# Deep Learning as an investigation tool in Neuroimaging

Arna Ghosh

Integrated Program in Neuroscience

McGill University

Montreal, Quebec

July 4, 2019

A thesis submitted to McGill University
in partial fulfillment of the requirements of the degree of Master of Science

## DEDICATION

This thesis is dedicated to my parents.

## ACKNOWLEDGEMENTS

My sincere appreciation to my supervisors, Marie-Hélène Boudrias and Georgios D. Mitsis, who have provided me with the guidance, direction, and support towards my training and research development. Special thanks go out to all of the subjects and patients who took part in our research. Their involvement, curiosity, enthusiasm, and endless patience were an essential part of the studies.

Many thanks to my labmates and friends who made the laboratory a fantastic working environment: Fabien Dal Maso, Nikolaos Dimitriou, Michalis Kassinopoulos, Kyriaki Kostoglou, Gregory Laredo, Dylan Mann-Krizsnik, Prokopis Prokopiou, Hamidreza Sadreazami, Alba Xifra-Porxas, Xuanteng Yan. Special thanks to Dylan for helping me write the French version of the abstract.

I would like to thank all the people from the IPN, notably: Alexander DeGuise, Dhabisha Kohilanathan, Vivian Omune, Mirko Sablich, and Katherine Vanka for providing ongoing advice and guidance. I would also like to thank all our collaborators at the Jewish Rehabilitation Hospital, notably: Bennet Desormeau and Marc Roig for their expertise and efforts in the EEG data collection. Numerous thanks to my committee members, Stefanie Blain-Moraes and Karim Jerbi. Their expertise and advice were extremely valuable for the successful completion of the projects.

I would also like to extend my sincere gratitude to the funding agencies that supported me throughout the course of my degree, notably: Healthy Brains for Healthy Lives, MITACS, the McGill University Faculty of Medicine and the Québec Bio-Imaging Network

## CONTRIBUTION TO ORIGINAL KNOWLEDGE

1. **Project 1** – This study investigates the use of convolutional neural networks to reveal subtle discriminative patterns from neuroimaging data, specifically in the context of identifying exercise-induced EEG patterns during fixed-force handgrips. We trained a convolutional neural network to discriminate exercise-induced modulations of event-related desynchronization patterns in the EEG. Added to this, we developed a feature visualization technique, termed ccCAM, that revealed finer frequency bands of brain activity and specific brain regions that were modulated by exercise.
2. **Project 2** – This study investigates the role of dimensionality reduction techniques in the context of brain age prediction using a combination of structural and functional features of the brain. We applied principal component analysis and canonical correlation analysis in conjunction with machine learning models to predict the brain age of a cohort of healthy subjects. Canonical correlation analysis yielded the best performance and also revealed brain regions and features that reliably show age-related changes. In addition, we also developed a graph convolutional network-based approach that leverages the functional topology to improve prediction.

## CONTRIBUTION OF AUTHORS

The body of work presented in this thesis would not have been made possible without a close collaboration between myself and my M.Sc. supervisors, Marie-Hélène Boudrias and Georgios D. Mitsis. Listed below are the specific contributions:

1. **Marie-Hélène Boudrias** – For projects 1–2: formulated the research question, assisted in the interpretation of results, and helped with writing the manuscripts. For project 1: supervised the data collection and provided behavioral and imaging data from control and exercise group subjects. For project 2: supervised the MEG
2. **Georgios D. Mitsis** – For projects 1–2: helped with designing the analysis pipeline, assisted in doing the deep learning and machine learning analysis, provided advice on statistical analysis, assisted in the interpretation of results, and helped with writing the manuscripts.

The following list contains the other co-authors, in alphabetical order, who have greatly contributed to the manuscripts:

1. **Fabien Dal Maso** – Collected the data for project 1 and did the EEG data preprocessing. Provided advice on the data analysis pipeline and the statistical analysis for project 1.
2. **Marc Roig** – Designed the experiment protocol and supervised the data collection for project 1.
3. **Alba Xifra-Porxas** – Provided useful advice on the analysis pipeline for project 1. Provided the CAM-CAN data and did the MEG data preprocessing

for project 2. Also, helped with the machine learning analysis and interpretation of results for project 2.

## ABSTRACT

**Background:** Following the successful use of deep learning (DL) in the field of computer vision and natural language processing, Convolutional Neural Networks (CNN) and Recurrent Neural Networks (RNN) have been increasingly applied to neuroimaging studies to differentiate diseased population from healthy subjects. Deep learning models have been applied to datasets acquired from different modalities such as electroencephalography (EEG), Magnetic Resonance Imaging (MRI) and magnetoencephalography (MEG). Most of the previous studies have focussed on decoding stimulus from neural responses or identifying diseased-vs.-healthy condition-specific neural responses. Although DL models have outperformed competing methods in these contexts, the usefulness of the trained models to improve our understanding of the underlying neural substrates remains limited. This is primarily due to the difficulty associated with visualizing and interpreting the predictive features learned by DL models.

**Goal:** Our main aim was to leverage DL models to reveal subtle discriminative brain patterns from neuroimaging data and subsequently develop methods to visualize these patterns. We investigated the applicability of DL models to 2 specific neuroimaging studies: (1) an EEG dataset obtained from a small number of participants to understand the add-on effects of acute cardiovascular exercise on motor learning in healthy subjects, (2) a large-scale MRI and MEG study to characterize the process of healthy brain aging.

**Methods:** This thesis is comprised of two distinct experiments. (1) We leveraged

the end-to-end learning ability of CNNs to investigate the differences in EEG activity between an exercise and a control group from data collected while subjects performed isometric handgrips. Subsequently, we developed a method to visualize the task-specific, discriminatory EEG patterns between the two groups. (2) We investigated machine learning models coupled with several multivariate associativity techniques to predict the age of healthy individuals from their T1-weighted MRI and resting-state MEG data. We also explored Graph Convolutional Networks (GCNs) and CNNs to incorporate the topological information of the data into the prediction models.

**Results:** We found that CNNs can be reliably trained on a dataset collected from a relatively small number of participants using an adversarial training strategy. In addition, we were able to identify relevant frequency bands and brain regions that were modulated by exercise. A previous analysis from our group using more standard statistical analysis concluded that beta band activity in the range of 15-29 Hz originating from sensorimotor areas was modulated by exercise. Using the CNN-based pipeline on the same dataset, we observed (i) a finer frequency band within the beta-band that was modulated by exercise, as well as (ii) a significant modulation of the event-related desynchronization in this frequency band located in bilateral sensorimotor cortices and contralateral prefrontal regions to the moving hand. Therefore, our approach demonstrates the feasibility of identifying subtle discriminative features in a completely data-driven manner using deep learning. In the context of brain age prediction, we found that the subcortical regions were more reliable predictors of age as compared to cortical regions. In addition, we also observed how the functional

organization of the brain changed with age.

**Significance:** Our approach demonstrated the feasibility of identifying subtle discriminative features in a completely data-driven manner using DL. We believe these results hold a significant contribution to the methodological advances for small-scale neuroimaging studies where a small number of subjects are traditionally tested, e.g. – neurorehabilitation. Our findings from the brain age prediction expanded upon previous work in the field and provided useful insights into the brain areas that are reliably affected with age. In addition, the graph network architecture demonstrated ways to include the topology of the brains functional organization while analyzing age-related effects on a diverse set of neuroimaging features. Collectively, the findings and methods presented in this thesis demonstrated the wide scope of using DL models in the analysis of various modalities of neuroimaging data.

# ABRÉGÉ

**Contexte:** Suite à l'utilisation réussie de l'apprentissage en profondeur (DL) dans le domaine de la vision par ordinateur et du traitement du langage naturel, les réseaux de neurones à convolution (CNN) et les réseaux de neurones récurrents (RNN) ont été de plus en plus utilisés dans les études en neuroimagerie afin de différencier les populations en santé de celles présentant une maladie neurodégénérative. Des modèles d'apprentissage en profondeur ont été appliqués à des ensembles de données acquis à partir de différentes modalités telles que l'électroencéphalographie (EEG), l'imagerie par résonance magnétique (IRM) et la magnétoencéphalographie (MEG). La plupart des études antérieures se sont concentrées sur le décodage du stimulus à partir des réponses neuronales ou l'identification de réponses neuronales spécifiques à un état pathologique. Bien que les modèles DL aient surperformé les méthodes concurrentes dans ces contextes, l'utilité des modèles conus pour améliorer notre compréhension des substrats neuronaux sous-jacents reste limitée. Ceci est principalement dû à la difficulté associée à la visualisation et à l'interprétation des fonctions prédictives apprises par les modèles DL.

**Objectif:** Notre objectif principal était dexploiter des modèles DL pour révéler des caractéristiques cérébrales subtiles discriminantes à partir de données de neuroimagerie, puis de développer des méthodes permettant de les visualiser. Nous avons étudié l'applicabilité des modèles DL dans le cadre de deux études spécifiques utilisant la neuroimagerie soient: (1) un ensemble de données EEG obtenu à partir d'un

nombre restreint de participants afin d'étudier les effets additifs de l'exercice cardiovasculaire aigu sur l'apprentissage moteur auprès de sujets en santé, et (2) une étude à partir de données dIRM et de MEG à grande échelle pour caractériser le processus du vieillissement cérébral sain.

**Méthodes:** Cette thèse est composée de deux expériences distinctes. (1) Nous avons exploité la capacité dapprentissage de bout en bout des CNN pour étudier les différences dans les activités EEG entre un groupe assujetti à lexercice cardiovasculaire et un groupe contrôle. Les données ont été collectées pendant que les sujets exécutaient des contractions manuelles isométriques. Par la suite, nous avons développé une méthode permettant de visualiser les formes EEG pouvant discriminer les deux groupes et étant spécifiques à la tâche. (2) Nous avons étudié des modèles dapprentissage machine associés à divers techniques dassociativité à variables multiples pour prédire lâge des individus en bonne santé à partir de leurs données dIRM pondérées en T1 et de MEG à létat de repos. Nous avons également exploré les réseaux convolutionnels de graphes (GCN) et CNN pour incorporer les informations topologiques des données dans les modèles de prévision.

**Résultats:** Nous avons constaté que les CNN peuvent être entraînés de manière fiable sur un ensemble de données collecté auprès d'un nombre relativement restreint de participants à l'aide d'une stratégie de formation contradictoire. De plus, nous avons pu identifier des bandes de fréquences pertinentes et des régions du cerveau modulées par l'exercice. Une analyse antérieure de notre groupe utilisant une analyse statistique plus standard a conclu que l'activité de la bande bêta dans linterval de 15 à 29 Hz provenant de zones sensorimotrices était modulée par l'exercice. En utilisant

le pipeline basé sur les CNN sur le même ensemble de données, nous avons observé (i) une bande de fréquence plus fine modulée par l'exercice dans la bande bêta, ainsi que (ii) une modulation significative de la désynchronisation liée à l'événement dans cette bande de fréquence située dans les cortex sensorimoteurs bilatéraux et les régions préfrontales controlatérales à la main en mouvement. Par conséquent, notre approche démontre la faisabilité d'identifier des caractéristiques discriminantes subtiles d'une manière entièrement basée sur les données en utilisant l'apprentissage en profondeur. Dans le contexte de la prédiction de l'âge du cerveau, nous avons constaté que les régions sous-corticales étaient des prédicteurs plus fiables de l'âge en comparaison aux régions corticales. De plus, nous avons également observé comment l'organisation fonctionnelle du cerveau changeait avec l'âge.

**Importance:** Notre approche a démontré la faisabilité d'identifier des caractéristiques discriminantes subtiles d'une manière entièrement basée sur les données en utilisant la DL. Nous pensons que ces résultats apportent une contribution significative aux avancées méthodologiques pour les études en neuroimagerie à petite échelle dans lesquelles un nombre habituellement restreint de sujets est testé, par exemple la neuroréhabilitation. Nos conclusions tirées de la prédiction de l'âge du cerveau complémentent celles détudes antérieurs se penchant sur la même question et ont fourni des informations utiles sur les zones du cerveau systématiquement affectées par l'âge. En outre, larchitecture du réseau convolutionnel de graphes met en évidence des moyens pour inclure la topologie de lorganisation fonctionnelle du cerveau, tout en analysant les effets liés à lâge basés sur un ensemble divers de quantités caractéristiques issues de la neuroimagerie. Ensemble, les résultats et les méthodes

présentés dans cette thèse ont démontré le large champ dutilisation des modèles DL dans lanalyse de diverses modalités de données de neuroimagerie.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Figure page

# CHAPTER 1
# INTRODUCTION

One of the major aims of neuroimaging is to reveal insights about the function and structure of the brain from non-invasive measurements. Studying the living brain using various neuroimaging modalities provides insight into the functioning of the brain in typical as well as diseased conditions. Existing neuroimaging modalities leverage different physiological characteristics to reflect properties of the brain structure or function. Currently used popular imaging modalities include magnetic resonance imaging (MRI), functional MRI, diffusion tensor imaging (DTI), positron emission tomography (PET), electroencephalography (EEG), and magnetoencephalography (MEG), among others.

A great majority of neuroimaging studies performed so far, have relied on the *a priori* selection of features from the recorded structural or functional data [11]. However, this approach may yield sub-optimal feature selection and eventually prevent the detection of subtle discriminative patterns present in the data. With the increasing popularity of neuroimaging techniques to study the brain, the amount of digital data has exponentially risen. This has allowed the use of several "big data" techniques in extracting relevant information from the data [172]. Such techniques range from associative techniques to graph theoretic methods and machine learning models to improve the understanding of functional and structural aspects of the human brain in healthy and diseased connditions.

Data-driven approaches such as deep learning (DL) allow the identification of the optimal discriminative features in a given dataset. Specifically, Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) have been applied to computer vision and speech processing datasets with great success [61, 81, 92, 189]. DL architectures excel at finding hidden, low-dimensional features by discovering statistical regularities in high-dimensional training data, and can do so in a relatively unsupervised fashion [13, 59, 70, 96, 178, 179]. These approaches have also been used in the context of neuroimaging to identify features for structural and functional MRI data analysis [123] and EEG data decoding [12, 137, 139, 164]. Although the field of neuroimaging has seen a recent surge of papers involving the use of DL, the applicability of such models is primarily restricted to the classification of data segments into known categories. The usefulness of trained DL models to improve our understanding of the neural substrates underlying observed behavior is less straightforward, primarily due to difficulty associated with the visualization and interpretation of the feature space learnt by DL models.

The **main objective** of the present study is to develop a DL pipeline suited to most neuroimaging data analysis pipelines and subsequently analyze the features learnt by the DL architecture. We intend to establish a framework for future researchers to use DL architectures as data analysis tools for neuroimaging data. To this end, we will design a CNN-based analysis pipeline for the analysis of EEG data. This pipeline will subsequently be applied to EEG recordings collected to study the effect of acute cardiovascular exercise on motor learning in healthy subjects. We will also present a novel technique to identify the optimal feature set learnt by the

trained CNN, thus extending the CNN's usability beyond classification tasks. The resulting optimal feature set would yield insights into the neurophysiological substrates underlying the positive effects of exercise. Thereafter, we will compare the CNN-yielded optimal feature sets to prior literature in the field of cardiovascular exercise to exhibit the potential of CNNs to act as a powerful data analysis tool that can be used in a relatively unsupervised setting for EEG analysis.

We will further **explore machine and deep learning pipelines** to other neuroimaging modalities, namely MRI and MEG, to estimate the biological age of subjects from their neuroimaging data. The discrepancy between the reported and estimated age will subsequently be compared for healthy and diseased populations to be used as potential biomarker of neurological disorders. The ability of the multivariate associativity techniques, like canonical correlation analysis, in conjunction with machine learning models as well as Graph Convolutional Networks to characterize the complex process of brain aging demonstrates the usability of such methods to other applications across the neuroimaging domain.

# CHAPTER 2
# BACKGROUND

## 2.1 Deep Learning for neuroimaging

Following the success of DL in other fields, CNNs and RNNs have been increasingly used in neuroimaging to differentiate diseased population from a cohort of healthy subjects. The application of DL models have spanned several neuroimaging modalities, including EEG [12, 137, 139, 164], functional and structural MRI [50, 123] as well as MEG [24]. Most of the previous studies have aimed to either decode stimulus from neural responses or compare stimulus encoding in neural responses to those in different processing layers of a deep network. Studies involving DL for MRI have also focused on automated segmentation of the MR images. Some recent applications of DL models include inferring latent dynamics from simultaneously recorded, single-trial neural spiking data [118]. Although DL has outperformed competing methods in most of these paradigms, the usefulness of the trained models to improve our understanding of neural substrates remains limited. This is primarily due to the difficulty associated with the visualization and interpretation of predictive features learnt by the DL models.

### 2.1.1 DL architectures for EEG decoding

Machine learning techniques have been widely used for extracting information from EEG recordings, and therefore play an important role in several EEG-based research and application areas. For instance, machine learning forms a central tool for

EEG-based Brain Computer Interface (BCI) systems for clinical applications [114]. Such systems have aided, for example, people with severe paralysis to communicate [116] and to control telepresence robots [168]. Such systems have also been applied to facilitate stroke rehabilitation [126] and have the potential to be used in the treatment of epilepsy [55]. Furthermore, machine learning techniques have been increasingly recognized as important tools for neuroimaging data analysis pipelines that are aimed to answer neuroscientific queries [37, 88, 94, 150].

Although traditional machine learning methods have allowed impressive progress in EEG decoding and BCI design, there is considerable possibility of improvement with respect to several important aspects of information extraction from EEG, including decoding accuracy and interpretability. Hence, innovations in the area of machine learning have found continued applicability in improving current EEG decoding methods and BCI designs. Correspondingly, recent studies have investigated the potential of CNN and RNN for EEG decoding in the context of BCI design [22, 109, 139, 158, 183], cognitive load classification [12, 63] and seizure prediction [7, 137, 164]. Most of these studies aim to learn a DL network that is optimized for a specific application using a sufficiently large cohort of subjects. A recent study leverages the power of auto-encoders [85] to learn feature representation from noisy EEG data collected from limited number of subjects [154].

However, a core limitation of all the above mentioned studies lies in the interpretability of the feature representation learnt by the trained DL network. Consequently, this limits the use of DL to improve our understanding of the neural substrates underlying the observed behavior. Although previous studies have presented

certain feature visualization methods to address this issue, these methods have limited applicability in varied neuroimaging studies. Therefore, a major objective of this study is to fill this gap in existing literature by devising a method to reliably visualize the features learnt by the DL network trained on EEG data.

### 2.1.2 Visualizing the predictive feature space

CNNs have achieved superior performance in EEG decoding problems as compared to other *apriori* feature-based machine learning methods. But their lack of decomposability into intuitive and understandable components makes them hard to interpret [103]. The importance of interpretability carries paramount importance in neuroimaging analysis. It allows the researcher to understand the salient features in the input data that allows a deep network to exhibit superior performance. Therefore, this allows the researcher to reliably use CNNs as an investigation tool in neuroimaging. Although there has been some efforts in developing methodologies to create feature visualizations from trained CNNs, those methods are not widely applicable to varied neuroimaging studies.

Schirrmeister et al. described a correlation-based analysis to interpret the spectral features in EEG data that determine the performance of a trained CNN [139]. Their method involved correlating the input power of a particular frequency band to the activation of an artificial neuron in a layer of the CNN. Neurons showing high correlation values were considered to be capturing the information in that particular frequency band. Although the authors successfully used this method for interpreting the spectral features that were captured by each filter in a layer of the CNN, the aforesaid method could be considered as a feature validation approach. The authors

had to select the spectral feature *apriori* to observe its contribution to the CNN's performance.

Some studies have relied on techniques used in computer vision literature, namely the deep dream algorithm [107, 108], to infer activation maps. The activation maps depict the features that were considered "salient" by the CNN for prediction. However, the deep dream algorithm is devised for input data with a uniform feature scale, which is not the case for EEG data. The pink noise characteristic of the EEG spectrum is well-known, whereby the lower frequencies have higher power and higher frequencies have lower power [44]. Thereby, the deep dream algorithm would yield mis-leading results by showing higher feature importance for lower frequency spectral features. This conjecture is consistent with findings in our data whereby a CNN was trained to identify the modulatory effect of cardiovascular exercise on the EEG power spectrum while subjects performed a fixed force hand-grip task.

Given the existing caveats in visualizing features for a CNN trained on neuroimaging data, we developed a novel methodology called the cue-combination for Class Activation Map (ccCAM) to generate activation maps signifying the salient features in input data. We believe that using ccCAM along with existing CNN architectures will extend the usability of DL as a reliable data analysis tool for neuroimaging researchers.

## 2.2 Effect of cardiovascular exercise on motor learning

A single bout of cardiovascular exercise, when performed in close temporal proximity to a session of visuo-motor skill practice has been shown to facilitate motor

memory consolidation [35, 130]. The positive effects of exercise on motor memory have been associated with a variety of events at the molecular and systems level. These involve increased concentration of neurotrophin molecules such as brain-derived neurotrophic factors (BDNF), vascular endothelial growth factor (VEGF) and insulin-like growth factor (IGF-1). These factors in-turn mediate downstream effects like neurogenesis and synaptogenesis which form the basis of events underlying neuroplasticity. Increased corticospinal excitability has also been observed during the memory consolidation period and is thought to facilitate synaptic transmission between neuronal networks involved in motor skill practice [117]. However, the precise contribution of distinct brain areas and networks associated with the positive effects of exercise on motor memory consolidation remain largely unknown.

EEG has been widely used to study the electrical activity originating in different brain areas while the subject is performing a task. The EEG signal arises from synchronized postsynaptic potentials of neurons that generate electrophysiological oscillations in different frequency bands. During movement, the EEG signal power spectrum within the alpha (8–12 Hz) and beta (15–29 Hz) range decreases in amplitude and this is thought to reflect increased excitability of neurons in the sensorimotor areas [33, 115, 122, 135]. This phenomenon is termed Event-Related Desynchronization (ERD). Previous studies have reported that alpha and beta-band ERD patterns are modulated during motor skill learning [16, 74, 191]. There is also converging evidence toward an association between cortical oscillations in the motor cortex and neuroplasticity events underlying motor memory consolidation [16, 124]. A recent

EEG study described the oscillatory patterns in brain's electrical activity while subjects performed a handgrip task [35]. Using a time-frequency decomposition-based pipeline, a significant decrease in post-exercise beta-band ERD in EEG electrodes over the sensorimotor area in both hemispheres was found. Additionally, changes in beta-band ERD were associated with better skill retention 24 hr after motor practice. These results suggest that changes in oscillatory patterns take place when motor learning is combined with acute exercise and that some of these brain changes have implications on skill retention.

### 2.3 Brain age prediction

The human brain is known to change across the adult lifespan. This process of *Brain aging* underlies the gradual decline in cognitive performance. Although aging-induced changes are not necessarily pathological, but neurodegenerative disorders have shown to be more likely with increasing age [1]. The wide range of onset ages for age-associated brain disorders indicates that the effects of aging on the brain structure and function vary greatly among individuals. Therefore, understanding the process of brain aging and identifying biomarkers for characterizing the same are vital to improve the detection of early-stage neurodegeneration and age-related cognitive decline.

The effect of aging on brain structure, like cortical thinning, have been robustly identified in previous studies [71, 155]. Similarly, age-related differences in brain function have been demonstrated through studies involving functional connectivity [36, 39]. It is important to establish the trajectories of these changes over the lifespan by means of longitudinal studies to obtain a basis for characterizing clinically relevant

deviations [128, 192]. Brain-based age prediction offers a useful tool to identify personalized biomarkers of age-related brain changes and thereby promises to predict future cognitive impairments by indicating deviations from typical structural and functional development.

Prediction of brain age has been widely explored in literature. Most studies have trained a prediction model on healthy subjects and then used these models to predict brain age in independent clinical samples. By comparing the estimated age of the individual with their chronological age, conclusions about age-typical and atypical brain development can be drawn. For instance, if the *brain-predicted age* is greater than the reported chronological age, it is thought to be an indicator of aberrant accumulation of age-related changes to the brain. The degree of this added brain aging has been more frequently quantified by subtracting chronological age from brain-predicted age. This approach has been extended to the context of several brain-related disorders including Alzheimer's [51, 58], traumatic brain injury [28], schizophrenia [89, 141], HIV [31], epilepsy [119], Down's syndrome [27] and diabetes [52]. Interestingly, the utility of prediciting brain-age is not just limited to understanding neurological disorders. Some studies have shown positive influence of meditation [105] and increased education and physical exercise [151] on the brain age.

Aforementioned studies have shown that the structure and function of one's brain can be both positively and negatively influenced to reflect their age. Therefore, the problem at hand is an interesting one and has far-reaching consequences given the usability of the brain-predicted age as a biomarker to study the impact of various

factors on typical brain development trajectory. Although there's a trend in brain-age prediction to shift from correlative to predictive approach primarily based on machine learning [20, 54, 121, 176, 185], it is important from neuroscience point-of-view to understand the underlying neurological factors influencing these predictions. Added to this, a major roadblock to clinical applications of the machine learning models is their interpretability and reliability [93].

One of the most popular machine learning models used for brain age prediction is the Gaussian Process Regression (GPR) models [127]. GPRs have been popularly used on T1-MRI data from a cohort of healthy subjects but the lack of interpretability has restricted further neuroscientific insights into the process of healthy aging. This has further restricted researchers to draw conclusions on atypical brain aging in the context of the aforementioned brain-related disorders. We sought to establish the contribution of dimensionality reduction techniques, specifically Principal Component Analysis (PCA) [80] and Canonical Correlation Analysis (CCA) [165], on the performance of GPRs and other regression models. Further analysis using PCA or CCA enabled us to observe major areas of the brain that contribute to the prediction of brain age.

Recent studies have aimed to combine both structural and functional brain data to characterize the aging process. These studies have primarily used structural and functional MRI data [23, 102]. With an increase in rich datasets and such high dimensionality data, machine and deep learning algorithms have been explored to predict brain age from neuroimaging data [29]. In this context, we applied dimensionality reduction techniques in conjunction with machine learning models as well as

Graph Convolutional Networks in the context of brain age prediction from MRI and MEG data. We intend to use the features extracted by these models as biomarkers for age-associated health problems, including brain-related disorders.

# CHAPTER 3
# EFFECTS OF CARDIOVASCULAR EXERCISE ON NEURAL CORRELATES OF MOTOR LEARNING

## 3.1 Abstract

Cardiovascular exercise is known to promote the consolidation of newly acquired motor skills. Previous studies seeking to understand the neural correlates underlying motor memory consolidation that is modulated by exercise, have relied so far on using traditional statistical approaches for *a priori* selected features from neuroimaging data, including EEG. With recent advances in machine learning, data-driven techniques such as deep learning have shown great potential for EEG data decoding for brain-computer interfaces, but have not been explored in the context of exercise. Here, we present a novel Convolutional Neural Network (CNN)-based pipeline for analysis of EEG data to study the brain areas and spectral EEG measures modulated by exercise. To the best of our knowledge, this work is the first one to demonstrate the ability of CNNs to be trained in a limited sample size setting. Our approach revealed discriminative spectral features within a refined frequency band (27–29 Hz) as compared to the wider beta bandwidth (15–30 Hz), which is commonly used in data

---

**This work is currently submitted as:**
Ghosh A, Dal Maso F, Roig M, Mitsis GD, Boudrias MH. Unfolding the effects of acute cardiovascular exercise on neural correlates of motor learning using Convolutional Neural Networks.

analyses, as well as corresponding brain regions that were modulated by exercise. These results indicate the presence of finer EEG spectral features that could have been overlooked using conventional hypothesis-driven statistical approaches. Our study thus demonstrates the feasibility of using deep network architectures for neuroimaging analysis, even in small-scale studies, to identify robust brain biomarkers and investigate neuroscience-based questions.

## 3.2 Introduction

A single bout of cardiovascular exercise, when performed in close temporal proximity to a session of visuomotor skill practice has been shown to facilitate motor memory consolidation [35, 130]. The positive effects of exercise on motor memory have been associated with a variety of events at the molecular and systems level. These involve an increased concentration of neurotrophin molecules such as brain-derived neurotrophic factors (BDNF), vascular endothelial growth factor (VEGF) and insulin-like growth factor (IGF-1). In turn, these factors mediate downstream effects like neurogenesis and synaptogenesis, which form the basis of events underlying neuroplasticity. Increased corticospinal excitability has also been observed during the memory consolidation period and is thought to facilitate synaptic transmission between neuronal networks involved in motor skill practice [117]. However, the precise contribution of distinct brain areas and networks associated with the positive effects of exercise on motor memory consolidation remain largely unknown. Understanding how the brain is altered by exercise could hold the key to designing therapies that could optimize the neurophysiological changes associated with motor memory consolidation for varied purposes [68].

Electroencephalography (EEG) is a technique used to study the electrical activity originating in different brain areas. The EEG signal arises from synchronized postsynaptic potentials of neurons that generate electrophysiological oscillations in different frequency bands. During movement, the EEG signal power spectrum within the alpha (8–12 Hz) and beta (15–29 Hz) range decreases in amplitude and this is thought to reflect increased excitability of neurons in the sensorimotor areas [33, 115, 122, 135]. This phenomenon is termed Event-Related Desynchronization (ERD). Various studies have reported that alpha- and beta-band ERD patterns are modulated during motor skill learning [16, 74, 191]. There is also converging evidence towards an association between cortical oscillations in the motor cortex and neuroplasticity events underlying motor memory consolidation [16, 124]. Using EEG, we recently described the oscillatory patterns in brain electrical activity while subjects performed a handgrip task [35]. Using a time-frequency decomposition-based analysis pipeline, we found a significant decrease in post-exercise beta-band ERD in EEG electrodes located over the sensorimotor area in both hemispheres. Additionally, changes in beta-band ERD were associated with better skill retention 24 h after motor practice. These results suggest that changes in brain oscillatory patterns occur when motor learning is combined with acute exercise and that some of these changes have implications for skill retention. However, as these inferences were drawn from a hypothesis-driven approach, whereby standard EEG frequency bands from pre-selected electrodes were considered, the existence of more subtle, fine-scale neurophysiological features that are modulated by a single bout of exercise cannot be excluded.

Neural network models, particularly deep learning (DL) models, have been successful in identifying optimal discriminative features in a given dataset [98]. Convolutional Neural Networks (CNN) and Recurrent Neural Networks (RNN) have been applied to computer vision and speech processing datasets [61, 81, 92, 188] with great success. These approaches have also been used in the context of neuroimaging to identify features for structural and functional magnetic resonance imaging data analysis [123] and EEG data decoding [12, 140, 164] among others.

CNNs are artificial neural networks that can learn low-level patterns in a given dataset by using convolution operations as a key component. CNN architectures may range from shallow architectures with just one convolutional layer [2], deep CNNs with multiple sequential convolutional layers [92] to very deep architectures with over 1000 layers [66]. CNNs have an edge over other machine learning models as they are well suited for end-to-end learning, i.e. learning from raw data without any *a priori* feature selection and they can exploit hierarchical structures that may be present in the data.

Although the field of EEG signal decoding has recently seen a surge of papers [166, 184, 186] involving the use of DL, the applicability of DL models has been primarily restricted to the classification of EEG data segments into known categories. The usefulness of CNNs to improve our understanding of the neural substrates underlying observed behaviors is less straightforward, primarily due to the difficulty associated with the visualization and interpretation of the feature space learned by DL architectures, *e.g.* CNNs. For instance, Schirrmeister *et al.* ([140]) proposed a systematic CNN framework for EEG decoding, including the impact of various

architectural considerations on decoding performance. However, they presented a feature validation approach to understand which *a priori* selected features were used by the CNN rather than a feature discovery approach.

In this context, the aim of the present study was to develop a data-driven approach for studying the positive effects of exercise on motor learning by investigating the EEG-based ERD patterns during an isometric motor grip execution in healthy young subjects. We aimed to identify specific EEG spectral features modulated by exercise and further investigate if these features were related to skill retention performance. The subjects performed a repetition of isometric handgrips before and after a session of intense cycling exercise (exercise group) or rest for the same period (control group). In addition to isometric handgrips, subjects also practiced a new motor tracking task with their dominant hand in close proximity to the exercise or rest session. To identify the neurophysiological substrates underlying the positive effects of exercise, we used a CNN-based deep network architecture to identify exercise-induced changes in neural activity from EEG signals recorded during the handgrip task. Since neural networks are known to be universal function approximators [73], with the capability of identifying linear as well as non-linear boundaries in high-dimensional data spaces, this allowed us to differentiate the exercise and control groups in the EEG time-frequency data space. The training was carried out in a hierarchical structure – initially in the time-frequency domain and subsequently for topographical maps of ERD pattern. Visualizing the features after each stage of training allowed us to identify frequency bands as well as the corresponding brain regions modulated by the positive effects of acute exercise on motor learning.

Moreover, the majority of previous related DL studies used datasets comprising of hundreds of subjects for training purposes [123, 138]. Therefore, one of our main goals was to develop a DL-based method that is suitable for neuroimaging studies with smaller subject numbers, which is frequently the case. To this end, we added a regularizer (adversary component) to the CNN, which prevented the latter from learning subject-specific features, thus favoring the identification of group-specific features. The proposed approach revealed that the CNN-extracted features were strongly correlated to the improvement in motor learning scores. Visualizing these features revealed finer frequency bands and corresponding brain regions where ERD patterns were modulated by exercise. Therefore, the proposed analysis provided observational evidence for the identified frequency band-related ERD to be associated with the positive effects of exercise on motor memory consolidation.

### 3.3 Dataset

The experiment and data collection are described in detail in [35]. Briefly, 25 right-handed healthy subjects were recruited and assigned to the Control (CON, $n = 13$ subjects) or Exercise (EXE, $n = 12$ subjects) groups in matched blocks. The blocks of subjects were created with similar age, gender, body mass index, working and episodic memory as well as cardiorespiratory fitness. All subjects signed a written consent form according to the research protocol that complied with the recommendations of the declaration of Helsinki for investigation of human participants and was approved by our local ethics committee (CRIR-1134-0116).

Each subject reported to the laboratory on four occasions as described in [35]. Visit 1 required the participants to go through a Graded Exercise Test (GXT), which

was used to determine their cardiorespiratory fitness. Visit 2 was conducted at least 48 hrs after the GXT to avoid potential long-term effects of exercise on memory [14, 72]. EEG recordings were collected at baseline (before the exercise session) while subjects performed repetitions of visually cued isometric handgrips with their dominant right hand using a hand clench dynamometer (Biopac, Goleta, CA, USA). Each contraction was maintained for 3.5 sec at 15% of each participant's maximum voluntary contraction (MVC). This was followed by a 3 to 5 sec rest period. The baseline assessment was followed by the practice of a visuomotor tracking task (skill acquisition), which was used to calculate the motor learning score of each subject. Following the training period, participants were randomly assigned to two groups. The EXE group performed a bout of high-intensity interval cycling of 15 min, while the CON group rested on the cycle ergometer for the same amount of time. EEG recordings similar to baseline were repeated 30, 60 and 90 min after the exercise or rest period. During visits 3 and 4, two blocks of the visuomotor tracking task were performed 8 and 24 hrs after the exercise or rest period.

EEG activity was recorded using a 64-channel ActiCap cap (BrainVision, Munich, Germany) with electrode locations arranged according to the 10–20 international system. The electrical conductive gel was inserted at each electrode site to keep impedances below 5 kΩ. EEG signals were referenced to the FCz electrode and sampled at 2500 Hz.

### 3.4 Methods

The analysis pipeline was first applied to the time and frequency domain data without incorporating spatial information. Subsequently, it was applied to the data

obtained by creating topographical maps corresponding to the distribution of activity within specific frequency bands across the cortex. The entire pipeline consisted of 3 segments, i.e.– Preprocessing, CNN training and cue-combination for Class Activation Map (ccCAM) generation.

### 3.4.1 Time-Frequency (TF) maps

#### Preprocessing

EEG data preprocessing was similar to that described previously [35] and was performed using the Brainstorm Matlab toolbox [159]. The preprocessing pipeline is summarized in Figure 3.1. Briefly, EEG signals were bandpass-filtered between 0.5 Hz and 55 Hz and average-referenced. The data were visually inspected and signal segments with artifacts were rejected. Independent component analysis (ICA) was subsequently applied (total number of components: 20) and eye-blink related components were rejected based on their topography and time signatures [38]. The resulting data were epoched with respect to the period of time (3.5 sec) corresponding to the appearance of the visual cue that triggered the initiation of the isometric handgrips (n = 50/subject). Finally, each epoch was visually inspected and those containing artifacts were manually removed. EEG electrodes with atypical power spectrum density were interpolated using spherical splines.

Morlet wavelet (wave number = 7) coefficients between 1 to 55 Hz with 1 Hz resolution were extracted to obtain time-frequency decompositions of the EEG data [161]. The time-frequency data for each electrode were normalized with respect to their spectral power before the start of the grip event, using a window of 0.5 sec. An average over all trials was then calculated in order to obtain a single time-frequency

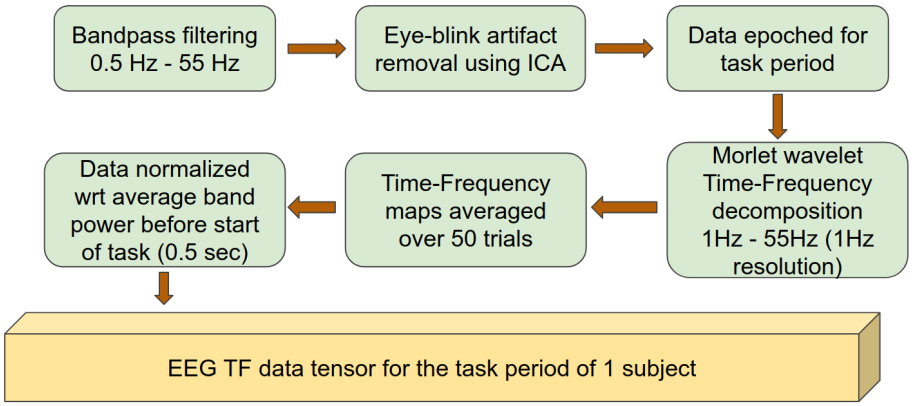


Figure 3.1: Data preprocessing pipeline outlining the preparation of data to be used as input to the deep network.

map for each electrode. Subsequent analysis was performed on the EEG recording segment corresponding to 0.5–3.5 sec after the appearance of the visual cue, i.e. during the handgrip task. Finally, one time-frequency map was obtained for each electrode for each session (baseline, 30 min, 60 min or 90 min after exercise) and for each subject.

**CNN training**

The proposed CNN architecture is shown in Figure 3.2. The preprocessed data for each session was rearranged to form 2D matrices comprising of the frequency spectra for all electrodes at a given time instant t. Each matrix had a dimension of $64 \times 55$ (64 electrodes $\times$ 55 frequency bands). For training the network, a pair of matrices was used – the first corresponding to time point t from the baseline session and the second corresponding to the same time point t from the post-intervention session. Each pair was labeled based on the respective group allocation (EXE or CON). Structuring the data in this fashion allowed the network to take into account

the inter-subject variability in baseline measures and therefore did not require the experimenter to adopt techniques for normalizing the EEG signal from the post-intervention session with respect to the baseline session. Thus, the network was expected to capture the EEG features that were modulated by the effects of acute exercise.

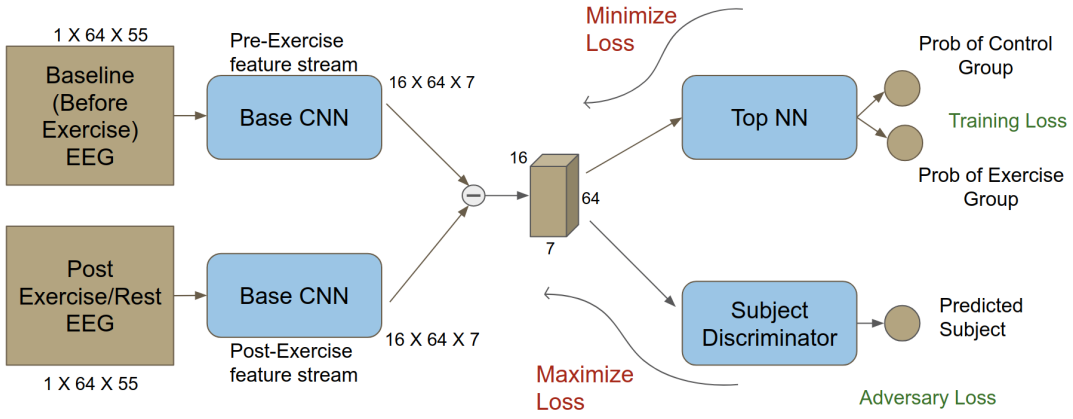


Figure 3.2: Modified deep network architecture with an adversary component (bottom right). The adversary component makes it feasible to use CNNs for identifying subject-invariant features in neuroimaging studies with a limited number of subjects. Dimensions corresponding to the obtained TF maps are also shown.

*Dataset Notation*:- $\boldsymbol{B}$ and $\boldsymbol{A}$ represent the entire data tensor at baseline and post-intervention respectively. Each data tensor consists of data matrices from all 25 subjects and timepoints. For subject $s$, the goal was to classify whether the tuple containing the matrices $B_t^s$ and $A_t^s$ (where $t$ denotes timepoint) belongs to the EXE or CON groups. Matrices $B_t^s$ and $A_t^s$ were arranged so that they belonged to the set $R^{64X55}$, where 64 is the number of electrodes and 55 is the number of frequency bands.

To identify EEG features modulated by exercise, we used a deep CNN that was trained to discriminate between EEG data from EXE and CON group. The network architecture is similar to the one described in [3]. Features from matrices $B_t^s$ and $A_t^s$ were extracted using a network termed Base CNN. The difference between the obtained feature vectors was passed to a discriminator network, termed Top NN, to predict the correct group to which each pair belonged to. The schematic view of the architecture is shown in Figure 3.2 and the details for each network's architecture are provided in Tables S1 and S2 in Supplementary Material respectively. Using a sampling frequency of 2500 Hz and a time period of interest of 3 sec duration, each tuple input to the CNN was of the form ($B_t^s$,$A_t^s$) where timepoint $t \in$ [1,7500] and subject $s \in$ [1,25].

The convolutions performed in the Base CNN were with respect to the frequency direction and not the electrode (sensor) dimension. This is done because the frequency dimension was by definition arranged in terms of increasing frequencies, as opposed to the electrode dimension, which was not arranged in terms of the spatial locations of the electrodes in a meaningful manner. Consequently, the features extracted by the Base CNN corresponded to the frequency bands significantly affected by exercise. Therefore, all convolutional filters in the Base CNN were implemented as $1 \times n$ 2D filters, where $n$ is the extent of the filter in the frequency domain. The same holds for the Max-Pooling layers.

Initially, a network that did not include an adversary loss component (Figure A1.a from Appendix A) was used; however, it was found that this network was able to learn subject-specific features as opposed to subject-invariant, exercise-related

features. This is illustrated in Figure A2 (Appendix A) and Table 3.2. In most neuroimaging studies examining the effect of exercise, the number of participants scanned was relatively low [62, 75, 129], which typically prevents deep networks from learning subject-invariant features. To address this issue, we followed a domain adaptation approach. Specifically, each subject was considered as a separate domain comprising of subject-specific features along with subject-invariant, exercise-related features. Since our goal was to learn features mainly related to the effect of exercise on the consolidation of motor memory, we incorporated the domain confusion strategy [169] to train the network, thus adding the subject discriminator as an adversary (Figure 3.2 – bottom right). Specifically, we included this network in parallel to the Top NN with similar model capacity (see Table A3 in Appendix A for architecture details).

*Network Architecture Notation*:- The feature extractor operation and parameters of the Base CNN are denoted as $f_{\theta_f}$ and $\theta_f$ respectively, the Top NN feature discrimination operator and its parameters by $h_{\theta_t}$ and $\theta_t$ respectively, while the subject discrimination operator and its parameters by $h_{\theta_s}$ and $\theta_s$ respectively. The input tuple is denoted by $x$ and its corresponding group and subject labels by $y_g$ and $y_s$ respectively. We used the Negative Log Likelihood (NLL) loss for each classifier with the Adam optimizer [84] in Torch [32] for training the network. The Subject Discriminator was trained to minimize the subject prediction loss given by –

$$J_s(\theta_s, \theta_f) = -[\sum_{i=1}^{m} log h_{\theta_s}^{(y_s^{(i)})}(f_{\theta_f}(x^{(i)}))] \tag{3.1}$$

The Top NN was trained to minimize the group prediction loss given by –

$$J_g(\theta_t, \theta_f) = -[\sum_{i=1}^{m} log h_{\theta_t}^{(y_g^{(i)})}(f_{\theta_f}(x^{(i)}))] \tag{3.2}$$

We trained the feature extractor (Base CNN) to extract features that would be agnostic to the originating subject; therefore, the target distribution for the subject prediction network had a uniform distribution. Hence, we used the domain confusion loss [169] over the gradient reversal layer [57]. We also used the Kullback-Leibler (KL) divergence from the uniform distribution over 25 classes (25 subjects) as our loss metric. Conclusively, the Base CNN was trained to minimize the loss given by –

$$J_f(\theta_f, \theta_t, \theta_s) = -[\sum_{i=1}^{m} log h_{\theta_t}^{(y_g^{(i)})}(f_{\theta_f}(x^{(i)}))] + \lambda[\sum_{i=1}^{m} KL(U, h_{\theta_s}(f_{\theta_f}(x^{(i)}))] \tag{3.3}$$

where $KL(P, Q)$ denotes the KL divergence between distributions $P$ & $Q$, $U$ denotes the uniform distribution, $m$ denotes the total number of training examples, and $\lambda$ is a hyperparameter that determines the weight for the subject discrimination regularizer. Here, we used a 80-20 split of the data set, whereby 80% was used for training and 20% was used for validation.

**ccCAM**

An innovative contribution of the present work is the development of a novel method for the visualization of the features that guide the proposed network's decision. Although well-known techniques used in the computer vision literature include Global Average Pooling (GAP) [190] and grad-CAM [144], they are not well suited

to the neurophysiological paradigm considered here. For instance, GAP requires averaging the activations of each filter map, i.e. each channel of the extracted feature tensor. This leads to loss of information related to electrode positions, as convolutions were performed only in the frequency domain. Consequently, we were unable to obtain adequate classification accuracy ($\approx 56\%$) with a GAP layer in the network. Grad-CAM is sensitive to the absolute scale of variability in the features in the input data and, as a result, it yielded results that were biased towards frequency bands with higher power-values, i.e. the lower frequency bands (¡10 Hz).

Given the above limitations in existing analytic methods, we used the linear cue-combination theory applied in human perception studies [47] to develop a method that improves the interpretability of the network's decisions. Let us consider, for example, a CNN with only 2 channels, i.e. filter maps, in the final feature tensor extracted after convolutions. Each of these filter maps preserves the spatial and/or semantic structure of the input data. Each of these filter maps acts as a cue to the network's classifier layers, denoted as $c_1$ and $c_2$. These cues guide the networks prediction. If we denote the desired class label as $y_1$ and assuming $c_1$ and $c_2$ to be independent to each other, we can use Bayes' Theorem to write –

$$P(y_1|c_1, c_2) = \frac{P(c_1, c_2|y_1)P(y_1)}{P(c_1, c_2)} = \frac{P(c_1|y_1)P(c_2|y_1)P(y_1)}{P(c_1)P(c_2)} = \frac{P(y_1|c_1)P(y_1|c_2)}{P(y_1)} \tag{3.4}$$

If the likelihood for predicting $y_1$ due to cue $c_i$ is Gaussian with mean $\mu_i$ and variance $\sigma_i^2$, the maximum likelihood estimate (MLE) yields the combined cue, denoted by $c^*$, that summarizes the important features on which the network bases

its decisions. Therefore, the combined cue, $c^*$, is the desired Class Activation Map (CAM).

$$c^* = \sum_{i=1}^{2} w_i \mu_i \text{ where } w_i = \frac{1/\sigma_i^2}{\sum_{i=1}^{2} 1/\sigma_i^2} \tag{3.5}$$

Since the network is trained, $\mu_i = c_i$. To calculate the values of $\sigma_i$, we used the NLL loss values. The NLL loss when a cue was removed from the network provided an estimate of the $\sigma$ associated with that specific cue, as shown in Equation 3.6.

$$\begin{aligned}
\epsilon &= -logP(y_1|c_1, c_2) \\
&= -logP(y_1|c_1) - logP(y_1|c_2) + logP(y_1) \text{ [From eq 3.4]} \\
\epsilon_1 &= \epsilon|_{c_1=0} = -logP(y_1|c_1 = 0) - logP(y_1|c_2) + logP(y_1) \\
\epsilon_1 - \epsilon &= logP(y_1|c_1) - logP(y_1|c_1 = 0) \\
&= \frac{\mu_1^2}{2\sigma_1^2} \\
\text{Therefore, } \frac{1}{\sigma_1^2} &= \frac{2(\epsilon_1 - \epsilon)}{\mu_1^2}
\end{aligned} \tag{3.6}$$

$\sigma_i$ is estimated over the entire dataset as shown in Equation 3.7.

$$\frac{1}{\sigma_i^2} = \sum_{j=1}^{m} \frac{2[(\epsilon_i - \epsilon)]^{(j)}}{[\mu_i^2]^{(j)}} \tag{3.7}$$

Using the estimated $\sigma_i$, the CAM corresponding to the correct class for each input was obtained. Since in the present case $\mu_i$ corresponded to a 2D matrix, the denominator in Equation 3.7 was replaced by the mean-squared value of the corresponding matrix. A summary of the process of generating a CAM is outlined in Figure 3.3.

The resulting CAM was a 2D matrix with each row corresponding to an electrode and each column corresponding to a frequency band. To obtain the contribution of each frequency band in determining the correct class, we averaged the CAM along the row dimension to yield a vector corresponding to the importance of each frequency band power across the whole brain. This vector was generated at each timepoint t, depicting the importance values of each frequency band at timepoint t. Since the subject was exerting a fixed force during the entire time period of the EEG segment that was considered, we expected important frequency bands to exhibit pronounced differences at all timepoints. Consequently, we considered a frequency band to be reliably modulated by exercise only if the corresponding CAM value was high for all timepoints. Therefore, we followed a procedure similar to the bootstrapping technique to estimate the reliability of each frequency band [45, 112, 111]. Specifically, we calculated the bootstrapped ratio (BSR) of the CAM values for each band and subject by dividing the mean of CAM values across time by its standard deviation. The obtained BSR values were subsequently group-averaged to extract frequency bands that contained features characteristic to each group (CON and EXE).

### 3.4.2 Topographical maps

Topographical maps were created using the frequency bands obtained from the ccCAM corresponding to the TF-maps following the same procedure as described in [12]. Specifically, the 3D position of the electrodes on the EEG cap was projected to a 2D space to construct a $64 \times 64$ image. The image intensities corresponded to the respective electrodes spectral power within the frequency band of interest at time point $t$. Since this procedure yielded a sparse matrix, cubic interpolation was used to

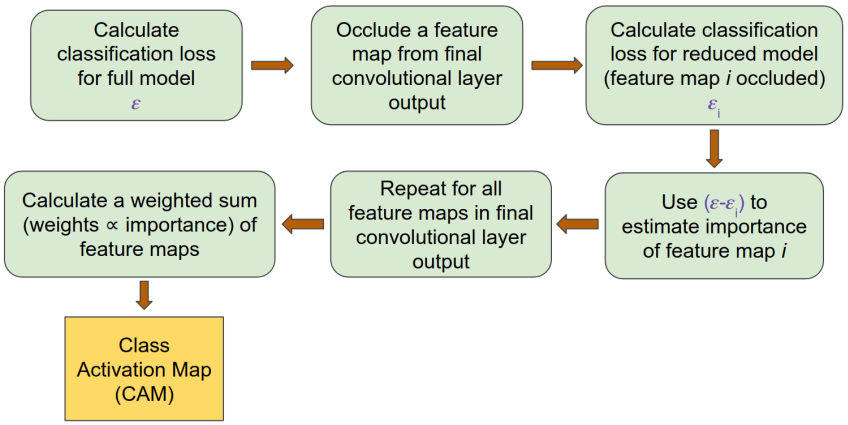


Figure 3.3: ccCAM generation pipeline outlining the procedure for obtaining CAMs from the trained deep network.

obtain a continuous image depicting the distribution of activity within each frequency band over the entire scalp. A total of three such matrices were concatenated to form a $3 \times 64 \times 64$ tensor corresponding to activity maps at three consecutive timepoints (say $t$, $t+1$ and $t+2$ respectively). The entire data tensor for a given subject was created by taking non-overlapping time windows. Hence, the total number of tensors for each subject was equal to 2500.

Similar to the analysis of TF-maps, we trained a CNN-based network to classify each data tensor into the CON and EXE groups. Since the inputs were 2D image tensors, we used 2D convolutional filters in the Base CNN (see Tables A4, A5 and A6 in Appendix A for more details). Following training, ccCAM was applied to obtain CAMs for each subject at each time instant during task execution.

### 3.4.3 Statistical analysis for TF curves and Topographical maps

To estimate the statistically significant frequency bands in the resultant ccCAM maps for the two groups, we employed one-way ANOVA on a point-by-point basis.

This yielded a series of $p$-values, each corresponding to a frequency bin in the ccCAM map. Since these tests were not mutually independent, we chose the Simes method [136, 147] for correcting the obtained $p$-values as done in [104]. This method was better suited to our problem compared to the standard Bonferroni correction, as several highly correlated statistics were involved. The method is based on the ordered $p$-values in a sliding window of length $L$ with the $p$-value we want to correct being in the center ($L = 2w + 1$). The corrected $p$-value is given by

$$p'_j = \min_{1 \leq i \leq L} \frac{L p_{(i)}}{i} \tag{3.8}$$

Given the employed convolutional kernel sizes for the TF and topographical maps, we chose the corresponding window lengths for the Simes method to be 5 for TF maps and $3 \times 3$ for the topographical maps. All points with corrected $p$-values less than 0.05 were considered to be significant.

## 3.5 Results

The results presented here illustrate the differences between the baseline and 90 min post-exercise/rest datasets. The network architecture details for each type of data (TF and Topographical) map are presented in the Supplementary Material, along with details regarding the chosen hyperparameters.

### 3.5.1 Comparison to alternative EEG decoding methods

We compared the performance of the proposed CNN to alternative machine learning methods on a leave-two-out cross-validation strategy by using 23 subjects for training the model and 2 subjects (1 from CON and 1 from EXE) for validating the model performance on unseen data. Initially, we compared the proposed CNN

architectures to random forests (RF) [18] and support vector machines (SVM) [152] given the same input data. The proposed CNN with a subject adversary component clearly outperformed RF and SVM (Table 3.1). A feature importance analysis showed that RF was unable to account for the pink noise characteristics of EEG in the frequency domain [44] and therefore assigned more importance to lower frequencies. To control for this, the data were normalized by multiplying power in each frequency bin by the frequency value itself before being fed to the RF but this did not yield significant improvement.

Subsequently, the data were binned into pre-specified frequency bins, as done widely in EEG decoding techniques [11], and used as input in a RF. Although this improved the RF prediction performance, the resulting RF prediction accuracy was lower than the obtained CNN prediction accuracy (Table 3.1). Interestingly, the RF prediction accuracy was very similar to an adversary-less CNN. We also compared the performance using the common spatial pattern technique, which is widely used in EEG decoding for brain-computer interfaces [6]. However, the prediction accuracy was far below the CNN prediction accuracy. Overall, these results imply that the CNN architecture with subject adversary was able to generalize better across subjects compared to the alternative, widely used EEG decoding methods. The CNN leave-two-out prediction accuracy was equal to **65.4%**. The summary of the results is presented in Table 3.1.

### 3.5.2 Time-Frequency maps

We observed that the features extracted by the Base CNN, without any subject prediction regularizer, were able to perfectly identify the subject corresponding from

Table 3.1: Comparison of CNN performance with alternative machine learning methods

| **Method** | **Average validation accuracy over 10 folds** |
|---|---|
| Random Forests | 58.94 % |
| Random Forests with frequency bands | 60.13 % |
| Random Forests with Common Spatial Patterns | 45.94 % |
| CNN without baseline normalization architecture | 56.59 % |
| CNN without subject adversary | 60 % |
| **CNN with subject adversary** | **65.4** % |

the obtained time-frequency patterns. As the subject discriminator regularization was given more weight by increasing $\lambda$, the Base CNN learned to extract features that were agnostic to the originating subject. However, for very high $\lambda$ values, the extracted features could not be used to discriminate the EXE and CON groups, suggesting that the Base CNN was unable to learn any discriminative feature. The loss values obtained post-training for four different values of $\lambda$ are shown in Table 3.2. The choice of an optimal value for $\lambda$ depends on two factors – group prediction accuracy and subject prediction accuracy. To identify subject-invariant features, we sought for a value of $\lambda$ that achieved good group prediction accuracy and poor subject prediction accuracy (*i.e.*, a good tradeoff between the two prediction accuracies). According to this procedure, the model corresponding to $\lambda = 13$ was used for the ccCAM generation. The average loss over a batch for subject prediction was around 2.6 (Table 3.2), which roughly predicted the correct subject with probability of $\frac{1}{13}$ (Since $-log(\frac{1}{13}) \approx 2.6$). The group prediction accuracy was **99.984%** (**99.969%** for CON and **100%** for EXE) when evaluated using a train-test split strategy following

Table 3.2: Variation of Loss values with $\lambda$ after training network on TF maps.

| $\lambda$ | Group prediction loss (NLL) | Subject prediction loss (NLL) | KL divergence loss from Uniform distribtion |
|---|---|---|---|
| 0 | $\approx 0$ | $\approx 0$ | $\approx 0.3$ |
| 10 | $\approx 0.1$ | $\approx 1.5$ | $\approx 0.07$ |
| 13 | $\approx 0.4$ | $\approx 2.6$ | $\approx 0.004$ |
| 15 | $\approx 0.68$ | $\approx 3.2$ | $\approx 0.0002$ |

80-20 split of all timepoints from all subjects. Since we aimed to obtain a model that best explained the recorded EEG data, we focused on timepoint-level accuracy, instead of subject-level accuracy as before (Table 3.1). Therefore, we trained on 80% of timepoints from all subjects and validated the model performance on 20% unseen timepoints from all subjects. For $\lambda = 13$, the extracted features achieved excellent group prediction, while all subjects in the group were predicted with roughly equal probability (CON and EXE consisted of 13 and 12 subjects respectively).

As one of the main goals of this study was to identify the frequency bands that contained significant discriminative information between the CON and EXE groups, we calculated the BSR of the CAM values for each frequency bin as described above. The difference obtained using ccCAM BSRs is shown in Figure 3.4. The bold lines denote the group-mean and the shaded regions span 1 standard error over all subjects in the group. The two plots are significantly different within the band 23–33 Hz. A statistical significance analysis between the two curves with Simes correction revealed that the band 27–29 Hz lies below the significance threshold of 0.05. The uncorrected and corrected $p$-values for each frequency bin are shown in Figure 3.4. The aforementioned band lies within the beta-band and agrees with

findings in [35], where beta-band desynchronization was found to be significantly modulated by exercise. Note that Figure 3.4 corresponds to the differences between the 90min and baseline EEG recordings. In addition, the features extracted by the CNN were found to be strongly correlated to the motor skill retention improvement between the 8 hr and 24 hr session (Figure 3.5).

The CNN was trained to classify the EXE and CON groups given the baseline EEG and 90 min post-exercise EEG recordings. To observe how the discriminative features evolved over time (30 min and 60 min post-exercise), we used the t-SNE algorithm [106] to obtain a lower dimensional (2D) representation of the feature vectors extracted by the trained CNN for all sessions and subjects. Figure 3.6 shows the corresponding plot with each point representing one session from one subject, color-coded with respect to session and group. The differences between the features extracted from the 90 min post-exercise and baseline sessions were more pronounced as compared to the EXE group as compared to the CON group. Also, the difference between the two groups becomes prominent from the 30 min post-exercise session itself, with the feature trajectories being considerably different for the two groups.

### 3.5.3 Topographical Maps

Topographical maps were created to study the distribution of the activity within the 23–33 Hz frequency band across the cortex. Since we used a convolutional filter of size 5 in the BaseCNN for TF-maps, we selected the wider 23–33 Hz instead of the finer 27–29 Hz band, and projected this band-limited activity on the cortex to obtain topographical maps. After training a network to classify the CON and EXE groups from the resulting topographical maps, a classification accuracy of **98.70%** (**98.94%**

for CON and **98.43%** for EXE) was obtained for $\lambda = 5$. Generating ccCAMs for the topographical maps revealed brain areas where the activity was notably different between the CON and EXE groups. The BSR values were obtained to estimate the reliability as in the case of TF maps and the results are shown in Figure 3.7. Statistical analysis to highlight areas with significantly different activity within the 23–33 Hz band is shown in Figure 3.7. Notable areas that exhibited significant differences include the contralateral and ipsilateral sensorimotor areas, the contralateral prefrontal area and the occipital areas. The observed differences in occipital areas are not surprising as the occipital cortex is primarily responsible for visual information processing and the task under consideration is a visuomotor task. Therefore, these differences could also highlight a change in functional connectivity between the sensorimotor and occipital cortices. Findings in sensorimotor and prefrontal areas are in strong agreement with those reported in Dal Maso *et al.* and therefore suggest that the proposed method was able to identify the modulatory effect of exercise [35]. In addition, the topographical features extracted by the CNN were strongly correlated to the motor skill retention improvement between the 8 hr and 24 hr sessions (Figure 3.8), indicating a possible association of these features to the observed motor learning improvements in the EXE group [35].

### 3.6 Discussion

In the present work, we used a novel DL approach to investigate the effects of acute cardiovascular exercise on brain activity during the early stages of motor memory consolidation while subjects performed isometric handgrips. The proposed methodological approach does not require the specification of *a priori* features and

included a novel feature visualization technique to highlight the neurophysiological patterns modulated by exercise. Our approach addressed existing caveats related to the application of DL architectures, such as CNNs, to EEG data recorded from a relatively small number of subjects by means of three novel contributions:

1. We used two parallel feature extraction streams to identify informative features from EEG data before and after a session of cardiovascular exercise and subsequently characterize the modulatory effect on these baseline-corrected features rather than on the raw EEG data;
2. We incorporated a subject prediction adversary component in the network architecture to learn subject-invariant, group-related features instead of subject-specific features;
3. We developed a novel feature visualization method, termed cue-combination for Class Activation Map (ccCAM).

Previous studies using DL in the context of structural and functional neuroimaging [12, 123, 140, 164] have been primarily restricted to classification tasks and have relied on a large cohort of subjects for training. The interpretation of the results obtained from these studies and their usefulness for investigating neuroscientific questions has been less straightforward, primarily due to the difficulty associated with the visualization and interpretation of the feature space learned by the employed DL architectures. Schirrmeister *et al.* proposed a feature validation approach to understand which *a priori* selected features were given importance by a CNN that was trained to decode imagined or executed tasks from raw EEG [140]. Specifically, their method relied on calculating correlations between the output of each unit or layer of

the network and EEG power in specific frequency bands. This allowed them to verify if the network used the power in *a priori* decided frequency bands for its predictions. Additional studies involving CNNs for EEG decoding have extended the deep dream algorithm [107, 108] for identifying discriminative features in EEG time segments [125]. However, this algorithm was found to be sensitive to the scale of features and therefore may not be applicable to a diverse range of neuroimaging studies. The proposed ccCAM methodology is a feature discovery technique that is less sensitive to the feature scale and thereby allows extending the applicability of CNNs beyond classification tasks. Furthermore, our study introduced a novel adversary component to prevent the CNN from exclusively learning subject-specific features. This allowed the DL pipeline to learn group-specific features from the EEG data and therefore improved the generalizability of the CNN from a limited cohort size. Overall, to the best of our knowledge, our study is the first to investigate neuroscience-based questions, specifically related to motor learning, from a limited cohort size using a DL approach.

Our analysis using CNNs outperformed alternative machine learning and EEG decoding methods with regards to differentiating between previously unseen EXE and CON subjects, thus indicating the improved generalization properties of the proposed method. The proposed ccCAM method was able to identify a finer frequency band (27–29 Hz) that was modulated by exercise, as opposed to the entire beta band reported in [35], without using any prior knowledge. Applying ccCAM to the topographical patterns of the electrophysiological activity between 23–33 Hz revealed the specific brain areas (the contralateral and ipsilateral sensorimotor areas,

contralateral prefrontal area and occipital areas) that were mostly impacted by exercise in agreement to previous studies [35]. Correlation analysis between the extracted features and the motor learning scores yielded further evidence that these features are plausible neurophysiological substrates underlying the positive effects of exercise on motor learning.

### 3.6.1 Comparison to alternative EEG decoding methods

A comparison of the proposed DL pipeline to alternative machine learning methods, such as RF showed that it yielded improved generalization performance to unseen subjects (Table 3.1). In turn, this suggests that the DL pipeline was able to learn features that were specific to the exercise intervention. Specifically, the performance of RF was close to that of CNNs without an adversary. This suggests the importance of the adversary component in learning subject-invariant, group-specific features. A feature importance analysis for the RF models indicated that higher importance was assigned to features corresponding to lower frequency bands, which are more generally associated with cognitive behaviors [78]. Therefore, the absence of an adversary component possibly caused the examined machine learning models to assign higher importance scores to these subject-specific features, which were more likely reflective of the cognitive state of the subject while performing the task. In turn, this may explain their more limited generalizability to unseen subjects.

The CNN leave-two-out prediction accuracy was equal to **65.4%**, which was lower compared to the accuracy achieved by CNNs in computer vision or motor imagery EEG datasets [140]. We believe that the main reason for this is the number of subjects and that adding more subjects would yield improved accuracy. Furthermore,

we did not compare the classification accuracy of CNN architectures that have been popularly used for EEG decoding because those networks were designed specifically for raw time-series data [97, 140]. The main focus of our work was to train a CNN that could discriminate the two groups from their spectral EEG patterns before and after an exercise session.

### 3.6.2 Time-Frequency maps

The ccCAM approach was able to identify the frequency bands that were significantly different in terms of ERD patterns between the EXE and CON groups (Figure 3.4). Specifically, the frequency band between 27–29 Hz, which is a subset of the wider beta band (15–29 Hz) typically used in previous studies related to motor activity-associated ERD, yielded the most significant differences. This finding is in agreement with [35] and implies that decreased neural excitability was needed to perform the handgrip task after exercise. The $p$-value calculated from the obtained time-frequency data within the 27–29 Hz frequency band was equal to **0.0044**. This suggests that if the band of interest in [35] had been chosen to be 27–29 Hz, instead of the entire beta-band (15–29 Hz), similar statistically significant results would have been obtained. The standard error values in the ccCAM reliability plots (Figure 3.4) indicate that the frequency band modulated is variable among subjects. Overall, the proposed DL-based analysis was able to identify a narrower frequency band modulated by exercise without using any prior knowledge. This finding has particular importance for the development of targeted interventions, such as non-invasive stimulation protocols, whereby it is desirable to modulate the frequency band of interest without interfering with other frequencies.

The trajectory of the extracted features was significantly different for the CON and EXE groups (Figure 3.6), indicating that the effects of exercise were observable during the consolidation period. However, it cannot be concluded from these results whether these changes are short- or long-term. Measuring the brain activity of the subjects at a later timepoint (e.g. 24 hours later) may reveal further insights about the evolutionary patterns of exercise-induced changes during the retention period.

An important goal of the present study was to identify whether there was an association between the identified brain features and the learning scores of each subject. The features indeed were found to be strongly correlated to the motor learning improvement assessed by the difference between the skill retention scores at 8 and 24 hours after motor practice (Figure 3.5). As reported by Dal Maso *et al.* [35], these effects probably indicate the effect of sleep on motor memory consolidation [83]. This strongly suggests that the features extracted by the CNN corresponded to exercise-induced changes while carrying predictive information about the subjects motor skill retention abilities following a period of sleep.

### 3.6.3 Topographical Maps

The distribution of the discriminatory 23–33 Hz band power across the brain was found to be localized instead of being widespread across the brain (Figure 3.7), suggesting that activity in specific brain regions was modulated by exercise. Notable areas with differential activity included the contralateral and ipsilateral sensorimotor areas, contralateral prefrontal area and occipital areas. These results are in strong agreement with the inferences drawn in [35]. Interestingly, the analysis pipeline employed in the present work uncovered the areas that were differentially activated in

the two groups, whereas the same inferences required the use of 3 different metrics in the standard analysis performed by Dal Maso *et al.* Specifically, in [35] the differences in ERD were observed in the sensorimotor areas, differential functional connectivity was observed among the sensorimotor and occipital areas and the ERD in contralateral prefrontal area electrode activity was shown to have strong correlations to motor score. In comparison, using only the ERD values from the baseline and post-exercise sessions, the proposed DL approach was able to identify the brain regions modulated by exercise without any *a priori* knowledge.

Similar to what was reported in Dal Maso *et al.*, we found reduced ERD in the sensorimotor areas as well as the contralateral prefrontal area, suggesting a more efficient use of neural substrates involved in motor memory consolidation [175]. Consequently, these results indicated that after exercise, reduced neural excitability in these areas was required to perform the fixed force handgrip task. As discussed in Dal Maso *et al.*, the observed decrease in ERD could also be indicative of a reduction in gamma-aminobutyric acid (GABA) inhibitory activity due to exercise [148]. Given that the task was a visuomotor one, we also expected changes in the visual areas. Dal Maso *et al.* reported these changes by using functional connectivity analysis within the beta band. However, using a finer band revealed ERD changes in these areas and these could be indicative of changes in visual attention activity and perception. The range of beta band activity commonly used in visual studies is 15–25 Hz and is thought to be a carrier for visual attention, whereas gamma band (30–60 Hz) activity is thought to be responsible for visual synchronization and perception [182]. Interestingly, the range of frequencies uncovered by our analysis included high beta

activity and low gamma activity. Therefore, the proposed DL pipeline highlighted a modulation of these properties by a session of exercise. As we used a visuomotor task, it is difficult to delineate the modulation effects in visual areas from motor areas and thereby draw conclusive inferences about findings in visual areas. The significantly different activity in the prefrontal area aligns well with prior literature demonstrating the role of the dorsolateral prefrontal cortex in motor memory consolidation [56]. Therefore, it is plausible that acute cardiovascular exercise promotes the efficient distribution of neural resources in the prefrontal area [41], thus reducing neural demands of cognitive processes that underlie the consolidation of motor memory. Interestingly, the CNN-extracted topographical features are well correlated to motor learning improvement measured after a period of sleep (Figure 3.8), i.e. once the memory has been well consolidated [131]. Taken together, the identified features are indicative of the resultant consolidated memory and not necessarily behavior during the memory consolidation period.

The presented results provide observational evidence for the extracted features to be considered as candidate neurophysiological substrates underlying motor memory consolidation. As argued by Tonegawa *et al.* [167], observational studies demonstrate a correlation between specific neural activity and behavior and therefore act as preliminary evidence for establishing causality. Future studies need to target loss-of-function and gain-of-function experiments to establish a causation link between neurophysiological mechanisms and motor memory consolidation. The proposed method achieved the identification of more specific, narrower frequency band activity modulated by cardiovascular exercise that was unique to each subject. To establish a

stronger causal link between frequency band-related ERD and the positive effects of exercise, future studies could modulate activity in these subject-specific bands by non-invasive electrical stimulation techniques and subsequently assess motor learning behavior. Our analysis pipeline also has clinical relevance as it could potentially be used to estimate the efficacy of rehabilitation strategies for individual subjects, e.g. stroke patients. Neurophysiological features modulated by the intervention could be extracted using the presented DL pipeline and subsequently used to obtain improvements in performance. Therefore, the proposed methodology yields significant potential for designing of patient-specific neurorehabilitation therapies that can significantly improve upon a one-size-fits-all approach.

### 3.6.4 Limitations

A major limitation of the CNN-based approach lies in computational demands. CNNs require more time and specialized hardware, namely GPUs, to train. In our case, CNNs were trained on a GeForce GTX 960 and required around 7 hours to train. The training demands and consequently required time would increase with more subjects. Therefore, the operational cost of using the presented methodology is expected to be higher as compared to alternative state-of-art methods.

## 3.7 Conclusion

The present work introduces a deep learning architecture for the analysis of EEG data and shows promising results in terms of discriminating the effects of an acute bout of high-intensity exercise/rest in close temporal proximity to performing a motor learning task on the brain activity of participants. The proposed approach outperformed alternative machine learning methods in terms of the classification of

CON and EXE subjects. Importantly, it also enabled us to visualize the features learned by deep networks such as CNNs, which may in turn yield better feature interpretation. The results are in general agreement with those reported in a previous study using standard statistical analysis using *a priori* selected features on the same dataset [35], with our analysis revealing a narrower, more specific frequency band associated with exercise-induced changes. In addition, our method revealed localized regions of the differential activity. Therefore, our approach demonstrates the feasibility of identifying subtle discriminative features in a completely data-driven manner using deep learning.

The proposed method is not restricted to the EEG modality and dataset described here. Hence, it paves the way for applying similar methods to other neuroimaging datasets of differing cohort sizes. This, in turn, yields promise for using deep learning as a tool towards the identification of neurophysiological changes associated with a variety of neurological disorders and ultimately lead to the design of optimized and individualized intervention strategies.

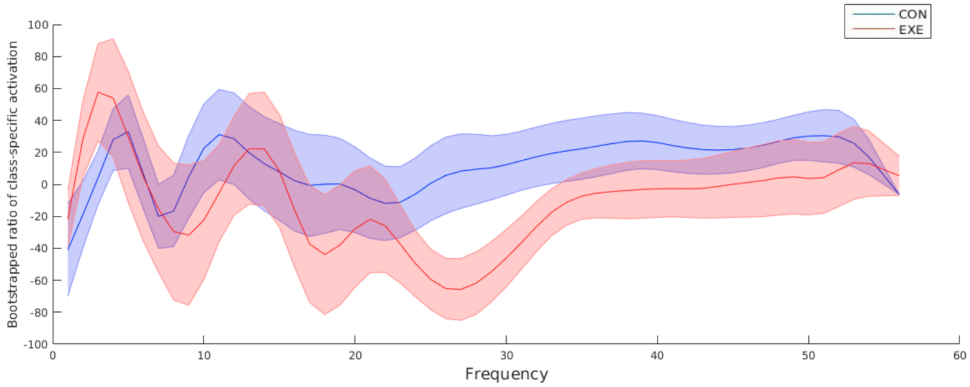


(a) Bootstrap ratio of TF map ccCAM averaged over electrodes & subjects showing discriminative frequencies. Bootstrapping is done using ccCAM values obtained for each timepoint to establish the reliability of the ccCAM values during the task period. The two groups exhibited different BSRs within the range 23–33 Hz.

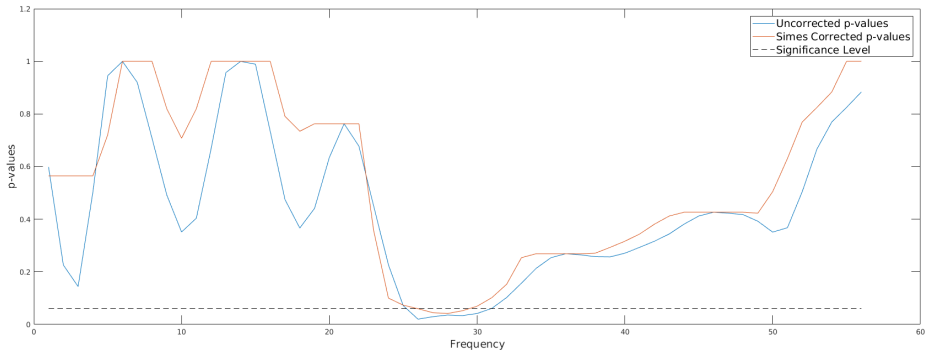


(b) Uncorrected and Simes corrected $p$-values corresponding to the difference between the bootstrapped ratio (BSR) values of the ccCAM TF maps obtained in CON and EXE groups. The difference in the BSR values between 27–29 Hz was found to be statistically significant ($p$ <0.05).

Figure 3.4: TF map ccCAM averaged over electrodes & subjects showing discriminative frequencies.

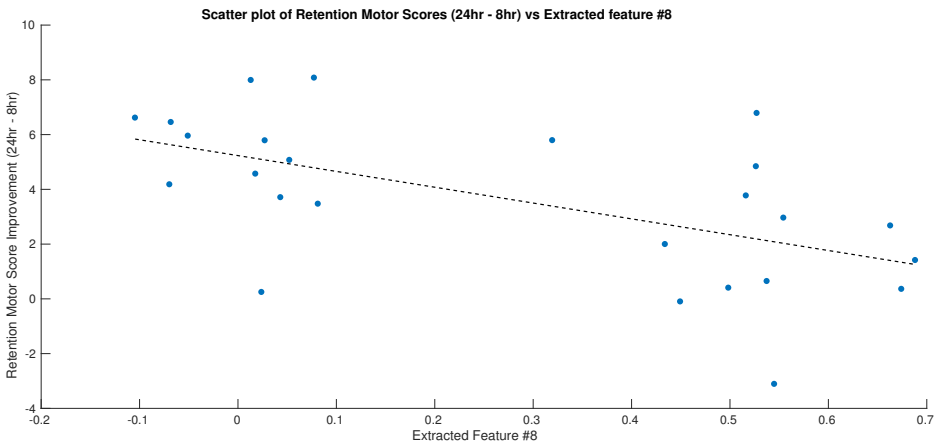


Figure 3.5: Scatter plot showing the motor skill retention improvement and a representative feature extracted by the final layer of CNN (Layer 3 of TopNN) from TF maps. The dotted line shows a linear fit between the two variables. The extracted feature is strongly correlated to the motor skill retention (correlation coefficient = -0.57).

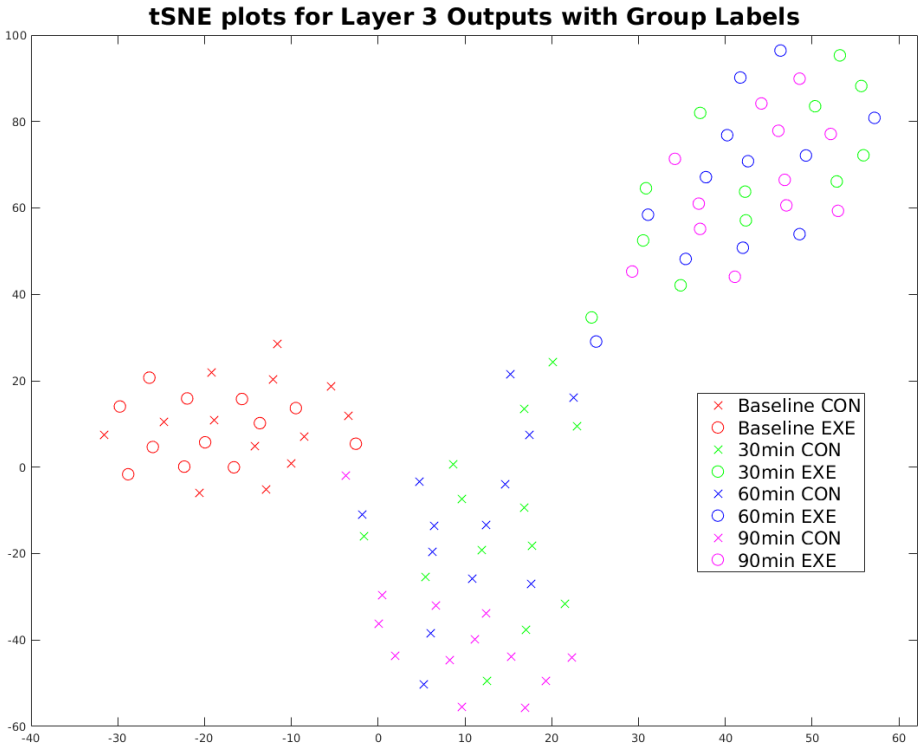


Figure 3.6: t-SNE plot of Top NN Layer 3 outputs (time-averaged) for the TF maps for all subjects and sessions. The EXE subjects (circles) move further away from baseline with time compared to the CON group (crosses), indicating exercise-induced changes on the underlying electrophysiological signals. Also, the feature trajectories across time are very different for the two groups starting from the 30 min session, which indicates that exercise-induced changes surface right after the session of acute exercise.

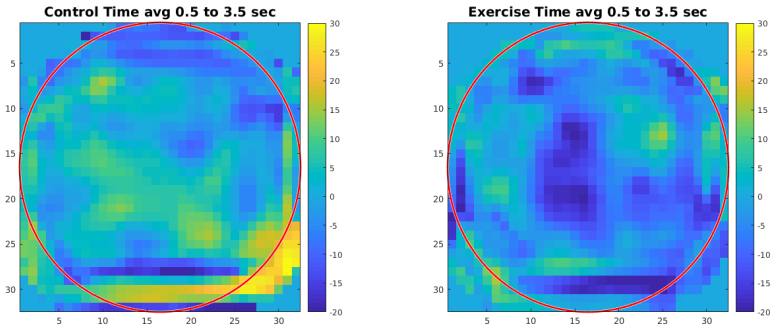


(a) Bootstrapped ratio of the topographical map of ccCAM values averaged over subjects in the CON and EXE groups showing regions with difference in activity before and 90min after rest/exercise. Bootstrapping was performed using ccCAM values obtained for each timepoint to establish the reliability of the ccCAM values during the task period.

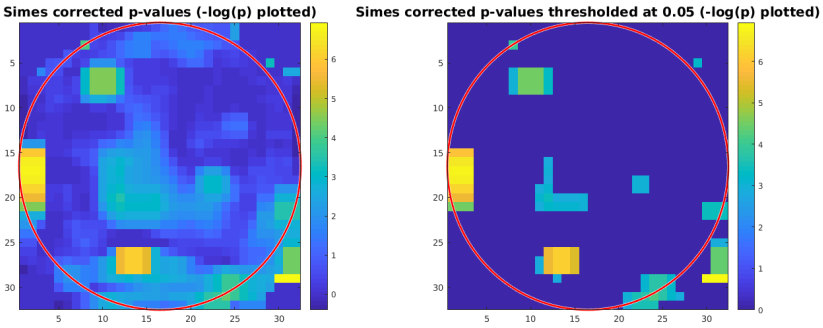


(b) Topographical maps of significantly different ccCAM bootstrapped ratio values between the CON and EXE groups. Instead of $p$-values, $-log(p)$ is color coded to delineate significant regions (yellow) more clearly. Non-significant regions are shown in blue. Significantly different activity was observed over the sensorimotor, occipital and frontal areas.

Figure 3.7: Topographical map ccCAM averaged over subjects showing brain areas with discriminative activity.

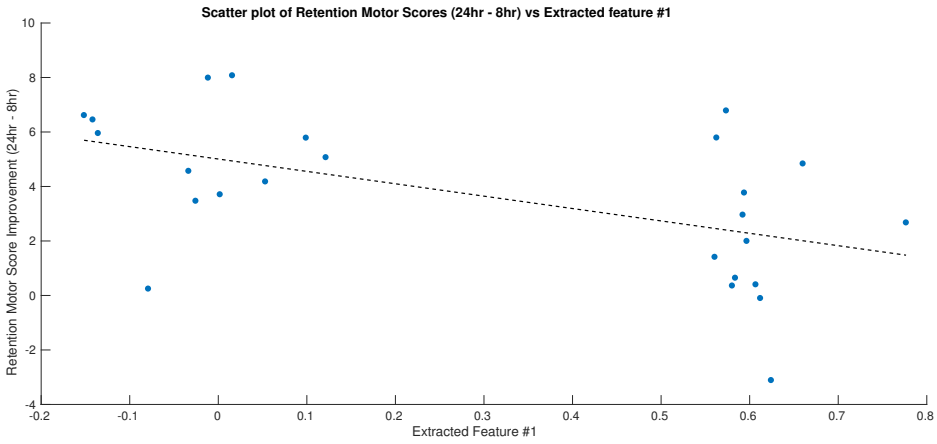


Figure 3.8: Scatter plot showing the motor skill retention improvement and a representative feature extracted by the final layer of CNN (Layer 3 of TopNN) from topographical maps. The dotted line shows a linear fit between the two variables. The extracted feature is strongly correlated to the motor skill retention (correlation coefficient = -0.53).

# CHAPTER 4
# BRAIN AGE PREDICTION

## 4.1 Abstract

The human brain changes with age and these age-related changes are known to act as biomarkers for several brain-related disorders. The prediction of brain age from T1-weighted MRI images therefore presents the scope for clinical applications. This study presents various analysis techniques involving dimensionality reduction techniques and regression models for prediction of brain age. Although Principal component analysis (PCA) degrades prediction performance, Canonical correlation analysis (CCA) enhances it when used along with Gaussian Process regression models. This combination proves to be the most accurate model for predicting brain age in the CAM-CAN dataset. The analysis also includes exploring the contribution of each brain region in determining the component that maximally covaries with brain age by using the CCA loading values. In agreement with previous studies, the loading values indicate brain atrophy all throughout the brain with strongest effect in grey matter of subcortical regions and corpus callosum fibre tracts. The analysis

---

**This work is to be submitted as:**
Xifra-Porxas A*, Ghosh A*, Mitsis GD, Boudrias MH. Combining structural MRI images and MEG recordings for Biological brain age prediction.

pipeline also reveals the contribution of the cerebellum and therefore inspires future studies to investigate these areas.

## 4.2 Introduction

The human brain is known to change across the adult lifespan. The process of *Brain aging* underlies the gradual decline in cognitive performance. Although aging-induced changes are not necessarily pathological, the risk of developing neurodegenerative disorders increases with increasing age [1]. The wide range of age-associated brain disorders indicates that the effect of aging on the brain structure and function vary greatly among individuals. In fact, diseases such as Alzheimer's and Schizophrenia are thought to be the results of processes associated with accelerated brain aging [87, 149]. Therefore, a better understanding of brain aging as well as better ways to identify biomarkers of healthy aging carrying vital importance will contribute to improve the detection of early-stage neurodegeneration and predict age-related cognitive decline.

One promising approach to identify individual differences in brain aging relies on the use of neuroimaging data to accurately predict the chronological or biological age of healthy individuals. In that context, machine learning techniques (ML) have proven to be a promising tool to 'learn' correspondence between patterns in the structural or functional brain features and the age label [43, 53]. ML models can formulate high-dimensional regression boundaries, with features derived from neuroimaging data as input variables to predict the biological age. When these models are trained on large training datasets with a large number of subjects, they can

generalize sufficiently well on unseen or 'novel' subjects. This provides the opportunity to deploy ML models at the population level and use the predicted age as a biomarker of the brain aging process.

Most studies have explored the use of ML on data obtained from neuroimaging techniques to quantify atypical brain development in diseased populations. A common practice entails training a ML-based prediction model on healthy subjects and subsequently using it to predict brain age in clinical samples, *e.g.* in the context of Alzheimer's population. In this context, if the brain-predicted age is greater than the reported chronological age for an individual, it is thought to be an indicator of aberrant accumulation of age-related changes to the brain. The degree of these atypical age-related brain changes has been more frequently quantified by subtracting chronological age from brain-predicted age. This approach has been applied in the case of several brain-related disorders including Alzheimer's [51, 58], traumatic brain injury [28], schizophrenia [89, 141], HIV [31], epilepsy [119], Down's syndrome [27] and diabetes [52]. Interestingly, the utility of predicting brain-age has also been extended beyond understanding neurological disorders. For instance, studies have reported a positive influence of meditation [105] and increased education and physical exercise [151] on the brain age. Recent work has also shown a relationship between the brain-predicted age difference and specific cognitive functions, namely visual attention, cognitive flexibility, and semantic verbal fluency [17].

The studies mentioned above have mainly focused on predicting the brain age based on structural Magnetic Resonance images (MRI), with most studies using T1-weighted MR images. This is due to the fact that large MR-based open datasets

are available, which has allowed researchers to train and validate their models on a large number of subjects. However, it is well known from the aging literature that in addition to structural alterations, changes in brain functional connectivity, described as the similarity between activity in two brain regions, are also closely related to cognitive aging [95]. In line with this, some studies have used ML to predict brain age from functional MRI (fMRI) data, specifically on fMRI-derived connectivity matrices [43, 100, 102, 177]. Few studies have also looked at changes in the functional signatures of the brain as captured by EEG, a modality that offers good temporal resolution, and subsequently build a brain age prediction model based on the temporal and spectral information of brain activity [4, 157]. A detailed overview of different neuroimaging modalities and ML methods that have been used to predict the brain age is presented in [26]. However, the aforementioned studies investigate the age-related structural and functional brain changes in isolation. In this study, we aimed to combine both structural and functional information to predict the brain age in a cohort of healthy subjects.

*Going Beyond Prediction*

Brain-age prediction methods have been primarily based on ML [20, 54, 121, 176, 185]. However, it is important from a neuroscience point-of-view to understand the underlying factors influencing these predictions. In addition to this, a major roadblock to clinical applications of ML models is their explainability and reliability [93]. As argued by Kriegeskorte *et al.*, decoding models can reveal the presence of information in fine-grained multivariate information [90]. The authors highlighted

the difficulties and confounds associated with interpreting weights in a linear decoding model and consequently suggest the use of multivariate techniques to identify the informative predictors. Based on these inferences, we sought to improve the explainability of some of the previously used ML models by applying correlative multivariate statistical methods. Specifically, we wanted to focus on dimensionality reduction techniques to identify neuroimaging features that significantly contributed to the performance of regression models like Gaussian Process Regressors [127].

*Including Topology to improve prediction*

Previous works have used neuroimaging features from different brain regions as independent features to predict the age using ML algorithms. However, these methods do not incorporate the underlying topology of the data, such as the neighborhood information of each voxel or information about the functionally connected regions. More recent work has focused on using Convolutional Neural Networks (CNN) for MR images [29, 69, 79] or functional connectivity patterns [100, 170] from fMRI data to capture the spatial topology of the data. However, the functional topology information has not yet been used for brain age prediction.

A functional connectivity matrix represents a graph structure arrangement of the regions of the brain. Therefore, directly using CNN with input as the matrices could lead to misleading conclusions. CNN is designed to capture the topology in Euclidean spaces. On the contrary, a graph is inherently a non-Euclidean structure. This has led to modifications in the CNN structure to capture the graph neighborhood of each region and using that information to predict the age of preterm infants

[82]. Although some recent work has focused on using graph-theory measures combined with ML techniques to predict brain age [42], none so far have explored the neighborhood information present in the graph structure. Therefore, in this study, we focused on the use of deep learning techniques to 'learn' the optimal features by utilizing the graph topology information. We explored the utility of Graph Convolutional Networks (GCN) [86] for age prediction. The GCN framework allows for graph nodes to have their respective features and afterward transforms these features by taking into account the graph topology. We developed a GCN framework to combine both structural and functional features of the cortical regions to predict the brain age. To investigate the functional brain changes, we relied on MEG to obtain a better temporal and spatial resolution of brain activity as compared to EEG. Specifically, we used the MEG-derived functional connectivity information to form the graph structure of the brain for each subject and used the electrophysiological and structural features of each brain region as features for the graph nodes. In this work, we aim to explore the role of the functional topology of the brain in age prediction using GCN for MEG-derived functional connectivity information.

## 4.3 Methods

### 4.3.1 Dataset

We used the open-access dataset from Cambridge Center for Aging Neuroscience (CAM-CAN) [145, 162], from which we used the T1-weighted MRI and the resting state data acquired with MEG. Briefly, the dataset consisted of 652 subjects (male/female = 322/330, mean age = 54.3 $\pm$ 18.6, age range 18-88 years). The MR images were acquired from a 3 T Siemens TIM Trio scanner with a 32-channel head

coil. The images were acquired using a MPRAGE sequence with TR = 2250 ms, TE = 2.99 ms, Flip angle = 9°, Field of View = 256 × 240 × 192 $mm^3$ and voxel size = 1 × 1 × 1 $mm^3$. The resting-state MEG data were recorded for 9 min using a 306-channel Elekta Neuromag Vectorview (102 magnetometers and 204 planar gradiometers) at a sampling rate of 1kHz. Shafto *et al.* and Taylor *et al.* present a detailed description of the dataset and the data acquisition protocols [145, 162].

### 4.3.2 Neuroimaging Data processing

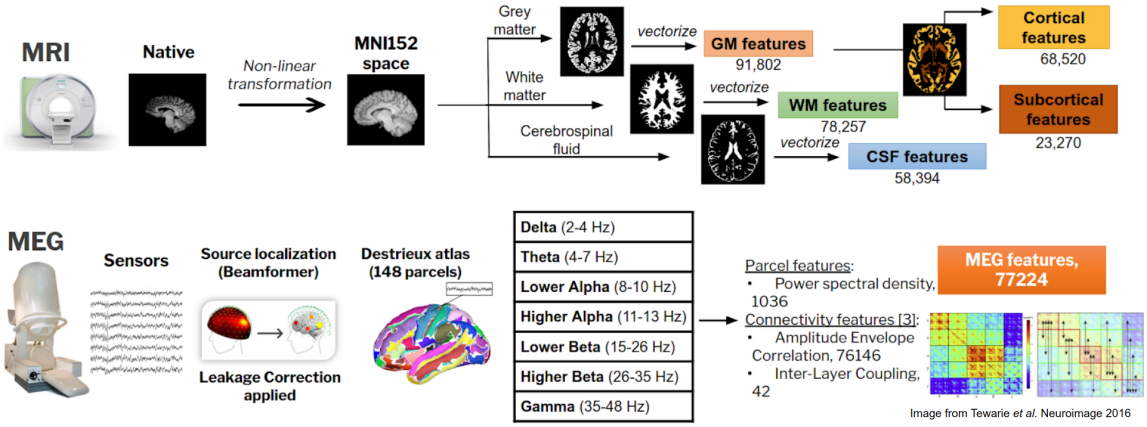


Figure 4.1: Feature Extraction pipeline for MRI and MEG data

A summary of the entire feature extraction process is presented in Fig. 4.1. The processing of T1-weighted MR images followed the pipeline presented in Cole *et al.* [29] and was implemented in FSL [77]. Briefly, all images were registered to the MNI152 space (2mm resolution) by using the brain extraction pipeline followed by a non-linear spatial transformation. Next, we segmented the images into Grey Matter (GM), White Matter (WM), and Cerebrospinal Fluid (CSF). The GM maps were further segmented into cortical and subcortical regions to delineate the effects of aging on these regions. The resultant images were vectorized and subsequently

z-scored to obtain a feature vector for each subject. This process resulted in a feature matrix where each row consists of normalized feature maps for a single subject.

For the MEG data, we first convert the data from sensor space to the source space using the Linearly Constrained Minimum Variance Beamformer technique [173, 174]. The sources were constrained to the cortical regions of the brain. Further, we parcellated the cortex into 148 parcels using the Destrieux atlas [40]. For each of the cortical parcels, the power spectrum density (PSD) for the entire resting state was calculated for 7 frequency bands, namely Delta (2–4 Hz), Theta (5–7 Hz), lower Alpha (8–10 Hz), higher Alpha (11-13 Hz), lower Beta (15–25 Hz), higher Beta (26–35 Hz), and Gamma (36–48 Hz). In addition to the PSD features, amplitude envelope correlation in each frequency band was used to estimate the functional connectivity relationship among cortical parcels. Thus, we had a functional connectivity matrix for each of the 5 frequency bands – Delta, Theta, Alpha (8–13 Hz), Beta (15–35 Hz) and Gamma. Inter-layer Coupling (ILC) was calculated from the functional connectivity matrices to estimate the similarity of the connectivity profile across frequency bands [163]. This feature extraction process is summarized in Fig. 4.1. Some subjects were removed during this process due to unsatisfactory preprocessing results. Therefore, we had a pruned dataset with 612 subjects.

### 4.3.3 Gaussian Process Regression

Gaussian Process regression (GPR) has been widely used in predicting chronological age from T1-MRI data [27, 28, 29, 30, 31]. The GP approach, which is a non-parametric approach, finds a distribution over possible functions that are consistent with the data. The main assumption underlying GPs is that any finite subset

of the available data must follow a multivariate Gaussian distribution. The prior belief about the relationship between variables is decided by the definition of these multivariate Gaussian distributions to generate a model that represents the observed variance in the available data. Multivariate Gaussian distributions can reflect local patterns of covariance between individual data points. Therefore, a combination of multiple such distributions in a GP can model non-linear relationships and is more flexible than conventional parametric models, which rely on fitting global models. In this work, GPs were used to predict age from neuroimaging data, which was a regression problem.

### 4.3.4 Similarity metric

The GPR method was implemented using the scikit-learn toolbox [120] in Python. The feature vectors obtained were input to the GPR model for training. We used a k-fold cross-validation strategy, and each random split of the dataset consisted of 500 subjects in the training set and the rest in the validation or testing set. Following [29], we represented the data as a $N \times N$ similarity matrix (N being the number of subjects in training set). The similarity between any two subjects was calculated using the dot product of their corresponding feature vectors. We also tried a different similarity metric, namely the cosine similarity, but it yielded comparable performance. Therefore, each subject was represented as a 500-length vector containing similarity values corresponding to each of the 500 training subjects. The model performance was evaluated using mean absolute error (MAE) and coefficient of determination ($R^2$) of the prediction, using 20-fold cross-validation.

We also compared the performance of GPRs with Support Vector Regression (SVR) models [10], which are more commonly used in the machine learning literature for regression problems. We observed that GPR outperformed SVR. Therefore, GPR was chosen to be the regression model for all further analyses.

However, the use of a similarity metric entails the following issues:

1. The training set needs to have enough subjects to sample the spectrum of healthy aging completely
2. The predictions are based on how similar a test subject is to each of the training subjects. Thus it is difficult to visualize the neuroimaging features that guide the brain age prediction

To avert the issues mentioned above, we used dimensionality reduction techniques. Specifically, we used Principal Component Analysis (PCA) and Canonical Correlation Analysis (CCA) to identify the features that contribute to age prediction.

### 4.3.5 Principal Component Analysis (PCA)

PCA is a linear dimensionality reduction technique using Singular Value Decomposition of the data to project it to a lower dimensional space [80]. It is widely used to decompose a multivariate dataset in a set of successive orthogonal components that explain the maximum amount of the variance. The principal components obtained from this analysis correspond to the maximal modes of variation and hence correspond to the most prominently changing features in the dataset. Often, the number of principal components to be used is decided using the knee rule. We used PCA to project the feature matrix to a lower dimensional space and subsequently predict the age using a GPR model.

### 4.3.6 Canonical Correlation Analysis (CCA)

Canonical Correlation Analysis (CCA) is another dimensionality reduction technique that finds latent variables to model the covariance in input and output variables [165]. Like PCA, CCA is also a linear dimensionality reduction technique but uses Singular Value Decomposition of the covariance matrix instead of the input variance matrix.

In the present problem scenario, the inputs are the neuroimaging feature matrix and the brain age vector. Therefore, CCA should retrieve a linear combination of the neuroimaging features that are maximally correlated to the brain age. We used CCA to project the feature vector along this direction and subsequently use the projection values to predict brain age using GPR. Like previous analysis cases, the performance is judged using the MAE and $R^2$ values following 20-fold cross-validation.

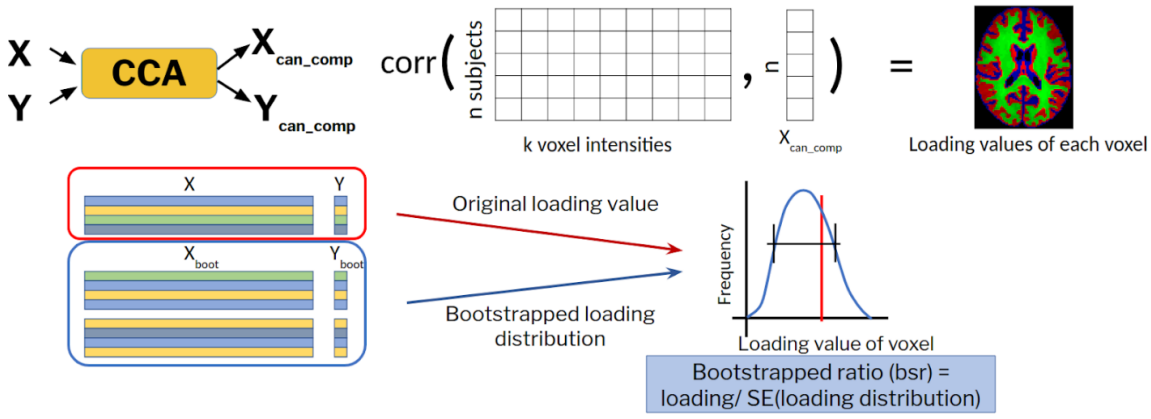


Figure 4.2: Figure illustrating the calculation of loadings and Bootstrapped ratio (BSR) of loadings from a CCA model

CCA also yields a loading vector that specifies the contribution of each feature to the CCA component. We used these loading values to estimate the contribution

of each feature to the brain age prediction and thereby understand which regions of the brain exhibit maximal age-related changes. To estimate the reliability of these loading values, we used the bootstrapped ratio by repeating the CCA analysis for 1000 bootstrapped samples of the dataset chosen at random with replacement [46, 111, 112]. The bootstrapped ratio (BSR) of the loading values indicates which areas reliably contribute to the brain age prediction, thus increasing the overall reliability of the prediction systems. The schematic for generating the BSR of loadings values is illustrated in Fig. 4.2.

### 4.3.7 Rich Club Organization

Another aim of the study was to include topology to improve prediction. To do so, we focused on exploiting the functional topology of the brain. The functional connectivity matrix in each frequency band of the MEG data yielded a graph topology of the cortical regions. Therefore, we created a graph for each subject and each frequency band where the nodes would correspond to the cortical parcels, and the edge weights would be the connectivity values in the respective frequency band. To obtain a sparsely connected graph, we binarized the connectivity matrix to have the top 5 percent weights for each subject and frequency band. Hence, each graph had a connectivity density of 5 percent.

To observe if the graph topology changes with age, we calculated the rich club coefficients for each functional graph [171]. Rich-clubs are known to be variations of small world network organization, with disproportionately dense interconnections and different topological properties [113, 171]. The rich-clubs in the brain effectively define the top-level structure of the brain and enable effective communication

throughout the brain. Therefore, we wanted to observe if the graph topology remained stable with aging or if it showed significant changes in its "rich-clubness". Consequently, we calculated the rich-club coefficients of the binary graphs following [171] and subsequently obtained normalized rich-club coefficients using the HQS-normalization [187].

We compared the rich-club coefficients as a function of node degree across different age groups. The entire dataset of 612 subjects was ordered in increasing order of age and then divided into six age groups, consisting of 102 subjects each. Apart from the rich-club coefficients, we also compared the distribution of node degrees and shortest path length across age groups for binary graphs of each frequency band.

### 4.3.8 Graph Convolutional Networks (GCN)

A key challenge from a machine learning perspective is to encode the high-dimensional, non-Euclidean information about the graph structure into a feature vector. Traditional machine learning approaches have relied on summary graph statistics (e.g. degrees or clustering coefficients) [15], kernel functions [180] or hand-crafted features to measure local neighbourhood structures [101]. However, these approaches are limited because the features used are inflexible. Specifically, they cannot adapt during the learning process, and designing these features can be a time-consuming process requiring domain-specific knowledge. In contrast, *representation learning* approaches aim to learn node embeddings that encode graph structure.

GCN, a recent development in the deep learning literature, is a type of *representation learning* approach that uses shared parameters between nodes and leverages

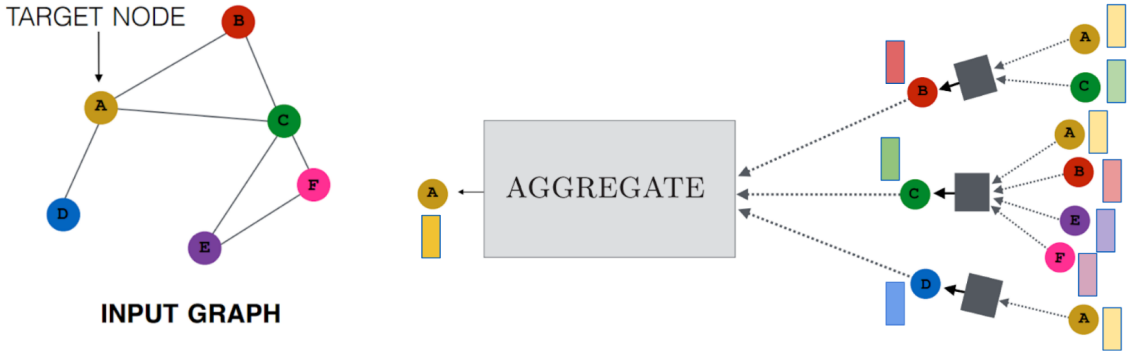


Figure 4.3: An overview of the embedding generation process in a 2-layer GCN. To generate an embedding for node $A$, the model aggregates messages from $A$s local graph neighbors (i.e., $B$, $C$, and $D$), and in turn, the messages coming from these neighbors are based on information aggregated from their respective neighborhoods, and so on. The boxes next to the nodes represent their embedding vectors. Darker color indicates a higher level of embedding generated by the GCN. Figure modified from [65].

node attributes during the embedding generation process. An overview of the embedding generation process is shown in Fig. 4.3. Depending on the type of aggregator function used to incorporate the neighborhood information, the GCN model can learn different information from the local graph neighborhood that aids end-to-end learning. Two types of aggregator functions that are commonly used are the mean aggregator and the pooling aggregator functions. GCN uses the mean aggregator [86] whereas a variant of GCN, namely the GraphSAGE, uses the pooling aggregator [64]. The two aggregator functions and their implications on the learning paradigm are discussed in detail in [64].

For this work, we combined the functional and structural features and used GCNs to predict age. As mentioned before, each subject had a distinct graph for each frequency band wherein each node of the graph represented a cortical parcel.

We used the PSD features, the cortical area, thickness & volume measures, and the mean & standard deviation of MRI intensities of voxels of the respective cortical parcel as the node attributes or features. Therefore, the dimensionality of the node features was 12 (7 for PSD features, 3 for the cortical area, thickness & volume, and 2 for the mean & standard deviation of MRI voxel intensities).

## 4.4 Results

### 4.4.1 Dimensionality reduction and GPR

We compared the performance of all models using a 10-fold cross-validation approach with random train and test splits of the dataset in each fold. Firstly, to compare the performance of different dimensionality reduction techniques with the similarity metric presented in [29], we used only the MRI voxel intensities as input to different models. A summary of the results is shown as a bar plot in Fig. 4.4.

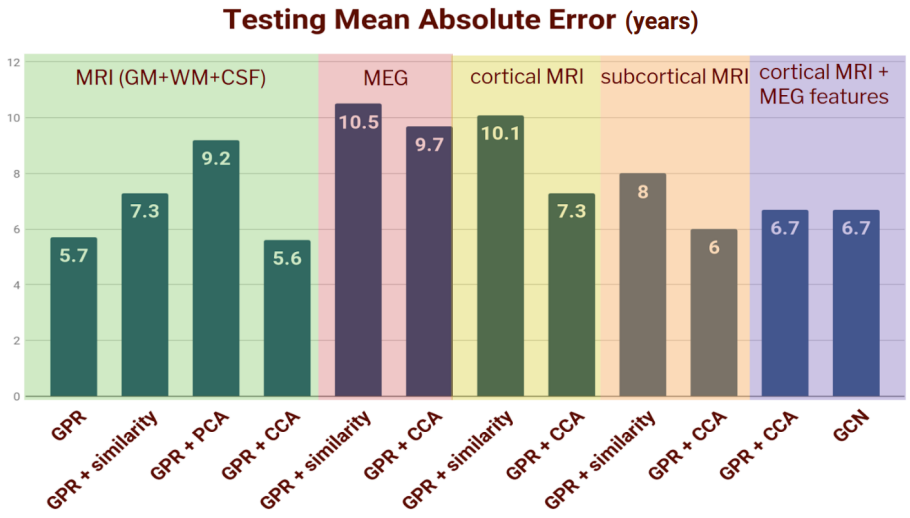


Figure 4.4: Brain age prediction Mean absolute error (MAE) in years for different methods and data modalities.

Using only MRI voxel intensities as the input, CCA turned out to be the best model for age prediction with MAE of 5.57 yrs. PCA degraded performance significantly, and the use of the similarity metric yielded worse performance (with MAE of 7.3 yrs) as compared to using GPR on raw features. The failure of the similarity metric to yield the best performance indicated the variability of features across subjects in the dataset. The presence of lesser subjects, as compared to [29], led to incomplete sampling of the aging subspace and hence yielded worse performance. Since CCA coupled with GPR was the best model, we compared this method for other data modalities.

### 4.4.2 Adding the Functional features from MEG data

The use of MEG features alone did not yield satisfactory performance with MAE of 9.7 yrs for the GPR+CCA model (Table B.2). However, the MEG features only contained information from the cortex, whereas the MRI intensities were from both cortical and subcortical regions. To delineate the effect of cortical from subcortical areas, we compared the performance of these features separately. Using the similarity metric on cortical MRI features yielded similar performance as MEG features, the MAE being 10.1 yrs and 10.5 yrs respectively (Table B.1). Also, there was a striking difference between the performance of models using subcortical MRI features and those using cortical features. Subcortical features performed better with MAE of 6 yrs for the GPR+CCA model, thus indicating a substantial contribution of subcortical regions towards age prediction.

To assess the utility of combining modalities and exploring the problem of age prediction from a multi-modal approach, we used CCA with GPR on cortical MRI

| Aggregator Method | Connectivity Information | MAE (yrs) |
|---|---|---|
| - | None | 7.21 ± 0.14 |
| Mean (GCN) | Delta | 7.14 ± 0.16 |
| Mean (GCN) | Theta | 7.19 ± 0.19 |
| Mean (GCN) | Alpha | 7.28 ± 0.21 |
| Mean (GCN) | Beta | 7.22 ± 0.19 |
| Mean (GCN) | Gamma | 7.18 ± 0.17 |
| Pooling (GraphSAGE) | Delta | **6.69 ± 0.13** |
| Pooling (GraphSAGE) | Theta | 6.82 ± 0.15 |
| Pooling (GraphSAGE) | Alpha | 6.94 ± 0.16 |
| Pooling (GraphSAGE) | Beta | 6.84 ± 0.15 |
| Pooling (GraphSAGE) | Gamma | 6.80 ± 0.14 |

Table 4.1: Comparison of age prediction by GCN models with different neighbourhood pooling functions. Mean Absolute Error (MAE) calculated over the testing set and averaged over 10-folds.

and MEG features. This model performed better than individual modalities (MAE of 6.69 yrs) but was worse than using subcortical MRI features (Table B.2). GCNs yielded similar performance as that of the GPR+CCA model (MAE of 6.72 yrs). However, a critical difference between the two is that the GCN model uses a linear regression model on the learned features to predict age. Therefore, we compared GCN models to similar models that do not have the connectivity information to identify if including the topology contributed to improvement in age prediction. These results are presented in Table 4.1.

### 4.4.3 Impact of Graph Topology

The rich-club coefficients were not significantly different across the different age groups for any frequency band. All groups showed a similar trend of normalized rich-club coefficients with increasing node degree, with all groups showing rich club behavior above node degree of 17. The results were similar for the degree as well as

shortest path distribution in the binary graphs of all frequency bands across ages. These results indicated that the graph topology in the binary graphs did not change significantly with age. Therefore, we aimed to leverage this topology to learn feature representations for each node and subsequently, predict age using GCNs. The graph metrics guided the structure of GCNs (and its variant, *i.e.* GraphSAGE). Since all age groups showed a peak shortest path length of 3, the GCN architectures consist of 3 layers. Each layer of GCN incorporates information from 1-hop, and therefore, a 3-layer GCN can incorporate information from 3-hop neighbors. Further details about the GCN architecture and comparison of different graph metrics across age groups are presented in Appendix B.

### 4.4.4 CCA loadings

We also identified which brain regions correspond to maximal age-related changes. The CCA loadings offered a window to assess the reliability of each neuroimaging feature to age prediction, thus indicating which features showed the most reliable age-related changes. The histogram of BSR values of loadings of voxel intensities and the top 15 percent BSR values of loadings of each of GM, WM, and CSF are shown in Fig. 4.5. The histogram of BSR values indicated that GM and WM voxels showed more reliable age-related changes as compared to CSF (histogram peak for GM and WM being around -300 and -400 respectively, whereas histogram peak for CSF is around -100). Most of the loadings were negative, thus indicating a decreased voxel intensity with increasing age. Also, the top 15 percent of BSR values are confined to subcortical regions, thus supporting our earlier results that subcortical regions yielded better age prediction. These areas are shown in Fig. 4.5. The

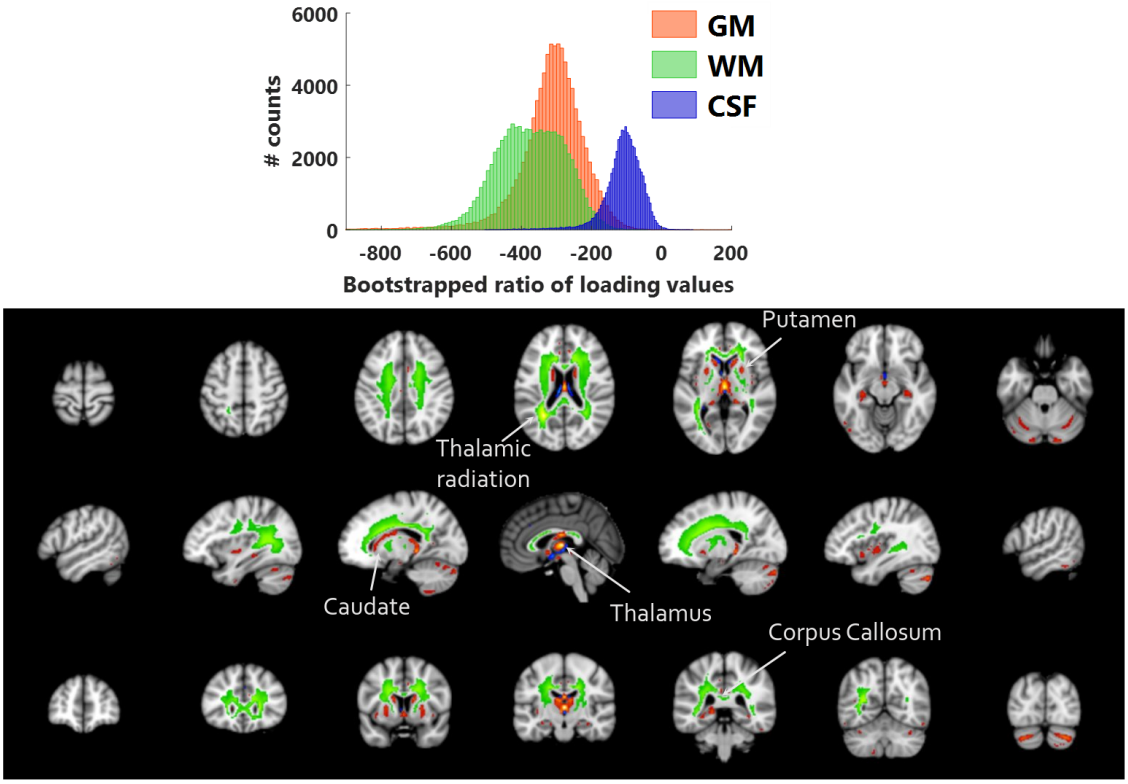


Figure 4.5: CCA loadings BSR values for the T1-weighted MR images. Top image shows the distribution of BSR values for GM, WM and CSF voxels. Bottom image shows the regions with top 15% BSR values, highlighting that most reliable voxels for brain age prediction were restricted to subcortical regions.

highlighted (red) GM areas were localized in important subcortical structures like Putamen, Thalamus, and the Caudate nucleus. Most of the highlighted (green) WM voxels were confined to the corpus callosum, thus indicating that the corpus callosum reflected most age-related changes among WM voxels. Another essential structure among WM voxels is the Thalamic radiation. Our analysis pipeline also uncovered certain regions in the cerebellum that showed age-related effects reliably across the dataset.

We visualized the CCA loadings for the PSD features, which are shown in Fig. 4.6. We observed different regions showing age-related effects in various frequency bands. The low-frequency bands, namely Delta and Theta, showed maximal age-related effects in the frontal areas. The alpha band showed maximal age-related effects in the visual and motor areas. Higher frequency bands, namely Beta and Gamma, showed increasing PSD values with age in motor areas. Contrary to MRI loadings, wherein most of the loading values were negative, PSD loadings were both positive and negative. These results are in agreement with previous literature studying the difference in resting-state PSD values in various frequency bands between healthy young and elderly populations [21, 34, 67, 76, 99, 132, 181].

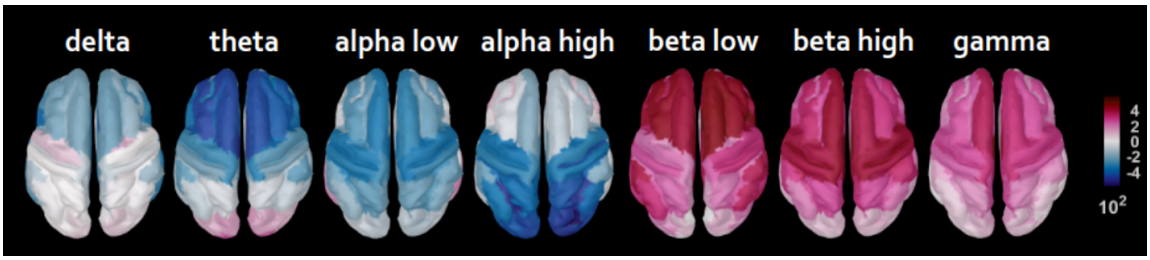


Figure 4.6: CCA loadings BSR values across different cortical regions depicting regions with PSD values that are correlated or anti-correlated with age for each frequency band.

## 4.5 Discussion

In this study, we aimed to leverage multimodal neuroimaging data to predict the age in a cohort of healthy subjects with age ranging from 18–88 years. The major contributions of our study are summarized below:

1. We applied dimensionality reduction techniques in conjunction with ML and found that the CCA+GPR model worked best for all data modalities. Furthermore, using PCA was detrimental to the performance of GPR.
2. We visualized the regions that exhibited the age-related changes and found that subcortical MRI intensities showed age-related changes more reliably than cortical MRI intensities. We also showed the age-related changes in the spectral features of various cortical regions, as observed using the MEG data.
3. We developed a GCN framework to incorporate the graph topology in the functional data and thereafter predict age. This framework allowed combining structural and functional features for cortical regions and thereby, improved age prediction.

### 4.5.1 Dimensionality reduction and GPR

We used T1-weighted MR images and resting-state MEG data to develop a brain-age prediction framework that uses both structural and functional information of the brain. We restricted our analysis to cortical sources of the MEG data and thereby had functional information from the cortical regions only. Since the goal was to predict the age, the desired MAE for the perfect model would be 0 yr. Owing to subject variability and the ill-conditioning of the problem, precisely the definition of "healthy" subjects, we did not expect to achieve a MAE of 0 yr. To get an estimate of the chance level of age prediction, we used predictions from a random model with no training. Irrespective of the modality of data used, the chance level of MAE was around **16.738 yrs** and $R^2$ was around **-0.034**. These values act as a baseline to judge the performance of various models. All models, irrespective of the modality

of data, performed better than chance level thus indicating that all neuroimaging features that were used had some age-related effects.

Following [29], GPRs outperformed other regression models, like Support Vector Regression (SVR) and Random Forest regression models. These experiments were done on the similarity metric-derived feature vectors. Based on these results, GPRs were used as the regression model of choice for age prediction for all further experiments. A vital contribution of this work is the role of dimensionality reduction techniques in age prediction. PCA is one of the most commonly used dimensionality reduction techniques in neuroimaging. However, PCA degraded performance in the current scenario. We feel that the primary reason underlying this is the subject variability in neuroimaging features. Since PCA yields components that are maximally varying in the dataset, these could be aligned to directions of subject variability in the dataset instead of being age-related changes. Hence, it degrades the age prediction performance. Instead, CCA improves performance by yielding the component that maximally covaries with age. Since CCA is a multivariate associativity technique relying on Singular Value Decomposition, it was able to identify features that covary with age. Following the success of CCA, we also tried deep CCA (DCCA) [5] to learn a non-linear combination of features that maximally covary with age. But deep CCA was not numerically stable and hence, was not explored further.

A major drawback of using GPR as a regression model is that the prediction is reasonably good only for points within the bounds of the training data. Hence, the age prediction for subjects that are younger than the youngest training subject (or older than the oldest training subject) is erroneous. Therefore, the model returns

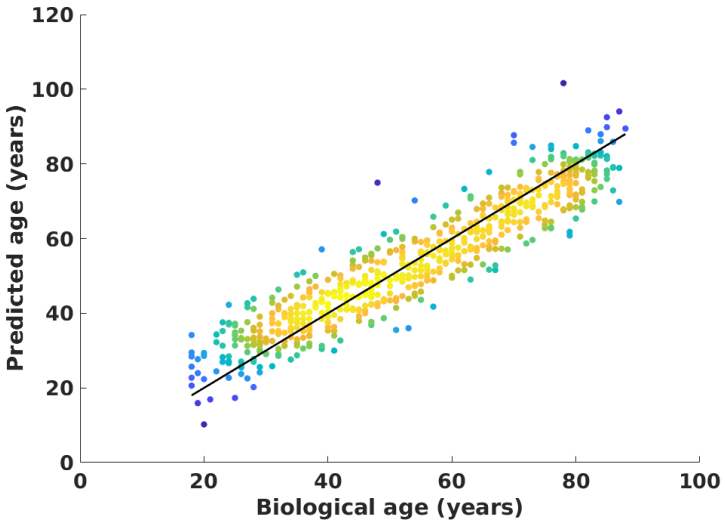


Figure 4.7: Plot of predicted age vs biological age for CCA+GPR model using concatenated MRI and MEG features.

biased age predictions for the youngest and oldest subjects. Fig. 4.7 shows the plot of predicted age vs. biological age for the CCA+GPR model using all neuroimaging features. The predictions for the youngest subjects is higher than the biological age, whereas the trend is reversed for the oldest subjects.

### 4.5.2 Adding Functional features from MEG data

We studied the contribution of spectral features and functional connectivity features obtained from resting-state MEG recordings towards age prediction. Among the spectral features, we used the relative PSD in 7 frequency bands. We used Amplitude Envelope Correlation (AEC) and ILC to characterize the functional connectivity information [163]. However, ILC values did not significantly contribute to age prediction with the MAE values being very close to that of the random model (MAE of 14.7 yrs). PSD performed much better than AEC values and combining

both these features yielded the best performance with a CCA+GPR model (MAE of 9.7 yrs). This model was worse compared to the CCA+GPR model using MRI intensities. However, the similarity+GPR model using only cortical MRI intensities yielded comparable performance to the one using PSD and AEC values. These results indicated that a possible reason for the inferior performance of models using MEG features was the absence of subcortical information.

### 4.5.3 CCA loadings

CCA also returns loading values for each input feature and therefore, improves model explainability. Using the BSR of loading values for MRI images, we found that most of the voxel intensities were negatively correlated with age. These findings are in agreement with previous studies that have studied cortical thinning with age [48, 71, 134, 155]. However, the absence of the strongest negative correlations in the cortex was possibly due to improper alignment of sulci and gyri to the standard brain in MNI152 space. We further explored the contribution of cortical area, thickness, and volume measures to age prediction. However, models that used these metrics failed to perform as well as models that used subcortical MRI intensities for age prediction. These results indicated that subcortical regions are more reliable predictors of age, possibly due to less subject variability within the same age group [133]. The subcortical structures that most reliably showed age-related changes are some of the most important structures for information relay across the brain and several cognitive functions [60, 142, 146]. Several stereological and MRI studies have explored the atrophy in these regions with age [19, 25, 49, 91, 110, 156]. Our results are in strong agreement with the inferences from these studies and also provide insight

into the age-related changes in these regions compared to the cortical areas. The WM regions highlighted in Fig. 4.5 are mostly confined to the corpus callosum and the thalamic radiation. These results are in strong agreement with previous studies exploring bimanual coordination in humans that have reported atrophy in corpus callosum fiber tracts [9, 143, 153]. Although CSF voxels did not exhibit high BSR values of loadings when compared to their GM and WM counterparts, including CSF improved model performance. CSF information possibly indicated changes in brain volume and ventricle sizes that helped brain age prediction.

Among the MEG features used, PSD features yielded the best performance. Adding the functional connectivity information helped age prediction slightly. Hence, we visualized the loadings values for only the PSD features. Several studies have investigated the relationship between resting-state power levels of various frequency bands and the healthy aging process. The reported results have a fair bit of controversy in several bands, primarily due to different resting-state protocols, *i.e.* eyes-closed or eyes-open. Regarding the slower waves (0.5–7 Hz), most studies have reported a decrease in power among older adults as compared to their younger counterparts [21, 34, 99, 181]. The CCA loadings for the Delta and Theta bands reveal a similar relationship with age, the more focal effect being restricted to the frontal regions. These results possibly indicate declining cognitive performance with age [78]. There is much more consensus about changes in beta power with age. Several studies have shown an increased beta power at rest [21, 67, 76, 132]. Using higher and lower beta bands, we could identify the spatial extent of this effect in each of

the bands. The lower beta band is restricted to frontal regions, indicating changes in cognitive processing during motor task performance [76].

On the contrary, the higher beta effects are restricted to the motor cortex indicating differences in motor task performance with age [67, 132]. Changes in alpha power are more controversial as compared to other bands. Several studies reported no significant changes in alpha power with age [21, 67], whereas other studies have reported a decrease in alpha [76]. Since alpha activity indicates changes in attention [21], the differences in findings could be related to the resting-state protocol [160]. The CCA loadings indicated a strong effect of age on the higher alpha activity in the occipital cortex. However, further investigation is required to conclude if this effect is due to subjects falling asleep during the eyes-closed resting state session [160].

### 4.5.4 Impact of Graph Topology

Among the GCN models, GraphSAGE performs better than the basic GCN variant. GraphSAGE can "learn" the neighborhood statistics that aid the task of age prediction, whereas GCN gives equal importance to all neighbors [65]. This behavior has close relations to various routing strategies in biological graphs, namely random walks and percolation [8]. GCN can be thought of using the random walk communication strategy as it uses all connections to learn feature representations of nodes. GraphSAGE can be thought of using a more percolation-like routing strategy, whereby it uses certain connections to aggregate information about the neighborhood. The ability to adaptively aggregate information from each node's neighborhood gives GraphSAGE models the power of inductive learning [64], thus improving the models to generalize better on unseen subjects. Further investigation

into the feature space learned by the models could provide more insight into the topological information learned by these graph-based models. In this work, we only focused on binary graphs. But the methodology can be easily extended to weighted graphs.

### 4.5.5 Limitations

A severe limitation of the problem of brain age prediction is the use of biological age as a surrogate for brain age. Although we used a cohort of healthy subjects, the brain age is known to depend on various other factors, like education [151]. In this work, we ignored all lifestyle factors and aimed to predict the biological age from neuroimaging features. Besides, we used a single model to predict the brain age for both males and females. These factors contribute to the biological age labels being noisy version of the "true" brain age of each subject. Also, we restricted our analysis to the cortical information from the functional data, which limits the investigation of age-related changes into the functional activity of deep brain structures. Like other deep learning architectures, GCN-based approaches require more time and specialized hardware, namely GPUs, to train. We used a GeForce GTX 960, and each cross-validation fold required about 30 minutes of training time. Therefore, the operational cost of deploying GCN-based methods is higher than other ML-based approaches.

## 4.6 Conclusion

In this study, we used a combination of structural and functional brain information to predict brain age in a cohort of healthy subjects. We showed that dimensionality reduction techniques could be used to simultaneously improve brain

age prediction and identify key neuroimaging features that show age-related effects. Specifically, we showed that using CCA in conjunction with GPR yielded the best model for age prediction. PCA degraded prediction performance. We also showed that subcortical structures exhibited the most reliable age-related effects in MRI features. We added functional features from resting-state MEG to further improve the age prediction obtained using only cortical information.

Additionally, we presented a framework to leverage the functional graph topology and perform brain age prediction using GCN. We found that including the topological neighborhood information improved age prediction. However, these models did not outperform models that used GPR. Therefore, further work could focus on using GPR in conjunction with GCN. In the current work, we explored the problem of age prediction only in the context of healthy subjects. The validation of our proposed analysis pipeline to predict age for a diseased population remains to be explored. This validation is essential to establish the potential for the identification of biomarkers relevant to atypical brain aging processes, and thereby as an early marker for neurological disorders.

# CHAPTER 5
# CONCLUSIONS & SIGNIFICANCE

The first project introduces a deep learning architecture for the analysis of EEG data. The proposed CNN outperformed alternative machine learning techniques with regards to differentiating between previously unseen Exercise and Control group subjects, based on their event-related desynchronization information. Importantly, we could also visualize the features learned by the CNN, which in turn, yields better interpretation of their classification basis. The results are in general agreement with those reported in a previous study using standard statistical analysis on *a priori* selected features from the same dataset [35], with our analysis revealing a narrower, more specific frequency band associated with exercise-induced changes. Also, our method revealed localized regions of the differential activity. Therefore, our approach demonstrates the feasibility of identifying subtle discriminative features in an entirely data-driven manner, using deep learning. These results are expected to contribute to the methodological advances for small-scale neuroimaging studies, such as those performed in the context of neurorehabilitation, where a small cohort of subjects/patients is traditionally tested. The findings are also expected to promote research towards the identification of neurophysiological changes associated with a variety of neurological disorders and ultimately lead to the design of optimized and individualized intervention strategies.

The second project presents a brain age prediction framework using structural and functional neuroimaging data. The results indicate the effectiveness of dimensionality reduction techniques coupled with regression models to improve brain age prediction. Interestingly, CCA improved the predictor performance, whereas PCA degraded the performance when used in conjunction with GPR. Therefore, our analysis provides conclusive evidence that the age-related brain changes are not maximally varying changes in structural features of the brain. Subsequently, we identified brain regions and features that reliably show age-related changes. Specifically, the structural features from subcortical areas showed more prominent age-related changes. Also, we showed that the proposed GCN-based framework improved prediction performance by leveraging the functional topology in the MEG data. Further work will attempt to establish the validation of these methods in the context of diseased populations. This work is expected to have significant contributions in the field of aging neuroscience to improve the understanding of the aging process and clinical neuroscience to establish biomarkers for atypical brain aging.

# APPENDIX

# Appendix A

## Effects of Cardiovascular Exercise on Neural Correlates of Motor Learning

### Network Architecture

**Notation**:- *Conv* denotes the 2D Spatial Convolutional layer. *ReLU* denotes the Rectified Linear Unit Layer that adds non-linearity to the network. *MaxPool* denotes the 2D Spatial Max Pooling layer. *FullyConn* denotes a Fully Connected layer, also known as the linear layer of the network.

### TF maps

| Layer | Type | Maps and Neurons | Filter Size |
|---|---|---|---|
| 0 | Input | 1M × 64 × 55N | - |
| 1 | Conv | 6M × 64 × 28N | 1×5 |
| 2 | ReLU | 6M × 64 × 28N | - |
| 3 | MaxPool | 6M × 64 × 14N | 1×2 |
| 4 | Conv | 16M × 64 × 14N | 1×5 |
| 5 | ReLU | 16M × 64 × 14N | - |
| 6 | MaxPool | 16M × 64 × 7N | 1×2 |

Table A.1: Network architecture used for EEG feature extraction network (Base CNN). The output of the network is a tensor of dimensions $16 \times 64 \times 7$.

| Layer | Type | Maps and Neurons | Filter Size |
|---|---|---|---|
| 0 | Input | 16M × 64 × 7N | - |
| 1 | Flatten | 7168N | - |
| 2 | Dropout (p=0.5) | - | - |
| 3 | FullyConn | 8N | 1×1 |
| 4 | ReLU | 8N | - |
| 5 | FullyConn | 2N | 1×1 |

Table A.2: Network architecture used for group discrimination network (Top NN). The output of the network is a vector of dimension 2, values corresponding to the probability that the data tuple belongs to particular class.

| Layer | Type | Maps and Neurons | Filter Size |
|---|---|---|---|
| 0 | Input | 16M × 64 × 7N | - |
| 1 | Flatten | 7168N | - |
| 2 | Dropout (p=0.5) | - | - |
| 3 | FullyConn | 8N | 1×1 |
| 4 | ReLU | 8N | - |
| 5 | FullyConn | 25N | 1×1 |

Table A.3: Network architecture used for subject discrimination network (adversary). The output of the network is a vector of dimension 25, values corresponding to the probability that the data tuple belongs to particular subject.

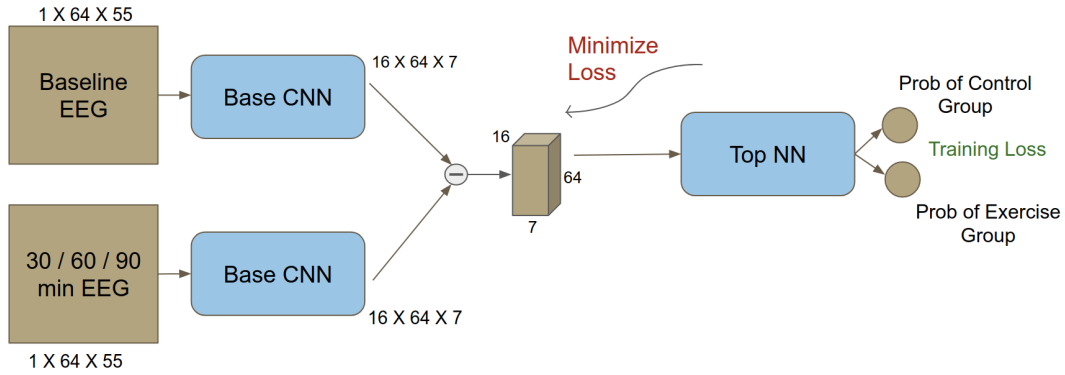


(a) Basic Architecture without adversary

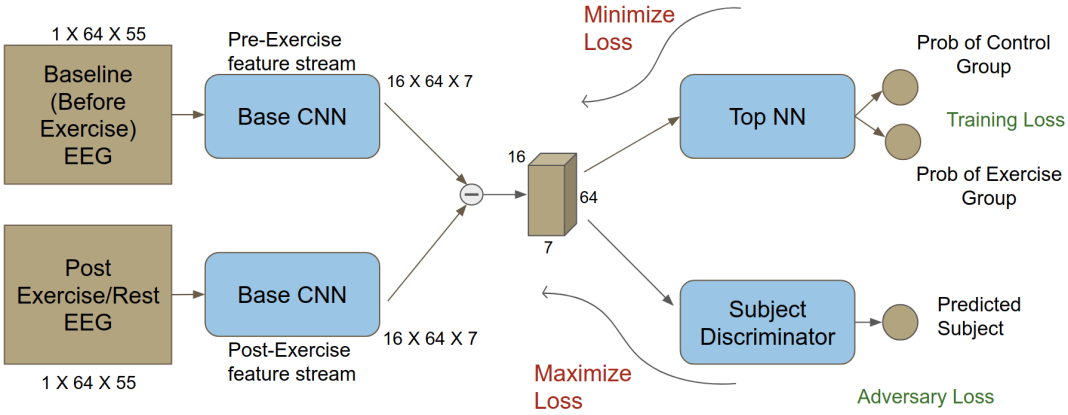


(b) Modified Architecture with adversary to avoid subject discrimination

Figure A.1: Deep Network Architecture. The initial choice of architecture (without any adversary) gives good subject prediction accuracy from features extracted by the Base CNN. Therefore, a subject discriminator of roughly the same model capacity as the Top NN is added. The subject discrimination acts as a regularizer while training and avoids the Base CNN from learning subject specific features.

**Topographical maps**

| Layer | Type | Maps and Neurons | Filter Size |
|---|---|---|---|
| 0 | Input | 3M × 64 × 64N | - |
| 1 | Conv | 16M × 32 × 32N | 5×5 |
| 2 | ReLU | 16M × 32 × 32N | - |
| 3 | MaxPool | 16M × 16 × 16N | 2×2 |
| 4 | Conv | 32M × 16 × 16N | 5×5 |
| 5 | ReLU | 32M × 16 × 16N | - |
| 6 | MaxPool | 32M × 8 × 8N | 2×2 |
| 7 | Conv | 64M × 8 × 8N | 3×3 |
| 8 | ReLU | 64M × 8 × 8N | - |
| 9 | MaxPool | 64M × 4 × 4N | 2×2 |

Table A.4: Network architecture used for EEG feature extraction network (Base CNN). The output of the network is a tensor of dimensions $64 \times 4 \times 4$.

| Layer | Type | Maps and Neurons | Filter Size |
|---|---|---|---|
| 0 | Input | 64M × 4 × 4N | - |
| 1 | Flatten | 1024N | - |
| 2 | Dropout (p=0.5) | - | - |
| 3 | FullyConn | 8N | 1×1 |
| 4 | ReLU | 8N | - |
| 5 | FullyConn | 2N | 1×1 |

Table A.5: Network architecture used for group discrimination network (Top NN). The output of the network is a vector of dimension 2, values corresponding to the probability that the data tuple belongs to particular class.

| Layer | Type | Maps and Neurons | Filter Size |
|---|---|---|---|
| 0 | Input | 64M × 4 × 4N | - |
| 1 | Flatten | 1024N | - |
| 2 | Dropout (p=0.5) | - | - |
| 3 | FullyConn | 8N | 1×1 |
| 4 | ReLU | 8N | - |
| 5 | FullyConn | 25N | 1×1 |

Table A.6: Network architecture used for subject discrimination network (adversary). The output of the network is a vector of dimension 25, values corresponding to the probability that the data tuple belongs to particular subject.

### Training curves

### Time-Frequency Maps

| Hyperparameter | Value |
|---|---|
| Learning Rate | 0.001 |
| Learning Rate Decay | 0.0001 |
| Weight Decay | 0.001 |

Table A.7: List of hyperparameters used for training the networks on TF maps.

### Topographical Maps

| Hyperparameter | Value |
|---|---|
| Learning Rate | 0.001 |
| Learning Rate Decay | 0.001 |
| Weight Decay | 0.03 |

Table A.8: List of hyperparameters used for training the networks on Topographical maps.

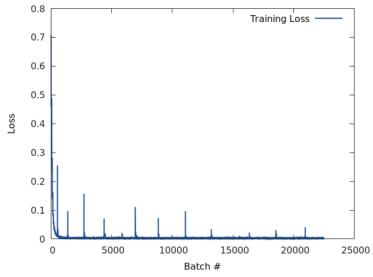


(a) Group prediction loss ($\lambda = 0$)

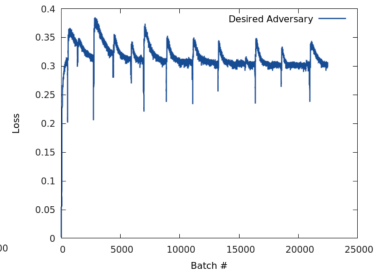


(b) KL divergence ($\lambda = 0$)

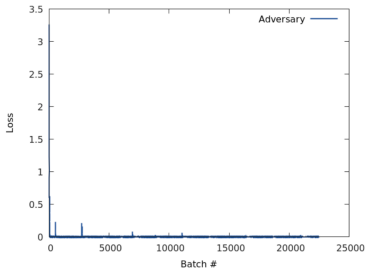


(c) Subject prediction loss ($\lambda = 0$)

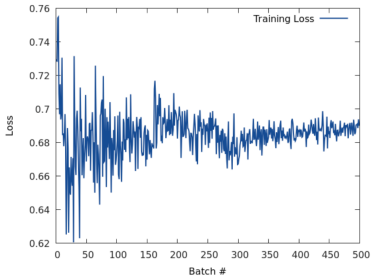


(d) Group prediction loss ($\lambda = 15$)

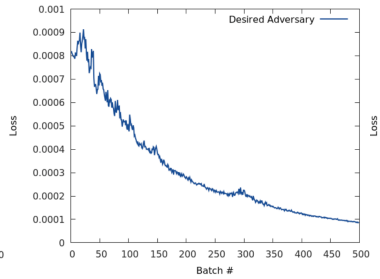


(e) KL divergence ($\lambda = 15$)

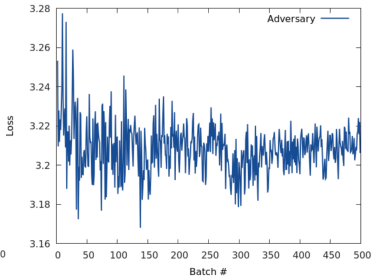


(f) Subject prediction loss ($\lambda = 15$)

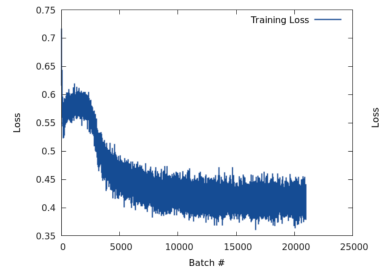


(g) Group prediction loss ($\lambda = 13$)

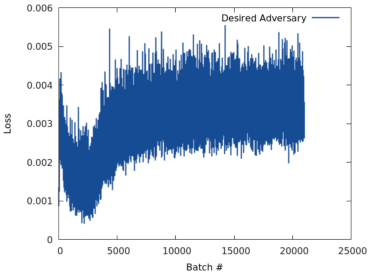


(h) KL divergence ($\lambda = 13$)

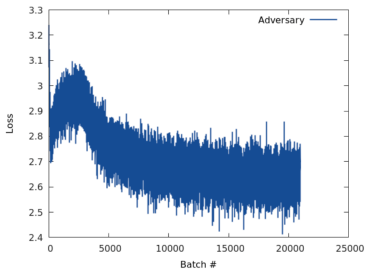


(i) Subject prediction loss ($\lambda = 13$)

Figure A.2: Time-Frequency Maps Training curves for three different weight values to the subject predictor regularizer.

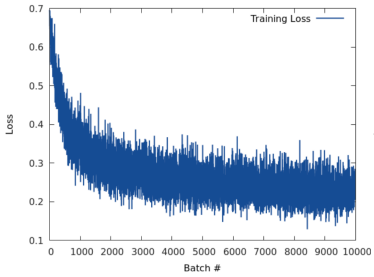


(a) Group prediction loss ($\lambda = 0$)

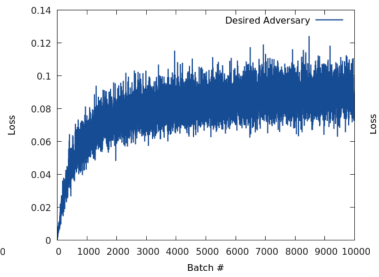


(b) KL divergence ($\lambda = 0$)

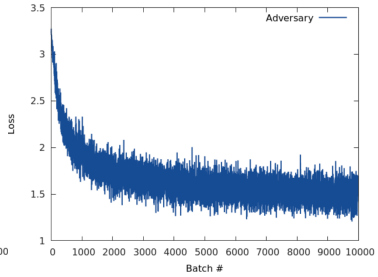


(c) Subject prediction loss ($\lambda = 0$)

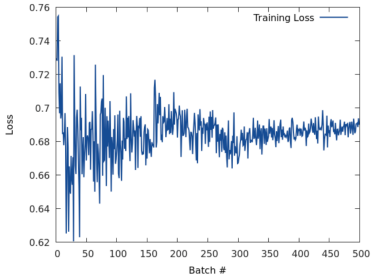


(d) Group prediction loss ($\lambda = 10$)

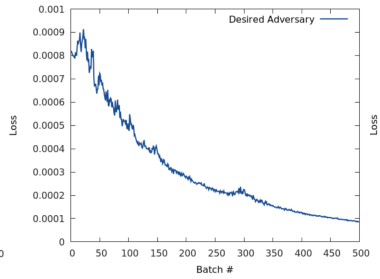


(e) KL divergence ($\lambda = 10$)

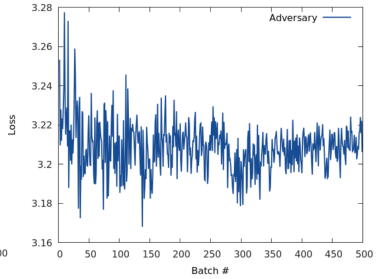


(f) Subject prediction loss ($\lambda = 10$)

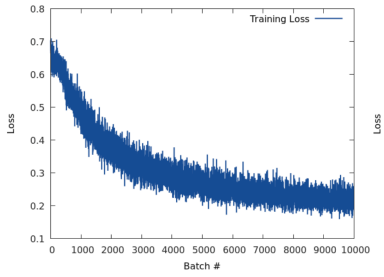


(g) Group prediction loss ($\lambda = 5$)

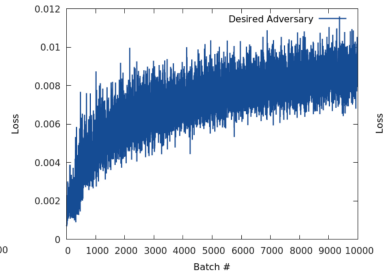


(h) KL divergence ($\lambda = 5$)

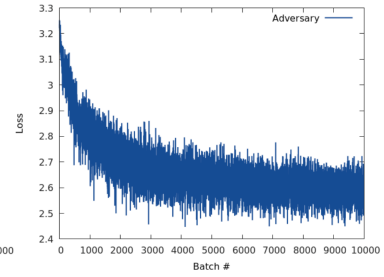


(i) Subject prediction loss ($\lambda = 5$)

Figure A.3: Topographical Maps Training curves for three different weight values to the subject predictor regularizer.

# Appendix B

## Brain Age Prediction

### Dimensionality reduction and GPR

| Model | Input Feature | MAE (yrs) | $R^2$ |
|---|---|---|---|
| GPR | WM | 6.82 ± 0.12 | 0.78 |
| GPR | GM | 6.42 ± 0.1 | 0.81 |
| GPR | cortical GM | 7.29 ± 0.13 | 0.76 |
| GPR | subcortical GM | 6.05 ± 0.11 | 0.83 |
| GPR | GM+WM+CSF | 5.64 ± 0.07 | 0.85 |
| GPR | subcortical GM+WM+CSF | 5.52 ± 0.06 | 0.86 |
| GPR+similarity | WM | 11.07 ± 0.48 | 0.49 |
| GPR+similarity | GM | 8.28 ± 0.31 | 0.70 |
| GPR+similarity | cortical GM | 10.1 ± 0.41 | 0.57 |
| GPR+similarity | subcortical GM | 8.04 ± 0.29 | 0.71 |
| GPR+similarity | GM+WM+CSF | 7.3 ± 0.14 | 0.77 |
| GPR+similarity | subcortical GM+WM+CSF | 7.29 ± 0.12 | 0.76 |
| GPR+PCA | WM | 11.71 ± 0.51 | 0.38 |
| GPR+PCA | GM | 8.63 ± 0.34 | 0.66 |
| GPR+PCA | cortical GM | 10.18 ± 0.4 | 0.54 |
| GPR+PCA | subcortical GM | 8.77 ± 0.32 | 0.65 |
| GPR+PCA | GM+WM+CSF | 9.23 ± 0.39 | 0.62 |
| GPR+CCA | WM | 6.77 ± 0.15 | 0.78 |
| GPR+CCA | GM | 6.37 ± 0.11 | 0.81 |
| GPR+CCA | cortical GM | 7.27 ± 0.16 | 0.77 |
| GPR+CCA | subcortical GM | 5.97 ± 0.08 | 0.83 |
| GPR+CCA | GM+WM+CSF | **5.57 ± 0.06** | **0.86** |
| GPR+CCA | subcortical GM+WM+CSF | **5.45 ± 0.07** | **0.86** |

Table B.1: Comparison of age prediction by GPR models with different dimensionality reduction techniques for GM, WM and CSF voxel intensities. Mean Absolute Error (MAE) and Coefficient of determination ($R^2$) are calculated over the testing set and averaged over 10-folds.

**Adding functional information from MEG data**

| Model | Input Feature | MAE (yrs) | $R^2$ |
|---|---|---|---|
| GPR | PSD | 9.75 ± 0.32 | 0.55 |
| GPR | AEC | 11.15 ± 0.43 | 0.45 |
| GPR | ILC | 14.47 ± 0.55 | 0.15 |
| GPR | PSD+AEC+ILC | 9.74 ± 0.3 | 0.57 |
| GPR+PCA | PSD | 10.68 ± 0.38 | 0.47 |
| GPR+PCA | AEC | 13.55 ± 0.51 | 0.20 |
| GPR+PCA | ILC | 14.60 ± 0.57 | 0.14 |
| GPR+PCA | PSD+AEC+ILC | 12.79 ± 0.48 | 0.28 |
| GPR+CCA | PSD | 12.75 ± 0.56 | 0.24 |
| GPR+CCA | AEC | 11.11 ± 0.4 | 0.46 |
| GPR+CCA | ILC | 15.88 ± 0.75 | 0.04 |
| GPR+CCA | PSD+AEC+ILC | 9.67 ± 0.28 | 0.58 |
| GPR | PSD+AEC+ILC+cATV | 8.37 ± 0.29 | 0.67 |
| GPR+PCA | PSD+AEC+ILC+cATV | 12.29 ± 0.48 | 0.33 |
| GPR+CCA | PSD+AEC+ILC+cATV | 8.28 ± 0.26 | 0.70 |
| GPR | PSD+AEC+ILC+MRI | 5.61 ± 0.08 | 0.85 |
| GPR+PCA | PSD+AEC+ILC+MRI | 7.65 ± 0.18 | 0.73 |
| GPR+CCA | PSD+AEC+ILC+MRI | **5.28 ± 0.07** | **0.87** |

Table B.2: Comparison of age prediction by GPR models with different dimensionality reduction techniques for MEG features, MEG+cortical area,thickness & volume, and MEG+MRI features. Mean Absolute Error (MAE) and Coefficient of determination ($R^2$) are calculated over the testing set and averaged over 10-folds.

**Rich Club Organization**

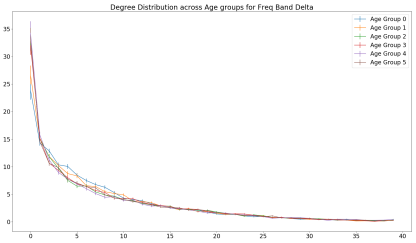


(a) Delta Band

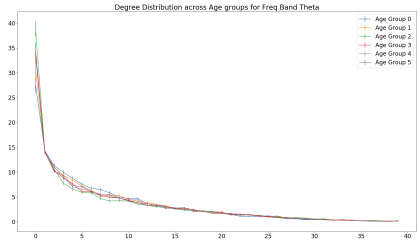


(b) Theta Band

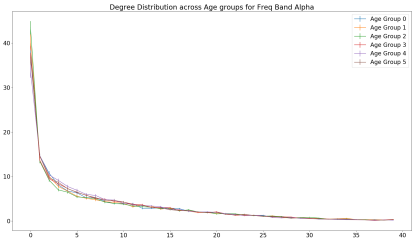


(c) Alpha Band

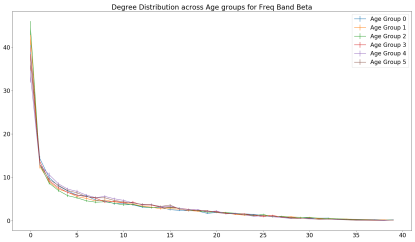


(d) Beta Band

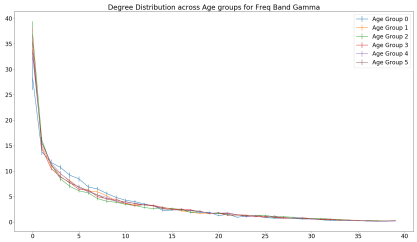


(e) Gamma Band

Figure B.1: Degree distribution for the graph generated by binarizing the FC matrix of each frequency band. All frequency bands show a tailed distribution without any differences across age groups.

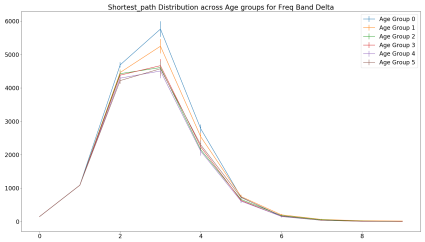


(a) Delta Band

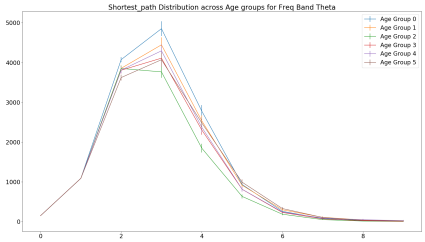


(b) Theta Band

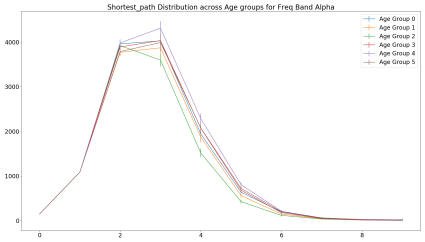


(c) Alpha Band

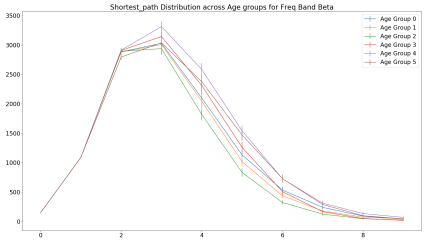


(d) Beta Band

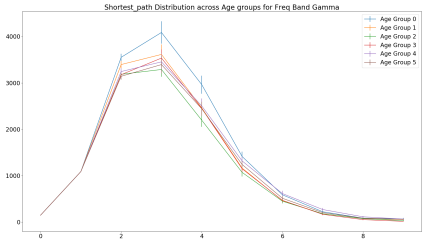


(e) Gamma Band

Figure B.2: Distribution of shortest path length for the graph generated by binarizing the FC matrix of each frequency band. All frequency bands show a peak at 3 without any differences across age groups.

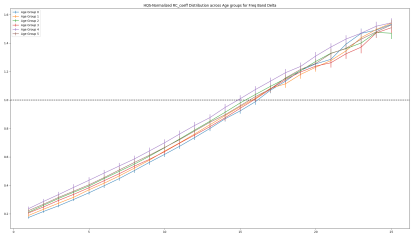

(a) Delta Band

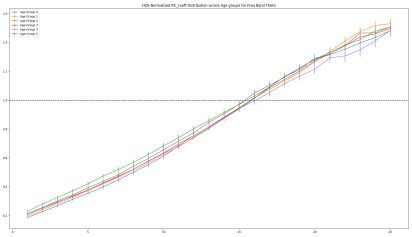

(b) Theta Band

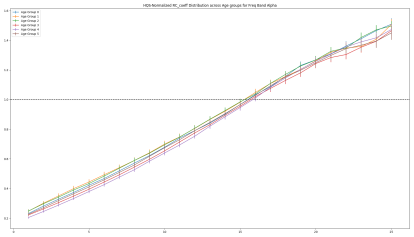

(c) Alpha Band

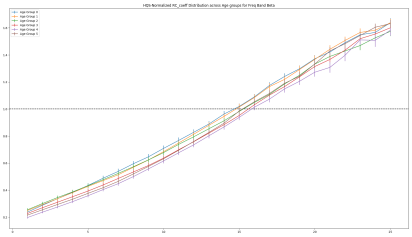

(d) Beta Band

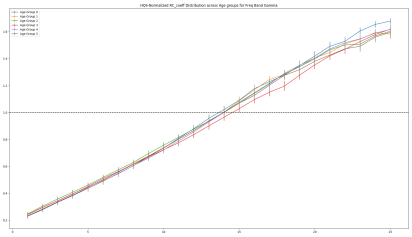

(e) Gamma Band

Figure B.3: HQS-normalized RC coefficients plotted as a function of node degree for the graph generated by binarizing the FC matrix of each frequency band. All frequency bands show a RC organization above node degree = 17 without any differences across age groups.

**GCN architecture**

Since the length of shortest path distribution peaks at 3, we used all GCN models with 3 layers – 2 hidden layer and 1 output layer. Following the feature transformation by GCN models, the features from all nodes were vectorized and passed through a fully connected layer: $1776 \rightarrow 1$ (each graph had 148 nodes and 12 features for each, hence size of feature vector $= 148 \times 12 = 1776$).

| Model | Hidden Layers | Hidden dimension | Input/Output dimension | Architecture |
|---|---|---|---|---|
| GCN | 2 | 36 | 12 | $12 \rightarrow 36 \rightarrow 36 \rightarrow 12$ |
| GraphSAGE | 2 | 36 | 12 | $12 \rightarrow 36 \rightarrow 36 \rightarrow 12$ |

Table B.3: Model Architecture for GCN and GraphSAGE.

| Hyperparameter | Value |
|---|---|
| Optimizer | Adam |
| Learning Rate | 0.006 |
| Learning Rate Decay | 0.003 |
| Weight Decay | 0.0001 |
| Dropout | 0.75 |
| Batch size | 16 |
| FC threshold | 95 percentile |

Table B.4: List of hyperparameters used for training the GCN networks. Input features were – PSD, cortical area, thickness & volume, mean & standard deviation of MRI intensities.

References


[1] Alison Abbott. A problem for our age. *Nature*, 475(7355):S2, 2011.

[2] Ossama Abdel-Hamid, Abdel-rahman Mohamed, Hui Jiang, Li Deng, Gerald Penn, and Dong Yu. Convolutional neural networks for speech recognition. *IEEE/ACM Transactions on audio, speech, and language processing*, 22(10):1533–1545, 2014.

[3] Pulkit Agrawal, Joao Carreira, and Jitendra Malik. Learning to see by moving. In *Computer Vision (ICCV), 2015 IEEE International Conference on*, pages 37–45. IEEE, 2015.

[4] Obada Al Zoubi, Chung Ki Wong, Rayus Tiberius Kuplicki, Hung-wen Yeh, Ahmad Mayeli, Hazem Refai, Martin Paulus, and Jerzy Bodurka. Predicting age from brain eeg signals–a machine learning approach. *Frontiers in aging neuroscience*, 10:184, 2018.

[5] Galen Andrew, Raman Arora, Jeff Bilmes, and Karen Livescu. Deep canonical correlation analysis. In *International conference on machine learning*, pages 1247–1255, 2013.

[6] Kai Keng Ang, Zheng Yang Chin, Haihong Zhang, and Cuntai Guan. Filter bank common spatial pattern (fbcsp) in brain-computer interface. In *Neural Networks, 2008. IJCNN 2008.(IEEE World Congress on Computational Intelligence). IEEE International Joint Conference on*, pages 2390–2397. IEEE, 2008.

[7] Andreas Antoniades, Loukianos Spyrou, Clive Cheong Took, and Saeid Sanei. Deep learning for epileptic intracranial eeg data. In *Machine Learning for Signal Processing (MLSP), 2016 IEEE 26th International Workshop on*, pages 1–6. IEEE, 2016.

[8] Andrea Avena-Koenigsberger, Bratislav Misic, and Olaf Sporns. Communication dynamics in complex brain networks. *Nature Reviews Neuroscience*, 19(1):17, 2018.

[9] Ashley S Bangert, Patricia A Reuter-Lorenz, Christine M Walsh, Anna B Schachter, and Rachael D Seidler. Bimanual coordination and aging: neurobehavioral implications. *Neuropsychologia*, 48(4):1165–1170, 2010.

[10] Debasish Basak, Srimanta Pal, and Dipak Chandra Patranabis. Support vector regression. *Neural Information Processing-Letters and Reviews*, 11(10):203–224, 2007.

[11] Ali Bashashati, Mehrdad Fatourechi, Rabab K Ward, and Gary E Birch. A survey of signal processing algorithms in brain–computer interfaces based on electrical brain signals. *Journal of Neural engineering*, 4(2):R32, 2007.

[12] Pouya Bashivan, Irina Rish, Mohammed Yeasin, and Noel Codella. Learning representations from eeg with deep recurrent-convolutional neural networks. *arXiv preprint arXiv:1511.06448*, 2015.

[13] Yoshua Bengio, Pascal Lamblin, Dan Popovici, and Hugo Larochelle. Greedy layer-wise training of deep networks. In *Advances in neural information processing systems*, pages 153–160, 2007.

[14] NC Berchtold, G Chinn, M Chou, JP Kesslak, and CW Cotman. Exercise primes a molecular memory for brain-derived neurotrophic factor protein induction in the rat hippocampus. *Neuroscience*, 133(3):853–861, 2005.

[15] Smriti Bhagat, Graham Cormode, and S Muthukrishnan. Node classification in social networks. In *Social network data analytics*, pages 115–148. Springer, 2011.

[16] Tjeerd W Boonstra, Andreas Daffertshofer, Michael Breakspear, and Peter J Beek. Multivariate time–frequency analysis of electromagnetic brain activity during bimanual motor learning. *Neuroimage*, 36(2):370–377, 2007.

[17] Rory Boyle, Lee Jollans, Laura M Rueda-Delgado, Rossella Rizzo, Gorsev G Yener, Jason P McMorrow, Silvin P Knight, Daniel Carey, Ian H Robertson, DD Emek-Savaş, et al. Brain-predicted age difference score is related to specific cognitive functions: A multi-site replication analysis. *BioRxiv*, page 652867, 2019.

[18] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.

[19] O Bugiani, S Salvarani, F Perdelli, GL Mancardi, and A Leonardi. Nerve cell loss with aging in the putamen. *European neurology*, 17(5):286–291, 1978.

[20] Danilo Bzdok. Classical statistics and statistical learning in imaging neuroscience. *Frontiers in neuroscience*, 11:543, 2017.

[21] Jeremy B Caplan, Monica Bottomley, Pardeep Kang, and Roger A Dixon. Distinguishing rhythmic from non-rhythmic brain activity during rest in healthy neurocognitive aging. *NeuroImage*, 112:341–352, 2015.

[22] Hubert Cecotti and Axel Graser. Convolutional neural networks for p300 detection with application to brain-computer interfaces. *IEEE transactions on pattern analysis and machine intelligence*, 33(3):433–445, 2011.

[23] Andrea Cherubini, Maria Eugenia Caligiuri, Patrice Péran, Umberto Sabatini, Carlo Cosentino, and Francesco Amato. Importance of multimodal mri in characterizing brain tissue and its potential application for individual age prediction. *IEEE journal of biomedical and health informatics*, 20(5):1232–1239, 2016.

[24] Radoslaw Martin Cichy, Aditya Khosla, Dimitrios Pantazis, Antonio Torralba, and Aude Oliva. Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence. *Scientific reports*, 6:27755, 2016.

[25] CE Coffey, WE Wilkinson, LA Parashos, SAR Soady, RJ Sullivan, LJ Patterson, GS Figiel, MC Webb, CE Spritzer, and WT Djang. Quantitative cerebral anatomy of the aging human brain a cross-sectional study using magnetic resonance imaging. *Neurology*, 42(3):527–527, 1992.

[26] James Cole, Katja Franke, and Nicolas Cherbuin. Quantification of the biological age of the brain using neuroimaging. 2018.

[27] James H Cole, Tiina Annus, Liam R Wilson, Ridhaa Remtulla, Young T Hong, Tim D Fryer, Julio Acosta-Cabronero, Arturo Cardenas-Blanco, Robert Smith, David K Menon, et al. Brain-predicted age in down syndrome is associated with beta amyloid deposition and cognitive decline. *Neurobiology of aging*, 56:41–49, 2017.

[28] James H Cole, Robert Leech, David J Sharp, and Alzheimer's Disease Neuroimaging Initiative. Prediction of brain age suggests accelerated atrophy after traumatic brain injury. *Annals of neurology*, 77(4):571–581, 2015.

[29] James H Cole, Rudra PK Poudel, Dimosthenis Tsagkrasoulis, Matthan WA Caan, Claire Steves, Tim D Spector, and Giovanni Montana. Predicting brain age with deep learning from raw imaging data results in a reliable and heritable biomarker. *NeuroImage*, 163:115–124, 2017.

[30] James H Cole, Stuart J Ritchie, Mark E Bastin, MC Valdés Hernández, S Muñoz Maniega, Natalie Royle, Janie Corley, Alison Pattie, Sarah E Harris, Qian Zhang, et al. Brain age predicts mortality. *Molecular psychiatry*, 2017.

[31] James H Cole, Jonathan Underwood, Matthan WA Caan, Davide De Francesco, Rosan A van Zoest, Robert Leech, Ferdinand WNM Wit, Peter Portegies, Gert J Geurtsen, Ben A Schmand, et al. Increased brain-predicted aging in treated hiv disease. *Neurology*, 88(14):1349–1357, 2017.

[32] Ronan Collobert, Koray Kavukcuoglu, and Clément Farabet. Torch7: A matlab-like environment for machine learning. In *BigLearn, NIPS workshop*, number EPFL-CONF-192376, 2011.

[33] Nathan E Crone, Diana L Miglioretti, Barry Gordon, Jeffrey M Sieracki, Michael T Wilson, Sumio Uematsu, and Ronald P Lesser. Functional mapping of human sensorimotor cortex with electrocorticographic spectral analysis. i. alpha and beta event-related desynchronization. *Brain: a journal of neurology*, 121(12):2271–2299, 1998.

[34] Tarrant DR Cummins and Simon Finnigan. Theta power is reduced in healthy cognitive aging. *International Journal of Psychophysiology*, 66(1):10–17, 2007.

[35] Fabien Dal Maso, Bennet Desormeau, Marie-Hélène Boudrias, and Marc Roig. Acute cardiovascular exercise promotes functional changes in cortico-motor networks during the early stages of motor memory consolidation. *NeuroImage*, 174:380–392, 2018.

[36] Jessica S Damoiseaux, CF Beckmann, EJ Sanz Arigita, F Barkhof, Ph Scheltens, CJ Stam, SM Smith, and SARB Rombouts. Reduced resting-state brain activity in the default network in normal aging. *Cerebral cortex*, 18(8):1856–1864, 2007.

[37] Koel Das, Barry Giesbrecht, and Miguel P Eckstein. Predicting variations of perceptual performance across individuals from neural activity using pattern classifiers. *Neuroimage*, 51(4):1425–1437, 2010.

[38] Arnaud Delorme and Scott Makeig. Eeglab: an open source toolbox for analysis of single-trial eeg dynamics including independent component analysis. *Journal of neuroscience methods*, 134(1):9–21, 2004.

[39] Emily L Dennis and Paul M Thompson. Functional brain connectivity using fmri in aging and alzheimers disease. *Neuropsychology review*, 24(1):49–62, 2014.

[40] Christophe Destrieux, Bruce Fischl, Anders Dale, and Eric Halgren. Automatic parcellation of human cortical gyri and sulci using standard anatomical nomenclature. *Neuroimage*, 53(1):1–15, 2010.

[41] Arne Dietrich. Transient hypofrontality as a mechanism for the psychological effects of exercise. *Psychiatry research*, 145(1):79–83, 2006.

[42] Stavros I Dimitriadis and Christos I Salis. Mining time-resolved functional brain graphs to an eeg-based chronnectomic brain aged index (cbai). *Frontiers in human neuroscience*, 11:423, 2017.

[43] Nico UF Dosenbach, Binyam Nardos, Alexander L Cohen, Damien A Fair, Jonathan D Power, Jessica A Church, Steven M Nelson, Gagan S Wig, Alecia C Vogel, Christina N Lessov-Schlaggar, et al. Prediction of individual brain maturity using fmri. *Science*, 329(5997):1358–1361, 2010.

[44] G Dumermuth and L Molinari. Spectral analysis of the eeg. *Neuropsychobiology*, 17(1-2):85–99, 1987.

[45] Bradley Efron, Elizabeth Halloran, and Susan Holmes. Bootstrap confidence levels for phylogenetic trees. *Proceedings of the National Academy of Sciences*, 93(23):13429–13429, 1996.

[46] Bradley Efron and Robert Tibshirani. Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical science*, pages 54–75, 1986.

[47] Marc O Ernst and Martin S Banks. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870):429, 2002.

[48] Anders M Fjell, Lars T Westlye, Inge Amlien, Thomas Espeseth, Ivar Reinvang, Naftali Raz, Ingrid Agartz, David H Salat, Doug N Greve, Bruce Fischl, et al. High consistency of regional cortical thinning in aging across multiple samples. *Cerebral cortex*, 19(9):2001–2012, 2009.

[49] Anders M Fjell, Lars T Westlye, Håkon Grydeland, Inge Amlien, Thomas Espeseth, Ivar Reinvang, Naftali Raz, Dominic Holland, Anders M Dale, Kristine B Walhovd, et al. Critical ages in the life course of the adult brain: nonlinear subcortical aging. *Neurobiology of aging*, 34(10):2239–2247, 2013.

[50] Vladimir Fonov, Andrew Doyle, Alan C Evans, and D Louis Collins. Neuromtl iseg challenge methods. *BioRxiv*, page 278465, 2018.

[51] Katja Franke and Christian Gaser. Longitudinal changes in individual brainage in healthy aging, mild cognitive impairment, and alzheimers disease. *GeroPsych*, 2012.

[52] Katja Franke, Christian Gaser, Brad Manor, and Vera Novak. Advanced brainage in older adults with type 2 diabetes mellitus. *Frontiers in aging neuroscience*, 5:90, 2013.

[53] Katja Franke, Gabriel Ziegler, Stefan Klöppel, Christian Gaser, Alzheimer's Disease Neuroimaging Initiative, et al. Estimating the age of healthy subjects from t1-weighted mri scans using kernel methods: exploring the influence of various parameters. *Neuroimage*, 50(3):883–892, 2010.

[54] John DE Gabrieli, Satrajit S Ghosh, and Susan Whitfield-Gabrieli. Prediction as a humanitarian and pragmatic contribution from human cognitive neuroscience. *Neuron*, 85(1):11–26, 2015.

[55] Kais Gadhoumi, Jean-Marc Lina, Florian Mormann, and Jean Gotman. Seizure prediction for therapeutic devices: A review. *Journal of neuroscience methods*, 260:270–282, 2016.

[56] Joseph M Galea, Neil B Albert, Thomas Ditye, and R Chris Miall. Disruption of the dorsolateral prefrontal cortex facilitates the consolidation of procedural skills. *Journal of cognitive neuroscience*, 22(6):1158–1164, 2010.

[57] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016.

[58] Christian Gaser, Katja Franke, Stefan Klöppel, Nikolaos Koutsouleris, Heinrich Sauer, Alzheimer's Disease Neuroimaging Initiative, et al. Brainage in mild

cognitive impaired patients: predicting the conversion to alzheimers disease. *PloS one*, 8(6):e67346, 2013.

[59] Ian Goodfellow, Honglak Lee, Quoc V Le, Andrew Saxe, and Andrew Y Ng. Measuring invariances in deep networks. In *Advances in neural information processing systems*, pages 646–654, 2009.

[60] Jessica A Grahn, John A Parkinson, and Adrian M Owen. The cognitive functions of the caudate nucleus. *Progress in neurobiology*, 86(3):141–155, 2008.

[61] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks. In *Acoustics, speech and signal processing (icassp), 2013 ieee international conference on*, pages 6645–6649. IEEE, 2013.

[62] Boris Gutmann, Andreas Mierau, Thorben Hülsdünker, Carolin Hildebrand, Axel Przyklenk, Wildor Hollmann, and Heiko Klaus Strüder. Effects of physical exercise on individual resting state eeg alpha peak frequency. *Neural plasticity*, 2015, 2015.

[63] Mehdi Hajinoroozi, Zijing Mao, Tzyy-Ping Jung, Chin-Teng Lin, and Yufei Huang. Eeg-based prediction of driver's cognitive performance by deep convolutional neural network. *Signal Processing: Image Communication*, 47:549–555, 2016.

[64] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, pages 1024–1034, 2017.

[65] William L Hamilton, Rex Ying, and Jure Leskovec. Representation learning on graphs: Methods and applications. *arXiv preprint arXiv:1709.05584*, 2017.

[66] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.

[67] Elizabeth Heinrichs-Graham and Tony W Wilson. Is an absolute level of cortical beta suppression required for proper movement? magnetoencephalographic evidence from healthy aging. *Neuroimage*, 134:514–521, 2016.

[68] Henk T Hendricks, Jacques van Limbeek, Alexander C Geurts, and Machiel J Zwarts. Motor recovery after stroke: a systematic review of the literature. *Archives of physical medicine and rehabilitation*, 83(11):1629–1637, 2002.

[69] Paul Herent, Simon Jegou, Gilles Wainrib, and Thomas Clozel. Brain age prediction of healthy subjects on anatomic mri with deep learning: going beyond with an" explainable ai" mindset. *bioRxiv*, page 413302, 2018.

[70] Geoffrey E Hinton and Ruslan R Salakhutdinov. Reducing the dimensionality of data with neural networks. *science*, 313(5786):504–507, 2006.

[71] Larson J Hogstrom, Lars T Westlye, Kristine B Walhovd, and Anders M Fjell. The structure of the cerebral cortex across adult life: age-related patterns of surface area, thickness, and gyrification. *Cerebral cortex*, 23(11):2521–2530, 2012.

[72] Michael E Hopkins, F Caroline Davis, Michelle R VanTieghem, Paul J Whalen, and David J Bucci. Differential effects of acute and regular physical exercise on cognition and affect. *Neuroscience*, 215:59–68, 2012.

[73] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.

[74] S Houweling, Andreas Daffertshofer, Bob W van Dijk, and Peter J Beek. Neural changes induced by learning a challenging perceptual-motor task. *Neuroimage*, 41(4):1395–1407, 2008.

[75] Chung-Ju Huang, Chin-Wen Huang, Yu-Jung Tsai, Chia-Liang Tsai, Yu-Kai Chang, and Tsung-Min Hung. A preliminary examination of aerobic exercise effects on resting eeg in children with adhd. *Journal of attention disorders*, 21(11):898–903, 2017.

[76] Lena Hübner, Ben Godde, and Claudia Voelcker-Rehage. Older adults reveal enhanced task-related beta power decreases during a force modulation task. *Behavioural brain research*, 345:104–113, 2018.

[77] Mark Jenkinson, Christian F Beckmann, Timothy EJ Behrens, Mark W Woolrich, and Stephen M Smith. Fsl. *Neuroimage*, 62(2):782–790, 2012.

[78] Ole Jensen and Claudia D Tesche. Frontal theta activity in humans increases with memory load in a working memory task. *European journal of Neuroscience*, 15(8):1395–1399, 2002.

[79] Haitao Jiang, Jiajia Guo, Hongwei Du, Jinzhang Xu, and Bensheng Qiu. Transfer learning on t1-weighted images for brain age estimation. 2019.

[80] Ian Jolliffe. Principal component analysis. In *International encyclopedia of statistical science*, pages 1094–1096. Springer, 2011.

[81] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1725–1732, 2014.

[82] Jeremy Kawahara, Colin J Brown, Steven P Miller, Brian G Booth, Vann Chau, Ruth E Grunau, Jill G Zwicker, and Ghassan Hamarneh. Brainnetcnn: convolutional neural networks for brain networks; towards predicting neurodevelopment. *NeuroImage*, 146:1038–1049, 2017.

[83] Bradley R King, Kerstin Hoedlmoser, Franziska Hirschauer, Nina Dolfen, and Genevieve Albouy. Sleeping on the motor engram: the multifaceted nature of sleep-related motor memory consolidation. *Neuroscience & Biobehavioral Reviews*, 80:1–22, 2017.

[84] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[85] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[86] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.

[87] Brian Kirkpatrick, Erick Messias, Philip D Harvey, Emilio Fernandez-Egea, and Christopher R Bowie. Is schizophrenia a syndrome of accelerated aging? *Schizophrenia bulletin*, 34(6):1024–1032, 2007.

[88] André Knops, Bertrand Thirion, Edward M Hubbard, Vincent Michel, and Stanislas Dehaene. Recruitment of an area involved in eye movements during mental arithmetic. *Science*, 324(5934):1583–1585, 2009.

[89] Nikolaos Koutsouleris, Christos Davatzikos, Stefan Borgwardt, Christian Gaser, Ronald Bottlender, Thomas Frodl, Peter Falkai, Anita Riecher-Rössler, Hans-Jürgen Möller, Maximilian Reiser, et al. Accelerated brain aging in schizophrenia and beyond: a neuroanatomical marker of psychiatric disorders. *Schizophrenia bulletin*, 40(5):1140–1153, 2013.

[90] Nikolaus Kriegeskorte and Pamela K Douglas. Interpreting encoding and decoding models. *Current opinion in neurobiology*, 55:167–179, 2019.

[91] K Ranga Krishnan, Mustafa M Husain, WM McDonald, P Murali Doraiswamy, GS Figiel, OB Boyko, EH Ellinwood, and CB Nemeroff. In vivo stereological assessment of caudate volume in man: effect of normal aging. *Life sciences*, 47(15):1325–1329, 1990.

[92] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.

[93] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533*, 2016.

[94] Zeb Kurth-Nelson, Marcos Economides, Raymond J Dolan, and Peter Dayan. Fast sequences of non-spatial state representations in humans. *Neuron*, 91(1):194–204, 2016.

[95] Valentina La Corte, Marco Sperduti, Caroline Malherbe, François Vialatte, Stéphanie Lion, Thierry Gallarda, Catherine Oppenheim, and Pascale Piolino. Cognitive decline and reorganization of functional connectivity in healthy aging: the pivotal role of the salience network in the prediction of age and cognitive performances. *Frontiers in aging neuroscience*, 8:204, 2016.

[96] Hugo Larochelle, Yoshua Bengio, Jérôme Louradour, and Pascal Lamblin. Exploring strategies for training deep neural networks. *Journal of machine learning research*, 10(Jan):1–40, 2009.

[97] Vernon J Lawhern, Amelia J Solon, Nicholas R Waytowich, Stephen M Gordon, Chou P Hung, and Brent J Lance. Eegnet: a compact convolutional neural network for eeg-based brain–computer interfaces. *Journal of neural engineering*, 15(5):056013, 2018.

[98] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436, 2015.

[99] Vera Maria Leirer, Christian Wienbruch, Stephan Kolassa, Winfried Schlee, Thomas Elbert, and Iris-Tatjana Kolassa. Changes in cortical slow wave activity in healthy aging. *Brain imaging and behavior*, 5(3):222–228, 2011.

[100] Hongming Li, Theodore D Satterthwaite, and Yong Fan. Brain age prediction based on resting-state functional connectivity patterns using convolutional neural networks. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 101–104. IEEE, 2018.

[101] David Liben-Nowell and Jon Kleinberg. The link-prediction problem for social networks. *Journal of the American society for information science and technology*, 58(7):1019–1031, 2007.

[102] Franziskus Liem, Gaël Varoquaux, Jana Kynast, Frauke Beyer, Shahrzad Kharabian Masouleh, Julia M Huntenburg, Leonie Lampe, Mehdi Rahim, Alexandre Abraham, R Cameron Craddock, et al. Predicting brain-age from multimodal imaging data captures cognitive impairment. *Neuroimage*, 148:179–188, 2017.

[103] Zachary C Lipton. The mythos of model interpretability. *arXiv preprint arXiv:1606.03490*, 2016.

[104] Charalambos Loizides, Achilleas Achilleos, Gian Domenico Iannetti, and Georgios D Mitsis. Assessment of nonlinear interactions in event-related potentials elicited by stimuli presented at short interstimulus intervals using single-trial data. *Journal of neurophysiology*, 113(10):3623–3633, 2015.

[105] Eileen Luders, Nicolas Cherbuin, and Christian Gaser. Estimating brain age using high-resolution pattern recognition: younger brains in long-term meditation practitioners. *Neuroimage*, 134:508–513, 2016.

[106] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.

[107] Aravindh Mahendran and Andrea Vedaldi. Understanding deep image representations by inverting them. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5188–5196, 2015.

[108] Aravindh Mahendran and Andrea Vedaldi. Visualizing deep convolutional neural networks using natural pre-images. *International Journal of Computer Vision*, 120(3):233–255, 2016.

[109] Ran Manor and Amir B Geva. Convolutional neural network for multi-category rapid serial visual presentation bci. *Frontiers in computational neuroscience*, 9:146, 2015.

[110] Patrick L McGeer, Edith G McGeer, and Joane S Suzuki. Aging and extrapyramidal function. *Archives of neurology*, 34(1):33–35, 1977.

[111] Anthony Randal McIntosh and Nancy J Lobaugh. Partial least squares analysis of neuroimaging data: applications and advances. *Neuroimage*, 23:S250–S263, 2004.

[112] AR McIntosh, FL Bookstein, James V Haxby, and CL Grady. Spatial pattern analysis of functional brain images using partial least squares. *Neuroimage*, 3(3):143–157, 1996.

[113] Bratislav Mišić, Olaf Sporns, and Anthony R McIntosh. Communication efficiency and congestion of signal traffic in large-scale brain networks. *PLoS computational biology*, 10(1):e1003427, 2014.

[114] Saba Moghimi, Azadeh Kushki, Anne Marie Guerguerian, and Tom Chau. A review of eeg-based brain-computer interfaces as access pathways for individuals with severe disabilities. *Assistive Technology*, 25(2):99–110, 2013.

[115] Christa Neuper and Gert Pfurtscheller. Event-related dynamics of cortical rhythms: frequency-specific features and functional correlates. *International journal of psychophysiology*, 43(1):41–58, 2001.

[116] Femke Nijboer, EW Sellers, Jürgen Mellinger, Mary Ann Jordan, Tamara Matuz, Adrian Furdea, Sebastian Halder, Ursula Mochty, DJ Krusienski, TM Vaughan, et al. A p300-based brain–computer interface for people with amyotrophic lateral sclerosis. *Clinical neurophysiology*, 119(8):1909–1916, 2008.

[117] Fatemeh Ostadan, Carla Centeno, Jean-Felix Daloze, Mira Frenn, Jesper Lundbye-Jensen, and Marc Roig. Changes in corticospinal excitability during consolidation predict acute exercise-induced off-line gains in procedural memory. *Neurobiology of learning and memory*, 136:196–203, 2016.

[118] Chethan Pandarinath, Daniel J OShea, Jasmine Collins, Rafal Jozefowicz, Sergey D Stavisky, Jonathan C Kao, Eric M Trautmann, Matthew T Kaufman, Stephen I Ryu, Leigh R Hochberg, et al. Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature methods*, page 1, 2018.

[119] Heath R Pardoe, James H Cole, Karen Blackmon, Thomas Thesen, Ruben Kuzniecky, Human Epilepsy Project Investigators, et al. Structural brain changes in medically refractory focal epilepsy resemble premature brain aging. *Epilepsy research*, 133:28–32, 2017.

[120] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[121] Francisco Pereira, Tom Mitchell, and Matthew Botvinick. Machine learning classifiers and fmri: a tutorial overview. *Neuroimage*, 45(1):S199–S209, 2009.

[122] Gert Pfurtscheller, Bernard Graimann, Jane E Huggins, Simon P Levine, and Lori A Schuh. Spatiotemporal patterns of beta desynchronization and gamma synchronization in corticographic data during self-paced movement. *Clinical neurophysiology*, 114(7):1226–1236, 2003.

[123] Sergey M Plis, Devon R Hjelm, Ruslan Salakhutdinov, Elena A Allen, Henry J Bockholt, Jeffrey D Long, Hans J Johnson, Jane S Paulsen, Jessica A Turner, and Vince D Calhoun. Deep learning for neuroimaging: a validation study. *Frontiers in neuroscience*, 8:229, 2014.

[124] B Pollok, D Latz, V Krause, M Butz, and A Schnitzler. Changes of motor-cortical oscillations associated with motor learning. *Neuroscience*, 275:47–53, 2014.

[125] Michel JAM Putten, Sebastian Olbrich, and Martijn Arns. Predicting sex from brain rhythms with deep learning. *Scientific reports*, 8(1):3069, 2018.

[126] Ander Ramos-Murguialday, Doris Broetz, Massimiliano Rea, Leonhard Läer, Özge Yilmaz, Fabricio L Brasil, Giulia Liberati, Marco R Curado, Eliana Garcia-Cossio, Alexandros Vyziotis, et al. Brain–machine interface in chronic

stroke rehabilitation: a controlled study. *Annals of neurology*, 74(1):100–108, 2013.

[127] Carl Edward Rasmussen. Gaussian processes in machine learning. In *Advanced lectures on machine learning*, pages 63–71. Springer, 2004.

[128] Naftali Raz and Karen M Rodrigue. Differential aging of the brain: patterns, cognitive correlates and modifiers. *Neuroscience & Biobehavioral Reviews*, 30(6):730–748, 2006.

[129] Caroline V Robertson and Frank E Marino. Prefrontal and motor cortex eeg responses and their relationship to ventilatory thresholds during exhaustive incremental exercise. *European journal of applied physiology*, 115(9):1939–1948, 2015.

[130] Marc Roig, Sasja Nordbrandt, Svend Sparre Geertsen, and Jens Bo Nielsen. The effects of cardiovascular exercise on human memory: a review with meta-analysis. *Neuroscience & Biobehavioral Reviews*, 37(8):1645–1666, 2013.

[131] Marc Roig, Richard Thomas, Cameron S Mang, Nicholas J Snow, Fatemeh Ostadan, Lara A Boyd, and Jesper Lundbye-Jensen. Time-dependent effects of cardiovascular exercise on memory. *Exercise and sport sciences reviews*, 44(2):81–88, 2016.

[132] Holly E Rossiter, Emma M Davis, Ella V Clark, Marie-Hélène Boudrias, and Nick S Ward. Beta oscillations reflect changes in motor cortex inhibition in healthy ageing. *Neuroimage*, 91:360–365, 2014.

[133] Jamie A Ruffing, Jeri W Nieves, Marsha Zion, Susan Tendy, Patricia Garrett, Robert Lindsay, and Felicia Cosman. The influence of lifestyle, menstrual function and oral contraceptive use on bone mass and size in female military cadets. *Nutrition & metabolism*, 4(1):17, 2007.

[134] David H Salat, Randy L Buckner, Abraham Z Snyder, Douglas N Greve, Rahul SR Desikan, Evelina Busa, John C Morris, Anders M Dale, and Bruce Fischl. Thinning of the cerebral cortex in aging. *Cerebral cortex*, 14(7):721–730, 2004.

[135] Riitta Salmelin, M Hámáaláinen, M Kajola, and R Hari. Functional segregation of movement-related rhythmic activity in the human brain. *Neuroimage*, 2(4):237–243, 1995.

[136] Sanat K Sarkar and Chung-Kuei Chang. The simes method for multiple hypothesis testing with positively dependent test statistics. *Journal of the American Statistical Association*, 92(440):1601–1608, 1997.

[137] R Schirrmeister, Lukas Gemein, Katharina Eggensperger, Frank Hutter, and Tonio Ball. Deep learning with convolutional neural networks for decoding and visualization of eeg pathology. In *Signal Processing in Medicine and Biology Symposium (SPMB), 2017 IEEE*, pages 1–7. IEEE, 2017.

[138] R Schirrmeister, Lukas Gemein, Katharina Eggensperger, Frank Hutter, and Tonio Ball. Deep learning with convolutional neural networks for decoding and visualization of eeg pathology. In *Signal Processing in Medicine and Biology Symposium (SPMB), 2017 IEEE*, pages 1–7. IEEE, 2017.

[139] Robin Tibor Schirrmeister, Jost Tobias Springenberg, Lukas Dominique Josef Fiederer, Martin Glasstetter, Katharina Eggensperger, Michael Tangermann, Frank Hutter, Wolfram Burgard, and Tonio Ball. Deep learning with convolutional neural networks for eeg decoding and visualization. *Human brain mapping*, 38(11):5391–5420, 2017.

[140] Robin Tibor Schirrmeister, Jost Tobias Springenberg, Lukas Dominique Josef Fiederer, Martin Glasstetter, Katharina Eggensperger, Michael Tangermann, Frank Hutter, Wolfram Burgard, and Tonio Ball. Deep learning with convolutional neural networks for eeg decoding and visualization. *Human brain mapping*, 38(11):5391–5420, 2017.

[141] Hugo G Schnack, Neeltje EM Van Haren, Mireille Nieuwenhuis, Hilleke E Hulshoff Pol, Wiepke Cahn, and René S Kahn. Accelerated brain aging in schizophrenia: a longitudinal pattern recognition study. *American Journal of Psychiatry*, 173(6):607–616, 2016.

[142] Tamas Sefcsik, Dezso Nemeth, Karolina Janacsek, Ildiko Hoffmann, Jeff Scialabba, Peter Klivenyi, Geza Gergely Ambrus, Gabor Haden, and Laszlo Vecsei. The role of the putamen in cognitive functionsa case study. *Learning & Perception*, 1(2):215–227, 2009.

[143] Rachael D Seidler, Jessica A Bernard, Taritonye B Burutolu, Brett W Fling, Mark T Gordon, Joseph T Gwin, Youngbin Kwak, and David B Lipps. Motor control and aging: links to age-related brain structural, functional, and biochemical effects. *Neuroscience & Biobehavioral Reviews*, 34(5):721–733, 2010.

[144] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. *See https://arxiv. org/abs/1610.02391 v3*, 7(8), 2016.

[145] Meredith A Shafto, Lorraine K Tyler, Marie Dixon, Jason R Taylor, James B Rowe, Rhodri Cusack, Andrew J Calder, William D Marslen-Wilson, John Duncan, Tim Dalgleish, et al. The cambridge centre for ageing and neuroscience (cam-can) study protocol: a cross-sectional, lifespan, multidisciplinary examination of healthy cognitive ageing. *BMC neurology*, 14(1):204, 2014.

[146] S Murray Sherman and RW Guillery. The role of the thalamus in the flow of information to the cortex. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 357(1428):1695–1708, 2002.

[147] R John Simes. An improved bonferroni procedure for multiple tests of significance. *Biometrika*, 73(3):751–754, 1986.

[148] Amaya M Singh and W Richard Staines. The effects of acute aerobic exercise on the primary motor cortex. *Journal of Motor Behavior*, 47(4):328–339, 2015.

[149] Jasper D Sluimer, Wiesje M van der Flier, Giorgos B Karas, Ronald van Schijndel, Josephine Barnes, Richard G Boyes, Keith S Cover, Sílvia D Olabarriaga, Nick C Fox, Philip Scheltens, et al. Accelerating regional atrophy rates in the progression from normal aging to alzheimers disease. *European radiology*, 19(12):2826, 2009.

[150] Dustin E Stansbury, Thomas Naselaris, and Jack L Gallant. Natural scene statistics account for the representation of scene categories in human visual cortex. *Neuron*, 79(5):1025–1034, 2013.

[151] Jason Steffener, Christian Habeck, Deirdre O'Shea, Qolamreza Razlighi, Louis Bherer, and Yaakov Stern. Differences between chronological and brain age are related to education and self-reported physical activity. *Neurobiology of aging*, 40:138–144, 2016.

[152] Ingo Steinwart and Andreas Christmann. *Support vector machines*. Springer Science & Business Media, 2008.

[153] George E Stelmach, Paul C Amrhein, and Noreen L Goggin. Age differences in bimanual coordination. *Journal of Gerontology*, 43(1):P18–P23, 1988.

[154] Sebastian Stober, Avital Sternin, Adrian M Owen, and Jessica A Grahn. Deep feature learning for eeg recordings. *arXiv preprint arXiv:1511.04306*, 2015.

[155] Andreas B Storsve, Anders M Fjell, Christian K Tamnes, Lars T Westlye, Knut Overbye, Hilde W Aasland, and Kristine B Walhovd. Differential longitudinal changes in cortical thickness, surface area and volume across the adult life span: regions of accelerating and decelerating change. *Journal of Neuroscience*, 34(25):8488–8498, 2014.

[156] Edith V Sullivan, Margaret Rosenbloom, Kathleen L Serventi, and Adolf Pfefferbaum. Effects of age and sex on volumes of the thalamus, pons, and cortex. *Neurobiology of aging*, 25(2):185–192, 2004.

[157] Haoqi Sun, Luis Paixao, Jefferson T Oliva, Balaji Goparaju, Diego Z Carvalho, Kicky G van Leeuwen, Oluwaseun Akeju, Robert J Thomas, Sydney S Cash, Matt T Bianchi, et al. Brain age from the electroencephalogram of sleep. *Neurobiology of aging*, 74:112–120, 2019.

[158] Yousef Rezaei Tabar and Ugur Halici. A novel deep learning approach for classification of eeg motor imagery signals. *Journal of neural engineering*, 14(1):016003, 2016.

[159] François Tadel, Sylvain Baillet, John C Mosher, Dimitrios Pantazis, and Richard M Leahy. Brainstorm: a user-friendly application for meg/eeg analysis. *Computational intelligence and neuroscience*, 2011:8, 2011.

[160] Enzo Tagliazucchi and Helmut Laufs. Decoding wakefulness levels from typical fmri resting-state data reveals reliable drifts between wakefulness and sleep. *Neuron*, 82(3):695–708, 2014.

[161] Catherine Tallon-Baudry, Olivier Bertrand, Claude Delpuech, and Jacques Pernier. Stimulus specificity of phase-locked and non-phase-locked 40 hz visual responses in human. *Journal of Neuroscience*, 16(13):4240–4249, 1996.

[162] Jason R Taylor, Nitin Williams, Rhodri Cusack, Tibor Auer, Meredith A Shafto, Marie Dixon, Lorraine K Tyler, Richard N Henson, et al. The cambridge centre for ageing and neuroscience (cam-can) data repository: structural and functional mri, meg, and cognitive data from a cross-sectional adult lifespan sample. *Neuroimage*, 144:262–269, 2017.

[163] Prejaas Tewarie, Arjan Hillebrand, Bob W van Dijk, Cornelis J Stam, George C O'neill, Piet Van Mieghem, Jil M Meier, Mark W Woolrich, Peter G Morris, and Matthew J Brookes. Integrating cross-frequency and within band functional networks in resting-state meg: a multi-layer network approach. *Neuroimage*, 142:324–336, 2016.

[164] Pierre Thodoroff, Joelle Pineau, and Andrew Lim. Learning robust features using deep learning for automatic seizure detection. In *Machine Learning for Healthcare Conference*, pages 178–190, 2016.

[165] Bruce Thompson. Canonical correlation analysis. *Encyclopedia of statistics in behavioral science*, 2005.

[166] Marleen C Tjepkema-Cloostermans, Rafael CV de Carvalho, and Michel JAM van Putten. Deep learning for detection of focal epileptiform discharges from scalp eeg recordings. *Clinical neurophysiology*, 129(10):2191–2196, 2018.

[167] Susumu Tonegawa, Xu Liu, Steve Ramirez, and Roger Redondo. Memory engram cells have come of age. *Neuron*, 87(5):918–931, 2015.

[168] Luca Tonin, Tom Carlson, Robert Leeb, and José del R Millán. Brain-controlled telepresence robot by motor-disabled people. In *Engineering in Medicine and Biology Society, EMBC, 2011 Annual International Conference of the IEEE*, pages 4227–4230. IEEE, 2011.

[169] Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. Simultaneous deep transfer across domains and tasks. In *Computer Vision (ICCV), 2015 IEEE International Conference on*, pages 4068–4076. IEEE, 2015.

[170] Pál Vakli, Regina J Deák-Meszlényi, Petra Hermann, and Zoltán Vidnyánszky. Transfer learning improves resting-state functional connectivity pattern analysis using convolutional neural networks. *GigaScience*, 7(12):giy130, 2018.

[171] Martijn P Van Den Heuvel and Olaf Sporns. Rich-club organization of the human connectome. *Journal of Neuroscience*, 31(44):15775–15786, 2011.

[172] John Darrell Van Horn and Arthur W Toga. Human neuroimaging as a big data science. *Brain imaging and behavior*, 8(2):323–331, 2014.

[173] Barry D Van Veen and Kevin M Buckley. Beamforming: A versatile approach to spatial filtering. *IEEE assp magazine*, 5(2):4–24, 1988.

[174] Barry D Van Veen, Wim Van Drongelen, Moshe Yuchtman, and Akifumi Suzuki. Localization of brain electrical activity via linearly constrained minimum variance spatial filtering. *IEEE Transactions on biomedical engineering*, 44(9):867–880, 1997.

[175] Bernadette van Wijk, Peter J Beek, and Andreas Daffertshofer. Neural synchrony within the motor system: what have we learned so far? *Frontiers in human neuroscience*, 6:252, 2012.

[176] Gael Varoquaux and Bertrand Thirion. How machine learning is shaping cognitive neuroimaging. *GigaScience*, 3(1):28, 2014.

[177] Svyatoslav Vergun, Alok Deshpande, Timothy B Meier, Jie Song, Dana L Tudorascu, Veena A Nair, Vikas Singh, Bharat B Biswal, Mary Elizabeth Meyerand, Rasmus M Birn, et al. Characterizing functional connectivity differences in aging adults using machine learning on resting state fmri data. *Frontiers in computational neuroscience*, 7:38, 2013.

[178] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pages 1096–1103. ACM, 2008.

[179] Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, and Pierre-Antoine Manzagol. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of machine learning research*, 11(Dec):3371–3408, 2010.

[180] S Vichy N Vishwanathan, Nicol N Schraudolph, Risi Kondor, and Karsten M Borgwardt. Graph kernels. *Journal of Machine Learning Research*, 11(Apr):1201–1242, 2010.

[181] Eleni L Vlahou, Franka Thurm, Iris-Tatjana Kolassa, and Winfried Schlee. Resting-state slow wave power, healthy aging and cognitive performance. *Scientific reports*, 4:5101, 2014.

[182] Andrzej Wróbel et al. Beta activity: a carrier for visual attention. *Acta neurobiologiae experimentalis*, 60(2):247–260, 2000.

[183] Huijuan Yang, Siavash Sakhavi, Kai Keng Ang, and Cuntai Guan. On the use of convolutional neural networks and augmented csp features for multi-class

motor imagery of eeg signals classification. In *Engineering in Medicine and Biology Society (EMBC), 2015 37th Annual International Conference of the IEEE*, pages 2620–2623. IEEE, 2015.

[184] Roy Yannick, Banville Hubert, Albuquerque Isabela, Gramfort Alexandre, Faubert Jocelyn, et al. Deep learning-based electroencephalography analysis: a systematic review. *arXiv preprint arXiv:1901.05498*, 2019.

[185] Tal Yarkoni and Jacob Westfall. Choosing prediction over explanation in psychology: Lessons from machine learning. *Perspectives on Psychological Science*, 12(6):1100–1122, 2017.

[186] Ye Yuan, Guangxu Xun, Qiuling Suo, Kebin Jia, and Aidong Zhang. Wave2vec: Deep representation learning for clinical temporal data. *Neurocomputing*, 324:31–42, 2019.

[187] Andrew Zalesky, Alex Fornito, and Ed Bullmore. On the use of correlation as a measure of network connectivity. *Neuroimage*, 60(4):2096–2106, 2012.

[188] Xiang Zhang and Yann LeCun. Text understanding from scratch. *arXiv preprint arXiv:1502.01710*, 2015.

[189] Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text classification. In *Advances in neural information processing systems*, pages 649–657, 2015.

[190] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Computer Vision and Pattern Recognition (CVPR), 2016 IEEE Conference on*, pages 2921–2929. IEEE, 2016.

[191] P Zhuang, C Toro, J Grafman, P Manganotti, L Leocani, and M Hallett. Event-related desynchronization (erd) in the alpha frequency during development of implicit and explicit learning. *Electroencephalography and clinical neurophysiology*, 102(4):374–381, 1997.

[192] Gabriel Ziegler, Robert Dahnke, Lutz Jäncke, Rachel Aine Yotter, Arne May, and Christian Gaser. Brain structural trajectories over the adult lifespan. *Human brain mapping*, 33(10):2377–2389, 2012.

## KEY TO ABBREVIATIONS