

COMPUTATIONAL METHODS FOR ANALYSIS OF RESTING STATE FUNCTIONAL
CONNECTIVITY AND THEIR APPLICATION TO STUDY OF AGING

Harini Eavani

A DISSERTATION

in

Bioengineering

Presented to the Faculties of the University of Pennsylvania in Partial
Fulfillment of the Requirements for the Degree of Doctor of Philosophy

2016

Supervisor of Dissertation

Graduate Group Chair

Christos Davatzikos,
Professor of Radiology,
University of Pennsylvania

Jason A. Burdick,
Professor of Bioengineering,
University of Pennsylvania

DISSERTATION COMMITTEE

Paul Yushkevich, Professor of Radiology, University of Pennsylvania

Danielle Bassett, Assistant Professor of Bioengineering, University of Pennsylvania

Ted Satterthwaite, Assistant Professor of Psychiatry, University of Pennsylvania

Bruce Fischl, Professor of Radiology, Harvard Medical School

Stephen Smith, Professor of Biomedical Engineering, Oxford University

To my family.

Acknowledgements

I would like to thank my advisor, Christos, for his guidance and support during my tenure as a graduate student. I appreciate his immense knowledge and insight in a very broad range of topics, which are extremely valuable in the highly interdisciplinary field of computational neuroscience. He gave me considerable freedom to think and work independently, which was crucial to the synthesis of novel ideas. I have learned a great deal from his lucid writing style and presenting complex research ideas to any audience. Most of all, he invests a lot of time and resources in his protégés' careers, and is genuinely interested in seeing them grow into good scientists. This is a rare attribute that sets him apart from everyone else.

I would also like to offer my gratitude to my collaborators Ted Satterthwaite, Lori Beason-Held and Susan Resnick whose help shaped large parts of this thesis. My shared experiences with fellow graduate students Aoyan, Maria, Erdem and Ke made my time in the lab memorable. I would also like to thank my colleagues Nicolas, Aris, Spyros, Mohamed, Michael, Guray, Jimit and Mark who have helped me with my research projects time and time again.

Above all, this thesis is a testament to the unwavering support of my family and their ability to put my needs above all else. I hope I can return the favor, at least to some measure, some day.

ABSTRACT

COMPUTATIONAL METHODS FOR ANALYSIS OF RESTING STATE FUNCTIONAL CONNECTIVITY AND THEIR APPLICATION TO STUDY OF AGING

Harini Eavani

Christos Davatzikos

The functional organization of the brain and its variability over the life-span can be studied using resting state functional MRI (rsfMRI). It can be used to define a “macro-connectome” describing functional interactions in the brain at the scale of major brain regions, facilitating the description of large-scale functional systems and their change over the lifespan. The connectome typically consists of thousands of links between hundreds of brain regions, making subsequent group-level analyses difficult. Furthermore, existing methods for group-level analyses are not equipped to identify heterogeneity in patient or otherwise affected populations.

In this thesis, we incorporated recent advances in sparse representations for modeling spatial patterns of functional connectivity. We show that the resulting Sparse Connectivity Patterns (SCPs) are reproducible and capture major directions of variance in the data. Each SCP is associated with a scalar value that is proportional to the average connectivity within all the regions of that SCP. Thus, the SCP framework provides an interpretable basis for subsequent group-level analyses.

Traditional univariate approaches are limited in their ability to detect heterogeneity in diseased/aging populations in a two-group comparison framework. To address this issue, we developed a Mixture-Of-Experts (MOE) method that combines unsupervised modeling of mixtures of distributions with supervised learning of classifiers, allowing discovery of multiple disease/aging phenotypes and the affected individuals associated with each pattern.

We applied our methods to the Baltimore Longitudinal Study of Aging (BLSA), to find multiple advanced aging phenotypes. We built normative trajectories of functional and structural brain aging, which were used to identify individuals who seem resilient to aging, as well as individuals who show advanced signs of aging. Using MOE, we discovered five distinct patterns of advanced aging. Combined with neuro-cognitive data, we were able to further characterize one group as consisting of individuals with early-stage dementia. Another group had focal hippocampal atrophy, yet had higher levels of connectivity and somewhat higher cognitive performance, suggesting these individuals were recruiting their cognitive reserve to compensate for structural losses. These results demonstrate the utility of the developed methods, and pave the way for a broader understanding of the complexity of brain aging.

Contents

Acknowledgements	iii
Abstract	iv
Contents	vi
List of Tables	x
List of Figures	xi
1 Introduction	1
1.1 Overview	1
1.2 Aims of this thesis	2
1.3 Significance	5
1.4 Innovation	7
1.5 rsfMRI data	9
1.5.1 Datasets	9
1.5.2 Node definition	11
1.5.3 Computing the functional connectome	13
1.6 Organization of this thesis	13

2	Identifying Sparse Connectivity Patterns	14
2.1	Introduction	14
2.2	Literature review	15
2.3	The Sparse Learning approach	18
2.3.1	Model formulation	18
2.3.2	Optimization strategy	21
2.3.3	Model selection	21
2.3.4	SCP visualization	22
2.4	Experiments on simulated data	23
2.4.1	Generation of simulated data	23
2.4.2	Performance on simulated data	24
2.5	SCPs computed from rsfMRI data	29
2.5.1	Areal graph nodes	29
2.5.2	GraSP parcels	37
2.6	Extensions	38
2.6.1	Hierarchical Sparse Learning	38
2.6.2	Discriminative Sparse Learning	43
2.7	Discussion	48
3	Capturing Heterogeneity using Mixture of Experts	52
3.1	Introduction	52
3.2	Literature review	54
3.3	Mixture-of-Experts: Formulation and optimization	57
3.3.1	The expert model	58
3.3.2	The mixture model	59
3.3.3	The joint model	60

3.3.4	Optimization strategy	60
3.3.5	Prediction of test cases	61
3.3.6	Model selection	61
3.4	Experiments on simulated data	63
3.4.1	Simulated data	63
3.4.2	MOE choice of parameters and results	65
3.5	Discussion	68
4	Analyzing the effect of aging on functional connectivity	71
4.1	Introduction	71
4.2	Literature review	72
4.3	Effect of aging on functional connectivity	74
4.4	Functional brain aging trajectory	79
4.5	Resilient and advanced aging	80
4.6	Discovering heterogeneity in advanced aging	84
4.7	Discussion	89
5	Final Remarks and Future Work	94
5.1	Future work	95
5.2	Software	98
	Appendices	101
A	rsfMRI data pre-processing	102
A.1	Time-series pre-processing	102
A.2	Spatial Alignment	103
A.3	Controlling for motion confound	103
A.4	Global Signal Regression	104

B BLSA: Cognitive Data	106
Bibliography	108

List of Tables

2.1 Accuracy and AUC values for discriminative Sparse Learning, and other methods	48
3.1 Table comparing ten-fold cross-validation accuracy for MOE method vs. Gaussian-kernel SVM, for four 2-D simulated cases.	68

List of Figures

1.1	Figure illustrating computation of the functional connectome	2
1.2	Visualization of the functional connectome as a correlation matrix (left), an abstract graph representation (middle) and mapped to corresponding regions in the brain (right)	3
1.3	Left: Areal Graph nodes, as defined in Power et al. (2011). Right: Common group-parcellation of GM regions obtained by running GraSP (Honnorat et al., 2015) on BLSA data	12
2.1	Schematic 1 illustrating our method. Each subject specific correlation matrix Σ_n is approximated by a non-negative sum of sparse rank one matrices $\mathbf{b}_k \mathbf{b}_k^T$, or Sparse Connectivity Patterns (SCPs).	19
2.2	Schematic 2 illustrating our method. Each subject specific correlation matrix Σ_n is approximated by a product of three factors - the group-common matrix \mathbf{B} , subject specific matrix \mathbf{C}_n and \mathbf{B}^T	20
2.3	Visualization of SCPs. Opposing colors reflect anti-correlation between regions.	22

2.4	Simulated connectome design consisting of 50 nodes. (a) All nodes and edges. Nodes are common to all individuals, different sets of edges are inactive in different individuals. Edges in blue indicate anti-correlation. (b) Average correlation matrix across all individuals. (c) Correlation matrices of five randomly chosen individuals.	25
2.5	Cross-validation results for simulated data: Plots of the mean square error (Eqn. 2.3) vs. number of SCPs K (left), and sparsity level λ (right).	26
2.6	Performance on simulated data. The basis vectors identified by the Sparse Learning approach, InfoMap, Spatial and Temporal ICA, shown as node-weights (\mathbf{b}_k), compared against the ground-truth.	27
2.7	Performance on simulated data. Accuracy, measured using normalized inner-product.	28
2.8	Individual-level coefficients estimated by Sparse Learning, projected down to 2-D space, shown along with the ground truth.	29
2.9	Cross-validation results for PNC data: Plots of the mean square error (Eqn. 2.3) vs. number of SCPs K (left), and sparsity level λ (right).	30
2.10	Ten SCPs computed in the Areal Graph node space.	31
2.11	Figure illustrating the heterogeneity of the data sample captured by SCPs. The color indicates the extent to which each SCP is present in a individual.	33
2.12	Reproducibility, data-fit, spatial overlap and temporal correlation measured for SCPs obtained from Areal Graph nodes.	34

2.13 SCPs 1, 2, 5, and 9 computed from the PNC dataset(left), HCP+fBIRN dataset (middle) and BLSA dataset (right) using Areal Graph nodes. The inner product value (computed in the Areal Graph node space) for each comparison is also shown.	38
2.14 SCPs 1-5 generated from the PNC data (left) and BLSA data (right) using GraSP parcels.	39
2.15 SCPs 6-10 generated from the PNC data (left) and BLSA data (right) using GraSP parcels. Note that SCPs 9 and 10 do not match between the two datasets.	40
2.16 A primary SCP and some of its associated secondary SCPs computed from the PNC dataset using GraSP parcels.	41
2.17 Split-sample error and reproducibility for secondary SCPs computed from the PNC dataset.	42
2.18 SVM weight vector for l_2 -SVM (left) and l_1 -SVM (right) for discriminating children from young adults using the PNC dataset.	43
2.19 Schematic illustrating the discriminative Sparse Learning framework. Panel to the left describes the SCP identification term, which factorizes connectivity matrices Σ_n of each individual n into a set of common SCPs $\mathbf{B} = [\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_K]$ and its associated coefficients. Panel to the right illustrates a linear SVM, which uses the total absolute connectivity values $diag(\mathbf{B}^T \Sigma_n \mathbf{B})$ of all SCPs as input features to classify two groups, resulting in the hyperplane w	45

2.20	Top two SCPs that contribute to classification, based on their hyper-plane weight. Corresponding graphs (bottom) plot total connectivity within SCP for each individual vs. individual age. Uni-variate p-value scores comparing total connectivity between two groups are also shown.	47
3.1	An illustration that shows heterogeneity in the affected group, relative to a reference group. Non-linear classifiers can implicitly model non-linearity but estimating the boundary (dashed lines) in high dimensional spaces is difficult. Mixture-of-Experts can approximate the non-linear curve with a piece-wise linear boundary (red and green lines) and find subgroups associated with each line (red and green points).	55
3.2	Four 2-D simulated cases used to evaluate the performance of the method. Hyper-planes and sub-groups obtained using MOE are also shown using a different color for each hyper-plane and associated subgroup.	64
3.3	Variation in four performance measures for $K = \{1, 2, \dots, 5\}$, for each of the four simulated cases. From top-left, clockwise: Accuracy, Maximum inner-product, cluster separation and cluster reproducibility. Results from each case is plotted in a different color.	66
3.4	Variation in four performance measures for $K = 3, C, \lambda = \{2^{-3}, 2^{-2}, \dots, 2^{10}\}$, for Case 3. From top-left, clockwise: Accuracy, Maximum inner-product, cluster separation and cluster reproducibility.	67
4.1	Age distribution for the subjects used in this study. For each age bin, counts for male and female participants are shown in different colors.	74

4.2	Scatter plot showing the relationship between age and motion (measured using MRD), which is statistically insignificant. The color indicates the number of time-points that were removed during the scrubbing procedure.	74
4.3	(Above) Connectivity within the DM vs. DA SCP is significantly decreased with age. (Below) Connectivity within the medial temporal areas is significantly increased with age	75
4.4	Connectivity within the visual vs. DM areas (above) and motor areas (below) is significantly decreased with age.	76
4.5	Eight secondary SCPs of the motor SCP, correlation between corresponding coefficients and age, and associated p-value.	78
4.6	Brain Age Index of each individual plotted against their age. The Mean Absolute Error (MAE) of prediction, as well as the correlation coefficient between the BAI and age is provided.	79
4.7	Primary and associated secondary SCPs that were found to be significant after permutation testing. Although not all primary SCPs are significant (faded images are not), they are shown for reference. . . .	81
4.8	Functional and structural Brain Age Indices (BAIs) plotted against each individual's age. Functional BAI was computed from SCP coefficients and local functional coherence measures; structural BAI was computed from average GM density values. The aging trajectories (solid lines) show the expected brain age. The BAI Residual for each individual in each modality can be computed by subtracting their expected brain age from their BAI.	82

4.9	The DM vs. DA SCP alone shows significantly decreased connectivity in the advanced agers relative to the resilient agers after correcting for age.	82
4.10	Univariate differences in local functional coherence and GM density between resilient and advanced agers.	83
4.11	Summary of significant differences in GM density, functional coherence and connectivity for all five groups of advanced agers relative to resilient agers. As before, changes to GM density and coherence at the level of ROIs are shown using p-value maps overlaid on a template image (first two columns on right). SCPs whose coefficients were significantly different between groups are shown in the last column. The color overlay for SCPs indicates patterns of correlated (or anti-correlated) regions.	85
4.12	Functional and structural BAI Residual (= Expected Brain age - BAI) plotted in a two-dimensional grid. The group membership of each individual is reflected in the color, resilient agers are shown as white points.	87
4.13	Differences in cognitive performance between advanced group 1 and resilient agers, for eight cognitive domains. Mean and 95% confidence intervals of the difference estimates are shown as errorbars. CVLT: California Verbal Learning Task. BVRT: Benton Visual Retention Test. CRDROT: Card Rotation Test. FLULET: Letter Fluency. FLUCAT: Category Fluency. TRATS: Trail Making Test Part A. TRBTS: Trail Making Test Part B. DSST: Digit Symbol Test.	88
5.1	Command-line interface of SCPLearn	99

5.2 Command-line interface of MOE 100

A.1 SCPs computed without global signal regression, from the PNC dataset, using the Areal Graph nodes. Only SCP 1 has a significant area of negative correlation - the Dorsal Attention vs. Default mode anti-correlation pattern. The other nine SCPs showed only positive correlation. Of these, the familiar patterns were the sensori-motor (8, 9 and 10), visual (4, 7) and cingulo-opercular networks (2). 105

Chapter 1

Introduction

1.1 Overview

Magnetic Resonance Imaging (MRI) provides a non-invasive, in-vivo method of studying the human brain. In the clinic, these three-dimensional images of the brain are used to identify gross abnormalities visible to the naked eye (lesions, tumours, abscesses). In research applications, MRI is acquired from many research participants, in order to understand the subtle effects of development, aging or illness on brain function and structure which cannot be observed by visual inspection.

Resting state functional MRI (rsfMRI) (Biswal et al., 1995; Raichle et al., 2001) is acquired when the individual is at rest inside the scanner. It measures the blood oxygenation level in different regions of the brain. Higher oxygen levels imply higher neural activity, therefore fMRI is assumed to be a surrogate measure of brain function. Multiple 3-D images are acquired, typically once every couple of seconds. In this manner, we can obtain a time-series of varying blood oxygenation level at every location in the brain. This “spontaneous” brain activity can be used to measure the amount of co-activation, or *functional connectivity*, between any two regions

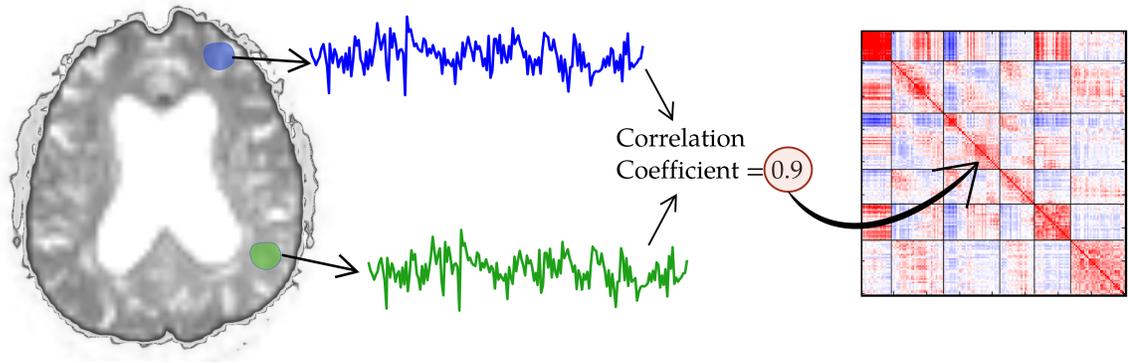


Figure 1.1: Figure illustrating computation of the functional connectome

in the brain by quantifying the amount of synchrony between corresponding activity. This synchrony is popularly quantified by computing the Pearson correlation coefficient between the corresponding time-series, as illustrated in Figure 1.1. In this manner, one can build a whole-brain *functional connectome* that captures a time-averaged picture of the functional interactions between all brain regions. Figure 1.2 shows the various ways in which the functional connectome is visually represented.

A growing number of projects (Satterthwaite et al., 2014; Van Essen et al., 2012; Shock et al., 1984) are acquiring rsfMRI data for a large number of individuals with the aim of investigating normal inter-subject variability and variation over the lifespan. To keep pace with the scale of such projects, there is an urgent need to develop novel methodology that can capture complex, whole-brain normative patterns of age-related changes in a data-driven manner, while ensuring interpretability.

1.2 Aims of this thesis

The functional connectome of each individual is very high dimensional, however, this data is acquired on far fewer number of subjects. This entails the need for a di-

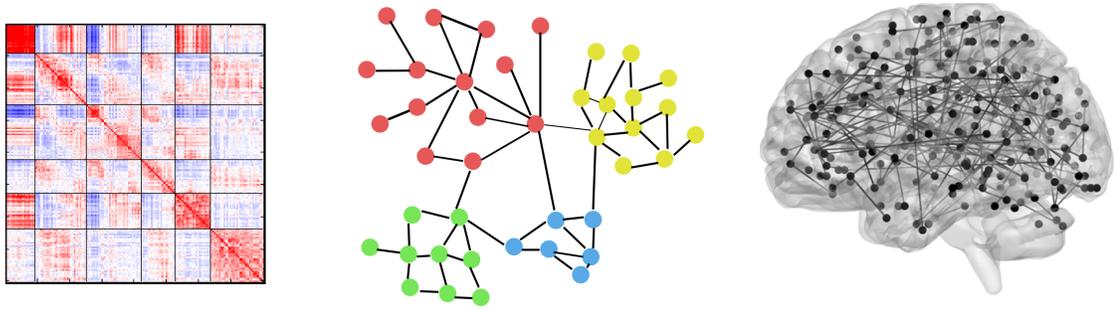


Figure 1.2: Visualization of the functional connectome as a correlation matrix (left), an abstract graph representation (middle) and mapped to corresponding regions in the brain (right)

dimensionality reduction method that is based on functional connectivity data. Existing approaches for dimensionality reduction divide the brain into spatially smaller patterns, varyingly referred to as components, sub-graphs or sub-networks. Some of these existing approaches use the rsfMRI time-series as input, which is not ideal for connectome data. Others assume spatial or temporal separation of the patterns, which is not suited to describe functional systems in the brain. As we describe in the following chapter, we exploited recent advances in the mathematics of sparse modeling to develop a methodological framework that decomposes the complex functional connectome into parsimonious *Sparse Connectivity Patterns* (SCPs). The description of our *Sparse Learning* method and resulting SCPs is provided in Chapter 2 of this thesis.

In rsfMRI studies, univariate statistical approaches or multi-variate classification methods are often employed to identify connections that can predict diseased state of an individual or are strongly correlated with age. Such linear models inherently assume that there is a single pattern of change that linearly scales with disease severity or age in all individuals. However many of the conditions under study manifest in a heterogeneous manner across individuals, forming a spectrum that is

composed of a variety of external symptoms and cognitive changes. In Chapter 3 of this thesis, we present a *Mixture-Of-Experts* (MOE) method that explicitly models and captures heterogeneous patterns of change in the affected group relative to a reference group of controls.

Increased life expectancy has resulted in a growing proportion of older people across the globe. This has led to a greater prevalence of aging related mental illnesses. In order to better understand functional disruptions of such disorders, it is crucial to gain insight into how the functional organization in the brain changes during the normal aging process. Although, on average, cognitive abilities decrease with age, some individuals age well, while others decline faster. Some of these individuals could be at a greater risk for developing clinical pathology later in life. Therefore, in addition to understanding the effect of aging on function, it is important to identify these high-risk individuals, as well as potentially heterogeneous patterns of aging. Towards this goal, we applied the methodological approaches developed and tested in chapters 2 and 3 to rsfMRI data acquired as a part of the *Baltimore Longitudinal Study of Aging* (BLSA). Multiple patterns of aging related changes in connectivity are reported and discussed in Chapter 4.

These aims are summarized below:

1. Develop and test a method based on sparse representations for the identification of resting-state connectivity networks
2. Develop and test a mixture of experts method for identification of heterogeneity in affected populations
3. Apply the developed methods to the BLSA dataset (normal adults of age 50-96 years) to find heterogeneous effects of aging on brain structure and function

1.3 Significance

In order to make rsfMRI applicable to large heterogeneous populations, we require automated data-driven analysis approaches that can capture the complexity of functional connectivity patterns within the data. We summarize the significance of our work vis-a-vis existing methods below, and discuss them in greater detail in individual chapters.

In Aim one, we proposed a sparsity-based dimensionality reduction algorithm that is complementary to existing graph-based or time-series based methods. The significance of this aim is described below:

1. Seed-based correlation approaches require prior knowledge of a stable seed location and do not directly quantify inter-subject variability.
2. Graph-based approaches (Power et al., 2011) that identify networks based on correlation strengths limit their analysis to strong positive correlations based on an ad-hoc thresholding operation, ignoring the *positive-semi-definite* property of the correlation matrix. Negative and weak edges correlations could be informative, especially if considered collectively as part of a distributed pattern.
3. Graph-based methods are limited to identifying patterns in a group-averaged manner (Power et al., 2011), thereby losing valuable information about inter-individual variability.
4. Time-series based approaches, such as spatial (Beckmann et al., 2005) and temporal ICA (Smith et al., 2012) have been successful in producing stable and reproducible spatial components (especially spatial ICA). However (1) these approaches are not directly applicable to connectome data and (2) do

not directly encode inter-individual variability in connectivity. This entails novel methods that can directly explain the variance in connectivity data.

Aim two proposed a mixture of experts framework to address the drawbacks of linear and non-linear supervised learning approaches, which are enumerated below:

4. Multi-variate classification methods are needed to find subtle, distributed patterns of aging related changes in the connectome. However, linear methods assume a single pattern of aging or disease-related change, which is unrealistic.
5. While kernel based methods can be used to model non-linear effects, the pattern of change itself is not explicitly computed; kernel approaches do not provide any information about the features that contribute to classification. This motivates the need for developing methods that model heterogeneity in the data explicitly.

As a part of the third aim, we applied the methods developed in Aims 1 and 2 to understand the effects of aging on functional connectivity. Previous studies that investigated aging effects used traditional analysis approaches, whose major limitations are described below:

6. Past studies investigating the effects of aging used mass-univariate approaches to identify functional connections that significantly correlated with age. However, traditional analysis approaches make it difficult to gauge the extent and severity of aging for each individual.
7. In prior studies of aging, small sample sizes limited researchers to assume aging to be a homogeneous process that affects all older individuals uniformly.

Using mass-univariate methods, an overall decrease in functional connectivity, especially in the default mode and motor regions, has been found consistently. However, a single pattern of change does not explain the wide variation in cognitive abilities seen among the elderly. These studies did not stratify older populations in terms of their patterns of change in brain structure and function.

1.4 Innovation

The major contributions of this thesis are summarized as follows:

1. The developed Sparse Learning method for the analysis of correlation matrices does not require the removal of weak and negative correlation values. This model represents the data as a combination of Sparse Connectivity Patterns, while retaining the positive semi-definite nature of correlation matrices in the representation. The P.S.D assumption constrains the degrees of freedom in the model resulting in more stable, reliable networks.
2. We incorporated inter-individual variability in the strength of SCPs in our modeling strategy. Such a model allocates a scalar value per pattern in each individual that reflects the average connectivity of all the regions within that particular pattern. This is a clear advancement over seed-based correlation approaches, which do not consider individual level variations.
3. Sparse decompositions are being greatly favored in the signal processing community, as evidence of parsimony in nature as well as in the functioning of the brain is abundant. We used spatial sparsity as the driving assumption for the definition of underlying sub-networks. We incorporate recent advances made

in sparse representations of covariance matrices (Sivalingam et al., 2011; Sra and Cherian, 2011; Soufiani and Airoidi, 2012) in our modeling strategy.

4. The MOE framework for identification of heterogeneous patterns of change explicitly models two crucial pieces of information that we would like to discover in the data - (1) multiple patterns that differentiate two groups and (2) sub-groups of individuals associated with each pattern of change. The MOE method has a complexity level that is in-between that of linear and non-linear methods. Thus, it can capture more information than a linear method, while at the same time is not as complex as the kernel-based method, and by design, has better interpretability.
5. With data from the BLSA study, we used multi-variate methods to pool information across all connections to predict each individual's age. Using multi-variate techniques, we identified those connections that highly contribute to the prediction, similar to uni-variate techniques. More importantly, we built *aging trajectories* of functional connectivity, similar to growth charts that are used for children at the pediatrician's office. Each individual is assigned a *Brain Aging Index* (BAI), that measures the functional age of that individual. We then identified a subset of individuals whose BAI is higher than their expected BAI. These *advanced agers* could be at a higher risk for developing signs of dementia as they grow older.
6. We applied the MOE method to the advanced aging individuals (with *resilient agers* as a reference) and identified sub-groups of advanced agers with heterogeneous patterns of functional and structural change. To the best of our knowledge, we are the first to identify advanced aging patterns and understand the heterogeneity present in advanced aging from a purely data-driven

perspective.

1.5 rsfMRI data

1.5.1 Datasets

MRI data used in this thesis was acquired as a part of two major studies. Details of the study participants and MR acquisition parameters are described in the following section. Appendix A details the pre-processing pipeline that was used to prepare rsfMRI data prior to generation of functional connectivity matrices.

1.5.1.1 Baltimore Longitudinal Study of Aging (BLSA)

The Baltimore Longitudinal Study on Aging (BLSA) (Shock et al., 1984) is the longest running study on aging at the National Institute on Aging (NIA). It is a longitudinal study that acquires comprehensive cognitive assessments and MR brain imaging measurements (and many others for physical health), from research volunteers.

A subset of the data from this study was used for methodology testing in Chapters 2 and 3. A more extensive analysis of this data is described in Chapter 4.

Participants

As of February 1, 2015, BLSA had acquired 780 rsfMRI scans from 567 participants. For this thesis, we considered only baseline (first visit) scans of participants selected based on the following criteria: (1) participant was older than 50 years at time of scan (2) had low head motion during the acquisition, measured using Mean Relative Displacement - $MRD < 0.2\text{mm}$ (Satterthwaite et al., 2012) and (3) participants

were “normal” at time of scan; i.e., did not meet criteria for onset of Mild Cognitive Impairment or Alzheimer’s disease. This resulted in 400 subjects in the age range 50 – 96 years (72.5 ± 9.4 years).

MRI Acquisition

Images were acquired at the NIA clinical research facility on a Philips Achieva 3T MRI scanner, with an in plane resolution of 3×3 mm, slice thickness of 4 mm, TR/TE=2000/30s and total scan duration of 6 minutes.

1.5.1.2 Philadelphia Neuro-developmental Cohort (PNC)

The Philadelphia Neuro-developmental Cohort (PNC) is a large scale study to understand the normal and abnormal developmental processes in the human brain. In addition to neuroimaging, participants also received cognitive and psychiatric assessments.

Subsets of this dataset were used for methodological testing in Chapter 2.

Participants

MRI data was acquired on 1445 participants as a part of this study. Of these, rsfMRI was acquired on 1275 participants. We excluded 426 subjects based on bad data quality or abnormal cognitive and/or psychiatric assessments, resulting in a final tally of 849 rsfMRI scans.

MRI Acquisition

All data were acquired on the same scanner (Siemens Tim Trio 3 Tesla; 32 channel head coil) using the same imaging sequences. Blood oxygen level dependent

(BOLD) fMRI was acquired with the following parameters: 124 volumes, TR 3000 ms, TE 32 ms, flip angle 90° , FOV 192x192 mm, matrix 64X64, slice thickness/gap 3mm/0mm, effective voxel resolution 3.0x3.0x3.0mm. During the resting-state scan, a fixation cross was displayed as images were acquired. Subjects were instructed to stay awake, keep their eyes open, fixate on the displayed crosshair, and remain still.

1.5.2 Node definition

A crucial aspect of SCP estimation is node definition. The high spatial dimensionality of fMRI data makes voxel-wise correlation matrices computationally infeasible for many approaches, hence most studies resort to dimensionality reduction, often through the use of apriori defined ROIs or through functional parcellation schemes. In this thesis, we used two sets of node definitions; a set of ROIs defined from a meta-analysis of fMRI studies (Power et al., 2011), and a data-driven rsfMRI parcellation defined by running the GraSP method (Honnorat et al., 2015) on each dataset. Details are given below.

1.5.2.1 Spherical ROI Definition (Areal Graph)

We used 264 node locations defined in (Power et al., 2011) (“Areal Graph”) for some of our methodology validation experiments in Chapter 2. These nodes were defined exclusively based on fMRI. Of these nodes, 151 were non-overlapping 10mm diameter spheres identified based on a meta-analysis of task-fMRI based studies (Dosenbach et al., 2006). The remaining 193 were cortical patches obtained by functional connectivity mapping using resting state fMRI (Cohen et al., 2008). Figure 1.3 (left) shows the spatial extent of Areal Graph nodes in one axial plane.

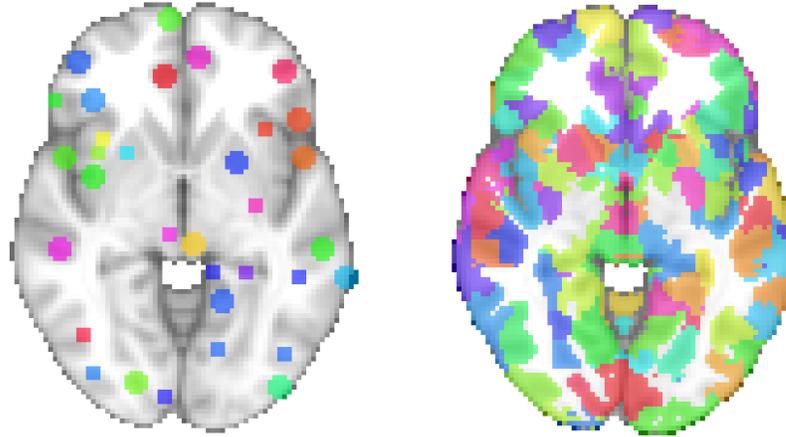


Figure 1.3: Left: Areal Graph nodes, as defined in Power et al. (2011). Right: Common group-parcellation of GM regions obtained by running GraSP (Honnorat et al., 2015) on BLSA data

1.5.2.2 Data-driven Parcellation (GraSP Parcels)

Using the Areal Graph described in the previous section limits the analysis to well-established functional foci of activity. Alternately, one could use widely available anatomical cortical atlases (Desikan et al., 2006; Van Essen, 2005) that delineate region boundaries based on cortical macro-structure, i.e., sulci and gyri, and average data within anatomical regions. While this approach incorporates information from all regions of the cortex, averaging observations within these anatomical boundaries might cause averaging across functional boundaries, which is not ideal. It is therefore important that we define a parcellation based on the same rsfMRI dataset that is being analyzed. Therefore, we used GRaSP (Honnorat et al., 2015), which is a data-driven method used for parcellating the grey matter based on local functional connectivity of the voxels. Run independently on the PNC and the BLSA datasets, GraSP provided a study-based parcellation of 583 and 744 parcels respectively. An example of the parcellation result on the BLSA data with 583 parcels can be seen in Figure 1.3 (right), in one axial plane.

1.5.3 Computing the functional connectome

After the nodes are defined in a common anatomical space, for each individual, the average time-series \mathbf{x}_i within each node i is computed. In this thesis, we define functional connectivity r_{ij} between any pair of nodes to be the Pearson's correlation coefficient between the corresponding time-series, as defined below:

$$r_{ij} = \frac{(\mathbf{x}_i - \bar{\mathbf{x}}_i)^T (\mathbf{x}_j - \bar{\mathbf{x}}_j)}{\|\mathbf{x}_i - \bar{\mathbf{x}}_i\|_2 \|\mathbf{x}_j - \bar{\mathbf{x}}_j\|_2} \quad (1.1)$$

where \bar{x} denotes mean of the time-series.

The functional connectivity value r_{ij} lies between -1 and 1 . If two nodes have a negative correlation value, they are said to be *anti-correlated*. The functional connectivity between all pairs of nodes can be represented in a correlation matrix Σ_i for each individual i . This matrix is symmetric and *positive-semi-definite*.

1.6 Organization of this thesis

The two main methodological contributions of this thesis are described in chapters 2 and 3. Chapter 2 details the Sparse Learning method, its performance on simulated data, and SCPs found using real data. In Chapter 3, we describe the MOE approach and its performance using synthetic data. Chapter 4 describes the main application of this thesis to aging, with detailed analysis of the BLSA dataset, including identification of heterogeneous sub-types of advanced aging. Chapter 5 summarizes all the contributions of this thesis, and discusses future work.

Chapter 2

Identifying Sparse Connectivity Patterns

2.1 Introduction

The functional connectome provides a rich source of information on the functional organization of the human brain. It has also been demonstrated that functional connectivity is altered in psychiatric and neurodegenerative illnesses such as schizophrenia (Venkataraman et al., 2012) and Alzheimer's (Greicius et al., 2004). An accurate description of the brain's functional connectome, and its variability across individuals, is a critical prerequisite for understanding both normal brain function and its aberrations in disease.

A multi-variate dimensionality reduction algorithm can be used to address both issues - (a) it can provide a parsimonious representation of the functional connectome in terms of a set of spatial patterns and (b) differential presence of these patterns explains inter-subject variability, i.e., each individual's connectome can be reduced to a set of scalar values with these patterns as the common reference.

We exploit recent advances in the mathematics of sparse modeling to develop a methodological framework aiming to understand complex resting-state fMRI connectivity data. By favoring patterns that explain the data via a relatively small number of participating brain regions, we obtain a parsimonious representation of brain function in terms of multiple “Sparse Connectivity Patterns” (SCPs), such that differential presence of these SCPs explains inter-subject variability.

2.2 Literature review

rs-fMRI connectivity has been used to delineate major functional brain systems (Biswal et al., 1995; Fox et al., 2006; Vincent et al., 2008), based on prior knowledge of a “seed” region of interest. Given a seed region, the average time-series of the seed is correlated with the time-series of every other voxel in the brain. The resulting significant correlations form a *correlation map* with respect to that seed. For example, using the Posterior Cingulate Cortex (PCC) as the seed region results in a correlation map which delineates regions belonging to the default mode, as well as regions of the Dorsal Attention (DA) system that are known to be anti-correlated with it. However, this method cannot be applied in a data-driven manner as it requires prior knowledge of a stable seed location. It does not directly quantify inter-subject variability.

Data-driven approaches have been developed to identify such functional patterns. They can be divided into two major categories - (a) graph partitioning approaches, which use correlation matrices as input, and (b) time-series based approaches, which are applied directly to time-series data. Although time-series approaches do not directly describe connectome data, we describe them here, as prior studies have used patterns computed from time-series for subsequent analysis of

connectivity information.

Graph partitioning approaches, such as InfoMap (Rosvall and Bergstrom, 2008) assume that any region of the brain can belong to only one brain network. This approach was applied to rsfMRI in Power et al. (2011). Retaining only high positive correlation values, the authors identified multiple spatially separated networks, or “sub-graphs”, whose regions are strongly correlated, on average, across individuals. However, InfoMap limits its analysis to strong positive correlations, while removing negative and weak edges from the graph that could be informative (Fox et al., 2005; Keller et al., 2013). In Yeo et al. (2011), a node’s functional connectivity to all other nodes was used as input to a clustering algorithm. Similar to InfoMap, this approach also provided spatially segregated networks.

Alternative approaches addressing some of these issues have been proposed in other fields. The notion of “link communities” introduced in Ahn et al. (2010) is elegantly able to handle overlaps by assigning unique membership to edges rather than nodes, naturally resulting in multiple assignments per node. Approaches like correlation clustering (Bansal et al., 2004) and the Potts model based approach proposed in Traag and Bruggeman (2009) are partitioning approaches which allow negative values. Since most of these methods are used to analyze social networks, they interpret negative edge links as repulsion, and hence attempt to assign negatively connected groups to different communities. While this may be appropriate for social networks, in resting state fMRI, highly negative edges imply strong anti-correlation - meaning that despite opposing phase information, these nodes express the same information, since they are strongly statistically dependent. Allocating anti-correlated regions to the same network can provide interesting new insights into the functional organization of the brain.

Alternately, time-series based approaches such as spatial or temporal Indepen-

dent Component Analysis (ICA) are applied directly to the time-series to obtain a set of basis, where each vector is a set of weights, one for each node. ICA incorporates higher-order moments to reveal patterns that are maximally independent. However, time-series based approaches do not directly explain the variance in connectome data. Spatial ICA is widely applied to rsfMRI data to obtain spatially independent components, commonly referred to as “Intrinsic Connectivity Networks (ICNs)” (Calhoun et al., 2003). In practice, ICNs found using spatial ICA are usually non-overlapping. To address this issue of non-overlap, the study by Smith et al. (2012) applied temporal ICA to rsfMRI data and found multiple functional brain networks, or “Temporal Functional Modes (TFMs)”. Although this is a significant advancement, these networks have been identified on the basis of independent temporal behavior, i.e., lacking between-network interactions, which is contrary to the notion that brain systems often act in concert during complex cognitive functioning.

A major disadvantage of connectivity based approaches is their inability to directly quantify inter-individual variability in functional connectivity, requiring additional post-processing and analysis. An important source of variation across individuals is the average strength of patterns. This could possibly be due to the extent to which (how much and for how long) that functional unit is recruited in each individual, or as an indicator of life span changes or disease state. There are studies that have found strong relationships between the clinical variable of interest and the strength of such intrinsic rsfMRI networks (von dem Hagen et al., 2012; Mayer et al., 2011; He et al., 2007a). It is important to build a model that can capture these inter-individual variations.

2.3 The Sparse Learning approach

Motivated by models of neuronal activity (Vinje and Gallant, 2000), we propose the use of spatial sparsity to drive identification of patterns of connectivity. In a neuronal sparse coding system, information is encoded by a small number of synchronous neurons that are selective to a particular property of the stimulus (e.g. edges of a particular orientation within a visual stimulus). Multiple such spatial patterns of neurons constitute a sparse neural basis which acts in concert in response to the stimulus. A nearly infinite number of stimuli can be parsimoniously encoded by varying the proportion in which these patterns are combined.

Extending this idea to rsfMRI, we assume that the observed spontaneous activity arises from the concerted activity of multiple “Sparse Connectivity Patterns (SCPs)” that encode system-level function, similar to sparse codes that are present at the level of neurons. Each SCP consists of a small set of spatially distributed, functionally synchronous brain regions, forming a basic pattern of co-activation. These SCPs capture the range of resting functional connectivity patterns in the brain, although they do not necessarily need to be present in each individual or subsets of individuals. Using spatial sparsity as a constraint, we learn the identity of these SCPs and the strength of their presence in each individual, revealing the variability in the dataset.

2.3.1 Model formulation

The objective of our method is to find SCPs consisting of functionally synchronous regions whose connectivity values co-vary across individuals, and are smaller than the whole-brain connectome. The information content within any one of these SCPs is relatively low, since all the nodes within an SCP are correlated, and express the

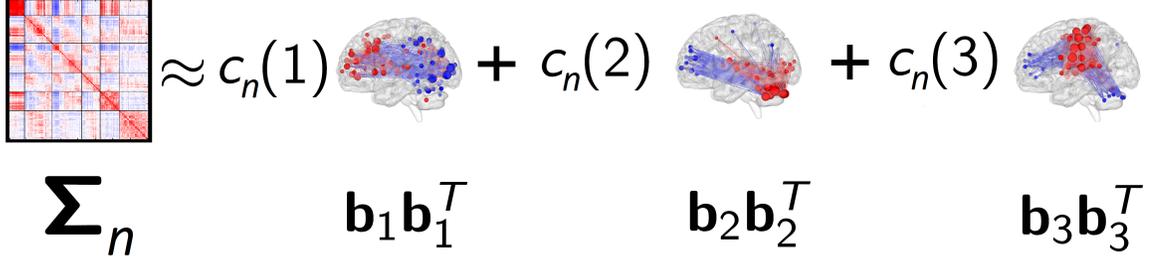


Figure 2.1: Schematic 1 illustrating our method. Each subject specific correlation matrix Σ_n is approximated by a non-negative sum of sparse rank one matrices $\mathbf{b}_k \mathbf{b}_k^T$, or Sparse Connectivity Patterns (SCPs).

same information. Furthermore, we assume that if a set of ROIs belong to such a pattern, then, in a set of normal subjects, inter subject variability is introduced by the extent to which each system is recruited in a subject. Thus, nodes are assigned to an SCP if the strength of the edges between them co-vary across subjects. To summarize, SCPs would have the following properties - (1) large number of edges with zero weights, or sparsity (2) low information content - or rank deficiency and (3) an associated scalar SCP *coefficient*, whose value is variable across individuals. Our method takes as input correlation matrices and finds SCPs that satisfy all properties.

A schematic diagram illustrating our method is shown in Figure 2.1. The input to our method is size $P \times P$ correlation matrices $\Sigma_n \succeq 0$, one for each subject n , $n = 1, 2, \dots, N$. We would like to find smaller SCPs common to all the subjects, such that a non-negative combination of these SCPs approximates the full-correlation matrix Σ_n , for each subject n . Each of these K SCPs can be represented by a vector of node-weights \mathbf{b}_k , where $-1 \preceq \mathbf{b}_k \preceq 1$, $\mathbf{b}_k \in \mathbf{R}^P$. Each vector \mathbf{b}_k reflects the membership of the nodes to the sub-network k . If $|\mathbf{b}_k(i)| > 0$, node i belongs to the sub-network k , and if $\mathbf{b}_k(i) = 0$ it does not. If two nodes in \mathbf{b}_k have the same sign, then they are positively correlated and opposing sign reflects anti-correlation. Thus, the rank-one matrix $\mathbf{b}_k \mathbf{b}_k^T$ reflects the correlation behavior of SCP k . In addition, we

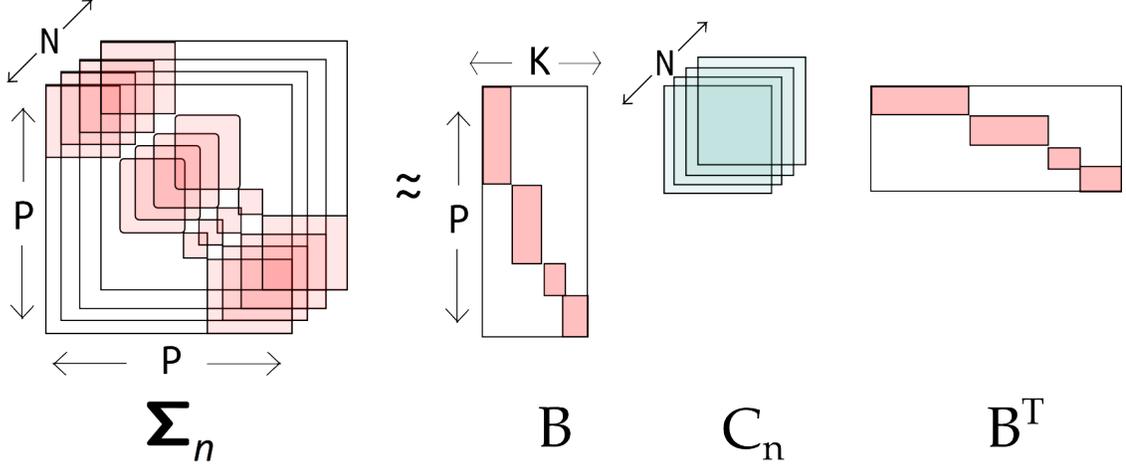


Figure 2.2: Schematic 2 illustrating our method. Each subject specific correlation matrix Σ_n is approximated by a product of three factors - the group-common matrix \mathbf{B} , subject specific matrix \mathbf{C}_n and \mathbf{B}^T .

constrain these SCPs to be much smaller than the whole-brain network by restricting the l_1 -norm of \mathbf{b}_k to not exceed a constant value λ .

We would like to approximate each matrix $\{\Sigma_n\}_{n=1}^N$ by a non-negative combination of SCPs $\mathbf{B} = [\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_K]$, as shown in Figure 2.2. Thus, we want

$$\Sigma_n \approx \sum_{k=1}^K c_n(k) \mathbf{b}_k \mathbf{b}_k^T = \mathbf{B} \text{diag}(\mathbf{c}_n) \mathbf{B}^T \triangleq \hat{\Sigma}_n$$

$$\|\mathbf{b}_k\|_1 \leq \lambda, \quad -1 \leq \mathbf{b}_k(i) \leq 1, \quad \mathbf{c}_n \geq 0 \quad (2.1)$$

where $\text{diag}(\mathbf{c}_n)$ denotes a diagonal matrix with values $\mathbf{c}_n \in \mathbf{R}_+^K$ along the diagonal. Thus, each subject n is associated with a vector of K subject-specific measures \mathbf{c}_n which are non-negative and reflect the relative contribution of each SCP to the whole-brain functional network in the n -th subject.

We quantify the approximation between Σ_n and $\hat{\Sigma}_n$ using the frobenius norm. Note that there is an ambiguity in amplitude between the two factors - if \mathbf{b}_k and $c_n(k)$ is a solution, $\alpha \mathbf{b}_k$ and $c_n(k)/\alpha^2$ is also a solution for any scalar α . To prevent

this, we fix the maximum value in each SCP to unity; i.e., $\max_i |\mathbf{b}_k(i)| = 1$.

Bringing the objective and the constraints together, we have the following optimization problem w.r.t the unknowns \mathbf{B} and $\mathbf{C} = [\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_N]$:

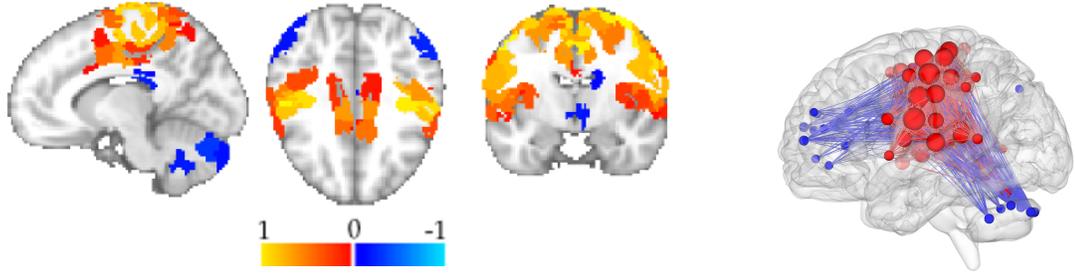
$$\begin{aligned}
& \underset{\mathbf{B}, \mathbf{C}}{\text{minimize}} \quad \sum_{n=1}^N \left\| \boldsymbol{\Sigma}_n - \mathbf{B} \text{diag}(\mathbf{c}_n) \mathbf{B}^T \right\|_F^2 \\
& \text{subject to} \\
& \quad \|\mathbf{b}_k\|_1 \leq \lambda, \quad k = 1, \dots, K, \\
& \quad -1 \leq \mathbf{b}_k(i) \leq 1, \quad \max_i |\mathbf{b}_k(i)| = 1, \quad i = 1, \dots, P \\
& \quad \mathbf{c}_n \geq 0, \quad n = 1, \dots, N
\end{aligned} \tag{2.2}$$

2.3.2 Optimization strategy

The objective function in the proposed model is non-convex w.r.t both unknown variables \mathbf{B} and \mathbf{C} . We use the method of alternating minimization to solve for \mathbf{B} and \mathbf{C} . At each iteration a local minimum is obtained using projected gradient descent (Batmanghelich et al., 2012). Such a procedure converges to a local minimum.

2.3.3 Model selection

The free parameters of the model are the number of SCPs K and the sparsity level of each SCP λ . As values of K and λ are increased, the approximation error is reduced; however beyond a certain value of K it is likely that the model is over-fit to the data; i.e., the SCPs computed by the algorithm are possibly used to explain noisy (unwanted) variations in individual subjects. Hence we resort to cross-validation in order to avoid over-fitting. Using a grid search, for each value of the parameters K and λ repeated two-fold cross-validation is performed, and the value at which



(a) Node-based volumetric rendering of \mathbf{b}_k . Overlay color indicates the membership value of each region in \mathbf{b}_k .

(b) Edge-based rendering of $\mathbf{b}_k \mathbf{b}_k^T$. Size of nodes indicate value in \mathbf{b}_k .

Figure 2.3: Visualization of SCPs. Opposing colors reflect anti-correlation between regions.

there is no gain in generalizability (drop in error) is chosen to be the operating point. This provides us with SCPs that might generalize better across data. The cross-validation measure is the error computed on the test dataset relative to the variance in the test data, defined as follows:

$$\text{Test Error} = \frac{\sum_{n=1}^{N^{test}} \left\| \Sigma_n^{test} - \mathbf{B}^{train} \text{diag}(\mathbf{c}_n^{test}) (\mathbf{B}^{train})^T \right\|_F^2}{\sum_{n=1}^{N^{test}} \left\| \Sigma_n^{test} - \bar{\Sigma}^{test} \right\|_F^2} \quad (2.3)$$

where $\bar{\Sigma}^{test}$ is the subject-averaged correlation matrix of the test dataset.

2.3.4 SCP visualization

Node-based visualization

Recall that each SCP \mathbf{b}_k is of size P , where P is the number of nodes, defined either using the Areal Graph node definition or GraSP parcellation. For visualization purposes, SCPs \mathbf{b}_k with large spatial extent are mapped to the voxels in the template space. In the case of the Areal Graph node definition, we performed dual-

regression, similar to the procedure described in Smith et al. (2012) to map the SCPs from nodes to voxels. For nodes defined using GraSP, the membership values of each parcel within an SCP is directly mapped to the spatial extent of that parcel. These voxel-wise maps can be visualized using overlays in volumetric space in FSLView (Jenkinson et al., 2012), as shown in Figure 2.3a or projected onto the cortical surface using Caret Visualization software (Van Essen et al., 2001).

Edge-based visualization

Smaller SCPs can be mapped to the brain using their edge-based representation $\mathbf{b}_k \mathbf{b}_k^T$, rendered using BrainNet Viewer (Xia et al., 2013) as shown in Figure 2.3b.

2.4 Experiments on simulated data

2.4.1 Generation of simulated data

In order to illustrate the behaviour of our method, we generated a synthetic dataset with forty instances (or individuals). Our connectome design is illustrated in Figure 2.4. Each subject is associated with a underlying connectome configuration consisting of fifty nodes, as shown in Figure 2.4a. This connectome is designed in such a way that it has eight SCPs, with SCP size varying between three and ten nodes. Some of these SCPs are overlapping, with overlap size varying between one and three nodes. The strength of each SCP varies across subjects in a binary fashion, i.e, in each individual, SCPs are either “active/on” or “inactive/off”. In other words, an SCP is inactive in a individual when all the edges/correlation strengths of that SCP are close to zero for that particular individual. Furthermore, individuals are categorized into three groups based on whether SCPs are active or inactive, simulating

inter-individual variability.

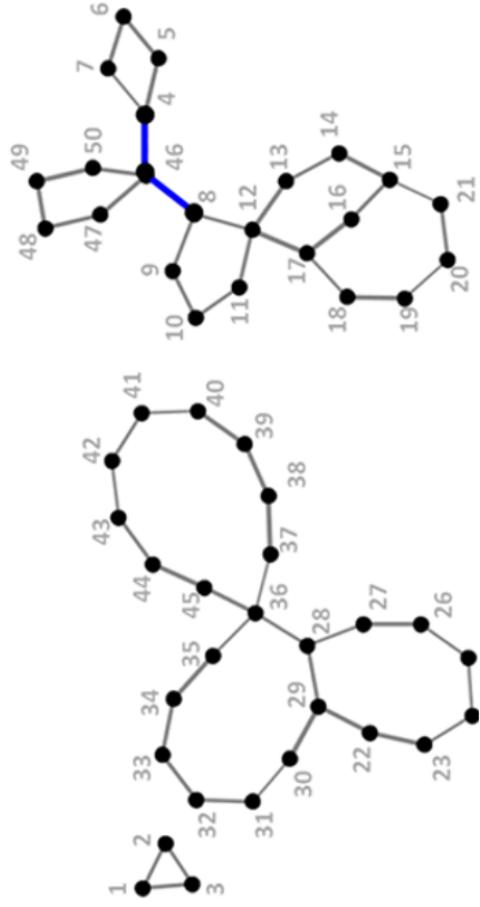
We input this connectome design into the simulation software NetSim (Smith et al., 2011), which simulates BOLD time-series at each node. Each time-series has 120 time-points and a TR value of 3 seconds, making each dataset 6 minutes long, similar to standard clinical rsfMRI acquisitions. In addition to inter-individual variability introduced by differential activation of SCPs we also include small random perturbation of edge strength for all edges. Random external inputs are input to some of the nodes. Thermal (white) noise is added to the output time-series at each node.

Correlation matrices are computed from the simulated time-series for all forty individuals, which form the input to our method. The average correlation matrix is shown in Figure 2.4b. The matrices shown in Figure 2.4c correspond to the correlation values computed from the time-series for five randomly chosen individuals. The average correlation matrix, thresholded to obtain varying levels of edge density was used as input to InfoMap. Absolute values of correlation were used as input to InfoMap. The clustering assignments output by InfoMap are converted into a set of binary basis vectors, one for each sub-graph. Concatenated time-series data were used as input for Temporal ICA as well as Spatial ICA. In case of ICA, dimensionality reduction was performed by running PCA first, as is routinely done in fMRI-ICA literature (Smith et al., 2012).

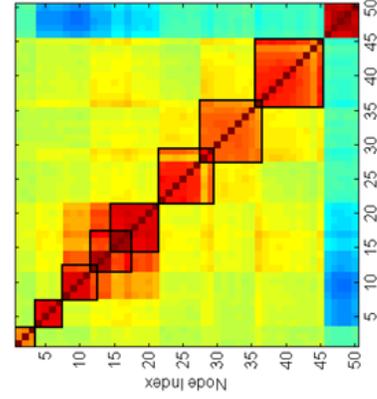
2.4.2 Performance on simulated data

The output of the cross-validation experiments are shown in the plots in Figures 2.5, which shows the variation of the cross-validated mean square error as the free parameters, K and λ , are varied. It is clear that the MSE saturates beyond $\lambda = 0.2$. Choice of K is somewhat unclear. Using $K = 8, \lambda = 0.2$ as the operating point, we

a



b



c

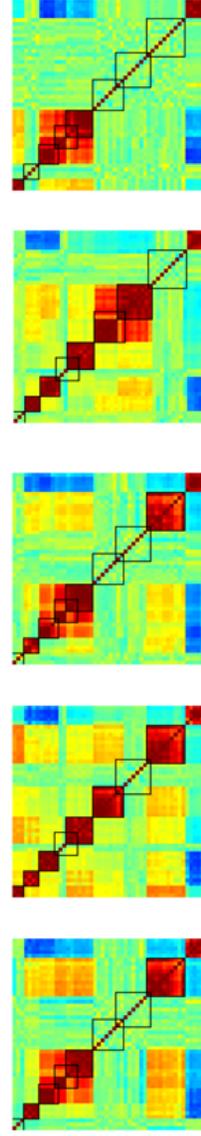


Figure 2.4: Simulated connectome design consisting of 50 nodes. (a) All nodes and edges. Nodes are common to all individuals, different sets of edges are inactive in different individuals. Edges in blue indicate anti-correlation. (b) Average correlation matrix across all individuals. (c) Correlation matrices of five randomly chosen individuals.

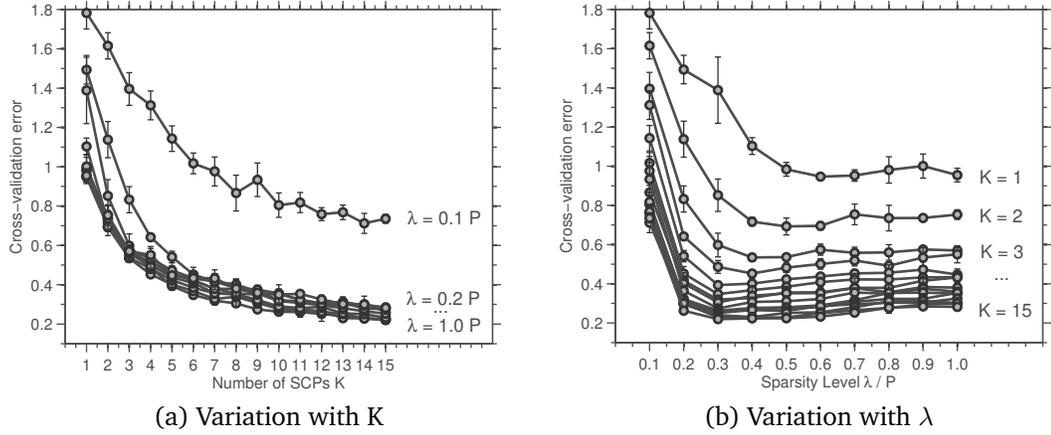


Figure 2.5: Cross-validation results for simulated data: Plots of the mean square error (Eqn. 2.3) vs. number of SCPs K (left), and sparsity level λ (right).

computed the basis vectors for the simulated data using all four methods. (Note: For the sampling of edge-densities used to compute results for InfoMap, we were able to obtain $K = 7$ communities followed by $K = 9$. As $K = 7$ has greater accuracy, we display those results.)

The first image in Figure 2.6 shows the simulated ground-truth as a vector of node-weights. The results from the four methods are shown next to the ground-truth. Each column displays a basis vector \mathbf{b}_k . It is easy to see that SCPs computed using Sparse Learning are closest to the ground-truth.

In order to quantify the performance of the algorithm on simulated data, we compare the set of SCPs output by Sparse Learning \mathbf{B} with the ground truth \mathbf{B}^{true} . Before the comparison we first perform a one-to-one matching between the two sets of vectors using the Hungarian Algorithm (Munkres, 1957). The sign of some of the vectors in \mathbf{B} is reversed, if necessary. We use the normalized inner product (cosine of the angle between vectors) to compare these paired set of vectors.

Figure 2.7 shows the accuracy of SCPs, sub-graphs, ICNs and TFM for vary-

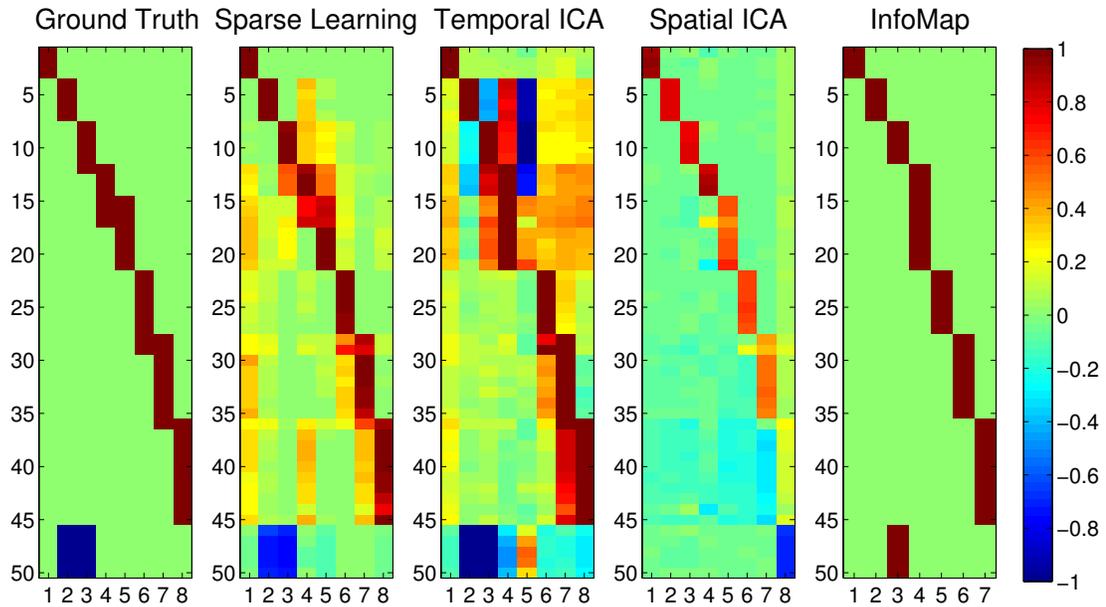


Figure 2.6: Performance on simulated data. The basis vectors identified by the Sparse Learning approach, InfoMap, Spatial and Temporal ICA, shown as node-weights (b_k), compared against the ground-truth.

ing K . When compared with the ground-truth, SCPs show slightly higher accuracy than the three other methods, for all values of K ($p < 0.05$, compared at $K=8$, using a two-sided t-test for Sparse Learning vs. all other methods). Temporal ICA comes a close second, as it is able to capture many of the positive/negative correlations but results in a denser basis. Both Spatial ICA and InfoMap produce non-overlapping ICNs/sub-graphs - and as the value of K is increased, these components get smaller/more fragmented, leading to a drop in accuracy, as seen in the graph.

Finally, Figure 2.8 displays the individual level coefficients estimated by Sparse Learning, along with the ground truth. This result shows that Sparse Learning is able to capture the heterogeneity at the level of individuals, since the clustering of the three groups is retained to a large extent in the estimation.

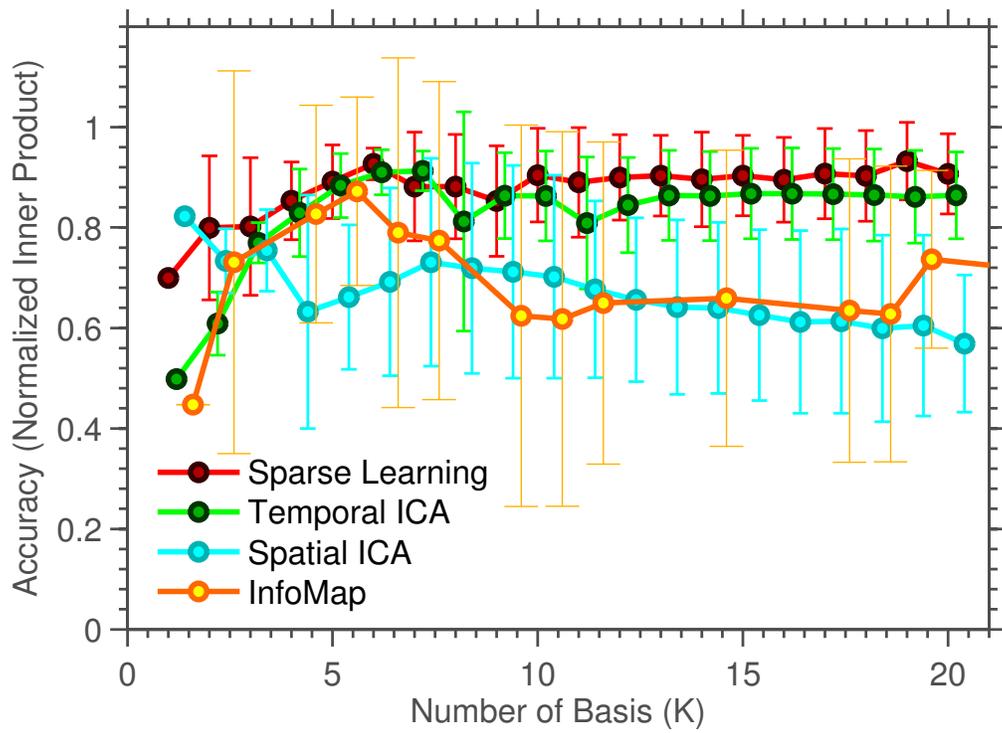


Figure 2.7: Performance on simulated data. Accuracy, measured using normalized inner-product.

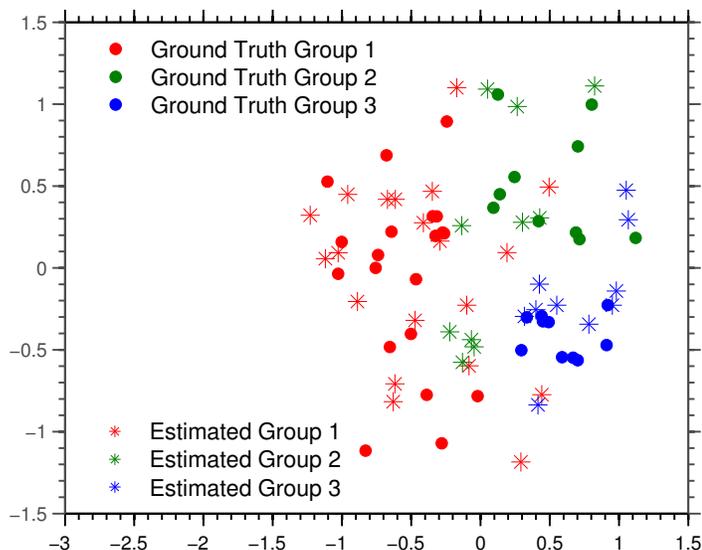


Figure 2.8: Individual-level coefficients estimated by Sparse Learning, projected down to 2-D space, shown along with the ground truth.

2.5 SCPs computed from rsfMRI data

2.5.1 Areal graph nodes

Having shown that Sparse Learning performs better on synthetic data that has spatial and temporal overlaps, we next compared performance of the three methods using resting state data from 130 healthy, young adults between the ages 19 to 22 years, acquired as a part of the Philadelphia Neuro-developmental Cohort (PNC) (Satterthwaite et al., 2014). We used the Areal Graph nodes to define one correlation matrix for each individual, as described in Chapter 1.

Figure 2.9 plots the variation of the cross-validated error as K and λ is varied. From these figures, it is clear that the MSE saturates around $\lambda = 0.3$. However, the “knee” of the graph plotting variation with K is unclear.

Sparse Learning was run on the entire sample of 130 individuals with the values

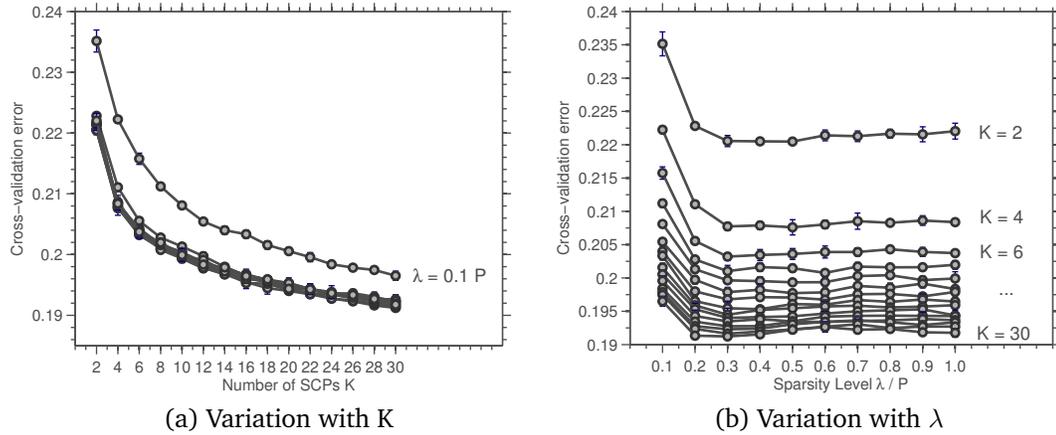


Figure 2.9: Cross-validation results for PNC data: Plots of the mean square error (Eqn. 2.3) vs. number of SCPs K (left), and sparsity level λ (right).

$K = 10, \lambda = 0.3$. These ten SCPs are shown in Figure 2.10. We describe them in detail below, and compare to existing knowledge of the spatial extent and behavior of known task-processing, attention and control systems.

Dorsal Attention SCP Figure 2.10 shows the first SCP defined by the anterior middle temporal area (aMT), superior parietal lobule (SPL), intra parietal sulcus (IPS) and the frontal eye fields (FEF)(shown in red), which are known to be part of the Dorsal Attention (DA) system (Corbetta and Shulman, 2002). These regions are anti-correlated with the middle temporal gyrus (MTG), inferior parietal lobule (IPL), medial pre-frontal cortex(mPFC), posterior cingulate cortex (PCC) and anterior frontal operculum, which are part of the default-mode (DM) system (Raichle et al., 2001)(shown in blue).

Executive Control SCPs SCPs 2, 3 and 4 predominantly show executive task-control system (red) anti-correlated with different aspects of the DM system (blue). SCP 4 shows the Salience system (Seeley et al., 2007) consisting of dorsal anterior cingulate cortex (dACC) along with anterior insula and the anterior pre-frontal cortex.

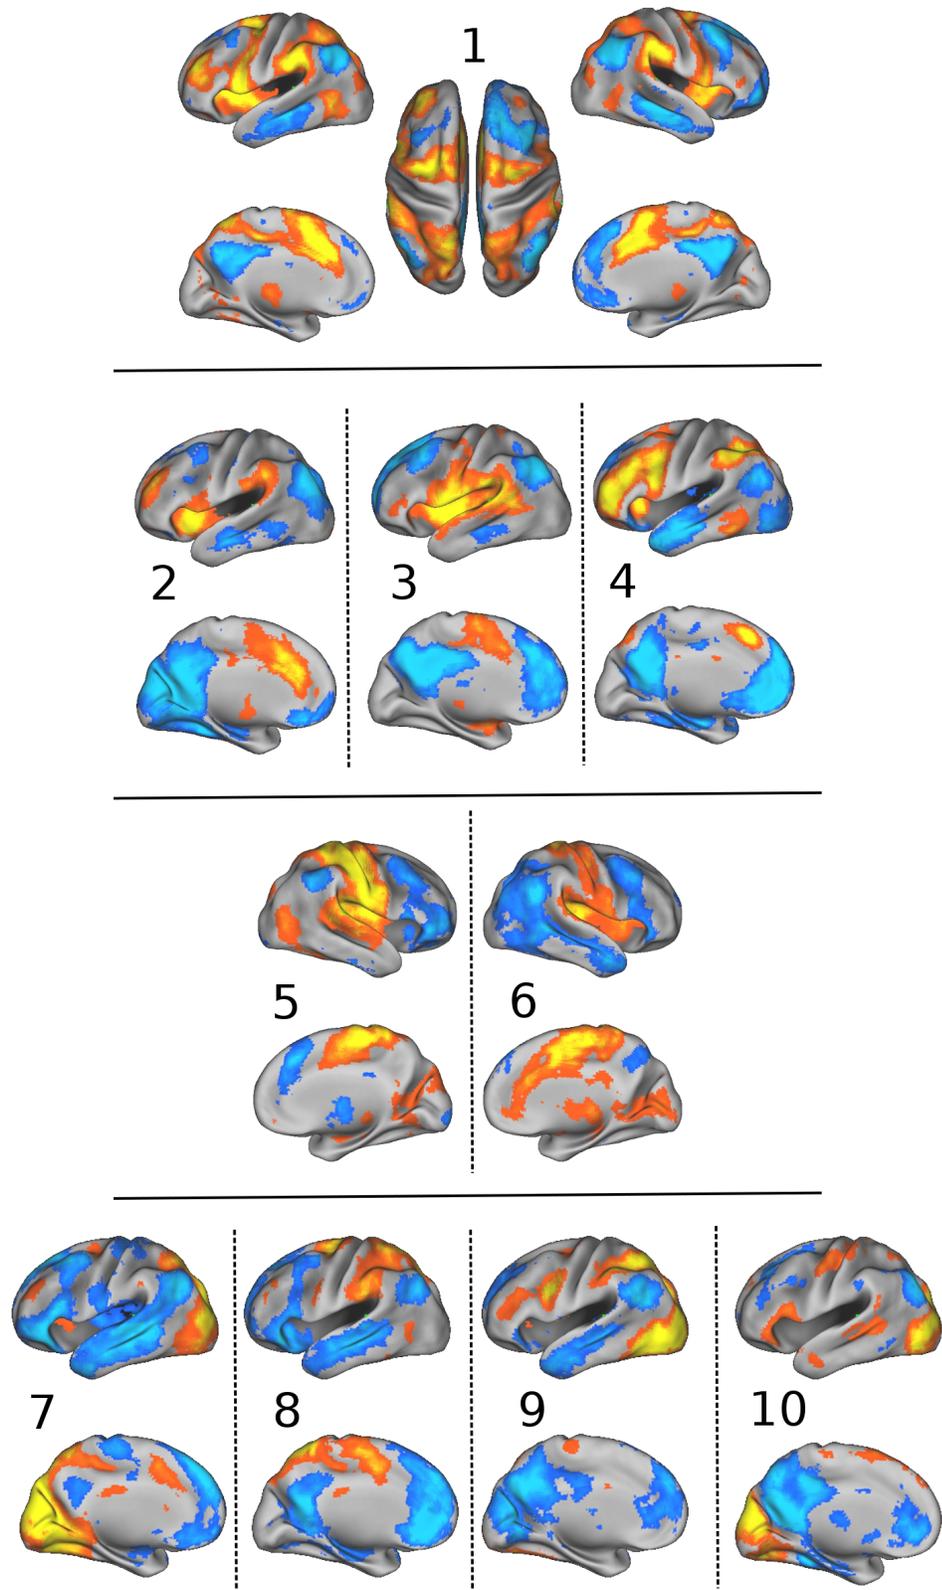


Figure 2.10: Ten SCPs computed in the Areal Graph node space.

The anti-correlated DM regions include the IPL, PCC and vmPFC. Regions from the operculum, insula, temporal-parietal junction (TPJ), inferior frontal gyrus (IFG) and the dACC dominate SCP 5, with anti-correlations to PCC and dmPFC. This SCP consists of the Cingulo-Opercular (COP) system (Dosenbach et al., 2007) which is known to de-activate the DM. SCP 6 consists of the aPFC, aI, IPL, and MT, which form the Fronto Parietal task-control system, anti-correlating with the inferior MTG, IPL, PCC, mPFC and PHC.

Motor SCPs SCPs 5 and 6 exhibit contributions from the sensori-motor, auditory and visual areas. Both SCPs show the pre central (prCG) and post central gyrus (poCG). In SCP 5 the motor areas positively correlate with the superior temporal gyrus (STG) and posterior insula. The positive correlations in SCP 6 are more anterior within the insula, and a large extent of the cingulum. Anti-correlated regions include the FP system (aPFC, IPL, aI, ACC) in SCP 5 and aspects of the DM system in SCP 6.

Visual SCPs SCPs 7, 8, 9 and 10 show four types of connectivity patterns involving the visual areas. SCP 7 covers the entire visual system, including the medial visual, lateral visual and higher visual (dorsal attention) areas. SCP 8 shows the higher visual areas alone. The visual areas are anti-correlated with the DM system in both the SCPs. SCPs 9 and 10 shows contributions from areas in the lower levels of the visual hierarchy; the FEF and the prCS are less dominant, while including the lateral visual areas, which are involved in higher level visual task-processing. Concomitant with moving down along the hierarchy, we observe changes to the anti-correlated regions - the involvement of the mPFC is greatly reduced, but the anti-correlation the posterior cingulate is retained.

Overlap between SCPs The SCPs described above are clearly overlapping, mainly enabling the description of multiple relationships between a functional system and

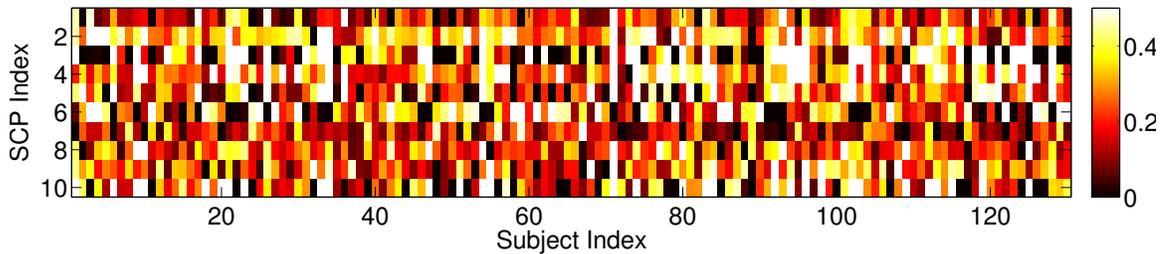


Figure 2.11: Figure illustrating the heterogeneity of the data sample captured by SCPs. The color indicates the extent to which each SCP is present in a individual.

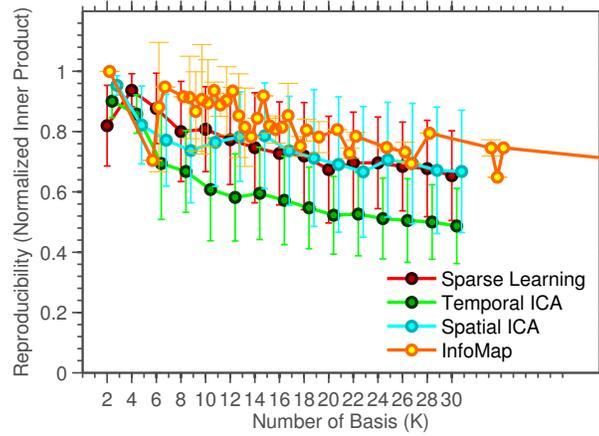
other systems. We note that the PCC and the IPL contribute to most of the SCPs, which were identified by a prior study as one of the central hubs of connectivity in the brain (Buckner et al., 2009).

Inter-individual variability Differential presence of the SCPs explains inter individual variability in functional connectivity. Figure 2.11 shows the strength of presence of each SCP in every individual. The Sparse Learning approach exploits this variability; a heterogeneous distribution of the samples in this lower-dimensional space allows robust identification of the SCPs.

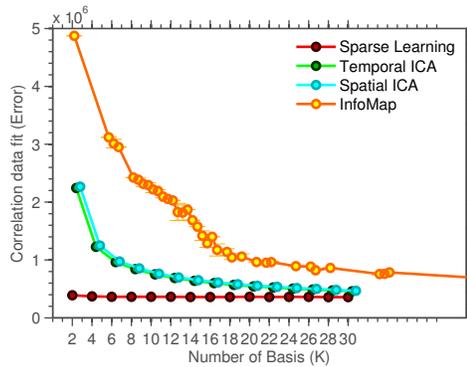
Reproducibility

We evaluate the performance of our algorithm as well as Infomap and ICA based on repeated split-sample reproducibility. Reproducibility was evaluated for $K = 2, 4, \dots, 30$. In the case of InfoMap, the edge-density was varied between 2% and 40%. This provided sub-graphs varying in number from 4 upto 60, although not equally spaced. Similar to our earlier experiments involving simulated data, we quantify the comparison between sub-samples using the normalized inner product, averaged across basis vectors.

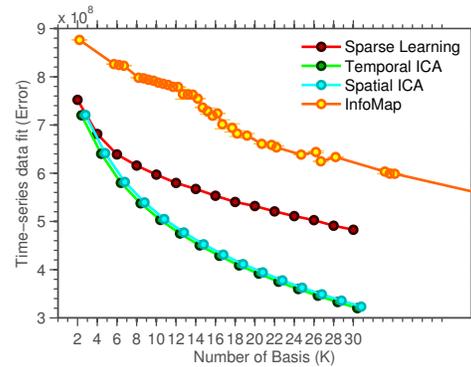
The reproducibility of the results is shown in Figure 2.12a, computed for values of $K \in 2, 4, \dots, 30$. For $K = 10$, the reproducibility of our results is 0.80 ± 0.09 ,



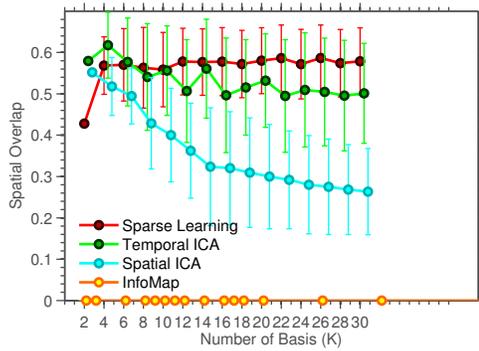
(a) Reproducibility



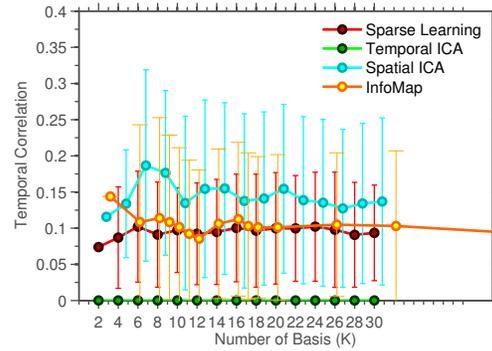
(b) Correlation data-fit (Error)



(c) Time-series data-fit (Error)



(d) Spatial Overlap



(e) Temporal Correlation

Figure 2.12: Reproducibility, data-fit, spatial overlap and temporal correlation measured for SCPs obtained from Areal Graph nodes.

compared to 0.86 ± 0.15 for InfoMap, 0.79 ± 0.07 for Spatial ICA and 0.60 ± 0.12 for Temporal ICA. InfoMap shows comparable reproducibility with sparse learning ($p < 0.4$, computed using two-group t-test). Spatial ICA components are as reproducible as Sparse Learning ($p < 0.8$), while Temporal ICA performs significantly worse ($p < 10^{-4}$) in terms of reproducibility.

Data fit

In addition to reproducibility, the data fit (approximation error) of all the methods to the data was also compared. Given that Sparse Learning and InfoMap use correlation values as input, and ICA methods use time-series as input, we evaluated the approximation error for both types of input. Let \mathbf{B} denote the set of basis vectors output by any of the four methods. Let $\mathbf{Y}_n \in \mathbf{R}^{P \times T}$ and $\mathbf{X}_n \in \mathbf{R}^{K \times T}$ be the time-series of the n th individual and basis respectively. In the case of Sparse Learning and InfoMap, the basis-specific time-series \mathbf{X}_n can be computed by regressing the basis \mathbf{B} against the individual-specific time-series \mathbf{Y}_n . Using these values, the correlation data-fit measure is computed as follows:

$$\text{Correlation data fit (Error)} = \sum_{n=1}^N \left\| \left\| \boldsymbol{\Sigma}_n - \hat{\boldsymbol{\Sigma}}_n \right\|_F \right\|^2 \quad (2.4)$$

where

$$\begin{aligned} \hat{\boldsymbol{\Sigma}}_n &= \mathbf{B} \text{diag}(\mathbf{c}_n) \mathbf{B}^T && \text{for Sparse Learning} \\ &= \mathbf{B} \mathbf{X}_n \mathbf{X}_n^T \mathbf{B}^T && \text{for InfoMap, Spatial and Temporal ICA} \end{aligned} \quad (2.5)$$

Similarly the time-series data fit for all four methods is defined as

$$\text{Time-series data fit (Error)} = \sum_{n=1}^N \|\mathbf{Y}_n - \mathbf{B} \mathbf{X}_n\|_F^2 \quad (2.6)$$

The correlation data-fit measure is the same as the objective function that is optimized in the Sparse Learning method. Obviously, Sparse Learning is expected to have the best correlation data-fit (lowest error). Similarly, as the ICA methods operate on the time-series as input, they are likely to outperform Sparse Learning and InfoMap with respect to time-series data fit. It is more interesting to note how the correlation methods compare with respect to time-series data fit, and vice versa.

Figure 2.12b and 2.12c show the correlation data-fit and the time-series data-fit of all the methods for varying values of K . InfoMap performs poorly in terms of data-fit for both correlation as well as time-series data. As expected, Sparse Learning has the best correlation data-fit, while the ICA methods provide the best time-series data fit. This is one of the primary reasons why Sparse Learning is a better choice for describing connectome data, as it explains a greater amount of variance in the connectivity data.

Spatial Overlap and Temporal Correlation

Finally, to further understand the behavior of the methods under consideration, we computed the degree to which the estimated basis vectors are spatially overlapping and temporally correlated. These values are computed as follows:

$$\text{Spatial Overlap} = \sum_{i,j=1, i \neq j}^K \frac{|\mathbf{b}_i|^T |\mathbf{b}_j|}{\|\mathbf{b}_i\|_2 \|\mathbf{b}_j\|_2} \quad (2.7)$$

$$\text{Temporal Correlation} = \sum_{n=1}^N \sum_{i,j=1, i \neq j}^K \frac{\mathbf{x}_{ni}^T \mathbf{x}_{nj}}{\|\mathbf{x}_{ni}\|_2 \|\mathbf{x}_{nj}\|_2} \quad (2.8)$$

where x_{ni} denotes the time-series associated with the i th basis in the n th individual.

As expected, InfoMap has no spatial overlap in its basis. Spatial ICA shows decreasing overlap with increasing K , as shown in Figure 2.12d. Sparse Learning has the highest spatial overlap. Figure 2.12e shows the variation of the average temporal correlation with K . Temporal ICA has no temporal correlation while Spatial ICA has the highest temporal correlation. Of the four methods, only Sparse Learning has non zero spatial overlap as well as temporal correlation.

Reproducibility across datasets

To further evaluate the reproducibility of the presented results, we used two alternate datasets - (1) 131 scans pooled together from the pilot data of the Human Connectome Project (HCP) (Smith et al., 2012) and the BrainScope datasets from the functional Biomedical Informatics Research Network (fBIRN) (Fox et al., 2005, 2007) and (2) 400 scans from the BLSA dataset, described in chapter 1. All the scans were pre-processed using the same pipeline that was used for the PNC dataset, as detailed above.

We found that the SCPs are reasonably reproduced in the alternate datasets, with an average inner-product 0.65 ± 0.07 for HCP+fBIRN and 0.62 ± 0.10 for BLSA. A side-by-side comparison of four SCPs computed from the PNC and alternate datasets is shown in Figure 2.13.

2.5.2 GraSP parcels

In addition to the SCPs computed using Areal Graph nodes, we also generated SCPs for nodes defined by GraSP parcels, for both the PNC and BLSA datasets. Recall that GraSP parcels have full GM coverage, as opposed to Areal Graph nodes which

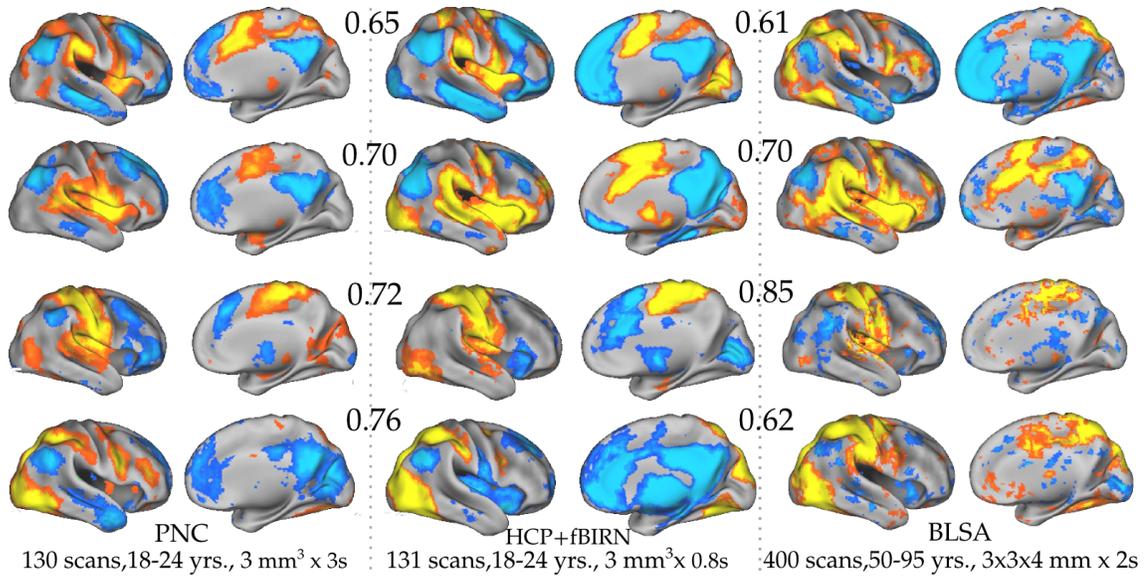


Figure 2.13: SCPs 1, 2, 5, and 9 computed from the PNC dataset(left), HCP+fBIRN dataset (middle) and BLSA dataset (right) using Areal Graph nodes. The inner product value (computed in the Areal Graph node space) for each comparison is also shown.

span only major functional activity peaks in the cortex, and sample some areas more densely than others. The SCPs for the PNC and BLSA datasets are shown side-by-side in Figures 2.14 and 2.15.

Visual inspection of Figures 2.14 and 2.15 shows that the two sets of SCPs are remarkably similar despite vast differences in scanner acquisition protocols, demographics and parcellation definition.

2.6 Extensions

2.6.1 Hierarchical Sparse Learning

Sparse Learning can be extended to a hierarchical framework with multiple levels, where each SCP is split into multiple smaller SCPs. This provides greater spatial specificity at the lower levels, and more SCP coefficients which can be used as

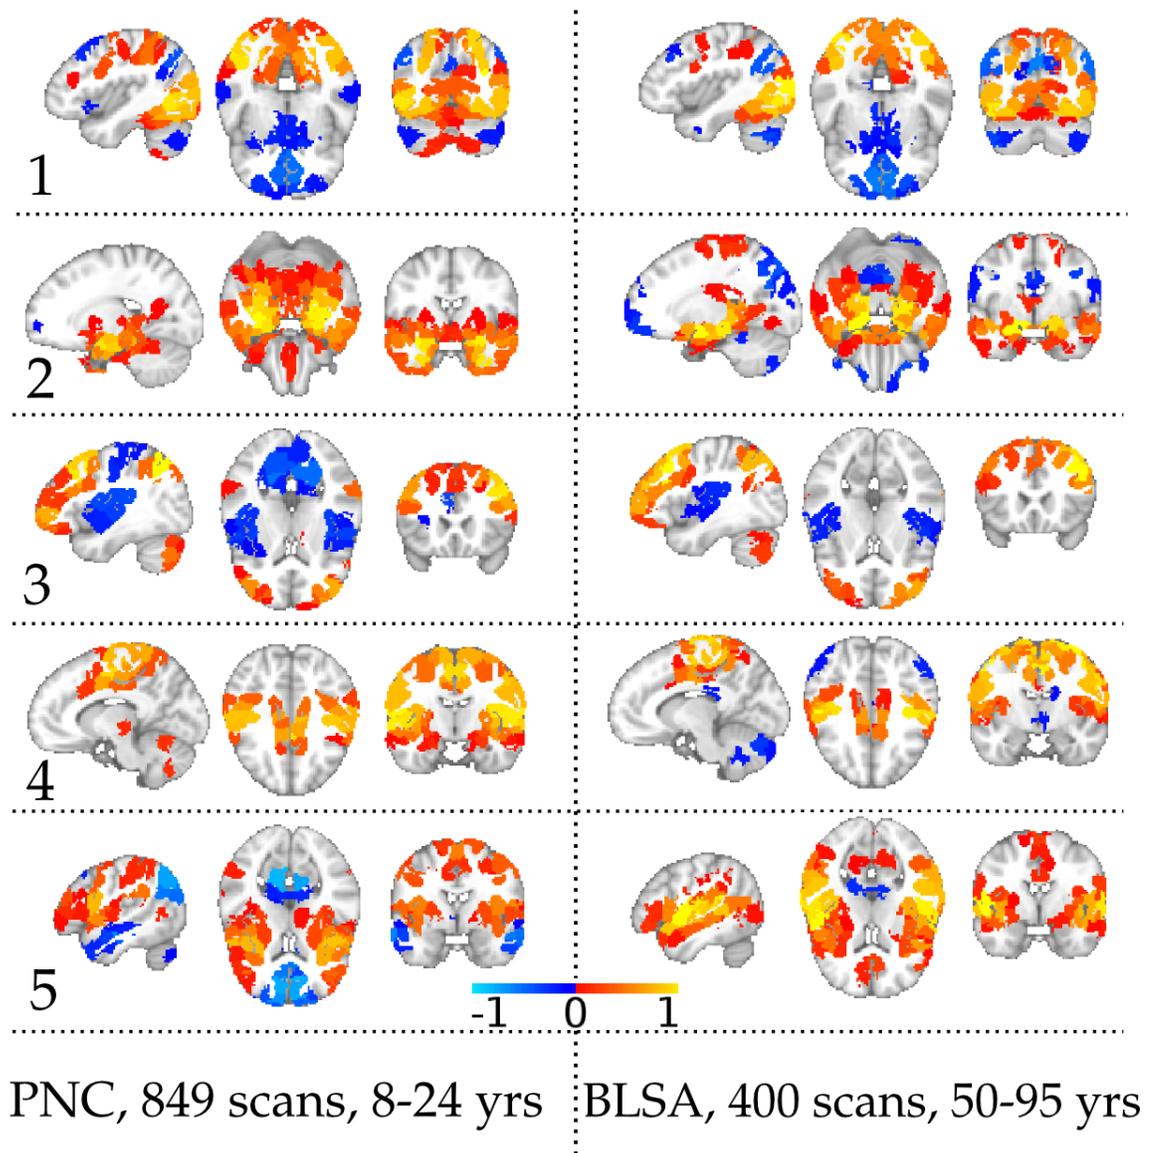


Figure 2.14: SCPs 1-5 generated from the PNC data (left) and BLSA data (right) using GraSP parcels.

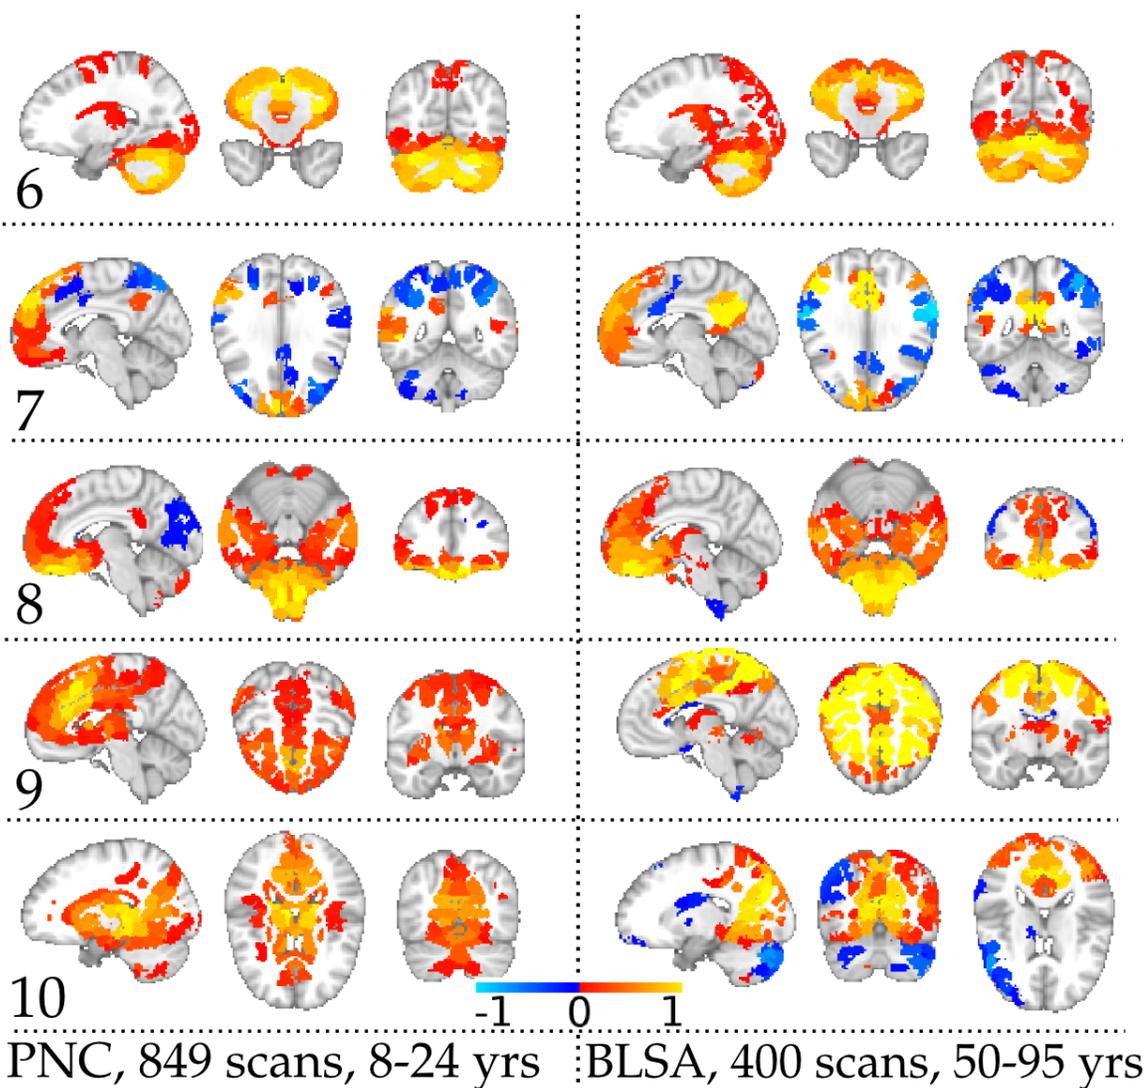


Figure 2.15: SCPs 6-10 generated from the PNC data (left) and BLSA data (right) using GraSP parcels. Note that SCPs 9 and 10 do not match between the two datasets.

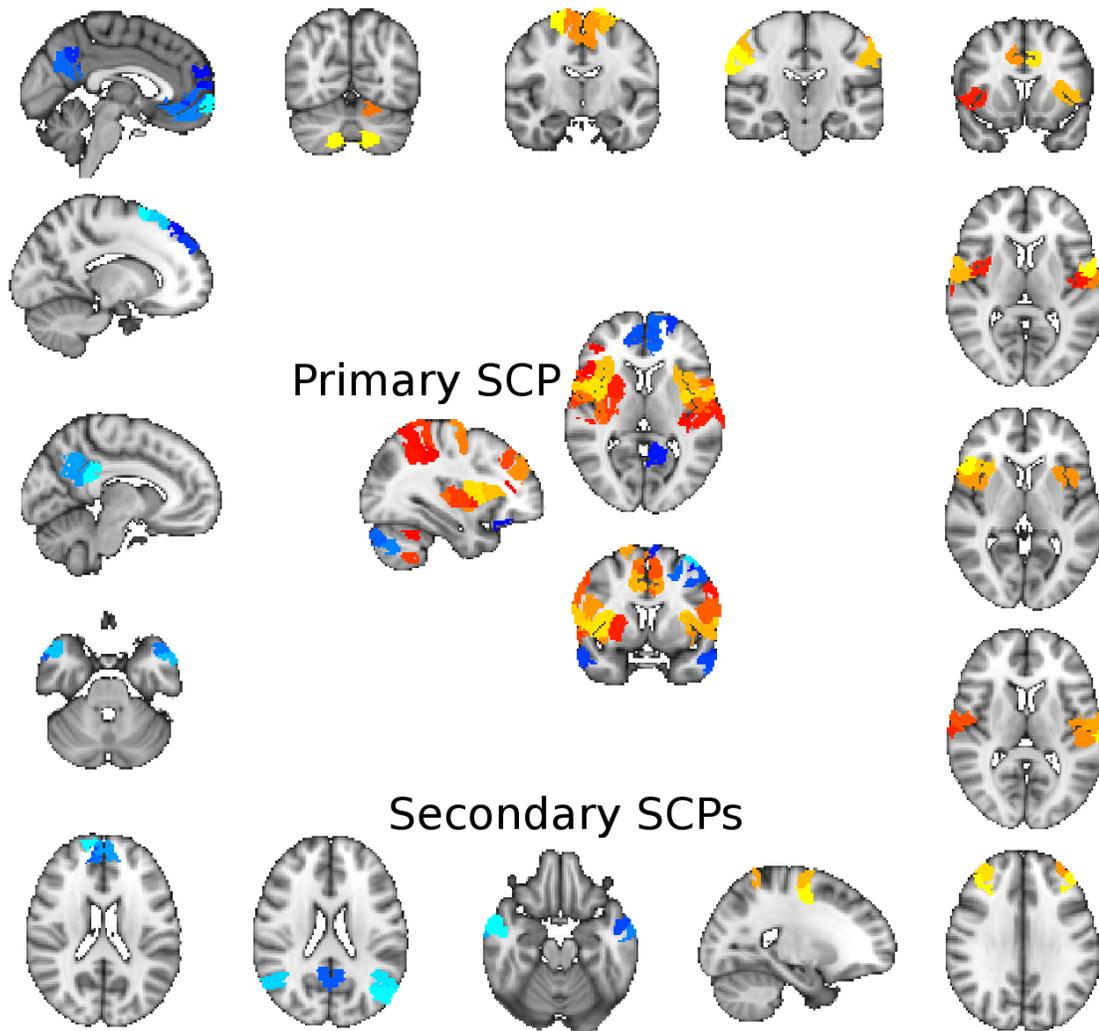


Figure 2.16: A primary SCP and some of its associated secondary SCPs computed from the PNC dataset using GraSP parcels.

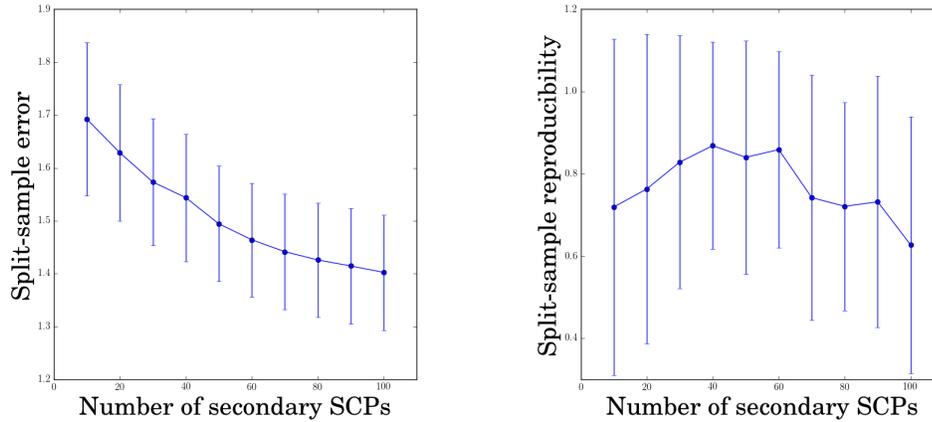


Figure 2.17: Split-sample error and reproducibility for secondary SCPs computed from the PNC dataset.

features for group-level analyses. For our experiments, we considered two levels. In the first level, “primary SCPs” are defined by using the original correlation matrices as input. In the second level, “secondary SCPs” are defined by re-applying the method to each primary SCP; i.e., only the rows and columns of the correlation matrix belonging to the primary SCP are used. At the secondary level, we used the maximum permissible sparsity value of $\lambda = 2$ (atleast two regions in each SCP). An example is shown in Figure 2.16.

Split-sample error and Reproducibility of Secondary SCPs

Similar to the evaluations performed for primary SCPs in Section 2.5.1, for a fixed set of primary SCPs, we computed the split-sample error and reproducibility as the number of secondary SCPs was varied. Figure 2.17 plots both measures. The reproducibility increases as as the number of SCPs is increased, and peaks between $K = 40 - 60$. The split-sample error reduces as the number of SCPs is increased;

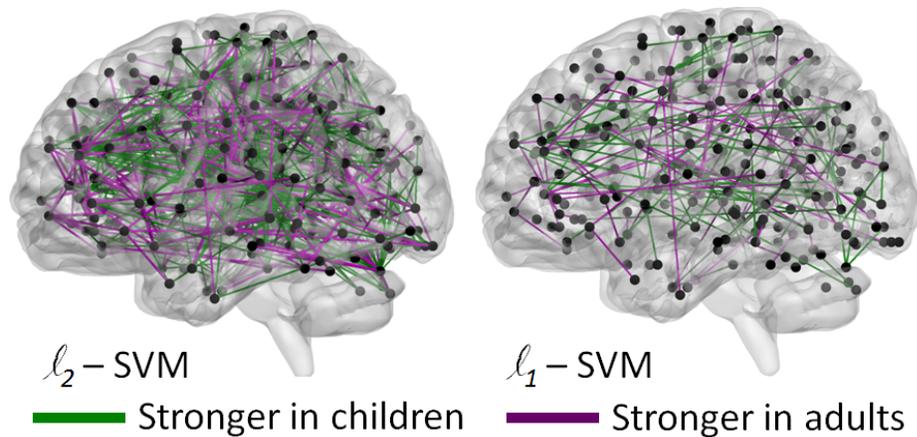


Figure 2.18: SVM weight vector for l_2 -SVM (left) and l_1 -SVM (right) for discriminating children from young adults using the PNC dataset.

however the graph does not have a clear “knee”.

2.6.2 Discriminative Sparse Learning

Many studies that examine changes in pair-wise connectivity use multi-variate methods, such as SVMs, with the vectorized correlation matrices as features (Satterthwaite et al., 2013). Applied to classification, SVM returns a list of connections or edges that are discriminative of the two groups. For example, the hyperplane weights for classification between 91 children (10.38 ± 1.01 yrs.) vs. 84 younger adults (20.21 ± 0.84 yrs.) from the PNC dataset are shown in Figure 2.18. Clearly, these high-dimensional patterns are very difficult to parse and interpret, requiring additional post-processing, such as thresholding. To address this issue, l_1 -regularized SVM can be used, which produces a sparse discriminative pattern (Figure 2.18, right). However, it is known to ignore features that are highly correlated (and therefore redundant for classification), which could be neuro-biologically relevant.

The Sparse Learning framework described in the previous sections can be used

to reduce dimensionality prior to the classification task. However a purely generative method seeks to represent functional activity in general, and does not aim specifically to find networks that relate to the classification task. To directly obtain patterns that are affected by the condition being studied, i.e., the classification task, we provide an extension of the current Sparse Learning framework to incorporate two-group classification, resulting in functionally meaningful SCPs that are discriminative of the two groups. Our method is akin to previous work in structural MRI (Batmanghelich et al., 2012), presenting a joint generative-discriminative formulation. Thus, this method reduces the high dimensionality of the connectivity data to small sets of regions, whose average correlation contributes to the classification task.

A schematic diagram illustrating our method is shown in Figure 2.19. As before, the input to the method is size $P \times P$ correlation matrices $\Sigma_n \succeq 0$. In addition, for the classification task, we know the associated binary group membership y_n , for each individual n , $n = 1, 2, \dots, N$. We would like to find SCPs common to all the individuals, such that the total connectivity within each network contributes to the two-group classification. Our formulation jointly optimizes two objectives: (1) Identification of SCPs (2) Learning discriminative SCPs.

The Sparse Learning objective function defined in Eqn. 2.2 is the *generative* term that identifies SCPs given the correlation data. We will call this generative objective function $\mathcal{G}(\mathbf{B}, \mathbf{C})$. Given an SCP \mathbf{b}_k , the scalar value $\mathbf{b}_k^T \Sigma_n \mathbf{b}_k$ measures the total absolute correlation between all the regions within the SCP for a given individual n . Computed for all SCPs, the K-dimensional vector $diag(\mathbf{B}^T \Sigma_n \mathbf{B})$ serves as the individual-specific measure that can be used in a multi-variate SVM framework. We use the squared hinge loss and l_2 regularization for the K-dimensional SVM

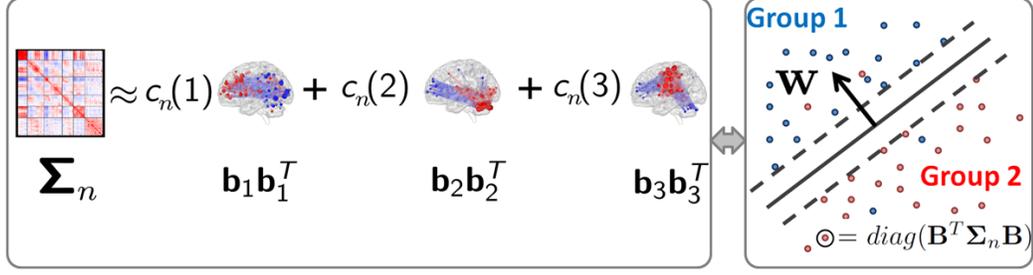


Figure 2.19: Schematic illustrating the discriminative Sparse Learning framework. Panel to the left describes the SCP identification term, which factorizes connectivity matrices Σ_n of each individual n into a set of common SCPs $\mathbf{B} = [\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_K]$ and its associated coefficients. Panel to the right illustrates a linear SVM, which uses the total absolute connectivity values $\text{diag}(\mathbf{B}^T \Sigma_n \mathbf{B})$ of all SCPs as input features to classify two groups, resulting in the hyperplane \mathbf{w} .

hyperplane \mathbf{w} . The cost function for the discriminative term is:

$$\mathcal{D}(\mathbf{B}, \mathbf{w}) = \sum_{n=1}^N (1 - y_n \mathbf{w}^T \text{diag}(\mathbf{B}^T \Sigma_n \mathbf{B}))_+^2 \quad (2.9)$$

where y_n is the binary group label for individual n , and the subscript $+$ denotes the positive part of the argument.

Bringing the terms in Eqns. 2.2 and 2.9 together, along with the l_2 regularizer for \mathbf{w} , we have the optimization problem

$$\begin{aligned} & \underset{\mathbf{B}, \mathbf{C}, \mathbf{w}}{\text{minimize}} \quad (1 - \mu) \mathcal{G}(\mathbf{B}, \mathbf{C}) + \mu \mathcal{D}(\mathbf{B}, \mathbf{w}) + \|\mathbf{w}\|_2^2 \\ & \text{subject to} \\ & \quad -1 \leq \mathbf{b}_k \leq 1, \quad \|\mathbf{b}_k\|_1 \leq \lambda, \quad k = 1, \dots, K \\ & \quad \mathbf{c}_n \geq 0, \quad n = 1, \dots, N \end{aligned} \quad (2.10)$$

where μ is the relative weighting fraction between the two terms. A value of $\mu = 0$ produces purely generative SCPs. We use alternating minimization to iteratively solve for \mathbf{B} , \mathbf{C} and \mathbf{w} . As before, we use a projected gradient method (Batmanghe-

lich et al., 2012) for \mathbf{B} and \mathbf{C} and the *libSVM* solver for \mathbf{w} (Chang and Lin, 2011). The parameter μ , which controls the trade-off between the two terms, will be linearly increased from a value of 0 to 1 during the iterative process. This ensures that the SCPs generated during the first few iterations are mainly generative, which tend to be more stable.

The free parameters of the proposed method are the number of SCPs K , and the sparsity level λ . Using grid search, for every pair of values in $K \in \{10, 20, \dots\}$ and $\lambda \in \{0.01, 0.02, \dots, 0.1\} * P$, we will use repeated five-fold cross-validation to find the optimal set of parameters.

For the children vs. younger adults classification task, the results of the cross-validation provided an operating point of $K = 50, \lambda = 0.03P$ (roughly 10 nodes per SCP). The cross-validation accuracy saturates at higher values of K . For these values, the proposed method gave an average classification accuracy of $76.3 \pm 7.08\%$ between children vs. young adults. The two most discriminative SCPs are shown in Figure 2.20, using edge based rendering Xia et al. (2013).

We compared our method with four alternate approaches: (1) l_2 -regularized l_2 -loss linear SVM (Chang and Lin, 2011) with pair-wise correlation values (2) l_1 -regularized l_2 -loss linear SVM (3) Principal Component Analysis (PCA), followed by classification (4) Independent Component Analysis (ICA), followed by classification. The first and second methods are purely discriminative, as they do not perform dimensionality reduction. The third and fourth methods use unsupervised network identification methods, followed by classification using the total absolute connectivity values as features. The number of components K (in PCA and ICA) and the cost parameter for the SVM was chosen using cross-validation.

The classification performance for all the methods is reported in Table 2.1. The un-supervised PCA and ICA methods perform poorly. Of the three methods, l_2 -SVM

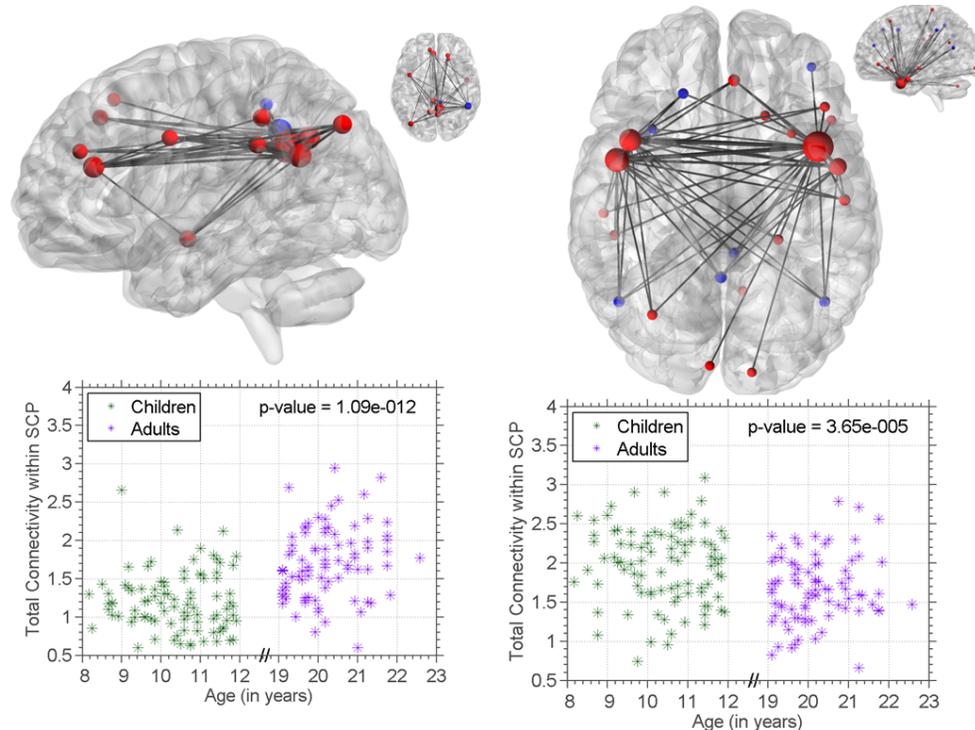


Figure 2.20: Top two SCPs that contribute to classification, based on their hyperplane weight. Corresponding graphs (bottom) plot total connectivity within SCP for each individual vs. individual age. Uni-variate p-value scores comparing total connectivity between two groups are also shown.

provides a slightly better performance compared to the proposed method, although the difference in accuracies between the two methods is insignificant ($p = 0.0625$). The marginally higher accuracy provided by the SVM is due to the 1000-fold increase in the number of features used, leading to a complete loss of interpretability as illustrated in Figure 2.18. The weight vector for l_1 -SVM is also shown in the same figure. While the l_1 penalty does dramatically reduce the number of features used, it does not necessarily alleviate the issue of non-interpretability. As explained earlier, strongly correlated features (connections) that are redundant to the classification are dropped. In contrast, the generative term \mathcal{G} within the proposed method tends to retain these features by allocating them to the same SCP. Thus, a whole-

Method	Acc (%)	AUC	# features
Disriminative Sparse Learning	76.3 ± 7.08	0.85 ± 0.06	50
l_2 -SVM	79 ± 1.45	0.87 ± 0.01	34716
l_1 -SVM	74 ± 2.25	0.81 ± 0.01	139
PCA+SVM	65.2 ± 2.30	0.67 ± 0.01	180
ICA+SVM	64.8 ± 1.42	0.71 ± 0.03	210

Table 2.1: Accuracy and AUC values for disriminative Sparse Learning, and other methods

brain discriminative pattern is split into multiple SCPs based on the dependencies between the connection strengths.

2.7 Discussion

The findings presented in this chapter are obtained from a connectivity-based modeling approach that identifies SCPs based on important observations - (1) Not all brain regions participate in a given SCP; (2) SCPs consist of those regions which are functionally correlated (or anti-correlated); and (3) SCPs consist of those regions whose connections co-vary across individuals, i.e., in other words, SCP coefficients capture inter-individual variability. Furthermore, the Sparse Learning framework can be extended to provide SCPs with greater spatial specificity using a hierarchical decomposition, or, with an additional discriminative term, capture only those patterns of connections relevant to a classification task.

Observations based on simulated experiments

Based on the simulated ground truth experiments described in Section 2.4, compared to Sparse Learning, both InfoMap and ICA are also able to capture many of the correlated/anti-correlated relationships between regions. However, as the sim-

ulation results show, their inherent methodological constraints result in networks with lower accuracy (for InfoMap and Spatial ICA). InfoMap and Spatial ICA are able to identify strongly correlated sets of nodes, while avoiding the points of overlap.

The results shown in Figure 2.6 also illustrate a potential limitation of the proposed method; when the network sizes are unbalanced, such as in the simulation design used here, using the same sparsity level λ is inappropriate. Picking a value of λ less than size of the largest SCP leads to those SCP getting truncated, and is sub-optimal. On the other hand, picking a value equal to the size of the largest community (= 20% in the simulated case) leads to noisy assignments in the smallest community, i.e, SCP 1, as seen in Figure 2.6. This limitation is similar to that of the edge density parameter used in InfoMap, where lower edge density leads to smaller communities. In practice, however, these false (noisy) assignments that are incorporated in the basis are generally very weak (have low absolute values).

Observations from SCPs computed using rsfMRI

SCPs are computed in a data-driven manner, without requiring knowledge of a “seed” region of interest. For example, SCP 1 in Figure 2.10 (from Areal Graph) and SCP 7 in Figure 2.15 (from GraSP parcels) show anti-correlation between the Dorsal Attention system and the Default Mode. This anti-correlation is a well-known finding (Fox et al., 2005), found using seed-based correlation. SCPs also capture other functional systems such as the visual system (SCP 7 in Figure 2.10 and SCP 1 in Figure 2.14), Fronto-parietal (SCP 4 in Figure 2.10 and SCP 3 in Figure 2.14) and Cingulo-Opercular (SCP 3 in Figure 2.10 and SCP 5 in Figure 2.14), along with their anti-correlations.

Not surprisingly, the SCPs computed using Areal Graph and using GraSP par-

cellation are not entirely similar. Node definition and density plays a major role in delineating whole-brain distributed patterns. The Areal graph nodes are based on a-priori knowledge of foci of functional activity. Nodes are not uniformly distributed across the cortex, there is greater nodal density in areas which are known to activate for certain tasks with visual and auditory inputs (Dosenbach et al., 2006). Therefore, it is likely that patterns computed using the Areal Graph reinforce our a-priori knowledge of functional systems. The GraSP method is data-driven, meaning that parcels are defined solely based on the underlying rsfMRI-based *geodesic distances*. This means that GraSP is agnostic to location, size and shape of the parcels (nodes). Due to different node definitions, some of the resulting SCPs computed using GraSP parcels (Figure 2.14, 2.15) are different from the Areal Graph SCPs (Figure 2.10). For example, SCP 2 (Figure 2.14) spans the medial temporal areas and SCP 8 (Figure 2.15) spans the medial dorsal frontal areas. These SCPs were not found using the Areal Graph nodes. Nevertheless, for both types of node definition, reproducibility of the SCPs is reasonably high, both within and across datasets (Figures 2.10, 2.13, 2.14 and 2.15).

Note that the SCPs presented in this chapter were computed from data generated using a pre-processing pipeline that included global signal regression. This pre-processing step is contentious, as it has been known to increase the number of negatively correlated nodes (Saad et al., 2012). We re-ran sparse learning on data that retained the global signal; however, visual inspection showed that delineation of major functional systems was poor. These results are summarized in Appendix A.4.

SCPs for group-level analyses

As the results from simulated data (Figure 2.8) and rsfMRI data (Figure 2.11) show, SCP coefficients can quantify the strength of presence of an SCP at the level of individuals. This ability provides us with a framework for group-level analyses. This advantage cannot be found in seed-based correlation methods or graph-partitioning approaches. SCP coefficient values c_i are comparable across individuals within the same SCP. Univariate SCP-wise analyses can be performed to look for differences between healthy and patient groups. SCP coefficients can be used to cluster individuals into latent sub-groups, based on connectivity strength. The joint generative-discriminative framework can also be used to find connectivity patterns that are changed in patients with a specific diagnosis, or normal development/aging-related effects.

Thus, performing dimensionality reduction by exploiting inter-individual variability facilitates group-level analyses. In addition, the availability of fairly large datasets allows investigation of more complex hypotheses - for example, aging is a highly heterogeneous process that co-exists with the onset of multiple pathological processes causing neuro-degeneration. This is accompanied by wide variation in cognitive performance in the elderly. Lower dimensional SCP coefficients along with large sample sizes allows the investigation of such a hypothesis in a data-driven manner. In the next chapter, we will describe a *Mixture of Experts* (MOE) method that models multiple patterns of change within a classification framework. The Sparse Learning framework as well as MOE was used to investigate aging using the BLSA dataset; this is described in Chapter 4.

Chapter 3

Capturing Heterogeneity using Mixture of Experts

3.1 Introduction

A growing number of projects (Shock et al., 1984; Satterthwaite et al., 2014; Van Essen et al., 2012; Biswal et al., 2010) and consortia (Di Martino et al., 2014) are collecting MR-based neuroimaging data from a large number of individuals to investigate the complex patterns of brain change associated with non-pathological and pathological processes, such as effects of development, aging, injury or disease. While MRI has been successfully used to understand functional and structural disruptions, in many studies, the main objective is to compare two groups of subjects, i.e. between normal controls and patients, or younger and older, with the assumption that the specific condition affects all subjects in a uniform fashion. In other words, each affected subject is assumed to possess the same pattern of abnormality, because only a common denominator consistently present across subjects is found. This approach conflicts with what is observed in clinical assessments, which point to

inherently multi-dimensional symptoms or cognitive changes (Ylikoski et al., 1999) that reflect a broad “spectrum” of changes associated with disease or developmental and maturational processes. Machine learning tools provide a great opportunity for investigating the heterogeneity of patterns of brain change associated with various diseases and processes, which have been for the most part ignored in previous studies.

In this chapter, we extend the analytical framework of two-group comparisons, where a diseased or otherwise affected group is compared to a relatively more normal reference group. We propose the application of a method that combines unsupervised clustering and supervised learning of classifiers to identify heterogeneity of brain changes in the affected group. Our main assumption is that the affected group was subjected to a heterogeneous underlying pathological or non-pathological process. Thus, the affected group consists of multiple subgroups, each of which has a different pattern of group differences, relative to the reference group. We assume that normal variation in the brain in the reference group evolves into potentially multiple patterns of abnormality or differences in the affected group, which are presumably caused by a variety of underlying potential pathological processes. As illustrated in Figure 3.1, in the space of multi-dimensional MRI data, the affected group “deviates” from the reference group along many different directions. We are interested in (1) capturing these heterogeneous patterns of group differences and, (2) identifying subgroups within the affected group that are associated with each pattern of group difference.

We propose the use of a Mixture-of-Experts (MOE) framework (Jacobs et al., 1991) to capture heterogeneous patterns of brain change. The MOE framework was initially proposed for vowel discrimination within speech recognition (Jacobs et al., 1991) and later, as a fast and efficient alternative to “kernel” SVMs (Ladicky

and Torr, 2011; Fu et al., 2010). The MOE method combines unsupervised clustering with supervised classification to approximate the non-linear boundary that separates the two classes with a piece-wise linear separating boundary. Thus, it provides us the identification of the subgroups as well as the multivariate patterns that discriminate each subgroup from the reference group. The data is modeled using a mixture of distributions, by assigning a soft subgroup membership to each subject in the affected group. The linear boundary between each affected subgroup and the reference group can be found using a linear classifier, such as a linear Support Vector Machine (linear-SVM). We describe the MOE method in detail in the methods section. We thoroughly validate the MOE method using multiple simulation cases and four validation measures used to quantify its performance; these results follow the methods section.

In the following section we describe the MOE formulation, the optimization strategy and the model validation steps. The validation of the performance of the method using simulated data is described in Section 3.4. We discuss the advantages and limitations of our method in Section 3.5.

3.2 Literature review

Prior to the application of MOE in this study, there were few studies that aimed to capture heterogeneity using machine learning tools. For example, in Song et al. (2010), the authors first propose clustering the subjects within each group, followed by supervised learning. In Sabuncu et al. (2009), the authors propose a joint clustering-coregistration algorithm, which can compute data-driven templates that summarize the different modes in the population. It may be possible to identify heterogeneity in the affected group by applying such a purely unsupervised cluster-

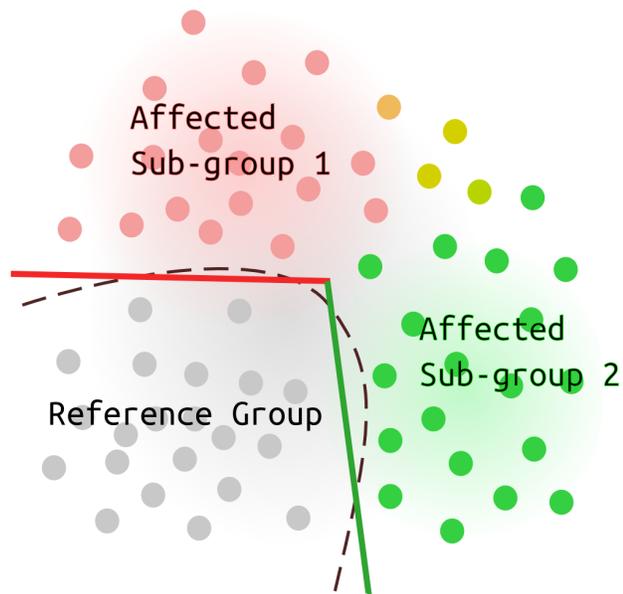


Figure 3.1: An illustration that shows heterogeneity in the affected group, relative to a reference group. Non-linear classifiers can implicitly model non-linearity but estimating the boundary (dashed lines) in high dimensional spaces is difficult. Mixture-of-Experts can approximate the non-linear curve with a piece-wise linear boundary (red and green lines) and find subgroups associated with each line (red and green points).

ing method to the affected group alone. However, in medical imaging data, two-group differences are often small and subtle, and can be nearly orthogonal to the dominant direction of variance. Running a purely unsupervised clustering method produces clusters along this direction, which may not be relevant to the problem we are attempting to solve - to find heterogeneity in the discriminating boundary that reflects the underlying pathologic process. Hence our attempt to solve this problem involves using the reference group as an “anchor”. In other words, we assume that the reference group is transformed by an underlying process which affects the reference group in multiple different ways. We would like to identify these multiple directions of deviation from the reference group. In our proposed method, allowing the reference group to be equally shared amongst all sub-groups is an indirect manner of modeling this deviation from the reference group.

While kernel SVMs can also successfully model non-linear separation boundaries between groups, such as the one shown in Figure 1, they suffer from a major limitation in neuroimaging applications, namely the lack of interpretability of the results. In a kernel-based method, the data is implicitly projected into a higher dimensional space prior to being classified and the non-linear separating boundary in the original feature space is not explicitly computed. This limitation was tackled in multiple papers (Golland, 2001; Fan et al., 2007; Rasmussen et al., 2011). Golland (2001) find reflections of each support vector of the fitted kernel-SVM on the other side of the separating boundary. The difference between the support vector and its reflection is the local discriminant direction in the feature space, which is interpreted in terms of the changes to the input data. Similarly, in COMPARE (Fan et al., 2007), Fan et. al average these local discriminant directions across all support vectors, resulting in a single group difference map. In Rasmussen et al. (2011), the importance of each feature for the classifier is estimated by computing the average

extent to which predictions get perturbed when the data points are perturbed. This results in a single sensitivity map for the non-linear classifier. However, none of these approaches targeted the identification of subgroups in the population based on the discriminant direction of change in comparison to the reference group. An easily interpretable and not overly complex way of representing this heterogeneity is necessary for the clinical adoption of such methods.

Concurrently with MOE, other methods were proposed for the same objective, but had different approaches. In Dong et al. (2015), the authors proposed the use of a purely generative method where the reference and affected groups are modeled as distributions in a high-dimensional space, such that the multiple disease effects are different point cloud transformations from the reference to the affected groups. In Varol et al. (2015), authors proposed the use of a maximum margin polytope method to identify variations in disease using T1-structural data. This formulation uses maximal margin classification objectives only, and hence is purely discriminative; it relies on the distance to the hyperplane to determine the grouping of the subjects. The proposed MOE approach is a well-established approach which combines both - an expert (classifier/regressor) with a mixture model (clustering), and hence falls under the umbrella of generative-discriminative methods. Thus, MOE not only models the distance from the hyperplane but also the natural clustering that may be present within the affected group, for identification of heterogeneity.

3.3 Mixture-of-Experts: Formulation and optimization

Consider a binary classification problem with data $\mathbf{x}_i \in \mathbf{R}^D$ obtained from $i = 1, 2, \dots, N$ subjects. Each subject is associated with a binary label $y_i \in \{-1, 1\}$, -1 for the reference group and $+1$ for the affected group. We assume that the dis-

criminative direction is not constant across the feature space. In other words, the group difference is heterogeneous due to multiple processes that might affect brain structure and function in different ways. This heterogeneity can be modeled using multiple piece-wise linear hyperplanes. Our objective is to learn the multiple discriminant patterns of abnormality along with subgroups of affected subjects corresponding to each pattern. We propose to model this heterogeneity with a piece-wise linear boundary with K segments. Each segment is a hyperplane \mathbf{w}^k , which is interpretable in terms of the discriminative/affected features in each subgroup k .

3.3.1 The expert model

Let $\mathbf{m}_i = [m_i^1, m_i^2, \dots, m_i^K]$, $m_i \in [0, 1]$, $\sum_{k=1}^K m_i^k = 1$ indicate the relative membership of subject i to group k . Recall that along the discriminative direction, subjects deviate in multiple different directions due to underlying process, away from the reference group. Therefore we associate the reference subjects with all K hyperplanes, i.e., if $y_i = -1$, $m_i^k = 1/K \forall k$. If the membership values were known for all subjects, the k^{th} linear-SVM hyperplane $\mathbf{w}^k \in \mathbf{R}^D$ can be learned by solving the following optimization problem (Bishop et al., 2006):

$$\underset{\mathbf{w}^k}{\text{minimize}} \frac{1}{2} \|\mathbf{w}^k\|_2 + C \sum_{i=1}^N m_i^k (1 - y_i (\mathbf{w}^k)^T \mathbf{x}_i)_+ \quad (3.1)$$

where $x_+ = \max(0, x)$.

The above optimization problem is the standard formulation of the ℓ_1 -loss, ℓ_2 -regularized SVM in its primal form, with the membership values m_i^k acting as sample weights. The user-defined SVM cost parameter C controls the extent to which misclassified points are penalized. Note that the intercept of the SVM hyperplane has been subsumed into the variable \mathbf{w}^k by appending a constant value to all data

points.

3.3.2 The mixture model

The unknown group membership values m_i can be obtained by jointly optimizing the SVM objective function (above) with a data clustering objective. In this paper, we learn these subgroups using Fuzzy-C-Means (Bezdek et al., 1984). Each one of K subgroups is associated with a centroid \mathbf{d}^k . The mixture model is formulated as an optimization problem, as follows:

$$\begin{aligned}
 & \underset{\{\mathbf{d}^k\}_k, \{m_i^k\}_{i,k}}{\text{minimize}} && \sum_{k=1}^K \sum_{i=1}^N (m_i^k)^a \|\mathbf{x}_i - \mathbf{d}^k\|_F^2 \\
 & \text{subject to} && \\
 & && \sum_{k=1}^K m_i^k = 1, \quad m_i^k \in [0, 1] \quad n = 1, \dots, N
 \end{aligned} \tag{3.2}$$

As explained in (Bezdek et al., 1984), the “fuzziness” coefficient m controls the “softness” of the membership assignments. When $m = 1$, the memberships are binary. As $m \rightarrow \infty$, the membership values tend to $1/K$, where K is the number of sub-groups.

In our model, we set the “fuzziness” coefficient a to a value of 2. We make this choice because when the cluster centroids (and other variables) are known, for $a = 2$, the minimization problem becomes a quadratic programming problem, which is convex.

3.3.3 The joint model

Bringing the optimization problems 3.1 and 3.2 together, we get

$$\begin{aligned} & \underset{\{\mathbf{w}^k\}_k, \{m_i^k\}_{i,k}}{\text{minimize}} \quad \sum_{k=1}^K \left\{ \frac{1}{2} \|\mathbf{w}^k\|_2 + C \sum_{i=1}^N m_i^k (1 - y_i (\mathbf{w}^k)^T \mathbf{x}_i)_+ + \lambda \sum_{i=1}^N (m_i^k)^a \|\mathbf{x}_i - \mathbf{d}^k\|_F^2 \right\} \\ & \text{subject to} \\ & \quad \sum_{k=1}^K m_i^k = 1, \quad m_i^k \in [0, 1] \quad n = 1, \dots, N \end{aligned} \quad (3.3)$$

where λ is a user-defined value that controls the trade-off between the cost of classification and clustering. The other user-defined parameters are the number of groups K , and the SVM cost-value C .

3.3.4 Optimization strategy

We use alternating minimization to solve for the cluster centroids \mathbf{d}^k , membership values m_i^k and the SVM hyperplanes \mathbf{w}^k . Note that, from the joint optimization problem in Eq. 3.3, only the membership values m_i^k are common to both the classification and the clustering problems. Knowing these values would allow us to decompose the joint optimization problem into $K + 1$ smaller problems, which are all convex: K weighted-SVM objectives with weighted samples, and solving for cluster centroids. We use libSVM with weighted samples (Fan et al., 2008) to solve each of the K SVM objectives, and the cluster centroids can be updated as follows:

$$\mathbf{d}^k = \frac{\sum_{i=1}^N (m_i^k)^a \mathbf{x}_i}{\sum_{i=1}^N (m_i^k)^a} \quad (3.4)$$

When the values of variables \mathbf{w}^k and \mathbf{d}^k are known, the joint problem can be decomposed into N optimization problems, one for each subject. Each of these N

problem can be solved for the membership values $m_i^1, m_i^2, \dots, m_i^K$ of subject i , as follows:

$$\begin{aligned} & \underset{\{m_i^k\}_k}{\text{minimize}} \sum_{k=1}^K \left\{ C e_i^k m_i^k + \lambda \|\mathbf{x}_i - \mathbf{d}^k\|_F^2 (m_i^k)^a \right\} \\ & \text{subject to} \end{aligned} \tag{3.5}$$

$$\sum_{k=1}^K m_i^k = 1, \quad m_i^k \in [0, 1]$$

where e_i^k is the slack value ($= 1 - y_i(\mathbf{w}^k)^T \mathbf{x}_i$) for the i^{th} individual in the k^{th} SVM. When $a = 2$, the objective function is quadratic and convex, and the constraints are linear.

The optimization strategy alternately solves for \mathbf{d}^k , m_i^k and \mathbf{w}^k until convergence. Convexity of each of the sub-problems guarantees its convergence to a local minimum.

3.3.5 Prediction of test cases

We use ten-fold cross-validation to evaluate the fit of the model to the data. Given a test subject \mathbf{x}^* with an unknown label y^* , first the unknown membership values \mathbf{m}^* are obtained by solving the clustering objective alone with \mathbf{x}^* as the data. Then the label y^* is determined as the sign of the weighted combination of individual SVM predictions as follows:

$$y^* = \text{sign} \left(\sum_{k=1}^K m_k^* \text{sign}((\mathbf{w}^k)^T \mathbf{x}^*) \right) \tag{3.6}$$

3.3.6 Model selection

We use four summary measures to quantify the performance of the method:

1. Cross-validated accuracy: Estimated labels are compared with the known labels of the left out fold. This testing is repeated for all folds for multiple runs to obtain an averaged cross-validation accuracy.
2. Maximum pair-wise inner product: For $K > 1$, we compute the maximum normalized inner product between all pairs of hyperplanes, as follows:

$$r_{\mathbf{w}} = \max \left\{ \frac{(\mathbf{w}^k)^T \mathbf{w}^l}{\|\mathbf{w}^k\|_2 \|\mathbf{w}^l\|_2}, \forall k, l \in \{1, 2, \dots, K\} \right\} \quad (3.7)$$

This value measures the minimum extent to which the hyperplanes have rotated away from each other.

3. Cluster Reproducibility: For $K > 1$, we evaluate the reproducibility of the subgroups across repeated runs of the proposed method. We use the Adjusted-Rand Index (ARI) for fuzzy cluster assignments, as defined in Brouwer (2009). The ARI is a scalar value between $[-1, 1]$ which measures the extent to which two fuzzy cluster assignments are similar, after adjusting for chance. An ARI value of $+1$ denotes perfect reproducibility, 0 indicates that some subjects have the same fuzzy membership solely due to chance and -1 indicates disagreement among all pairs of memberships.
4. Cluster Separation Index: For $K > 1$, we evaluate the extent to which the K clusters are separated across repeated runs of the proposed method. We use the Bezdek Partition Coefficient (BPC) (Bezdek, 1981; Dave, 1996) which provides a scalar value between $[0, 1]$ for each fuzzy clustering assignment. A value of 1 indicates full cluster separation, i.e., cluster assignments for all subjects are binary. A value of 0 indicates no separation, i.e., cluster assignments for all subjects are equal to $1/K$, where K is the number of clusters.

The user-defined parameters of the MOE method are the number of subgroups K , the SVM cost value C and the classification-clustering trade-off parameter λ . We use a grid-based search to find those parameters for which the four above mentioned measures (accuracy, maximum inner-product, cluster reproducibility, cluster separation) are optimal.

3.4 Experiments on simulated data

3.4.1 Simulated data

We evaluated the ability of the proposed method in revealing underlying subgroups using four two-dimensional simulated datasets. These datasets reflected different heterogeneity patterns in the affected subgroups, defined as the subgroups deviating from the reference group:

Case 1: One affected subgroup with a single pattern of change relative to a reference group. Each group is modeled by an isotropic Gaussian distribution.

Case 2: Two affected subgroups, with heterogeneous patterns of change relative to a reference group. Each subgroup is modeled by an isotropic Gaussian distribution.

Case 3: Multiple affected subgroups, with heterogeneous patterns of change relative to a reference group. This is modeled using concentric circles

Case 4: Multiple affected subgroups, with heterogeneous patterns of change relative to a reference group. This is modeled using arcs of concentric circles

For each case, 200 points were simulated, with 100 affected and 100 reference subjects. In each of the four cases above, 20% of the data points were deliberately mis-labeled.

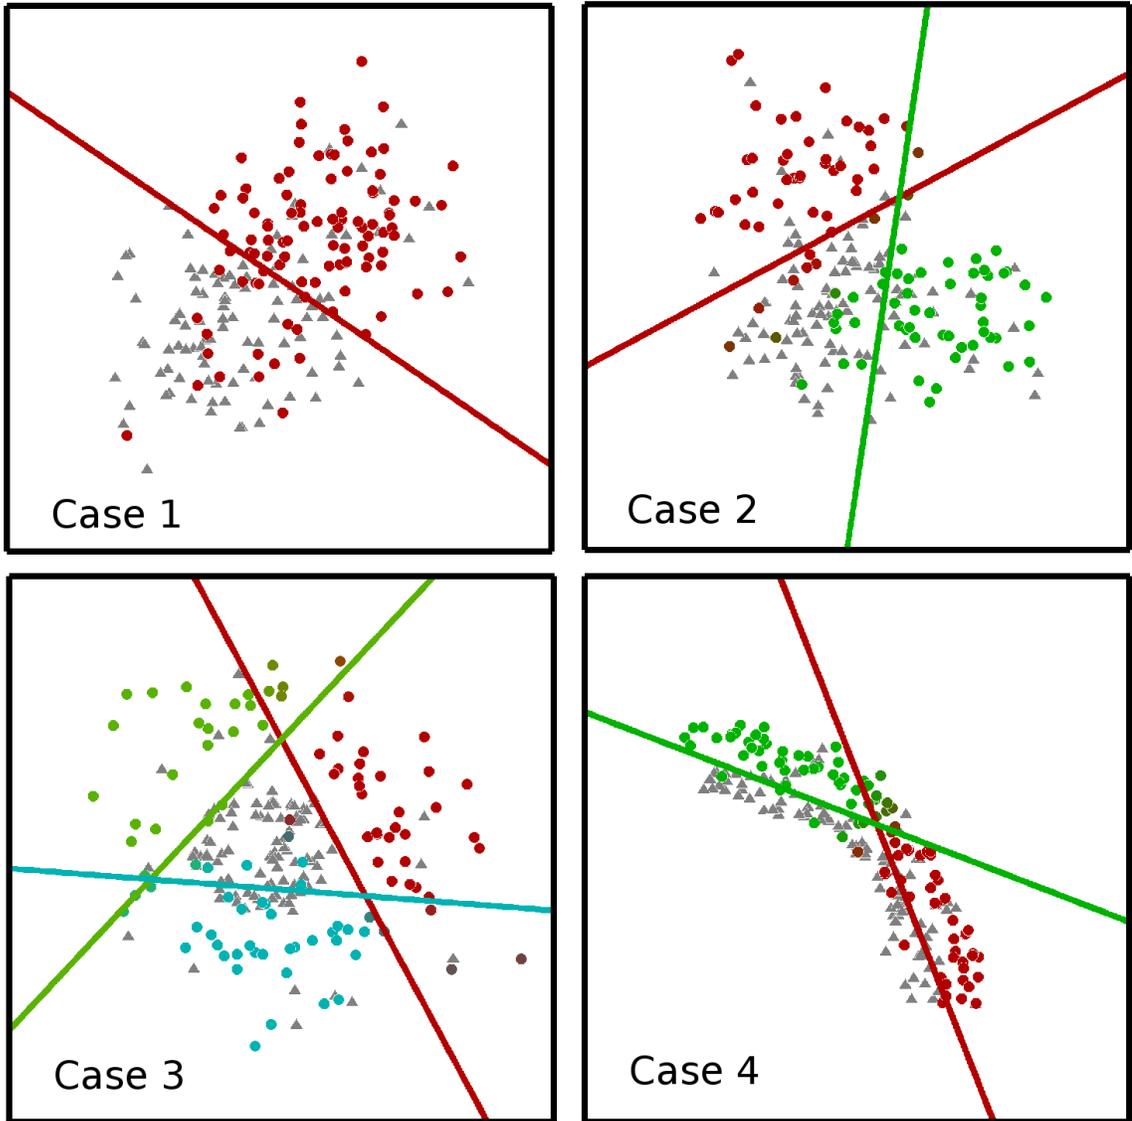


Figure 3.2: Four 2-D simulated cases used to evaluate the performance of the method. Hyper-planes and sub-groups obtained using MOE are also shown using a different color for each hyper-plane and associated subgroup.

3.4.2 MOE choice of parameters and results

We measured (1) ten-fold cross-validation accuracy (2) maximum normalized inner-product between the resulting $\binom{K}{2}$ pairs of hyperplanes (3) reproducibility of the resulting affected subgroups, measured using ARI and (4) separation of the resulting affected subgroups, measured using BPC. These results are shown in Figures 3.3 and 3.4. We expect a good model to provide high cross-validated accuracy, low inner-product (large angle between hyperplanes), high reproducibility of the resulting subgroups across runs, and high separability between subgroups.

From Figure 3.3, we observe that highest accuracy values are obtained for K values of 2, 3, and 2 for the cases 2, 3, and 4 respectively. For the same three values, the inner-product between the hyper-planes is at a minimum, suggesting that the hyperplanes have rotated to their maximum extent. Cluster reproducibility and separation are also high. Note that in the absence of heterogeneity, as simulated in Case 1, the fuzziness of the clustering algorithm prevents the creation of spurious noise-based clusters; the membership value of each affected subject will be close to $1/K$ and consequently, the resulting hyperplanes are almost identical, and cluster separation is low, as seen in Case 1.

For Case 3, and for $K = 3$, we used a grid-based search to evaluate the three measures for $C = \{2^{-3}, 2^{-2}, \dots, 2^{10}\}$ and $\lambda = \{2^{-3}, 2^{-2}, \dots, 2^{10}\}$. In all four plots, the parameter space is clearly split into two: $C \geq \lambda - 3$ and $C < \lambda - 3$. When $C \geq \lambda - 3$, the inner-product measure is low and the accuracy, cluster reproducibility and separation indices are high. Considering all these observations together, the values $K = 2$, $C = 2^{10}$ and $\lambda = 2^{-3}$ seem reasonable, with ten-fold cross-validation accuracy at $78 \pm 0.8\%$, normalized inner-product at 0.5 ± 0.01 , ARI at 0.75 ± 0.08 and BPC at 1.

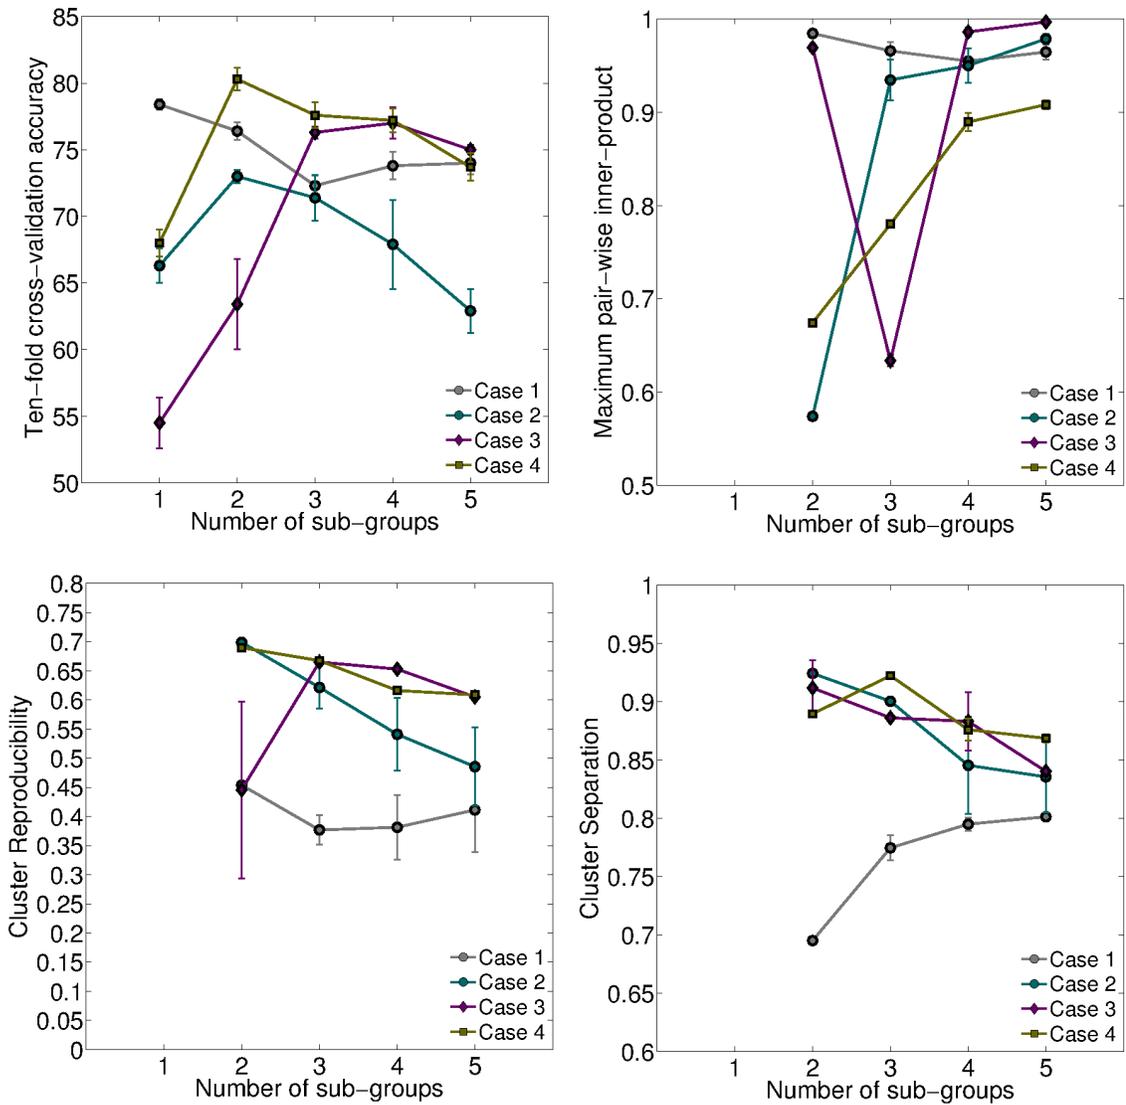


Figure 3.3: Variation in four performance measures for $K = \{1, 2, \dots, 5\}$, for each of the four simulated cases. From top-left, clockwise: Accuracy, Maximum inner-product, cluster separation and cluster reproducibility. Results from each case is plotted in a different color.

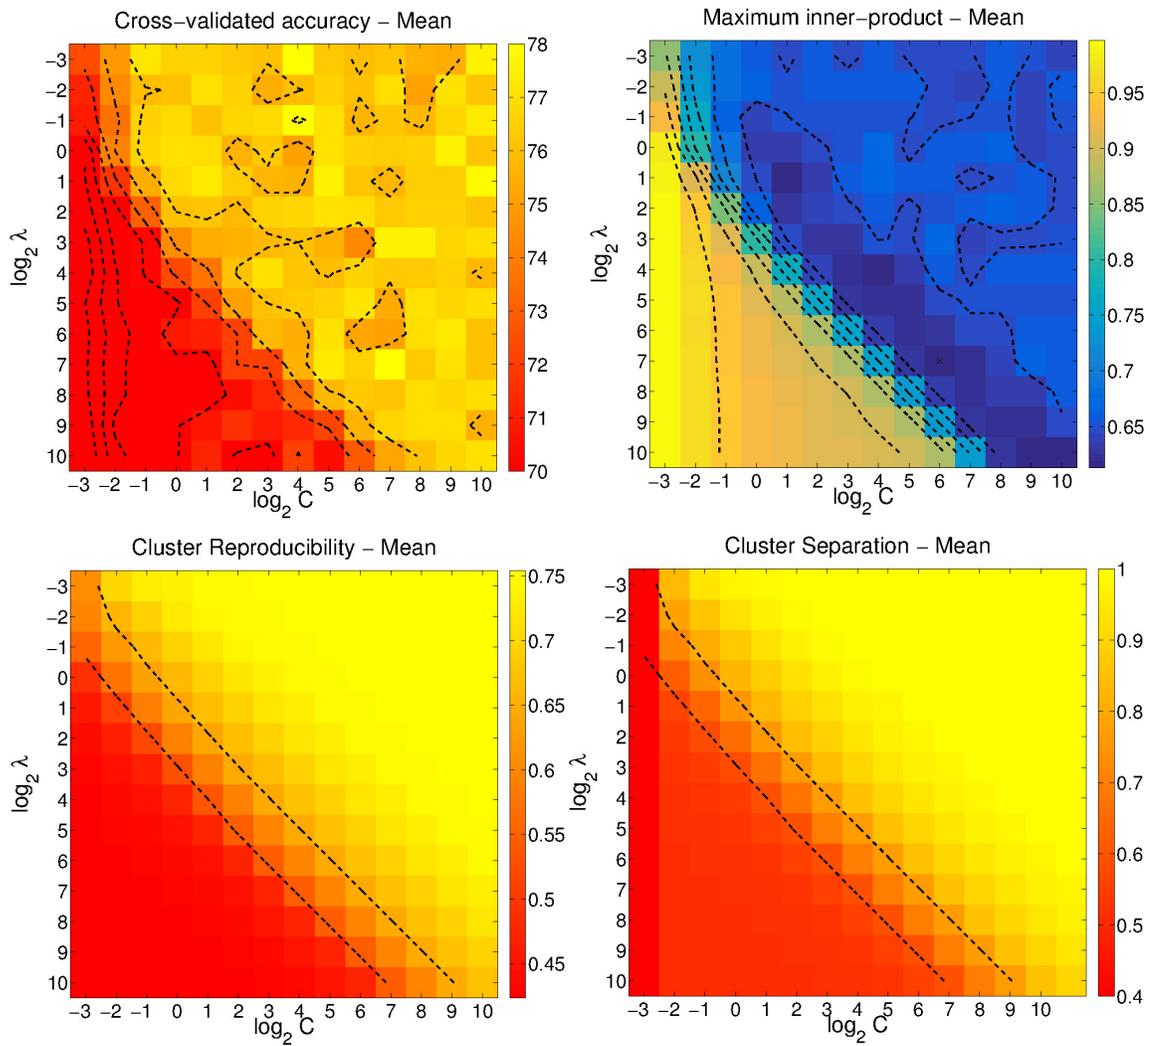


Figure 3.4: Variation in four performance measures for $K = 3$, $C, \lambda = \{2^{-3}, 2^{-2}, \dots, 2^{10}\}$, for Case 3. From top-left, clockwise: Accuracy, Maximum inner-product, cluster separation and cluster reproducibility.

Simulation Case	1	2	3	4
Gaussian-kernel SVM	79.0 ± 0.54	75.4 ± 1.08	86 ± 0.01	85.3 ± 0.27
MOE	79.0 ± 0.22	74.1 ± 0.74	77.9 ± 0.80	82.5 ± 1.35

Table 3.1: Table comparing ten-fold cross-validation accuracy for MOE method vs. Gaussian-kernel SVM, for four 2-D simulated cases.

Table 3.1 shows the ten-fold cross-validation accuracy of MOE model compared with that of an RBF-kernel SVM for each of the four simulated cases. In Cases 2-4, the accuracy of the MOE model is slightly lower than the non-linear model. This difference in accuracy is highest for Case 3, where the non-linearity is the greatest among all four cases. This decrease is expected, as a piece-wise linear model is used to approximate the non-linearity. The MOE model performed with a small drop in accuracy, when compared to non-linear models, while providing interpretable results and explicit assignment into a small number of subgroups. Such interpretability is very important for adoption of such methods in clinical environments.

3.5 Discussion

In this chapter, we proposed the use of a Mixture-of-Experts framework to capture diverse disease or degeneration patterns. This is a general framework that can be applied to any type of data (structural or functional MR data) using any expert (classification or regression-based) and any mixture model (K-Means, fully Gaussian mixtures and others). We used a linear-SVM along with fuzzy c-means clustering to identify multiple subgroups in the heterogeneous population, along with the associated abnormal connectivity pattern for each subgroup. We evaluated the resulting models based on four performance measures - cross-validated accuracy, maximum hyperplane inner-product, cluster reproducibility and cluster separation.

We tested the performance of the method using multiple 2-D simulation cases.

As we stated earlier, multiple linear SVMs are used to approximate the non-linear boundary between the two groups, hence there is some loss of accuracy. This loss in accuracy is traded for a richer description of the data in terms of multiple subgroups and linear hyperplanes associated with each subgroup, that are interpretable in terms of changes to the underlying features. Furthermore, as fuzzy membership values are used, in the absence of heterogeneity in the data, the fuzziness of the model prevents the generation of noisy clusters.

A methodological limitation to consider is that as the MOE model combines classification with a clustering objective, it inherits the drawbacks of clustering as well. High dimensional data, such as vectorized correlation matrices, is difficult to cluster due to the well known “curse of dimensionality” (Steinbach et al., 2004). The effectiveness of applying our method for high dimensional data needs to be investigated further. The Sparse Learning method used in this paper provides an interpretable connectivity-based set of SCP bases, while reducing data dimension, making it applicable for MOE analysis.

Another limitation to consider is the trade-off between higher accuracy and better interpretability. The MOE classifier can only *approximate* a non-linear boundary. Therefore, in terms of accuracy, it can only perform as well as the non-linear classifier, and no more. In fact, in both simulated and real rsfMRI data, the MOE classifier slightly under performs the non-linear classifier. Our motivation behind proposing the MOE classifier as an alternative is for the sake of interpretability, to discover heterogeneous two-group differences.

Our method primarily looks for variation in discriminating direction, and may not directly relate to disease/aging severity. Individuals that are further away from the hyperplane have greater magnitude of changes relative to the reference group, therefore has a more severe disease/aging effect. Thus severity can be estimated by

calculating the distance of each subject from the discriminating hyperplane. Such an approach has been used in prior studies to quantify severity of the disorder, such as Alzhiemers' and Mild Cognitive Impairment (Davatzikos et al., 2009; Clark et al., 2012) and Autism Spectrum Disorder (Ingalhalikar et al., 2011). Currently, MOE provides information about heterogeneity alone. In addition to the sub-group memberships, severity along each heterogenous aging direction is also potentially very useful; this will be investigated in the future.

A possible extension of this method is to incorporate heterogeneity in the reference group as well, considering that the variability in a typical reference “normal” group can be quite high. Using such an extension, each estimated affected group is associated with its closest set of reference subjects. This is advantageous, as the associated hyperplane weights and associated two-group p-values are more specific to the disease/aging heterogeneity. On the other hand, the hyperplanes are no longer comparable across sub-groups, as the reference group is different for each affected sub-group. The advantages and trade-offs of such an extension need to be investigated further.

In the next chapter we apply MOE to rsfMRI data acquired from normal older adults as a part of the BLSA study. Using SCP coefficients as well as estimates of GM density within each parcel, we define functional and structural brain aging trajectories. We use these trajectories to identify individuals whose imaging data suggests that their brain age is higher than expected; i.e., they are *advanced agers*. Similarly we identify individuals whose brain age is lower than expected; we refer to them as *resilient agers*. Using resilient agers as a reference, we apply the MOE method to identify heterogeneous patterns of advanced aging. Combined with cognitive data, we further characterize these sub-groups of advanced agers.

Chapter 4

Analyzing the effect of aging on functional connectivity

4.1 Introduction

The number of older people in the United States, who currently constitute nearly 15% of the population, is projected to increase dramatically over the next two decades. Normal aging is associated with increased risk of physical sickness and injury, declining mental health and greater risk of developing dementia. It reduces the quality of life and normal day-to-day functioning in older people. Therefore, it is imperative that we understand the effects of normal aging on brain structure and function.

Age is known to be significant factor that determines inter-individual variability in functional connectivity (Biswal et al., 2010). This makes rsfMRI an appealing non-invasive method to understand functional disruptions due to aging. To this end, we used MRI data from 400 cognitively normal participants of the BLSA study to obtain functional correlates of aging. In addition, we used multi-variate regression

to build a normative brain aging trajectory that captures an average snapshot of aging related changes. Using this trajectory, we computed a personalized Brain Aging Index (BAI) of functional brain health for each individual in the study.

It is also important to note that aging is probably a heterogeneous process, given the large amount of variance present in cognitive performance of older individuals, even though they are cognitively normal. As Rowe and Kahn (1987) observed in their paper, using the term “normal” to describe these individuals connotes lack of risk towards aging-related pathology. Although normal, it is possible that multiple pathologic processes could have initiated a diverse set of functional and structural changes in the brain, without showing external signs or symptoms of cognitive decline. As we describe in the second half of this chapter, we used both functional and structural MRI data from the BLSA to identify individuals who show advanced signs of aging, as well as those that seem resilient to it (“successful” aging). Furthermore, we identify heterogeneous patterns of functional and structural change in advanced aging relative to resilient aging. We think that some of these advanced agers are at a greater risk for dementia in the future.

The following section reviews results obtained from prior imaging studies of aging. Subsequent sections report and discuss results from our analysis of the BLSA dataset.

4.2 Literature review

There is a large body of literature investigating aging effects on resting state functional connectivity (Ferreira and Busatto, 2013) using univariate statistics. Of these, reduced connectivity between regions of the Default Mode (DM) has been most consistently reported. In (Wu et al., 2011) and Andrews-Hanna et al. (2007), the

authors used the Posterior Cingulate Cortex (PCC) as seed to delineate other regions of the Default Mode, and investigate changes in its connectivity with age. Reduced connectivity was also observed in Dorsal Attention (DA) system, which is anti-correlated with the DM (Tomasi and Volkow, 2012), as well as in the motor cortex (Langan et al., 2010). When comparing older adults with younger adults, the authors in (Wu et al., 2007) found that the functional connectivity of the motor network is reduced in the rest state. Significant differences were also observed in functional connectivity of the visual cortex between young and older individuals (Yan et al., 2011).

There are far fewer studies that apply multi-variate machine learning based approaches to understand functional disruptions due to aging. In Meier et al. (2012), the authors used vectorized correlation matrices as input to a Support Vector Machine (SVM), which was used to classify young adults from older adults. The authors applied a significant amount of post-processing to the resulting hyperplane weights to discover that connections within the motor and cingulo-opercular areas played a large role in classification. Similarly, in Vergun et al. (2013), the authors applied Support Vector Regression (SVR) to vectorized correlation matrices to predict age of individuals in the age range 19-85 years.

In this study, we investigated the effect of aging on functional connectivity by looking for significant correlations between SCP coefficients computed from the data and age. We used imaging features from both functional and structural MRI as input to multivariate regression in order to build brain aging trajectories. By comparing an individual's expected brain age from their predicted age, we can identify individuals who are advanced agers relative to their chronological age. (A similar approach was used in Franke et al. (2012) for individuals with Mild Cognitive Impairment (MCI) in order to identify those who are most at risk to progress to

Alzheimer’s disease.) In the following sections, we describe the structural and functional brain phenotypes associated with normal and advanced aging.

4.3 Effect of aging on functional connectivity

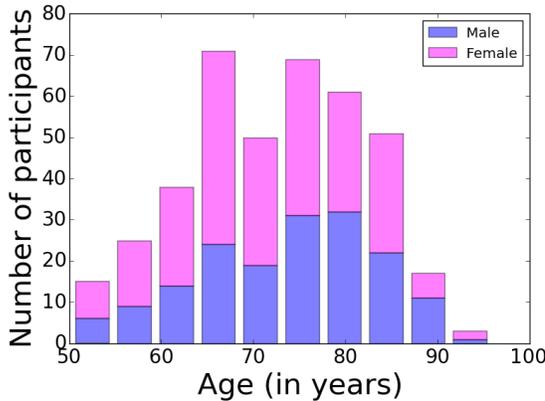


Figure 4.1: Age distribution for the subjects used in this study. For each age bin, counts for male and female participants are shown in different colors.

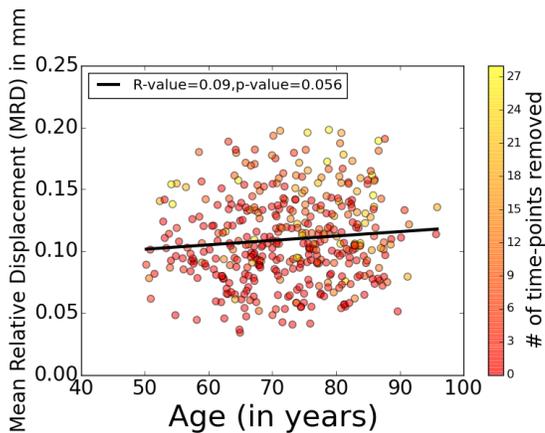


Figure 4.2: Scatter plot showing the relationship between age and motion (measured using MRD), which is statistically insignificant. The color indicates the number of time-points that were removed during the scrubbing procedure.

We used BLSA first-visit data from 400 participants; their age and sex distribution is shown in Figure 4.1. Age is not significantly correlated with motion in this study, see Figure 4.2. For details about rsfMRI pre-processing, see Appendix A.

Using the Sparse Learning framework we computed SCPs in a hierarchical manner with 10 primary SCPs and 10×50 secondary SCPs. These measurements were RI analyses.

Recall that SCP coefficients are subject-specific; they are proportional to the average connectivity between the regions that belong to the SCP for a particular subject. For each of the primary and secondary SCPs, we evaluated the correlation and associated p-value between each of the coefficients and individual ages. P-values are corrected for multiple comparisons using Benjamini-Hochberg False

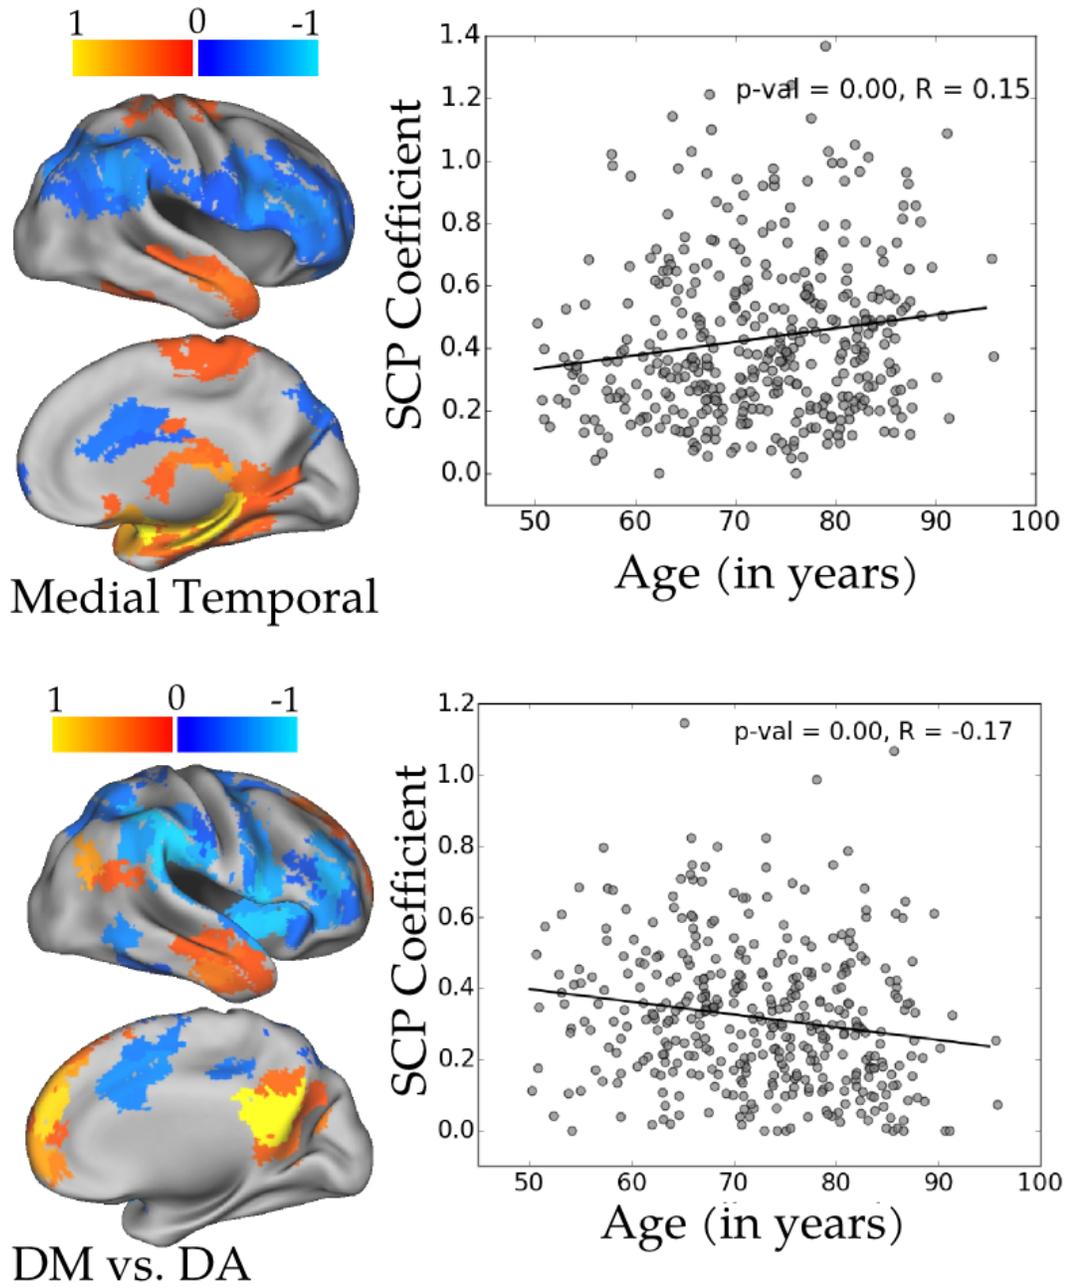


Figure 4.3: (Above) Connectivity within the DM vs. DA SCP is significantly decreased with age. (Below) Connectivity within the medial temporal areas is significantly increased with age

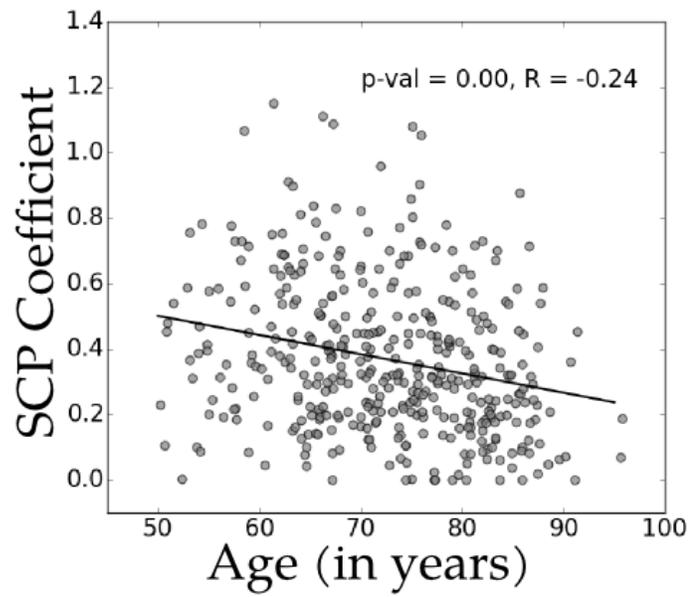
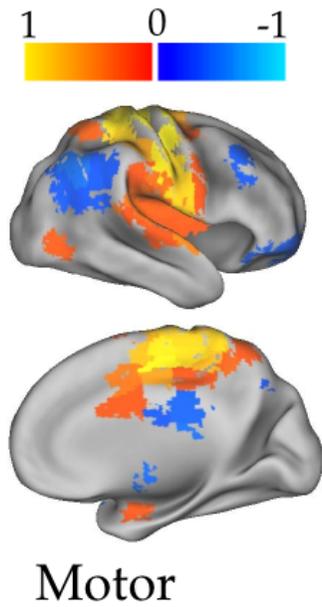
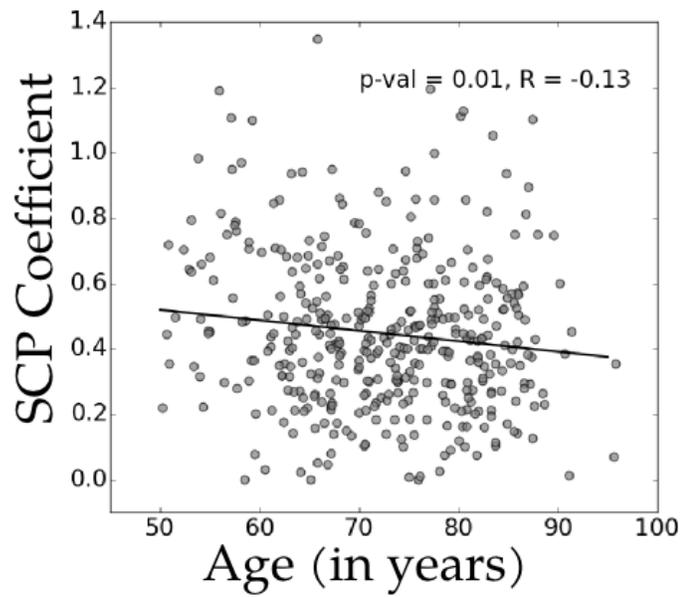
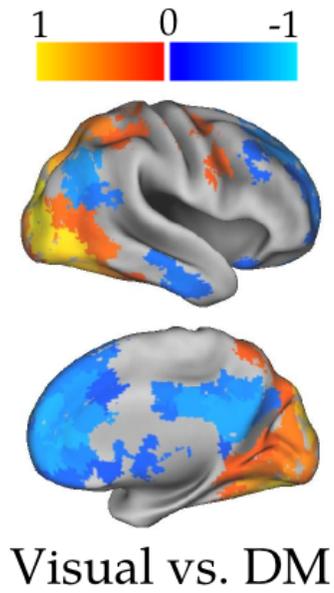


Figure 4.4: Connectivity within the visual vs. DM areas (above) and motor areas (below) is significantly decreased with age.

(Benjamini and Hochberg, 1995).

We found that four of the ten primary SCPs showed significant correlation with age ($p < 0.05$, corrected). Figure 4.3 shows the DM vs. DA SCP, whose average connectivity is significantly reduced with age. We found that connectivity within the motor region, as well as the visual areas is reduced (Figure 4.4). The only primary SCP to show increased connectivity with age involves mainly the medial temporal regions (with weak anti-correlation in the frontal and parietal areas).

Associated secondary SCPs reflect the trends in the primary SCP, but provide greater spatial specificity. For example, eight of the secondary SCPs associated with the primary motor SCP (shown in Figure 4.4) show significant negative correlation with age. Each of these SCPs is bilateral, spanning from the operculum to the most superior motor regions. Many other secondary SCPs are also significantly correlated with age.

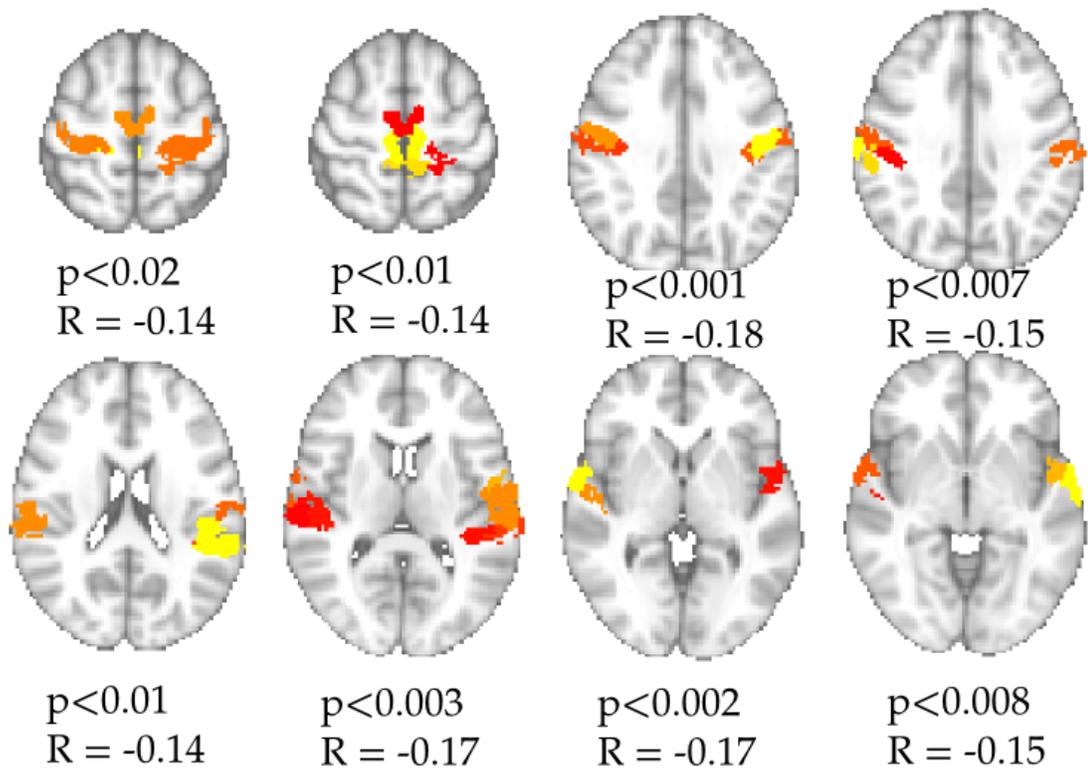


Figure 4.5: Eight secondary SCPs of the motor SCP, correlation between corresponding coefficients and age, and associated p-value.

4.4 Functional brain aging trajectory

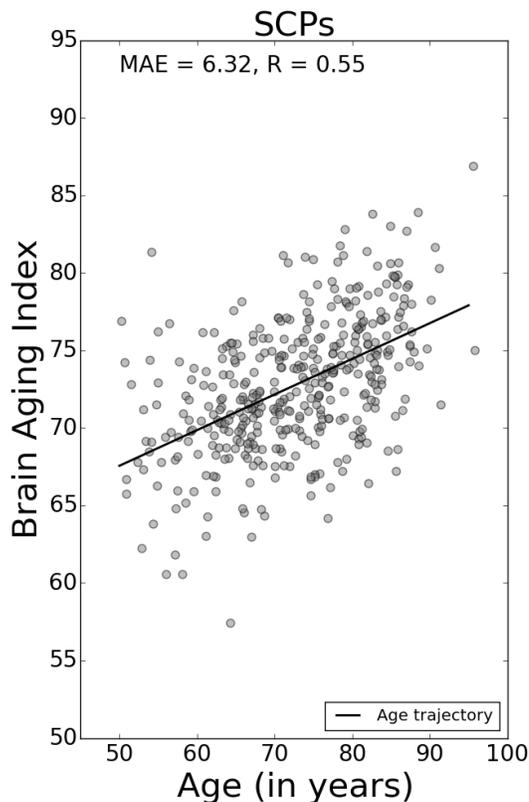


Figure 4.6: Brain Age Index of each individual plotted against their age. The Mean Absolute Error (MAE) of prediction, as well as the correlation coefficient between the BAI and age is provided.

For each participant, the SCP coefficients for all 510 SCPs form a feature “vector” which was input to the Support Vector Regression (SVR) method. SVR is a widely used supervised learning algorithm that generalizes Support Vector Machines (SVM) to continuous outcome variables. Within a ten-fold cross-validation framework, we used SVR to predict an individual’s age using the SCP coefficient vector. The SVR predicted age of each individual summarizes the SCP data into one value, while preserving the variance contributed by age. We call this value the “Brain Aging Index”(BAI) which tracks the changes in functional connectivity with age. The BAI of all 400 participants is plotted against their age in Figure

4.6. Individuals that are above the trajectory are those whose patterns of functional connectivity changes are worse than normal. Similarly, individuals who are below the trajectory are aging better than expected, with respect to their brain function.

We ran permutation tests to find the SCPs that significantly contribute to the SVR model. In each of 1000 tests, we randomly permuted the ages of all 400 individuals. An SVR was trained with the permuted ages as input along with the

original SCP features. In this manner we obtained a distribution for the SVR weight of each feature under the null hypothesis that the features are not predictive of age. By comparing the observed SVR weight with the distribution under the null hypothesis, we computed a p-value for each feature. Figure 4.7 shows all SCPs that were found to be significant. The medial temporal and motor primary SCPs play a major role in predicting the functional BAI.

4.5 Resilient and advanced aging

It is widely known that clinical dementia in the elderly is often the end stage of a long process of pathological brain changes (Morris, 2004). Onset of pathology is believed to occur perhaps decades before individuals start to present a decline in their cognitive faculties, at which point therapies become ineffective. However, brain changes with aging are complex and highly heterogeneous throughout the population, potentially reflecting the heterogeneity of the underlying pathologic processes and individuals resilience to them. Elucidating this heterogeneity, and ultimately relating it to cognitive resilience and vulnerability, is of central interest in studies of aging and early dementia (Jack et al., 2010). It is also important from the perspective of constructing bio-markers that identify older individuals at risk for clinical progression.

To be able to identify individuals who are aging better or worse than average, we used other MR-based features in addition to functional connectivity. From rsfMRI data, we used Regional Homogeneity (Zang et al., 2004) which measures the average local functional coherence within each parcel of the parcellation. Furthermore, we also incorporated average GM density values computed from corresponding T1-MRI data for each individual, for each of the parcels. Local GM density was re-

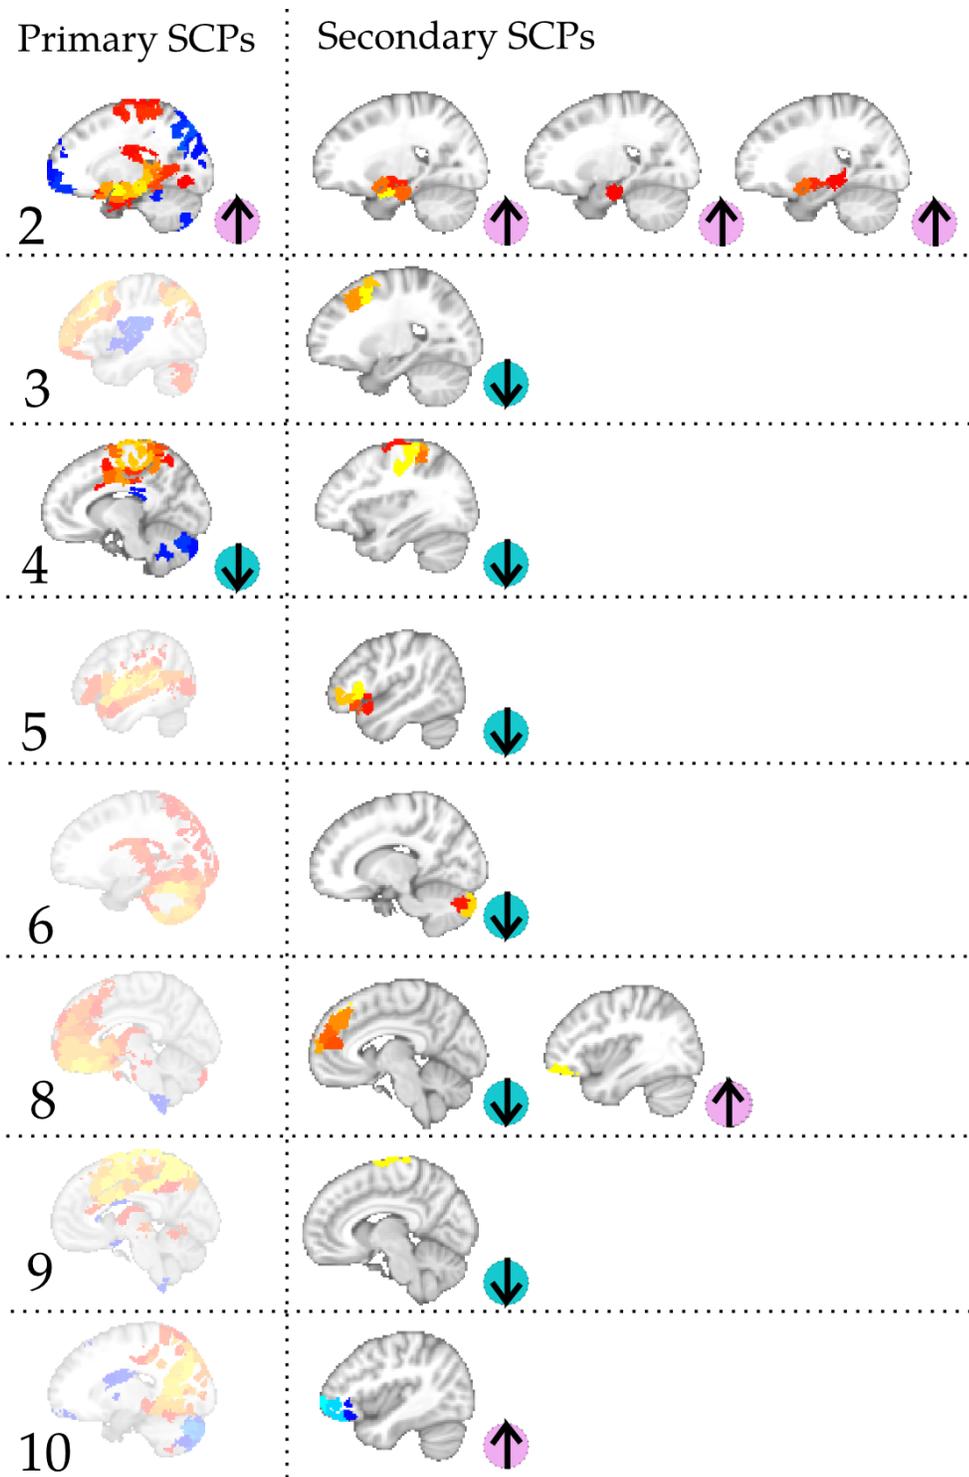


Figure 4.7: Primary and associated secondary SCPs that were found to be significant after permutation testing. Although not all primary SCPs are significant (faded images are not), they are shown for reference.

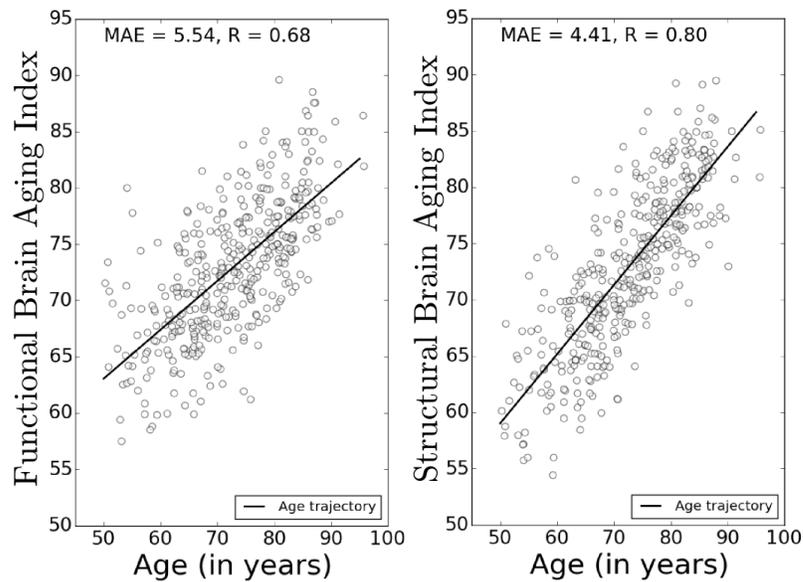


Figure 4.8: Functional and structural Brain Age Indices (BAIs) plotted against each individual's age. Functional BAI was computed from SCP coefficients and local functional coherence measures; structural BAI was computed from average GM density values. The aging trajectories (solid lines) show the expected brain age. The BAI Residual for each individual in each modality can be computed by subtracting their expected brain age from their BAI.

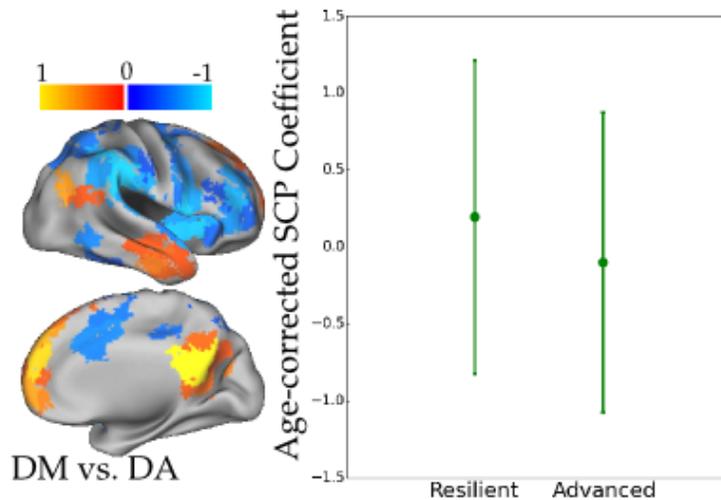


Figure 4.9: The DM vs. DA SCP alone shows significantly decreased connectivity in the advanced agers relative to the resilient agers after correcting for age.

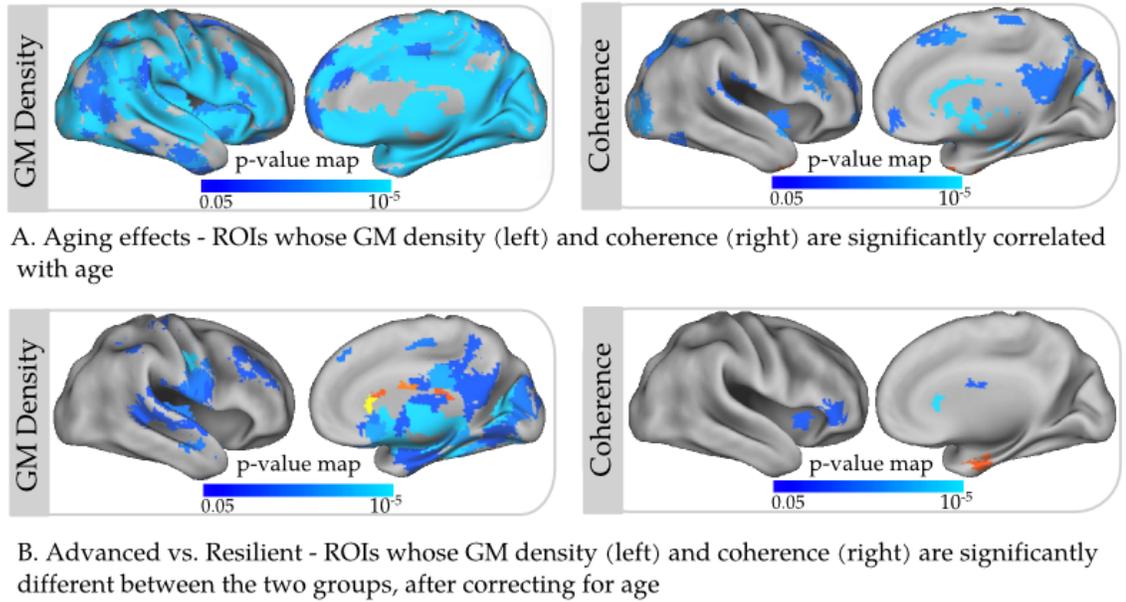


Figure 4.10: Univariate differences in local functional coherence and GM density between resilient and advanced agers.

gressed out of the local coherence estimates.

Since aging-related structural and functional brain changes might not occur concurrently (Jack et al., 2009), we applied SVR (as described in the previous section) separately for functional and structural MRI features. Both the functional and structural aging trajectories are shown in Figure 4.8. Note that local functional coherence provides information that is complementary to that of functional connectivity, therefore the functional BAI prediction is better with both sets of features ($R = 0.68$, $MAE = 5.54$, Figure 4.8), when compared to prediction with SCP features alone ($R = 0.55$, $MAE = 6.32$, Figure 4.8). Features generated from T1-MRI predict age with greater accuracy ($R = 0.80$, $MAE = 4.41$) than features from rsfMRI.

Using both these BAI values, we identified individuals who were aging very well, or *Resilient* agers, as those whose predicted BAI was lower than their expected BAI,

for both functional and structural data. We considered all other individuals to be advanced agers, as at least one of their aging patterns was worse than normal. In order to understand effects of advanced aging beyond the population-based aging effect (which was reported in the previous section), age was regressed out from all our MRI measurements.

Unlike aging effects which were observed in many of the primary and secondary SCPs, significant reduction in functional connectivity was only found in the DM vs. DA primary SCP (Figure 4.9) for the advanced group, relative to the resilient group. Other age-related reductions in connectivity, such as in the visual regions, were not reduced further in advanced aging. Aging-related patterns and two-group differences (resilient vs. advanced) for the local functional coherence and GM density features are summarized in Figure 4.10. Unlike aging effects in GM density, which are global, advanced agers showed spatially localized GM atrophy in the bilateral thalamus, hippocampus, amygdala, fronto-orbital cortex, precuneus and insula. Bilateral caudate, thalamus and anterior insula also showed significant decreases in coherence in advanced individuals relative to resilient individuals.

4.6 Discovering heterogeneity in advanced aging

Age-regressed data was also used as input to MOE for identifying subsets of advanced agers with diverse aging patterns. MOE with five groups of advanced agers had the highest ten-fold cross-validation accuracy ($74.0 \pm 5\%$). These five associated groups are reproducible (Adjusted Rand: 0.49 ± 0.20) and well separated (Bezdek Partition Coefficient: 0.76 ± 0.21).

There were no significant age or sex differences between individuals in the five groups. We found significant patterns of functional and structural change ($p < 0.05$,

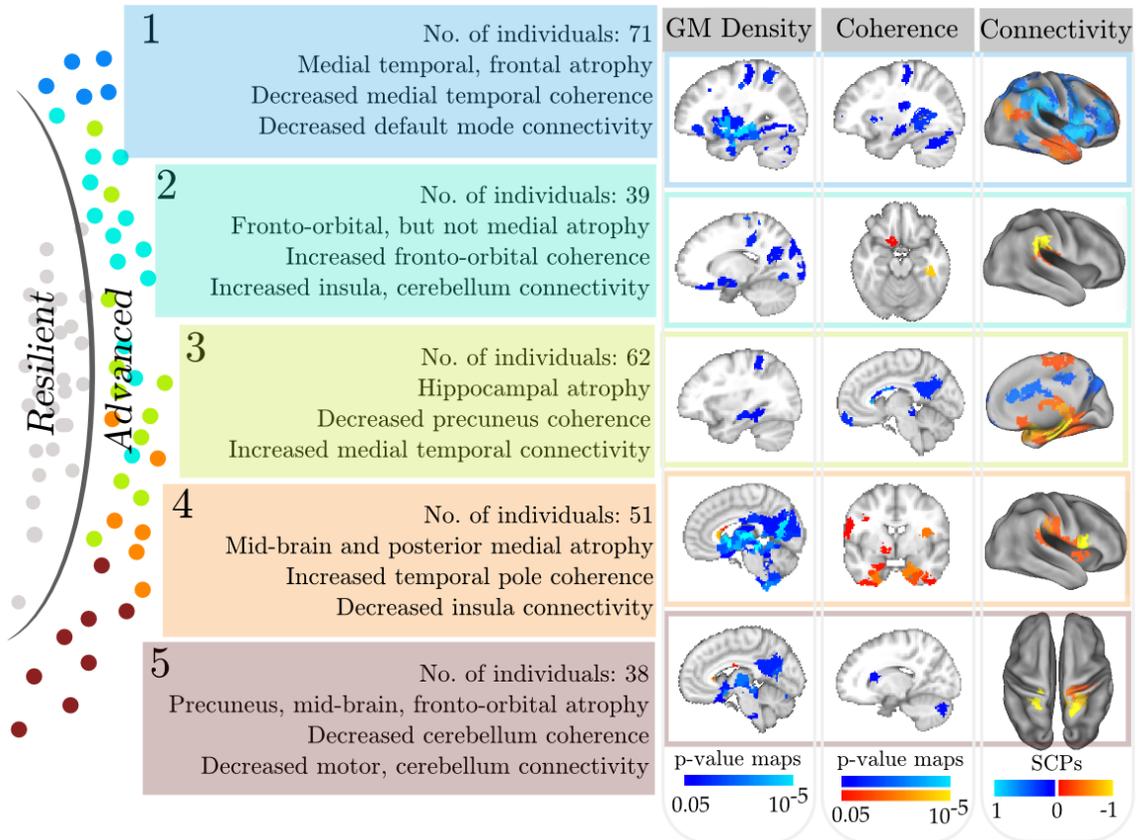


Figure 4.11: Summary of significant differences in GM density, functional coherence and connectivity for all five groups of advanced agers relative to resilient agers. As before, changes to GM density and coherence at the level of ROIs are shown using p-value maps overlaid on a template image (first two columns on right). SCPs whose coefficients were significantly different between groups are shown in the last column. The color overlay for SCPs indicates patterns of correlated (or anti-correlated) regions.

corrected for multiple comparisons) in the five groups relative to the resilient agers. These changes are summarized in Figure 4.11. Note that these changes are in addition to the normal aging effect that was observed across all the individuals (Figures 4.3,4.4,4.5,4.10).

The first and largest group consist of 71 advanced agers who show significant GM atrophy in the medial temporal lobe (MTL), including the hippocampus, parahippocampal gyrus, amygdalae bilaterally. Thalamus, caudate, superior temporal gyrus, inferior frontal GM, superior pre-central gyrus and cerebellum also showed significant atrophy. Coherence was significantly reduced in the Hippocampus, thalamus and anterior insula. Functional connectivity was reduced in the DM and DA regions.

With 39 individuals, group 2 had GM atrophy in the occipital fusiform and fronto-orbital regions but not in the MTL or mid-brain. It showed increased coherence in left temporal fusiform gyrus and increased connectivity in the bilateral supramarginal gyrus.

Group 3 comprised 62 individuals and has a remarkable pattern of decreases and increases: it had focused bilateral atrophy in the hippocampus, thalamus and superior temporal gyrus. Decreased coherence was found in the in the precuneus and left hippocampus. More importantly, it showed significantly increased connectivity in the MTL relative to the resilient group.

Group 4 included 51 individuals had GM atrophy in the entire mid brain and MTL, but not beyond these regions. Interestingly, it showed significantly increased coherence in the bilateral temporal poles, bilateral fronto-orbital regions and brain stem. Connectivity was significantly decreased between the bilateral opercular regions.

Group 5 with 38 individuals had the least amount of differences among all five

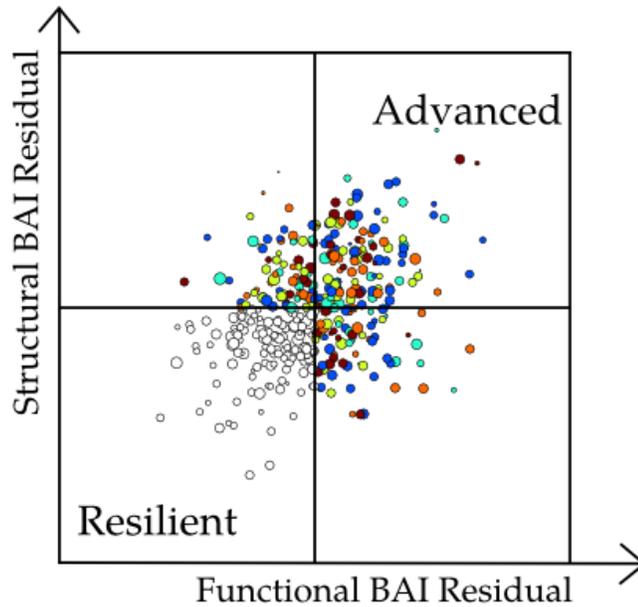


Figure 4.12: Functional and structural BAI Residual (= Expected Brain age - BAI) plotted in a two-dimensional grid. The group membership of each individual is reflected in the color; resilient agers are shown as white points.

groups. It showed thalamus, precuneus and hippocampal atrophy, decreased coherence in the posterior cerebellar regions and decreased connectivity in the cerebellum, motor and insula regions.

To further investigate if the patterns of change in each group differed only by a scale factor, we plotted the BAI residuals of all individuals in a two-dimensional plot, as shown in Figure 4.12. Clearly, the five groups are randomly distributed in this plot. Clear lack of clustering along either BAI direction suggests that these groups do not reflect different severity levels in advanced aging, but rather, present heterogeneous patterns of change, which were only revealed by the full multi-variate MOE analysis.

We compared the concurrent cognitive performance and cognitive rates of change between the six groups (five advanced + resilient) using data 10 years prior to

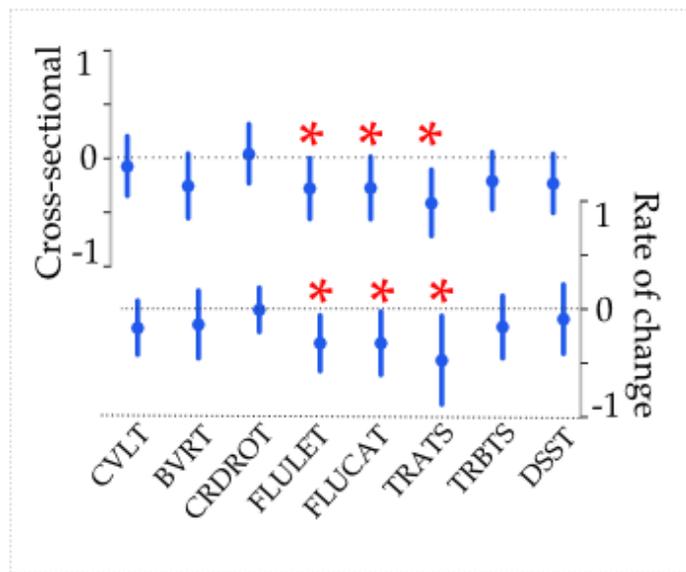


Figure 4.13: Differences in cognitive performance between advanced group 1 and resilient agers, for eight cognitive domains. Mean and 95% confidence intervals of the difference estimates are shown as errorbars. CVLT: California Verbal Learning Task. BVRT: Benton Visual Retention Test. CRDROT: Card Rotation Test. FLULET: Letter Fluency. FLUCAT: Category Fluency. TRATS: Trail Making Test Part A. TRBTS: Trail Making Test Part B. DSST: Digit Symbol Test.

the imaging visits. We considered participants' performance in eight cognitive domains (for details please see Appendix B). There were no significant differences between resilient and advanced agers (all five groups combined); but both the concurrent performance and rate of change was significantly lower for advanced group 1, compared to the resilient agers, in domains measuring verbal fluency and visual processing ($p < 0.05$, uncorrected). These estimates are shown in Figure 4.13.

We also observed that only one out of five advanced groups - group 3 showed somewhat better concurrent cognitive performance relative to resilient group ($p > 0.05$). Although insignificant ($p > 0.1$), individuals in group 3 also had higher educational levels on average (17.6 ± 0.38 years) compared to all others in the study (16.99 ± 0.150 years).

4.7 Discussion

A variety of unknown genetic and environmental factors may be associated with disease onset in older adults, leading to a decades-long process that may result in cognitive decline and dementia. Preclinical variants of age-related brain diseases, such as Alzheimers disease (AD), Lewy body disease and cerebro-vascular disease, are superimposed on the normal aging process. Separating normal from pathological aging in the preclinical phases of these disease states is a challenging task. Yet, it is of vital importance that we gain a better understanding of variations in brain aging so that we may identify distinct disease processes underlying accelerated brain changes to aid in development and application of potential therapeutic interventions.

In this chapter, we showed the application of the sparse learning framework to a population study of aging. Similar to other cross-sectional imaging studies, using

uni-variate analysis, we determined the effects of aging in a group of 400 adults in the age range of 50-96 years, all of whom were cognitively performing within the normal testing limits at the time of scan. We found that connectivity within the DM and DA regions is reduced with aging, consistent with earlier reports (Wu et al., 2011; Andrews-Hanna et al., 2007). More recently, increased hippocampal connectivity with aging was also reported (Salami et al., 2014), a finding we were able to replicate - the SCP coefficient associated with the MTL SCP showed a positive correlation with age (Figure 4.3). We also found that connectivity in the motor and visual areas degraded with age.

Note that the structural and functional effects of aging that we presented in Figures 4.3,4.4, may be overestimated, because some of the individuals in the study are likely to be in the preclinical stages of various neurodegenerative diseases (Burgmans et al., 2009). Individuals in preclinical stages of disease may show greater structural and functional brain changes (Pacheco et al., 2015; Beason-Held et al., 2013). Thus, our analyses present a typical view of aging, which is an average picture that combines resilient and advanced agers. In this study, we sought to identify heterogeneous subgroups of advanced agers whose patterns of brain structure and function is more advanced than the age-specific average.

Two group comparison of advanced vs. resilient aging revealed that in addition to the degeneration seen in normal aging, advanced agers were found to have spatially localized atrophy in the mid brain and medial temporal areas. Furthermore, advanced aging seemed to be preferentially affecting the connectivity between DM and DA regions. Other age-related reductions in connectivity, such as in the visual regions, were not reduced further in advanced aging.

Using MOE, we identified five distinct groups with heterogeneous patterns of functional and structural changes associated with advanced aging. All five groups

showed varying extent and severity of structural atrophy. Two groups had decreased functional connectivity and coherence alongside the atrophy, while others showed opposing effects in functional changes.

Combining all evidence, group 1 had the most dominant pattern, with significant GM atrophy, reduced functional coherence as well as connectivity. Results from neuro-cognitive data show that this group has significantly advanced cognitive decline in many domains. Preferential degeneration in the MTL, decreased coherence and significant losses in connectivity in the DM regions, are generally consistent with changes that have been previously reported in aging and AD studies. Although normal aging itself may lead to global GM atrophy and functional decline, preferential loss of mid-brain and hippocampal GM density and reduced DM connectivity is uniquely observed in Alzheimer's like dementia (Seeley et al., 2009).

Group 3 was very interesting, in that it displayed imaging characteristics consistent with very early AD, including focal hippocampal GM atrophy, reduced posterior cingulate/precuneus coherence, and relatively increased functional connectivity of the MTL. Similar patterns have been previously reported in individuals with early stages of Mild Cognitive Impairment Bai et al. (2008); Sperling (2011); Dickerson et al. (2005). The increased cognitive performance and education level in this group suggest that their functional hyperconnectivity could be a result of compensatory mechanisms that initiate as pathological processes set in Sperling (2011), and/or of higher levels of cognitive reserve (Rentz et al., 2010).

Group 4 shows significant increases in bilateral orbito-frontal and temporal pole regions. Increased coherence has been reported in studies investigating early stages of Alzheimer's, but in other regions in the brain (He et al., 2007b). Reduced motor and cerebellar function was limited to individuals in group 5, which also showed signs of advanced mid brain and precuneus atrophy. Reduced brain function in

these regions could possibly indicate impaired motor ability, which is prevalent among the elderly.

In this study, we sought to identify advanced agers with patterns of brain structure and function that are worse than age-specific averages. The results show that there are several distinct patterns associated with advanced brain aging in cognitively normal individuals. The groups of advanced agers are not significantly different in terms of age or sex distribution, indicating individual differences in patterns of brain aging. Moreover, our findings suggest that pathological changes do not universally manifest after a certain age. Instead, differences may be due to unknown underlying biological and genetic mechanisms, some of which are independent of age (Ritchie and Kildea, 1995; Nelson et al., 2011). The most striking findings of this study are that three of the five advanced aging groups demonstrate atrophy and altered functional integrity in regions that show similar characteristics in age-related disease, particularly AD. Continued follow-up of these individuals will help to further separate the effects of age from disease.

This is the first exploratory study that attempts to understand the heterogeneity present in normal aging. We used advanced multivariate techniques to identify advanced aging individuals who are beginning to show signs of structural and functional brain disruption, possibly due to the development of pathological processes. Such an analysis was made possible due to the large sample size of the BLSA dataset and recent advances in machine learning and pattern recognition technology.

The patterns of brain aging that we report here provide preclinical imaging-based biomarkers of neuro-degeneration that may be relevant to a variety of disease processes. With continued follow-up we will determine whether the subgroups identified here show different long-term cognitive outcomes. With validation in other cohorts, we hope that the use of these types of imaging biomarkers will fa-

facilitate selection of individuals for therapeutic interventions to delay onset or slow progression of neurodegenerative disease. The presented results significantly enhance our understanding of aging, and hopefully serve as a milestone for future studies.

Chapter 5

Final Remarks and Future Work

The novel computational methods presented in this thesis complement existing analytical approaches for functional connectivity research. The methodological contributions of this thesis are two fold. First, we developed a sparsity-based matrix factorization approach to find patterns of connectivity that co-vary across individuals. In addition to SCPs, Sparse Learning provides individual-level estimates of connectivity that can be used for comparing connectomes.

Second, we extended the popular two-group SVM classification framework to work with heterogeneity in the affected group relative to a reference group. The MOE method finds sub-groups in the affected group based on (1) distinct patterns of change between the affected sub-groups relative to the reference group and (2) underlying clustering patterns within the affected group. Thus, MOE combines supervised and unsupervised machine learning, to fill a major methodological gap in MR imaging studies for identification of not one, but multiple distinct disease or aging patterns. These methods can be applied to any rsfMRI study to elucidate complex and heterogeneous changes in the functional connectome.

Sparse Learning was applied to rsfMRI data acquired as a part of the BLSA to

find effects of aging on brain function. Using a purely data-driven analysis pipeline, we discovered a negative association between age and connections between the DM vs. DA, visual and motor regions; and a positive association for connections between regions within the medial temporal lobe. The advantage of a dimensionality reduction method was evident in the results; SCPs provided an interpretable basis using which aging effects on connectivity could be better understood.

We used both methods to disentangle the heterogeneity seen in advanced aging. To discover this heterogeneity, we used features extracted from sMRI, along with SCP Coefficients, to build normative brain aging trajectories of structure and function. Using these trajectories as a reference, we were able to identify individuals who showed signs of advanced aging. Using MOE, we were able to discover a remarkable amount of diversity among the advanced agers. A subset of advanced agers showed significantly reduced GM density and functional connectivity, along with cognitive decline. Another subset had focal hippocampal GM atrophy, but showed signs of functional compensation, by increased recruitment of the bilateral medial temporal lobe. This is the first study that has identified imaging based bio-markers that capture precursors to the multitude of pathologic processes that develop with aging. Our research paves the way to developing translational diagnostic tools that can provide personalized indices of brain health in the elderly.

5.1 Future work

The methods described in this paper, along with pre-processing and parcellation performed prior to it, constitute a complete, data-driven analytical framework for resting state connectivity research. These contributions lay the groundwork for investigating open questions in rsfMRI connectivity research and its clinical applica-

tions in disease, development and aging. We propose three specific future objectives that can be accomplished using this framework below:

1. Functional Connectivity Dynamics

In this thesis, we assumed that the extent of coupling between any two regions in the brain was static; a single correlation matrix was computed for every individual, which was used as input for all subsequent analysis. In recent years, multiple studies have found evidence for variations in connectivity over time within a recording session; since then, dynamics of functional connectivity has emerged as a very active research area (Chang and Glover, 2010; Sakoğlu et al., 2010; Kiviniemi et al., 2011).

A popular technique to investigate changes in functional connectivity is to use sliding time windows within a recording period (scan) and compute multiple correlation matrices for each individual. This provides a “time-series” of correlations which can then be explored to characterize distinct “brain-states”. Using sliding-time windows increases the number of correlations to be analyzed by a factor of ten, therefore dimensionality reduction is necessary to parse this type of data. Earlier studies have used PCA or K-Means to characterize these brain-states (Leonardi et al., 2013; Allen et al., 2012). As we demonstrated SCPLearn’s successful use in analyzing average correlation matrices, applying the same to sliding window correlation matrices might lead to potentially informative results. Such an analysis would provide SCPs that capture sets of correlated regions whose connections covary not only across individuals, but also in time; this variation with time is quantified in the SCP coefficients. It might be interesting to note how temporally-delineated SCPs differ from static SCPs reported in this thesis.

However, sliding window analysis suffers from some limitations. Apart from issues pertaining to contribution of non-stationary noise and effect of window length, a primary concern is that it cannot distinguish between dynamic amplitude (first-order), and dynamic correlation (second-order) (Hutchison et al., 2013). In an exploratory paper, we proposed the use of Hidden Markov Models (HMMs) combined with Sparse Learning to model rsfMRI time-series (Eavani et al., 2013). The HMM is a state-space model that assumes that each individual's rsfMRI time-series is composed of latent discrete "brain states" such that each state is defined by a distinct mean vector (amplitude) and covariance matrix (connectivity). We can test for the presence of dynamics in connectivity by checking if the data supports modeling of more than one brain-state, using a split-sample validation framework. Using this approach, we were able to find evidence for six brain states based on data-fit; however, further experiments are needed to check for reproducibility within and across studies.

2. Heterogeneity for continuous outcomes

In Chapter 3, we demonstrated that by combining an SVM classifier with clustering, we can capture heterogeneity in an affected group relative to a reference group. This idea can be easily extended to a regression framework that uses Support Vector Regression (SVR), in order to model non-linear or multiple regression curves. In the non-linear case, one can build a piece-wise linear aging trajectory, with each linear segment associated with a different age range, where the change-points are determined automatically by the algorithm. Alternately, it can be used to build not one but multiple aging trajectories spanning the entire age range (for example, slow-aging and rapid-aging trajectories) and identify individuals associated with each.

3. Functional connectivity changes over the lifespan

Correlation values are limited in range $[-1, 1]$ and are comparable across studies irrespective of acquisition protocols (barring systematic biases such as motion effects). This advantage facilitates the study of connectome changes across the lifespan, by combining data across studies such as the PNC and BLSA. SCPs and coefficients learned from the entire age range provide compact representations of these changes. MOE-regression can be applied to SCP coefficients to capture non-linear connectivity changes with age - for example, increased connectivity during maturation and decline during aging. In this regard, the remarkable reproducibility seen between SCPs obtained for the PNC and the BLSA dataset (Figures 2.14,2.15) is encouraging.

5.2 Software

SCPLearn is open-source software, implemented primarily in MATLAB, with a command line interface. For argument parsing and nifti I/O operations, python wrapper code is used. For ease of use, the software takes as input either rsfMRI voxel time-series data in NIFTI file format, or ROI time-series data in MATLAB's "mat" file format. The list of NIFTI or mat files is provided as input using a text file. It also takes as required input, the node definitions as an atlas in NIFTI space. This atlas must have the ROI regions numbered sequentially and must be in the same space as the individual data. SCPLearn outputs SCPs mapped to the atlas space in NIFTI file format, along with SCP coefficients indexed against individual filenames as a comma separated values (csv) file. A snapshot of the SCPLearn command-line is shown in 5.1.

MOE is also available as open-source software. It is written entirely in Python

```

scplearn--
Generates Sparse Connectivity Patterns from resting state fMRI data

Usage: scplearn [OPTIONS]

Required Options:
[-d --data]          Specify the text file with list of nifti files (required)
                    *** OR ***
                    Specify the text file with list of mat files (required)
                    Each mat file must contain a variable named 'ts'
                    'ts' must be a matrix of time-series, size (# of ROIs X # of timepoints)

[-m --mask]          Specify the nifti ROI/parcel/atlas file (required)
[-p --prefix]        Specify the prefix of the output file (required)

Options:
[-t --type]          Specify the data type of input with either "matlab" or "nii".
[-s --sparsity]      Specify the sparsity constraint as positive value. Default = nROIs/10.
[-n --numberOfSCPs] Specify number of primary SCPs. Default = 10.
[-r --pruning]       Specify the pruning threshold as a value between [0,1]. Default = 0.7.
[-l --levels]        Specify number of secondary SCPs. Default = 50.

[-o --outputDir]    The output directory to write the results. Defaults to the location of the input file
[-w --workingDir]    Specify a working directory. By default a tmp dir is created and used
[-u --usage | -h --help] Display this message
[-v --verbose]       Verbose output
[-V --Version]       Display version information

```

Figure 5.1: Command-line interface of SCPLearn

language. The command line interface provides options for an input csv file, which contains ROI-level data, and other user-defined parameters. A snapshot is shown in 5.2. It outputs a separate csv file for the hyperplane weights and membership values, which are indexed against ROI labels and individual IDs respectively.

```

moe --
Runs MOE-based heterogeneity classification on spreadsheet data
Works on any set of features - Average regional volume/density/diffusion/connectivity,
                              clinical scores, cognitive data, etc
*** For best performance, pick optimal -c and -t values based on cross-validation ***

Usage: moe [OPTIONS]

Required Options:
[-d --data]          Specify the spreadsheet with list of ROIs (required)
[-p --prefix]       Specify the prefix of the output file (required)
[-H --header]       Specify the header column that has the discrete labels (required)
[-l --label]        Specify the label for the heterogenous group (required)

Options:
[-n --nSVMs]        Specify the number of SVMs as a number. Default = 2.
[-c --cost]         Specify the SVM cost value. Default = 1.
[-t --tradeoff]     Specify the classification vs. clustering tradeoff. Default = 0.1.
[-f --folds]        Specify the number of cross-validation folds to run. Default = 0. (not run)

[-o --outputDir]    The output directory to write the results. Defaults to the location of the input file
[-w --workingDir]   Specify a working directory. By default a tmp dir is created and used
[-u --usage | -h --help] Display this message
[-v --verbose]      Verbose output
[-V --Version]      Display version information

```

Figure 5.2: Command-line interface of MOE

Appendices

Appendix A

rsfMRI data pre-processing

A.1 Time-series pre-processing

The first few volumes of the functional time-series were removed to allow signal stabilization. Functional images were slice-time corrected and re-aligned using MCFLIRT (Jenkinson et al., 2002). Structural images were skull-stripped using BET (Smith, 2002) and segmented using FAST (Zhang et al., 2001); mean white matter (WM) and cerebro-spinal fluid (CSF) signals were extracted from the tissue segments generated for each subject. Confound regression (Satterthwaite et al., 2012) included these 6 standard motion parameters, the WM signal, the CSF signal, and the global signal (i.e., 9 parameters total), as well as the temporal derivative, quadratic term, and temporal derivative of the quadratic of each (36 regressors total). Notably, in order to avoid a mismatch in the frequency domain (Hallquist et al., 2013), both the confound matrix and the time-series data was simultaneously band-pass filtered to retain signals between 0.01-0.08 Hz using AFNI's 3dBandpass utility (Cox, 1996).

A.2 Spatial Alignment

Subject-level BOLD images were co-registered to the T1 image using FLIRT (Jenkinson et al., 2012). Whole-head T1 images were registered to the Montreal Neurologic Institute 152 template using DRAMMS deformable registration (Ou et al., 2011) for the BLSA dataset, and ANTs diffeomorphic SyN registration (Avants et al., 2008) for the PNC dataset.

A.3 Controlling for motion confound

Head motion during acquisition of rsfMRI scans is known to affect functional connectivity in a systematic manner (Power et al., 2012). This is problematic in studies of development and aging, as younger children and older adults generally move more, making motion a nuisance confounder (Mowinckel et al., 2012). In this study, in order to mitigate the effects of motion, we incorporated three corrective steps in various stages of the pre-processing pipeline. First, we restricted our analysis to only those subjects with a summary motion value of less than 0.2mm. Second, for the BLSA study, we performed a “scrubbing” procedure that discards all volumes whose summary motion and BOLD signal variance values are higher than a selected threshold (Power et al., 2012). Third, the global signal was regressed out of the voxel-wise data, as it is known to be a good surrogate measure for the effect of motion and other physiological effects on BOLD signal (Satterthwaite et al., 2012).

A.4 Global Signal Regression

SCPs presented in this thesis were obtained after removing the baseline global signal from each subject's data, which facilitates the delineation of functional systems by removing the confounding effects of motion and other non-neuronal sources of noise (Fox et al., 2009). On the other hand, many researchers argue that performing Global Signal Regression (GSR) on rsfMRI data removes relevant signal and tends to increase the number of negatively correlated nodes (Saad et al., 2012). It is unclear if the high reproducibility of our results reflects true signal, or a systematic artifact induced due to global signal regression. Hence we re-ran Sparse Learning on the PNC dataset for the Areal Graph nodes, but with the global signal retained. SCPs continued to be reproducible, however they were substantively different. Of the ten SCPs computed, only one SCP had a significant areas of negative correlation - the Dorsal Attention vs. Default mode anti-correlation pattern. The other nine SCPs showed only positive correlations. Of these, the familiar SCP patterns were the sensori-motor, visual and cingulo-opercular networks (see Figure A.1, and compare to Figure 2.10).

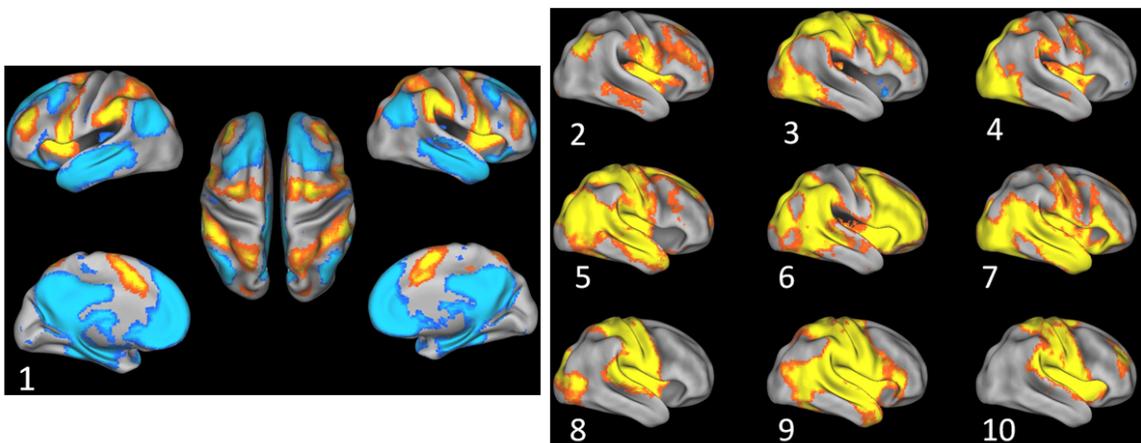


Figure A.1: SCs computed without global signal regression, from the PNC dataset, using the Areal Graph nodes. Only SCP 1 has a significant area of negative correlation - the Dorsal Attention vs. Default mode anti-correlation pattern. The other nine SCs showed only positive correlation. Of these, the familiar patterns were the sensori-motor (8, 9 and 10), visual (4, 7) and cingulo-opercular networks (2).

Appendix B

BLSA: Cognitive Data

BLSA participants receive a battery of cognitive tests at every visit. Participants are assessed in eight cognitive domains, listed below:

1. California Verbal Learning Task (CVLT) was used to assess verbal learning and memory. Higher values indicate better performance.
2. Benton Visual Retention Test (BVRT) quantifies figural memory and visuo-constructional ability. Lower values indicate better performance.
3. CARD Rotation Test (CRT) measure the ability to mentally manipulate figures. Higher values indicate better performance.
4. Letter Fluency (FLULET) measures phonemic fluency. Higher values indicate better performance.
5. Category Fluency (FLUCAT) measures semantic fluency. Higher values indicate better performance.
6. Trail Making Test Part A (TRATS) was used as an indicator of visual attention and processing speed. Lower values indicate better performance.

7. Trail Making Test Part B (TRBTS) was used to evaluate executive function. Lower values indicate better performance.
8. Digit Symbol Test (DSST). Higher values indicate better performance.

As BLSA is an ongoing longitudinal study, we have cognitive data from multiple cognitive assessments for each participant over up to a 19 year period prior to, and concurrent with, the time of scan. However, participants have varying numbers of visits and cognitive assessments. Therefore, every participant's cognitive performance in each of the eight domains is summarized using three measurements. For each of the $3 * 8 = 24$ measurements, we evaluate whether the subgroups obtained using the MOE model (using the SCP data alone) showed differences.

For comparing sub-groups of advanced agers reported in Chapter 4, Section 4.6, we used linear mixed models to compare among groups - (1) the concurrent (cross-sectional) cognitive performance and (2) cognitive rates of change (slope) using 10 year data prior to the imaging visits. The time of follow up is coded reversely, that is using the time of imaging visit as origin (interval = 0), and all the visits prior have the negative interval (-1, -2, -3 years, etc.). Coding in this manner allows us to compare the concurrent cognitive performance and cognitive rates of change in one single model. We additionally adjusted for sex. (Age is not adjusted, since it is balanced across groups.) The predictors include sex, group, interval and interaction terms - sex*group, and group*interval. F-tests were used to test the null hypothesis that all the groups are equal.

Bibliography

- Ahn, Y.-Y., Bagrow, J. P., and Lehmann, S. (2010). Link communities reveal multi-scale complexity in networks. *Nature*, 466(7307):761–764.
- Allen, E. A., Damaraju, E., Plis, S. M., Erhardt, E. B., Eichele, T., and Calhoun, V. D. (2012). Tracking whole-brain connectivity dynamics in the resting state. *Cerebral Cortex*.
- Andrews-Hanna, J. R., Snyder, A. Z., Vincent, J. L., Lustig, C., Head, D., Raichle, M. E., and Buckner, R. L. (2007). Disruption of large-scale brain systems in advanced aging. *Neuron*, 56(5):924–935.
- Avants, B. B., Epstein, C. L., Grossman, M., and Gee, J. C. (2008). Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Medical image analysis*, 12(1):26–41.
- Bai, F., Zhang, Z., Yu, H., Shi, Y., Yuan, Y., Zhu, W., Zhang, X., and Qian, Y. (2008). Default-mode network activity distinguishes amnesic type mild cognitive impairment from healthy aging: a combined structural and resting-state functional mri study. *Neuroscience letters*, 438(1):111–115.
- Bansal, N., Blum, A., and Chawla, S. (2004). Correlation clustering. *Machine Learning*, 56(1-3):89–113.

- Batmanghelich, N., Taskar, B., and Davatzikos, C. (2012). Generative-discriminative basis learning for medical imaging. *Medical Imaging, IEEE Transactions on*, 31(1):51–69.
- Beason-Held, L. L., Goh, J. O., An, Y., Kraut, M. A., O'Brien, R. J., Ferrucci, L., and Resnick, S. M. (2013). Changes in brain function occur years before the onset of cognitive impairment. *The Journal of Neuroscience*, 33(46):18008–18014.
- Beckmann, C. F., DeLuca, M., Devlin, J. T., and Smith, S. M. (2005). Investigations into resting-state connectivity using independent component analysis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1457):1001–1013.
- Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 289–300.
- Bezdek, J. C. (1981). *Pattern recognition with fuzzy objective function algorithms*. Kluwer Academic Publishers.
- Bezdek, J. C., Ehrlich, R., and Full, W. (1984). Fcm: The fuzzy c-means clustering algorithm. *Computers & Geosciences*, 10(2):191–203.
- Bishop, C. et al. (2006). *Pattern recognition and machine learning*, volume 4. springer New York.
- Biswal, B., Zerrin Yetkin, F., Haughton, V. M., and Hyde, J. S. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar mri. *Magnetic resonance in medicine*, 34(4):537–541.
- Biswal, B. B., Mennes, M., Zuo, X.-N., Gohel, S., Kelly, C., Smith, S. M., Beckmann, C. F., Adelstein, J. S., Buckner, R. L., Colcombe, S., et al. (2010). Toward dis-

- covery science of human brain function. *Proceedings of the National Academy of Sciences*, 107(10):4734–4739.
- Brouwer, R. K. (2009). Extending the rand, adjusted rand and jaccard indices to fuzzy partitions. *Journal of Intelligent Information Systems*, 32(3):213–235.
- Buckner, R. L., Sepulcre, J., Talukdar, T., Krienen, F. M., Liu, H., Hedden, T., Andrews-Hanna, J. R., Sperling, R. A., and Johnson, K. A. (2009). Cortical hubs revealed by intrinsic functional connectivity: mapping, assessment of stability, and relation to alzheimer’s disease. *The Journal of Neuroscience*, 29(6):1860–1873.
- Burgmans, S., van Boxtel, M. P., Vuurman, E. F., Smeets, F., Gronenschild, E. H., Uylings, H., and Jolles, J. (2009). The prevalence of cortical gray matter atrophy may be overestimated in the healthy aging brain. *Neuropsychology*, 23(5):541.
- Calhoun, V., Adali, T., Hansen, L., Larsen, J., and Pekar, J. (2003). Ica of functional mri data: an overview. In *in Proceedings of the International Workshop on Independent Component Analysis and Blind Signal Separation*. Citeseer.
- Chang, C. and Glover, G. H. (2010). Time–frequency dynamics of resting-state brain connectivity measured with fmri. *Neuroimage*, 50(1):81–98.
- Chang, C.-C. and Lin, C.-J. (2011). LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- Clark, V. H., Resnick, S. M., Doshi, J., Beason-Held, L. L., Zhou, Y., Ferrucci, L., Wong, D. F., Kraut, M. A., and Davatzikos, C. (2012). Longitudinal imaging pattern analysis (spare-cd index) detects early structural and functional

- changes before cognitive decline in healthy older adults. *Neurobiology of aging*, 33(12):2733–2745.
- Cohen, A. L., Fair, D. A., Dosenbach, N. U., Miezin, F. M., Dierker, D., Van Essen, D. C., Schlaggar, B. L., and Petersen, S. E. (2008). Defining functional areas in individual human brains using resting functional connectivity mri. *Neuroimage*, 41(1):45.
- Corbetta, M. and Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature reviews neuroscience*, 3(3):201–215.
- Cox, R. W. (1996). Afni: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical research*, 29(3):162–173.
- Davatzikos, C., Xu, F., An, Y., Fan, Y., and Resnick, S. M. (2009). Longitudinal progression of alzheimer’s-like patterns of atrophy in normal older adults: the spare-ad index. *Brain*, 132(8):2026–2035.
- Dave, R. N. (1996). Validating fuzzy partitions obtained through c-shells clustering. *Pattern Recognition Letters*, 17(6):613–623.
- Desikan, R., Ségonne, F., Fischl, B., Quinn, B., Dickerson, B., Blacker, D., Buckner, R., Dale, A., Maguire, R., Hyman, B., et al. (2006). An automated labeling system for subdividing the human cerebral cortex on mri scans into gyral based regions of interest. *Neuroimage*, 31(3):968–980.
- Di Martino, A., Yan, C.-G., Li, Q., Denio, E., Castellanos, F. X., Alaerts, K., Anderson, J. S., Assaf, M., Bookheimer, S. Y., Dapretto, M., et al. (2014). The autism brain imaging data exchange: towards a large-scale evaluation of the intrinsic brain architecture in autism. *Molecular psychiatry*, 19(6):659–667.

- Dickerson, B., Salat, D., Greve, D., Chua, E., Rand-Giovannetti, E., Rentz, D., Bertram, L., Mullin, K., Tanzi, R., Blacker, D., et al. (2005). Increased hippocampal activation in mild cognitive impairment compared to normal aging and ad. *Neurology*, 65(3):404–411.
- Dong, A., Honnorat, N., Gaonkar, B., and Davatzikos, C. (2015). Chimera: Clustering of heterogeneous disease effects via distribution matching of imaging patterns.
- Dosenbach, N., Fair, D., Miezin, F., Cohen, A., et al. (2007). Distinct brain networks for adaptive and stable task control in humans. *Proceedings of the National Academy of Sciences*, 104(26):11073.
- Dosenbach, N. U., Visscher, K. M., Palmer, E. D., Miezin, F. M., Wenger, K. K., Kang, H. C., Burgund, E. D., Grimes, A. L., Schlaggar, B. L., Petersen, S. E., et al. (2006). A core system for the implementation of task sets. *Neuron*, 50(5):799.
- Eavani, H., Satterthwaite, T. D., Gur, R. E., Gur, R. C., and Davatzikos, C. (2013). Unsupervised learning of functional network dynamics in resting state fmri. In *Information Processing in Medical Imaging*, pages 426–437. Springer.
- Fan, R.-E., Chang, K.-W., Hsieh, C.-J., Wang, X.-R., and Lin, C.-J. (2008). Liblinear: A library for large linear classification. *The Journal of Machine Learning Research*, 9:1871–1874.
- Fan, Y., Shen, D., Gur, R. C., Gur, R. E., and Davatzikos, C. (2007). Compare: classification of morphological patterns using adaptive regional elements. *Medical Imaging, IEEE Transactions on*, 26(1):93–105.
- Ferreira, L. K. and Busatto, G. F. (2013). Resting-state functional connectivity in normal brain aging. *Neuroscience & Biobehavioral Reviews*, 37(3):384–400.

- Fox, M. D., Corbetta, M., Snyder, A. Z., Vincent, J. L., and Raichle, M. E. (2006). Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems. *Proceedings of the National Academy of Sciences*, 103(26):10046–10051.
- Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., and Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences of the United States of America*, 102(27):9673–9678.
- Fox, M. D., Snyder, A. Z., Vincent, J. L., and Raichle, M. E. (2007). Intrinsic fluctuations within cortical systems account for intertrial variability in human behavior. *Neuron*, 56(1):171–184.
- Fox, M. D., Zhang, D., Snyder, A. Z., and Raichle, M. E. (2009). The global signal and observed anticorrelated resting state brain networks. *Journal of neurophysiology*, 101(6):3270–3283.
- Franke, K., Luders, E., May, A., Wilke, M., and Gaser, C. (2012). Brain maturation: predicting individual brainage in children and adolescents using structural mri. *Neuroimage*, 63(3):1305–1312.
- Fu, Z., Robles-Kelly, A., and Zhou, J. (2010). Mixing linear svms for nonlinear classification. *Neural Networks, IEEE Transactions on*, 21(12):1963–1975.
- Golland, P. (2001). Discriminative direction for kernel classifiers. In *Advances in neural information processing systems*, pages 745–752.
- Greicius, M. D., Srivastava, G., Reiss, A. L., and Menon, V. (2004). Default-mode network activity distinguishes alzheimer’s disease from healthy aging: evidence from functional mri. *Proceedings of the National Academy of Sciences of the United States of America*, 101(13):4637–4642.

- Hallquist, M. N., Hwang, K., and Luna, B. (2013). The nuisance of nuisance regression: Spectral misspecification in a common approach to resting-state fmri preprocessing reintroduces noise and obscures functional connectivity. *NeuroImage*.
- He, B. J., Snyder, A. Z., Vincent, J. L., Epstein, A., Shulman, G. L., and Corbetta, M. (2007a). Breakdown of functional connectivity in frontoparietal networks underlies behavioral deficits in spatial neglect. *Neuron*, 53(6):905–918.
- He, Y., Wang, L., Zang, Y., Tian, L., Zhang, X., Li, K., and Jiang, T. (2007b). Regional coherence changes in the early stages of alzheimer’s disease: a combined structural and resting-state functional mri study. *Neuroimage*, 35(2):488–500.
- Honorat, N., Eavani, H., Satterthwaite, T., Gur, R., Gur, R., and Davatzikos, C. (2015). Grasp: Geodesic graph-based segmentation with shape priors for the functional parcellation of the cortex. *NeuroImage*, 106:207–221.
- Hutchison, R. M., Womelsdorf, T., Allen, E. A., Bandettini, P. A., Calhoun, V. D., Corbetta, M., Della Penna, S., Duyn, J. H., Glover, G. H., Gonzalez-Castillo, J., et al. (2013). Dynamic functional connectivity: promise, issues, and interpretations. *Neuroimage*, 80:360–378.
- Ingalhalikar, M., Parker, D., Bloy, L., Roberts, T. P., and Verma, R. (2011). Diffusion based abnormality markers of pathology: toward learned diagnostic prediction of asd. *Neuroimage*, 57(3):918–927.
- Jack, C. R., Knopman, D. S., Jagust, W. J., Shaw, L. M., Aisen, P. S., Weiner, M. W., Petersen, R. C., and Trojanowski, J. Q. (2010). Hypothetical model of dynamic biomarkers of the alzheimer’s pathological cascade. *The Lancet Neurology*, 9(1):119–128.

- Jack, C. R., Lowe, V. J., Weigand, S. D., Wiste, H. J., Senjem, M. L., Knopman, D. S., Shiung, M. M., Gunter, J. L., Boeve, B. F., Kemp, B. J., et al. (2009). Serial pib and mri in normal, mild cognitive impairment and alzheimer's disease: implications for sequence of pathological events in alzheimer's disease. *Brain*, page awp062.
- Jacobs, R. A., Jordan, M. I., Nowlan, S. J., and Hinton, G. E. (1991). Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87.
- Jenkinson, M., Bannister, P., Brady, M., and Smith, S. (2002). Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage*, 17(2):825–841.
- Jenkinson, M., Beckmann, C. F., Behrens, T. E., Woolrich, M. W., and Smith, S. M. (2012). Fsl. *Neuroimage*, 62(2):782–790.
- Keller, C. J., Bickel, S., Honey, C. J., Groppe, D. M., Entz, L., Craddock, R. C., Lado, F. A., Kelly, C., Milham, M., and Mehta, A. D. (2013). Neurophysiological investigation of spontaneous correlated and anticorrelated fluctuations of the bold signal. *The Journal of Neuroscience*, 33(15):6333–6342.
- Kiviniemi, V., Vire, T., Remes, J., Elseoud, A. A., Starck, T., Tervonen, O., and Nikkinen, J. (2011). A sliding time-window ica reveals spatial variability of the default mode network in time. *Brain connectivity*, 1(4):339–347.
- Ladicky, L. and Torr, P. (2011). Locally linear support vector machines. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 985–992.
- Langan, J., Peltier, S. J., Bo, J., Fling, B. W., Welsh, R. C., and Seidler, R. D. (2010). Functional implications of age differences in motor system connectivity. *Frontiers in systems neuroscience*, 4.

- Leonardi, N., Richiardi, J., Gschwind, M., Simioni, S., Annoni, J.-M., Schluiep, M., Vuilleumier, P., and Van De Ville, D. (2013). Principal components of functional connectivity: A new approach to study dynamic brain connectivity during rest. *NeuroImage*, 83:937–950.
- Mayer, A. R., Mannell, M. V., Ling, J., Gasparovic, C., and Yeo, R. A. (2011). Functional connectivity in mild traumatic brain injury. *Human brain mapping*, 32(11):1825–1835.
- Meier, T. B., Desphande, A. S., Vergun, S., Nair, V. A., Song, J., Biswal, B. B., Meyerand, M. E., Birn, R. M., and Prabhakaran, V. (2012). Support vector machine classification and characterization of age-related reorganization of functional brain networks. *Neuroimage*, 60(1):601–613.
- Morris, J. (2004). Early-stage and preclinical alzheimer disease. *Alzheimer disease and associated disorders*, 19(3):163–165.
- Mowinckel, A. M., Espeseth, T., and Westlye, L. T. (2012). Network-specific effects of age and in-scanner subject motion: a resting-state fmri study of 238 healthy adults. *Neuroimage*, 63(3):1364–1373.
- Munkres, J. (1957). Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial & Applied Mathematics*, 5(1):32–38.
- Nelson, P. T., Head, E., Schmitt, F. A., Davis, P. R., Neltner, J. H., Jicha, G. A., Abner, E. L., Smith, C. D., Van Eldik, L. J., Kryscio, R. J., et al. (2011). Alzheimer’s disease is not “brain aging”: neuropathological, genetic, and epidemiological human studies. *Acta neuropathologica*, 121(5):571–587.
- Ou, Y., Sotiras, A., Paragios, N., and Davatzikos, C. (2011). Dramms: Deformable

- registration via attribute matching and mutual-saliency weighting. *Medical image analysis*, 15(4):622–639.
- Pacheco, J., Goh, J. O., Kraut, M. A., Ferrucci, L., and Resnick, S. M. (2015). Greater cortical thinning in normal older adults predicts later cognitive impairment. *Neurobiology of aging*, 36(2):903–908.
- Power, J. D., Barnes, K. A., Snyder, A. Z., Schlaggar, B. L., and Petersen, S. E. (2012). Spurious but systematic correlations in functional connectivity mri networks arise from subject motion. *Neuroimage*, 59(3):2142–2154.
- Power, J. D., Cohen, A. L., Nelson, S. M., Wig, G. S., Barnes, K. A., Church, J. A., Vogel, A. C., Laumann, T. O., Miezin, F. M., Schlaggar, B. L., and Petersen, S. E. (2011). Functional network organization of the human brain. *Neuron*, 72(4):665 – 678.
- Raichle, M., MacLeod, A., Snyder, A., Powers, W., Gusnard, D., and Shulman, G. (2001). A default mode of brain function. *Proceedings of the National Academy of Sciences*, 98(2):676–682.
- Rasmussen, P. M., Madsen, K. H., Lund, T. E., and Hansen, L. K. (2011). Visualization of nonlinear kernel models in neuroimaging by sensitivity maps. *NeuroImage*, 55(3):1120–1131.
- Rentz, D. M., Locascio, J. J., Becker, J. A., Moran, E. K., Eng, E., Buckner, R. L., Sperling, R. A., and Johnson, K. A. (2010). Cognition, reserve, and amyloid deposition in normal aging. *Annals of neurology*, 67(3):353–364.
- Ritchie, K. and Kildea, D. (1995). Is senile dementia “age-related” or “ageing-related” - evidence from meta-analysis of dementia prevalence in the oldest old. *The Lancet*, 346(8980):931–934.

- Rosvall, M. and Bergstrom, C. T. (2008). Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences*, 105(4):1118–1123.
- Rowe, J. W. and Kahn, R. L. (1987). Human aging: usual and successful. *Science*, 237(4811):143–149.
- Saad, Z. S., Gotts, S. J., Murphy, K., Chen, G., Jo, H. J., Martin, A., and Cox, R. W. (2012). Trouble at rest: how correlation patterns and group differences become distorted after global signal regression. *Brain connectivity*, 2(1):25–32.
- Sabuncu, M. R., Balci, S. K., Shenton, M. E., and Golland, P. (2009). Image-driven population analysis through mixture modeling. *Medical Imaging, IEEE Transactions on*, 28(9):1473–1487.
- Sakoğlu, Ü., Pearlson, G., Kiehl, K., Wang, Y., et al. (2010). A method for evaluating dynamic functional network connectivity and task-modulation: application to schizophrenia. *Magnetic Resonance Materials*, 23(5):351–366.
- Salami, A., Pudas, S., and Nyberg, L. (2014). Elevated hippocampal resting-state connectivity underlies deficient neurocognitive function in aging. *Proceedings of the National Academy of Sciences*, 111(49):17654–17659.
- Satterthwaite, T. D., Elliott, M. A., Gerraty, R. T., Ruparel, K., Loughead, J., Calkins, M. E., Eickhoff, S. B., Hakonarson, H., Gur, R. C., Gur, R. E., et al. (2012). An improved framework for confound regression and filtering for control of motion artifact in the preprocessing of resting-state functional connectivity data. *Neuroimage*.
- Satterthwaite, T. D., Elliott, M. A., Ruparel, K., Loughead, J., Prabhakaran, K., Calkins, M. E., Hopson, R., Jackson, C., Keefe, J., Riley, M., Mentch, F. D.,

- Sleiman, P., Verma, R., Davatzikos, C., Hakonarson, H., Gur, R. C., and Gur, R. E. (2014). Neuroimaging of the philadelphia neurodevelopmental cohort. *NeuroImage*, 86:544–553.
- Satterthwaite, T. D., Wolf, D. H., Ruparel, K., Erus, G., Elliott, M. A., Eickhoff, S. B., Gennatas, E. D., Jackson, C., Prabhakaran, K., Smith, A., et al. (2013). Heterogeneous impact of motion on fundamental patterns of developmental changes in functional connectivity during youth. *NeuroImage*, 83:45–57.
- Seeley, W. W., Crawford, R. K., Zhou, J., Miller, B. L., and Greicius, M. D. (2009). Neurodegenerative diseases target large-scale human brain networks. *Neuron*, 62(1):42–52.
- Seeley, W. W., Menon, V., Schatzberg, A. F., Keller, J., Glover, G. H., Kenna, H., Reiss, A. L., and Greicius, M. D. (2007). Dissociable intrinsic connectivity networks for salience processing and executive control. *The Journal of neuroscience*, 27(9):2349–2356.
- Shock, N. W. et al. (1984). Normal human aging: The baltimore longitudinal study of aging.
- Sivalingam, R., Boley, D., Morellas, V., and Papanikolopoulos, N. (2011). Positive definite dictionary learning for region covariances. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 1013–1019. IEEE.
- Smith, S., Miller, K., Moeller, S., Xu, J., Auerbach, E., Woolrich, M., Beckmann, C., Jenkinson, M., Andersson, J., Glasser, M., et al. (2012). Temporally-independent functional modes of spontaneous brain activity. *Proceedings of the National Academy of Sciences*, 109(8):3131–3136.

- Smith, S. M. (2002). Fast robust automated brain extraction. *Human brain mapping*, 17(3):143–155.
- Smith, S. M., Miller, K. L., Salimi-Khorshidi, G., Webster, M., Beckmann, C. F., Nichols, T. E., Ramsey, J. D., and Woolrich, M. W. (2011). Network modelling methods for fmri. *Neuroimage*, 54(2):875–891.
- Song, Y., Cai, W., Huang, H., Zhou, Y., Feng, D., Wang, Y., Fulham, M., and Chen, M. (2010). Large margin local estimate with applications to medical image classification.
- Soufiani, H. A. and Airoidi, E. M. (2012). Graphlet decomposition of a weighted network. *arXiv preprint arXiv:1203.2821*.
- Sperling, R. (2011). The potential of functional mri as a biomarker in early alzheimer’s disease. *Neurobiology of aging*, 32:S37–S43.
- Sra, S. and Cherian, A. (2011). Generalized dictionary learning for symmetric positive definite matrices with application to nearest neighbor retrieval. *Machine Learning and Knowledge Discovery in Databases*, pages 318–332.
- Steinbach, M., Ertöz, L., and Kumar, V. (2004). The challenges of clustering high dimensional data. In *New Directions in Statistical Physics*, pages 273–309. Springer.
- Tomasi, D. and Volkow, N. D. (2012). Aging and functional brain networks. *Molecular psychiatry*, 17(5):549–558.
- Traag, V. and Bruggeman, J. (2009). Community detection in networks with positive and negative links. *Physical Review E*, 80(3):036115.
- Van Essen, D. (2005). A population-average, landmark-and surface-based (pals) atlas of human cerebral cortex. *Neuroimage*, 28(3):635–662.

- Van Essen, D. C., Drury, H. A., Dickson, J., Harwell, J., Hanlon, D., and Anderson, C. H. (2001). An integrated software suite for surface-based analyses of cerebral cortex. *Journal of the American Medical Informatics Association*, 8(5):443–459.
- Van Essen, D. C., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T., Bucholz, R., Chang, A., Chen, L., Corbetta, M., Curtiss, S., et al. (2012). The human connectome project: a data acquisition perspective. *Neuroimage*, 62(4):2222–2231.
- Varol, E., Sotiras, A., and Davatzikos, C. (2015). Disentangling disease heterogeneity with max-margin multiple hyperplane classifier. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015*, pages 702–709. Springer.
- Venkataraman, A., Whitford, T. J., Westin, C.-F., Golland, P., and Kubicki, M. (2012). Whole brain resting state functional connectivity abnormalities in schizophrenia. *Schizophrenia research*, 139(1):7–12.
- Vergun, S., Deshpande, A. S., Meier, T. B., Song, J., Tudorascu, D. L., Nair, V. A., Singh, V., Biswal, B. B., Meyerand, M. E., Birn, R. M., et al. (2013). Characterizing functional connectivity differences in aging adults using machine learning on resting state fmri data. *Frontiers in computational neuroscience*, 7.
- Vincent, J. L., Kahn, I., Snyder, A. Z., Raichle, M. E., and Buckner, R. L. (2008). Evidence for a frontoparietal control system revealed by intrinsic functional connectivity. *Journal of neurophysiology*, 100(6):3328–3342.
- Vinje, W. E. and Gallant, J. L. (2000). Sparse coding and decorrelation in primary visual cortex during natural vision. *Science*, 287(5456):1273–1276.
- von dem Hagen, E. A., Stoyanova, R. S., Baron-Cohen, S., and Calder, A. J. (2012). Reduced functional connectivity within and between 'social' resting state networks in autism spectrum conditions. *Social cognitive and affective neuroscience*.

- Wu, J.-T., Wu, H.-Z., Yan, C.-G., Chen, W.-X., Zhang, H.-Y., He, Y., and Yang, H.-S. (2011). Aging-related changes in the default mode network and its anti-correlated networks: a resting-state fmri study. *Neuroscience letters*, 504(1):62–67.
- Wu, T., Zang, Y., Wang, L., Long, X., Hallett, M., Chen, Y., Li, K., and Chan, P. (2007). Aging influence on functional connectivity of the motor network in the resting state. *Neuroscience letters*, 422(3):164–168.
- Xia, M., Wang, J., and He, Y. (2013). Brainnet viewer: A network visualization tool for human brain connectomics. *PloS one*, 8(7):e68910.
- Yan, L., Zhuo, Y., Wang, B., and Wang, D. J. (2011). Loss of coherence of low frequency fluctuations of bold fmri in visual cortex of healthy aged subjects. *The open neuroimaging journal*, 5(1).
- Yeo, B. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., Roffman, J. L., Smoller, J. W., Zöllei, L., Polimeni, J. R., et al. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of neurophysiology*, 106(3):1125–1165.
- Ylikoski, R., Ylikoski, A., Keskivaara, P., Tilvis, R., Sulkava, R., and Erkinjuntti, T. (1999). Heterogeneity of cognitive profiles in aging: successful aging, normal aging, and individuals at risks for cognitive decline. *European Journal of Neurology*, 6(6):645–652.
- Zang, Y., Jiang, T., Lu, Y., He, Y., and Tian, L. (2004). Regional homogeneity approach to fmri data analysis. *Neuroimage*, 22(1):394–400.
- Zhang, Y., Brady, M., and Smith, S. (2001). Segmentation of brain mr images

through a hidden markov random field model and the expectation-maximization algorithm. *Medical Imaging, IEEE Transactions on*, 20(1):45–57.