Abou Bekr Belkaid University
Tlemcen Algeria
جامعة أبي بكر بلقايد
تلمسان الجزائر

PEOPLE'S DEMOCRATIC REPUBLIC OF ALGERIA
الجمهورية الجزائرية الديمقراطية الشعبية

MINISTRY OF HIGHER EDUCATION AND SCIENTIFIC RESEARCH
وزارة التعليم العالي و البحث العلمي

FACULTY OF TECHNOLOGY

DEPARTMENT OF BIOMEDICAL ENGINEERING

MASTER THESIS

Submitted in partial fulfillment of the requirements for the Master's degree:

In: Biomedical Engineering
Specialty: Biomedical and Hospital Informatics

By:

Bensalem Zakaria Mohab Eddine

Theme

Individual Brain Identification via Functional Connectivity: From Classical Methods to Graph Attention Networks

Thesis publically defended on June 22, 2026 in Tlemcen before a jury composed of:

Pr. Boukli Hacene Ismail	University of Tlemcen	President
Dr. Kazi Tani Lamia	University of Tlemcen	Examiner
Pr. Feroui Amel	University of Tlemcen	Supervisor
Dr. Saim Meriem	University of Tlemcen	Co-Supervisor
Dr. Chiali Hakim	CHU of Tlemcen	Invited Member

Academic Year 2025/2026

"We can only see a short distance ahead, but we can see plenty there that needs to be done."

— Alan Turing

Acknowledgements

First and foremost, I am deeply grateful to Allah, whose mercy and guidance have carried me through every step of this journey.

This thesis is wholeheartedly dedicated to my beloved parents, brothers, and sister, who have been my constant source of inspiration, giving me strength when I thought of giving up, and who continually provide their unwavering moral, spiritual, emotional, and financial support.

To my friends and colleagues, thank you for making this journey so special. Your support, friendship, and the memories we shared are absolute treasures to me.

I am especially grateful to my supervisors, **Dr. Saim Meriem** and **Pr. Feroui Amel**, for their guidance and support throughout this work. Dr. Saim Meriem proposed the research topic, provided constant technical guidance, and was always available with detailed and constructive feedback. Pr. Feroui Amel enriched this work with her experience in research and her valuable academic insight. I am deeply thankful to both for their patience and trust in my abilities.

I would also like to express my sincere gratitude to all the professors who have taught and guided me throughout these five years, for their knowledge, mentorship, and commitment to our academic formation.

Finally, I am thankful to everyone who has contributed to this work, whether directly or indirectly. I hope that this research will make a meaningful contribution to the field of neuroscience and beyond.

Abstract

Understanding the human brain as an individualised biological system has emerged as a central goal of contemporary neuroscience. This research investigates brain fingerprinting, the phenomenon whereby an individual's functional connectome serves as a stable and unique biological signature. Despite high identification accuracies achieved using linear metrics such as Pearson correlation, current methods often fail to capture complex non-linear dependencies and discard vital non-Euclidean graph topology by flattening the connectome into vectors.

This thesis proposes a systematic, three-stage experimental framework to address these limitations, evaluating models of increasing expressive capacity: a Pearson correlation baseline, classical machine learning classifiers formulated as pairwise verification tasks, and deep Siamese architectures utilising Graph Neural Networks (GNNs) governed by an InfoNCE contrastive loss. Experiments are conducted on resting-state fMRI data from the HCP-Young Adult and HCP-Aging cohorts, with atlas resolution (Schaefer-100 vs. Schaefer-400) varied systematically across both datasets.

Results demonstrate that supervised discriminative approaches significantly outperform unsupervised baselines. The Siamese Graph Attention Network (GAT) is the strongest performer across all conditions, achieving 100% identification accuracy on HCP-Aging at fine-grained resolution. Its learned attention maps provide neurobiological interpretability by naturally prioritising the frontoparietal (Control) network, independently corroborating established theories on inter-individual variability. Model rankings remain consistent across cohorts, with performance differences attributable to demographic and acquisition factors. Finer parcellations consistently benefit graph-based models while penalising linear classifiers, providing empirical guidance for future methodological choices.

By establishing a robust, topologically-aware, and interpretable framework for brain fingerprinting, this work contributes to the advancement of individualised clinical tools and precision medicine.

Keywords: Brain fingerprinting, Functional connectivity, Graph Neural Networks, Siamese architecture, Contrastive learning, Precision medicine.

Résumé

Comprendre le cerveau humain en tant que système biologique individualisé est devenu un objectif central des neurosciences contemporaines. Cette recherche étudie l'empreinte cérébrale, phénomène par lequel le connectome fonctionnel d'un individu constitue une signature biologique stable et unique. Malgré les taux d'identification élevés obtenus avec des métriques linéaires telles que la corrélation de Pearson, les méthodes actuelles peinent à capturer les dépendances non linéaires complexes et négligent la topologie de graphe non euclidienne en aplatissant le connectome en vecteurs.

Cette thèse propose un cadre expérimental systématique en trois étapes pour surmonter ces limitations, évaluant des modèles de capacité expressive croissante : une ligne de base basée sur la corrélation de Pearson, des classificateurs classiques formulés comme des tâches de vérification par paires, et des architectures Siamoises profondes utilisant des réseaux de neurones de graphes (GNN) gouvernés par une fonction de perte contrastive InfoNCE. Les expériences sont conduites sur des données d'IRMf de repos issues des cohortes HCP-Young Adult et HCP-Aging, avec la résolution de l'atlas (Schaefer-100 vs. Schaefer-400) variée systématiquement sur les deux jeux de données.

Les résultats démontrent que les approches discriminatives supervisées surpassent significativement les méthodes non supervisées. Le réseau Siamois à attention de graphe (GAT) est le plus performant dans toutes les conditions, atteignant une précision d'identification de 100% sur HCP-Aging à résolution fine. Ses cartes d'attention offrent une interprétabilité neurobiologique en priorisant naturellement le réseau frontopariétal (Control), corroborant les théories établies sur la variabilité inter-individuelle. Les classements des modèles restent cohérents entre cohortes, et des parcellisations plus fines bénéficient systématiquement aux modèles basés sur les graphes.

En établissant un cadre robuste, topologiquement conscient et interprétable pour l'empreinte cérébrale, ce travail contribue à l'avancement des outils cliniques individualisés et de la médecine de précision.

Mots-clés : Empreinte cérébrale, Connectivité fonctionnelle, Réseaux de neurones de graphes, Architecture Siamoise, Apprentissage contrastif, Médecine de précision.

ملخص

أصبح فهم الدماغ البشري بوصفه نظاماً بيولوجياً فردياً هدفاً محورياً في علم الأعصاب المعاصر. تبحث هذه الدراسة في ظاهرة البصمة الدماغية، وهي الظاهرة التي يُشكّل فيها الكونيكتوم الوظيفي للفرد توقيعاً بيولوجياً مستقراً وفريداً من نوعه. ورغم دقة التعرف العالية التي تحققها المقاييس الخطية كارتباط بيرسون، إلا أن الأساليب الحالية كثيراً ما تعجز عن استيعاب التبعيات غير الخطية المعقدة، فضلاً عن إهمالها لطوبولوجيا الرسم البياني غير الإقليدية عند تسطيح الكونيكتوم وتحويله إلى متجهات.

تقترح هذه الأطروحة إطاراً تجريبياً منهجياً من ثلاث مراحل لمعالجة هذه القيود، يُقيّم نماذج ذات قدرة تعبيرية متصاعدة: خط أساس يعتمد على ارتباط بيرسون، ومصنفات تعلم آلي كلاسيكية مُصاغة كمهام تحقق زوجية، وبنيات سيامية عميقة تستخدم شبكات العصبونات البيانية (GNN) الموجهة بدالة خسارة تبايُنية (InfoNCE). أُجريت التجارب على بيانات التصوير بالرنين المغناطيسي الوظيفي في حالة الراحة المستمدة من مجموعتي HCP-Young Adult وHCP-Aging، مع تغيير دقة الأطلس التشريحي (Schaefer-100 مقابل Schaefer-400) بصورة منهجية عبر كلتا المجموعتين.

أظهرت النتائج أن الأساليب التمييزية الخاضعة للإشراف تتفوق بشكل ملحوظ على خطوط الأساس غير الخاضعة للإشراف. وتُعدّ شبكة الانتباه البياني السيامية (GAT) الأفضل أداءً في جميع الظروف، إذ تبلغ دقة التعرف 100% على مجموعة HCP- Aging عند استخدام الدقة العالية. وتوفر خرائط الانتباه المكتسبة قابلية تفسير عصبي بيولوجي من خلال إيلائها الأولوية تلقائياً لشبكة التحكم الجبهية الجدارية، مما يعزز بصورة مستقلة النظريات الراسخة حول التباين بين الأفراد. وتبقى تصنيفات النماذج متسقة عبر المجموعتين، في حين تعود فوارق الأداء إلى عوامل ديموغرافية وتقنية. كما تُفيد التقسيمات الأكثر دقة باستمرار النماذج القائمة على الرسوم البيانية، بينما تؤثر سلباً على المصنفات الخطية، مما يوفر توجيهاً تجريبياً للخيارات المنهجية المستقبلية.

من خلال إرساء إطار متين ومدرك للطوبولوجيا وقابل للتفسير لتحقيق البصمة الدماغية، يُسهم هذا البحث في تطوير الأدوات السريرية الفردية والطب الدقيق.

الكلمات المفتاحية :البصمة الدماغية، الكونيكتوم الوظيفي، شبكات العصبونات البيانية، البنية السيامية، التعلم التبايني، الطب الدقيق.

Contents

Acknowledgements **3**

Abstract **4**

List of Abbreviations **13**

Introduction **16**

1 Theoretical Background **21**

1.1 Introduction 21

1.2 Brain Anatomy and Functional Organisation 21

1.2.1 Macroscopic Structure and Functional Specialisation 22

1.3 Neuroimaging: fMRI and the BOLD Signal 23

1.3.1 The BOLD Signal and Neurovascular Coupling 23

1.3.2 Resting-State fMRI 24

1.4 Brain Networks, Parcellation, and Atlases 26

1.4.1 Structural and Functional Connectivity 26

1.4.2 Nodes, Edges, and the FC Matrix 27

1.4.3 Parcellation Schemes and Atlases 29
1.5 Brain Fingerprinting 30
1.5.1 Individual Functional Connectivity as a Biometric 30
1.5.2 Factors Governing Identifiability 30
1.5.3 Brain Fingerprinting and Precision Medicine 31
1.6 Conclusion 31

2 Related Work and Methodology 34

2.1 Introduction 34
2.2 Related Work 34
2.2.1 Classical Functional Connectome Fingerprinting 34
2.2.2 Machine Learning Approaches to Connectome Fingerprinting 38
2.2.3 Deep Learning Approaches for Brain Fingerprinting 40
2.2.4 Brain Fingerprinting for Precision Medicine and Clinical Applications 43
2.2.5 Research Positioning and Proposed Approach 46
2.2.6 Conclusion 46
2.3 Methodology 48
2.4 Data and Preprocessing 49
2.4.1 Cohorts and Acquisition 49
2.4.2 Brain Parcellations 50
2.4.3 Functional Connectivity Estimation 51
2.4.4 Quality Control and Cohort Filtering 51
2.4.5 Subject Level Train / Test Split 51

2.5 Stage 1: Pearson Correlation Baseline 52
2.5.1 Identification Protocol 52
2.5.2 Vectorised Computation 52
2.6 Stage 2: Classical Machine Learning Baselines 53
2.6.1 Preprocessing: Standardisation and PCA Compression 53
2.6.2 Pairwise Feature Construction 54
2.6.3 Linear Discriminant Analysis 54
2.6.4 Linear Support Vector Machines 54
2.6.5 RBF Kernel SVM 55
2.6.6 AdaBoost 55
2.6.7 Hyperparameter Summary 55
2.6.8 Unified Identification Protocol 56
2.7 Stage 3: Deep Siamese Architectures 56
2.7.1 Siamese Twin Framework 57
2.7.2 The InfoNCE Contrastive Objective 57
2.7.3 Method 1: Siamese MLP 58
2.7.4 Functional Connectivity Graph Construction 60
2.7.5 Method 2: Siamese GCN 61
2.7.6 Method 3: Siamese GAT with Attention Interpretability 63
2.7.7 Common Training Protocol 64
2.7.8 Test Time Evaluation Protocol 65
2.8 Evaluation Metrics 65

2.9 Conclusion 66

3 Results and Discussion **68**

3.1 Introduction 68

3.2 Stage 1: Pearson Correlation Baseline 69

3.2.1 HCP-Young Adult 69

3.2.2 HCP-Aging 69

3.3 Stage 2: Classical Machine Learning 70

3.3.1 Results on HCP-Young Adult 70

3.3.2 Results on HCP-Aging 72

3.3.3 Cross-Cohort Comparison 74

3.4 Stage 3: Deep Siamese Architectures 75

3.4.1 Results on HCP-Young Adult 76

3.4.2 Results on HCP-Aging 77

3.4.3 Cross-Cohort Comparison 77

3.5 Effect of Atlas Resolution 78

3.6 Neurobiological Interpretability: GAT Attention Maps 79

3.7 Conclusion 83

Conclusion and Future Research Directions **84**

List of Tables

2.1 Hyperparameters for the Stage 2 classical ML experiments. 56

3.1 Classical ML results on HCP-Young Adult (TOP-1 and ROC-AUC, %). . . . 70

3.2 Classical ML results on HCP-Aging (TOP-1 and ROC-AUC, %). 72

3.3 Siamese deep learning results on HCP-YA (TOP-1 and ROC-AUC). 76

3.4 Siamese deep learning results on HCP-Aging (TOP-1 and ROC-AUC). . . . 77

List of Figures

1.1 Lateral view of the four cortical lobes. 23

1.2 The canonical haemodynamic response function (HRF). 24

1.3 Spatial maps of six major resting-state networks. 26

1.4 Functional connectivity matrix for a representative HCP-YA participant. . . 28

2.1 Three-stage experimental framework overview. 49

2.2 Stage 3 Siamese deep learning framework. 58

2.3 Siamese MLP encoder architecture. 59

2.4 Siamese GCN encoder architecture. 61

2.5 Siamese GAT encoder architecture. 63

3.1 GAT attention maps across cohorts and atlas resolutions. 81

List of Abbreviations

fMRI Functional Magnetic Resonance Imaging

rs-fMRI Resting-State Functional Magnetic Resonance Imaging

BOLD Blood-Oxygen-Level-Dependent

FC Functional Connectivity

HRF Haemodynamic Response Function

DMN Default Mode Network

FPN Frontoparietal Network

MRI Magnetic Resonance Imaging

EPI Echo-Planar Imaging

HCP Human Connectome Project

HCP-YA Human Connectome Project Young Adult

HCP-Aging Human Connectome Project Aging

ABCD Adolescent Brain Cognitive Development

REST1 Resting-State fMRI Session 1 (probe)

REST2 Resting-State fMRI Session 2 (gallery)

ML Machine Learning

DL Deep Learning

PCA Principal Component Analysis

ICA Independent Component Analysis

LDA Linear Discriminant Analysis

SVM Support Vector Machine

RBF Radial Basis Function

MLP Multi-Layer Perceptron

GNN Graph Neural Network

GCN Graph Convolutional Network

GAT Graph Attention Network

InfoNCE Info Noise-Contrastive Estimation Loss

k-NN k-Nearest Neighbour

ReLU Rectified Linear Unit

ELU Exponential Linear Unit

BN Batch Normalisation

DP Differential Power

TOP-1 Top-1 Identification Accuracy

ROC-AUC Receiver Operating Characteristic – Area Under the Curve

ICC Intraclass Correlation Coefficient

PFN Personalized Functional Network

fpFC Fingerprint of Functional Connectivity

ADHD Attention Deficit Hyperactivity Disorder

ASD Autism Spectrum Disorder

NMF Non-Negative Matrix Factorization

SDL Sparse Dictionary Learning

ConvAE Convolutional Autoencoder

CVAE Conditional Variational Autoencoder

LSTM Long Short-Term Memory

Introduction

Understanding the human brain as an **individualised biological system** has emerged as a central ambition of contemporary neuroscience. Historically, the dominant paradigm in functional neuroanatomy was **strict localisationism**—the view that discrete cognitive capacities reside entirely within anatomically circumscribed, independent brain regions. However, the availability of modern whole-brain neuroimaging has shifted the empirical focus toward a **network-based perspective**, revealing that most complex, higher-order cognitive processes depend fundamentally on the coordinated activity of spatially distributed regions that operate as large-scale, interconnected systems.

Functional magnetic resonance imaging (fMRI) serves as the primary non-invasive tool for observing these dynamics by measuring the **blood-oxygen-level-dependent (BOLD) signal**. This signal provides an indirect proxy for neural activity through neurovascular coupling, reflecting synaptic input and intracortical processing. A particularly consequential application is **resting-state fMRI (rs-fMRI)**, which records spontaneous low-frequency BOLD fluctuations in the absence of explicit tasks. These fluctuations reveal the brain's intrinsic functional architecture, characterised as **functional connectivity (FC)**—the statistical dependence of activity time series across different regions.

The Connectome as a Biological Signature

The discovery that an individual's functional connectome is sufficiently stable and unique to serve as a **biological identifier** has introduced the phenomenon of **brain fingerprinting**. For FC to serve as a viable biometric, it must satisfy two critical statistical properties: **temporal stability** (reliability across different time points) and **distinctiveness** (high variability between individuals). Research has established that while primary sensory and motor networks are relatively stereotyped across the population, association networks—specifically

the **frontoparietal (Control) and default mode networks**—exhibit the highest inter-individual variability and are most informative for identification.

Despite high identification accuracies achieved in foundational research using linear similarity measures like **Pearson correlation**, several critical **methodological bottlenecks** persist. Traditional methods often rely on hand-crafted linear metrics that may fail to capture the complex, non-linear dependencies latent within high-dimensional connectivity data. Furthermore, most classical machine learning approaches require flattening the connectome into a one-dimensional vector, a process that strips away the **non-Euclidean graph topology** and modular organisation of the brain, potentially discarding vital identifying information encoded in the network's spatial arrangement.

Research Objectives and Methodology

This research proposes a systematic experimental framework to advance individual identification by evaluating models of increasing expressive capacity under identical experimental conditions. The study is structured around three primary modelling stages:

- **Stage 1:** A zero-training **Pearson correlation baseline** that replicates foundational fingerprinting protocols.
- **Stage 2:** Five **classical machine learning classifiers** (e.g., LDA, SVM, AdaBoost) formulated as a pairwise verification task over a compressed feature space.
- **Stage 3: Deep Siamese architectures** utilising **Graph Neural Networks (GCN and GAT)** and Multi-Layer Perceptrons (MLP) governed by an **InfoNCE contrastive loss**.

The methodology further addresses **resolution sensitivity** by systematically varying atlas granularity—comparing the **Schaefer-100 and Schaefer-400** parcellations—and assesses **cross-cohort consistency** by applying the framework to both the **HCP Young Adult** and **HCP Aging** datasets.

Key Findings and Interpretability

The empirical results confirm that supervised discriminative approaches substantially outperform unsupervised baselines, demonstrating that learning a decision boundary extracts more identity-relevant structure from the connectome. Notably, the **Siamese Graph Attention Network (GAT)** emerged as the strongest performer in every experimental condition. On the HCP-Aging dataset at a fine-grained resolution, the GAT achieved **100% identification accuracy**. The GAT's advantage lies in its use of learned, edge-specific **attention weights**, which allow the model to prioritise the most distinctive functional connections for identity discrimination.

Beyond raw accuracy, the GAT provides **interpretability through attention maps**. These maps revealed that the model naturally prioritises the **frontoparietal (Control) network** for identification, providing computational support for the neurobiological primacy of these regions without requiring explicit supervision on network labels.

Clinical Significance

Establishing high identification accuracy in normative populations is a necessary precondition for the eventual application of these methods to **precision medicine**. Individual FC profiles can serve as **biomarkers** to predict behavioural phenotypes (such as fluid intelligence or attention), stratify patients in heterogeneous conditions like schizophrenia or autism, and track longitudinal changes in brain health. By establishing a robust, topologically-aware framework for brain fingerprinting, this research contributes to the broader effort of moving neuro-diagnostics from group-level averages toward **individualised clinical tools**.

Structure of the Research

The remainder of this thesis is organised as follows:

- **Chapter 1: Theoretical Background** establishes the neurobiological and methodological foundations, tracing the shift toward network-based neuroscience and introducing the conceptual framework of connectome fingerprinting.

- **Chapter 2: Related Work and Methodology** surveys existing literature across three generations of fingerprinting and details the three-stage experimental framework, data acquisition, and preprocessing pipelines.
- **Chapter 3: Results and Discussion** presents the empirical evaluation, characterising how atlas resolution and cohort impact performance, highlighting the dominance of the Siamese GAT, and providing neurobiological validation through attention map analysis.

1

Chapter 1

Theoretical Background

Chapter 1

Theoretical Background

1.1 Introduction

Understanding the human brain as an individualised biological system has emerged as one of the central ambitions of contemporary neuroscience. Decades of post-mortem anatomical study and lesion-based inference have given way to non-invasive neuroimaging techniques, most notably functional magnetic resonance imaging (fMRI), that permit direct observation of large-scale neural dynamics in living individuals. A key insight arising from this literature is that the brain operates not through isolated regions but through the coordinated activity of spatially distributed networks, whose organisation varies meaningfully and reproducibly across people. This inter-individual variability in functional network architecture, far from constituting noise, has proven sufficiently stable and distinctive to serve as a biological identifier, a phenomenon termed brain fingerprinting, with profound implications for precision medicine. This chapter introduces the anatomical and functional organisation of the brain, the principles underlying fMRI and the BOLD signal, the network-theoretic formalisation of connectivity, and the concept of brain fingerprinting together with its clinical significance.

1.2 Brain Anatomy and Functional Organisation

The brain's functional complexity arises from the interplay between its large-scale anatomical organisation and the specialised roles of its constituent regions. This section outlines the macroscopic structure of the cerebral cortex and key subcortical components, framing the

shift from strict localisation towards a network-based understanding of brain function.

1.2.1 Macroscopic Structure and Functional Specialisation

Historically, the macroscale characterisation of the human brain has relied on partitioning the cerebral cortex into geographically and anatomically distinct regions. Structurally, the human cerebral cortex is divided into four principal lobes—frontal, parietal, temporal, and occipital—each broadly associated with distinct functional domains Kandel et al. (2000). Specifically, the frontal lobe subserves motor control, executive function, working memory, and decision-making; meanwhile, the parietal lobe integrates somatosensory and spatial information. In contrast, the temporal lobe primarily supports auditory processing, language, and memory encoding, whereas the occipital lobe is almost exclusively devoted to visual processing Kandel et al. (2000).

Beyond this cortical mantle, subcortical structures remain equally indispensable to whole-brain coordination. The thalamus acts as the principal sensory and motor relay nucleus to the cortex, the hippocampus and amygdala mediate memory consolidation and emotional processing, and the basal ganglia regulate voluntary movement and reward-based learning Kandel et al. (2000); Tian et al. (2020). Collectively, these regions have historically been mapped via localised boundaries, as illustrated in Figure 1.1.

Consequently, for much of the nineteenth and early twentieth centuries, the dominant paradigm in functional neuroanatomy was strict localizationism—the view that discrete cognitive capacities reside entirely within anatomically circumscribed, independent brain regions. However, contemporary neuroscience has substantially revised this rigid picture. While certain specialized areas indisputably display strong functional preferences, most complex, higher-order cognitive processes depend fundamentally on the coordinated activity of spatially distributed regions that operate as large-scale, interconnected systems Sporns (2016).

Indeed, attention and cognitive control depend on distributed frontoparietal interactions, just as self-referential thought and memory consolidation depend on a distributed default mode network spanning the medial prefrontal cortex and posterior cingulate Yeo et al. (2011); Raichle et al. (2001). Ultimately, the availability of modern whole-brain neuroimaging has shifted the empirical focus away from asking where isolated activity occurs, toward uncovering how different regions dynamically co-activate, communicate, and share information across the entire neural topography Fox et al. (2005).

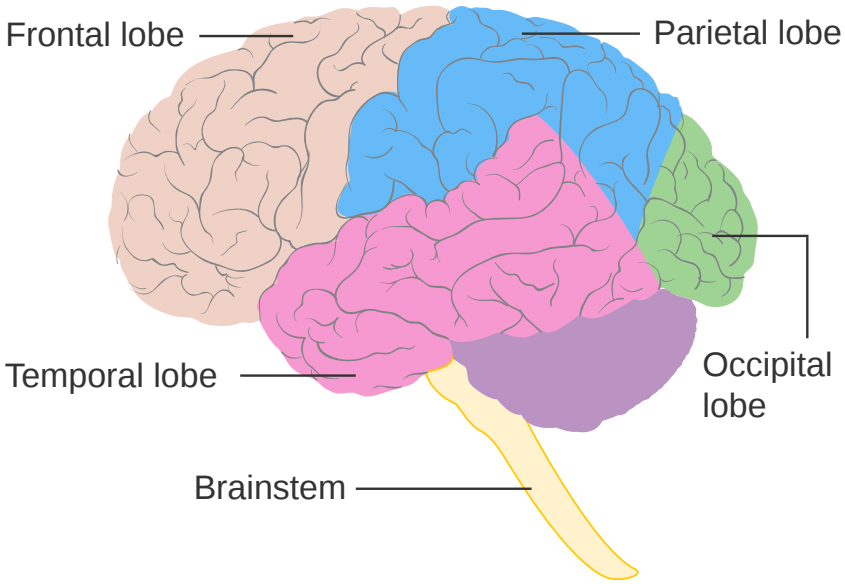


Figure 1.1: Lateral view of the human cerebral hemisphere showing the four cortical lobes colour-coded: frontal (peach), parietal (blue), temporal (pink), and occipital (green). Adapted from Cancer Research UK (2014).

1.3 Neuroimaging: fMRI and the BOLD Signal

1.3.1 The BOLD Signal and Neurovascular Coupling

Functional magnetic resonance imaging (fMRI) measures brain activity indirectly through the blood-oxygen-level-dependent (BOLD) contrast mechanism, first characterised by Ogawa et al. (1990). The key physical insight is that deoxygenated haemoglobin is paramagnetic: it distorts the local magnetic field and accelerates T2* relaxation, thereby reducing the MR signal. Neuronal activity increases local metabolic demand, triggering a haemodynamic response in which cerebral blood flow rises substantially, delivering oxygenated blood in excess of metabolic consumption. As a net result, the local concentration of deoxygenated haemoglobin falls, T2* lengthens, and the BOLD signal increases. This neurovascular coupling was systematically characterised by Logothetis et al. (2001), who demonstrated through simultaneous electrophysiology and fMRI that the BOLD signal correlates most strongly with local field potentials (reflecting synaptic input and intracortical processing) rather than with

single-unit spiking output.

The temporal dynamics of the haemodynamic response are captured by the haemodynamic response function (HRF): a canonical waveform that begins approximately one to two seconds after a brief neural event, peaks around five to six seconds, and returns to baseline over roughly twenty seconds, preceded by a brief undershoot. It is worth emphasising that the BOLD signal is an indirect and temporally sluggish proxy for neural activity, and its interpretation requires careful consideration of potential confounds including physiological noise (cardiac and respiratory fluctuations), head motion, and regional differences in neurovascular coupling.

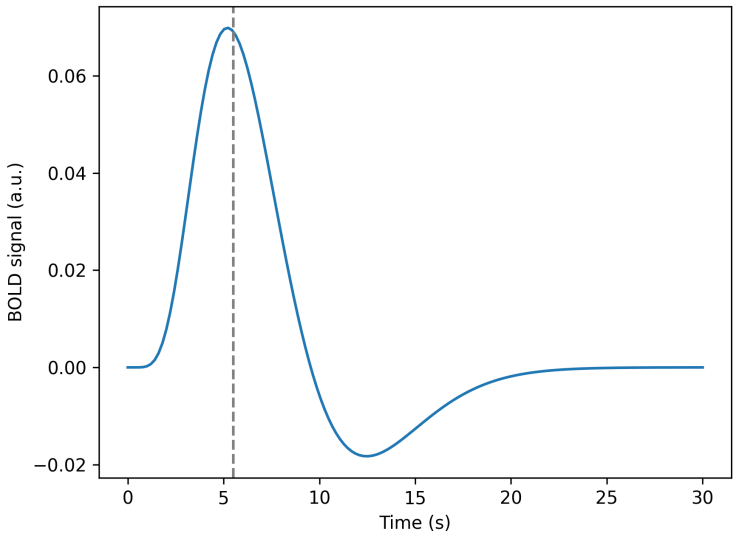


Figure 1.2: The canonical haemodynamic response function (HRF), describing the BOLD signal change following a brief neural event. The response peaks at approximately 5–6 s post-stimulus and is followed by a brief undershoot before returning to baseline.

1.3.2 Resting-State fMRI

A particularly consequential application of BOLD fMRI was introduced by Biswal et al. (1995), who observed that in the absence of any explicit task, low-frequency fluctuations (approximately 0.01–0.1 Hz) in the BOLD signal of the left primary motor cortex were highly correlated with those of the right motor cortex and adjacent sensorimotor regions. This finding implied that the brain maintains coherent patterns of spontaneous activity reflecting

its intrinsic functional architecture, independent of external stimulation. The method became known as resting-state fMRI (rs-fMRI), and the term *functional connectivity* was adopted to describe these temporal correlations.

In a typical rs-fMRI acquisition, participants lie at rest for five to twenty minutes while their spontaneous BOLD fluctuations are recorded. The data require careful preprocessing to remove confounds: slice-timing correction, motion realignment, temporal band-pass filtering, and the regression of nuisance signals (including head motion parameters, white matter signal, and cerebrospinal fluid signal) are standard steps Power et al. (2014). Motion artefacts in particular inflate apparent short-range connectivity while suppressing long-range connections, introducing systematic biases that can corrupt the individual FC profile if unaddressed Power et al. (2014).

Subsequent work established that resting-state functional connectivity reflects genuine large-scale functional architecture. Fox and colleagues demonstrated that spontaneous BOLD fluctuations co-vary with task-evoked activations and deactivations Fox et al. (2005), and a number of canonical resting-state networks have since been reproducibly identified across independent datasets and analytical methods. These include the default mode network (DMN), active during mind-wandering and self-referential thought and suppressed during externally directed tasks Raichle et al. (2001); the frontoparietal network (FPN), associated with executive control and working memory; the salience network, centred on the anterior insula and dorsal anterior cingulate; and dorsal and ventral attention, sensorimotor, visual, and auditory networks Smith et al. (2009). The reliability and neurobiological significance of these networks have made rs-fMRI the dominant tool for characterising individual differences in brain organisation.

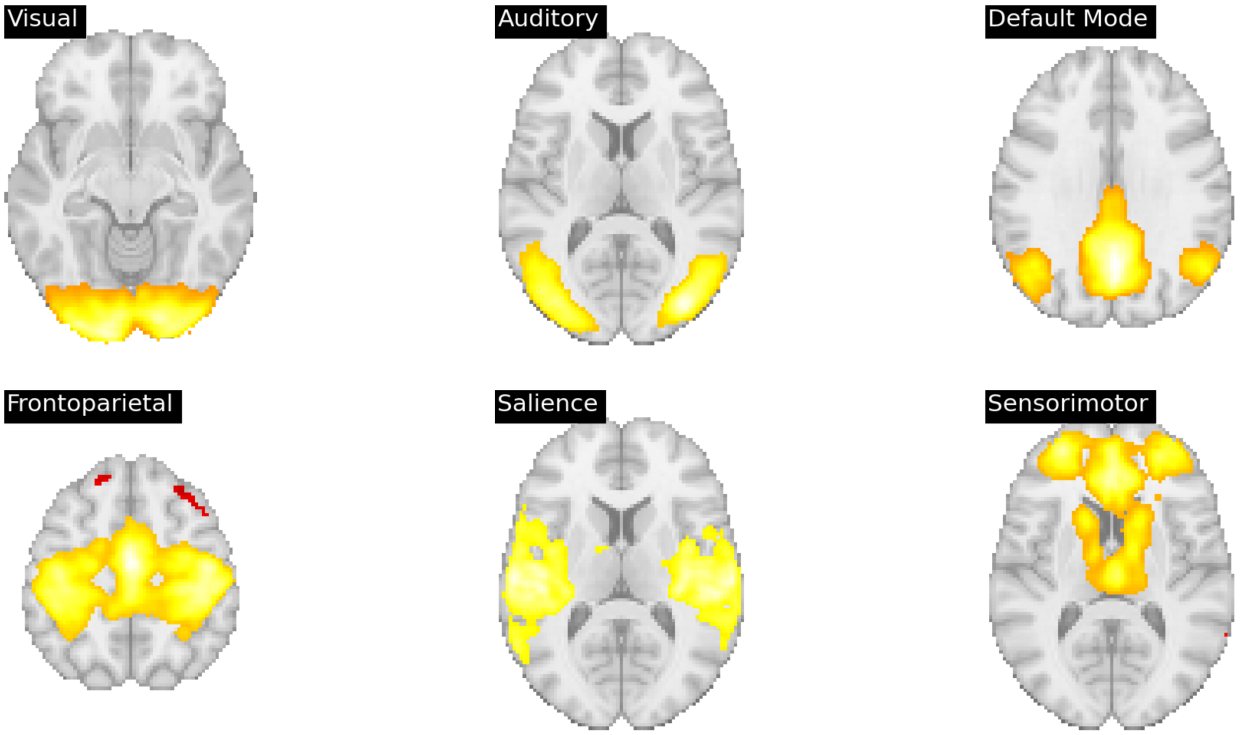


Figure 1.3: Spatial maps of six major resting-state networks identified from rs-fMRI data via independent component analysis, displayed as axial slices on the MNI152 standard brain. Networks shown include the visual, auditory, default mode (DMN), frontoparietal (FPN), salience, and sensorimotor networks. Maps are thresholded at $z > 3.0$. Derived from Smith et al. (2009).

1.4 Brain Networks, Parcellation, and Atlases

1.4.1 Structural and Functional Connectivity

Within network neuroscience, it is important to distinguish between two complementary notions of brain connectivity, each reflecting a different level of neural organisation. **Structural connectivity** refers to the physical wiring of the brain: the white matter fibre tracts that directly connect grey matter regions, measurable via diffusion-weighted MRI and tractography. Structural connectivity is relatively stable across short time scales and represents a set of anatomical constraints on information flow.

Functional connectivity (FC) refers to the statistical dependence of activity time series across brain regions, most commonly measured as the Pearson correlation of BOLD signals in rs-fMRI Friston et al. (1993). Functional connectivity captures something distinct from structural connectivity: two regions may show strong temporal correlation without a direct anatomical connection, provided they receive common input or are linked via polysynaptic pathways.

1.4.2 Nodes, Edges, and the FC Matrix

The application of graph theory to the brain begins with a choice of representation Bullmore and Sporns (2009). Brain regions define the nodes of a network; the statistical dependence between their activity time series defines the edges. For a parcellation of N regions, Pearson correlation of BOLD time series yields an $N \times N$ symmetric functional connectivity matrix, where entry (i, j) contains the correlation between the time series of regions i and j.

The choice of Pearson correlation is motivated by its simplicity, interpretability, and robustness, though it captures only linear dependencies and is sensitive to indirect connections that inflate apparent coupling between region pairs that are not directly functionally related. Partial correlation, which conditions on all other regions, can mitigate the indirect-connection problem, but is sensitive to estimation noise when the number of regions approaches the number of time points Smith et al. (2011).

As illustrated in Figure 1.4, the functional connectivity matrix encodes the Pearson correlation between the resting-state BOLD time series of discrete brain regions. This specific representation utilizes the Schaefer-100 cortical parcellation extended with 16 subcortical regions from the Tian atlas Tian et al. (2020), yielding an $N \times N$ matrix where $N = 116$. To highlight the modular architecture of the brain, the 100 cortical regions are ordered according to their Yeo-7 network assignment (Visual, Somatomotor, Dorsal Attention, Salience/Ventral Attention, Limbic, Control, and Default Mode), followed by the subcortical regions. Visual inspection of the matrix reveals that within-network correlations, which appear as prominent blocks along the main diagonal, are systematically stronger than the off-diagonal between-network correlations.

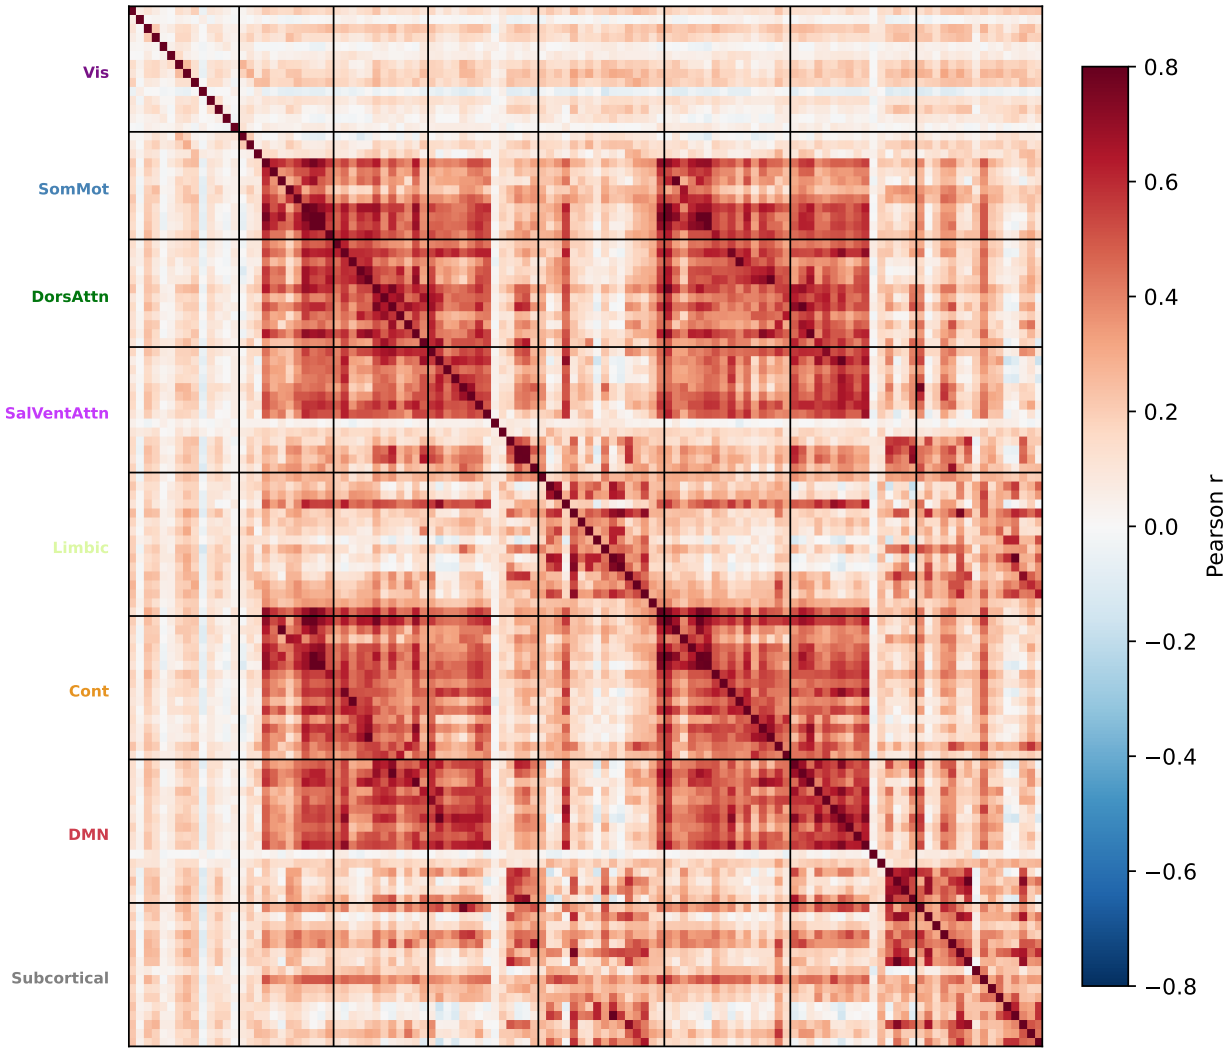


Figure 1.4: Functional connectivity matrix for a representative HCP-YA participant.

1.4.3 Parcellation Schemes and Atlases

The granularity and definition of parcellation nodes are among the most consequential methodological choices in any FC study. A parcellation divides the brain into non-overlapping regions serving as network nodes, and the properties of these regions (their boundaries, spatial extent, and functional homogeneity) directly influence the FC matrices derived from them.

Functionally defined atlases derive parcels directly from resting-state data, grouping voxels with similar FC profiles into regions. The Yeo 7- and 17-network parcellations Yeo et al. (2011) are prominent examples, providing a widely-used reference frame for assigning brain regions to large-scale functional systems.

The Schaefer atlas Schaefer et al. (2018a) represents a sophisticated hybrid approach to brain mapping. Unlike traditional methods that rely exclusively on either local architectural boundaries or global clustering, this atlas is derived from a local-global gradient methodology applied to large-scale resting-state fMRI datasets. Specifically, the model simultaneously optimizes for two competing criteria: local cortical topology and global connectivity patterns. It detects abrupt spatial gradients in functional connectivity to define sharp, contiguous anatomical borders, while simultaneously ensuring that vertices within a single parcel share highly similar whole-brain interaction profiles. A major advantage of this parcellation is its inherently hierarchical design, as it strictly aligns the resulting parcel boundaries with the canonical Yeo 7-network and 17-network frameworks Yeo et al. (2011). Consequently, each discrete region is intrinsically assigned to a well-established macroscale functional system (such as the Default Mode, Frontoparietal, or Salience networks), providing a seamless analytical bridge between localized topological nodes and whole-brain network dynamics.

Fine-grained parcellations preserve spatial detail but substantially increase FC matrix dimensionality and may introduce noise in small regions whose time series are poorly estimated. Coarser parcellations yield more stable, lower-dimensional representations at the cost of spatial resolution.

1.5 Brain Fingerprinting

1.5.1 Individual Functional Connectivity as a Biometric

The notion that individuals are reliably distinguishable on the basis of their brain connectivity patterns fundamentally treats the functional connectome as a biological signature. For any biological trait to serve as a viable biometric identifier, it must satisfy two critical statistical properties: temporal stability (low intra-subject variability across different time points) and distinctiveness (high inter-subject variability within a population). While early functional neuroimaging focused almost exclusively on extracting group-level commonalities, subsequent theoretical frameworks revealed that the residual "noise" discarded in group averages actually contained highly structured, idiosyncratic features unique to individual brains.

1.5.2 Factors Governing Identifiability

Not all brain regions contribute equally to fingerprint distinctiveness. The edges most informative for identification are concentrated in the frontoparietal and default mode networks Finn et al. (2015), association networks that support higher cognitive functions and exhibit substantial inter-individual variability Mueller et al. (2013). Primary sensory and motor networks, which are more stereotyped across individuals, contribute less to identification.

Several factors modulate identification accuracy. **Scan duration** is positively associated with accuracy: longer acquisitions yield more reliable FC estimates, and performance has been shown to increase systematically as scan length increases Finn et al. (2015, 2017). **Atlas resolution** also plays a role: finer-grained parcellations generally retain more individual-specific information, though the relationship is not simply monotonic and depends on the stability of regional time series estimates. **Data quality** (including head motion, signal-to-noise ratio, and the adequacy of preprocessing) is critical; motion artefacts inflate short-range connections and suppress long-range ones, introducing systematic biases that corrupt the individual FC profile Power et al. (2012).

Inter-individual variability in FC is itself partly heritable: twin studies have demonstrated that monozygotic twins show more similar FC profiles than dizygotic twins, indicating a genetic contribution to the connectivity fingerprint Miranda-Dominguez et al. (2018). At the same time, FC patterns are shaped by experience, learning, and developmental history,

making the fingerprint a product of both nature and nurture, a property that is directly relevant to its potential role as a longitudinal biomarker.

1.5.3 Brain Fingerprinting and Precision Medicine

The clinical relevance of brain fingerprinting stems from the bridge it provides between individual brain organisation and individual outcomes. If FC profiles are stable and distinctive enough to identify individuals, they may also serve as biomarkers predicting clinically relevant characteristics. Finn and colleagues demonstrated that edges most informative for identification also predicted individual fluid intelligence scores Finn et al. (2015). Subsequent work has extended this logic to sustained attention Rosenberg et al. (2016) and a wide range of behavioural phenotypes including memory, self-control, and socioeconomic status Ooi et al. (2022).

In the context of neurological and psychiatric conditions, many brain disorders (including schizophrenia, depression, autism spectrum disorder, and dementia) are characterised by alterations in large-scale functional connectivity Buckner et al. (2013). However, these disorders are highly heterogeneous: patients sharing a clinical diagnosis may differ substantially in their symptoms, cognitive profiles, and treatment responses. A precision medicine approach would use individual FC profiles to stratify patients into biologically meaningful subgroups, identify those most likely to benefit from a given intervention, or track longitudinal changes in connectivity as a quantitative treatment biomarker Woo et al. (2017).

The present work contributes to this broader programme by developing and evaluating methods for individual identification from functional connectivity networks in a healthy adult population. Establishing high identification accuracy in this normative setting is a necessary precondition for eventual application to clinical populations, where the fingerprinting signal may be weaker or more variable due to disease-related disruptions in network architecture. The following chapters describe the methodological pipeline, the empirical results obtained, and the directions for future work that they suggest.

1.6 Conclusion

This chapter has established the neurobiological and methodological foundations requisite for understanding the human functional connectome. By tracing the paradigm shift from

strict anatomical localizationism to a network-based perspective, it was highlighted how the brain operates as a highly coordinated, distributed system. Resting-state fMRI provides a non-invasive window into this intrinsic architecture, allowing for the quantification of functional connectivity across the entire neural topography. Furthermore, the application of sophisticated parcellation schemes enables the systematic reduction of high-dimensional neuroimaging data into biologically meaningful, network-aligned nodes.

Crucially, the theoretical framework of connectome fingerprinting demonstrates that these functional networks are not merely generalizable models of human cognition, but contain highly distinct, individual-specific signatures. Having established the conceptual viability of using functional connectivity as a reliable biometric identifier, the subsequent chapter will explore the state-of-the-art computational methodologies. Specifically, it will examine how the field has transitioned from traditional statistical correlation metrics toward advanced machine learning and deep learning architectures to optimize the extraction and matching of these complex neural fingerprints.

2

Chapter 2

Related Work and Methodology

Chapter 2

Related Work and Methodology

2.1 Introduction

Brain fingerprinting has evolved rapidly from a proof-of-concept observation into a rich methodological landscape encompassing classical statistics, machine learning, and deep learning. Navigating this landscape requires both a critical understanding of what has been achieved and a clear rationale for the choices made in any new contribution. This chapter provides both. The first part surveys the existing literature, organised around three successive methodological generations and their clinical extensions. The second part describes the methodology adopted in this work, detailing the pipeline through which functional connectivity data are processed, modelled, and evaluated.

2.2 Related Work

2.2.1 Classical Functional Connectome Fingerprinting

The notion that an individual's brain connectivity profile constitutes a stable, identifying signature analogous to a biological fingerprint was brought into focus by the seminal work of Finn et al. (2015). Using resting-state fMRI data from 126 participants drawn from the Human Connectome Project (HCP) Q2 release, whole-brain functional connectivity (FC) matrices were constructed via Pearson correlation across 268 cortical and subcortical regions

of interest. Rather than treating the connectome as a population-level summary, each subject's FC matrix was treated as a unique, individual-level phenotype. By computing the pairwise similarity between a probe session and a gallery of target sessions belonging to different individuals, identification accuracies of 92.9% to 94.4% were achieved for rest-to-rest session pairs Finn et al. (2015). Strikingly, when the analysis was restricted to frontoparietal networks, which are associated with higher-order cognition and executive function, identification accuracy climbed to approximately 99%, underscoring that not all edges contribute equally to the fingerprint Finn et al. (2015). To formalise this heterogeneity, the concepts of Differential Power (DP) and Group Consistency were introduced, quantifying respectively the degree to which each edge discriminates between individuals and the degree to which it is reliably expressed across the group Finn et al. (2015). A notable limitation of this foundational study, however, was that the two sessions compared were separated by only a single day, leaving open the question of whether such fingerprints would persist across longer time intervals Finn et al. (2015).

That question was addressed directly by Horien et al. (2019), whose investigation of functional connectome stability drew upon four independent longitudinal datasets, including a cohort of 93 participants from the University of Pittsburgh, with interscan intervals ranging from three months to three years. Pearson correlation remained the primary similarity measure, and identification rates remained significantly above chance across all datasets and all time horizons, with subject-specific patterns detectable even when scans were separated by up to three years Horien et al. (2019). These findings reframed the functional connectome not as a momentary state but as a durable trait-level feature of neural organisation, one that persists through the slow biological changes accompanying early adulthood. The robustness of this conclusion was, however, tempered by several methodological constraints: all datasets relied exclusively on resting-state acquisitions, and differences in acquisition protocols across sites precluded direct quantitative comparisons of identification rates between samples Horien et al. (2019).

Building upon these foundational findings, the quest for optimal identifiability was further advanced by Amico and Goñi (2018). Evaluating data from 100 unrelated HCP subjects, they noted that classical whole-brain resting-state identifiability, while robust at over 90%, declined significantly during cross-task or task-to-rest evaluations, where it traditionally ranged between 54% and 87% Amico and Goñi (2018). To address this vulnerability, they introduced a principal component analysis (PCA) based "connectivity eigenmode" approach to reconstruct functional connectomes. By projecting the connectome into a lower-dimensional eigenspace and filtering out components associated with noise or transient task-induced vari-

ance, they effectively purified the underlying subject-specific functional signature. This eigenmode reconstruction yielded a substantial increase in cross-task identifiability, demonstrating that extracting the optimal latent subspace of the connectome is crucial for isolating the core individual fingerprint across varying cognitive states Amico and Goñi (2018).

A complementary perspective emerged from the Midnight Scan Club (MSC), a precision imaging study in which ten healthy adults each contributed more than five hours of resting-state fMRI data across multiple sessions. By accumulating data a t the scale necessary for reliable within-subject estimation, Gordon et al. (2017) demonstrated that individual FC organisation is qualitatively distinct from what is recoverable through group-average approaches, a distinction that is systematically obscured when sparse data are pooled across participants. Stability analyses revealed that correlation matrices required between 30 and 90 minutes of scan time before reaching an acceptable reliability threshold ($r > 0.85$), raising substantive concerns about the adequacy of single-session acquisitions in standard neuroimaging protocols Gordon et al. (2017). While measures like modularity and participation coefficient proved either statistically invariant across subjects or unreliable, global efficiency and specific topological network arrangements exhibited pronounced inter-individual differences, reinforcing the view that network topology itself carries identifying information Gordon et al. (2017). The primary limitation of this work was necessarily its sample size: with only ten subjects, the study could not address whether the observed precision was generalisable, and one participant's data exhibited reduced reliability attributed to drowsiness during scanning Gordon et al. (2017).

While the studies above relied on static FC matrices averaged across entire scan sessions, evidence began to accumulate that the temporal dynamics of connectivity carry distinct and complementary identifying information. Liu et al. (2018) demonstrated this using a sliding-window dynamic functional connectivity (DFC) framework applied to 105 HCP participants, in which DFC strength, stability, and variability were each evaluated as candidate fingerprint features. Identification accuracies of 90.5% to 91.4% were achieved using DFC strength, while DFC variability, which captures the moment-to-moment fluctuations in coupling, was found to predict individual differences in cognitive flexibility with a correlation of r = 0.49, establishing a meaningful link between dynamic network properties and higher cognitive function Liu et al. (2018). This chronnectome framework was subsequently extended by Ghaffari et al. (2025), who employed a more principled decomposition of connectivity dynamics: rather than relying on sliding windows of fixed length, brain activity was segmented into seven metastable states across 215 HCP participants, and the concatenated state-specific connectomes yielded an identification accuracy of 97.21%, substantially

outperforming static methods, which plateaued at approximately 88–91% in the same sample Ghaffari et al. (2025). Both approaches are nonetheless subject to sensitivity regarding window length and parcellation scheme, and the neurobiological interpretation of discrete metastable states remains an active area of theoretical debate Ghaffari et al. (2025).

A separate line of inquiry has questioned whether the reliance on Pearson correlation as the universal similarity measure constitutes a methodological bottleneck in its own right. A geometry-aware alternative was proposed by Uddin et al. (2026), who argued that FC matrices are symmetric positive semi-definite matrices inhabiting a curved, non-Euclidean space known as a Riemannian manifold, and that standard measures such as Pearson correlation fail to respect this geometry. The Alpha-Z Bures-Wasserstein divergence addresses this by operating directly on this manifold, with its Alpha-Z parameterisation introducing additional flexibility that makes it robust to the rank deficiency inherent to high-resolution parcellations. Applied to 428 unrelated participants from the HCP 1200 release, this divergence achieved near-ceiling identification rates above 0.94 at the 400-parcel level and 100Euclidean distance, and the standard Bures-Wasserstein distance Uddin et al. (2026). However, the neurobiological underpinnings of this improvement are still unknown, and it is still unclear whether geometry-aware metrics can be practically interpreted in a clinical setting Uddin et al. (2026).

The choice of similarity measure was examined more broadly by Liu et al. (2025), whose benchmarking study evaluated 239 distinct pairwise interaction statistics spanning covariance, partial correlation, mutual information, and spectral coherence across 326 unrelated HCP subjects. Performance was assessed using the identifiability score introduced by Amico and Goñi (2018), which quantifies how much more similar an individual's FC profiles are to each other across sessions than to those of other subjects. Precision-based statistics, which incorporate inverse covariance estimation to control for indirect functional influences, achieved identifiability scores exceeding 2.1 on this metric, outperforming all other families of measures, and were also associated with the strongest observed structure-function coupling ($R^2 = 0.25$) Liu et al. (2025).

Crucially, the question of test-retest reliability was addressed in parallel by Jiang et al. (2023), whose systematic evaluation of over 524,000 combinations of methodological parameters across 41 HCP participants with repeated acquisitions demonstrated that whole-brain parcellations incorporating subcortical regions (such as the wbLGP-458 atlas) yielded substantially higher individual reliability (mean ICC = 0.671) than cortical-only schemes (Jiang et al., 2023).

Despite their methodological diversity, these classical studies converge on a common conclusion: individual functional connectivity profiles are stable, distinctive, and reliably identifiable across sessions. Yet they also reveal a shared limitation, namely the ceiling imposed by linear similarity measures, which motivates the transition towards machine learning frameworks capable of uncovering the non-linear, higher-order structure embedded within connectivity matrices.

2.2.2 Machine Learning Approaches to Connectome Fingerprinting

Once the foundational literature established that functional connectomes exhibit stable, subject-specific characteristics, the methodological challenge shifted towards treating these massive connectivity arrays as high-dimensional inputs for formal machine learning (ML) frameworks. The complexity of this undertaking was sharply illustrated by Shahhosseini and Miranda (2022), whose systematic evaluation of functional connectivity estimation methods on a single HCP participant demonstrated that the choice of statistical metric profoundly alters the inferred network topology. By comparing cross-correlation, partial correlation, Independent Component Analysis (ICA), and Principal Component Analysis (PCA), it was shown that summary measures such as mean network degree could vary dramatically, ranging from 69.5 down to 1.31, depending entirely on the estimation approach Shahhosseini and Miranda (2022). This extreme sensitivity underscores why extracting the 86,320-dimensional upper-triangle vector of a high-resolution FC matrix is non-trivial; it requires analytical strategies capable of distilling meaningful variance from an overwhelmingly large and structurally complex feature space.

To navigate the curse of dimensionality inherent to functional neuroimaging data, researchers initially turned to linear subspace projections to identify the latent regimes where individual fingerprints are most distinctly encoded. A prominent example is the Graph Embedding for Functional Fingerprinting (GEFF) framework proposed by Abbas et al. (2020), which concatenated vectorized FC matrices from 100 HCP subjects to perform a group-level PCA decomposition. By projecting individual connectomes into this learned eigenspace and computing cosine similarity, the GEFF approach achieved 100% identification accuracy for rest-to-rest session pairs; within-learning-task identification using resting-state and one additional task reached 98%, while cross-task generalisation to entirely unseen task types with the same two-task configuration achieved approximately 90% Abbas et al. (2020). However,

the reliance on a group-level eigenspace introduced a significant operational limitation: the entire representational space must be recomputed whenever new subjects are added to the gallery, rendering the approach computationally prohibitive for continuous, open-set verification scenarios.

An alternative dimensionality reduction strategy focused on isolating single, stable neural architectures rather than distributed eigenspaces. Liu and Zhang (2024) applied single-subject spatial ICA without temporal centering across a massive cohort of over 2,000 participants spanning the HCP Development, Aging, and Young Adult datasets. By isolating the component with the highest explained variance, termed the Maximal Independent Component (MIC), and demonstrated that this single spatial map could serve as a highly reliable individual identifier. The MIC achieved a 100% identification rate in both the child and elderly cohorts, outperforming traditional whole-connectome fingerprinting methods Liu and Zhang (2024). Curiously, the MIC exhibited lower stability in the young adult cohort, suggesting that the dominance of this baseline neural code may fluctuate across the lifespan, and highlighting a potential limitation of linear ICA assumptions when modelling the complex dynamics of the fully mature brain.

While projection methods focused on mapping the feature space, a parallel line of research reframed the fingerprinting problem as a supervised, pairwise classification task. In this paradigm, rather than computing raw similarities, ML models are trained to explicitly verify whether two input vectors belong to the same individual or to different individuals. Shojaee et al. (2019) established this framework using data from 1,003 HCP healthy young adults, formally comparing the discriminative capacity of several classical classifiers against the traditional correlation baseline. The results demonstrated significant variability in algorithmic suitability: while k-Nearest Neighbours achieved an identification accuracy of 76.16%, Decision Trees and Naive Bayes performed poorly, yielding accuracies of 55.51% and 46.97%, respectively Shojaee et al. (2019). Despite these varying performance levels, the conceptual shift towards supervised verification proved critical, though it was noted that optimising these classical models for such high-dimensional data incurred substantial computational complexity.

The dependency of classical classifiers on specific preprocessing pipelines was systematically quantified by Hannum et al. (2023), who evaluated 16 distinct functional connectivity construction methods alongside five ML architectures to decode both subject identity and cognitive state. Across a sample of 865 HCP participants and eight cognitive states, Linear Discriminant Analysis (LDA) emerged as exceptionally robust, utilising generalised singular

value decomposition to circumvent dimensionality issues, and achieved up to 99% accuracy for both subject identification and task decoding Hannum et al. (2023). Conversely, the performance of Support Vector Machines (SVM) was found to be highly volatile, yielding accuracies that ranged dramatically from 17% to 61% depending entirely on the chosen pipeline, underscoring the brittle nature of kernel methods when applied to unfiltered connectome data Hannum et al. (2023). A persistent concern across these highly optimised, classical ML pipelines is their potential tendency to overfit the pristine acquisition standards of the HCP protocol, leaving their generalisability to noisier, clinical datasets largely uncertain.

Beyond discrete identification, matrix factorization techniques have been successfully applied to map inter-individual heterogeneity along a continuous spectrum. Keller et al. (2023) utilised spatially-regularized non-negative matrix factorization (NMF) to define Personalized Functional Networks (PFNs) in a massive cohort of 6,972 children from the ABCD study Casey et al. (2018). By shifting the focus from edge-weights to the spatial topography of the networks themselves, ridge regression models trained on these personalized layouts successfully predicted general cognition, achieving accuracies of $r = 0.41$ in the discovery sample and $r = 0.45$ in the replication sample Keller et al. (2023). While powerful, this approach was constrained by its reliance on a single timepoint and the deliberate exclusion of subcortical regions from the topographic analysis, limiting its utility for whole-brain longitudinal tracking.

Despite the substantial improvements in identification accuracy and behavioural prediction offered by classical machine learning algorithms, these approaches share a fundamental architectural limitation: by flattening the connectome into a one-dimensional vector, they strip away the complex, non-Euclidean graph topology of the brain, a deficiency that necessitates a transition towards deep learning frameworks explicitly designed to operate on interconnected networks.

2.2.3 Deep Learning Approaches for Brain Fingerprinting

As the limitations of feature engineering and linear matrix operations became apparent, the field pivoted towards representation learning, leveraging deep neural networks to autonomously discover discriminative hierarchical features within raw functional connectivity data. The feasibility of applying artificial neural networks to connectome fingerprinting was initially demonstrated through the use of shallow feedforward architectures. Sarar et al. (2021) introduced the CorrNN and NormNN models, specifically designed to bypass the com-

plexity of recurrent neural networks (RNNs) when processing temporal neuroimaging data. Evaluated on a sample of 100 HCP participants, these purely feedforward networks achieved state-of-the-art identification accuracies exceeding 99.5%, remarkably doing so using scan segments as short as 20 seconds Sarar et al. (2021). This early success established that deep models could extract highly specific individual signatures from dramatically less data than required by classical correlation methods, though the massive parameter count necessary to process such short, dense data spans highlighted an ongoing computational challenge.

The capacity of neural networks to discover latent identifying structures was further advanced by Lee and Lee (2024), who applied a weight-sparsity deep neural network (DNN) across a large cohort of 870 HCP participants, corroborated by an independent dataset of 15 healthy adults. This work introduced the conceptual framework of the "fingerprint of FC" (fpFC), an intrinsic, stable neural mode explicitly learned by the network's hidden layers. By extracting this fpFC representation, the network identified 870 HCP participants with 93.3% accuracy and 300 individuals with 97.1% accuracy, both using 15-second time-varying FC windows, decisively outperforming shallow machine learning counterparts including Support Vector Machines and Logistic Regression Lee and Lee (2024). However, this framework was strictly formulated as a closed-set classification problem, limiting its utility to verifying individuals already present within the training cohort, and raising persistent concerns regarding the risk of overfitting the highly parameterized network to specific dataset characteristics.

To move beyond discrete label classification and toward generative representation learning, researchers began employing autoencoder architectures to explicitly separate shared human connectome features from individual-specific variability. Ghosh et al. (2025) developed a hybrid framework utilizing a Convolutional Autoencoder (ConvAE) coupled with Sparse Dictionary Learning (SDL) to process resting-state and task-based data from 339 HCP participants. By isolating the residual connectome, specifically the portion of the functional matrix not reconstructed by the shared-feature ConvAE, and applying SDL to extract discriminative structural atoms, the model achieved approximately a 10% improvement over baseline group-averaged functional connectivity models Ghosh et al. (2025). The primary constraint of this advancement was its focus on task-evoked fMRI, necessitating further validation across more diverse, clinical populations where active task performance cannot be guaranteed.

The generative separation of shared versus individual features was further refined by Lu et al. (2024), who replaced the deterministic convolutional approach with a probabilistic Conditional Variational Autoencoder (CVAE) evaluated on 997 HCP subjects. In this

framework, the CVAE actively models the commonalities shared across the group, while the residual functional connectomes are subsequently processed via SDL to aggressively amplify inter-subject variability. This synergistic approach yielded a near-perfect 99.7% identification accuracy for resting-state pairs, representing one of the highest benchmarks achieved for static connectivity Lu et al. (2024). Nevertheless, the architecture remains constrained by its reliance on static, flattened FC matrices, failing to account for the spatial topography and dynamic fluctuations of the living brain.

The persistent structural limitation of treating brain networks as flat vectors eventually necessitated the adoption of Graph Neural Networks (GNNs), a class of architectures explicitly designed to operate on non-Euclidean, topological data. A comprehensive 2025 systematic review by Mohammad et al. (2025), analysing 52 distinct studies, underscored that GNNs, encompassing Graph Convolutional Networks (GCNs), Graph Attention Networks (GATs), and Dynamic Graph CNNs (DGCNNs), are uniquely positioned to capture the non-linear, hierarchical patterns embedded within functional connectomes. By respecting the modularity and node connectivity of the brain, these architectures have demonstrated promising diagnostic performance across a range of psychiatric classification tasks Mohammad et al. (2025). The review noted, however, that these topological models frequently suffer from "black-box" interpretability issues and carry a substantial risk of overfitting when trained on the relatively small neuroimaging sample sizes typical of clinical settings.

The integration of graph-aware deep learning into the specific domain of connectome fingerprinting was pioneered by Shojaee et al. (2019), who introduced a Siamese Graph Convolutional Network applied to 1,003 HCP participants. By structuring the problem as a pairwise evaluation, simultaneously processing two distinct brain graphs through shared GCN weights to output a similarity score, their "Functional 64-16-4" model achieved a fingerprinting accuracy of 99.89% Shojaee et al. (2019). This milestone definitively demonstrated that graph-aware deep learning outperforms all classical ML baselines by constructing higher-order, non-linear features directly from the connectivity topology. However, this framework was formulated as a binary classification task, outputting a discrete "same" or "different" label, which limits its flexibility for continuous tracking and requires a computationally intensive training process spanning 3 to 4 hours per cross-validation fold Shojaee et al. (2019).

The transition from discrete classification to continuous spatial embedding marks the current frontier of fingerprinting methodology, a shift exemplified by the deep metric learning framework proposed by Xu et al. (2026). Termed Metric-BolT, this approach applied Fused Window Transformers to longitudinal data from 1,325 developing adolescents (ages 9–14)

within the ABCD cohort. By optimising a Triplet Margin Loss function (a training objective that minimises the distance between scans of the same individual while maximising the distance between scans of different individuals in the embedding space), the network was trained to pull intra-subject scans together in a high-dimensional embedding space while pushing inter-subject scans apart Xu et al. (2026). This open-set metric learning paradigm proved exceptionally robust, maintaining a 97.6% identification accuracy within a single session and preserving an 86.6% accuracy across a 4-year developmental interval Xu et al. (2026). This architectural evolution, from the rigid binary classification framing of early GCNs to the flexible, open-set metric embedding of contrastive models, directly informs the methodological trajectory of the present thesis, establishing the theoretical foundation for adopting a Siamese Graph Neural Network governed by contrastive loss to map stable, long-term connectome geometries.

2.2.4 Brain Fingerprinting for Precision Medicine and Clinical Applications

The scientific significance of connectome fingerprinting extends well beyond the identification task itself. As the field has matured, it has become increasingly apparent that the stable individual representations uncovered by fingerprinting methods encode clinically and behaviourally meaningful information about the person, a shift in framing with profound implications for personalized medicine. Ramduny and Kelly (2025) formalised this distinction in a comprehensive review spanning datasets ranging from 10 to 3,900 participants, including the HCP, ABCD, and UK Biobank cohorts, by articulating the critical dissociation between connectome distinctiveness and behavioural relevance. Evidence suggests that high identification accuracy, while necessary, is not a sufficient condition for clinical utility; a fingerprint can be highly unique and yet fail to explain meaningful variance in cognitive or psychiatric outcomes Ramduny and Kelly (2025). The absence of a physiological ground truth for functional connectivity further complicates the interpretation of these representations, and the review underscored that reproducibility must ultimately be coupled with demonstrated validity before fingerprinting can transition into a clinical tool.

Establishing that validity has required linking individual connectivity profiles directly to cognitive and behavioural phenotypes. A landmark study by Keller et al. (2023) pursued this objective at scale, applying spatially-regularized non-negative matrix factorization to define Personalized Functional Networks (PFNs) in 6,972 children from the ABCD study.

The spatial topography of these personalized networks within association cortex, varying systematically along the sensorimotor-association axis, predicted general cognitive ability with accuracies of $r = 0.41$ in the discovery sample and $r = 0.45$ in replication, demonstrating that individual network layout itself carries predictive signal beyond what raw edge weights encode Keller et al. (2023). This work motivated a broader conceptual framework: individual variability in network organisation is not noise to be averaged away but a structured biological signal that reflects cognitive capacity and, by extension, clinical vulnerability.

This individual-differences logic was extended into a formal normative deviation paradigm by Mohammad et al. (2025), whose synthesis of precision neurodiversity reviewed evidence suggesting that neurological conditions such as ADHD and autism spectrum disorder should be characterised not by group averages but by each patient's deviation from a healthy population norm. Drawing on PFNs, Connectome-based Predictive Modeling (CPM), and topological deviation indices, the review synthesised findings indicating that individualized parcellation of the Default Mode Network has been shown to improve autism classification performance, with reported AUC values of 0.87 compared to 0.74 for standard atlas-based approaches, and highlighted the identification of distinct neurobiological subtypes within the ADHD population that are invisible to conventional diagnostic criteria Mohammad et al. (2025). The clinical implication underscored by this synthesis is significant: if a healthy connectome baseline can be established for an individual, subsequent deviations from that baseline under disease conditions become a principled, quantitative signal for diagnosis and treatment monitoring.

The perturbation of individual connectome geometries by neuropathological processes has been investigated most directly in oncological settings. Colpo et al. (2025) constructed an Integrated Connectivity matrix, concatenating functional, structural, and microstructural connectivity, for 41 glioma patients, and demonstrated that lateralised gliomas disrupt networks in both hemispheres, with functional connectivity measures proving more sensitive than structural measures for detecting abnormalities in regions distant from the lesion Colpo et al. (2025). This work establishes a precedent for using integrated, multi-modal connectivity profiles as a basis for characterising individual-level neuropathological disruption, representing a step towards connectivity-derived biomarkers of disease severity and trajectory, even in the absence of a matched healthy control group.

In the domain of psychiatric neuroscience, dynamic and graph-structured connectivity representations have been applied to the automatic classification of neurodevelopmental disorders. Tang et al. (2024) demonstrated that a GNN-LSTM architecture incorporating

dynamic functional connectivity, estimated using Ledoit-Wolf regularisation within sliding windows, and a dynamic graph pooling mechanism achieved diagnostic accuracies of 80.4% on ABIDE I and 79.63% on ABIDE II for autism spectrum disorder. These results indicate that the spatio-temporal connectivity fingerprint of ASD carries discriminative information that static methods risk missing, though the exclusive reliance on fMRI data, without integration of genetic or environmental covariates, limits the biological completeness of the characterisation Tang et al. (2024).

Perhaps the most direct demonstration of clinical translation is found in the work of Tripathi et al. (2024), who applied connectome fingerprinting to presurgical motor mapping in patients with brain tumors. By training models on resting-state data from 208 HCP participants, the framework learned to predict individual task-evoked motor activations from connectivity alone, achieving prediction accuracy on par with within-subject task fMRI stability, and successfully estimated eloquent motor cortex locations in several patients who were unable to perform active tasks Tripathi et al. (2024). This application reframes resting-state fingerprinting as a practical clinical instrument: the subject's unique connectivity signature, when measured at rest, encodes sufficient information to reconstruct their individual functional organisation, even in the presence of pathological network disruption.

Looking further towards a preventative horizon, Hermens et al. (2023) proposed deploying brain fingerprinting at population scale to identify youth at risk for mental health disorders before symptom onset. Drawing on the LABS longitudinal study of early adolescents aged 12 and referencing the broader ABCD cohort of approximately 11,800 participants, Hermens et al. (2023) argued that the uniqueness of the functional connectome in early adolescence predicted psychological distress outcomes four months later, suggesting that the fingerprint encodes vulnerability information that precedes clinical presentation Hermens et al. (2023). The ethical and logistical dimensions of universal adolescent neuroimaging, including the risk of biotyping stigma and the prohibitive cost of population-scale MRI, remain substantial barriers that the authors themselves acknowledged. The pharmacological dimension of precision medicine was addressed in a review by Williams and Whitfield Gabrieli (2025), who synthesised evidence from major antidepressant trials including EMBARC, iSPOT-D, and CAN-BIND, demonstrating that individualised brain circuit scores referenced against healthy normative distributions predicted antidepressant response with greater than 75% accuracy, with dorsolateral prefrontal cortex activity norms correctly identifying 63% of remitters Williams and Whitfield Gabrieli (2025). This body of evidence collectively establishes the clinical regulatory context within which connectome-derived biomarkers are gradually finding institutional legitimacy.

2.2.5 Research Positioning and Proposed Approach

In light of the gaps identified in the current state of the art, particularly the absence of controlled cross-method comparisons, the poorly characterised sensitivity of model families to parcellation resolution and cohort composition, and the limited interpretability of learned representations, this work proposes a systematic experimental framework to advance the field of functional connectome fingerprinting.

Our proposed approach directly addresses these shortcomings by evaluating models of increasing expressive capacity under identical experimental conditions. The methodology is structured around three main objectives:

- **First**, to establish a fair cross-method comparison, we evaluate three stages of modelling capacity, namely a Pearson correlation baseline, classical machine learning classifiers, and deep Siamese architectures, under a shared cohort, parcellation, train/test split, and evaluation protocol, so that any performance difference can be attributed to the model alone.
- **Second**, to characterise resolution sensitivity, we systematically vary atlas granularity across a coarse (Schaefer-100 + Tian, $N = 116$) and a fine (Schaefer-400 + Tian, $N = 416$) parcellation, examining how different model families respond to the increased dimensionality and topological detail of finer parcellations.
- **Third**, to assess cross-cohort consistency, we apply the full experimental framework to both the HCP Young Adult and HCP Aging datasets and examine whether model rankings, performance trends, and attention-based interpretations remain stable across two cohorts that differ substantially in age range, scanner platform, and inter-subject variability. This comparison characterises the robustness of each modelling approach to population-level variation rather than claiming out-of-distribution generalisation.

2.2.6 Conclusion

This chapter has outlined the current landscape of functional connectome fingerprinting, highlighting both its remarkable progress and its persistent limitations. Despite advances spanning correlation-based methods, classical machine learning, and deep graph architectures, challenges related to cross-method comparability, resolution sensitivity, cohort generalisability, and model interpretability remain unresolved.

In response to these challenges, we have introduced a structured experimental framework designed to advance individual identification from functional connectivity through three stages of increasing modelling capacity, systematic variation of atlas resolution, and evaluation across two HCP cohorts differing in age composition and inter-subject variability to assess cross-cohort consistency. The next chapters detail the implementation of these methods, the preprocessing pipeline, and the evaluation strategies used to assess their performance and biological interpretability.

2.3 Methodology

The work proceeds along three experimental stages of increasing modelling capacity. Stage 1 reproduces the original fingerprinting protocol of Finn et al. (2015) as a zero training reference. Stage 2 introduces five classical machine learning classifiers in a pairwise verification formulation, replacing the raw Pearson criterion with learned decision functions over a compressed feature space. Stage 3 then turns to deep Siamese architectures that learn an embedding space directly from the connectivity data, with three encoder variants: a multi layer perceptron, a graph convolutional network, and a graph attention network. All three stages share the same cohort, the same identity preserving train / test split, and the same evaluation protocol, so that any performance gap between them can be attributed to the model alone.

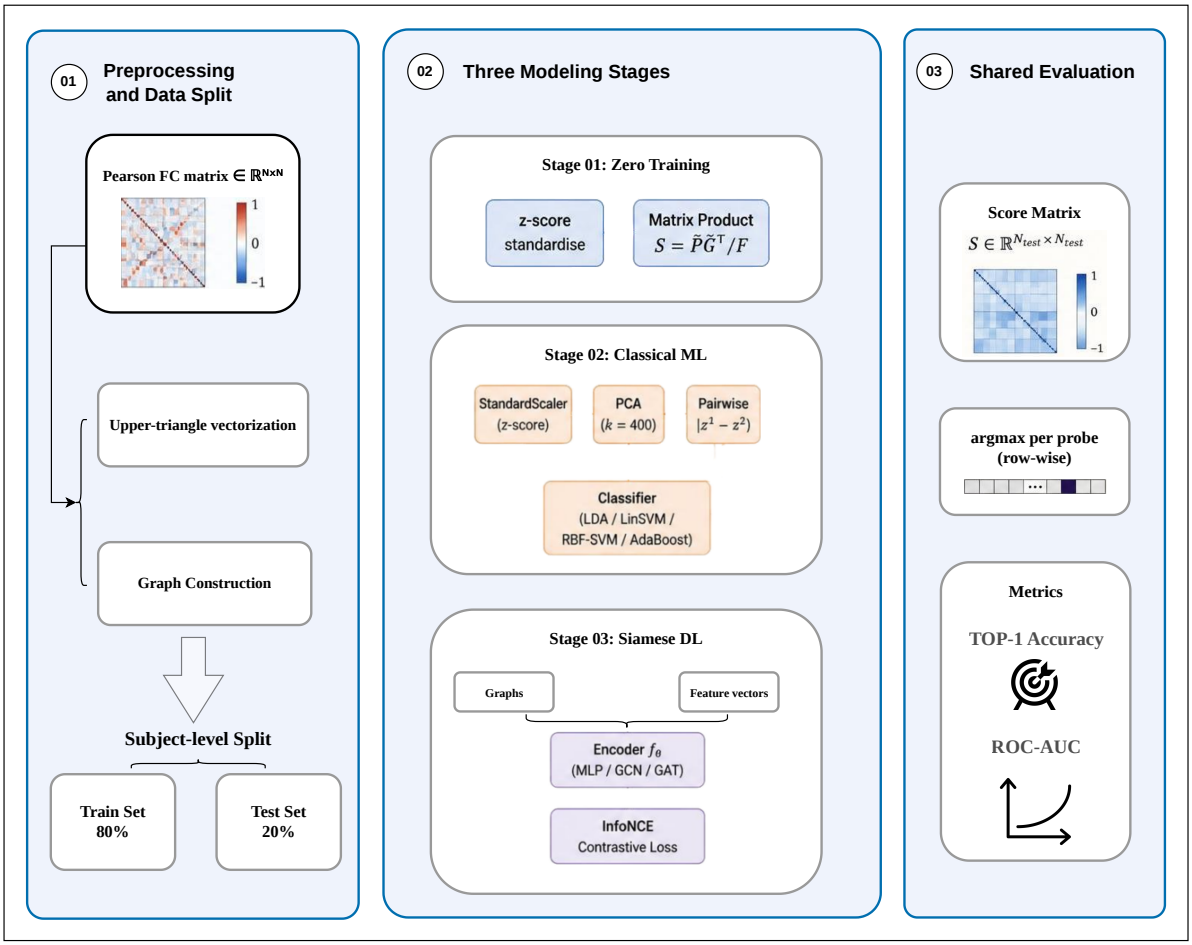


Figure 2.1: Overview of the three-stage experimental framework. (01) Preprocessing: Pearson FC matrices are vectorised and converted into graphs; subjects are partitioned into 80% training and 20% test sets. (02) Three modelling stages: Pearson correlation baseline (Stage 1), classical ML classifiers with PCA compression (Stage 2), and Siamese deep learning architectures with InfoNCE loss (Stage 3). (03) Shared evaluation via row-wise argmax on a score matrix, with performance reported as TOP-1 accuracy and ROC-AUC.

2.4 Data and Preprocessing

2.4.1 Cohorts and Acquisition

Two open-access cohorts are used. The Human Connectome Project Young Adult release (HCP-YA, Van Essen et al., 2013) supplies resting-state fMRI for healthy adults aged 22–35,

collected as part of a large-scale neuroimaging initiative that acquired multimodal imaging and extensive behavioural data from 1,200 participants. Scanning was performed on a customised Siemens 3T Connectome scanner with a 32-channel head coil using a multiband EPI sequence at 2 mm isotropic resolution and a repetition time of 720 ms. Each resting-state session consisted of two 14-minute sub-scans acquired in opposite phase-encoding directions (left-right and right-left), which were demeaned and variance-normalised prior to concatenation, yielding approximately 2,400 time points per session. The Human Connectome Project Aging release (HCP-Aging, Bookheimer et al., 2019) extends the same acquisition family to a lifespan cohort spanning 35 to 100 years of age, scanned on a Siemens Prisma scanner with sequences harmonised across four sites at a repetition time of 800 ms. Each HCP-Aging resting-state session similarly comprised two 6.5-minute sub-scans acquired in opposite phase-encoding directions (anterior-posterior and posterior-anterior), demeaned, variance-normalised, and concatenated. Both releases provide minimally preprocessed time series through the HCP pipelines of Glasser et al. (2013), including ICA-FIX denoising for the resting-state runs (Salimi-Khorshidi et al., 2014). Precomputed Pearson correlation matrices for both cohorts are drawn from the Tijhuis et al. repository (Tijhuis et al., 2022), derived from the minimally preprocessed rs-fMRI data. Two scanning sessions per subject (REST1 and REST2) are required for the fingerprinting task: REST1 acts as the probe session, REST2 as the gallery. After quality filtering for complete REST1/REST2 session pairs, 1,018 subjects were retained for HCP-YA and 720 for HCP-Aging. The HCP-YA's combination of large sample size, long scan durations, high-quality preprocessing, and open data distribution makes it the most widely used dataset in brain fingerprinting research, ensuring that results reported here are directly comparable with the broader literature.

2.4.2 Brain Parcellations

Each subject's preprocessed time series is averaged within the parcels of a cortical atlas extended by a subcortical complement. The cortical surface is parcellated using the Schaefer atlas (Schaefer et al., 2018b), which provides a family of fine-grained parcellations derived from a gradient-weighted Markov random field fit to resting-state connectivity in 1,489 subjects. Two granularities are evaluated: Schaefer-100 (100 cortical parcels) and Schaefer-400 (400 parcels). Subcortical coverage is added through the Tian S1 scale atlas (Tian et al., 2020), which provides 16 subcortical regions defined on functional gradients of the basal ganglia, thalamus, hippocampus, and amygdala. The two combinations used in this thesis are therefore Schaefer-100 + Tian S1 ($N = 116$ regions) and Schaefer-400 + Tian S1 ($N = 416$

regions). The atlas selected per dataset is dictated by the Stage 2 sensitivity analysis (see Section 3.3) and fixed once for Stage 3.

2.4.3 Functional Connectivity Estimation

For a given session, the parcel-level time courses $x_i(t) \in \mathbb{R}^T$ $(i = 1, \ldots, N)$ are pairwise correlated using the Pearson product-moment coefficient (Pearson, 1895):

$$\mathrm{FC}_{ij} = \frac{\sum_t \big(x_i(t) - \bar{x}_i\big)\big(x_j(t) - \bar{x}_j\big)}{\sqrt{\sum_t \big(x_i(t) - \bar{x}_i\big)^2} \cdot \sqrt{\sum_t \big(x_j(t) - \bar{x}_j\big)^2}}, \tag{2.1}$$

which yields a symmetric matrix $\mathrm{FC} \in [-1, 1]^{N \times N}$ with unit diagonal. The diagonal entries are self-correlations of value 1 and carry no discriminative information, while the symmetry of the matrix means that the lower and upper triangles are mirror images of each other. Only the upper triangle is therefore retained, flattened into a one-dimensional vector $v \in \mathbb{R}^F$ with

$$F = \binom{N}{2} = \frac{N(N-1)}{2}, \qquad v = \big(\mathrm{FC}_{ij}\big)_{1 \le i < j \le N}. \tag{2.2}$$

For Schaefer-400 + Tian ($N = 416$) this yields $F = 86{,}320$ features per session per subject; for Schaefer-100 + Tian ($N = 116$) the same operation gives $F = 6{,}670$ features.

2.4.4 Quality Control and Cohort Filtering

A subject is retained only when both REST1 and REST2 connectivity matrices load successfully and pass shape and finiteness checks for both atlases in use. The intersection across atlases enforces an identical subject pool across granularities, so that the train / test split is shared across all experiments. For HCP-YA this filtering yields 1,018 subjects; for HCP-Aging, 720 subjects.

2.4.5 Subject Level Train / Test Split

An 80% / 20% split is drawn once per dataset at the subject level using a fixed random seed (seed = 42), producing 814 training and 204 test subjects for HCP-YA, and 576 training and 144 test subjects for HCP-Aging. Splitting at the subject level (rather than session or

scan level) is essential: both REST1 and REST2 of the same subject must fall on the same side of the partition, or the model can encounter test identities during training and inflate performance through identity leakage. The same indices are reused across atlases and across all three stages, so every comparison in the results chapter is computed on the same set of 204 (or 144) test identities.

2.5 Stage 1: Pearson Correlation Baseline

2.5.1 Identification Protocol

The first stage replicates the connectome fingerprinting protocol of Finn et al. (2015) on the train / test split defined above. No model is fitted. For each REST1 probe vector $v_i^{(1)}$ in the test set, a similarity score is computed against every REST2 gallery vector $v_j^{(2)}$, and the predicted identity is the index of the most similar gallery entry. This anchors the thesis to the literature it builds on and establishes a zero training reference point against which the supervised stages are interpreted.

2.5.2 Vectorised Computation

Pearson correlation between two vectors is equivalent to the cosine similarity of their z-scored versions, which allows the full $N_{\text{test}} \times N_{\text{test}}$ score matrix to be computed in a single matrix product. Defining the row wise standardisation

$$\tilde{v} = \frac{v - \mu_v}{\sigma_v + \varepsilon}, \qquad \varepsilon = 10^{-8}, \tag{2.3}$$

where μ_v and σ_v are the empirical mean and standard deviation of v and ε is a small numerical stability constant, the score matrix is

$$S = \frac{1}{F}\,\tilde{P}\,\tilde{G}^{\top}, \qquad \tilde{P}, \tilde{G} \in \mathbb{R}^{N_{\text{test}} \times F}, \tag{2.4}$$

with $\tilde{P}$ stacking the z-scored REST1 probe vectors and $\tilde{G}$ the z-scored REST2 gallery vectors. The predicted identity for probe i is then

$$\hat{y}_i = \arg\max_j \; S_{ij}. \tag{2.5}$$

2.6 Stage 2: Classical Machine Learning Baselines

The second stage formulates fingerprinting as a pairwise verification task. Five classifiers are trained on labelled pairs of compressed connectivity vectors and asked to predict whether the two vectors originate from the same subject. At test time, the classifiers are used as scoring functions inside the same probe / gallery argmax protocol as in Stage 1, so the comparison with the Pearson baseline is strict.

2.6.1 Preprocessing: Standardisation and PCA Compression

The raw 86,320 dimensional feature vector of the Schaefer-400 atlas is prohibitively large for kernel methods, and combined with the 814 training subjects of HCP-YA it sits firmly in the curse of dimensionality regime. Two sequential preprocessing steps reduce it to a manageable size while preventing information leakage between the splits.

Standardisation. Each feature dimension is normalised to zero mean and unit variance using a `StandardScaler` fitted on the concatenation of the REST1 and REST2 training vectors. The scaler is then applied to both sessions of all subjects. Fitting on the training set only ensures that the test identities never influence the per feature statistics; using both sessions to fit the scaler makes the standardisation symmetric between probe and gallery.

PCA dimensionality reduction. Principal Component Analysis (Hotelling, 1933) projects the standardised vectors onto the k directions of maximum variance:

$$z = W^{\top}(v - \mu), \qquad W \in \mathbb{R}^{F \times k}, \tag{2.6}$$

where the columns of W are the top-k eigenvectors of the training data covariance matrix. Three values $k \in \{300, 400, 500\}$ form a sensitivity analysis on the compression level, with $k = 400$ retained as the primary operating point based on the empirical performance reported in Chapter 3. The PCA basis is fitted on the REST1 training vectors only, following the convention of the fingerprinting literature (Finn et al., 2015; Shojaee et al., 2019): this prevents the basis from being shaped by the REST2 distribution before the model is trained, and keeps the REST2 gallery a clean held out signal at test time.

2.6.2 Pairwise Feature Construction

Given two compressed connectivity vectors $z_i^{(1)}$ and $z_j^{(2)}$, the classifier input is the element wise absolute difference

$$x_{ij} = \left| z_i^{(1)} - z_j^{(2)} \right| \in \mathbb{R}^k_{\geq 0}, \qquad y_{ij} = \mathbb{1}[i = j]. \tag{2.7}$$

The absolute value is taken so that the representation is symmetric: swapping the two sessions does not change x_{ij}, which mirrors the symmetry of the underlying verification problem. All 814 same subject training pairs are kept as positives ($y = 1$). Negatives ($y = 0$) are sampled from the set of different subject pairs with a fixed negative to positive ratio of $r = 10$, giving 8,140 negatives per training run. Class weights are balanced during training to compensate for this imbalance.

2.6.3 Linear Discriminant Analysis

Linear Discriminant Analysis (Fisher, 1936) models each class as a multivariate Gaussian with class specific mean and a shared covariance matrix. The optimal projection direction maximises Fisher's criterion, the ratio of between class scatter to within class scatter:

$$w^* = \arg\max_w \frac{w^\top S_B\, w}{w^\top S_W\, w}, \tag{2.8}$$

with $S_B = (\mu_1 - \mu_0)(\mu_1 - \mu_0)^\top$ and $S_W = \sum_{c \in \{0,1\}} \sum_{x \in \mathcal{C}_c} (x - \mu_c)(x - \mu_c)^\top$. Because LDA pools all training samples to estimate class means and the shared covariance, it is less sensitive to noise near the decision boundary than discriminative margin based methods.

2.6.4 Linear Support Vector Machines

The linear SVM (Cortes and Vapnik, 1995) finds a hyperplane $w^\top x + b = 0$ that maximises the margin between the two classes, with slack variables tolerating soft errors:

$$\min_{w,b,\xi} \ \tfrac{1}{2}\|w\|^2 + C \sum_{i=1}^{n} \xi_i \quad \text{s.t.} \quad y_i \left(w^\top x_i + b\right) \geq 1 - \xi_i, \quad \xi_i \geq 0. \tag{2.9}$$

Two settings of the regularisation parameter are evaluated: $C = 1.0$ (default margin width) and $C = 0.1$ (stronger L2 penalty on $\|w\|$). The second variant tests whether constraining the model further helps in the high dimensional pairwise difference space.

2.6.5 RBF Kernel SVM

The kernel SVM (Boser et al., 1992) replaces the linear inner product with a Gaussian radial basis function:

$$K(x_i, x_j) = \exp\Big(-\gamma \, \|x_i - x_j\|^2\Big), \tag{2.10}$$

with $\gamma = 1/(k \cdot \mathrm{Var}(X))$ (the `scale` setting in scikit-learn), which adapts to the feature count and overall variance automatically. Probability calibration is enabled via Platt scaling (Platt, 1999), so the output is a calibrated posterior $P(y = 1 \mid x)$ that serves as the matching score at identification time.

2.6.6 AdaBoost

AdaBoost (Freund and Schapire, 1997) trains a sequence of weak classifiers $h_t(x)$ ($t = 1, \ldots, T$), each fitted on a reweighted version of the training set that emphasises the examples misclassified by previous learners. The ensemble output is a weighted majority vote:

$$H(x) = \mathrm{sign}\left(\sum_{t=1}^{T} \alpha_t \, h_t(x)\right), \qquad \alpha_t = \tfrac{1}{2} \ln \frac{1 - \epsilon_t}{\epsilon_t}, \tag{2.11}$$

where ϵ_t is the weighted error of h_t. The ensemble size is fixed at $T = 100$ decision stumps, which is the canonical default and produces a stable bias / variance tradeoff for the dimensionalities considered here.

2.6.7 Hyperparameter Summary

Table 2.1 consolidates the hyperparameters used across all five classifiers. Defaults are deliberately preferred over a fully blown hyperparameter search at this stage: the role of the classical baselines is to characterise how much can be recovered from off the shelf models, against which the Stage 3 deep methods can be judged.

Parameter	Value	Justification
`N_COMPONENTS` (PCA)	300 / 400 / 500	Sensitivity sweep; $k = 400$ chosen as primary
`NEG_RATIO`	10	Balances pair distribution; bounds the memory footprint
C (Linear SVM)	1.0	Default soft margin width
C (L2 SVM)	0.1	Stronger penalty to test overfitting hypothesis
γ (RBF)	`scale`	Adapts to feature count and variance automatically
T (AdaBoost)	100	Standard ensemble size for stumps
Random seed	42	Ensures full reproducibility of split and pair sampling

Table 2.1: Hyperparameters for the Stage 2 classical ML experiments.

2.6.8 Unified Identification Protocol

At test time the classifiers are reused inside the same probe / gallery argmax loop introduced in Stage 1. For each REST1 probe subject i, a difference vector x_{ij} is constructed against every REST2 gallery j and the classifier produces a score: the calibrated posterior $P(y = 1 \mid x_{ij})$ for the RBF SVM, and the signed decision value $f(x_{ij})$ for Linear SVM and LDA. AdaBoost is scored through the unnormalised ensemble margin $\sum_t \alpha_t h_t(x_{ij})$. The predicted identity is again the gallery index with the highest score (Eq. 2.5), so that Stage 1 and Stage 2 share an identical evaluation harness.

2.7 Stage 3: Deep Siamese Architectures

Stage 3 replaces the pairwise verification setup of Stage 2 with a learnt embedding space. A neural encoder f_θ maps each connectivity representation into a low dimensional embedding such that two sessions of the same subject are close, and sessions of different subjects are far apart. Three encoder variants are studied: a multi layer perceptron (Method 1), a graph convolutional network (Method 2), and a graph attention network (Method 3). The three methods share the same Siamese twin framework, the same InfoNCE contrastive objective, the same training loop, and the same evaluation protocol, so that any difference in performance is attributable to the encoder architecture.

2.7.1 Siamese Twin Framework

A Siamese network (Bromley et al., 1993; Chopra et al., 2005) comprises two encoder branches that share weights. Given two inputs u and v, the network produces two embeddings $z_u = f_\theta(u)$ and $z_v = f_\theta(v)$, and a similarity head scores the pair. Weight sharing makes the two branches functionally identical, which has two consequences. The mapping is symmetric: swapping the inputs swaps the embeddings without changing the similarity. And the entire network can be optimised with one set of parameters, halving the gradient footprint compared to two independent encoders.

The similarity head is fixed across all three methods. The encoder output is L_2 normalised onto the unit hypersphere, so that the inner product between two embeddings equals their cosine similarity:

$$\mathrm{sim}(z_u, z_v) = \frac{z_u^\top z_v}{\|z_u\|\,\|z_v\|} = z_u^\top z_v \quad \text{when} \quad \|z_u\| = \|z_v\| = 1. \tag{2.12}$$

Cosine similarity is a natural choice for fingerprinting because it captures the geometry of the embedding direction while ignoring its magnitude.

2.7.2 The InfoNCE Contrastive Objective

The training objective is the InfoNCE loss (van den Oord et al., 2018), a contrastive objective popularised in self supervised representation learning (Chen et al., 2020). For a mini batch of B subjects, each subject contributes the two embeddings of its two sessions, $z_i^{(1)}$ and $z_i^{(2)}$. For row i, the positive comparison is the pair of its own two sessions and the $B-1$ negatives are the second session embeddings of all other subjects in the batch:

$$\mathcal{L}_{\rightarrow} = -\frac{1}{B}\sum_{i=1}^{B} \log \frac{\exp\!\big({z_i^{(1)}}^\top z_i^{(2)}/\tau\big)}{\sum_{j=1}^{B} \exp\!\big({z_i^{(1)}}^\top z_j^{(2)}/\tau\big)}, \tag{2.13}$$

where $\tau > 0$ is a temperature that controls how sharply the softmax distinguishes the positive from the negatives. The loss is then symmetrised by repeating the same computation with the roles of $z^{(1)}$ and $z^{(2)}$ swapped, and averaging:

$$\mathcal{L}_{\text{InfoNCE}} = \tfrac{1}{2}\big(\mathcal{L}_{\rightarrow} + \mathcal{L}_{\leftarrow}\big). \tag{2.14}$$

The temperature is fixed at $\tau = 0.09$, and the batch size is $B = 64$ subjects, giving $B(B-1) = 4{,}032$ implicit negative comparisons per batch without ever materialising an explicit pair list. This batch level contrastive structure is what enables InfoNCE to scale to thousands of identities without combinatorial pair generation (Chen et al., 2020).

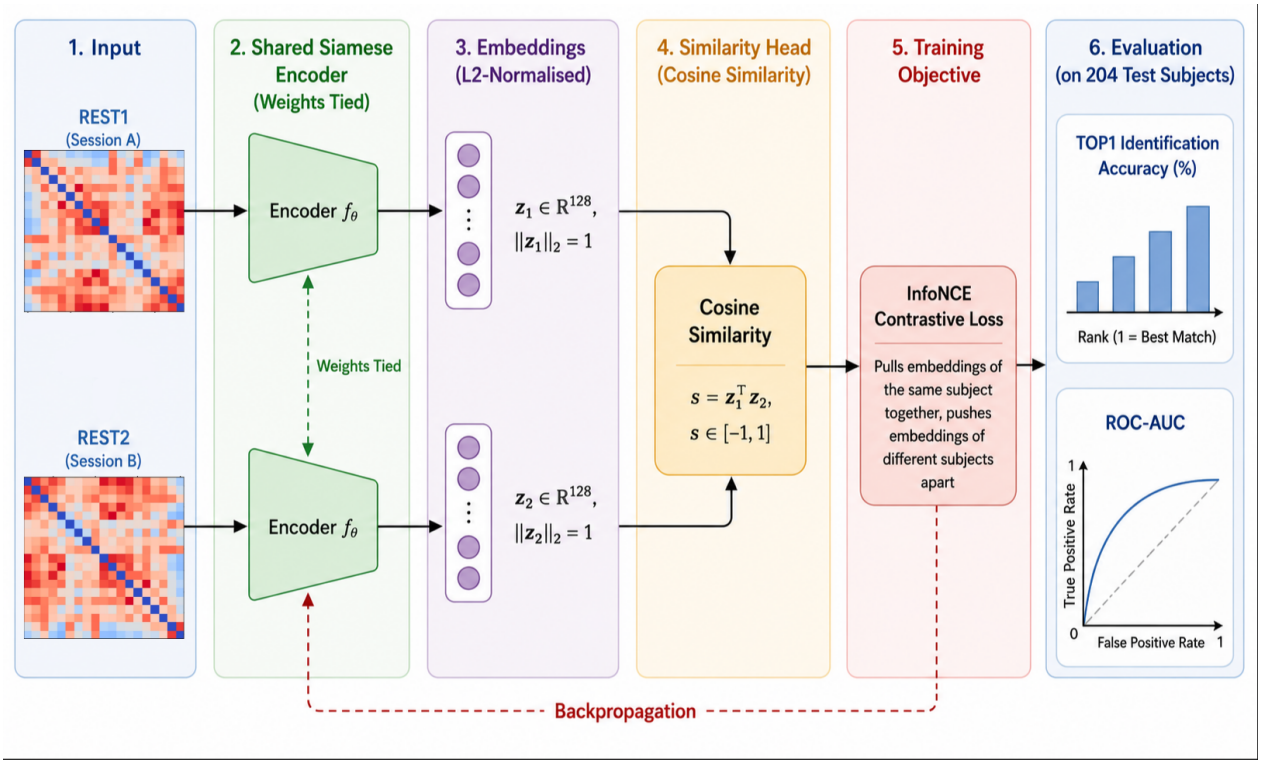


Figure 2.2: Overview of the Stage 3 Siamese deep learning framework. Two sessions (REST1 and REST2) are passed through a shared encoder f_θ (MLP, GCN, or GAT), each mapped to an L2-normalised 128-dimensional embedding. A cosine similarity head scores each pair, and the InfoNCE contrastive loss pulls same-subject embeddings together while pushing different-subject embeddings apart. At test time, TOP-1 accuracy and ROC-AUC are evaluated on 204 held-out subjects (HCP-YA) or 144 (HCP-Aging).

2.7.3 Method 1: Siamese MLP

The first deep method uses a multi layer perceptron as the encoder (Sarar et al., 2021; Lee and Lee, 2024). The input is the PCA compressed connectivity vector of Section 2.6.1, which keeps the input dimensionality of the MLP directly comparable to the Stage 2 classifiers and prevents the first linear layer from containing tens of millions of weights. The PCA target k is fixed per dataset at the best Stage 2 operating point ($k = 400$ for HCP-YA, $k = 500$ for HCP-Aging).

The encoder is a three layer feed forward network. The first two blocks each apply a linear transformation, batch normalisation (Ioffe and Szegedy, 2015), a ReLU activation, and dropout at rate $p = 0.1$, projecting the input first to a 512 dimensional hidden space and then refining it within that hidden space. The third layer is a linear projection to a 128 dimensional embedding followed by L_2 normalisation:

$$\begin{aligned} h_1 &= \mathrm{Dropout}\big(\mathrm{ReLU}(\mathrm{BN}(W_1 z))\big), \\ h_2 &= \mathrm{Dropout}\big(\mathrm{ReLU}(\mathrm{BN}(W_2 h_1))\big), \\ f_\theta(z) &= \frac{W_3 h_2}{\|W_3 h_2\|_2}, \end{aligned} \tag{2.15}$$

with $W_1 \in \mathbb{R}^{512\times k}$, $W_2 \in \mathbb{R}^{512\times 512}$, and $W_3 \in \mathbb{R}^{128\times 512}$. The embedding dimension of 128 is chosen as a compromise between expressivity and the number of test identities to discriminate (204 for HCP-YA, 144 for HCP-Aging): a 128 dimensional unit hypersphere has ample volume to separate a few hundred points while remaining cheap to train.

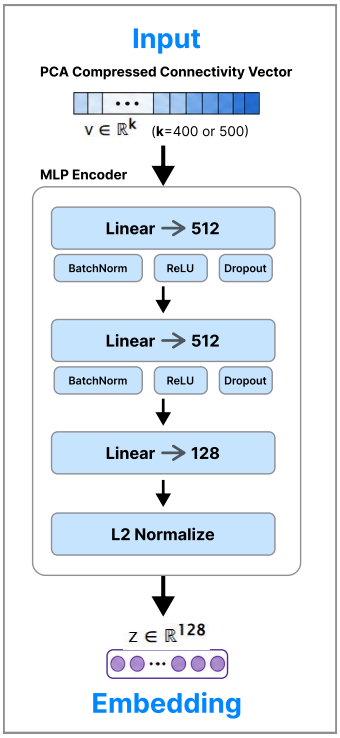


Figure 2.3: Siamese MLP encoder. Input: PCA-compressed vector $v \in \mathbb{R}^k$ ($k = 400$ for HCP-YA, $k = 500$ for HCP-Aging). Two hidden layers project to 512 dimensions (linear $\to$ BN $\to$ ReLU $\to$ dropout); a third layer projects to 128 dimensions followed by L2 normalisation.

2.7.4 Functional Connectivity Graph Construction

The two graph based methods (Sections 2.7.5 and 2.7.6) operate on graph objects rather than flat vectors. Each connectivity matrix is therefore converted into a graph $G = (\mathcal{V}, \mathcal{E}, X, W)$, in which the N brain regions are the nodes and the edges encode the strongest functional links between them. Graph construction is a one off preprocessing step run once per subject and per session, with the resulting graphs serialised to disk so that the GCN and GAT models can iterate over them without paying the construction cost on every epoch.

Node features. The feature vector of node i is the i-th row of the connectivity matrix with the diagonal element removed, giving an $(N-1)$-dimensional connectivity profile that describes how region i relates to every other region in the brain. Removing the diagonal removes the constant self correlation of value 1, which carries no discriminative information. The resulting node feature matrix is $X \in \mathbb{R}^{N \times (N-1)}$.

Edges (sparse k-NN on the absolute connectivity). The fully dense graph implied by the connectivity matrix is uninformative for message passing because every node is connected to every other node. The graph is therefore sparsified through a k-nearest neighbour rule on the absolute Pearson correlation:

$$\mathcal{N}_k(i) = \underset{j \neq i}{\arg\,\text{top-k}} \; |\mathrm{FC}_{ij}|, \tag{2.16}$$

with $k = 10$. The absolute value is taken because strong anticorrelations are as informative for fingerprinting as strong positive correlations: both indicate a tight functional relationship between two regions, whether synchronised or antagonistic. The raw k-NN graph is directed (region i may be in the top-k of region j without the reverse), so it is symmetrised by taking the union of (i, j) and (j, i) for every retained edge. After symmetrisation each node has at least $k = 10$ neighbours, with hub regions accumulating more.

Edge weights. Each retained edge (i, j) carries a scalar weight $w_{ij} = |\mathrm{FC}_{ij}|$, which the GCN consumes directly inside its message passing rule. The GAT discards these weights and learns its own edge importance through attention; the edge index is still required, since attention is computed only over the sparsified neighbourhood.

Normalisation. The symmetric normalisation $\hat{A} = D^{-1/2} A D^{-1/2}$ that the GCN rule expects (Kipf and Welling, 2017) is applied on the fly inside the convolutional layer of PyTorch Geometric (Fey and Lenssen, 2019), so only the raw symmetrised edge index and the unnormalised edge weights are stored on disk.

2.7.5 Method 2: Siamese GCN

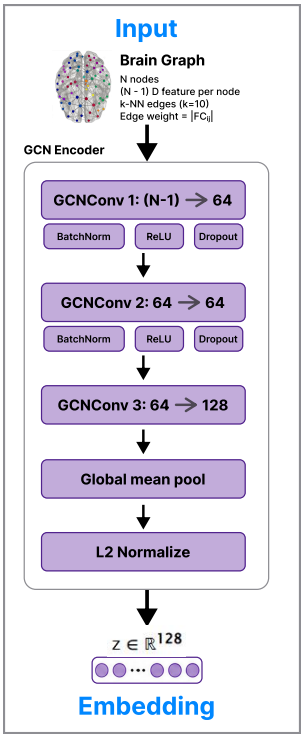


Figure 2.4: Siamese GCN encoder. Input: brain graph with N nodes, each carrying an $(N-1)$-dimensional connectivity profile, connected via a k-NN edge structure ($k = 10$) with edge weights $|\mathrm{FC}_{ij}|$. Three graph convolutional layers perform weighted neighbourhood aggregation: $(N-1) \to 64 \to 64 \to 128$ dimensions. Global mean pooling followed by L2 normalisation yields the final embedding $z \in \mathbb{R}^{128}$.

The Graph Convolutional Network (Kipf and Welling, 2017) replaces the dense linear layer of the MLP with a layer that aggregates information along the edges of each subject's connectivity graph. A single GCN layer with weights $\Theta \in \mathbb{R}^{d_{\text{in}} \times d_{\text{out}}}$ applies the propagation rule

$$H^{(\ell+1)} = \sigma\Big(\tilde{D}^{-1/2}\,\tilde{A}\,\tilde{D}^{-1/2}\,H^{(\ell)}\,\Theta^{(\ell)}\Big), \tag{2.17}$$

where $\tilde{A} = A + I_N$ is the adjacency matrix with added self loops, $\tilde{D}$ is the corresponding degree matrix, and σ is a non linear activation. The symmetric normalisation $\tilde{D}^{-1/2}\,\tilde{A}\,\tilde{D}^{-1/2}$ rescales each edge by the geometric mean of the degrees of its two endpoints, which prevents high degree nodes from dominating the aggregation. In the weighted variant used here, the entries of A are the edge weights $w_{ij} = |\mathrm{FC}_{ij}|$ produced by the graph builder, so the layer effectively performs a weighted neighbour average followed by a shared linear projection.

The encoder stacks three GCN layers. The first projects the $(N-1)$ dimensional node features into a 64 dimensional hidden space; the second keeps the same width and broadens the receptive field to two hop neighbourhoods; the third projects to the final 128 dimensional embedding. Each of the first two layers is followed by a ReLU activation and dropout at rate $p = 0.1$. The per node embeddings are then aggregated into a single graph level embedding by global mean pooling, which averages the embeddings across the N regions of the brain. This pooled embedding is L_2 normalised and fed into the shared cosine head defined in Section 2.7.1.

The GCN therefore learns a representation in which each region's embedding is informed both by its own connectivity profile (through the node feature input) and by the connectivity profiles of its strongest neighbours (through the two hop receptive field). Compared to the MLP, the inductive bias is shifted from "any feature can interact with any other" to "features that share strong functional connections are processed together", which is closer to how the brain is believed to organise its functional networks.

2.7.6 Method 3: Siamese GAT with Attention Interpretability

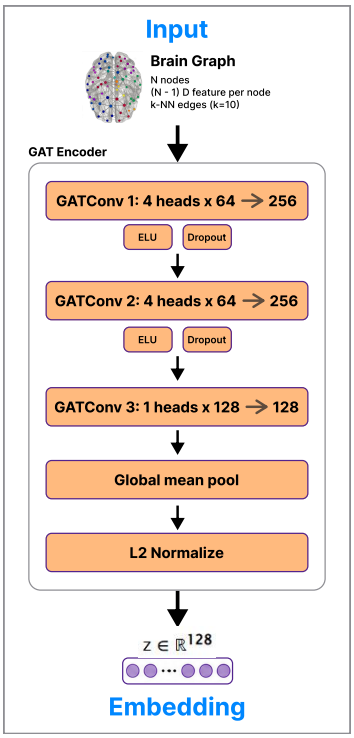


Figure 2.5: Siamese GAT encoder. Input: brain graph with N nodes, each carrying an $(N-1)$-dimensional connectivity profile, connected via a k-NN edge structure ($k = 10$). Three graph attention layers: layers 1–2 use 4-head attention (64 dims/head, concatenated to 256); layer 3 collapses to a single head projecting to 128 dimensions. Global mean pooling and L2 normalisation yield the final embedding $z \in \mathbb{R}^{128}$.

The Graph Attention Network (Veličković et al., 2018) replaces the fixed edge weighting of the GCN with attention coefficients learnt jointly with the encoder. For a node i with neighbour set $\mathcal{N}(i)$ (after the k-NN sparsification, plus a self loop that GATConv adds internally) and a single attention head with weight matrix Θ and attention vector a, the layer computes the unnormalised attention scores

$$e_{ij} = \mathrm{LeakyReLU}\Big(a^\top \left[\Theta h_i \,\|\, \Theta h_j\right]\Big), \quad j \in \mathcal{N}(i), \tag{2.18}$$

where $\|$ denotes concatenation of the two transformed node features. These scores are softmax normalised across the neighbours of i:

$$\alpha_{ij} = \frac{\exp(e_{ij})}{\sum_{j' \in \mathcal{N}(i)} \exp(e_{ij'})}, \tag{2.19}$$

so that the attention weights from i's neighbours sum to one. The output embedding for node i is a weighted sum of its neighbours' transformed features:

$$h_i' = \sigma\Big(\sum_{j \in \mathcal{N}(i)} \alpha_{ij}\, \Theta h_j \Big). \tag{2.20}$$

The full encoder uses three GATConv layers. The first two apply multi head attention with four heads, each producing a 64 dimensional projection that is concatenated across heads to give a 256 dimensional hidden representation. The third layer collapses the multi head representation back to a single head that outputs the 128 dimensional final embedding. Following the original GAT paper, the non linearity is ELU rather than ReLU, which gives a smoother gradient near zero and stabilises the attention computation when many of the unnormalised scores cluster around the same value. Dropout at $p = 0.1$ is applied both on the node features and inside the attention mechanism, which in the latter case has the effect of randomly dropping a fraction of the attention edges during training. The per node embeddings are pooled and normalised in the same way as for the GCN.

The two main differences between Method 2 and Method 3 are therefore the weighting and the heads. The GCN's weighted aggregation uses $|\mathrm{FC}_{ij}|$, which is fixed by the data and the same in every layer. The GAT learns α_{ij} from the data, independently in each layer and each head, and is therefore free to assign weights that are not proportional to the raw connectivity strength. This makes the attention coefficients interpretable in their own right: after training, the first layer attention can be averaged over heads and subjects, then collapsed onto the seven canonical resting state networks of Yeo et al. (2011) to produce a 7×7 attention matrix that summarises which network pairs the model relies on for identification.

2.7.7 Common Training Protocol

The three deep methods share the same optimisation setup so that comparisons between them isolate the encoder. A 10% subject level validation split is peeled off the training set ($N_{\mathrm{val}} = 81$ for HCP-YA, 57 for HCP-Aging) and used to monitor the InfoNCE loss after every epoch. The split is performed at the subject level, so both sessions of a held out subject fall in the validation set together. Optimisation uses Adam (Kingma and Ba, 2015) with a learning rate of 10^{-4} for the MLP and the GCN, and 5×10^{-4} for the GAT (chosen to compensate for the smaller gradient magnitudes that the attention softmax produces early in training). Weight decay is set to 10^{-5} in all three cases as a mild L2 regulariser. A cosine

annealing schedule (Loshchilov and Hutter, 2017) drives the learning rate from its initial value down to 10^{-6} over a maximum of 200 epochs. Training stops early if the validation loss fails to improve for 20 consecutive epochs, and the model weights are rolled back to the best epoch before evaluation. The random seed is fixed at 42 throughout (split generation, weight initialisation, batch shuffling, attention dropout) for reproducibility.

2.7.8 Test Time Evaluation Protocol

At evaluation time the trained encoder is run once over the REST1 test sessions and once over the REST2 test sessions, producing two embedding matrices $Z^{(1)}, Z^{(2)} \in \mathbb{R}^{N_{\text{test}} \times 128}$. Because the encoder ends in L_2 normalisation, the full $N_{\text{test}} \times N_{\text{test}}$ cosine similarity matrix is the plain matrix product $S = Z^{(1)} {Z^{(2)}}^{\top}$ and the predicted identity for probe i is $\hat{y}_i = \arg\max_j S_{ij}$, exactly as in Eq. 2.5. The same protocol is therefore used end to end across the three stages.

2.8 Evaluation Metrics

Two quantities summarise the performance of every method evaluated in the thesis.

TOP1 identification accuracy. The primary metric, directly comparable to Finn et al. (2015), counts the fraction of probe subjects whose own REST2 session ranks first among all gallery candidates:

$$\text{Acc} = \frac{1}{N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} \mathbb{1}\big[\hat{y}_i = i\big]. \tag{2.21}$$

Verification ROC-AUC. The same similarity matrix is also interpreted as a binary verification problem: each of the N_{test}^2 probe / gallery pairs is labelled positive when the two indices coincide (diagonal entries) and negative otherwise (off diagonal entries). The area under the resulting Receiver Operating Characteristic curve summarises the discrimination ability across all thresholds, and gives a metric that is insensitive to the precise scaling of the score function.

2.9 Conclusion

This chapter has detailed the complete methodological pipeline underlying the three-stage experimental framework introduced in this thesis. Starting from raw resting-state fMRI data, Pearson correlation matrices were derived using two atlas resolutions (Schaefer-100 + Tian and Schaefer-400 + Tian) and a strict subject-level train/test split was enforced to prevent identity leakage across all stages.

Stage 1 established a zero-training Pearson correlation baseline by replicating the original fingerprinting protocol of Finn et al. (2015), providing a reference point for all subsequent comparisons. Stage 2 reformulated identification as a supervised pairwise verification task, applying five classical classifiers over a PCA-compressed feature space and evaluating their sensitivity to atlas resolution and dimensionality. Stage 3 extended this to a deep contrastive learning paradigm using Siamese architectures with three encoder variants—MLP, GCN, and GAT—all trained with the InfoNCE objective and evaluated under the same probe/gallery argmax protocol.

Critically, the three stages share an identical cohort, split, and evaluation harness, so that any observed performance differences across methods are attributable solely to modelling capacity and architectural inductive bias. Performance is quantified using two metrics: TOP-1 identification accuracy and verification ROC-AUC. The results of applying this pipeline to both the HCP-YA and HCP-Aging cohorts are presented and analysed in the following chapter.

3

Chapter 3

Results and Discussion

Chapter 3

Results and Discussion

3.1 Introduction

This chapter presents and analyses the empirical results obtained by applying the three-stage experimental framework described in Chapter 2 to the HCP-Young Adult (HCP-YA) and HCP-Aging datasets. The evaluation is structured to answer four interconnected questions: (i) how does a zero-training Pearson correlation baseline perform, and does it motivate more expressive methods; (ii) how do supervised classical machine learning classifiers behave on each cohort individually; (iii) how do deep Siamese graph architectures compare to classical methods and to each other; and (iv) how does atlas resolution (coarse: Schaefer-100 + Tian, $N = 116$; versus fine: Schaefer-400 + Tian, $N = 416$) interact with modelling capacity across all stages.

The chapter is organised as follows. Section 3.2 establishes the Pearson correlation baseline (Stage 1) for both cohorts and contextualises it against published fingerprinting results. Section 3.3 presents the classical machine learning results (Stage 2), analysed separately for HCP-YA and HCP-Aging before a cross-cohort comparison. Section 3.4 presents the deep Siamese results (Stage 3), following the same dataset-by-dataset structure. Section 3.5 then synthesises the effect of atlas resolution across all methods and both datasets. Section 3.6 examines the neurobiological interpretability afforded by the GAT attention maps. Section 3.7 consolidates the principal findings across all stages.

3.2 Stage 1: Pearson Correlation Baseline

The Pearson correlation baseline requires no training: for each probe subject, the cosine similarity between its vectorised FC matrix and every gallery FC matrix is computed, and identity is assigned by nearest-neighbour matching. This replicates the protocol of Finn et al. (2015) and provides a reference point against which all supervised methods are compared.

3.2.1 HCP-Young Adult

On HCP-YA, the Pearson baseline achieves a TOP-1 accuracy of 68.14% at Atlas 100 and 87.75% at Atlas 400, with ROC-AUC values of 93.8% and 97.6% respectively. The Atlas 100 figure is substantially lower than the 92.9%–94.4% reported by Finn et al. (2015) on 126 HCP participants with a 268-region atlas. Two structural factors account for most of this gap. First, the test gallery here contains 204 subjects, considerably larger than the 126 used by Finn et al. (2015), which increases the number of distractors and lowers TOP-1 by construction. Second, the Schaefer-100 parcellation ($N = 116$ with subcortical Tian regions) differs from the 268-region Shen atlas used in the original study, altering the feature space from which similarity is computed. The improvement at Atlas 400 reflects the richer connectivity profile available at finer resolution, and is discussed in detail in Section 3.5.

These results confirm that Pearson correlation can recover individual identity at a meaningful rate without any supervision. However, the 68.14% TOP-1 at Atlas 100 leaves a large identification gap, particularly given the availability of labelled training data, and motivates the use of supervised discriminative approaches in Stages 2 and 3.

3.2.2 HCP-Aging

On HCP-Aging, the Pearson baseline reaches 94.44% TOP-1 at Atlas 100 (ROC-AUC = 98.4%) and 97.92% at Atlas 400 (ROC-AUC = 99.7%). These figures are markedly higher than their HCP-YA counterparts. The HCP-Aging cohort spans a wide age range (36–100+ years) (Bookheimer et al., 2019), which is likely to produce greater inter-subject variability in functional connectivity and therefore a stronger fingerprinting signal (Horien et al., 2019). The smaller test pool (144 subjects versus 204 for HCP-YA) also reduces the number of distractors, inflating TOP-1 by construction. These two factors jointly explain the higher

baseline performance on HCP-Aging and should be held in mind when comparing absolute accuracy figures across cohorts throughout this chapter.

The Atlas 100 figure of 94.44% falls within the range reported by Finn et al. (2015), providing a sanity check that the preprocessing pipeline and evaluation protocol are consistent with established practice despite the different parcellation and cohort composition.

3.3 Stage 2: Classical Machine Learning

Stage 2 reformulates individual identification as a supervised pairwise verification task. Each probe–gallery pair is featurised as the element-wise absolute difference of their PCA-compressed FC vectors, and a binary classifier predicts whether the pair belongs to the same subject. Five classifiers are evaluated: Linear SVM, Linear SVM with L_2 regularisation, RBF SVM, Linear Discriminant Analysis (LDA), and AdaBoost, each at three PCA dimensionalities (300, 400, and 500 components).

3.3.1 Results on HCP-Young Adult

Table 3.1: Classical ML results on HCP-Young Adult (TOP-1 and ROC-AUC, %).

Atlas	PCA	Classifier	TOP-1 (%)	ROC-AUC (%)
Schaefer-100	300	Linear SVM	87.75	97.8
		Linear SVM + L_2 reg.	88.24	97.8
		RBF SVM	87.75	96.8
		LDA	90.20	97.6
		AdaBoost	46.10	92.3
	400	Linear SVM	87.75	97.7
		Linear SVM + L_2 reg.	89.22	97.9
		RBF SVM	89.71	97.3
		LDA	90.20	97.8
		AdaBoost	42.10	92.8

Continued on next page

Table 3.1 – continued from previous page

Atlas	PCA	Classifier	TOP-1 (%)	ROC-AUC (%)
	500	Linear SVM	81.37	97.2
		Linear SVM + L_2 reg.	83.33	97.6
		RBF SVM	89.71	97.3
		LDA	89.71	97.6
		AdaBoost	42.65	92.3
Schaefer-400	300	Linear SVM	64.71	94.6
		Linear SVM + L_2 reg.	68.63	94.6
		RBF SVM	87.75	93.5
		LDA	83.82	94.3
		AdaBoost	54.41	91.4
	400	Linear SVM	70.59	95.2
		Linear SVM + L_2 reg.	76.96	95.9
		RBF SVM	90.69	94.5
		LDA	92.16	95.8
		AdaBoost	51.47	90.4
	500	Linear SVM	77.45	95.5
		Linear SVM + L_2 reg.	82.35	95.8
		RBF SVM	79.90	94.5
		LDA	72.55	95.4
		AdaBoost	55.39	90.2

All five supervised classifiers substantially outperform the Stage 1 Pearson baseline of 68.14% at Atlas 100, confirming that learning a discriminative boundary from labelled training pairs extracts considerably more identity-relevant structure than unsupervised similarity matching. The best single result on HCP-YA is LDA at Atlas 400, PCA = 400, which achieves 92.16% TOP-1 and 95.8% ROC-AUC, a gain of over 24 percentage points relative to the Atlas 100 Pearson baseline.

LDA is the strongest performer overall on HCP-YA. Its optimisation criterion, maximising between-class scatter relative to within-class scatter, aligns naturally with the fingerprinting objective, and it benefits from the PCA pre-compression that reduces the scatter matrices to a tractable rank. However, LDA degrades at PCA = 500 (89.71% at Atlas 100, 72.55% at

Atlas 400), which is consistent with the within-class scatter matrix becoming ill-conditioned as dimensionality grows close to the number of training pairs.

RBF SVM is the most robust classifier across PCA settings on HCP-YA, maintaining 89.71% TOP-1 at both PCA = 400 and PCA = 500 at Atlas 100. Its kernel implicitly maps the FC difference vector into a high-dimensional space where a nonlinear boundary is more stable than a linear one, explaining its resistance to the instabilities that affect LDA and Linear SVM at higher PCA.

AdaBoost consistently underperforms all other classifiers, sometimes falling below 50% TOP-1 while retaining moderate ROC-AUC values around 90–92%. This dissociation reflects AdaBoost's reliance on an ensemble of weak decision stumps, which cannot partition the high-dimensional identity space precisely enough for nearest-neighbour identification even when they produce a coarse population-level ranking.

Placing these results in context, Shojaee et al. (2019) reported k-NN identification at 76.16% on 1,003 HCP young adults, with Decision Trees and Naïve Bayes at 55.51% and 46.97% respectively. The LDA result of 92.16% here exceeds those figures by a considerable margin, consistent with the stronger alignment of the LDA objective with pairwise verification. Hannum et al. (2023) reported LDA up to 99% on 865 HCP participants, a ceiling not matched here; however, that study did not enforce a held-out subject split of the present size, making the task structurally easier. Their observation that SVM accuracy ranged from 17% to 61% depending on preprocessing is qualitatively consistent with the sensitivity of Linear SVM to atlas resolution seen here.

3.3.2 Results on HCP-Aging

Table 3.2: Classical ML results on HCP-Aging (TOP-1 and ROC-AUC, %).

Atlas	PCA	Classifier	TOP-1 (%)	ROC-AUC (%)
		Linear SVM	88.19	98.1
		Linear SVM + L_2 reg.	91.67	98.5
	300	RBF SVM	97.22	99.0
		LDA	97.22	99.1

Continued on next page

Schaefer-100

Table 3.2 – continued from previous page

Atlas	PCA	Classifier	TOP-1 (%)	ROC-AUC (%)
		AdaBoost	82.64	98.0
	400	Linear SVM	85.42	98.4
		Linear SVM + L_2 reg.	94.44	98.7
		RBF SVM	97.92	99.0
		LDA	97.22	99.1
		AdaBoost	80.56	97.0
	500	Linear SVM	84.03	97.6
		Linear SVM + L_2 reg.	89.58	98.6
		RBF SVM	98.61	99.06
		LDA	98.61	99.07
		AdaBoost	78.47	97.5
Schaefer-400	300	Linear SVM	56.94	96.2
		Linear SVM + L_2 reg.	66.67	97.2
		RBF SVM	98.15	98.2
		LDA	95.14	98.1
		AdaBoost	80.56	97.3
	400	Linear SVM	61.81	96.0
		Linear SVM + L_2 reg.	63.19	96.7
		RBF SVM	94.44	98.1
		LDA	62.50	97.1
		AdaBoost	84.03	97.4
	500	Linear SVM	77.78	97.3
		Linear SVM + L_2 reg.	84.03	97.5
		RBF SVM	97.97	98.0
		LDA	29.17	91.0
		AdaBoost	83.33	97.6

Classical classifiers perform strongly on HCP-Aging, surpassing their HCP-YA counterparts across virtually every configuration. At Atlas 100, RBF SVM and LDA both reach 97.22%–98.61% TOP-1, with ROC-AUC consistently above 99%. This confirms that the stronger inter-subject variability of the aging cohort provides a richer identity signal that

even linear methods can exploit effectively.

The most notable anomaly on HCP-Aging is the sharp collapse of LDA at Atlas 400, PCA = 500, where TOP-1 falls to 29.17% despite a ROC-AUC of 91.0%. This is not a failure of the underlying fingerprinting signal: RBF SVM reaches 97.97% under the same conditions, but rather a numerical instability: at $k = 500$ components with the expanded Atlas 400 feature space, the within-class scatter matrix used in the LDA projection becomes ill-conditioned, and the resulting discriminant direction is unreliable. This pathological sensitivity to PCA dimensionality is a well-known limitation of LDA in high-dimensional settings (**?**) and underscores the importance of selecting PCA = 400 as the operating point for this classifier.

RBF SVM is the most reliable classifier on HCP-Aging, maintaining competitive performance across both atlases and all three PCA settings. Its robustness to the expanded dimensionality of Atlas 400, where linear methods degrade, stems from the kernel's implicit nonlinear mapping, which finds a separating boundary in a higher-dimensional space without being overwhelmed by the raw feature count.

AdaBoost performs considerably better on HCP-Aging than on HCP-YA, occasionally exceeding 80% TOP-1, which reflects the stronger and more linearly separable identity signal present in the aging cohort. Its ROC-AUC consistently exceeds 97%, suggesting that the ensemble correctly ranks same-subject pairs but still struggles to achieve the precision required for 1-vs-N identification.

3.3.3 Cross-Cohort Comparison

HCP-Aging consistently produces higher identification accuracy than HCP-YA across every classifier and atlas configuration. Three structural factors jointly explain this pattern and should not be conflated.

First, the aging cohort spans a wider age range (36–100+ years) compared to the narrow HCP-YA range (22–35 years) (Bookheimer et al., 2019; Van Essen et al., 2013). Greater demographic heterogeneity translates into greater inter-subject variability in functional connectivity, producing a stronger and more distinctive fingerprinting signal for all methods.

Second, the test pool for HCP-Aging contains 144 subjects versus 204 for HCP-YA. Because TOP-1 accuracy is a 1-vs-N task, a smaller gallery has fewer distractors, making correct

identification structurally easier and inflating absolute accuracy figures by construction.

Third, the two cohorts differ in scanner hardware and acquisition site: HCP-YA was acquired on a single customised Siemens Connectome Skyra, while HCP-Aging used harmonised Siemens Prisma scanners across four sites (Van Essen et al., 2013; Bookheimer et al., 2019). Differences in signal-to-noise ratio and spatial resolution may alter the effective quality of the FC estimates in ways that are difficult to disentangle from the demographic factors above.

These three factors act in concert, and the present design does not permit isolating their individual contributions. What can be said is that model *rankings* are stable across cohorts: LDA and RBF SVM lead in both datasets, AdaBoost trails in both, and the ordering of classifiers within each atlas-PCA configuration is broadly consistent. This stability confirms that the relative merits of the methods are not cohort-specific artefacts.

3.4 Stage 3: Deep Siamese Architectures

Stage 3 replaces the closed-set multi-class objective of classical ML with a contrastive metric learning paradigm. Three Siamese encoders are evaluated (MLP, GCN, and GAT), all trained with the InfoNCE loss (van den Oord et al., 2018), which pulls same-subject embeddings together and pushes different-subject embeddings apart in a 128-dimensional unit hypersphere. Identity is recovered at test time by the same nearest-neighbour argmax protocol used in Stages 1 and 2, ensuring a fair comparison across all stages.

An important interpretive note: ROC-AUC and TOP-1 measure different things for these models. ROC-AUC summarises discrimination ability across all thresholds and remains high even when the embedding space is only partially organised. TOP-1 measures whether the nearest neighbour is the correct subject, which requires the embedding to be precisely structured. Meaningful architectural differences therefore emerge more clearly in TOP-1 than in ROC-AUC.

3.4.1 Results on HCP-Young Adult

Table 3.3: Siamese deep learning results on HCP-YA (TOP-1 and ROC-AUC).

Atlas	Model	TOP-1 (%)	ROC-AUC
Schaefer-100	Siamese MLP	60.78	0.9752
	Siamese GCN	33.82	0.9287
	Siamese GAT	62.25	0.9755
Schaefer-400	Siamese MLP	71.08	0.9851
	Siamese GCN	55.39	0.9686
	Siamese GAT	88.24	0.9966

The GAT is the strongest deep architecture on HCP-YA at both atlas resolutions, reaching 62.25% at Atlas 100 and 88.24% at Atlas 400. The latter represents a gain of 16 percentage points over the MLP (71.08%) and 33 points over the GCN (55.39%) under the same conditions. The GAT's advantage lies in its use of learned, edge-specific attention weights: rather than aggregating all neighbourhood information equally as the GCN does, the GAT learns to upweight the connections most consistently distinctive across sessions and subjects, performing implicit feature selection over the graph topology.

The GCN underperforms even the MLP on HCP-YA at both atlases (33.82% versus 60.78% at Atlas 100; 55.39% versus 71.08% at Atlas 400). This is a notable result. The MLP receives the full upper-triangle FC matrix as input, providing all $N(N-1)/2$ pairwise connectivity values, while the GCN operates only on the k-nearest-neighbour subgraph ($k = 10$), discarding edges outside the neighbourhood before the encoder is even invoked. Information is therefore lost upstream of the model, and the GCN's isotropic aggregation cannot recover it. The GAT partially compensates by learning to attend selectively within the available edges, but the GCN has no such mechanism.

A noteworthy finding on HCP-YA is that the best classical result (LDA at Atlas 400, PCA = 400: 92.16%) exceeds the best deep learning result (GAT at Atlas 400: 88.24%). This inversion is unexpected but interpretable. Classical classifiers operate on the full compressed FC vector and are trained directly on a large labelled dataset of subject pairs from the same cohort. The Siamese framework is designed for open-set generalisation, learning a universal similarity metric, which comes at a cost in closed-set accuracy when training data are abundant and the cohort is homogeneous.

3.4.2 Results on HCP-Aging

Table 3.4: Siamese deep learning results on HCP-Aging (TOP-1 and ROC-AUC).

Atlas	Model	TOP-1 (%)	ROC-AUC
Schaefer-100	Siamese MLP	90.97	0.9978
	Siamese GCN	84.03	0.9931
	Siamese GAT	98.61	0.9996
Schaefer-400	Siamese MLP	87.50	0.9951
	Siamese GCN	87.50	0.9942
	Siamese GAT	100.00	1.0000

On HCP-Aging, all three architectures improve substantially relative to HCP-YA, and the GAT achieves perfect identification at Atlas 400 (TOP-1 = 100%, ROC-AUC = 1.0000). This is the highest result in the entire study. At Atlas 100, the GAT reaches 98.61%, while MLP and GCN both exceed 84%.

Crucially, the situation reverses relative to HCP-YA: the GAT now surpasses the best classical method (RBF SVM at 98.61%) at Atlas 100 and decisively at Atlas 400, where the best classical result is 98.15% versus the GAT's perfect score. This suggests that the graph-based architecture, with its topological inductive bias, is better positioned to exploit the richer inter-subject variability present in the aging cohort, particularly when paired with the finer Atlas 400 parcellation.

The GCN recovers substantially on HCP-Aging relative to its HCP-YA performance, reaching 84.03% at Atlas 100 and 87.50% at Atlas 400. The stronger identity signal in the aging cohort provides enough discriminative structure in the k-NN subgraph for the GCN to achieve competitive results even without the attention mechanism.

3.4.3 Cross-Cohort Comparison

Model rankings are consistent across cohorts: GAT > MLP > GCN at Atlas 100 on both datasets, and GAT dominates at Atlas 400 in both cases. The absolute performance gap between cohorts is explained by the same three factors identified in Section 3.3.3: wider demographic heterogeneity, smaller test pool, and scanner differences. The consistency of

rankings confirms that the relative merits of the architectures are genuine and not cohort-specific.

Placing these results against published benchmarks, Shojaee et al. (2019) reported 99.89% identification with a Siamese GCN on 1,003 HCP subjects, and Lu et al. (2024) achieved 99.7% using a Conditional Variational Autoencoder with Sparse Dictionary Learning on 997 HCP subjects. Both figures exceed the 88.24% GAT result on HCP-YA, but neither study enforced a held-out subject split of the present size, making the identification task structurally easier. The GAT result of 100% on HCP-Aging was obtained under the same rigorous held-out evaluation protocol used throughout this work, confirming that the architecture is capable of matching the highest figures in the literature when cohort conditions are favourable.

3.5 Effect of Atlas Resolution

Atlas resolution is a systematic experimental variable that interacts differently with each model family. This section brings together the resolution effects observed across all three stages and both datasets to draw unified conclusions.

Pearson baseline. The Pearson baseline gains from Atlas 400 on both cohorts: TOP-1 rises from 68.14% to 87.75% on HCP-YA and from 94.44% to 97.92% on HCP-Aging. A finer parcellation provides a longer FC vector with more pairwise connectivity values, increasing the informativeness of the cosine similarity score. Because the baseline requires no training, it is not penalised by the increased dimensionality in the way that linear classifiers are.

Classical machine learning (Stage 2). The effect of atlas resolution on classical classifiers is classifier-dependent and does not follow a single direction. Linear SVM degrades sharply at Atlas 400 on both cohorts: on HCP-YA at PCA = 400, TOP-1 drops from 87.75% to 70.59%. The upper triangle of a 400-region FC matrix contains roughly 16 times more features than the 100-region equivalent; without a sufficient increase in training data, a linear decision boundary becomes harder to place reliably in this expanded space. LDA behaves similarly, except at PCA = 400 on HCP-YA where it peaks at 92.16% at Atlas 400, where its discriminant projection handles the finer resolution better than a raw linear boundary because it explicitly reduces dimensionality along the most class-discriminative directions.

RBF SVM is the notable exception among classical methods: it maintains or improves performance at Atlas 400 across most configurations, reaching 90.69% on HCP-YA and 98.15% on HCP-Aging. The kernel implicitly maps the expanded FC vector into a higher-dimensional space where a nonlinear boundary is more stable, and it benefits from the additional connectivity structure that the finer atlas provides.

Deep Siamese architectures (Stage 3). Graph-based methods (GCN and GAT) consistently benefit from Atlas 400 while the MLP shows mixed behaviour. On HCP-Aging, the GCN improves from 84.03% to 87.50% and the GAT from 98.61% to 100.00% as resolution increases. This is the reverse of what linear classifiers exhibit. Graph encoders exploit the parcellation topology directly: a finer atlas provides more nodes and denser k-NN edges for the message-passing and attention mechanisms to operate on, enriching the geometric structure of the embedding space without requiring a linear decision boundary to scale with the feature dimension. The MLP, which treats the FC matrix as an unstructured vector like classical classifiers, does not share this advantage: on HCP-Aging it actually drops from 90.97% to 87.50% at Atlas 400, suggesting that the increased input dimensionality imposes a cost that the contrastive objective alone cannot fully absorb.

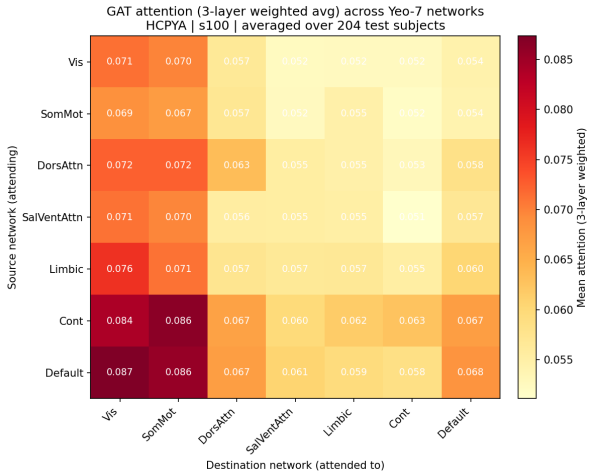
Summary. The resolution sensitivity of each method follows directly from its inductive biases. Methods that treat the connectome as a flat vector (linear classifiers, MLP) are hurt or unaffected by finer parcellations due to the curse of dimensionality. Methods that respect graph structure (GCN, GAT) benefit from finer parcellations because more nodes and edges provide a richer topological representation. This finding provides concrete methodological guidance: for future fingerprinting work using graph neural networks, finer atlas resolution is unambiguously beneficial; for linear or vector-based methods, the optimal resolution depends on the available training set size and should be treated as a hyperparameter.

3.6 Neurobiological Interpretability: GAT Attention Maps

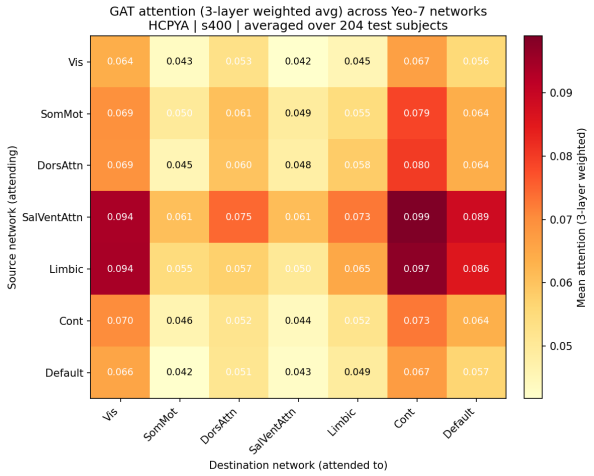
Beyond identification accuracy, the GAT architecture yields a qualitatively distinctive output: for each subject pair, each graph edge is assigned a learned attention coefficient that reflects how heavily the model weighted that functional connection for identity discrimina-

tion. Aggregating these coefficients across the test set and grouping edges by the Yeo-7 functional network membership of their endpoint nodes (Yeo et al., 2011) produces a 7×7 network-level attention matrix, where entry (i, j) represents the mean attention weight assigned to connections from source network i to destination network j, averaged across all test subjects, without any explicit supervision on network identity.

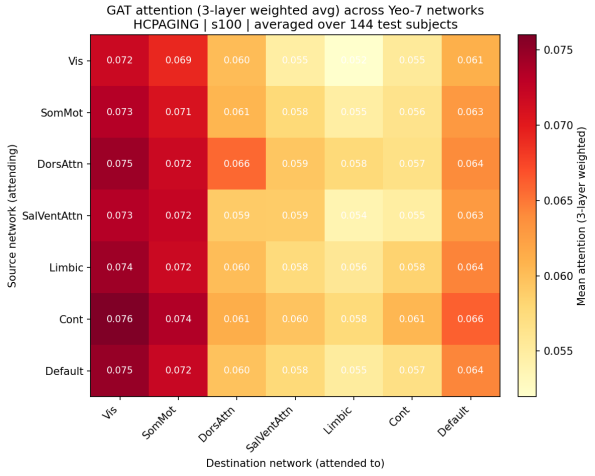
Figure 3.1 presents these matrices for all four experimental conditions: two cohorts (HCP-YA and HCP-Aging) by two atlas resolutions (Schaefer-100 and Schaefer-400).



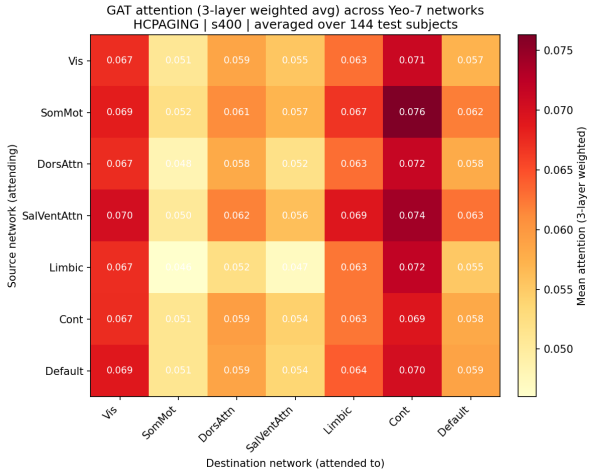
(a) HCP-YA, Schaefer-100



(b) HCP-YA, Schaefer-400



(c) HCP-Aging, Schaefer-100



(d) HCP-Aging, Schaefer-400

Figure 3.1: Mean GAT attention weights across Yeo-7 functional networks for all four experimental conditions. Each cell (i, j) is the 3-layer weighted average attention coefficient for edges from source network i to destination network j, averaged over all test subjects. Darker red indicates higher attention. At Schaefer-100, the Visual and SomMotor columns dominate in both cohorts; at Schaefer-400, the Control (frontoparietal) column becomes the primary attended destination, corroborating the differential power hierarchy of Finn et al. (2015).

Several patterns emerge clearly from Figure 3.1.

At Schaefer-100, the Visual and SomMotor networks consistently receive the highest attention as destination networks across both cohorts. In HCP-YA (panel a), the Control and Default networks are the strongest attending sources, with the Control×Vis (0.084)

and Default×Vis (0.087) cells reaching the highest values in the matrix. In HCP-Aging (panel c), a similar left-column dominance is observed, with DorsAttn×Vis (0.075), Cont×Vis (0.076) and Default×Vis (0.075) among the highest entries. This pattern suggests that at coarse resolution, the most identity-distinctive information the model can access is encoded in connectivity between higher-order networks and the Visual cortex, likely because sensory regions are anatomically compact and reliably parcellated even at low resolution.

At Schaefer-400, the pattern shifts markedly. In HCP-Aging (panel d), the Control column becomes the dominant attended destination across virtually all source networks, with SomMot×Cont (0.076) reaching the peak value in the matrix. In HCP-YA (panel b), the attention structure is more distributed, with SalVentAttn and Limbic emerging as strongly attending sources (0.094–0.099), and the Control column again prominent as a destination. The shift from Visual to Control network prominence as the dominant attended destination at Schaefer-400 is consistent across both cohorts and constitutes the principal interpretability finding of this work.

This resolution-dependent shift is neurobiologically interpretable. Frontoparietal regions are spatially distributed and finely subdivided; at a coarse parcellation, they are not resolved into sufficiently distinct nodes for the attention mechanism to leverage their inter-individual variability. The Visual cortex, anatomically compact and well-localised, remains discriminable even at low resolution. A finer parcellation restores the resolution of frontoparietal nodes and allows the attention mechanism to converge on the connectivity features these regions are known to carry.

The prominence of the Control network at Schaefer-400 is directly consistent with Finn et al. (2015), who demonstrated that restricting identification to frontoparietal connections alone raises accuracy to approximately 99% and assigned those regions the highest Differential Power scores across the whole brain. The GAT converges on the same network hierarchy without any explicit supervision on network labels, providing independent computational support for the neurobiological primacy of frontoparietal connectivity in individual identification. That this convergence replicates across two demographically distinct cohorts strengthens the claim that it reflects genuine neurobiological structure rather than dataset-specific overfitting.

It is important to note that this remains a qualitative observation. No quantitative alignment was computed between the learned attention weights and the differential power maps of Finn et al. (2015), and no statistical test of correspondence was performed. A formal comparison of these two importance measures represents a natural direction for future work.

3.7 Conclusion

Across all three stages of this study, individual brain identity is reliably recoverable from resting-state functional connectivity matrices, from the zero-training Pearson correlation baseline to classical supervised classifiers to deep graph attention networks. This consistency confirms that the fingerprinting signal is a genuine property of FC organisation rather than an artefact of any particular analytical choice.

Stage 1 establishes that unsupervised similarity matching already recovers identity at a meaningful rate (68.14%–97.92% depending on cohort and atlas), but leaves a substantial gap relative to supervised approaches, motivating the subsequent stages. Stage 2 shows that supervised classical classifiers close most of this gap, particularly LDA and RBF SVM, but are sensitive to atlas resolution and PCA dimensionality in ways that limit their ceiling. Stage 3 demonstrates that the Siamese GAT surpasses all other methods on HCP-Aging and matches or approaches the best classical results on HCP-YA, while additionally providing a neurobiologically interpretable attention map at no extra cost.

The perfect identification achieved by the GAT on HCP-Aging at Atlas 400 (TOP-1 = 100%, ROC-AUC = 1.0000) constitutes a proof of concept that individual brain identity is fully recoverable from functional connectivity under the right conditions. The convergence of the GAT attention mechanism on the frontoparietal network, mirroring the differential power hierarchy established by Finn et al. (2015), further suggests that the model is capturing genuine neurobiological structure rather than fitting noise.

Establishing high identification accuracy in normative populations is a necessary precondition for extending these methods to clinical settings, where individual FC profiles may serve as biomarkers for psychiatric and neurological conditions, track longitudinal changes in brain health, or stratify patients in heterogeneous cohorts. The framework developed here, topologically aware, contrastively trained, and interpretable, provides a principled foundation for that extension.

Conclusion and Future Research Directions

Summary of Contributions

This thesis set out to investigate whether increasingly expressive computational models could extract more identity-relevant structure from the functional connectome than traditional linear similarity measures. The work was structured around a systematic, three-stage experimental framework applied under identical conditions, enabling fair and principled comparison across a spectrum of modelling approaches.

The first stage established a zero-training Pearson correlation baseline, replicating foundational brain fingerprinting protocols and confirming that raw functional connectivity matrices carry sufficient individual-specific signal to achieve non-trivial identification rates. This baseline served as the lower bound against which all subsequent models were evaluated.

The second stage demonstrated that reformulating identification as a pairwise verification task and applying classical machine learning classifiers (including Linear Discriminant Analysis, Support Vector Machines, and AdaBoost) over a compressed feature space yielded substantial improvements over the unsupervised baseline. These results established that learning a discriminative decision boundary extracts considerably more identity-relevant structure from the connectome than hand-crafted similarity metrics alone.

The third and most significant stage showed that deep Siamese architectures governed by an InfoNCE contrastive loss further advanced identification performance, with the Graph Attention Network (GAT) emerging as the dominant architecture across every experimental condition. Crucially, the GAT achieved **100% Top-1 identification accuracy** on the HCP-Aging dataset at fine-grained atlas resolution, demonstrating that topologically-aware,

graph-based representations substantially outperform vector-based approaches that discard the non-Euclidean structure of the connectome.

Beyond raw performance, the framework contributed two additional dimensions of scientific value. First, the systematic comparison of Schaefer-100 and Schaefer-400 parcellations revealed the sensitivity of identification accuracy to atlas resolution, providing empirical guidance for future methodological choices. Second, the GAT's learned attention maps offered a form of neurobiological interpretability that emerged without explicit supervision: the model naturally prioritised the frontoparietal (Control) network, providing computational corroboration for established neuroscientific theories on the primacy of association networks in encoding inter-individual variability.

Limitations

Despite the promising results, several limitations must be acknowledged. The study relied exclusively on data from the Human Connectome Project, a highly controlled acquisition protocol with homogeneous young and aging adult cohorts scanned under near-identical conditions. Generalisation to multi-site, multi-scanner, or clinically acquired data (where noise profiles, acquisition parameters, and participant characteristics vary considerably) remains an open question.

Furthermore, the parcellations used in this work, while widely validated, introduce a degree of spatial smoothing that may obscure finer-grained connectivity patterns. The choice of atlas is itself a source of variance, and the framework was not evaluated across the full spectrum of available parcellation schemes. Additionally, the dataset size, though sufficient for the identification task, is modest by deep learning standards, raising questions about the scalability of the learned representations to larger and more heterogeneous populations.

Finally, while the attention maps provide valuable post-hoc interpretability, they do not constitute a causal explanation of which connectivity features are truly discriminative versus incidentally correlated with identity. Establishing causal interpretability remains a broader open challenge in neuroimaging-based deep learning.

Future Research Directions

The findings of this thesis open several promising avenues for future investigation, spanning methodological refinements, clinical translation, and broader scientific questions.

Multi-Modal and Dynamic Connectivity

The present framework operates exclusively on static, session-averaged functional connectivity matrices. A natural extension is to incorporate **dynamic functional connectivity**, capturing the time-varying reconfiguration of brain networks across a session. Sliding-window approaches and more principled decompositions of connectivity dynamics have shown complementary identifying information to static FC, and integrating both modalities within a unified Siamese architecture represents a compelling direction. Similarly, combining structural connectivity derived from diffusion MRI with functional connectivity in a **multi-modal graph** could provide richer and more stable individual signatures.

Parcellation-Agnostic Representations

Rather than committing to a fixed atlas, future work could explore **parcellation-agnostic approaches** that operate directly on voxel-level or surface-level signals, learning spatially adaptive groupings as part of the optimisation process. Emerging architectures such as graph transformers and self-supervised spatial encoders are particularly well-suited to this challenge, potentially circumventing the resolution sensitivity observed in this work.

Cross-Site and Clinical Generalisability

A critical step toward real-world applicability is evaluating the framework on **multi-site datasets** such as the UK Biobank or the Adolescent Brain Cognitive Development (ABCD) study, where scanner heterogeneity and demographic diversity pose greater challenges. Domain adaptation and harmonisation techniques (including ComBat and adversarial training) could be incorporated into the pipeline to improve robustness across acquisition sites.

Clinical Applications and Biomarker Discovery

The most impactful extension of this work lies in its translation toward **precision medicine**. The learned graph embeddings could be repurposed as continuous biomarkers for predicting behavioural phenotypes such as fluid intelligence, sustained attention, or socioeconomic indicators. In clinical populations, the same framework could be adapted to stratify patients with heterogeneous neurological and psychiatric conditions (including autism spectrum disorder (ASD), attention deficit hyperactivity disorder (ADHD), and schizophrenia) by measuring the deviation of an individual's connectivity fingerprint from a normative reference derived from healthy controls.

Explainability and Causal Inference

While the GAT attention maps provide an initial window into which connections drive identification, future work should pursue more rigorous **explainability methods**, including gradient-based attribution, integrated gradients, and causal graph analysis. Establishing which connectivity patterns are causally (rather than merely correlationally) linked to individual identity would strengthen the neurobiological credibility of the framework and guide targeted hypothesis testing in subsequent experimental studies.

Closing Remarks

This thesis has demonstrated that the functional connectome encodes individual identity in ways that go well beyond what linear similarity measures can capture. By systematically advancing from correlation baselines through classical machine learning to topologically-aware deep Siamese architectures, this work has shown that graph-based contrastive learning is a powerful and interpretable paradigm for brain fingerprinting. The perfect identification accuracy achieved by the Siamese GAT on the HCP-Aging cohort, combined with its neurobiologically meaningful attention maps, positions this framework as a strong foundation for future work at the intersection of network neuroscience, deep learning, and precision medicine.

The brain is, above all, an individual organ. Every connectome tells a story unique to its owner, and the tools to read that story are only beginning to mature.

Bibliography

Cancer Research UK. Diagram showing the lobes of the brain. Wikimedia Commons, 2014. URL `https://commons.wikimedia.org/wiki/File:Diagram_showing_the_lobes_of_the_brain_CRUK_308.svg`. Licensed under CC BY-SA 4.0.

Stephen M Smith, Peter T Fox, Karla L Miller, David C Glahn, P Mickle Fox, Clare E Mackay, Nicola Filippini, Kate E Watkins, Roberto Toro, Angela R Laird, et al. Correspondence of the brain's functional architecture during activation and rest. *Proceedings of the national academy of sciences*, 106(31):13040–13045, 2009.

Eric R Kandel, James H Schwartz, Thomas M Jessell, et al. *Principles of Neural Science*, volume 4. McGraw-hill New York, 2000.

Ye Tian, Daniel S Margulies, Michael Breakspear, and Andrew Zalesky. Topographic organization of the human subcortex unveiled with functional connectivity gradients. *Nature Neuroscience*, 23(11):1421–1432, 2020. doi: 10.1038/s41593-020-00711-6.

Olaf Sporns. *Networks of the Brain*. MIT press, 2016.

B. T. Thomas Yeo, Fenna M. Krienen, Jorge Sepulcre, Mert R. Sabuncu, Daniyar Lashkari, Marisa Hollinshead, Joshua L. Roffman, Jordan W. Smoller, Lilla Zöllei, Jonathan R. Polimeni, et al. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, 106(3):1125–1165, 2011. doi: 10.1152/jn.00338.2011.

Marcus E Raichle, Ann Mary MacLeod, Abraham Z Snyder, William J Powers, Debra A Gusnard, and Gordon L Shulman. A default mode of brain function. *Proceedings of the national academy of sciences*, 98(2):676–682, 2001.

Michael D Fox, Abraham Z Snyder, Justin L Vincent, Maurizio Corbetta, David C Van Essen, and Marcus E Raichle. The human brain is intrinsically organized into dynamic,

anticorrelated functional networks. *Proceedings of the National Academy of Sciences*, 102 (27):9673–9678, 2005.

Seiji Ogawa, Tso-Ming Lee, Alan R Kay, and David W Tank. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *proceedings of the National Academy of Sciences*, 87(24):9868–9872, 1990.

Nikos K Logothetis, Jon Pauls, Mark Augath, Torsten Trinath, and Axel Oeltermann. Neurophysiological investigation of the basis of the fmri signal. *nature*, 412(6843):150–157, 2001.

Bharat Biswal, F Zerrin Yetkin, Victor M Haughton, and James S Hyde. Functional connectivity in the motor cortex of resting human brain using echo-planar mri. *Magnetic resonance in medicine*, 34(4):537–541, 1995.

Jonathan D Power, Anish Mitra, Timothy O Laumann, Abraham Z Snyder, Bradley L Schlaggar, and Steven E Petersen. Methods to detect, characterize, and remove motion artifact in resting state fmri. *neuroimage*, 84:320–341, 2014.

Karl J Friston, Chris D Frith, Peter F Liddle, and Richard SJ Frackowiak. Functional connectivity: the principal-component analysis of large (pet) data sets. *Journal of Cerebral Blood Flow & Metabolism*, 13(1):5–14, 1993.

Ed Bullmore and Olaf Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature reviews neuroscience*, 10(3):186–198, 2009.

Stephen M Smith, Karla L Miller, Gholamreza Salimi-Khorshidi, Matthew Webster, Christian F Beckmann, Thomas E Nichols, Joseph D Ramsey, and Mark W Woolrich. Network modelling methods for fmri. *Neuroimage*, 54(2):875–891, 2011.

Alexander Schaefer, Ru Kong, Evan M. Gordon, Timothy O. Laumann, Xi-Nian Zuo, Avram J. Holmes, Simon B. Eickhoff, and B. T. Thomas Yeo. Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. *Cerebral Cortex*, 28 (9):3095–3114, 2018a. doi: 10.1093/cercor/bhx179.

Emily S. Finn, Xilin Shen, Dustin Scheinost, Monica D. Rosenberg, Marvin Chun, Xenophon Papademetris, and R. Todd Constable. Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. *Nature Neuroscience*, 18(11):1664–1671, 2015. doi: 10.1038/nn.4135.

Sophia Mueller, Danhong Wang, Michael D Fox, BT Thomas Yeo, Jorge Sepulcre, Mert R Sabuncu, Rebecca Shafee, Jie Lu, and Hesheng Liu. Individual variability in functional connectivity architecture of the human brain. *Neuron*, 77(3):586–595, 2013.

Emily S Finn, Dustin Scheinost, Daniel M Finn, Xilin Shen, Xenophon Papademetris, and R Todd Constable. Can brain state be manipulated to emphasize individual differences in functional connectivity? *NeuroImage*, 160:140–151, 2017.

Jonathan D Power, Kelly A Barnes, Abraham Z Snyder, Bradley L Schlaggar, and Steven E Petersen. Spurious but systematic correlations in functional connectivity mri networks arise from subject motion. *Neuroimage*, 59(3):2142–2154, 2012.

Oscar Miranda-Dominguez, Eric Feczko, David S Grayson, Hasse Walum, Joel T Nigg, and Damien A Fair. Heritability of the human connectome: A connectotyping study. *Network Neuroscience*, 2(02):175–199, 2018.

Monica D Rosenberg, Emily S Finn, Dustin Scheinost, Xenophon Papademetris, Xilin Shen, R Todd Constable, and Marvin M Chun. A neuromarker of sustained attention from whole-brain functional connectivity. *Nature neuroscience*, 19(1):165–171, 2016.

Leon Qi Rong Ooi, Jianzhong Chen, Shaoshi Zhang, Ru Kong, Angela Tam, Jingwei Li, Elvisha Dhamala, Juan Helen Zhou, Avram J Holmes, and BT Thomas Yeo. Comparison of individualized behavioral predictions across anatomical, diffusion and functional connectivity mri. *NeuroImage*, 263:119636, 2022.

Randy L Buckner, Fenna M Krienen, and BT Thomas Yeo. Opportunities and limitations of intrinsic functional connectivity mri. *Nature neuroscience*, 16(7):832–837, 2013.

Choong-Wan Woo, Luke J Chang, Martin A Lindquist, and Tor D Wager. Building better biomarkers: brain models in translational neuroimaging. *Nature neuroscience*, 20(3):365–377, 2017.

Corey Horien, Xilin Shen, Dustin Scheinost, and R Todd Constable. The individual functional connectome is unique and stable over months to years. *NeuroImage*, 189:676–687, 2019.

Enrico Amico and Joaquín Goñi. The quest for identifiability in human functional connectomes. *Scientific reports*, 8(1):8254, 2018.

Evan M Gordon, Timothy O Laumann, Adrian W Gilmore, Dillan J Newbold, Deanna J Greene, Jeffrey J Berg, Mario Ortega, Catherine Hoyt-Drazen, Caterina Gratton, Haoxin

Sun, et al. Precision functional mapping of individual human brains. *Neuron*, 95(4): 791–807, 2017.

Jin Liu, Xuhong Liao, Mingrui Xia, and Yong He. Chronnectome fingerprinting: Identifying individuals and predicting higher cognitive functions using dynamic brain connectivity patterns. *Human brain mapping*, 39(2):902–915, 2018.

Amin Ghaffari, Yufei Zhao, Xu Chen, Jason Langley, and Xiaoping Hu. Dynamic fingerprinting of the human functional connectome. *Brain Connectivity*, page 21580014261420882, 2025.

Md Kaosar Uddin, Nghi Nguyen, Huajun Huang, Duy Duong-Tran, and Jingyi Zheng. Divergence unveils further distinct phenotypic traits of human brain connectomics fingerprint. *iScience*, 29(1), 2026.

Zhen-Qi Liu, Andrea I Luppi, Justine Y Hansen, Ye Ella Tian, Andrew Zalesky, BT Thomas Yeo, Ben D Fulcher, and Bratislav Misic. Benchmarking methods for mapping functional connectivity in the brain. *Nature Methods*, 22(7):1593–1602, 2025.

Chao Jiang, Ye He, Richard F Betzel, Yin-Shan Wang, Xiu-Xia Xing, and Xi-Nian Zuo. Optimizing network neuroscience computation of individual differences in human spontaneous brain activity for test-retest reliability. *Network Neuroscience*, 7(3):1080–1108, 2023.

Yasaman Shahhosseini and Michelle F Miranda. Functional connectivity methods and their applications in fmri data. *Entropy*, 24(3):390, 2022.

Kausar Abbas, Enrico Amico, Diana Otero Svaldi, Uttara Tipnis, Duy Anh Duong-Tran, Mintao Liu, Meenusree Rajapandian, Jaroslaw Harezlak, Beau M Ances, and Joaquín Goñi. Geff: Graph embedding for functional fingerprinting. *NeuroImage*, 221:117181, 2020.

Wei Liu and Xuemin Zhang. Using independent component analysis to extract a cross-modality and individual-specific brain baseline pattern. *NeuroImage*, 303:120925, 2024.

Ali Shojaee, Kendrick Li, and Gowtham Atluri. A machine learning framework for accurate functional connectome fingerprinting and an application of a siamese network. In *International workshop on connectomics in neuroimaging*, pages 83–94. Springer, 2019.

Andrew Hannum, Mario A Lopez, Saúl A Blanco, and Richard F Betzel. High-accuracy machine learning techniques for functional connectome fingerprinting and cognitive state decoding. *Human Brain Mapping*, 44(16):5294–5308, 2023.

AS Keller, AR Pines, S Shanmugan, VJ Sydnor, Z Cui, MA Bertolero, R Barzilay, AF Alexander-Bloch, N Byington, A Chen, et al. Personalized functional brain network topography is associated with individual differences in youth cognition. nat. commun. 14 (1), 8411, 2023.

B.J. Casey, Tariq Cannonier, May I Conley, Alexandra O Cohen, Deanna M Barch, Mary M Heitzeg, Mary E Soules, Theresa Teslovich, Danielle V Dellarco, Hugh Garavan, et al. The adolescent brain cognitive development (abcd) study: Imaging acquisition across 21 sites. *Developmental Cognitive Neuroscience*, 32:43–54, 2018.

Gokce Sarar, Bhaskar Rao, and Thomas Liu. Functional connectome fingerprinting using shallow feedforward neural networks. *Proceedings of the National Academy of Sciences*, 118(15):e2021852118, 2021.

Juhyeon Lee and Jong-Hwan Lee. Discovering individual fingerprints in resting-state functional connectivity using deep neural networks. *Human Brain Mapping*, 45(1):e26561, 2024.

Sanjay Ghosh et al. Functional connectome fingerprinting using convolutional and dictionary learning. *arXiv preprint arXiv:2510.26120*, 2025.

Jiayu Lu, Tianyi Yan, Lan Yang, Xi Zhang, Jiaxin Li, Dandan Li, Jie Xiang, and Bin Wang. Brain fingerprinting and cognitive behavior predicting using functional connectome of high inter-subject variability. *NeuroImage*, 295:120651, 2024.

Suleiman Ibrahim Mohammad, Eman Raeed Azzam, Asokan Vasudevan, Sayed M Ismail, Hina Ayaz, and KDV Prasad. Precision neurodiversity: personalized brain network architecture as a window into cognitive variability. *Frontiers in Human Neuroscience*, 19: 1669431, 2025.

Rui Xu, Shuwan Zhao, Zhengyi Liu, Qianhui Jin, Zeyao Wang, Dongxu Liu, Kun Zhao, Suyu Zhong, Jiaying Zhang, Yong Liu, et al. Functional fingerprinting for the developing brain using deep metric learning. *Imaging Neuroscience*, 4:IMAG–a, 2026.

Jivesh Ramduny and Clare Kelly. Connectome-based fingerprinting: reproducibility, precision, and behavioral prediction. *Neuropsychopharmacology*, 50(1):114–123, 2025.

Maria Colpo, Erica Silvestri, Alessandro Salvalaggio, Diego Cecchin, Maurizio Corbetta, and Alessandra Bertoldo. Structural-functional fingerprinting for abnormalities investigation in glioma patients. *Scientific Reports*, 15(1):38404, 2025.

Jun Tang, Jie Chen, Miaojun Hu, Yao Hu, Zixi Zhang, and Liuming Xiao. Diagnosis of autism spectrum disorder (asd) by dynamic functional connectivity using gnn-lstm. *Sensors*, 25 (1):156, 2024.

Vaibhav Tripathi, Laura Rigolo, Bethany K Bracken, Colin P Galvin, Alexandra J Golby, Yanmei Tie, and David C Somers. Utilizing connectome fingerprinting functional mri models for motor activity prediction in presurgical planning: A feasibility study. Technical report, Wiley Online Library, 2024.

Daniel F Hermens, Colin Russo, Zack Shan, and Jim Lagopoulos. Brain fingerprinting: A promising future application for predicting mental illness. *Futures*, 152:103211, 2023.

Leanne M Williams and Susan Whitfield Gabrieli. Neuroimaging for precision medicine in psychiatry. *Neuropsychopharmacology*, 50(1):246–257, 2025.

Susan Y. Bookheimer, David H. Salat, Melanie Terpstra, Beau M. Ances, Deanna M. Barch, Randy L. Buckner, Gregory C. Burgess, Sandra W. Curtiss, Mirella Diaz-Santos, Jennifer Stine Elam, Bruce Fischl, Douglas N. Greve, Hannah A. Hagy, Michael P. Harms, Olivia M. Hatch, Trey Hedden, Cynthia Hodge, Kevin C. Japardi, Taylor P. Kuhn, Trinity K. Ly, Stephen M. Smith, Leah H. Somerville, Kâmil Uğurbil, Andre Van der Kouwe, David Van Essen, Essa Williams, Essa Yacoub, and Lily Angeles. The lifespan human connectome project in aging: An overview. *NeuroImage*, 185:335–348, 2019. doi: 10.1016/j.neuroimage.2018.10.009.

Matthew F. Glasser, Stamatios N. Sotiropoulos, J. Anthony Wilson, Timothy S. Coalson, Bruce Fischl, Jesper L. Andersson, et al. The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage*, 80:105–124, 2013. doi: 10.1016/j.neuroimage.2013.04.127.

Gholamreza Salimi-Khorshidi, Gwenaëlle Douaud, Christian F. Beckmann, Matthew F. Glasser, Ludovica Griffanti, and Stephen M. Smith. Automatic denoising of functional MRI data: combining independent component analysis and hierarchical fusion of classifiers. *NeuroImage*, 90:449–468, 2014. doi: 10.1016/j.neuroimage.2013.11.046.

F. Tijhuis, M. Schepers, E. Centeno, Maciel, B. de A.P.C., L. Douw, and Nobrega Santos, F. Human connectome project rs-fMRI repository. GitHub, 2022. URL `https://github.com/floristijhuis/HCP-rfMRI-repository`. Accessed 2024. Data available at https://doi.org/10.5281/zenodo.6770120.

Alexander Schaefer, Ru Kong, Evan M Gordon, Timothy O Laumann, Xi-Nian Zuo, Avram J Holmes, Simon B Eickhoff, and BT Thomas Yeo. Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity mri. *Cerebral cortex*, 28(9):3095–3114, 2018b.

Karl Pearson. Notes on regression and inheritance in the case of two parents. *Proceedings of the Royal Society of London*, 58:240–242, 1895. doi: 10.1098/rspl.1895.0041.

Harold Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of Educational Psychology*, 24(6):417–441, 1933. doi: 10.1037/h0071325.

Ronald A. Fisher. The use of multiple measurements in taxonomic problems. *Annals of Eugenics*, 7(2):179–188, 1936. doi: 10.1111/j.1469-1809.1936.tb02137.x.

Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20(3): 273–297, 1995. doi: 10.1007/BF00994018.

Bernhard E. Boser, Isabelle M. Guyon, and Vladimir N. Vapnik. A training algorithm for optimal margin classifiers. In *Proceedings of the 5th Annual Workshop on Computational Learning Theory (COLT)*, pages 144–152, 1992. doi: 10.1145/130385.130401.

John C. Platt. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. In Alexander J. Smola, Peter Bartlett, Bernhard Schölkopf, and Dale Schuurmans, editors, *Advances in Large Margin Classifiers*, pages 61–74. MIT Press, 1999.

Yoav Freund and Robert E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997. doi: 10.1006/jcss.1997.1504.

Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. Signature verification using a Siamese time delay neural network. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 6, pages 737–744, 1993.

Sumit Chopra, Raia Hadsell, and Yann LeCun. Learning a similarity metric discriminatively, with application to face verification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 539–546, 2005. doi: 10.1109/CVPR.2005.202.

Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018. URL `https://arxiv.org/abs/1807.03748`. arXiv:1807.03748.

Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *Proceedings of the International Conference on Machine Learning (ICML)*, volume 119, pages 1597–1607, 2020.

Sergey Ioffe and Christian Szegedy. Batch normalization: accelerating deep network training by reducing internal covariate shift. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 448–456, 2015.

Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017. URL `https://arxiv.org/abs/1609.02907`. arXiv:1609.02907.

Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019. URL `https://arxiv.org/abs/1903.02428`. arXiv:1903.02428.

Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018. URL `https://arxiv.org/abs/1710.10903`. arXiv:1710.10903.

Diederik P. Kingma and Jimmy Ba. Adam: a method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015. URL `https://arxiv.org/abs/1412.6980`. arXiv:1412.6980.

Ilya Loshchilov and Frank Hutter. SGDR: stochastic gradient descent with warm restarts. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017. URL `https://arxiv.org/abs/1608.03983`. arXiv:1608.03983.

David C. Van Essen, Stephen M. Smith, Deanna M. Barch, Timothy E. J. Behrens, Essa Yacoub, and Kâmil Ugurbil. The wu-minn human connectome project: An overview. *NeuroImage*, 80:62–79, 2013. doi: 10.1016/j.neuroimage.2013.05.041.