek HL, Biessels GJ, et al. Disruption of cerebral networks and cognitive impairment in Alzheimer disease. Neurology. 2013;80(15):1370-7. Available from: https://www.neurology.org/doi/abs/10.1212/WNL.0b013e31828c2ee5.

[27] de Haan W, Mott K, van Straaten ECW, Scheltens P, Stam CJ. Activity Dependent Degeneration Explains Hub Vulnerability in Alzheimer's Disease. PLOS Computational Biology. 2012 08;8(8):1-14. Available from: https://doi.org/10.1371/journal.pcbi.1002582.

[28] Chen G, Zhang HY, Xie C, Chen G, Zhang ZJ, Teng GJ, et al. Modular reorganization of brain resting state networks and its independent validation in Alzheimer's disease patients. Frontiers in Human Neuroscience. 2013;7. Available from: https://www.frontiersin.org/articles/10.3389/fnhum.2013.00456.

[29] ;. Available from: https://adni.loni.usc.edu/study-design/collaborative-studies/worldwide-adni/.

[30] The Preclinical Alzheimer Cognitive Composite. JAMA Neurol;71. Available from: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4439182/.

[31] Kueper JK, Speechley M, Montero-Odasso M. The Alzheimer's Disease Assessment Scale–Cognitive Subscale (ADAS-Cog): Modifications and Responsiveness in Pre-Dementia Populations. A Narrative Review. Journal of Alzheimer's Disease. 2018 Apr;63(2):423-44. Available from: https://www.medra.org/servlet/aliasResolver?alias=iospress&doi=10.3233/JAD-170991.

[32] Mukherjee S, Choi SE, Lee ML, Scollard P, Trittschuh H, Saykin AJ, et al. ADSP Phenotype Harmonization Consortium – Derivation of Cognitive Composite Scores.

[33] Jack CR, Shiung MM, Gunter JL, O'Brien PC, Weigand SD, Knopman DS, et al. Comparison of different MRI brain atrophy rate measures with clinical disease progression in AD. Neurology. 2004 Feb;62(4):591-600. Available from: https://www.neurology.org/doi/10.1212/01.WNL.0000110315.26026.EF.

[34] Ipiña IP, Kehoe PD, Kringelbach M, Laufs H, Ibañez A, Deco G, et al. Modeling regional changes in dynamic stability during sleep and wakefulness. NeuroImage. 2020;215:116833. Available from: https://www.sciencedirect.com/science/article/pii/S1053811920303207.

[35] Dertat A. Applied Deep Learning - Part 3: Autoencoders — towardsdatascience.com;. [Accessed 24-01-2024]. https://towardsdatascience.com/applied-deep-learning-part-3-autoencoders-1c083af4d798.

[36] Shen X, Finn ES, Scheinost D, Rosenberg MD, Chun MM, Papademetris X, et al. Using connectome-based predictive modeling to predict individual behavior from brain connectivity. Nature Protocols. 2017 Mar;12(3):506-18. Available from: https://www.nature.com/articles/nprot.2016.178.

[37] Setsompop K, Kimmlingen R, Eberlein E, Witzel T, Cohen-Adad J, McNab JA, et al. Pushing the limits of in vivo diffusion MRI for the Human Connectome Project. NeuroImage. 2013 Oct;80:220-33. Available from: https://linkinghub.elsevier.com/retrieve/pii/S1053811913005788.

[38] Softplus — PyTorch 2.3 documentation;. Available from: https://pytorch.org/docs/stable/generated/torch.nn.Softplus.html#softplus.

[39] Marinescu RV, Oxtoby NP, Young AL, Bron EE, Toga AW, Weiner MW, et al.. The Alzheimer's Disease Prediction Of Longitudinal Evolution (TADPOLE) Challenge: Results after 1 Year Follow-up. arXiv; 2021. ArXiv:2002.03419 [q-bio, stat]. Available from: http://arxiv.org/abs/2002.03419.

[40] R2 Score - scikit-learn Documentation;. Available from: https://scikit-learn/stable/modules/generated/sklearn.metrics.r2_score.html.

[41] Sunil G, Gowtham S, Bose A, Harish S, Srinivasa G. Graph neural network and machine learning analysis of functional neuroimaging for understanding schizophrenia. BMC Neuroscience. 2024 Jan;25(1):2. Available from: https://bmcneurosci.biomedcentral.com/articles/10.1186/s12868-023-00841-0.