

**SPATIO-TEMPORAL MIXED MODELS FOR
DIFFUSION TENSOR MAGNETIC RESONANCE
IMAGING**

by

John A. Scott

B.A., Sarah Lawrence College, 1998

A.M., Washington University in St. Louis, 2001

Submitted to the Graduate Faculty of
the Graduate School of Public Health in partial fulfillment
of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2008

UNIVERSITY OF PITTSBURGH
GRADUATE SCHOOL OF PUBLIC HEALTH

This dissertation was presented

by

John A. Scott

It was defended on

August 8, 2008

and approved by

Dissertation Advisor: Nicholas Lange, Sc.D.
Adjunct Associate Professor, Department of Biostatistics
Graduate School of Public Health
University of Pittsburgh
Associate Professor, Department of Psychiatry
Harvard Medical School
Harvard University

Committee Member: Roger Day, Sc.D.
Associate Professor, Department of Biostatistics
Graduate School of Public Health
University of Pittsburgh

Committee Member: Stewart Anderson, Ph.D.
Associate Professor, Department of Biostatistics
Graduate School of Public Health
University of Pittsburgh

Committee Member: Greg Siegle, Ph.D.
Associate Professor, Department of Psychiatry
School of Medicine
University of Pittsburgh

Committee Member: Raymond Cho, M.D.
Assistant Professor, Department of Psychiatry
School of Medicine
University of Pittsburgh

Copyright © by John A. Scott
2008

SPATIO-TEMPORAL MIXED MODELS FOR DIFFUSION TENSOR MAGNETIC RESONANCE IMAGING

John A. Scott, PhD

University of Pittsburgh, 2008

Diffusion tensor imaging (DTI) is a magnetic resonance imaging modality that provides useful *in vivo* information about the microstructure of human brain tissue, particularly the white matter structures that comprise the ‘wiring’ of the brain. DTI holds great promise for enhancing our understanding of white matter disorders, which comprise public health burdens in a variety of medical domains. Due to its relatively complex structure, however, extracting useful information from DTI data presents a number of statistical challenges. More effective statistical methodologies will improve the sensitivity of DTI data analyses and increase their clinical relevance, a goal of substantial public health significance.

In this dissertation, I propose a series of analytic approaches to DTI data analysis based on linear mixed effects models (LMEs). These models provide a number of advantages over several expedient DTI data analyses in current use. I demonstrate the applicability and advantages of my LME-based approaches in an analysis that compares white matter microstructure in a group of children and young adults with autism spectrum disorders (ASDs) to typically developing controls.

I first identify a class of LMEs for DTI data analyses for which closed-form maximum likelihood estimators of all parameters exist. By avoiding iteration, these models enable practitioners to perform exploratory and confirmatory analyses of large DTI datasets in clinically feasible time. This family of models incorporates group heterogeneity in variance-covariance structure.

I then compare the results of my approach with other approaches currently in practice in an analysis of white matter abnormalities associated with ASDs. I also introduce a data analytic framework that incorporates the entire multivariate tensor in a single analysis.

I last describe a unified likelihood-based approach to addressing reliability with a new estimator of a generalized intraclass correlation coefficient. I establish the robustness of this approach to model perturbations with a series of theoretical and simulation results and apply it to quantify local spatial reliability in the ASDs example.

TABLE OF CONTENTS

PREFACE	xiii
1.0 INTRODUCTION	1
1.1 WHITE MATTER ABNORMALITIES IN AUTISM	3
1.2 DIFFUSION TENSOR IMAGING	4
1.2.1 Applications Of DTI	5
1.2.2 DTI Procedures And Data Processing	5
1.2.2.1 Diffusion Weighted Imaging	6
1.2.2.2 Fitting Tensors To DWIs	7
1.2.2.3 Calculation Of Scalar Quantities From Tensors	10
1.3 DTI DATA ANALYSIS	13
1.3.1 Analyses Based On Regions Of Interest	14
1.3.2 Voxelwise Analyses	16
1.4 PROPOSED STATISTICAL APPROACHES FOR DTI	17
2.0 CLOSED-FORM SPATIO-TEMPORAL NEUROIMAGING MODELS	19
2.1 TENSOR ELEMENT DECORRELATION	20
2.1.1 Disadvantages Of Scalar Reduction	20
2.1.2 Tensor-Valued Normal Distribution	21
2.1.3 Tensor Element Correlation	22
2.2 CONDITIONS FOR THE EXISTENCE OF CLOSED-FORM ML ESTI- MATORS	24
2.2.1 Normal Estimation Under Additive Mean And Covariance	25
2.2.2 Closed-Form Estimation Of Mean Parameters	26

2.2.3	Closed-Form Estimation Of Covariance Parameters	27
2.3	CLOSED-FORM ML ESTIMATORS UNDER COMPOUND SYMMETRY	28
2.3.1	Notation	28
2.3.2	Overall Scalar Mean	29
2.3.3	Group Scalar Mean	34
2.3.4	Voxelwise Scalar Means	34
2.3.5	The General Case For Homogeneous Compound Symmetry	36
2.3.6	REML Estimation	39
2.3.7	Summary	40
2.4	CLOSED-FORM ML ESTIMATION UNDER HETEROGENEOUS VARI- ANCE	40
2.4.1	Group Heterogeneity In Residual Variance	41
2.4.2	Group Heterogeneity In Between-Subject Variance	43
2.4.2.1	Closed-Form Estimation Of Mean Parameters	43
2.4.2.2	Closed-Form Estimation Of Covariance Parameters	44
2.4.3	Combined Residual And Between-Subjects Variance Heterogeneity .	45
2.4.4	The General Case For Heterogeneous Compound Symmetry	47
2.4.5	Nonexistence Of Closed-Form Estimators And Implications	51
2.4.6	Three Or More Groups	51
2.4.7	Summary	52
2.5	CLOSED-FORM ML ESTIMATION FOR LONGITUDINAL MODELS .	53
2.5.1	Longitudinal Model With Random Intercept	53
2.5.2	Longitudinal Model With Random Slope	56
2.5.3	General Orthogonal Polynomial Random Effects Models	60
2.5.4	Multiple Groups And Variance Heterogeneity	61
2.6	CONCLUSIONS	62
2.6.1	Nonexistence Of Closed-Form Estimators	63
3.0	APPLICATIONS TO NEUROIMAGING STUDIES OF AUTISM . .	65
3.1	DATA DESCRIPTION	65
3.2	FRACTIONAL ANISOTROPY ANALYSES	66

3.2.1	Brain Regions Explored	68
3.2.1.1	Region 1	68
3.2.1.2	Region 2	69
3.2.1.3	Region 3	70
3.2.2	Data Analysis	70
3.2.2.1	Group Comparison Of Regional Means	71
3.2.2.2	Group Comparison Of Voxelwise Means	72
3.2.2.3	LME With Homogeneous Variance	73
3.2.2.4	LME With Heterogeneous Between-Subject Variance	74
3.2.2.5	LME With Combined Heterogeneous Variance	76
3.2.3	Results	77
3.2.3.1	Region 1	77
3.2.3.2	Region 2	80
3.2.3.3	Region 3	81
3.3	FULL TENSOR ANALYSES	81
3.3.1	Decorrelating Transformation	86
3.3.2	Data Analyses	87
3.3.2.1	Regional Multivariate Means	88
3.3.2.2	Multivariate Voxelwise Means	88
3.3.2.3	Multiple Linear Mixed Effects Models	89
3.3.3	Results	90
3.3.3.1	Region 1	93
3.3.3.2	Region 2	93
3.3.3.3	Region 3	96
3.4	DISCUSSION	99
3.4.1	LMEs Compared To Traditional Approaches	99
3.4.2	Full Tensor Analyses Compared To FA Analyses	101
3.4.3	Homogeneous Compared To Heterogeneous Variance Models	102
3.4.4	Conclusions	104
4.0	GENERALIZED ICC IN NEUROIMAGING MODELS	105

4.1	Reliability	105
4.1.1	Local Spatial Reliability	106
4.1.2	Other Applications	107
4.2	INTRACLASS CORRELATION COEFFICIENT BACKGROUND	107
4.2.1	ICC Definition	108
4.2.2	ICC Estimation	108
4.2.3	Generalized ICC	109
4.3	GENERALIZED ICC IN GROWTH CURVE MODELS	110
4.3.1	Notation And Assumptions	111
4.3.2	Preliminary Results	112
4.3.3	Invariance Of The Generalized ICC To Number Of Random Effects	113
4.3.4	Autoregressive Error	114
4.3.5	Simulations	115
4.4	GENERALIZED ICC APPLICATION	117
4.4.1	Methods	118
4.4.2	Results	119
4.5	CONCLUSIONS	125
	APPENDIX. PROOFS OF THEOREMS IN CHAPTER 4	129
A.1	PROOF OF THEOREM 4.3.4	129
A.2	PROOF OF THEOREM 4.3.6	130
	BIBLIOGRAPHY	133

LIST OF TABLES

3.1	Demographic and clinical characteristics of the sample	66
3.2	Likelihood ratio tests comparing FA analyses under Models 3–5	103
3.3	Likelihood ratio tests comparing full tensor analyses under Models 3–5	103
4.1	AR(1) simulation results 1	116
4.2	AR(1) simulation results 2	116
4.3	AR(1) simulation results 3	116
4.4	Mean and standard deviation of gICC in the midaxial slice	123

LIST OF FIGURES

3.1	FA mean for 80 ASDs and 40 control participants	67
3.2	ASDs mean FA - control mean FA	68
3.3	FA variance for 80 ASDs and 40 control participants	69
3.4	Mean FA in Region 1	69
3.5	Mean FA in Region 2	70
3.6	Mean FA in Region 3	71
3.7	Group difference t-statistics for FA in Region 1	78
3.8	Group difference p-values for FA in Region 1	79
3.9	Group difference t-statistics for FA in Region 2	82
3.10	Group difference p-values for FA in Region 2	83
3.11	Group difference t-statistics for FA in Region 3	84
3.12	Group difference p-values for FA in Region 3	85
3.13	Group tensor difference T-squared statistics in Region 1	91
3.14	Group tensor difference p-values in Region 1	92
3.15	Group tensor difference T-squared statistics in Region 2	94
3.16	Group tensor difference p-values in Region 2	95
3.17	Group tensor difference T-squared statistics in Region 3	97
3.18	Group tensor difference p-values in Region 3	98
4.1	gICC maps for FA, control participants	120
4.2	gICC maps for sRA, control participants	120
4.3	gICC maps for VR, control participants	121
4.4	gICC maps for FA, ASDs participants	121

4.5	gICC maps for sRA, ASDs participants	122
4.6	gICC maps for VR, ASDs participants	122
4.7	gICC maps comparison, midaxial slice	124
4.8	gICC ratio maps, midaxial slice	126

PREFACE

ACKNOWLEDGEMENTS

I would like to thank several people whose advice and assistance have been invaluable in completing this dissertation.

First, this work would not have been possible without the advice and support of my dissertation advisor, Dr. Nicholas Lange. I am extremely grateful for his guidance and mentorship.

I am also very grateful to my academic advisor, Dr. Roger Day, for his support and insightful comments about the dissertation and much else.

I would also like to thank my other dissertation committee members, Dr. Stewart Anderson, Dr. Greg Siegle and Dr. Raymond Cho. Their careful reading and suggestions have consistently improved this dissertation through successive drafts.

I would like to thank the National Institutes of Health, which provided my funding while I worked on this dissertation through grant MH73451, “Training Biostatisticians in Psychiatric Research.” I am especially grateful to Dr. Sati Mazumdar who, as P.I. of that training grant and *de facto* “third advisor,” has provided me with ongoing support and guidance throughout this process. Dr. Lange’s funding was provided by grant NS37483.

A number of people have been very generous with their time and resources in providing or assisting me with the neuroimaging data analyzed in Chapters 3 and 4. The data were collected under the auspices of Dr. Janet Lainhart’s lab at the Utah Autism Research Program, and I am very grateful to her for making the data available to me for this dissertation. And I’d like to thank Alyson Froelich for her help in ironing out data issues as they arose. Dr. Mariana Lazar of the Center for Biomedical Imaging at NYU provided considerable help in processing and transferring the data to me, and I am very grateful for her time and hard

work. And Dr. Derek Jones of Cardiff University was very helpful in providing materials for learning about DTI during the early stages of my research.

DEDICATION

This dissertation is dedicated to my beloved wife, Layla, without whom it would have been completely impossible. It is also dedicated to our son, Victor, without whom it frankly would have been a little easier, but much less fun.

1.0 INTRODUCTION

A large number of neurological and psychiatric conditions are associated with changes in the cerebral white matter. Such disorders include primarily white matter diseases, both demyelinating diseases (e.g. multiple sclerosis) and dysmyelinating diseases (e.g. leukodystrophies), as well as loss of white matter integrity secondary to trauma, ischemia or tumor. A number of psychiatric disorders are believed to involve disconnectivities between brain regions, and white matter abnormalities have been observed in conditions as varied as Alzheimer’s disease[96], obsessive-compulsive disorder[91], schizophrenia[48] and alcohol use[59]. Conversely, schizophrenia-like psychotic symptoms have been observed in a variety of white matter diseases[93]. Together, these diseases comprise a massive public health burden. Techniques which make it possible to identify and localize white matter abnormalities are critical to fully understanding the etiology and course of these illnesses.

Diffusion tensor imaging (DTI) is a relatively new magnetic resonance imaging modality that makes it possible to visualize white matter *in vivo*[57]. DTI has already shown great promise in illuminating the mechanisms underlying many neurological and psychiatric disorders. However, the statistical approaches to DTI data in practice to date have been driven to a very large degree by convenience and computational simplicity. The major drawbacks of common analytic approaches include disregarding within-subject variability, potential heterogeneity and autocorrelation, limiting sensitivity to mean effects across predetermined regions of interest, failing to adequately control for sources of variation, and inflexibility in detecting and handling group heterogeneity in variance structure (see Section 1.3). The potential consequences of a lack of a principled statistical approach to such data are twofold: (1) results reported in the literature may not be reproducible (false positives), and (2) legitimate white matter abnormalities in imaged brains may go undetected.

I propose a new approach to the statistical modeling of DTI data based on linear mixed effects (LME) models. The use of LMEs helps to overcome many of the potential shortcomings of standard approaches. Within-subject variability can be modeled in LMEs, rather than ignored, potentially accounting for spatial autocorrelation. By borrowing strength across multiple voxels, LMEs make it possible to detect voxel-scale effects while controlling for between-subject covariates. And it is possible to test for and incorporate group heterogeneity in variance structures with LMEs. The major disadvantage of LMEs relative to simpler statistical models for DTI data is that parameter estimation generally requires computationally-intensive iterative procedures, which may be infeasible with large neuroimaging datasets, especially in an exploratory, model-building analysis. Thus, there is a premium on identifying special cases of LMEs for which closed-form (no iteration) likelihood estimation is available.

In Chapter 1, I provide an overview of DTI and common statistical approaches to DTI data. I will introduce a class of linear mixed effects models with closed-form maximum likelihood estimation, including models allowing for group heterogeneity in covariance structure, in Chapter 2. I focus on cross-sectional and longitudinal models designed for two-group comparisons, as these are particularly relevant to identifying white matter abnormalities associated with psychiatric disorders (i.e., by comparing disease to control groups).

In Chapter 3, these models are illustrated with applications to a neuroimaging study comparing young adults with autism spectrum disorders to normally developing controls. I will contrast group comparison results obtained from closed-form LMEs with results obtained from typical DTI data analytic approaches.

Finally, I present an approach to evaluating reliability in neuroimaging studies based on a generalization of the intraclass correlation coefficient (Chapter 4). The focus here is not on group comparisons alone. The techniques introduced are closely tied statistically to the material in Chapters 2 and 3, and reliability considerations are relevant to a broad range of data analytic questions of concern to neuroimaging practitioners.

1.1 WHITE MATTER ABNORMALITIES IN AUTISM

Autism spectrum disorders (ASDs) encompass a range of pervasive developmental disorders (autism, Asperger’s disorder and pervasive developmental disorder not otherwise specified), all involving deficits in communication and social interaction, and repetitive behaviors. Based on the prevalence, difficulty of treatment, and degree of impairment to individuals with autism and their families (which can range from mild to very severe, and is often life-long), ASDs comprise a massive public health burden. It is notoriously difficult to obtain a complete picture of the prevalence of ASDs in the U.S. population, apparently due largely to ongoing changes in surveillance and diagnostic criteria. However, most recent estimates of the prevalence of ASDs have been in the vicinity of 60 cases per 10,000 children, with estimates of autism itself falling mostly in the range of 10 to 20 cases per 10,000[71].

In order to develop more effective treatments for autism, it is critically important to understand the neurological underpinnings of autism. The neurobiology of ASDs appears to be complex and is, as yet, poorly understood[82]. A recent review cited a number of findings of anatomical abnormalities in ASDs patients, the most consistent of which were increased brain volume, decreased numbers of Purkinje cells in the cerebellum, and cerebral cortex dysgenesis[26]. A number of functional magnetic resonance imaging (fMRI) studies have also found abnormalities in ASDs associated with functional activation of the brain in various tasks, including tasks involving social attribution, sentence comprehension and working memory, and face perception[49].

Some of the behavioral deficits characteristic of ASDs, as well as some of the functional neuroimaging studies of ASDs, seem to suggest underlying disconnectivity issues, and investigating white matter structures in ASDs patients is a natural development. However, the exact nature of these disconnectivities remains an open question as to date there have been only a handful of DTI studies comparing ASDs patients to controls to investigate white matter abnormalities *in vivo*. These studies have variously found white matter abnormalities associated with ASDs in the corpus callosum[2, 42], the superior temporal gyrus and temporal stem[58], and in the ventromedial prefrontal cortex and anterior cingulate gyrus[8]. More work is needed to pin down the areas of most consistent anatomical disconnectivity associ-

ated with ASDs, and more sophisticated modeling techniques than those typically employed by practitioners may increase the sensitivity of DTI to detect such abnormalities and reduce the number of spurious findings. Specifically, an LME-based approach will help reduce bias in estimates of group difference parameters due to covariance structure misspecification. LMEs will also make it possible to adequately control for important covariates by borrowing strength across voxels, while still retaining sensitivity to voxel-scale effects, unlike region of interest-based analyses.

1.2 DIFFUSION TENSOR IMAGING

Diffusion tensor imaging (DTI) is a magnetic resonance imaging technique designed to take advantage of the diffusion properties of water in human tissue[12, 57, 56]. Due to Brownian motion, water molecules in a purely liquid medium such as cerebrospinal fluid (CSF) will tend to diffuse according to a Gaussian distribution characterized by the viscosity and temperature of the medium[17]. Likewise, in a liquid medium that also contains tissue that is relatively sparse, loosely organized or highly permeable to water (such as the gray matter cells of the brain), water molecules will tend to diffuse equally in all directions on average; this is referred to as isotropic diffusion. However, in dense, highly organized, relatively impermeable tissue (such as the major white matter tracts of the brain), water molecules will preferentially diffuse parallel to the principal orientation of the tissue, and will show less diffusion perpendicular to the tissue; this is called anisotropic diffusion.

Because it is sensitive to anisotropic diffusion, DTI is ideally suited for imaging dense, highly organized tissue *in vivo*. Furthermore, since DTI makes it possible to measure both the degree of anisotropy and the primary directions of diffusion, the information recovered from DTI scans can also be used to trace the path of highly organized, fibrous tissue throughout the human body, a process called tractography[15].

1.2.1 Applications Of DTI

Although DTI has been used to image a variety of tissues including skeletal muscle[33], cardiac muscle[38], and kidney[79], its greatest impact has been in the *in vivo* imaging of white matter in the brain. DTI is sensitive to diffusion anisotropy due to intact axonal membranes and to myelination[17], with at least one recent study suggesting that it is possible to distinguish changes in anisotropy due to axonal membrane collapse from changes due to demyelination[23]. This sensitivity has made DTI invaluable for researching white matter diseases such as multiple sclerosis[34, 60, 62] and adrenomyeloneuropathy[29], as well as other neurological disorders that can have a profound impact on white matter structures, including amyotrophic lateral sclerosis (ALS)[1, 5, 21] and epilepsy[22, 92]. DTI has also proved useful in investigating axonal injury due to trauma[30, 81] or ischemic events[87], and in mapping axonal displacement due to brain tumors[83]. Researchers have also used DTI to examine the white matter connections underlying normal brain processes such as reading[18], language processing[10, 77] and cognitive control[61]. The technology is also beginning to see clinical use[89], including use in presurgical planning for brain tumors[97].

A wide range of psychiatric disorders also have white matter involvement which can be investigated with DTI[93, 59]. To cite a few examples, Kubicki et al. (2002) found reduced asymmetry in the diffusion anisotropy of the uncinate fasciculus in schizophrenia patients compared to controls[48]. Xie et al. (2005) found reduced diffusion anisotropy in bilateral cingulum bundles of patients with Alzheimer’s disease compared to controls[96]. And Szeszko et al. (2005) found reduced anisotropy in the white matter of the anterior cingulate gyrus in obsessive-compulsive patients relative to controls[91].

1.2.2 DTI Procedures And Data Processing

The processing stream for a DTI study involves several steps that I describe in detail in the following sections. Briefly, the process begins with the acquisition of diffusion weighted images (DWIs, described below). A tensor field (Section 1.2.2.2) can be estimated from these DWIs, and there is a wide variety of ways in which to visualize the tensor field. Many of these visualization approaches involve calculation of scalar quantities, including anisotropy indices

and indices of diffusivity. These scalars are also used directly for data analytic purposes. For the most part, analytic techniques that are meant to reflect the geometry of white matter structures rely on diagonalizing the diffusion tensors to obtain their spectral decomposition.

1.2.2.1 Diffusion Weighted Imaging Diffusion weighted imaging is a magnetic resonance imaging (MRI) technique that is sensitive to molecular diffusion. A thorough description of MRI was provided in the statistical literature by Lange (1996)[51]. Briefly, MRI depends on placing the subject to be scanned in a high-strength static magnetic field. This magnetic field causes hydrogen protons in the target tissue to align and precess about their axes at a known frequency. A second, much weaker, magnetic pulse is then applied to the target tissue, which causes some of the protons to enter a higher-energy state. After this pulse has been applied, the protons will gradually release this excess energy, which is detected by the scanner.

DWI is a modification of standard MRI that uses additional magnetic pulses applied at gradients to achieve sensitivity to diffusion. There are several methodologies for obtaining DWIs, the most common of which is the Stejskal-Tanner technique[7]. The Stejskal-Tanner method involves the application of identical rectangular magnetic gradient pulses. This has the effect of dephasing and then rephasing protons, and those protons that have moved subsequent to the first, dephasing gradient pulse will have experienced random changes in their spin phase at the time of the rephasing pulse[68]. This leads to a signal attenuation. On the other hand, for protons that have experienced no motion, the dephasing and rephasing pulses cancel each other, and there is no net change in signal.

A single pair of identical gradient pulses sensitizes the signal to diffusion in a single direction. The amount of diffusion is quantified by the apparent diffusion coefficient (ADC), which is related to the signal strength S by the Stejskal-Tanner relationship:

$$S = S_0 \exp\left(-\frac{\text{TE}}{T_2}\right) \exp\left(-\gamma^2 G^2 \delta^2 \left(\Delta - \frac{\delta}{3}\right) \text{ADC}\right), \quad (1.1)$$

where S_0 is the signal with no diffusion-encoding pulse, TE is the echo time, T_2 is the transverse relaxation time, γ is the gyromagnetic ratio (4,258 Hz/G for hydrogen nuclei), G is the amplitude and δ the duration of the gradient pulse, and Δ is the time interval between

the gradient pulses. The term $\gamma^2 G^2 \delta^2 (\Delta - \frac{\delta}{3})$ is denoted by b and referred to as the b -factor. The ADC can thus be calculated as:

$$\text{ADC} = -\frac{1}{b} \log \left(\frac{S}{S_0} \right). \quad (1.2)$$

A DWI is constituted by measurements of ADCs at each voxel in the scanned space. Although I will be concerned with diffusion tensor images rather than the diffusion weighted images which form their basis, it is worth noting that DWIs themselves have important clinical applications. In particular, DWIs make it possible to identify even very small ischemic lesions within minutes of a stroke[69].

1.2.2.2 Fitting Tensors To DWIs A single DWI is sensitive to diffusion in one direction. This property means that a single DWI is adequate for imaging isotropic diffusion, but is unable to identify anisotropic diffusion. Diffusion tensor imaging is an approach that uses multiple DWIs, acquired from at least six non-collinear directions, to obtain a more complete image of diffusion, including possible anisotropy. In DTI, instead of a single ADC, diffusion is quantified by a diffusion tensor — a 3×3 positive definite symmetric (p.d.s.) matrix[12]:

$$\mathbf{D} = \begin{pmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{xy} & D_{yy} & D_{yz} \\ D_{xz} & D_{yz} & D_{zz} \end{pmatrix}. \quad (1.3)$$

Here, D_{xx} , D_{yy} and D_{zz} represent the average diffusion (over molecules) along three orthogonal axes and D_{xy} , D_{xz} and D_{yz} represent the covariances between the diffusion in the x , y and z directions. This leads to an expansion of Equation 1.2 above:

$$\log \left(\frac{S}{S_0} \right) = -\mathbf{b} : \mathbf{D} \quad (1.4)$$

$$= -(b_{xx}D_{xx} + b_{yy}D_{yy} + b_{zz}D_{zz} + 2b_{xy}D_{xy} + 2b_{xz}D_{xz} + 2b_{yz}D_{yz}). \quad (1.5)$$

where \mathbf{b} is now a known symmetric 3x3 matrix specified by the researcher and $(\cdot : \cdot)$ denotes the generalized dot product, $\mathbf{x} : \mathbf{y} = \text{Vec}(\mathbf{x}) \cdot \text{Vec}(\mathbf{y})$. Analytical expressions for the elements

of the \mathbf{b} -matrix in terms of imaging parameters are available[65], and depend on the direction and strength of each gradient encoding pulse.

Equation 1.4 has six unknowns and requires at least six distinct images acquired with different diffusion-encoding gradient pulse directions (in addition to the S_0 image acquired with no gradient pulse) to estimate \mathbf{D} . In the case when there are DWIs acquired with gradients from exactly six directions available, there is a unique solution for \mathbf{D} . There have been several different estimation approaches proposed for the situation in which there are more than six diffusion-encoding directions available[46, 63], but the least squares and weighted least squares approaches (described below) are the most common.

I follow Kingsley's notation in describing least squares diffusion tensor estimation[46]. I assume there are N acquired images. The tensor elements and the log of the non-diffusion encoded S_0 image comprise a column vector:

$$\boldsymbol{\alpha} = (D_{xx} \ D_{yy} \ D_{zz} \ D_{xy} \ D_{xz} \ D_{yz} \ \log S_0)^T. \quad (1.6)$$

Each individual \mathbf{b} -matrix is included as a row in a design matrix \mathbf{B} , along with a column of 1s for the $\log S_0$ term:

$$\mathbf{B} = \begin{pmatrix} -b_{xx1} & -b_{yy1} & -b_{zz1} & -2b_{xy1} & -2b_{xz1} & -2b_{yz1} & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ -b_{xxN} & -b_{yyN} & b_{zzN} & -2b_{xyN} & -2b_{xzN} & -2b_{yzN} & 1 \end{pmatrix}. \quad (1.7)$$

And the observed log signal intensities are collected in a column matrix,

$$\mathbf{x} = (\log S_1 \ \log S_2 \ \dots \ \log S_N)^T. \quad (1.8)$$

Based on Equation 1.4, the system of equations can be expressed as

$$\mathbf{x} = \mathbf{B}\boldsymbol{\alpha} + \boldsymbol{\eta}, \quad (1.9)$$

where $\boldsymbol{\eta}$ is an error term.

Again, in the situation in which \mathbf{B} is invertible, with exactly six distinct diffusion-encoding gradients and one $\mathbf{b} = 0$ image, there is a unique solution:

$$\hat{\boldsymbol{\alpha}}_0 = \mathbf{B}^{-1}\mathbf{x}. \quad (1.10)$$

In the more general case, a least squares estimator is possible:

$$\hat{\boldsymbol{\alpha}}_{\text{ls}} = (\mathbf{B}^T \mathbf{B})^{-1} \mathbf{B}^T \mathbf{x}. \quad (1.11)$$

By the Gauss-Markov theorem, the least squares estimator is the best linear unbiased estimator (BLUE) when the errors are homoscedastic (that is, when $\text{Var}(\boldsymbol{\eta}) = \sigma^2 \mathbf{I}$) [84]. However, while it is generally assumed that the errors in the original signal units, S_1, \dots, S_N are homoscedastic, this implies that the errors are heteroscedastic in the log-signal units, $\log S_1, \dots, \log S_N$. Specifically, if $S_i \sim N(\exp(\mathbf{B}_i \boldsymbol{\alpha}), \sigma_i^2)$ then, by the delta method,

$$\log S_i \sim N(\mathbf{B}_i \boldsymbol{\alpha}, \frac{\sigma_i^2}{S_i^2}).$$

Thus, the BLUE of $\boldsymbol{\alpha}$ is the weighted least squares estimate:

$$\hat{\boldsymbol{\alpha}}_{\text{wls}} = (\mathbf{B}^T \boldsymbol{\Sigma}^{-1} \mathbf{B})^{-1} \mathbf{B}^T \boldsymbol{\Sigma}^{-1} \mathbf{x}, \quad (1.12)$$

where $\boldsymbol{\Sigma}$ is a diagonal matrix whose (i, i) th entry is $\frac{\sigma_i^2}{S_i^2}$. In most applications it is assumed that $\sigma_i^2 = \sigma^2$ for all i . Several authors have proposed additional methodologies for estimating D , including non-linear least squares [73] and Geman-McClure M-estimators [63], but such methods are not widely used.

Although the entries of the diffusion tensor are occasionally of interest in themselves (at least the diagonal entries, which represent diffusion along the coordinate axes), the most important properties of the tensor \mathbf{D} are more clearly seen via diagonalization, $\mathbf{D} = \mathbf{E} \boldsymbol{\Lambda} \mathbf{E}^T$, where

$$\mathbf{E} = \begin{pmatrix} \epsilon_{1x} & \epsilon_{2x} & \epsilon_{3x} \\ \epsilon_{1y} & \epsilon_{2y} & \epsilon_{3y} \\ \epsilon_{1z} & \epsilon_{2z} & \epsilon_{3z} \end{pmatrix}, \text{ and } \boldsymbol{\Lambda} = \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix}.$$

The columns of \mathbf{E} (denoted $\boldsymbol{\epsilon}_1$, $\boldsymbol{\epsilon}_2$ and $\boldsymbol{\epsilon}_3$) are the eigenvectors of \mathbf{D} , and λ_1 , λ_2 and λ_3 the corresponding eigenvalues. (It is customary in DTI processing to order the columns of \mathbf{E} such that $\lambda_1 \geq \lambda_2 \geq \lambda_3$, and I will adopt this convention throughout.) Since \mathbf{D} is p.d.s., it is always diagonalizable, with three positive eigenvalues. (It should be noted, however, that estimated tensors are sometime not positive definite due to noise and numerical errors.) A closed-form solution is available for the eigensystem of a 3×3 p.d.s. matrix [44].

The diagonalization of \mathbf{D} allows for a simple but powerful geometric interpretation of the tensor as an ellipsoid in \mathbb{R}^3 , with the eigenvectors determining the directions of the principal axes and the eigenvalues their radii. The surface of the diffusion ellipsoid corresponds to the isosurface of molecular diffusion after a fixed period of time.

1.2.2.3 Calculation Of Scalar Quantities From Tensors As discussed above, the diffusion tensor at any one voxel encodes information related to the magnitude, direction, and shape of diffusion at that location in the brain. However, the tensor is an awkward object from a data analytic perspective, in that it is multivariate and encodes structured geometric information. Thus, the first step in extracting usable information from a DTI study is often to calculate a small number of scalar-valued tensor functionals at each voxel representing important features of local diffusion. These scalars can then be used to produce images, or as the basis of voxelwise data analyses.

A large number of such scalars have been proposed in the literature, most of which fall into one of three not necessarily mutually exclusive categories: those reflecting tensor magnitude or diffusivity, those reflecting anisotropy, and those reflecting the shape of diffusion. I summarize some widely used examples from each of these categories in this section. In Chapter 4, I present an approach for comparing the reliability of scalar indices, thereby providing novel statistical guidance on which scalars should be considered for different analytic purposes.

The total amount of diffusion in a voxel reflects the density and water content of the underlying tissue[70]. In particular, diffusion is much greater in the fluid-filled ventricles than in either white or gray matter structures[76]. Total diffusion is represented by measures of the magnitude of the tensor, and several such measures have been proposed. The two most widely used measures are scalar multiples of one another: the trace of the tensor[44]:

$$\text{Tr}(\mathbf{D}) = D_{xx} + D_{yy} + D_{zz} \tag{1.13}$$

$$= \lambda_1 + \lambda_2 + \lambda_3, \tag{1.14}$$

and the mean diffusivity[16]:

$$\mathbf{D}_{av} = \frac{\text{Tr}(\mathbf{D})}{3}. \tag{1.15}$$

Another measure of diffusivity is the Frobenius norm of the tensor[31]:

$$\|\mathbf{D}\|_F = \sqrt{\text{Tr}(\mathbf{D}\mathbf{D}^T)} \quad (1.16)$$

$$= \sqrt{D_{xx}^2 + D_{yy}^2 + D_{zz}^2 + 2D_{xy}^2 + 2D_{xz}^2 + 2D_{yz}^2} \quad (1.17)$$

$$= \sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2} \quad (1.18)$$

Much of the focus to date of DTI studies has been in using DTI images to quantify diffusion anisotropy. Diffusion anisotropy is assumed to reflect the degree of organization of the underlying tissue, with compact white matter structures exhibiting greater anisotropy than more diffuse, non-compact white matter[67], gray matter structures showing considerably less anisotropy than white matter structures[70], and little to no anisotropy in cerebrospinal fluid[76]. Importantly, it is also believed that axonal bundles with greater myelination exhibit greater anisotropy than areas with less (or compromised) myelination, although the exact relative contributions of myelin and of axonal membranes to total diffusion anisotropy are unknown[17].

Among the many scalar indices representing diffusion anisotropy, the most widely used is the fractional anisotropy (FA), or the normalized standard deviation of the eigenvalues[16]:

$$\text{FA} = \frac{\sqrt{(\lambda_1 - \mathbf{D}_{\text{av}})^2 + (\lambda_2 - \mathbf{D}_{\text{av}})^2 + (\lambda_3 - \mathbf{D}_{\text{av}})^2}}{\frac{\|\mathbf{D}\|_F}{\sqrt{3}}} \quad (1.19)$$

$$= \sqrt{\frac{3}{2}} \sqrt{\frac{(\lambda_1 - \mathbf{D}_{\text{av}})^2 + (\lambda_2 - \mathbf{D}_{\text{av}})^2 + (\lambda_3 - \mathbf{D}_{\text{av}})^2}{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}}. \quad (1.20)$$

FA ranges from 0 to 1 with higher values reflecting increased anisotropy. In particular, FA = 0 reflects three equal eigenvalues, a spherical tensor, and FA = 1 reflects two eigenvalues equal to 0, a degenerate ellipsoid corresponding to a line segment.

Basser[16] also proposed using the relative anisotropy (RA), or the coefficient of variation of the eigenvalues:

$$\text{RA} = \frac{\sqrt{\text{Var}(\lambda)}}{E(\lambda)} \quad (1.21)$$

$$= \frac{\sqrt{(\lambda_1 - D_{\text{av}})^2 + (\lambda_2 - D_{\text{av}})^2 + (\lambda_3 - D_{\text{av}})^2}}{\sqrt{3}D_{\text{av}}}. \quad (1.22)$$

RA ranges from 0 to $\sqrt{2}$ with higher values reflecting increased anisotropy. Some authors prefer to use a scaled RA:

$$\text{sRA} = \frac{\text{RA}}{\sqrt{2}}, \quad (1.23)$$

which ranges from 0 to 1[45].

Another measure of anisotropy, the volume ratio (VR)[75], is equal to 1 minus the ratio of the volume of the diffusion ellipsoid to the volume of a sphere with radius equal to the mean diffusivity:

$$\text{VR} = 1 - \frac{\text{Det}(\mathbf{D})}{\mathbf{D}_{av}^3} \quad (1.24)$$

$$= 1 - \frac{27\lambda_1\lambda_2\lambda_3}{(\lambda_1 + \lambda_2 + \lambda_3)^3}. \quad (1.25)$$

VR also ranges from 0 to 1. As originally proposed by Pierpaoli[75], the volume ratio was equal to 1 minus VR as defined here, but I follow Kingsley[45] in reversing the direction so that higher values indicate increased anisotropy.

More recently, the geodesic anisotropy (GA) has been proposed[24], based upon the view that the space of diffusion tensors is best seen as a Riemannian manifold equipped with an affine-invariant Riemannian metric:

$$\text{GA} = \sqrt{(\log \lambda_1 - \overline{\log \lambda})^2 + (\log \lambda_2 - \overline{\log \lambda})^2 + (\log \lambda_3 - \overline{\log \lambda})^2}. \quad (1.26)$$

Unlike the other indices, GA ranges from 0 to ∞ , with higher values indicating increased anisotropy.

It is important to note that quantities such as these do not constitute independent measures of diffusion anisotropy. In some cases, anisotropy indices are relatively simple algebraic functions of one another. For instance,

$$\text{FA} = \text{sRA} \left[\frac{3}{2\text{sRA}^2 + 1} \right]^{\frac{1}{2}}.$$

A third class of scalar-valued tensor functionals are those which attempt to describe the shape of the diffusion ellipsoid. Westin[95] proposed three complementary indices, c_l , c_p , c_s , describing the degree to which the shape of the ellipsoid is linear, planar and spherical:

$$c_l = \frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2 + \lambda_3} \quad (1.27)$$

$$c_p = \frac{2(\lambda_2 - \lambda_3)}{\lambda_1 + \lambda_2 + \lambda_3} \quad (1.28)$$

$$c_s = \frac{3\lambda_3}{\lambda_1 + \lambda_2 + \lambda_3}. \quad (1.29)$$

These indices each range from 0 to 1 and they sum to 1.

Ennis (2006) proposed using the quantity $\text{mode}(\tilde{\mathbf{D}})$ to quantify shape:

$$\text{mode}(\tilde{\mathbf{D}}) = 3\sqrt{6} \det \left(\frac{\tilde{\mathbf{D}}}{\|\tilde{\mathbf{D}}\|_F} \right), \quad (1.30)$$

where $\tilde{\mathbf{D}}$ is the anisotropic part, or “deviatoric”, of \mathbf{D} , $\tilde{\mathbf{D}} = \mathbf{D} - \frac{\text{Tr}(\mathbf{D})}{3}\mathbf{I}$. $\text{mode}(\tilde{\mathbf{D}})$ ranges from -1 to 1, with -1 indicating planar shape, 0 indicating spherical shape and 1 indicating linear shape[31].

1.3 DTI DATA ANALYSIS

Although investigators have taken a variety of approaches to analyzing the results of DTI studies, there are two basic analytic strategies that are common enough to be considered canonical: the region of interest-based approach (ROI), and the voxelwise approach. The large majority of group comparisons of DTI data can be seen as a variation on one of these two basic analytic strategies.

For either approach, the first steps are generally a series of data pre-processing techniques. The most fundamental difficulty with DTI data analysis is the low signal-to-noise ratio (SNR) which is common to MRI procedures, due to noise from a variety of sources, both from the subject and the machine. Therefore, noise-correction steps are often employed prior to data analysis. Several authors have reviewed the nature and sources of noise in DTI, and have proposed a variety of approaches to mitigating the problem[9, 13, 11, 39, 43, 45, 94].

Unfortunately, some noise reduction techniques may result in loss of resolution of the image, and the effects of specific noise reduction techniques on data analysis results is generally unknown. However, a certain amount of noise reduction is considered critical to obtaining interpretable images, and data analysts are largely constrained to developing and employing methods conditional on these first pre-processing steps. For example, in the applications I present in Chapter 3, the data have been subjected to motion and eddy current distortion[63] correction prior to analysis.

The next step is often inter-subject registration: spatially transforming each subject's image into a common anatomical space. This step is critical for voxelwise analyses, but may not be necessary for ROI-based analyses. Again, registration induces distortion and reduces resolution in the resulting images, and the downstream effects on data analysis are difficult to quantify. Registration is particularly problematic between populations with possibly substantial differences in neuroanatomy. Although much work on registration from other neuroimaging modalities is applicable to DTI, there is a growing body of literature relating to registration of DTI images specifically[3, 40, 74, 80], including issues related to reorienting diffusion tensors in the target space.

1.3.1 Analyses Based On Regions Of Interest

In common practice, ROI-based analysis involves first drawing regions of interest in the brain manually, semi-automatically or automatically, generally on anatomical images co-registered to the DWIs, or extracting them automatically based on segmentation algorithms and/or spatial normalization to a standard anatomical atlas. Then tensor-based scalars, such as those described in Section 1.2.2.3, are calculated at each voxel in the ROI and averaged over the ROI for each subject. The mean values of these ROI-level average scalar measurements are then compared between groups using simple statistical tests (most frequently two-sample t-tests, but ANOVA, ANCOVA and non-parametric tests, including resampling-based procedures, are sometimes used instead).

For instance, three of five previously published DTI studies in ASDs used variations of this approach. Alexander et al. (2007) manually extracted four ROIs in the corpus callosum

from automatically generated white matter masks. They compared mean FA (Equation 1.19), \mathbf{D}_{av} (Equation 1.15), axial diffusivity (λ_1) and radial diffusivity ($\frac{\lambda_2+\lambda_3}{2}$) over these ROIs between participants with ASDs and controls using two-sample t-tests[2].

Lee et al. (2007) adopted a very similar approach, comparing mean FA, \mathbf{D}_{av} , axial diffusivity and radial diffusivity over four ROIs between ASDs participants and controls, this time using ANCOVA, controlling for age. They also employed a Bonferroni correction to adjust for the multiple ROIs (but not the multiple scalars) examined[58].

Ben Bashat et al. (2007) used a slightly different approach, in that their study only involved participants with ASDs and no controls. They manually extracted 18 ROIs from each subject and calculated average FA over these regions, along with two other indices, Prob and Diff, derived from q -space analysis[6]. They then compared the mean levels of these ROI averages to predicted normal population levels using one-tailed, one-sample t-tests. They used a modified Bonferroni method to adjust for multiple comparisons[19].

There are several strengths to the ROI-based approach. First is its simplicity: by marginalizing over regions, the complex information contained in a full DTI image can be reduced to a small number of scalar values for each subject, which are easily compared between groups. Another strength is interpretability: the ROIs are chosen to be anatomically meaningful, and the results of an ROI-based analysis are direct measures of overall group difference across these anatomically meaningful regions. Finally, ROI-based analyses do not necessarily require subjects to be co-registered, eliminating a pre-processing step that can entail substantial distortion and/or loss of resolution.

However, there are also corresponding weaknesses to the ROI-based approach. The simplicity that was mentioned as a strength is also a weakness: marginalizing over regions ignores within-region variability, and makes it quite possible to miss group differences that may exist only in subregions (or, conversely, to ascribe group differences to an entire ROI that belong properly only to a sub-region). Another weakness is that identification of ROIs generally requires strong prior knowledge of the boundaries of anatomical regions, and which regions are most likely to show differences. Manual segmentation of ROIs involves case-by-case subjective judgment on the part of the researcher. However, even under the assumption that specified ROIs are anatomically accurate and correspond exactly to important areas of

group differences, ignoring the within-subject variability across the ROI can produce inflated estimates of the precision of fixed effect estimates, potentially leading to spurious findings. Finally (and this is a weakness common to most DTI analyses), the selection of one or a few scalar indices to analyze leaves open the possibility that important information contained in the full diffusion tensor is being ignored.

1.3.2 Voxelwise Analyses

An important alternative to ROI-based analysis is voxelwise analysis. In a voxelwise analysis, subjects are spatially co-registered to a common anatomical space, and mean values of one or more DTI-derived scalars are compared between groups separately at each voxel. The statistical method used for the comparison at each voxel is again often, but not invariably, a t-test. Typically, because of the large number of comparisons, a threshold for declaring hypothesis tests significant is determined based on the number of comparisons and the spatial contiguity of voxels showing significant differences. Voxelwise analyses may involve the whole brain, or may themselves be limited to ROIs.

Two previous DTI studies in ASDs used voxelwise approaches. Barnea-Goraly et al. (2004) compared FA at each voxel of an image automatically masked to contain only white matter between ASDs and control participants. Each comparison was performed with a two-sample t-test, and the results were presented as Z statistics, thresholded based on spatial contiguity[8].

Keller et al. (2007) also performed voxelwise comparisons of FA. In their paper, each comparison was a random effects multiple regression (the authors are non-specific about which effects were random, but it is reasonable to assume that these were linear mixed effects models with random intercepts). At each voxel, they used a stepwise variable selection process to determine which covariates to include in the models. The final p-values were conservatively thresholded to adjust for multiple comparisons.

The greatest advantage of voxelwise analysis over ROI-based analysis is that it is possible to detect group differences that occur on a smaller scale than over an entire ROI, or that might overlap multiple ROIs in a way that would be missed by an approach that marginalizes

over ROI. Another advantage is that it is not necessary to impose prior anatomical knowledge on the analyses in terms of identifying the ROIs, although anatomical knowledge is certainly relevant in the interpretation of voxelwise analyses.

The disadvantages of voxelwise analyses are complementary to the advantages of ROI-based analyses. Inter-subject registration is required for voxelwise analyses, with consequent distortion and dependence of analytic results on the choice of registration algorithm. The results of voxelwise analyses may be difficult to interpret anatomically. Voxelwise analyses also present a very real multiple comparisons problem, and the choice of a thresholding approach is generally *ad hoc*. Finally, voxelwise analyses can be computationally burdensome, depending on the number of subjects, the number of voxels compared, and the complexity of the method used to compare groups at each voxel.

1.4 PROPOSED STATISTICAL APPROACHES FOR DTI

In this dissertation, I explore the application of linear mixed effects models with closed-form likelihood estimation to DTI data analyses. Chapters 2 and 3 are devoted to using LMEs to perform group comparisons of DTI data, while Chapter 4 is concerned with an LME-based approach to quantifying reliability. While all the examples considered involve DTI data, these methods are applicable to neuroimaging data more generally.

In Chapter 2, I describe a class of linear mixed-effects models for spatio-temporal data with closed-form maximum likelihood estimators. In Section 2.3, I discuss the circumstances under which closed-form estimation is available in cross-sectional models that assume homogeneous variance between groups. In Section 2.4, I relax the homogeneity assumption, and examine closed-form estimation in models which allow heterogeneous variance of multiple forms between groups. And in Section 2.5, I extend these cross-sectional models to include longitudinal effects.

Chapter 3 is devoted to applying the results of Chapter 2 to the analysis of a DTI study comparing children and young adults with ASDs to controls. In this chapter, I compare LMEs with closed-form estimation to more traditional ROI-based and voxelwise group comparisons.

In addition to analyses of FA, I also describe multivariate analyses of the whole tensor data.

Finally, in Chapter 4, I discuss generalizations to spatio-temporal models of the intraclass correlation coefficient as an index of reliability. I present two theorems identifying invariant properties of the generalized intraclass correlation coefficient (gICC) to model complexity within the class of growth-curve models, and conduct a series of simulations examining the performance of the gICC under the assumption of autoregressive errors in DTI data. I then provide an application of the gICC to the quantification of local spatial reliability in a sample of children and young adults with ASDs and controls.

2.0 CLOSED-FORM SPATIO-TEMPORAL NEUROIMAGING MODELS

Linear mixed-effects (LME) models have great potential for performing group-based inference on neuroimaging data, and have been productively employed in the analysis of fMRI data[52]. My work in this chapter is concerned with developing LMEs for the analysis of DTI data. In addition to making it possible to explore group differences in mean values of tensor elements, LMEs allow researchers to separate between-subject from within-subject variance components and thus borrow strength across multiple subjects in obtaining subject-specific estimated means. LME models also provide a convenient framework for exploring longitudinal data, for accommodating between-group heterogeneity in variance, and for incorporating spatial autocorrelation among neighboring voxels.

However, maximum likelihood estimation of LME models typically requires iteration. These can be costly in terms of computer time, especially for models with many parameters and large datasets, such as DTI scans. For instance, each whole brain DTI image in the dataset explored in Chapter 3 contains approximately 200,000 voxels. While confining attention to specific tissue types or to regions of interest leads to smaller datasets, the computation can still be prohibitive in exploratory, model-building analyses. Thus, it is important to identify subclasses of LME models appropriate to DTI datasets that will admit to closed-form ML estimators, and to derive the functional form of these estimators.

In this chapter, I first justify the use of univariate LME models for DTI tensor comparisons by deriving a data transformation that will remove the statistical dependence among the six tensor elements (Section 2.1). Next, I review the mathematical tools available to determine the circumstances under which a closed-form ML estimator will exist (Section 2.2). Then I will apply these tools to a series of models, first exploring the existence of closed-form ML estimators for successively more complicated mean structures for cross-sectional

data under simple variance assumptions (Section 2.3), and then exploring more complicated variance structures, including between-group variance heterogeneity (Section 2.4). Finally, I will explore the possibility of closed-form ML estimators in the case of longitudinal data (Section 2.5). Each of these developments will be illustrated by examples with a DTI study in the following chapter.

2.1 TENSOR ELEMENT DECORRELATION

Diffusion tensor imaging data analyses are considerably complicated by the multidimensional nature of the diffusion tensor estimated at each voxel. This multidimensionality in dependent variables is an added feature of complexity to DTI data relative to other neuroimaging modalities, on top of the ordinary multidimensionality due to the repeated spatial measurements represented by multiple voxels. Practitioners often try to work around this complexity by reducing the information in each tensor to a single scalar value, but this approach has disadvantages that will be outlined below. In this section, I will describe a simple linear transformation of tensor elements which serves to decorrelate the tensor elements under an assumption of rotational invariance. This serves as a justification for analyzing the transformed tensor elements independently of one another. This work is based on theoretical results for a tensor-valued multivariate normal distribution presented by Basser and Pajevic (2003)[14], but to the best of my knowledge this is a novel application of these results to data analytic considerations.

2.1.1 Disadvantages Of Scalar Reduction

As discussed in Section 1.2.2.3, one immediate solution to the problem of multidimensionality is to reduce the information contained in each diffusion tensor \mathbf{D} to a scalar $f(\mathbf{D}) \in \mathbb{R}$ and to then use this scalar as the dependent variable in analyses. The fractional anisotropy (FA, Equation 1.19) is the most popular choice for such scalar reduction. In addition to dimensional reduction, this approach has the advantage that such scalars are often directly

interpretable in more useful ways than the original tensor elements (for instance, as representing anisotropy, diffusivity or shape).

However, the disadvantages of this approach are also clear: the sheer number of candidate scalar indices as described in Section 1.2.2.3 points to the problem of identifying *which* scalar index best encapsulates the relevant information for a given problem. In many cases, it may also be that no single scalar reduction of the tensor adequately summarizes the diffusion properties across an entire brain or region for the purposes of group comparison.

We would like to be able to retain the full depth of information contained in the diffusion tensor without multiplying the dimensionality of our already spatially (and, in the case of longitudinal DTI data, temporally) extensive datasets. A natural approach would be to treat the tensor elements as independent observations and to essentially perform the same analysis six times - once for each distinct tensor element. However, the tensor elements are correlated with one another over repeated measurements, and this approach would ignore such dependence, leading to an overestimation of effective degrees of freedom and invalid inference.

2.1.2 Tensor-Valued Normal Distribution

Basser and Pajevic (2003) explored the algebraic structure of a tensor-valued normal distribution for DTI data[14]. By analogy with the density of the multivariate normal distribution on \mathbb{R}^n ,

$$p(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}^{-1}) \propto \exp\left(\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})' \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})\right),$$

they defined a tensor-valued normal distribution on the space of diffusion tensors as:

$$p(\mathbf{D}; \mathbf{D}^0, \mathbf{A}) \propto \exp\left(\frac{1}{2}(\mathbf{D} - \mathbf{D}^0) : \mathbf{A} : (\mathbf{D} - \mathbf{D}^0)\right), \quad (2.1)$$

where \mathbf{D}^0 is the mean diffusion tensor, \mathbf{A} is a fourth-order “precision” tensor analogous to the precision matrix $\boldsymbol{\Sigma}^{-1}$, and $(\cdot : \cdot)$ represents tensor contraction:

$$\mathbf{D} : \mathbf{A} : \mathbf{D} = \sum_{i,j,m,n} D_{ij} A_{ijmn} D_{mn}.$$

Although \mathbf{A} consists of $3^4 = 81$ elements, symmetry conditions (for instance, that $A_{ijmn} = A_{mnij}$) imply that there are at most 21 independent elements of \mathbf{A} . Basser and Pajevic demonstrated that this tensor-valued distribution was equivalent to a multivariate normal distribution on the six-dimensional vector of tensor elements, $\tilde{\mathbf{D}} = [D_{xx}D_{yy}D_{zz}D_{xy}D_{xz}D_{yz}]'$, with precision matrix Σ^{-1} made up of elements of \mathbf{A} :

$$\Sigma^{-1} = \begin{bmatrix} A_{xxxx} & A_{xxyy} & A_{xxzz} & 2A_{xxxy} & 2A_{xxxz} & 2A_{xxyz} \\ A_{xxyy} & A_{yyyy} & A_{yyzz} & 2A_{yyxy} & 2A_{yyxz} & 2A_{yyyz} \\ A_{xxzz} & A_{yyzz} & A_{zzzz} & 2A_{zzxy} & 2A_{zzxz} & 2A_{zzyz} \\ 2A_{xxxy} & 2A_{yyxy} & 2A_{zzxy} & 4A_{xyxy} & 4A_{xyxz} & 4A_{xyyz} \\ 2A_{xxxz} & 2A_{yyxz} & 2A_{zzxz} & 4A_{xyxz} & 4A_{xzzz} & 4A_{xzyz} \\ 2A_{xxyz} & 2A_{yyyz} & 2A_{zzyz} & 4A_{xyyz} & 4A_{xyyz} & 4A_{yzyz} \end{bmatrix}. \quad (2.2)$$

2.1.3 Tensor Element Correlation

Basser and Pajevic then considered the important special case in which \mathbf{A} is a general isotropic fourth-order tensor. This corresponds to the assumption that \mathbf{A} is rotationally invariant in the sense that its form is unchanged under any rotation, reflection or inversion of the coordinate system in which \mathbf{D} is measured. Note that this is a non-trivial assumption — indeed, Jones et al. (in press) showed that rotational invariance is dependent upon the gradient sampling scheme and the true but unknown value of the diffusion tensor[41]. Thus, there cannot be a design that is completely rotationally invariant across the whole brain. However, certain gradient encoding schemes will minimize rotational variance under general conditions, and popular choices of sampling schemes (dodecahedral and icosahedral) may satisfy this condition.

Under the assumption of rotational invariance, Σ^{-1} takes the particularly simple form:

$$\Sigma^{-1} = \begin{bmatrix} \zeta + 2\theta & \zeta & \zeta & 0 & 0 & 0 \\ \zeta & \zeta + 2\theta & \zeta & 0 & 0 & 0 \\ \zeta & \zeta & \zeta + 2\theta & 0 & 0 & 0 \\ 0 & 0 & 0 & 4\theta & 0 & 0 \\ 0 & 0 & 0 & 0 & 4\theta & 0 \\ 0 & 0 & 0 & 0 & 0 & 4\theta \end{bmatrix}. \quad (2.3)$$

In this form, Σ^{-1} can be diagonalized as $\Sigma^{-1} = \mathbf{E}\mathbf{\Lambda}\mathbf{E}'$, where

$$\mathbf{E} = \begin{bmatrix} \frac{1}{\sqrt{3}} & \frac{1}{\sqrt{2}} & \frac{1}{2\sqrt{2}} & 0 & 0 & 0 \\ \frac{1}{\sqrt{3}} & \frac{-1}{\sqrt{2}} & \frac{1}{2\sqrt{2}} & 0 & 0 & 0 \\ \frac{1}{\sqrt{3}} & 0 & \frac{-1}{\sqrt{2}} & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \text{ and } \mathbf{\Lambda} = \begin{bmatrix} 3\zeta + 2\theta & 0 & 0 & 0 & 0 & 0 \\ 0 & 2\theta & 0 & 0 & 0 & 0 \\ 0 & 0 & 2\theta & 0 & 0 & 0 \\ 0 & 0 & 0 & 4\theta & 0 & 0 \\ 0 & 0 & 0 & 0 & 4\theta & 0 \\ 0 & 0 & 0 & 0 & 0 & 4\theta \end{bmatrix}$$

. The columns of \mathbf{E} are the eigenvectors of Σ^{-1} (and hence of Σ), and the diagonal entries of $\mathbf{\Lambda}$ are the eigenvalues of Σ^{-1} (and hence their reciprocals are the eigenvalues of Σ). As a consequence of this diagonalization and the orthogonality of \mathbf{E} , if $\tilde{\mathbf{D}} \sim \text{MVN}(\tilde{\mathbf{D}}^0, \Sigma)$, then $\mathbf{E}'\tilde{\mathbf{D}} \sim \text{MVN}(\mathbf{E}'\tilde{\mathbf{D}}^0, \mathbf{\Lambda}^{-1})$.

Thus, \mathbf{E}' is a decorrelating transformation for the vector of tensor elements under the assumptions of normality of the diffusion tensor and rotational invariance of its precision tensor. We may therefore treat the transformed elements independently in subsequent analyses, and transform results back into the original space as necessary for interpretability.

However, it is also worth noting that the transformed tensor elements themselves have somewhat interesting interpretations in terms of diffusion parameters. The first transformed element,

$$\mathbf{E}'_1\tilde{\mathbf{D}} = \frac{1}{\sqrt{3}}(D_{xx} + D_{yy} + D_{zz}),$$

is proportional to the mean diffusivity, \mathbf{D}_{av} , and thus represents the total amount of diffusion in a given voxel. The second transformed element,

$$\mathbf{E}'_2 \tilde{\mathbf{D}} = \frac{1}{\sqrt{2}}(D_{xx} - D_{yy}),$$

represents the difference in diffusivity between the x and y directions, and thus is a simple anisotropy measure. And the third transformed element,

$$\mathbf{E}'_3 \tilde{\mathbf{D}} = \frac{1}{\sqrt{2}} \left(\frac{D_{xx} + D_{yy}}{2} - D_{zz} \right),$$

represents the difference in diffusivity between the mean of the x and y directions and the z direction, thus capturing anisotropy unaccounted for by the second transformed element. (Note that this interpretation diverges somewhat from that given by Basser and Pajevic, who exhibited a different diagonalization of Σ^{-1} .)

The remaining three elements of the transformed vector of tensor elements are the unchanged off-diagonal elements of the original diffusion tensor, \mathbf{D}_{xy} , \mathbf{D}_{xz} , and \mathbf{D}_{yz} .

2.2 CONDITIONS FOR THE EXISTENCE OF CLOSED-FORM ML ESTIMATORS

The general linear mixed-effects model can be written as:

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\boldsymbol{\gamma} + \boldsymbol{\epsilon}, \tag{2.4}$$

where \mathbf{Y} is an $N \times 1$ vector of observed responses, \mathbf{X} is an $N \times p$ fixed effects design matrix, $\boldsymbol{\beta}$ is a fixed but unknown $p \times 1$ vector of fixed effects parameters, \mathbf{Z} is an $N \times q$ random effects design matrix with $q \leq p$, $\boldsymbol{\gamma}$ is a vector of random effects parameters, with $\boldsymbol{\gamma} \sim \text{MVN}(\mathbf{0}, \mathbf{G})$, and $\boldsymbol{\epsilon}$ is a vector of residual errors, with $\boldsymbol{\epsilon} \sim \text{MVN}(\mathbf{0}, \mathbf{R})$ and $\boldsymbol{\gamma} \perp \boldsymbol{\epsilon}$. Under this model, $\mathbf{Y} \sim \text{MVN}(\mathbf{X}\boldsymbol{\beta}, \boldsymbol{\Sigma})$, where $\boldsymbol{\Sigma} = \mathbf{Z}\mathbf{G}\mathbf{Z}' + \mathbf{R}$.

The unknown parameters in the LME model are usually estimated via likelihood methods, which have well-known favorable asymptotic properties. However, likelihood estimation of LMEs typically requires iterative numerical procedures, such as Newton's method or the

EM algorithm[50]. While it is known that closed-form, non-iterative likelihood estimation is possible for certain important subclasses of the general LME model, such as balanced growth curves[54], such models may fail to capture the complexity of neuroimaging data. However, closed-form estimation would be especially desirable in models for neuroimaging data, due to the computational burden involved. In this section, I review theoretical tools for identifying conditions under which closed-form likelihood estimation is possible in LMEs. I then use these tools in subsequent sections to describe, in practical terms, classes of useful models for neuroimaging data for which closed-form estimation is possible.

2.2.1 Normal Estimation Under Additive Mean And Covariance

Anderson (1973) examined the case of estimating the parameters of a multivariate normal distribution, $\mathbf{Y} \sim \text{MVN}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ in the special case of a linear mean and covariance structure:

$$\begin{aligned}\boldsymbol{\mu} &= \mathbf{X}\boldsymbol{\beta}, \text{ and} \\ \boldsymbol{\Sigma} &= \sum_{g=0}^m \sigma_g \mathbf{G}_g,\end{aligned}$$

where $\boldsymbol{\beta}$ and $\boldsymbol{\sigma} = [\sigma_0 \dots \sigma_m]'$ are unknown, \mathbf{X} is known and each \mathbf{G}_g is a known, symmetric matrix[4]. He noted that, for M i.i.d. observations $\mathbf{Y}_1, \dots, \mathbf{Y}_M$, the likelihood is maximized over unknown parameters $\boldsymbol{\beta}$ and $\boldsymbol{\sigma}$ at the joint solution of the series of equations:

$$\mathbf{X}'\hat{\boldsymbol{\Sigma}}^{-1}\mathbf{X}\hat{\boldsymbol{\beta}} = \mathbf{X}'\hat{\boldsymbol{\Sigma}}^{-1}\bar{\mathbf{Y}} \quad (2.5)$$

$$\text{Tr} \left(\left(\left(\sum_{h=0}^m \hat{\sigma}_h \mathbf{G}_h \right)^{-1} \mathbf{G}_g \right) \right) = \text{Tr} \left(\left(\left(\sum_{h=0}^m \hat{\sigma}_h \mathbf{G}_h \right)^{-1} \mathbf{G}_g \left(\sum_{h=0}^m \hat{\sigma}_h \mathbf{G}_h \right)^{-1} \mathbf{C} \right) \right), \quad (2.6)$$

for $g = 0, \dots, m$, where

$$\mathbf{C} = \frac{1}{M} \sum_{\alpha=1}^M (\mathbf{Y}_\alpha - \hat{\boldsymbol{\mu}})(\mathbf{Y}_\alpha - \hat{\boldsymbol{\mu}})', \quad (2.7)$$

for $\hat{\boldsymbol{\mu}} = \mathbf{X}\hat{\boldsymbol{\beta}}$. Anderson observed that this arrangement suggests an iterative estimation procedure, in which an initial estimate is adopted for $\hat{\boldsymbol{\sigma}}$, which is substituted into Equation 2.5 to obtain an estimate of $\hat{\boldsymbol{\beta}}$, which is then substituted into Equation 2.6 to obtain a revised

estimate of $\hat{\sigma}$, and so on until convergence to the maximum likelihood estimates. He went on to derive asymptotic properties of this estimation procedure.

Szatrowski (1980) addressed the question of the general circumstances under which explicit, closed-form solutions exist for the maximum likelihood estimates in the models with additive mean and covariance structure introduced above[90]. That is, he provided necessary and sufficient conditions under which the iterative procedure discussed by Anderson converges in one iteration. He separately considered closed-form estimation for the mean and for the variance parameters.

2.2.2 Closed-Form Estimation Of Mean Parameters

Szatrowski observed that a closed-form expression for $\hat{\beta}$ exists if and only if Equation 2.5 doesn't depend on $\hat{\Sigma}$. In other words, since

$$\hat{\beta} = (\mathbf{X}'\hat{\Sigma}^{-1}\mathbf{X})^{-1}\mathbf{X}'\hat{\Sigma}^{-1}\bar{\mathbf{Y}},$$

a closed-form expression exists for $\hat{\beta}$ iff

$$(\mathbf{X}'\hat{\Sigma}^{-1}\mathbf{X})^{-1}\mathbf{X}'\hat{\Sigma}^{-1} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}', \quad (2.8)$$

in which case $\hat{\beta} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\bar{\mathbf{Y}}$, the ordinary least squares (OLS) solution. He then derived an algebraic condition equivalent to the independence of Equation 2.5 to $\hat{\Sigma}$, summarized here as a theorem:

Theorem 2.2.1 (Adapted from Szatrowski's Theorem 2[90]). *For a multivariate normal distribution with additive mean and covariance structure as described in Section 2.2.1, a closed-form solution exists for the maximum likelihood estimate $\hat{\beta}$ if and only if the p columns of \mathbf{X} are spanned by p eigenvectors of Σ . In this case, the closed-form ML estimator is the OLS estimator, $\hat{\beta} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\bar{\mathbf{Y}}$*

Although this condition is potentially difficult to verify in general, in the case of certain important covariance structures, it is possible to identify the spectral decomposition of the covariance matrix, and derive specific, practical conditions on the design matrix \mathbf{X} under which closed-form estimation of $\hat{\beta}$ is possible. I employ this line of argument several times in subsequent sections.

2.2.3 Closed-Form Estimation Of Covariance Parameters

Szatrowski also developed an analogous algebraic condition to that in Theorem 2.2.1 for the existence of closed-form ML estimates for covariance parameters. However, this condition is algebraically cumbersome and difficult to apply. Fortunately, in the important special case that the \mathbf{G} matrices are simultaneously diagonalizable, there is a considerably simplified condition under which closed-form estimates for the covariance parameters exist.

The matrices $\mathbf{G}_0, \dots, \mathbf{G}_m$ are said to be *simultaneously diagonalizable* if there exists an orthonormal matrix \mathbf{P} such that $\mathbf{P}\mathbf{G}_\alpha\mathbf{P}'$ is a diagonal matrix for each $\alpha = 0, \dots, m$. Symmetric matrices $\mathbf{G}_0, \dots, \mathbf{G}_m$ are simultaneously diagonalizable if and only if they form a commuting family of matrices (that is, if $\mathbf{G}_i\mathbf{G}_j = \mathbf{G}_j\mathbf{G}_i \forall i, j$)[37]. As will be seen, in all of the important cases I'll consider below, the \mathbf{G} matrices form commuting families and therefore are simultaneously diagonalizable.

When the \mathbf{G} matrices are simultaneously diagonalizable, Szatrowski gives the following condition for the existence of closed-form ML estimates for covariance parameters:

Theorem 2.2.2 (Adapted from Szatrowski's Theorem 5[90]). *Consider the multivariate normal distribution with additive mean and covariance structure as described in Section 2.2.1. Assume there is a closed-form solution for the ML estimate of $\boldsymbol{\beta}$ and that $\mathbf{P}\mathbf{G}_\alpha\mathbf{P}'$ is diagonal for some orthonormal \mathbf{P} for $\alpha = 0, \dots, m$. Then a closed-form solution exists for the ML estimate of $\boldsymbol{\sigma}$ if and only if the eigenvalues of $\boldsymbol{\Sigma}$ (i.e., the diagonal entries of $\mathbf{P}\boldsymbol{\Sigma}\mathbf{P}'$) consist of exactly $m + 1$ linearly independent combinations of $\sigma_0, \dots, \sigma_m$.*

When a closed-form expression for $\hat{\boldsymbol{\sigma}}$ does exist, it is given by

$$\hat{\boldsymbol{\sigma}} = [\text{Tr}(\mathbf{G}_g\mathbf{G}_h)]^{-1} (\text{Tr}(\mathbf{G}_g\mathbf{C})),$$

where $[\text{Tr}(\mathbf{G}_g\mathbf{G}_h)]$ is the $(m + 1) \times (m + 1)$ matrix whose $(i, j)^{\text{th}}$ entry is $\text{Tr}(\mathbf{G}_i\mathbf{G}_j)$ and $(\text{Tr}(\mathbf{G}_g\mathbf{C}))$ is the $(m + 1) \times 1$ column vector whose i^{th} entry is $\text{Tr}(\mathbf{G}_i\mathbf{C})$ for \mathbf{C} given by Equation 2.7.

These two theorems provide very powerful tools for identifying practical conditions under which closed-form estimates exist in important subclasses of LMEs, including models that are

potentially very useful for neuroimaging data. I make repeated reference to both theorems in the remainder of this chapter.

2.3 CLOSED-FORM ML ESTIMATORS UNDER COMPOUND SYMMETRY

In this section, I apply the tools introduced in Section 2.2 to the case of LME models with additive mean structures and simple covariance structures (compound symmetry). As discussed in Section 2.2, it is already known that closed-form expressions exist for mean and covariance parameters in many important models in this class. However, examining these familiar, simple cases in detail will provide motivating examples for applying the same tools to novel models in later sections. I first consider the simplest possible case, where a single scalar mean is fit across all voxels, and then the case with a separate mean fit at each voxel, both for one-group models and models with group-specific mean structures. Finally, I present a general result concerning the existence of closed-form solutions in this class of model.

2.3.1 Notation

Throughout, I fix one transformed tensor element to be the dependent variable, as treating the six transformed tensor elements as independent was justified in Section 2.1. I will assume that there are N subjects, n_1 of whom are from Group 1 (a control group, for instance), and n_2 from Group 2 (an ASDs group).

\mathbf{Y}_i will denote the column vector of V observations of the tensor element in question for subject i , $i = 1, \dots, N$, where V is the number of voxels in each coregistered DTI scan. \mathbf{Y} will denote the column vector of NV observations from all subjects stacked one after

another. That is,

$$\mathbf{Y} = \begin{bmatrix} \mathbf{Y}_1 \\ \mathbf{Y}_2 \\ \vdots \\ \mathbf{Y}_N \end{bmatrix}$$

2.3.2 Overall Scalar Mean

The simplest non-vacuous LME model for DTI data involves fitting a single scalar mean term across all voxels for all subjects, with a random intercept term allowing for subject-specific offsets from that mean. This model would be inadequate for whole-brain voxelwise analyses of DTI data, since the mean values of the tensor elements certainly depend on spatial location. However, the model might be appropriate for small-region ROI analyses, and is certainly an improvement over the simple averaging over regions that's often seen in practice (e.g., see Section 1.3.1).

More importantly for the present purposes, the overall scalar mean model serves as an instructive example of the process of proving the existence and deriving the form of the explicit ML estimators of the mean and variance parameters. Many of the derivations in this section will be generalizable to more complex models in subsequent sections.

The overall scalar mean model is given by

$$\mathbf{Y}_{NV \times 1} = \mathbf{1}_{NV \times 1} \mu + (\mathbf{I}_{N \times N} \otimes \mathbf{1}_{V \times 1}) \boldsymbol{\delta}_{N \times 1} + \boldsymbol{\epsilon}_{NV \times 1}, \quad (2.9)$$

where $\mathbf{1}$ is a column vector of ones, μ is the unknown, fixed mean parameter, $\boldsymbol{\delta} = [\delta_1, \dots, \delta_N]'$ is the vector of subject-specific random intercepts, with $\boldsymbol{\delta} \sim \text{MVN}(\mathbf{0}, \sigma^2 \mathbf{I})$, and $\boldsymbol{\epsilon} = [\epsilon_{11}, \epsilon_{12}, \dots, \epsilon_{1V}, \epsilon_{21}, \dots, \epsilon_{NV}]'$ is the vector of residual errors, with $\boldsymbol{\epsilon} \sim \text{MVN}(\mathbf{0}, \tau^2 \mathbf{I})$. The variance of \mathbf{Y} is given by

$$\begin{aligned} \boldsymbol{\Sigma} &= (\mathbf{I}_{N \times N} \otimes \mathbf{1}_{V \times 1}) \sigma^2 \mathbf{I}_{N \times N} (\mathbf{I}_{N \times N} \otimes \mathbf{1}_{V \times 1})' + \tau^2 \mathbf{I}_{NV \times NV} \\ &= \sigma^2 (\mathbf{I}_{N \times N} \otimes \mathbf{J}_{V \times V}) + \tau^2 \mathbf{I}_{NV \times NV}. \end{aligned} \quad (2.10)$$

In Szatrowski's framework for identifying the existence of explicit ML solutions, $\boldsymbol{\Sigma} = \sigma^2 \mathbf{G}_0 + \tau^2 \mathbf{G}_1$, where $\mathbf{G}_0 = \mathbf{I}_{N \times N} \otimes \mathbf{J}_{V \times V}$ (where $\mathbf{J}_{R \times C}$ is the matrix with R rows and C columns, all of whose entries are 1) and $\mathbf{G}_1 = \mathbf{I}_{NV \times NV}$.

To take a small numerical example, for two subjects with three voxels each ($N = 2$ and $V = 3$),

$$\mathbf{G}_0 = \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 \end{bmatrix}, \quad \mathbf{G}_1 = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \quad \text{and}$$

$$\mathbf{\Sigma} = \begin{bmatrix} \sigma^2 + \tau^2 & \sigma^2 & \sigma^2 & 0 & 0 & 0 \\ \sigma^2 & \sigma^2 + \tau^2 & \sigma^2 & 0 & 0 & 0 \\ \sigma^2 & \sigma^2 & \sigma^2 + \tau^2 & 0 & 0 & 0 \\ 0 & 0 & 0 & \sigma^2 + \tau^2 & \sigma^2 & \sigma^2 \\ 0 & 0 & 0 & \sigma^2 & \sigma^2 + \tau^2 & \sigma^2 \\ 0 & 0 & 0 & \sigma^2 & \sigma^2 & \sigma^2 + \tau^2 \end{bmatrix}.$$

Szatrowski gives conditions for the existence of a closed-form ML estimator for μ that depend on the span of the eigenvectors of $\mathbf{\Sigma}$. Since $\mathbf{\Sigma}$ is block diagonal, its eigensystem can be derived by combining the eigensystems from each block. A single compound symmetric block can be diagonalized as

$$\sigma^2 \mathbf{J} + \tau^2 \mathbf{I} = \mathbf{P} \mathbf{D} \mathbf{P}^{-1},$$

where

$$\mathbf{P} = \begin{bmatrix} 1 & \mathbf{1}_{1 \times (V-1)} \\ \mathbf{1}_{(V-1) \times 1} & -\mathbf{I}_{(V-1) \times (V-1)} \end{bmatrix}, \quad \text{and} \quad \mathbf{D} = \text{diag}(V\sigma^2 + \tau^2, \underbrace{\tau^2, \dots, \tau^2}_{(V-1) \text{ times}})$$

Thus $\mathbf{\Sigma}$ itself has two distinct eigenvalues: $V\sigma^2 + \tau^2$ (with multiplicity N) and τ^2 (with multiplicity $N(V-1)$). It can be diagonalized as

$$\mathbf{\Sigma} = (\mathbf{I}_{N \times N} \otimes \mathbf{P})(\mathbf{I}_{N \times N} \otimes \mathbf{D})(\mathbf{I}_{N \times N} \otimes \mathbf{P})^{-1}. \quad (2.11)$$

In the example with $N = 2, V = 3$,

$$\mathbf{I}_{2 \times 2} \otimes \mathbf{P} = \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & -1 & 0 & 0 & 0 & 0 \\ 1 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & -1 & 0 \\ 0 & 0 & 0 & 1 & 0 & -1 \end{bmatrix},$$

$$(\mathbf{I}_{2 \times 2} \otimes \mathbf{P})^{-1} = \frac{1}{3} \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & -2 & 1 & 0 & 0 & 0 \\ 1 & 1 & -2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & -2 & 1 \\ 0 & 0 & 0 & 1 & 1 & -2 \end{bmatrix}, \text{ and}$$

$$\mathbf{I}_{2 \times 2} \otimes \mathbf{D} = \begin{bmatrix} 3\sigma^2 + \tau^2 & 0 & 0 & 0 & 0 & 0 \\ 0 & \tau^2 & 0 & 0 & 0 & 0 \\ 0 & 0 & \tau^2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 3\sigma^2 + \tau^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & \tau^2 & 0 \\ 0 & 0 & 0 & 0 & 0 & \tau^2 \end{bmatrix}$$

The columns of $(\mathbf{I}_{2 \times 2} \otimes \mathbf{P})$, $[\mathbf{p}_{11} \ \mathbf{p}_{12} \ \dots \ \mathbf{p}_{1V} \ \mathbf{p}_{21} \ \dots \ \mathbf{p}_{2V}]$, are eigenvectors of Σ , and the diagonal entries of $\mathbf{I}_{2 \times 2} \otimes \mathbf{D}$ the corresponding eigenvalues. It is also immediate that the sum of eigenvectors with the same eigenvalue is itself an eigenvector, as $\Sigma \mathbf{x}_1 = \lambda \mathbf{x}_1$ and $\Sigma \mathbf{x}_2 = \lambda \mathbf{x}_2$ yields:

$$\begin{aligned} \Sigma(\mathbf{x}_1 + \mathbf{x}_2) &= \Sigma \mathbf{x}_1 + \Sigma \mathbf{x}_2 \\ &= \lambda \mathbf{x}_1 + \lambda \mathbf{x}_2 \\ &= \lambda(\mathbf{x}_1 + \mathbf{x}_2). \end{aligned}$$

Therefore, the vector $\mathbf{p}_{\cdot 1}$, defined as

$$\mathbf{p}_{\cdot 1} = \sum_{i=1}^N \mathbf{p}_{i1} = \mathbf{1}_{NV \times 1}, \quad (2.12)$$

is an eigenvector of Σ with eigenvalue $V\sigma^2 + \tau^2$. Since this one eigenvector is exactly equal to (and thus trivially spans) the one column of the fixed effects design matrix in model 2.9, we have satisfied the conditions given in Theorem 2.2.1, and thus a closed-form ML estimator for μ exists, and is equal to the OLS estimator,

$$\begin{aligned} \hat{\mu} &= (\mathbf{1}'\mathbf{1})^{-1} \mathbf{1}'\mathbf{Y} \\ &= \frac{1}{NV} \sum_{i,j} Y_{ij}, \end{aligned}$$

the simple average of the elements of \mathbf{Y} .

In order to determine whether a closed-form ML estimator exists for $\boldsymbol{\sigma} = (\sigma^2, \tau^2)$, we first note that the G matrices commute and so are simultaneously diagonalizable. The diagonalized form of Σ is the matrix $(\mathbf{I}_{2 \times 2} \otimes \mathbf{D})$ described above, which has two distinct entries: $V\sigma^2 + \tau^2$ and τ^2 . Therefore, by Theorem 2.2.2, a closed-form expression exists for $\hat{\boldsymbol{\sigma}}$.

The existence of a closed-form solution for $\hat{\boldsymbol{\sigma}}$, the vector of ML estimates of the variance parameters, is equivalent to the convergence in one iteration from any allowable starting point of the likelihood equation,

$$\left[\text{Tr}(\hat{\Sigma}^{-1} \mathbf{G}_g \hat{\Sigma}^{-1} \mathbf{G}_h) \right] \hat{\boldsymbol{\sigma}} = \left(\text{Tr}(\hat{\Sigma}^{-1} \mathbf{G}_g \hat{\Sigma}^{-1} \mathbf{C}) \right),$$

where:

1. $\left[\text{Tr}(\hat{\Sigma}^{-1} \mathbf{G}_g \hat{\Sigma}^{-1} \mathbf{G}_h) \right]$ denotes the $(m+1) \times (m+1)$ matrix with i, j^{th} entry equal to $\text{Tr}(\hat{\Sigma}^{-1} \mathbf{G}_i \hat{\Sigma}^{-1} \mathbf{G}_j)$,
2. $\left(\text{Tr}(\hat{\Sigma}^{-1} \mathbf{G}_g \hat{\Sigma}^{-1} \mathbf{C}) \right)$ denotes the column vector of length $m+1$ with i^{th} entry $\text{Tr}(\hat{\Sigma}^{-1} \mathbf{G}_i \hat{\Sigma}^{-1} \mathbf{C})$, and
3. $\mathbf{C} = (\mathbf{Y} - \mathbf{1}\hat{\mu})'(\mathbf{Y} - \mathbf{1}\hat{\mu})$ denotes the sample covariance matrix.

Since $\hat{\Sigma} = \mathbf{I}$ is an allowable starting point for the covariance structure under consideration, this is equivalent to the following solution for $\hat{\sigma}$:

$$\hat{\sigma} = [\text{Tr}(\mathbf{G}_g \mathbf{G}_h)]^{-1} (\text{Tr}(\mathbf{G}_g \mathbf{C})). \quad (2.13)$$

For the model under consideration in this section,

$$\begin{aligned} \text{Tr}(\mathbf{G}_1 \mathbf{G}_1) &= \text{Tr}(\mathbf{G}_0 \mathbf{G}_1) = \text{Tr}(\mathbf{G}_1 \mathbf{G}_0) = NV, \\ \text{Tr}(\mathbf{G}_0 \mathbf{G}_0) &= NV^2, \\ \text{Tr}(\mathbf{G}_0 \mathbf{C}) &= \sum_{i=1}^N \sum_{j=1}^V \sum_{k=1}^V (Y_{ij} - \hat{\mu})(Y_{ik} - \hat{\mu}), \text{ and} \\ \text{Tr}(\mathbf{G}_1 \mathbf{C}) &= \sum_{i=1}^N \sum_{j=1}^V (Y_{ij} - \hat{\mu})^2. \end{aligned}$$

Equation 2.13 yields:

$$\begin{aligned} \begin{bmatrix} \hat{\sigma}^2 \\ \hat{\tau}^2 \end{bmatrix} &= \begin{bmatrix} NV^2 & NV \\ NV & NV \end{bmatrix}^{-1} \begin{bmatrix} \text{Tr}(\mathbf{G}_0 \mathbf{C}) \\ \text{Tr}(\mathbf{C}) \end{bmatrix} \\ &= \frac{1}{(V-1)NV} \begin{bmatrix} 1 & -1 \\ -1 & V \end{bmatrix} \begin{bmatrix} \text{Tr}(\mathbf{G}_0 \mathbf{C}) \\ \text{Tr}(\mathbf{C}) \end{bmatrix} \\ &= \frac{1}{(V-1)NV} \begin{bmatrix} \text{Tr}(\mathbf{G}_0 \mathbf{C}) - \text{Tr}(\mathbf{C}) \\ V\text{Tr}(\mathbf{C}) - \text{Tr}(\mathbf{G}_0 \mathbf{C}) \end{bmatrix}. \end{aligned}$$

Thus,

$$\hat{\sigma}^2 = \frac{1}{(V-1)NV} \sum_{i=1}^N \sum_{j \neq k}^V (Y_{ij} - \hat{\mu})(Y_{ik} - \hat{\mu}), \quad (2.14)$$

the mean of the within-subject off-diagonal elements of \mathbf{C} , and

$$\begin{aligned} \hat{\tau}^2 &= \frac{1}{(V-1)NV} \left[V \sum_{i=1}^N \sum_{j=1}^V (Y_{ij} - \hat{\mu})^2 - \sum_{i=1}^N \sum_{j=1}^V \sum_{k=1}^V (Y_{ij} - \hat{\mu})(Y_{ik} - \hat{\mu}) \right] \\ &= \frac{1}{(V-1)NV} \left[(V-1) \sum_{i=1}^N \sum_{j=1}^V (Y_{ij} - \hat{\mu})^2 - \sum_{i=1}^N \sum_{j \neq k}^V (Y_{ij} - \hat{\mu})(Y_{ik} - \hat{\mu}) \right] \\ &= \frac{1}{NV} \sum_{i=1}^N \sum_{j=1}^V (Y_{ij} - \hat{\mu})^2 - \hat{\sigma}^2, \end{aligned} \quad (2.15)$$

the mean of the diagonal elements of \mathbf{C} minus the mean of the within-subject off-diagonal elements.

2.3.3 Group Scalar Mean

The next simple case involves fitting a separate scalar mean to each of the two groups:

$$\mathbf{Y}_{NV \times 1} = \mathbf{X}_{NV \times 2} \boldsymbol{\beta}_{2 \times 1} + (\mathbf{I}_{N \times N} \otimes \mathbf{1}_{V \times 1}) \boldsymbol{\delta}_{N \times 1} + \boldsymbol{\epsilon}_{NV \times 1}, \quad (2.16)$$

where $\mathbf{X} = [\mathbf{x}_1 \ \mathbf{x}_2]$ is the fixed effects design matrix defined by $\mathbf{x}_1 = \mathbf{1}_{NV \times 1}$ and $\mathbf{x}_2 = [\mathbf{0}_{Vn_1 \times 1} \ \mathbf{1}_{Vn_2 \times 1}]'$, $\boldsymbol{\beta} = [\beta_1 \ \beta_2]'$ is the vector of unknown fixed parameters, $\boldsymbol{\delta}$ is the vector of subject-specific random intercepts, with $\boldsymbol{\delta} \sim \text{MVN}(\mathbf{0}, \sigma^2 \mathbf{I})$, and $\boldsymbol{\epsilon}$ is the vector of residual errors, with $\boldsymbol{\epsilon} \sim \text{MVN}(\mathbf{0}, \tau^2 \mathbf{I})$.

The variance structure here is the same as for model 2.9, so in order to determine the existence of a closed-form ML estimator for $\boldsymbol{\beta}$, it suffices to demonstrate that \mathbf{x}_1 and \mathbf{x}_2 are contained in the span of two eigenvectors of $\boldsymbol{\Sigma}$ as given in Equation 2.10. This is immediate, as both \mathbf{x}_1 and \mathbf{x}_2 are each the sum of eigenvectors of $\boldsymbol{\Sigma}$ with the same eigenvalues, and thus are themselves eigenvectors:

$$\begin{aligned} \mathbf{x}_1 &= \sum_{i=1}^N \mathbf{p}_{i1} \\ \mathbf{x}_2 &= \sum_{i=n_1+1}^N \mathbf{p}_{i1}. \end{aligned}$$

The ML estimates of $\hat{\boldsymbol{\beta}}$ are thus given by the OLS solution, $(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y}$. Since a closed-form solution exists for $\boldsymbol{\beta}$, we may apply the same reasoning as in Section 2.3.2 and obtain the same closed-form expressions for $\hat{\sigma}^2$ and $\hat{\tau}^2$ given in Equations 2.14 and 2.15.

2.3.4 Voxelwise Scalar Means

A more realistic example for DTI data is a model which fits a separate mean at each voxel, allowing for spatial inhomogeneity in diffusion structure:

$$\mathbf{Y}_{NV \times 1} = (\mathbf{1}_{N \times 1} \otimes \mathbf{I}_{V \times V}) \boldsymbol{\beta}_{V \times 1} + (\mathbf{I}_{N \times N} \otimes \mathbf{1}_{V \times 1}) \boldsymbol{\delta}_{N \times 1} + \boldsymbol{\epsilon}_{NV \times 1}, \quad (2.17)$$

where $\boldsymbol{\beta} = [\beta_1 \ \dots \ \beta_V]'$ is a vector of unknown fixed parameters, $\boldsymbol{\delta} \sim \text{MVN}(\mathbf{0}, \sigma^2 \mathbf{I})$ is the vector of subject-specific random intercepts, and $\boldsymbol{\epsilon} \sim \text{MVN}(\mathbf{0}, \tau^2 \mathbf{I})$ is the vector of residual errors.

The variance structure, Σ , is the same as in the previous two examples (Equation 2.10), so to prove the existence of a closed-form for $\hat{\beta}$ it again suffices to find a set of V eigenvectors of Σ that span the columns of $\mathbf{X} = (\mathbf{1}_{N \times 1} \otimes \mathbf{I}_{V \times V})$.

In this case, the columns of \mathbf{X} are not eigenvectors of Σ , so it is necessary to exhibit a separate set of V eigenvectors that span the column space of \mathbf{X} . To this end, I extend the notation introduced in Equation 2.12, and let $\mathbf{p}_{.j} = \sum_{i=1}^N \mathbf{p}_{ij}$ for $j = 1, \dots, V$. Thus, each $\mathbf{p}_{.j}$ is the sum across blocks of corresponding columns of $(\mathbf{I}_{N \times N} \otimes \mathbf{P})$, the matrix of eigenvectors of Σ . And, since corresponding columns across blocks of $(\mathbf{I}_{N \times N} \otimes \mathbf{P})$ have equal eigenvalues ($V\sigma^2 + \tau^2$ for column 1, τ^2 for columns 2 to V), each $\mathbf{p}_{.j}$ is an eigenvector of Σ .

The claim is that the set of V eigenvectors of Σ , $\{\mathbf{p}_{.j}\}_{j=1}^V$ spans the column space of \mathbf{X} . Let $\mathbf{E} = [\mathbf{p}_{.1} \ \dots \ \mathbf{p}_{.V}]$, let $\mathbf{z}_1 = \frac{1}{V} \mathbf{1}_{V \times 1}$, and, for $2 \leq j \leq V$, let \mathbf{z}_j be the vector defined by $z_{jl} = \frac{1}{V} ((1 - \delta_{jl}) + (1 - V)\delta_{jl})$, where δ is Kronecker's δ . That is, for $j \geq 2$, \mathbf{z}_j is the column vector with $\frac{(1-V)}{V}$ in the i^{th} row and $\frac{1}{V}$ in all other rows.

With \mathbf{E} and \mathbf{z}_j defined above, $\mathbf{E}\mathbf{z}_j = \mathbf{x}_j$ for $j = 1, \dots, V$, and thus the V columns of \mathbf{X} are spanned by the V columns of \mathbf{E} , which are eigenvectors of Σ . For instance, with $N = 2$, $V = 3$,

$$\begin{aligned} \mathbf{E}\mathbf{z}_1 &= \begin{bmatrix} 1 & 1 & 1 \\ 1 & -1 & 0 \\ 1 & 0 & -1 \end{bmatrix} \begin{bmatrix} \frac{1}{3} \\ \frac{1}{3} \\ \frac{1}{3} \end{bmatrix} \\ &= \begin{bmatrix} 3(\frac{1}{3}) \\ \frac{1}{3} - \frac{1}{3} \\ \frac{1}{3} - \frac{1}{3} \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \\ &= \mathbf{x}_1, \end{aligned}$$

$$\begin{aligned}
\mathbf{Ez}_2 &= \begin{bmatrix} 1 & 1 & 1 \\ 1 & -1 & 0 \\ 1 & 0 & -1 \end{bmatrix} \begin{bmatrix} \frac{1}{3} \\ \frac{-2}{3} \\ \frac{1}{3} \end{bmatrix} \\
&= \begin{bmatrix} \frac{1}{3} - \frac{2}{3} + \frac{1}{3} \\ \frac{1}{3} - \frac{-2}{3} \\ \frac{1}{3} - \frac{1}{3} \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} \\
&= \mathbf{x}_2,
\end{aligned}$$

and similarly $\mathbf{Ez}_3 = \mathbf{x}_3$.

Thus, there is a closed-form expression for $\hat{\beta}$, which may again be found by OLS calculations. And we may once more apply Theorem 2.2.2 as in Sections 2.3.2 and 2.3.3 to obtain the same closed-form expressions for $\hat{\sigma}^2$ and $\hat{\tau}^2$ given in Equations 2.14 and 2.15.

The details are omitted, but we may also extend the closed-form estimation of model 2.17 to accommodate separate voxelwise mean structure per group, in exactly the same manner in which the closed-form estimation of model 2.9 was extended to accommodate model 2.16.

2.3.5 The General Case For Homogeneous Compound Symmetry

Next, I discuss general conditions under which closed-form ML estimates exist for the fixed effects parameters in the models with compound symmetric covariance structures described in the previous sections. Consider the class of models given by:

$$\mathbf{Y}_{NV \times 1} = \mathbf{X}_{NV \times r} \boldsymbol{\beta}_{r \times 1} + (\mathbf{I}_{N \times N} \otimes \mathbf{1}_{V \times 1}) \boldsymbol{\delta}_{N \times 1} + \boldsymbol{\epsilon}_{NV \times 1}, \quad (2.18)$$

where \mathbf{X} is an arbitrary fixed-effects design matrix, $\boldsymbol{\beta}$ is the vector of r fixed effects parameters, $\boldsymbol{\delta}$ is a random vector of subject-specific intercepts, with $\boldsymbol{\delta} \sim \text{MVN}(\mathbf{0}, \sigma^2 \mathbf{I})$ and $\boldsymbol{\epsilon} \sim \text{MVN}(\mathbf{0}, \tau^2 \mathbf{I})$ is a vector of residual errors.

The covariance structure is compound symmetric, as described in Equation 2.10, and the eigensystem of $\boldsymbol{\Sigma}$ was derived in Section 2.3.2. There are two eigenvalues, $V\sigma^2 + \tau^2$ (multiplicity N) and τ^2 (multiplicity $N(V - 1)$). $\boldsymbol{\Sigma}$ is of full rank NV , and so the spectral decomposition of $\boldsymbol{\Sigma}$ partitions \mathbb{R}^{NV} into two orthogonal subspaces: U_1 and U_2 , spanned by eigenvectors of $\boldsymbol{\Sigma}$ corresponding to $V\sigma^2 + \tau^2$ and τ^2 , respectively.

As discussed in Section 2.3.2, by Theorem 2.2.2, whenever a closed-form estimate exists for $\hat{\beta}$, closed-form estimates also exist for σ^2 and τ^2 (and are given by Equations 2.14 and 2.15). And, by Theorem 2.2.1, a necessary and sufficient condition for the existence of a closed-form estimate for $\hat{\beta}$ is that the r columns of \mathbf{X} be spanned by r eigenvectors of Σ or, equivalently, by r vectors in $U_1 \cup U_2$.

Based on these facts, we may draw several interesting conclusions:

Theorem 2.3.1. *For the model described in Equation 2.18, a closed-form expression exists for $\hat{\beta}$ whenever the columns of \mathbf{X} represent:*

1. *Between-subject covariates: columns of the form $\mathbf{x} = \mathbf{1}_{V \times 1} \otimes \boldsymbol{\alpha}_{N \times 1}$ for arbitrary $\boldsymbol{\alpha}$,*
2. *Within-subject spatial contrasts: columns $\mathbf{x} = [x_{11} \ x_{12} \ \dots \ x_{1V} \ x_{21} \ \dots \ x_{NV}]'$ with the property that $\sum_{j=1}^V x_{ij} = 0 \ \forall i$, and*
3. *Within-subject spatial effects that are the same within groups of subjects defined by between-subject covariates: columns of the form $\mathbf{x} = \mathbf{x}_{\alpha_{N \times 1}} \otimes \boldsymbol{\gamma}_{V \times 1}$ for arbitrary $\boldsymbol{\gamma}$, where \mathbf{x}_{α} is a between-subject covariate vector consisting of ones and zeros.*

Proof. First, for any column vector \mathbf{x} of the form described in part 1 (between-subject effects), $\mathbf{x} \in U_1$:

$$\begin{aligned} \Sigma \mathbf{x} &= \Sigma(\mathbf{1}_{V \times 1} \otimes \boldsymbol{\alpha}_{N \times 1}) \\ &= \mathbf{1}_{V \times 1} \otimes (V\sigma^2 + \tau^2)\boldsymbol{\alpha}_{N \times 1} \\ &= (V\sigma^2 + \tau^2)\mathbf{1}_{V \times 1} \otimes \boldsymbol{\alpha}_{N \times 1} \\ &= (V\sigma^2 + \tau^2)\mathbf{x}. \end{aligned}$$

Thus, \mathbf{x} is itself an eigenvector of Σ .

Next, for any column vector \mathbf{x} of the form described in part 2 (within-subject spatial contrasts), $\mathbf{x} \in U_2$:

$$\begin{aligned} \Sigma \mathbf{x} &= \Sigma[x_{11} \ x_{12} \ \dots \ x_{1V} \ x_{21} \ \dots \ x_{NV}]' \\ &= \tau^2[x_{11} \ x_{12} \ \dots \ x_{1V} \ x_{21} \ \dots \ x_{NV}]' \\ &= \tau^2\mathbf{x}. \end{aligned}$$

That is, again, \mathbf{x} is itself an eigenvector of Σ .

Finally, for part 3 above, consider an \mathbf{x} which represents a spatial effect that is the same for all subjects with values of one on a between-subject covariate represented by another column, $\mathbf{x}_\alpha = \mathbf{1}_{V \times 1} \otimes \boldsymbol{\alpha}_{N \times 1}$ of \mathbf{X} , where $\boldsymbol{\alpha}$ consists of only zeros and ones.

For each subject, i , let $\bar{x}_{i.} = \frac{1}{V} \sum_{j=1}^V x_{ij}$. Since \mathbf{x} is the same for all subjects for whom $\alpha_i = 1$, $\bar{x}_{i.}$ is also constant for these subjects, and thus:

$$\bar{\mathbf{x}}_W = \mathbf{1}_{V \times 1} \otimes [\bar{x}_{1.} \ \dots \ \bar{x}_{N.}]' \propto \mathbf{x}_\alpha.$$

Furthermore, $\mathbf{x} - \bar{\mathbf{x}}_W \in U_2$, as centering the spatial effect \mathbf{x} within subjects results in a vector that sums to zero for each subject. Putting this together, we have:

$$\begin{aligned} \boldsymbol{\Sigma} \mathbf{x} &= \boldsymbol{\Sigma}(\mathbf{x} - \bar{\mathbf{x}}_W) + \boldsymbol{\Sigma} \bar{\mathbf{x}}_W \\ &= \tau^2(\mathbf{x} - \bar{\mathbf{x}}_W) + (V\sigma^2 + \tau^2)\bar{\mathbf{x}}_W. \end{aligned}$$

Thus, \mathbf{x} is spanned by two eigenvectors, one of which was already “used” to span \mathbf{x}_α . Put another way, adding \mathbf{x} as a column to \mathbf{X} requires adding the single eigenvector $(\mathbf{x} - \bar{\mathbf{x}}_W)$ for the columns of \mathbf{X} to be spanned by eigenvectors of $\boldsymbol{\Sigma}$. \square

As a corollary to part 3 of this theorem, we have the following useful result:

Corollary 2.3.2. *When \mathbf{X} includes an overall intercept, $\mathbf{1}_{N \times 1}$ as one of its columns, then a closed-form estimate exists whenever the other columns of \mathbf{X} include any spatial effect that is the same for all subjects, as well as any member of the classes of vectors described in parts 1, 2 or 3 of Theorem 2.3.1.*

Proof. This is a consequence of the proof of part 3 of Theorem 2.3.1, with $\boldsymbol{\alpha}_{N \times 1} = \mathbf{1}_{N \times 1}$. \square

Thus, for compound symmetric models with an overall intercept, closed-form expressions exist for the ML estimates of mean and variance terms for a wide variety of mean structures. The only circumstance under which a closed-form expression may not exist is when one of the columns of \mathbf{X} contains subject-specific spatial covariates. However, even in this case, we may re-express such an effect as the sum of two effects, $\mathbf{x} = \bar{\mathbf{x}}_W + (\mathbf{x} - \bar{\mathbf{x}}_W)$, and obtain closed-form estimates of the two effects separately, as $\bar{\mathbf{x}}_W$ is a between-subject effect and $\mathbf{x} - \bar{\mathbf{x}}_W$ is a within-subject spatial contrast.

2.3.6 REML Estimation

Although ML estimation has many favorable asymptotic properties, ML estimates of variance components in LMEs are biased downwards. Thus, ML estimation may lead to spurious findings of group differences due to unrealistically low estimates of standard errors of fixed effects. Another standard approach to estimation of variance components in mixed models is Restricted Maximum Likelihood (REML), which produces unbiased estimates under very general conditions[35].

REML estimation differs from ML in its adjustment for loss of degrees of freedom due to the estimation of fixed effects. In the balanced models I have been considering in this and subsequent sections, this adjustment takes the form of a simple algebraic modification to the closed-form ML variance component estimates. The relationship between the ML and REML variance component estimates can be derived from Lange and Laird (1989)[54].

The ML estimate of τ^2 , $\hat{\tau}^2$, which was derived in Equation 2.15, is related to the REML estimate, $\tilde{\tau}^2$, by (Equation 11 from Lange and Laird):

$$\tilde{\tau}^2 = \frac{N(V - m)}{N(V - m) - f(w - m)} \hat{\tau}^2, \quad (2.19)$$

where m is the number of random effects, f is the number of between-subject effects and w is the number of within-subject effects. Note that $m = 1$ for all the models under consideration in this section, but will increase with variance heterogeneity and longitudinal designs in subsequent sections.

The ML estimate of σ^2 , $\hat{\sigma}^2$ (Equation 2.14) is related to the REML solution, $\tilde{\sigma}^2$, by (Equation 14 of Lange and Laird):

$$\tilde{\sigma}^2 = \frac{N}{N - f} \hat{\sigma}^2. \quad (2.20)$$

In the context of exploratory model-building for DTI data, I am interested in using ML estimation to fit the LMEs discussed in this and subsequent sections, in order to use likelihood ratio significance testing to compare the fits of nested models. However, once a final model has been settled on, it will be advisable to refit the model using REML estimation in order to obtain unbiased estimates of the variance parameters.

2.3.7 Summary

I have shown that closed-form expressions exist for the maximum likelihood estimators for mean and variance parameters in a series of simple models with compound symmetric covariance structures. The mean structures considered included a single overall scalar mean term fit to all voxels for subjects in a single group, a separate mean fit to each voxel for subjects in a single group, and two-group variants of these structures. In each model, there was a random intercept term and residuals were assumed to be i.i.d.

I also generalized these results, and showed that closed-form estimates exist whenever the fixed effects include only between-subject effects, within-subject spatial contrasts and group-specific within-subject spatial effects (where the group is defined by a between-subject effect). Also, in the special case in which the model includes an overall intercept term, closed-form estimates exist for within-subject spatial effects that are the same for all subjects.

The one case in which closed-form estimation is not possible is for spatially-varying covariates (that is, effects that change both between- and within-subject). However, this limitation can be overcome by centering the effect, breaking it into distinct between-subject and within-subject spatial contrast effects.

The next step is to broaden the class of models under consideration by allowing group heterogeneity in covariance parameters (either in the between-subject variance, the residual variance, or both).

2.4 CLOSED-FORM ML ESTIMATION UNDER HETEROGENEOUS VARIANCE

In this section, I examine the extent to which the results in Section 2.3 can be extended to models with between-group heterogeneity in covariance structure. I consider heterogeneity in residual variance, heterogeneity in between-subject variance, and simultaneous residual and between-subjects heterogeneity.

2.4.1 Group Heterogeneity In Residual Variance

Suppose we adopt the model described in Equation 2.16 in Section 2.3.3:

$$\mathbf{Y}_{NV \times 1} = \mathbf{X}_{NV \times 2} \boldsymbol{\beta}_{2 \times 1} + (\mathbf{I}_{N \times N} \otimes \mathbf{1}_{V \times 1}) \boldsymbol{\delta}_{N \times 1} + \boldsymbol{\epsilon}_{NV \times 1}, \quad (2.21)$$

where $\mathbf{X} = [\mathbf{x}_1 \ \mathbf{x}_2]$ is the fixed effects design matrix defined by $\mathbf{x}_1 = \mathbf{1}_{NV \times 1}$ and $\mathbf{x}_2 = [\mathbf{0}_{Vn_1 \times 1} \ \mathbf{1}_{Vn_2 \times 1}]'$, $\boldsymbol{\beta} = [\beta_1 \ \beta_2]'$ is the vector of unknown fixed parameters, $\boldsymbol{\delta} \sim \text{MVN}(\mathbf{0}, \sigma^2 \mathbf{I})$ is the vector of subject-specific random intercepts, but with the modified assumption that $\boldsymbol{\epsilon} \sim \text{MVN}(\mathbf{0}, \text{diag}(\tau_1^2 \mathbf{I}_{n_1 \times n_1}, \tau_2^2 \mathbf{I}_{n_2 \times n_2}))$. That is, that there is a different residual variance parameter for group 1 (τ_1^2) than for group 2 (τ_2^2). Such a model might be appropriate in practice if there is reason to think that the amount of within-subject noise is greater in one group than in the other, while the amount of underlying biological variability between subjects is comparable in both. For instance, perhaps the two groups were scanned on different equipment, or participants from one group had greater head motion in the scanner than the other group.

In this model, $\boldsymbol{\Sigma} = \tau_1^2 \mathbf{G}_0 + \tau_2^2 \mathbf{G}_1 + \sigma^2 \mathbf{G}_2$, where $\mathbf{G}_0 = \text{diag}(\mathbf{I}_{n_1 \times n_1}, \mathbf{0}_{n_2 \times n_2}) \otimes \mathbf{I}_{V \times V}$, $\mathbf{G}_1 = \mathbf{I}_{NV \times NV} - \mathbf{G}_0$, and $\mathbf{G}_2 = \mathbf{I}_{N \times N} \otimes \mathbf{J}_{V \times V}$. Each subject-specific block of $\boldsymbol{\Sigma}$ has a compound symmetric structure, as described in Equation 2.10. The eigenvalues of the blocks corresponding to group 1 are $V\sigma^2 + \tau_1^2$ (multiplicity 1) and τ_1^2 (multiplicity $V - 1$), while the eigenvalues of the blocks for group 2 participants are $V\sigma^2 + \tau_2^2$ (multiplicity 1) and τ_2^2 (multiplicity $V - 1$). Thus, $\boldsymbol{\Sigma}$ itself has four distinct but not linearly independent eigenvalues: $V\sigma^2 + \tau_1^2$ (multiplicity n_1), $V\sigma^2 + \tau_2^2$ (multiplicity n_2), τ_1^2 (multiplicity $n_1(V - 1)$) and τ_2^2 (multiplicity $n_2(V - 1)$).

$\boldsymbol{\Sigma}$ can be diagonalized much as in Equation 2.11:

$$\boldsymbol{\Sigma} = (\mathbf{I}_{N \times N} \otimes \mathbf{P})(\text{diag}(\mathbf{I}_{n_1 \times n_1} \otimes \mathbf{D}_1, \mathbf{I}_{n_2 \times n_2} \otimes \mathbf{D}_2))(\mathbf{I}_{N \times N} \otimes \mathbf{P})^{-1}. \quad (2.22)$$

\mathbf{P} is the same here as in Section 2.3.2, \mathbf{D}_1 is the $V \times V$ diagonal matrix with $V\sigma^2 + \tau_1^2$ as the entry in row 1, column 1 and τ_1^2 along the rest of the diagonal, and \mathbf{D}_2 is the $V \times V$

diagonal matrix with $V\sigma^2 + \tau_2^2$ as the entry in row 1, column 1 and τ_2^2 along the rest of the diagonal. I extend the notation developed in Equation 2.12, and define:

$$\mathbf{p}_{\cdot 1n_1} = \sum_{i=1}^{n_1} \mathbf{p}_{i1} = \begin{bmatrix} \mathbf{1}_{n_1 V \times 1} \\ \mathbf{0}_{n_2 V \times 1} \end{bmatrix}, \text{ and } \mathbf{p}_{\cdot 1n_2} = \sum_{i=n_1+1}^N \mathbf{p}_{i1} = \begin{bmatrix} \mathbf{0}_{n_1 V \times 1} \\ \mathbf{1}_{n_2 V \times 1} \end{bmatrix}$$

These are each eigenvectors of Σ (with eigenvalues $V\sigma^2 + \tau_1^2$ and $V\sigma^2 + \tau_2^2$, respectively), and they span the columns of \mathbf{X} :

$$\begin{aligned} \mathbf{x}_1 &= \mathbf{p}_{\cdot 1n_1} + \mathbf{p}_{\cdot 1n_2} \\ \mathbf{x}_2 &= \mathbf{p}_{\cdot 1n_2}. \end{aligned}$$

Thus, there is a closed-form solution to the maximum likelihood estimate $\hat{\beta}$, obtained from the least squares solution $(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y}$.

Note that in this case, there would be no closed-form expression for $\hat{\beta}$ if there were only a single mean term fit to both groups. If \mathbf{X} were just a single column of ones, no one eigenvector would span it. In other words, separating the means by group is a precondition for closed-form estimability of the mean structure when also separating residual variance by group.

The next step is to ascertain whether there is a closed-form expression for the maximum likelihood estimates $\hat{\tau}_1^2$, $\hat{\tau}_2^2$ and $\hat{\sigma}^2$. Since the \mathbf{G} matrices all commute with one another, they are simultaneously diagonalizable, and a closed-form expression exists for $\hat{\beta}$, so Theorem 2.2.2 applies. And, since there are four distinct eigenvalues of Σ but only three \mathbf{G} matrices, there are no closed-form expressions for the maximum likelihood estimates of the variance parameters.

2.4.2 Group Heterogeneity In Between-Subject Variance

The next model I consider incorporates a single variance parameter for residual variance across groups, but allows for heterogeneity in between-subjects variance (i.e. unequal variance parameters for the subject-specific random intercept term). This model would be applicable in situations in which subjects are scanned under similar conditions, but there is reason to believe that the degree of biological variability is greater in one group than in the other. This would be a reasonable assumption in many studies comparing a group of participants with some neuropathological condition to controls, where the range of variation will often be greater in the neuropathology group.

The model is:

$$\mathbf{Y}_{NV \times 1} = \mathbf{X}_{NV \times 2} \boldsymbol{\beta}_{2 \times 1} + (\boldsymbol{\Delta} \otimes \mathbf{1}_{V \times 1}) \boldsymbol{\delta}_{1_{N \times 1}} + ((\mathbf{I}_{N \times N} - \boldsymbol{\Delta}) \otimes \mathbf{1}_{V \times 1}) \boldsymbol{\delta}_{2_{N \times 1}} + \boldsymbol{\epsilon}_{NV \times 1}, \quad (2.23)$$

where $\boldsymbol{\Delta} = \text{diag}(\mathbf{I}_{n_1 \times n_1}, \mathbf{0}_{n_2 \times n_2})$, $\boldsymbol{\delta}_1 \sim \text{MVN}(\mathbf{0}, \sigma_1^2 \mathbf{I})$, $\boldsymbol{\delta}_2 \sim \text{MVN}(\mathbf{0}, \sigma_2^2 \mathbf{I})$, $\mathbf{X} = [\mathbf{x}_1 \ \mathbf{x}_2]$ is the fixed effects design matrix defined by $\mathbf{x}_1 = \mathbf{1}_{NV \times 1}$ and $\mathbf{x}_2 = [\mathbf{0}_{V n_1 \times 1} \ \mathbf{1}_{V n_2 \times 1}]'$, $\boldsymbol{\beta} = [\beta_1 \ \beta_2]'$ is the vector of unknown fixed parameters, and $\boldsymbol{\epsilon} \sim \text{MVN}(\mathbf{0}, \tau^2 \mathbf{I})$ is the vector of residual errors.

2.4.2.1 Closed-Form Estimation Of Mean Parameters In this model, $\boldsymbol{\Sigma} = \tau^2 \mathbf{G}_0 + \sigma_1^2 \mathbf{G}_1 + \sigma_2^2 \mathbf{G}_2$, where $\mathbf{G}_0 = \mathbf{I}_{NV \times NV}$, $\mathbf{G}_1 = \boldsymbol{\Delta}_{N \times N} \otimes \mathbf{J}_{V \times V}$ and $\mathbf{G}_2 = (\mathbf{I}_{N \times N} - \boldsymbol{\Delta}_{N \times N}) \otimes \mathbf{J}_{V \times V}$.

$\boldsymbol{\Sigma}$ can be diagonalized in a manner analogous to Equation 2.22:

$$\boldsymbol{\Sigma} = (\mathbf{I}_{N \times N} \otimes \mathbf{P})(\text{diag}(\mathbf{I}_{n_1 \times n_1} \otimes \mathbf{D}_1, \mathbf{I}_{n_2 \times n_2} \otimes \mathbf{D}_2))(\mathbf{I}_{N \times N} \otimes \mathbf{P})^{-1}.$$

Here, \mathbf{P} is as in previous sections, \mathbf{D}_1 is the $V \times V$ diagonal matrix with $V\sigma_1^2 + \tau^2$ in the first position and τ^2 in all other positions, and \mathbf{D}_2 is the $V \times V$ diagonal matrix with $V\sigma_2^2 + \tau^2$ in the first position and τ^2 in all other positions. Thus there are three distinct, linearly independent eigenvalues.

By the same argument given in Section 2.4.1, eigenvectors $\mathbf{p}_{.1n_1}$ and $\mathbf{p}_{.1n_2}$ span the two columns of \mathbf{X} , so the OLS solution is a closed-form expression for the maximum likelihood estimates as well.

2.4.2.2 Closed-Form Estimation Of Covariance Parameters With regard to the existence of closed-form solutions for the variance parameters, note that again all three \mathbf{G} matrices commute with one another ($\mathbf{G}_0 = \mathbf{I}$ commutes with any conformable matrix and $\mathbf{G}_1\mathbf{G}_2 = \mathbf{G}_2\mathbf{G}_1 = \mathbf{0}$). Therefore, the existence of exactly three eigenvalues for $\mathbf{\Sigma}$ guarantees the existence of closed-form solutions for $\hat{\tau}^2$, $\hat{\sigma}_1^2$ and $\hat{\sigma}_2^2$. These solutions can be derived from Equation 2.13.

First, note the values of the following quantities:

$$\begin{aligned}
\text{Tr}(\mathbf{G}_0\mathbf{G}_0) &= NV \\
\text{Tr}(\mathbf{G}_0\mathbf{G}_1) &= \text{Tr}(\mathbf{G}_1\mathbf{G}_0) = n_1V \\
\text{Tr}(\mathbf{G}_0\mathbf{G}_2) &= \text{Tr}(\mathbf{G}_2\mathbf{G}_0) = n_2V \\
\text{Tr}(\mathbf{G}_1\mathbf{G}_1) &= n_1V^2 \\
\text{Tr}(\mathbf{G}_2\mathbf{G}_2) &= n_2V^2 \\
\text{Tr}(\mathbf{G}_1\mathbf{G}_2) &= \text{Tr}(\mathbf{G}_2\mathbf{G}_1) = 0 \\
\text{Tr}(\mathbf{G}_0\mathbf{C}) &= \sum_{i=1}^N \sum_{j=1}^V (Y_{ij} - \hat{\mu})^2 \\
\text{Tr}(\mathbf{G}_1\mathbf{C}) &= \sum_{i=1}^{n_1} \sum_{j=1}^V \sum_{k=1}^V (Y_{ij} - \hat{\mu})(Y_{ik} - \hat{\mu}), \text{ and} \\
\text{Tr}(\mathbf{G}_2\mathbf{C}) &= \sum_{i=n_1+1}^{n_2} \sum_{j=1}^V \sum_{k=1}^V (Y_{ij} - \hat{\mu})(Y_{ik} - \hat{\mu}).
\end{aligned}$$

Then the ML estimates of the covariance parameters are given by:

$$\begin{aligned}
\begin{bmatrix} \hat{\tau}^2 \\ \hat{\sigma}_1^2 \\ \hat{\sigma}_2^2 \end{bmatrix} &= \begin{bmatrix} NV & n_1V & n_2V \\ n_1V & n_1V^2 & 0 \\ n_2V & 0 & n_2V^2 \end{bmatrix}^{-1} \begin{bmatrix} \text{Tr}(\mathbf{G}_0\mathbf{C}) \\ \text{Tr}(\mathbf{G}_1\mathbf{C}) \\ \text{Tr}(\mathbf{G}_2\mathbf{C}) \end{bmatrix} \\
&= \frac{1}{(V-1)NV} \begin{bmatrix} V & -1 & -1 \\ -1 & \frac{Vn_1+(V-1)n_2}{Vn_1} & \frac{1}{V} \\ -1 & \frac{1}{V} & \frac{Vn_2+(V-1)n_1}{Vn_2} \end{bmatrix} \begin{bmatrix} \text{Tr}(\mathbf{G}_0\mathbf{C}) \\ \text{Tr}(\mathbf{G}_1\mathbf{C}) \\ \text{Tr}(\mathbf{G}_2\mathbf{C}) \end{bmatrix}. \quad (2.24)
\end{aligned}$$

Thus, the residual variance, $\hat{\tau}^2$, is again the average of the diagonal elements of \mathbf{C} minus the mean of the within-subject off-diagonal elements, as in Equation 2.15.

The between-subjects variance terms $\hat{\sigma}_1^2$ and $\hat{\sigma}_2^2$ are linear combinations of the within-subject diagonal and off-diagonal elements of \mathbf{C} . Interestingly, although the model allows for heterogeneous within-subject covariance between groups, the ML estimates of these covariance terms for each group depend partly upon the empirical within-subject covariance for the *other* group as well.

For example, for the case where $n_1 = n_2 = \frac{N}{2}$, Equation 2.24 yields:

$$\hat{\sigma}_1^2 = \frac{1}{(V-1)NV} \left(2\text{Tr}(\mathbf{G}_1\mathbf{C}) + \frac{1}{V} (\text{Tr}(\mathbf{G}_2\mathbf{C}) - \text{Tr}(\mathbf{G}_1\mathbf{C})) - \text{Tr}(\mathbf{C}) \right). \quad (2.25)$$

This form of $\hat{\sigma}_1^2$ is intriguing. It demonstrates what we may have anticipated, that the estimate is similar in form to the estimate of σ^2 from the homogeneous variance model, but with the off-diagonal covariance terms from group 1 “counted twice,” and the residual variance based on the diagonal terms of \mathbf{C} for both groups. However, there is also a penalty term of sorts in the form of $\frac{1}{V} (\text{Tr}(\mathbf{G}_2\mathbf{C}) - \text{Tr}(\mathbf{G}_1\mathbf{C}))$. This means that, if the covariance terms are smaller in group 2 than in group 1, we will be shrinking the estimate of the group 1 covariance a little toward the group 2 covariance, and vice-versa.

2.4.3 Combined Residual And Between-Subjects Variance Heterogeneity

Finally, I consider closed-form estimation for the model with group heterogeneity both in residual and in between-subjects variance. This model would be appropriate in situations in which there is reason to believe both that there is more subject-to-subject biological variability and that there is more within-subject noise in one group (likely the group with neuropathology) than in the other group. The model is:

$$\mathbf{Y}_{NV \times 1} = \mathbf{X}_{NV \times 2} \boldsymbol{\beta}_{2 \times 1} + (\boldsymbol{\Delta} \otimes \mathbf{1}_{V \times 1}) \boldsymbol{\delta}_{1_{N \times 1}} + ((\mathbf{I}_{N \times N} - \boldsymbol{\Delta}) \otimes \mathbf{1}_{V \times 1}) \boldsymbol{\delta}_{2_{N \times 1}} + \boldsymbol{\epsilon}_{NV \times 1}, \quad (2.26)$$

where where $\boldsymbol{\Delta} = \text{diag}(\mathbf{I}_{n_1 \times n_1}, \mathbf{0}_{n_2 \times n_2})$, $\boldsymbol{\delta}_1 \sim \text{MVN}(\mathbf{0}, \sigma_1^2 \mathbf{I})$, $\boldsymbol{\delta}_2 \sim \text{MVN}(\mathbf{0}, \sigma_2^2 \mathbf{I})$, $\mathbf{X} = [\mathbf{x}_1 \ \mathbf{x}_2]$ is the fixed effects design matrix defined by $\mathbf{x}_1 = \mathbf{1}_{NV \times 1}$ and $\mathbf{x}_2 = [\mathbf{0}_{Vn_1 \times 1} \ \mathbf{1}_{Vn_2 \times 1}]'$, $\boldsymbol{\beta} = [\beta_1 \ \beta_2]'$ is the vector of unknown fixed parameters, and

$$\boldsymbol{\epsilon} \sim \text{MVN}(\mathbf{0}, \text{diag}(\tau_1^2 \mathbf{I}_{n_1 \times n_1}, \tau_2^2 \mathbf{I}_{n_2 \times n_2})).$$

The covariance matrix of \mathbf{Y} breaks down into four additive components in this model:

$$\boldsymbol{\Sigma} = \tau_1^2 \mathbf{G}_0 + \tau_2^2 \mathbf{G}_1 + \sigma_1^2 \mathbf{G}_2 + \sigma_2^2 \mathbf{G}_3,$$

where $\mathbf{G}_0 = \text{diag}(\mathbf{I}_{n_1 \times n_1}, \mathbf{0}_{n_2 \times n_2}) \otimes \mathbf{I}_{V \times V}$, $\mathbf{G}_1 = \mathbf{I}_{NV \times NV} - \mathbf{G}_0$, $\mathbf{G}_2 = \boldsymbol{\Delta}_{N \times N} \otimes \mathbf{J}_{V \times V}$ and $\mathbf{G}_3 = (\mathbf{I}_{N \times N} - \boldsymbol{\Delta}_{N \times N}) \otimes \mathbf{J}_{V \times V}$.

$\boldsymbol{\Sigma}$ can be diagonalized as:

$$\boldsymbol{\Sigma} = (\mathbf{I}_{N \times N} \otimes \mathbf{P})(\text{diag}(\mathbf{I}_{n_1 \times n_1} \otimes \mathbf{D}_1, \mathbf{I}_{n_2 \times n_2} \otimes \mathbf{D}_2))(\mathbf{I}_{N \times N} \otimes \mathbf{P})^{-1},$$

where \mathbf{P} is as defined in Section 2.4.2, \mathbf{D}_1 is the diagonal matrix with $V\sigma_1^2 + \tau_1^2$ in the first position and τ_1^2 in all other positions, and \mathbf{D}_2 is the diagonal matrix with $V\sigma_2^2 + \tau_2^2$ in the first position and τ_2^2 in all other positions. Thus, $\boldsymbol{\Sigma}$ has four distinct, linearly independent eigenvalues: $V\sigma_1^2 + \tau_1^2$ (multiplicity n_1), τ_1^2 (multiplicity $n_1(V - 1)$), $V\sigma_2^2 + \tau_2^2$ (multiplicity n_2) and τ_2^2 (multiplicity $n_2(V - 1)$).

As in the previous two sections, eigenvectors $\mathbf{p}_{.1n_1}$ and $\mathbf{p}_{.1n_2}$ span the two columns of \mathbf{X} , so we have the same OLS closed-form expression for $\hat{\boldsymbol{\beta}}$.

All four \mathbf{G} matrices commute with one another, and there are four distinct, linearly independent eigenvalues, so there are also closed-form expressions for the maximum likelihood estimators of the variance parameters. These are obtained from Equation 2.13 as (after rearranging the rows in a convenient fashion):

$$\begin{bmatrix} \hat{\tau}_1^2 \\ \hat{\sigma}_1^2 \\ \hat{\tau}_2^2 \\ \hat{\sigma}_2^2 \end{bmatrix} = \begin{bmatrix} n_1V & n_1V & 0 & 0 \\ n_1V & n_1V^2 & 0 & 0 \\ 0 & 0 & n_2V & n_2V \\ 0 & 0 & n_2V & n_2V^2 \end{bmatrix} \begin{bmatrix} \text{Tr}(\mathbf{G}_0\mathbf{C}) \\ \text{Tr}(\mathbf{G}_2\mathbf{C}) \\ \text{Tr}(\mathbf{G}_1\mathbf{C}) \\ \text{Tr}(\mathbf{G}_3\mathbf{C}) \end{bmatrix}. \quad (2.27)$$

Therefore, the ML covariance parameter estimates are exactly as given in Equations 2.15 and 2.14, but with all calculations performed *within-group*:

$$\hat{\sigma}_1^2 = \frac{1}{(V-1)n_1V} \sum_{i=1}^{n_1} \sum_{j \neq k}^V (Y_{ij} - \hat{\mu})(Y_{ik} - \hat{\mu}) \quad (2.28)$$

$$\hat{\tau}_1^2 = \frac{1}{n_1V} \sum_{i=1}^{n_1} \sum_{j=1}^V (Y_{ij} - \hat{\mu})^2 - \hat{\sigma}_1^2 \quad (2.29)$$

$$\hat{\sigma}_2^2 = \frac{1}{(V-1)n_2V} \sum_{i=n_1+1}^{n_2} \sum_{j \neq k}^V (Y_{ij} - \hat{\mu})(Y_{ik} - \hat{\mu}) \quad (2.30)$$

$$\hat{\tau}_2^2 = \frac{1}{n_2V} \sum_{i=n_1+1}^{n_2} \sum_{j=1}^V (Y_{ij} - \hat{\mu})^2 - \hat{\sigma}_2^2. \quad (2.31)$$

That is, this model entails complete separation of the covariance estimation by group.

2.4.4 The General Case For Heterogeneous Compound Symmetry

In Sections 2.4.1 - 2.4.3 above, I investigated the existence of closed-form maximum likelihood estimates for models with three different kinds of group heterogeneity in variance structure. However, the models I considered had extremely simple mean structures. In this section, I will present general conditions under which closed-form ML solutions exist for models with the heterogeneous variance structures considered above. The development is parallel to that in Section 2.3.5 for models with homogeneous compound symmetric variance structures.

Theorem 2.4.1. *For the models with covariance structures described in Equations 2.21, 2.23 and 2.26, a closed-form expression exists for $\hat{\beta}$ whenever the columns of \mathbf{X} consist of:*

1. *Group-specific between-subject effects: columns of the form $\mathbf{x} = \mathbf{G}_l(\mathbf{1}_{V \times 1} \otimes \boldsymbol{\alpha}_{N \times 1})$, where \mathbf{G}_l is an $NV \times NV$ diagonal matrix with 1s for group l and 0s for the other group.*

2a. *For Models 2.21 and 2.26, group-specific within-subject spatial contrasts (columns $\mathbf{x} = \mathbf{G}_l[x_{11} \ x_{12} \ \dots \ x_{1V} \ x_{21} \ \dots \ x_{NV}]'$ with the property that $\sum_{j=1}^V x_{ij} = 0 \ \forall i$, where \mathbf{G} is as in part 1 above.*

2b. *For Model 2.23, within-subject spatial contrasts that are not necessarily group specific (columns $\mathbf{x}_1 = [x_{11} \ x_{12} \ \dots \ x_{1V} \ x_{21} \ \dots \ x_{NV}]'$ with the property that $\sum_{j=1}^V x_{ij} = 0 \ \forall i$).*

3. *Group-specific within-subject spatial effects that are the same within groups of subjects defined by between-subject covariates: columns of the form $\mathbf{x} = \mathbf{G}_l(\mathbf{x}_{\alpha_{N \times 1}} \otimes \boldsymbol{\gamma}_{V \times 1})$ for arbitrary $\boldsymbol{\gamma}$, where \mathbf{x}_{α} is a between-subject covariate vector consisting of ones and zeros and \mathbf{G} is as in part 1 above.*

For Models 2.23 and 2.26, a closed-form solution exists for the ML estimates of the covariance parameters whenever a closed-form solution exists for $\hat{\boldsymbol{\beta}}$. A closed-form solution never exists for the ML estimates of the covariance parameters for Model 2.21.

Proof. The proof of this theorem is closely analogous to that of Theorem 2.3.1, but the details of the spectral decomposition need to be treated separately for the three cases of covariance structure.

Case 1: Heterogeneous Residual Variance

The spectral decomposition of the covariance matrix for Model 2.21 partitions \mathbb{R}^{NV} into four orthogonal subspaces:

1. U_1 of rank n_1 , spanned by eigenvectors corresponding to the eigenvalue $V\sigma^2 + \tau_1^2$.
2. U_2 of rank $n_1(V - 1)$, spanned by eigenvectors corresponding to the eigenvalue τ_1^2 .
3. U_3 of rank n_2 , spanned by eigenvectors corresponding to the eigenvalue $V\sigma^2 + \tau_2^2$.
4. U_4 of rank $n_2(V - 1)$, spanned by eigenvectors corresponding to the eigenvalue τ_2^2 .

Columns of type 1 above are in U_1 for $l = 1$ (group 1) and in U_3 for $l = 2$ (group 2).

Columns of type 2a above are in U_2 for $l = 1$ and U_4 for $l = 2$.

Columns of type 3 above can be decomposed by centering into a multiple of a column of type 1 already accounted for in the model, and a column in U_2 for $l = 1$ and U_4 for $l = 2$.

Case 2: Heterogeneous Between-Subjects Variance

The spectral decomposition of the covariance matrix for Model 2.23 partitions \mathbb{R}^{NV} into three orthogonal subspaces:

1. U_1 of rank n_1 , spanned by eigenvectors corresponding to the eigenvalue $V\sigma_1^2 + \tau^2$.
2. U_2 of rank n_2 , spanned by eigenvectors corresponding to the eigenvalue $V\sigma_2^2 + \tau^2$.
3. U_3 of rank $N(V - 1)$, spanned by eigenvectors corresponding to the eigenvalue τ^2 .

Columns of type 1 above are in U_1 for $l = 1$ (group 1) and in U_2 for $l = 2$ (group 2).

Columns of type 2b above are in U_3 .

Columns of type 3 above can be decomposed by centering into a multiple of a column of type 1 already accounted for in the model, and a column in U_3 .

Case 3: Heterogeneous Residual And Between-Subjects Variance

The spectral decomposition of the covariance matrix for Model 2.26 partitions \mathbb{R}^{NV} into four orthogonal subspaces:

1. U_1 of rank n_1 , spanned by eigenvectors corresponding to the eigenvalue $V\sigma_1^2 + \tau_1^2$.
2. U_2 of rank $n_1(V - 1)$, spanned by eigenvectors corresponding to the eigenvalue τ_1^2 .
3. U_3 of rank n_2 , spanned by eigenvectors corresponding to the eigenvalue $V\sigma_2^2 + \tau_2^2$.
4. U_4 of rank $n_2(V - 1)$, spanned by eigenvectors corresponding to the eigenvalue τ_2^2 .

Columns of type 1 above are in U_1 for $l = 1$ (group 1) and in U_3 for $l = 2$ (group 2).

Columns of type 2a above are in U_2 for $l = 1$ and U_4 for $l = 2$.

Columns of type 3 above can be decomposed by centering into a multiple of a column of type 1 already accounted for in the model, and a column in U_2 for $l = 1$ and U_4 for $l = 2$.

□

There are several interesting conclusions that follow immediately from this result. First, note that the list of permissible columns for closed-form estimation given in Theorem 2.4.1 is not exhaustive. Whenever $\mathbf{X}_{NV \times p}$ is spanned by p eigenvectors of $\mathbf{\Sigma}$, then so is $\mathbf{X}\mathbf{A}_{p \times p}$ for any invertible \mathbf{A} . In particular, this means that pairs of group-specific effects \mathbf{x}_1 and \mathbf{x}_2 may be reparameterized into an overall effect, $\tilde{\mathbf{x}}_1$ and a group specific deviation from that effect, \mathbf{x}_2 . For a model that also contains a group-specific intercept, this means that any group \times effect interactions are permissible for closed-form estimation. I summarize this as a corollary:

Corollary 2.4.2. *For the models described in Equations 2.21, 2.23 and 2.26, a closed-form expression exists for $\hat{\boldsymbol{\beta}}$ whenever the fixed-effects structure, \mathbf{X} , consists of group \times effect interactions.*

Second, the important special case of models with a voxelwise mean structure of the form

$$E(\mathbf{Y}_{NV \times 1}) = \left(\begin{bmatrix} \mathbf{1}_{n_1 \times 1} & \mathbf{0}_{n_1 \times 1} \\ \mathbf{1}_{n_2 \times 1} & \mathbf{1}_{n_2 \times 1} \end{bmatrix} \otimes \mathbf{I}_{V \times V} \right) \boldsymbol{\beta}_{2V \times 1},$$

does have a closed-form solution for $\hat{\boldsymbol{\beta}}$, even though it doesn't fall explicitly into the criteria outlined in Theorem 2.4.1 (which are sufficient but not necessary for the existence of closed-form estimates). To see why this is, let

$$\mathbf{X}_{NV \times V} = \left(\begin{bmatrix} \mathbf{1}_{n_1 \times 1} & \mathbf{0}_{n_1 \times 1} \\ \mathbf{1}_{n_2 \times 1} & \mathbf{1}_{n_2 \times 1} \end{bmatrix} \otimes \mathbf{I}_{V \times V} \right),$$

and

$$\tilde{\mathbf{X}}_{NV \times (V+2)} = \begin{bmatrix} \mathbf{1}_{NV \times 1} & \begin{bmatrix} \mathbf{0}_{n_1 V \times 1} \\ \mathbf{1}_{n_2 V \times 1} \end{bmatrix} & \mathbf{X} \end{bmatrix}.$$

By Theorem 2.4.1, a closed-form solution exists for the model with $\tilde{\mathbf{X}}$ as the fixed-effects structure (as the first two columns are of type 1 and the remaining V columns are of type 3 given the first two). However, since the first two columns of $\tilde{\mathbf{X}}$ are in the span of \mathbf{X} itself, it is not necessary to actually add them to the design in order to include two more eigenvectors in the spanning set.

Therefore there are V eigenvectors of $\boldsymbol{\Sigma}$ that span \mathbf{X} , and a closed-form solution exists for the model with fixed-effects structure given by \mathbf{X} . These may be exhibited explicitly in a development exactly parallel to the generalization of the scalar mean model in Section 2.3.2 to the voxelwise mean model in Section 2.3.4. This is summarized in the following corollary:

Corollary 2.4.3. *For the models described in Equations 2.21, 2.23 and 2.26, a closed-form expression exists for $\hat{\boldsymbol{\beta}}$ whenever the fixed-effects structure, \mathbf{X} , represents group-specific voxelwise means, as well as any other permissible effect for closed-form estimation as outlined in Theorem 2.4.1 and Corollary 2.4.2.*

2.4.5 Nonexistence Of Closed-Form Estimators And Implications

Of critical interest is what the preceding results imply about circumstances under which a closed-form solution is *not* available. A closed-form solution for $\hat{\beta}$ is unavailable whenever \mathbf{X} contains a column that represents an effect that is homogeneous across groups (with one exception: homogeneous spatial contrasts are estimable in closed-form for Model 2.23).

This has important implications for model-building and hypothesis testing for DTI data. Even in situations in which it is theoretically appropriate to model a fixed effect consistently across groups (for instance, a between-subject covariate that pertains to the external scanning circumstances and is independent of group membership), we are constrained to include interaction terms in order to retain closed-form estimation in heterogeneous variance models. Including unnecessary (from the standpoint of the subject matter) fixed effects terms entails inflating the variance of all other parameters in the model.

Thus, the gain in computational efficiency granted by closed-form estimation is offset by a loss in statistical power for the heterogeneous variance models discussed in this section. Whether the cost in power is worth the added efficiency will be a highly context-dependent question. For large N , small p studies, the loss in power of adding a small number of unwanted interaction terms may be negligible and well worth the gain in computational efficiency of the closed-form ML estimation. On the other hand, for studies with small sample sizes and/or a proliferation of important covariates, the loss in power may be substantial and the more practical choice be to use iterative algorithms or to assume that variance is homogeneous across groups.

2.4.6 Three Or More Groups

As a final note in this Section, all of the preceding results generalize in obvious ways to the case where participants come from more than two groups. This includes the formulas for closed-form covariance parameter estimates. For instance, for the case of three groups with

heterogeneous between-subject variance, Equation 2.24 generalizes to:

$$\begin{bmatrix} \hat{\tau}^2 \\ \hat{\sigma}_1^2 \\ \hat{\sigma}_2^2 \\ \hat{\sigma}_3^2 \end{bmatrix} = \begin{bmatrix} NV & n_1V & n_2V & n_3V \\ n_1V & n_1V^2 & 0 & 0 \\ n_2V & 0 & n_2V^2 & 0 \\ n_3V & 0 & 0 & n_3V^2 \end{bmatrix}^{-1} \begin{bmatrix} \text{Tr}(\mathbf{G}_0\mathbf{C}) \\ \text{Tr}(\mathbf{G}_1\mathbf{C}) \\ \text{Tr}(\mathbf{G}_2\mathbf{C}) \\ \text{Tr}(\mathbf{G}_3\mathbf{C}) \end{bmatrix}$$

$$\propto \begin{bmatrix} V & -1 & -1 & -1 \\ -1 & \frac{Vn_1+(V-1)(n_2+n_3)}{Vn_1} & \frac{1}{V} & \frac{1}{V} \\ -1 & \frac{1}{V} & \frac{Vn_2+(V-1)(n_1+n_3)}{Vn_2} & \frac{1}{V} \\ -1 & \frac{1}{V} & \frac{1}{V} & \frac{Vn_3+(V-1)(n_1+n_2)}{Vn_3} \end{bmatrix} \begin{bmatrix} \text{Tr}(\mathbf{G}_0\mathbf{C}) \\ \text{Tr}(\mathbf{G}_1\mathbf{C}) \\ \text{Tr}(\mathbf{G}_2\mathbf{C}) \\ \text{Tr}(\mathbf{G}_3\mathbf{C}) \end{bmatrix}$$

2.4.7 Summary

In this section, I have shown that closed-form ML estimators exist for the mean parameters in random intercept models with heterogeneous residual variance, between-subjects variance, or both, subject to the constraint that covariates must have group-specific effects, with the exception that non-group specific within-subject spatial contrasts are estimable in closed-form in the model with heterogeneous between-subjects variance and homogeneous residual variance. As such, closed-form estimation may entail a cost in statistical power if it constrains the data analyst to include interaction effects that would otherwise have been left out.

In this class of heterogeneous variance models, closed-form ML estimators only exist for the variance parameters in models with heterogeneous between-subjects variance (either with or without additional heterogeneity in residual variance). However, for models with heterogeneous residual variance but homogeneous between-subjects variance, only iterative estimation of the variance parameters is possible.

The next class of important models to consider for the purposes of identifying the existence of closed-form ML estimators are longitudinal models that will allow us to track changes in DTI parameters over time, either within a single group, or differentially between multiple groups.

2.5 CLOSED-FORM ML ESTIMATION FOR LONGITUDINAL MODELS

Longitudinal models for DTI data are of great potential importance for tracing the development of healthy and abnormal white matter anatomy over time. In addition, longitudinal DTI data share a common structure in many respect with fMRI data, so many of the results in this section apply equally to the analysis of fMRI studies. In this section, I will be applying the same principles introduced in Sections 2.3 and 2.4 to LME models for longitudinal data. I will explore longitudinal models for a single group, but the results will extend naturally to the case of two or more groups.

2.5.1 Longitudinal Model With Random Intercept

Suppose, as before, we have V voxels for each of N subjects with data coregistered to a common anatomical space, but that now each subject is also scanned on T occasions. Let Y_{ijk} denote the fixed transformed diffusion tensor element for subject i , voxel j , timepoint k . We will stack the observations across all subjects, voxels and timepoints into a single vector in which the timepoint index will “move fastest”, the subject index will “move slowest.” That is, $\mathbf{Y}_{ij} = [Y_{ij1} \dots Y_{ijT}]'$, $\mathbf{Y}_i = [\mathbf{Y}_{i1} \dots \mathbf{Y}_{iV}]'$, and $\mathbf{Y} = [\mathbf{Y}_1 \dots \mathbf{Y}_N]'$.

I will start with the simplest case, with an overall scalar mean fit across subjects, voxels and timepoints and a linear trend in the mean fit across timepoints, along with a subject-specific random offset from the overall mean. That is, let

$$\mathbf{Y}_{NVT \times 1} = \mathbf{X}_{NVT \times 2} \boldsymbol{\beta}_{2 \times 1} + (\mathbf{I}_{N \times N} \otimes \mathbf{1}_{V \times 1} \otimes \mathbf{1}_{T \times 1}) \boldsymbol{\delta}_{N \times 1} + \boldsymbol{\epsilon}_{NVT \times 1}, \quad (2.32)$$

where $\mathbf{X} = [\mathbf{x}_1 \mathbf{x}_2]'$ is the fixed effects design matrix given by $\mathbf{x}_1 = \mathbf{1}_{NVT \times 1}$ and

$$\mathbf{x}_2 = \underbrace{[1 \ 2 \ \dots \ T \ 1 \ \dots \ T]'}_{NV \text{ times}} - \underbrace{\left[\frac{T+1}{2} \ \dots \ \frac{T+1}{2} \right]'}_{NVT \text{ times}}.$$

$\boldsymbol{\beta} = [\beta_1 \ \beta_2]'$ is the vector of unknown fixed parameters, and $\boldsymbol{\delta} \sim \text{MVN}(\mathbf{0}, \sigma^2 \mathbf{I})$ and $\boldsymbol{\epsilon} \sim \text{MVN}(\mathbf{0}, \tau^2 \mathbf{I})$ as in Section 2.3. Thus, Σ is compound symmetric as in Equation 2.10.

This model is therefore a straightforward generalization of the models considered in Theorem 2.3.1, with the fixed slope effect serving as an extension of the class of within-subject contrast effects. It is instructive, however, to address closed-form estimation in this model specifically, as a preliminary step to generalizing further to random slope models and models with orthogonal polynomial random effects.

In order to determine whether there is a closed-form expression for $\hat{\boldsymbol{\beta}}$, we need to find two eigenvectors of $\boldsymbol{\Sigma}$ that span the two columns of \mathbf{X} . We already know that \mathbf{x}_1 is itself an eigenvector of $\boldsymbol{\Sigma}$. The design matrix column representing the (centered) linear effect across time is also an eigenvector of $\boldsymbol{\Sigma}$.

To demonstrate this, we note first that the off-diagonal blocks representing within-subject, between-voxel covariance are all equal to $\sigma^2 \mathbf{J}_{T \times T}$. The contribution of each of these blocks to $\boldsymbol{\Sigma} \mathbf{x}_2$ is:

$$\begin{aligned} \sigma^2 \mathbf{J}_{T \times T} \begin{bmatrix} (1 - \frac{T+1}{2}) \\ \vdots \\ (T - \frac{T+1}{2}) \end{bmatrix} &= \sigma^2 \begin{bmatrix} (1 + \dots + T) - (T \frac{T+1}{2}) \\ \vdots \\ (1 + \dots + T) - (T \frac{T+1}{2}) \end{bmatrix} \\ &= \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}. \end{aligned}$$

Thus, only the voxel-specific diagonal blocks, equal to $\tau^2 \mathbf{I}_{T \times T} + \sigma^2 \mathbf{J}_{T \times T}$, possibly contribute non-zero entries to the product $\boldsymbol{\Sigma} \mathbf{x}_2$. Specifically, the $kT + 1$ through $(k + 1)T$ entries of $\boldsymbol{\Sigma} \mathbf{x}_2$, for $k = 0, \dots, (NV - 1)$ are given by:

$$\begin{aligned} (\tau^2 \mathbf{I}_{T \times T} + \sigma^2 \mathbf{J}_{T \times T}) \begin{bmatrix} (1 - \frac{T+1}{2}) \\ \vdots \\ (T - \frac{T+1}{2}) \end{bmatrix} &= \tau^2 \mathbf{I}_{T \times T} \begin{bmatrix} (1 - \frac{T+1}{2}) \\ \vdots \\ (T - \frac{T+1}{2}) \end{bmatrix} \\ &= \tau^2 \begin{bmatrix} (1 - \frac{T+1}{2}) \\ \vdots \\ (T - \frac{T+1}{2}) \end{bmatrix}. \end{aligned}$$

which is a multiple of the $kT + 1$ through $(k + 1)T$ entries of \mathbf{x}_2 . Thus \mathbf{x}_2 is an eigenvector of Σ .

The calculations above are made clear with a numerical example. Let $N = 1$, $V = 2$, $T = 3$. Then:

$$\begin{aligned}
\Sigma \mathbf{x}_2 &= \begin{bmatrix} \tau^2 + \sigma^2 & \sigma^2 & \sigma^2 & \sigma^2 & \sigma^2 & \sigma^2 \\ \sigma^2 & \tau^2 + \sigma^2 & \sigma^2 & \sigma^2 & \sigma^2 & \sigma^2 \\ \sigma^2 & \sigma^2 & \tau^2 + \sigma^2 & \sigma^2 & \sigma^2 & \sigma^2 \\ \sigma^2 & \sigma^2 & \sigma^2 & \tau^2 + \sigma^2 & \sigma^2 & \sigma^2 \\ \sigma^2 & \sigma^2 & \sigma^2 & \sigma^2 & \tau^2 + \sigma^2 & \sigma^2 \\ \sigma^2 & \sigma^2 & \sigma^2 & \sigma^2 & \sigma^2 & \tau^2 + \sigma^2 \end{bmatrix} \begin{bmatrix} -1 \\ 0 \\ 1 \\ -1 \\ 0 \\ 1 \end{bmatrix} \\
&= \begin{bmatrix} -(\tau^2 + \sigma^2) + \sigma^2 - \sigma^2 + \sigma^2 \\ \sigma^2 - \sigma^2 + \sigma^2 - \sigma^2 \\ -\sigma^2 + (\tau^2 + \sigma^2) - \sigma^2 + \sigma^2 \\ -(\tau^2 + \sigma^2) + \sigma^2 - \sigma^2 + \sigma^2 \\ \sigma^2 - \sigma^2 + \sigma^2 - \sigma^2 \\ -\sigma^2 + (\tau^2 + \sigma^2) - \sigma^2 + \sigma^2 \end{bmatrix} \\
&= \begin{bmatrix} -\tau^2 \\ 0 \\ \tau^2 \\ -\tau^2 \\ 0 \\ \tau^2 \end{bmatrix} \\
&= \tau^2 \mathbf{x}_2.
\end{aligned}$$

Therefore, there is a closed-form expression for $\hat{\boldsymbol{\beta}}$, the OLS solution:

$$\begin{aligned}
\hat{\boldsymbol{\beta}} &= (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y} \\
&= \begin{bmatrix} NVT & 0 \\ 0 & NV\left(\frac{T^3-T}{12}\right) \end{bmatrix}^{-1} \mathbf{X}'\mathbf{Y} \\
&= \frac{1}{NV} \begin{bmatrix} \frac{\mathbf{x}'_1}{T} \\ \frac{12\mathbf{x}'_2}{T^3-T} \end{bmatrix} \mathbf{Y}.
\end{aligned} \tag{2.33}$$

And there are closed-form solutions for $\hat{\sigma}^2$ and $\hat{\tau}^2$, which are analogous to those given in Section 2.3:

$$\hat{\sigma}^2 = \frac{1}{(VT-1)NVT} \sum_{i=1}^N \sum_{(j,k) \neq (j',k')} (Y_{ijk} - \hat{\mu})(Y_{ij'k'} - \hat{\mu}), \tag{2.34}$$

the mean of the within-subject off-diagonal elements of \mathbf{C} , and

$$\hat{\tau}^2 = \frac{1}{NVT} \sum_{i=1}^N \sum_{j=1}^V \sum_{k=1}^T (Y_{ijk} - \hat{\mu})^2 - \hat{\sigma}^2. \tag{2.35}$$

2.5.2 Longitudinal Model With Random Slope

One level of added complexity from the longitudinal model with a subject-specific random intercept is the inclusion of a subject-specific random slope as well. The next step is to determine whether such an extension of the model in Section 2.5.1 will also admit to closed-form expressions for the ML estimators of the mean and covariance parameters.

The model is:

$$\begin{aligned}
\mathbf{Y}_{NVT \times 1} &= \mathbf{X}_{NVT \times 2} \boldsymbol{\beta}_{2 \times 1} + (\mathbf{I}_{N \times N} \otimes \mathbf{1}_{V \times 1} \otimes \mathbf{z}_{1T \times 1}) \boldsymbol{\delta}_{1N \times 1} \\
&\quad + (\mathbf{I}_{N \times N} \otimes \mathbf{1}_{V \times 1} \otimes \mathbf{z}_{2T \times 1}) \boldsymbol{\delta}_{2N \times 1} + \boldsymbol{\epsilon}_{NVT \times 1},
\end{aligned} \tag{2.36}$$

where $\mathbf{X} = [\mathbf{x}_1 \ \mathbf{x}_2]'$ is the fixed effects design matrix as defined in Section 2.5.1 and \mathbf{z}_1 and \mathbf{z}_2 are comprised of the first T entries of \mathbf{x}_1 and \mathbf{x}_2 , respectively.

As before, $\boldsymbol{\epsilon} \sim N(\mathbf{0}, \tau^2 \mathbf{I})$. $\boldsymbol{\delta}_1 \sim N(\mathbf{0}, \sigma_1^2)$ and $\boldsymbol{\delta}_2 \sim N(\mathbf{0}, \sigma_2^2)$. The covariance of \mathbf{Y} , $\boldsymbol{\Sigma}$ has the additive structure introduced in Section 2.2: $\boldsymbol{\Sigma} = \tau^2 \mathbf{G}_0 + \sigma_1^2 \mathbf{G}_1 + \sigma_2^2 \mathbf{G}_2$, where $\mathbf{G}_0 = \mathbf{I}$, $\mathbf{G}_1 = \mathbf{J}$ and

$$\mathbf{G}_2 = (\mathbf{I}_{N \times N} \otimes \mathbf{1}_{V \times 1} \otimes \mathbf{z}_{2V \times 1})(\mathbf{I}_{N \times N} \otimes \mathbf{1}_{V \times 1} \otimes \mathbf{z}_{2T \times 1})' \quad (2.37)$$

$$= \mathbf{I}_{N \times N} \otimes \mathbf{1}_{V \times 1} \otimes (\mathbf{z}_{2T \times 1} \mathbf{z}_{2T \times 1}'). \quad (2.38)$$

This is a block diagonal matrix with N blocks, each block equal to

$$\mathbf{J}_{V \times V} \otimes \left(\begin{bmatrix} 1 - \frac{T+1}{2} \\ \vdots \\ T - \frac{T+1}{2} \end{bmatrix} \begin{bmatrix} 1 - \frac{T+1}{2} & \dots & T - \frac{T+1}{2} \end{bmatrix} \right).$$

This latter matrix has as its $((kT + i), (lT + j))^{\text{th}}$ entry $(i - \frac{T+1}{2})(j - \frac{T+1}{2})$, for $0 \leq k, l < NV$ and $1 \leq i, j \leq T$. To take a simple numerical example, let $T = 3$ and $V = 2$. Then each of the N diagonal blocks of \mathbf{G}_2 are given by:

$$\begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \otimes \begin{bmatrix} 1 \\ 0 \\ -1 \end{bmatrix} \begin{bmatrix} 1 & 0 & -1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & -1 & 1 & 0 & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ -1 & 0 & 1 & -1 & 0 & 1 \\ 1 & 0 & -1 & 1 & 0 & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ -1 & 0 & 1 & -1 & 0 & 1 \end{bmatrix}.$$

And each of the N diagonal blocks of $\boldsymbol{\Sigma}$ itself are:

$$\begin{bmatrix} \tau^2 + \sigma_1^2 + \sigma_2^2 & \sigma_1^2 & \sigma_1^2 - \sigma_2^2 & \sigma_1^2 + \sigma_2^2 & \sigma_1^2 & \sigma_1^2 - \sigma_2^2 \\ \sigma_1^2 & \tau^2 + \sigma_1^2 & \sigma_1^2 & \sigma_1^2 & \sigma_1^2 & \sigma_1^2 \\ \sigma_1^2 - \sigma_2^2 & \sigma_1^2 & \tau^2 + \sigma_1^2 + \sigma_2^2 & \sigma_1^2 - \sigma_2^2 & \sigma_1^2 & \sigma_1^2 + \sigma_2^2 \\ \sigma_1^2 + \sigma_2^2 & \sigma_1^2 & \sigma_1^2 - \sigma_2^2 & \tau^2 + \sigma_1^2 + \sigma_2^2 & \sigma_1^2 & \sigma_1^2 - \sigma_2^2 \\ \sigma_1^2 & \sigma_1^2 & \sigma_1^2 & \sigma_1^2 & \tau^2 + \sigma_1^2 & \sigma_1^2 \\ \sigma_1^2 - \sigma_2^2 & \sigma_1^2 & \sigma_1^2 + \sigma_2^2 & \sigma_1^2 - \sigma_2^2 & \sigma_1^2 & \tau^2 + \sigma_1^2 + \sigma_2^2 \end{bmatrix}.$$

To satisfy the conditions of Theorem 2.2.1 for the existence of a closed-form ML solution for the mean parameters, $\boldsymbol{\beta}$, it suffices to show that \mathbf{x}_1 and \mathbf{x}_2 are themselves eigenvectors of $\boldsymbol{\Sigma}$. First, \mathbf{x}_1 :

$$\begin{aligned}\boldsymbol{\Sigma}\mathbf{x}_1 &= \tau^2\mathbf{G}_0\mathbf{1} + \sigma_1^2\mathbf{G}_1\mathbf{1} + \sigma_2^2\mathbf{G}_2\mathbf{1} \\ &= \tau^2\mathbf{1} + \sigma_1^2VT\mathbf{1} + \sigma_2^2\mathbf{G}_2\mathbf{1} \\ &= \tau^2\mathbf{1} + \sigma_1^2VT\mathbf{1} \\ &\propto \mathbf{x}_1,\end{aligned}$$

because the $(kT + i)^{\text{th}}$ entry of $\sigma_2^2\mathbf{G}_2\mathbf{1}$ is

$$\begin{aligned}\sigma_2^2V\left(i - \frac{T+1}{2}\right)\sum_{j=1}^T\left(j - \frac{T+1}{2}\right) &= \sigma_2^2V\left(i - \frac{T+1}{2}\right)\left(\left(\sum_{j=1}^T j\right) - T\frac{T+1}{2}\right) \\ &= 0\end{aligned}$$

Thus, \mathbf{x}_1 is an eigenvector of $\boldsymbol{\Sigma}$ with eigenvalue $\tau^2 + VT\sigma_1^2$. For \mathbf{x}_2 , we have:

$$\begin{aligned}\boldsymbol{\Sigma}\mathbf{x}_2 &= \tau^2\mathbf{G}_0\mathbf{x}_2 + \sigma_1^2\mathbf{G}_1\mathbf{x}_2 + \sigma_2^2\mathbf{G}_2\mathbf{x}_2 \\ &= \tau^2\mathbf{x}_2 + V\sigma_1^2\sum_{i=1}^T\left(i - \frac{T+1}{2}\right)\mathbf{1} + \sigma_2^2(\mathbf{x}_2\mathbf{x}_2')\mathbf{x}_2 \\ &= \tau^2\mathbf{x}_2 + \sigma_2^2\mathbf{x}_2(\mathbf{x}_2'\mathbf{x}_2) \\ &= \tau^2\mathbf{x}_2 + V\sigma_2^2\sum_{i=1}^T\left(i - \frac{T+1}{2}\right)^2\mathbf{x}_2 \\ &\propto \mathbf{x}_2\end{aligned}$$

Thus, \mathbf{x}_2 is also an eigenvector of $\boldsymbol{\Sigma}$, with eigenvalue $\tau^2 + V\sigma_2^2\sum_{i=1}^T\left(i - \frac{T+1}{2}\right)^2$. Therefore, the OLS estimate of $\boldsymbol{\beta}$ (given in Equation 2.33) is also the ML estimate.

Next, to establish the existence of closed-form solutions for τ^2 , σ_1^2 and σ_2^2 , we will first demonstrate that \mathbf{G}_0 , \mathbf{G}_1 , and \mathbf{G}_2 are simultaneously diagonalizable. This is clear: $\mathbf{G}_0 = \mathbf{I}$ commutes with every conformable matrix, and $\mathbf{G}_1\mathbf{G}_2 = \mathbf{G}_2\mathbf{G}_1 = \mathbf{0}$ (as the sum of any row or column of \mathbf{G}_2 is 0), so all three G matrices commute with each other, and are therefore simultaneously diagonalizable.

The diagonal entries of $\boldsymbol{\Sigma}$ in its diagonalized form are its eigenvalues. Two of these were exhibited above: $\tau^2 + VT\sigma_1^2$ corresponding to the N eigenvectors of \mathbf{G}_1 comprised of blocks

of \mathbf{x}_1 of length VT and $\tau^2 + V\sigma_2^2 \sum_{i=1}^T (i - \frac{T+1}{2})^2$ corresponding to the N eigenvectors of \mathbf{G}_2 comprised of blocks of \mathbf{x}_2 of length VT .

The remaining $(N-2)VT$ eigenvalues will all be equal to τ^2 . To see why this is, note that $\mathbf{\Sigma} = \mathbf{G}_0 + \mathbf{G}_1 + \mathbf{G}_2$ is of full rank NVT . \mathbf{G}_1 and \mathbf{G}_2 are each of rank N , and their column spaces are orthogonal, and spanned by linearly independent vectors comprised of blocks of \mathbf{x}_1 and \mathbf{x}_2 , respectively. The intersection of their null spaces is therefore of rank $(N-2)VT$. Therefore, the set comprised of \mathbf{x}_1 , \mathbf{x}_2 and any orthonormal basis for the intersection of the null spaces of \mathbf{G}_1 and \mathbf{G}_2 comprise a basis for the column space of $\mathbf{\Sigma}$. And, if \mathbf{a} is in the intersection of the null spaces of \mathbf{G}_1 and \mathbf{G}_2 then $\mathbf{\Sigma}\mathbf{a} = \tau^2\mathbf{G}_0\mathbf{a} + \sigma_1^2\mathbf{G}_1\mathbf{a} + \sigma_2^2\mathbf{G}_2\mathbf{a} = \tau^2\mathbf{a}$.

Therefore there are three linearly independent entries on the diagonal of $\mathbf{\Sigma}$ in its diagonalized form: $\tau^2 + VT\sigma_1^2$, $\tau^2 + V\sigma_2^2 \sum_{i=1}^T (i - \frac{T+1}{2})^2$ and τ^2 , and so there are closed-form maximum likelihood estimators for τ^2 , σ_1^2 and σ_2^2 , which can be derived from Equation 2.13, based on the following quantities:

$$\begin{aligned} \text{Tr}(\mathbf{G}_0\mathbf{G}_0) &= \text{Tr}(\mathbf{G}_0\mathbf{G}_1) = \text{Tr}(\mathbf{G}_1\mathbf{G}_0) = NVT \\ \text{Tr}(\mathbf{G}_1\mathbf{G}_1) &= N(VT)^2 \\ \text{Tr}(\mathbf{G}_0\mathbf{G}_2) &= \text{Tr}(\mathbf{G}_2\mathbf{G}_0) = \text{Tr}(\mathbf{G}_1\mathbf{G}_2) = \text{Tr}(\mathbf{G}_2\mathbf{G}_1) = 0 \\ \text{Tr}(\mathbf{G}_2\mathbf{G}_2) &= \text{Tr}([\mathbf{I}_{n \times n} \otimes (\mathbf{z}_2\mathbf{z}'_2)][\mathbf{I}_{n \times n} \otimes (\mathbf{z}_2\mathbf{z}'_2)]) \\ &= \text{Tr}(\mathbf{I} \otimes (\mathbf{z}_2\mathbf{z}'_2\mathbf{z}_2\mathbf{z}'_2)) \\ &= NV \left(\frac{T^3 - T}{12} \right) \text{Tr}(\mathbf{z}_2\mathbf{z}'_2) \\ &= NV \left(\frac{T^3 - T}{12} \right)^2 \end{aligned}$$

Thus the ML estimates of the covariance parameters are given by:

$$\begin{aligned} \begin{bmatrix} \hat{\tau}^2 \\ \hat{\sigma}_1^2 \\ \hat{\sigma}_2^2 \end{bmatrix} &= \begin{bmatrix} NVT & NVT & 0 \\ NVT & N(VT)^2 & 0 \\ 0 & 0 & NV \left(\frac{T^3 - T}{12} \right)^2 \end{bmatrix}^{-1} \begin{bmatrix} \text{Tr}(\mathbf{G}_0\mathbf{C}) \\ \text{Tr}(\mathbf{G}_1\mathbf{C}) \\ \text{Tr}(\mathbf{G}_2\mathbf{C}) \end{bmatrix} \\ &= \begin{bmatrix} \frac{1}{N(VT-1)} & \frac{-1}{NVT(VT-1)} & 0 \\ \frac{-1}{NVT(VT-1)} & \frac{1}{NVT(VT-1)} & 0 \\ 0 & 0 & \left(NV^2 \left(\frac{T^3 - T}{12} \right)^2 \right)^{-1} \end{bmatrix} \begin{bmatrix} \text{Tr}(\mathbf{G}_0\mathbf{C}) \\ \text{Tr}(\mathbf{G}_1\mathbf{C}) \\ \text{Tr}(\mathbf{G}_2\mathbf{C}) \end{bmatrix} \end{aligned} \quad (2.39)$$

2.5.3 General Orthogonal Polynomial Random Effects Models

In the previous two sections, I demonstrated the existence of closed-form ML estimators for spatio-temporal neuroimaging models with voxelwise means and linear slope effects over time. In Section 2.5.1, I considered the model with a random intercept only, whereas in Section 2.5.2, I considered the model with random intercept and slope effects. In practice, these are the most important cases, since linear longitudinal effects will often be of greatest interest to the researcher. However, it would be useful to have results on closed-form estimation for models with higher-order polynomial fixed and/or random effects.

Fortunately, the previous results generalize in a straightforward manner for orthogonal polynomial designs. Consider the extension of model 2.36 above to:

$$\begin{aligned} \mathbf{Y}_{NVT \times 1} &= \mathbf{X}_{NVT \times p} \boldsymbol{\beta}_{p \times 1} + (\mathbf{I}_{N \times N} \otimes \mathbf{1}_{V \times 1} \otimes \mathbf{z}_{1T \times 1}) \boldsymbol{\delta}_{1N \times 1} \\ &\quad + \dots + (\mathbf{I}_{N \times N} \otimes \mathbf{1}_{V \times 1} \otimes \mathbf{z}_{rT \times 1}) \boldsymbol{\delta}_{rN \times 1} + \boldsymbol{\epsilon}_{NVT \times 1}, \end{aligned} \quad (2.40)$$

where the columns of \mathbf{X} are orthogonal longitudinal effects, with $\mathbf{x}_1 = \mathbf{1}$ being an intercept effect. As before, each \mathbf{z}_i is comprised of one longitudinal segment (i.e., the first T elements) of \mathbf{x}_i . Assume $r \leq p$. The covariance matrix for this model can be decomposed as

$$\boldsymbol{\Sigma} = \tau^2 \mathbf{G}_0 + \sum_{i=1}^r \sigma_i^2 \mathbf{G}_i,$$

where $\mathbf{G}_0 = \mathbf{I}$, and

$$\mathbf{G}_i = \mathbf{I}_{N \times N} \otimes \mathbf{1}_{V \times 1} \otimes (\mathbf{z}_{iT \times 1} \mathbf{z}'_{iT \times 1})$$

for $i = 1, \dots, r$. $\boldsymbol{\Sigma}$ is of full rank NVT and its spectral decomposition partitions \mathbb{R}^{NVT} into $r + 1$ orthogonal subspaces, U_0, \dots, U_r . Each U_i for $i > 0$ is of rank N and is spanned by vectors representing corresponding within-subject longitudinal effects (i.e., each basis vector has $\mathbf{1}_{V \times 1} \otimes \mathbf{z}_i$ for one subject and zeros for other subjects). U_0 is of rank $(VT - r)N$ and is spanned by any orthonormal basis for the joint null space of the U_i , $i > 0$.

It is straightforward to show for this class of models that, whenever a closed-form expression exists for $\hat{\boldsymbol{\beta}}$, a closed-form expression also exists for $\hat{\boldsymbol{\sigma}}$. By applying Theorem 2.2.2,

it suffices to show that Σ has $r + 1$ linearly independent eigenvalues. This can be seen by noting that if \mathbf{u} is an eigenvector of \mathbf{G}_α for $\alpha > 0$, then:

$$\begin{aligned}\Sigma\mathbf{u} &= \tau^2\mathbf{G}_0\mathbf{u} + \sum_{i=1}^r \sigma_i^2\mathbf{G}_i\mathbf{u} \\ &= \tau^2\mathbf{u} + \sigma_\alpha^2\mathbf{G}_\alpha\mathbf{u},\end{aligned}$$

as \mathbf{u} is in the null space of \mathbf{G}_i for $i > 0, i \neq \alpha$. And any vector \mathbf{v} that is in the joint null space of \mathbf{G}_i for all $i > 0$ will be an eigenvector of Σ with eigenvalue τ^2 . Thus, Σ has $r + 1$ linearly independent eigenvalues, τ^2 and $\tau^2 + \sigma_i^2 V \mathbf{z}'_i \mathbf{z}_i$ for $i = 1, \dots, r$.

The question, then, is under which circumstances $\hat{\beta}$ has a closed-form estimator. When the columns of \mathbf{X} contain only orthogonal longitudinal effects, then we do have closed-form estimation, as each longitudinal effect is either in U_i for some $i > 0$, or it is in their joint null space (U_0) because of orthogonality. Either way, each orthogonal longitudinal effect is itself an eigenvector of Σ , and thus Theorem 2.2.1 applies.

For columns of \mathbf{X} beyond orthogonal longitudinal effects, the considerations of Theorem 2.3.1 apply: closed-form estimation is possible when the additional columns of \mathbf{X} are comprised of between-subject effects, within-subject spatial contrasts (or any within-subject effect that's orthogonal to the longitudinal effects in \mathbf{X} already accounted for above), or within-subject spatial effects that are specific to groups determined by between-subject effects.

Thus, closed-form estimation in the spatio-temporal models considered in this section is directly analogous to closed-form estimation in the cross-sectional spatial models considered in Section 2.3. This provides us considerable versatility in developing models for longitudinal neuroimaging data without compromising computational feasibility.

2.5.4 Multiple Groups And Variance Heterogeneity

The results in this section extend in a very straightforward manner to models with multiple groups, parallel to the development in Theorem 2.3.1. The extension of these longitudinal results to models with heterogeneous variance is complicated by the range of possibilities - in a model with r longitudinal random effects and a residual variance parameter, there are

2^{r+1} possible different combinations of heterogeneous and homogeneous variance parameters. However, it seems possible that general results in this direction would be obtainable by methods similar to those used in this chapter, and this could well be a fruitful avenue for future exploration.

2.6 CONCLUSIONS

In this chapter, I have investigated the possibility of applying LME models with closed-form likelihood estimation to the analysis of DTI data. While LMEs provide many potential advantages over the simpler approaches often used in practice for DTI group comparison, their applicability in practice may be hindered by computational burdens. Therefore, finding subclasses of LMEs with closed-form ML estimates is a potentially important step toward fostering the adoption of such models in DTI data analyses.

In Section 2.1, I described a decorrelating transformation that will allow the multivariate data contained in diffusion tensors to be treated with separate univariate analyses. I then reviewed theoretical tools described by Szatrowski (1980) for establishing the existence of closed-form ML estimation in Section 2.2.

In the simple case of LMEs with compound symmetric covariance structures (Section 2.3), I observed in Theorem 2.3.1 that closed-form ML solutions are available under general conditions on the fixed effects design. In particular, any purely between-subject and within-subject covariates can be included without requiring iterative estimation.

In Section 2.4, I expanded the class of models considered to include group heterogeneity in residual variance parameters (Section 2.4.1), between-subjects variance parameters (Section 2.4.2) or both residual and between-subjects variance parameters (Section 2.4.3). For the cases of heterogeneous between-subjects variance and of both heterogeneous between-subjects and residual variance, closed-form ML estimation is available for both mean and variance parameters under certain conditions on the fixed effects design detailed in Theorem 2.4.1. Specifically, closed-form estimation is possible in the model incorporating both heterogeneous between-subjects and residual variance if all the fixed effects are group-specific

or, equivalently, if every fixed effect in the design is supplemented with a group interaction effect. The conditions for the existence of closed-form estimation in the model with group heterogeneity in between-subjects variance only were relaxed in that within-subject spatial contrasts were not required to be group-specific.

In Section 2.5, I investigated the possibility of closed-form estimation for LMEs for longitudinal DTI data. I showed that closed-form estimation was possible for models with voxelwise mean effects and orthogonal polynomial random effects. Although longitudinal DTI studies have been infrequently seen in the literature to date, they hold great promise for investigating white matter development over time, and the results in this section indicate that it will be computationally to apply appropriate LMEs to the analyses of these data.

2.6.1 Nonexistence Of Closed-Form Estimators

Some of the most significant results in this chapter concern the identification of circumstances under which closed-form ML estimation is not possible. In the case of homogeneous compound symmetry, only the inclusion of uncentered spatially-varying covariates entails a loss of closed-form estimation for this simple covariance structure. However, spatially-varying covariates can be decomposed via centering into between- and within-subject components, both of which are estimable in closed-form. For instance, if we wanted to control for local gray matter volume in a DTI analysis assuming group homogeneity in variance parameters, it would be necessary to parameterize this as two effects in order to maintain closed-form estimation: a purely between-subject effect defined by each subject's global mean gray matter, and a purely within-subject spatial contrast effect consisting of local deviations from the mean gray matter.

In the case of group heterogeneous compound symmetric structures, I found that closed-form estimation of variance parameters was not possible in the case of heterogeneous residual variance alone. This means that practitioners who want to incorporate group heterogeneity in residual variance without resorting to iterative estimation procedures are constrained to include group heterogeneity in between-subjects variance as well, even if a homogeneity assumption is justified for between-subjects variance. Thus, for instance, if individuals from

two groups sampled from the same population are imaged on two different scanners, it would probably be desirable to allow for model flexibility in accommodating group heterogeneity in residual variance induced by different levels or quality of machine noise. In order to retain closed-form estimation in this model, however, it would be necessary also to fit separate between-subjects variance components for the two groups, even if this had little biological justification.

I also found that closed-form estimation in models with group heterogeneity in between-subjects variance or group heterogeneity in both between-subjects and residual variance requires substantial constraints on the fixed effects design. Specifically, all effects have to be group-specific or, equivalently, include group interactions, with the exception of spatial contrast effects in models that include heterogeneity only in between-subjects variance. This means that if a practitioner wanted to control for, say, age in a closed-form heterogeneous variance model, it would be necessary to include a group \times age interaction effect even if there was no reason to believe that age effects differed between the groups.

This requirement to add unnecessary effects to maintain closed-form likelihood estimation could lead to inflated estimates of the variance of fixed effects of interest, reducing the sensitivity of the models, particularly in studies with small samples or a large number of important covariates. Careful simulation studies would help determine the trade-offs between the gains in precision from accommodating group heterogeneity in variance and the losses in power from being constrained to include extraneous fixed effects, and this is an area of possible future investigation.

3.0 APPLICATIONS TO NEUROIMAGING STUDIES OF AUTISM

This chapter applies the results established in Chapter 2 to diffusion tensor imaging data analysis. I employ DTI data to compare white matter between individuals with ASDs and controls with the closed-form models identified in Chapter 2. I clarify the feasibility of these methods and discuss the incorporation of group heterogeneity in variance structure. I also compare these models with more conventional voxelwise and regionwise analyses (Section 1.3).

I first describe the sample and methods used to collect the data (Section 3.1). In Section 3.2, I perform analyses based on fractional anisotropy (FA) to compare the performance of the LMEs discussed in the previous chapter to standard methods for the most frequently analyzed dependent variable in DTI studies, FA. Then in Section 3.3, I perform group comparisons based on decorrelated whole-tensor data, as discussed in Section 2.1. The LME-based approach to these comparisons is compared to direct multivariate analogues of conventional voxelwise and regional approaches.

3.1 DATA DESCRIPTION

The data analyzed in this chapter consist of DTI scans of a sample of 80 children and young adults with a lifetime diagnosis of an ASDs and 40 typically-developing controls, all males. The data were collected under the auspices of the Utah Autism Research Program at the University of Utah (Janet Lainhart, M.D., P.I., NIH grant MH0808026). The sample and data collection methods have been described previously[2]. Basic demographic and clinical characteristics are provided below in Table 3.1. In the table, handedness refers to the score on

the Edinburgh Handedness Inventory, with scores ranging from -100 (completely left-handed) to 100 (completely right-handed)[72].

Table 3.1: Demographic and clinical characteristics of the sample

Variable	Control		ASDs	
	Mean	S.D.	Mean	S.D.
Age (years)	15.5	6.5	14.1	8.5
Handedness	70.1	35.4	66.5	50.5
Performance IQ	116.7	16.4	97.3	20.0
Verbal IQ	115.8	14.1	98.1	23.2
Full-scale IQ	118.9	15.5	96.2	22.3

As reported in the previous analysis by Alexander et al. (2007), diffusion weighted images were acquired for all participants on a Siemens Trio 3.0 Tesla Scanner[2]. The acquisition used twelve non-collinear diffusion encoding pulses with $b = 1000\text{s/mm}^2$ and a single unweighted pulse ($b = 0$), and the diffusion weighted images were corrected for head motion and eddy current distortion. The images were interpolated to obtain $2\text{mm} \times 2\text{mm} \times 2\text{mm}$ isotropic voxels. The diffusion tensor at each voxel was estimated via weighted least squares, as described in Section 1.2.2.2. The images were spatially transformed to a common anatomical space using the affine transformation option in the FLIRT utility of the FSL package. The orientation of the tensors was then recaptured by reorienting them in the new space using the procedure described by Alexander (2001)[3].

I performed all of the analyses described in the remainder of this chapter using custom code built in MATLAB[®], Version 7.4[64], unless otherwise noted.

3.2 FRACTIONAL ANISOTROPY ANALYSES

The first set of analyses I performed were intended to identify group differences in fractional anisotropy (FA, Equation 1.19) between the ASDs and control participants. FA is a measure

of the degree to which the diffusion at a given voxel is anisotropic, and is highest in areas of dense, well-organized white matter, and lowest in the fluid-filled ventricles. I chose this univariate dependent measure as it is the most commonly used scalar index in the DTI literature, although there are many other possible choices (see Section 1.2.2.3). Thus, it provides a useful benchmark with which to compare LME-based methods with standard approaches.

The midaxial slice (that is, the center of the brain as seen from above) of the mean FA image for the sample is shown in Figure 3.1. In this figure, lighter areas reflect regions where mean FA is higher, and dark areas reflect regions where mean FA is lower. For instance, the horseshoe-shaped structures at the top and bottom of the figure correspond to the most anterior (genu) and posterior (splenium) portions of the corpus callosum, respectively, where white matter is dense and highly organized. The black areas immediately below the genu correspond to the fluid-filled ventricles, and gray areas of the figure largely correspond to diffuse white matter in regions with substantial gray matter. Figure 3.2 shows the midaxial difference map for mean FA (ASDs - control). The control group has higher mean FA in the densest white matter structures, with mean differences as high as 0.12.

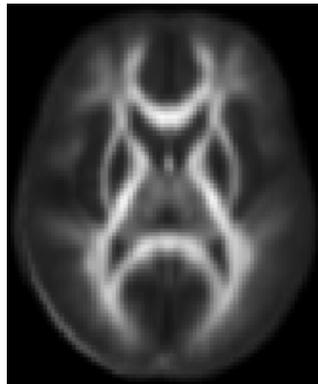


Figure 3.1: FA mean for 80 ASDs and 40 control participants

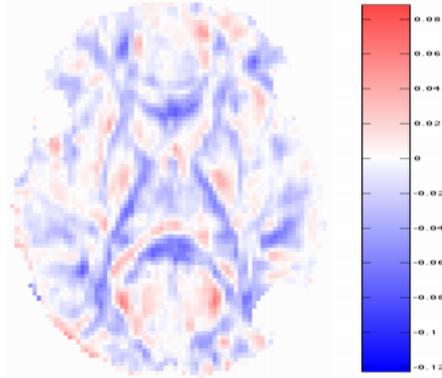


Figure 3.2: ASDs mean FA - control mean FA

3.2.1 Brain Regions Explored

In order to explore the performance of LME-based methods across a range of DTI data, I performed separate analyses on three regions of interest. These regions were chosen to be representative of different types of tissue. Also, the LMEs introduced in Chapter 2 all share a common assumption of spatial homogeneity in variance. Therefore, for demonstrative purposes I selected regions for which this assumption seemed plausible. Figure 3.3 shows the sample variance of FA at each voxel in the midaxial slice. The regions of highest variance are the most anterior and posterior portions of the corpus callosum (the genu and splenium, respectively), along with the extension of the splenium into the forceps major. Other white matter structures show middling variance, as do cortical areas containing both diffuse white matter and gray matter. Areas of primarily gray matter and the ventricles show much lower variance.

3.2.1.1 Region 1 The first region I selected is a 23×11 voxel (1012mm^2) rectangular region of the midaxial slice, symmetrically positioned around the sagittal midline, roughly corresponding to the splenium of the corpus callosum. This region represents a dense white matter structure, and an area of largely homogeneous, and high, variance. It was identified as an area of significant group difference in FA in the previous analysis by Alexander et al.

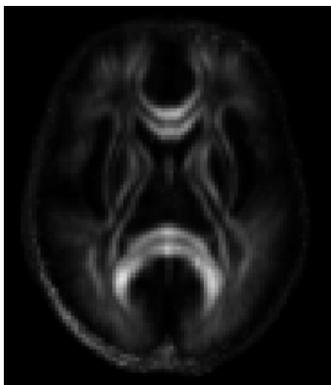


Figure 3.3: FA variance for 80 ASDs and 40 control participants

(2007)[2]. The region is outlined in Figure 3.4.

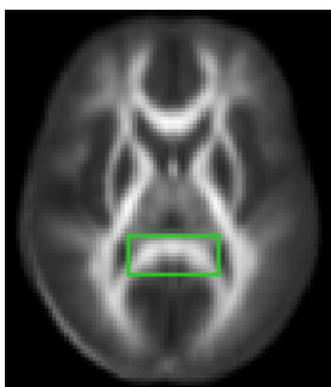


Figure 3.4: Mean FA in Region 1

3.2.1.2 Region 2 Region 2 (Figure 3.5) encompasses the genu. I again used a rectangular region of the midaxial slice, symmetric about the sagittal axis, this time encompassing 19×10 voxels, or 760mm^2 . This region was also found to display significant overall group difference in FA by Alexander et al.

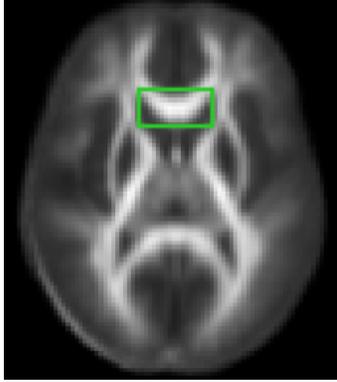


Figure 3.5: Mean FA in Region 2

3.2.1.3 Region 3 Region 3 (Figure 3.6) is a rectangular region taken from the frontal portion of a slice 4mm superior to Regions 1 and 2. This region encompasses portions of the anterior horns of the lateral ventricles, the septum pellucidum and some surrounding gray matter. There is also some compact white matter in the region: The medial anterior portion of the region overlaps the posterior portion of the genu, and the lateral posterior portions border the internal capsule. This region was chosen as a null case by which to informally judge the Type 1 error rate of the various models explored, as any white matter abnormalities detected in the fluid-filled ventricles are likely to be spurious findings. Choosing a brain region of considerable tissue heterogeneity was necessary in order to obtain a moderately-sized rectangular region with substantial amounts of CSF, and to maximize the image contrast.

3.2.2 Data Analysis

FA was calculated for each subject at each voxel in the regions of interest by diagonalizing the diffusion tensor at those voxels and substituting the observed values of λ_1 , λ_2 , and λ_3 into Equation 1.19. I then analyzed these FA data with five distinct approaches, here listed from least to greatest model complexity:

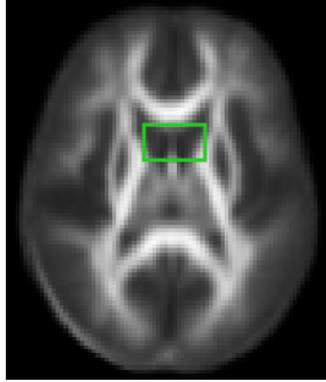


Figure 3.6: Mean FA in Region 3

1. A group comparison of regional means
2. A group comparison of voxelwise means
3. An LME with homogeneous between-subject variance and homogeneous residual variance
4. An LME with heterogeneous between-subject variance and homogeneous residual variance
5. An LME with heterogeneous between-subject variance and heterogeneous residual variance

3.2.2.1 Group Comparison Of Regional Means For the simplest analysis, I took a traditional ROI-based approach, as described in Section 1.3.1. For each subject, I first calculated the mean FA value across the regions of interest. This yielded a single scalar index for each region for each subject. I then compared the mean values for each region between ASDs and control participants with a two-sample t -test.

Let \bar{Y}_i denote the sample mean FA for subject i in a fixed region, and let C denote the set

of N_0 control participants and A denote the set of N_1 ASDs participants. I then calculated:

$$\begin{aligned}\bar{Y}_A &= \frac{1}{N_1} \sum_{i \in A} \bar{Y}_i \\ \bar{Y}_C &= \frac{1}{N_0} \sum_{i \in C} \bar{Y}_i \\ \sigma_A^2 &= \frac{1}{N_1 - 1} \sum_{i \in A} (\bar{Y}_i - \bar{Y}_A)^2 \\ \sigma_C^2 &= \frac{1}{N_0 - 1} \sum_{i \in C} (\bar{Y}_i - \bar{Y}_C)^2, \text{ and} \\ \sigma_p^2 &= \frac{(N_1 - 1)\sigma_A^2 + (N_0 - 1)\sigma_C^2}{N_0 + N_1 - 2}.\end{aligned}$$

The test statistic for the two-sample t -test is then:

$$t = \frac{\bar{Y}_A - \bar{Y}_C}{\sqrt{\sigma_p^2 \left(\frac{1}{N_0} + \frac{1}{N_1} \right)}}.$$

This value was compared to a t -distribution with $N_0 + N_1 - 2$ degrees of freedom to determine whether there was a significant regional difference in mean FA between ASDs and control participants.

3.2.2.2 Group Comparison Of Voxelwise Means The second method I used corresponds to the standard voxelwise approach introduced in Section 1.3.2. In this approach, a separate two-sample t -test was performed at each voxel in the regions of interest, yielding a separate t -statistic at each voxel. These can then be mapped directly for each region, or compared against a t -distribution with $N_0 + N_1 - 2$ degrees of freedom to form p-value maps.

3.2.2.3 LME With Homogeneous Variance The third method is a linear mixed effects model with a random intercept and homogeneous group variance, such as described in Section 2.3. The fixed effects portion of the model included a separate mean fit at each voxel and an ASDs-group deviation from the mean. The random effects portion of the model consisted of a region-wide subject-specific random intercept.

Specifically, suppose a region of interest has V voxels, denote the FA for subject i at voxel j by Y_{ij} , and let \mathbf{g} be an $N \times 1$ vector of group indicators ($g_i = 0$ for $i \in C$, $g_i = 1$ for $i \in A$). Then the model is:

$$Y_{ij} = \mu_j + g_i \gamma_j + \delta_i + \epsilon_{ij}, \quad (3.1)$$

In Equation 3.1, μ_j and γ_j are unknown, fixed parameters, $\delta_i \sim N(0, \sigma^2)$, and $\epsilon_{ij} \sim N(0, \tau^2)$. I will let $\boldsymbol{\beta} = [\boldsymbol{\mu}' \boldsymbol{\gamma}']'$ denote the entire vector of fixed effects parameters.

I estimated the model parameters according to the closed-form procedure described in Section 2.3. That is, letting

$$\mathbf{X}_{NV \times 2V} = [(\mathbf{1}_{N \times 1} \otimes \mathbf{I}_{V \times V}) (\mathbf{g}_{N \times 1} \otimes \mathbf{I}_{V \times V})], \quad (3.2)$$

the fixed effects parameters were given in standard form:

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{Y}. \quad (3.3)$$

The variance parameters were estimated according to Equations 2.14 and 2.15:

$$\begin{aligned} \hat{\sigma}^2 &= \frac{1}{(V-1)NV} \sum_{i=1}^N \sum_{j \neq k}^V (Y_{ij} - (\hat{\mu}_j + g_i \hat{\gamma}_j))(Y_{ik} - (\hat{\mu}_j + g_i \hat{\gamma}_j)) \\ \hat{\tau}^2 &= \frac{1}{NV} \sum_{i=1}^N \sum_{j=1}^V (Y_{ij} - (\hat{\mu}_j + g_i \hat{\gamma}_j))^2 - \hat{\sigma}^2. \end{aligned}$$

The estimated covariance matrix of the observations \mathbf{Y} is:

$$\hat{\boldsymbol{\Sigma}} = \mathbf{I}_{N \times N} \otimes (\hat{\sigma}^2 \mathbf{J}_{V \times V} + \hat{\tau}^2 \mathbf{I}_{V \times V}),$$

and its inverse is given in closed-form by:

$$\hat{\boldsymbol{\Sigma}}^{-1} = \frac{1}{V \hat{\tau}^2 \hat{\sigma}^2 + \hat{\tau}^4} \mathbf{I}_{N \times N} \otimes ((V \hat{\sigma}^2 + \hat{\tau}^2) \mathbf{I}_{V \times V} - \hat{\sigma}^2 \mathbf{J}_{V \times V}).$$

The covariance matrix of the fixed effects estimates, $\hat{\boldsymbol{\beta}}$ can thus be calculated as:

$$\text{Var}(\hat{\boldsymbol{\beta}}) = (\mathbf{X}'\hat{\boldsymbol{\Sigma}}^{-1}\mathbf{X})^{-1},$$

and a t -statistic can be calculated for the difference between ASDs and control participants on FA at each voxel by dividing the appropriate γ coefficient in Equation 3.1 by its standard error, the square root of the corresponding diagonal element of $\text{Var}(\hat{\boldsymbol{\beta}})$. These t -statistics can be mapped directly or compared to a t -distribution on $N_1 + N_2 - 2$ degrees of freedom to obtain p-value maps.

In order to compare this model with the other LME-based approaches via likelihood-ratio tests, I also calculated the maximized log-likelihood in closed-form as:

$$\begin{aligned} -2 \log L(\mathbf{Y}; \hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\Sigma}}) &= -2 \log \left(\frac{1}{(2\pi)^{N/2} |\hat{\boldsymbol{\Sigma}}|^{1/2}} \exp \left(-\frac{1}{2} (\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}})' \hat{\boldsymbol{\Sigma}}^{-1} (\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}) \right) \right) \\ &= N \log(2\pi) + \log(|\hat{\boldsymbol{\Sigma}}|) + \left((\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}})' \hat{\boldsymbol{\Sigma}}^{-1} (\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}) \right) \\ &= N \log(2\pi) + N \log(\tau^{2V} + V\sigma^2\tau^{2(V-1)}) + \left((\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}})' \hat{\boldsymbol{\Sigma}}^{-1} (\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}) \right). \end{aligned} \quad (3.4)$$

3.2.2.4 LME With Heterogeneous Between-Subject Variance The fourth analytic approach involved fitting an LME very similar to that in Section 3.2.2.3, but allowing for group heterogeneity in between-subjects variance. Closed-form estimation for such models was described in Section 2.4.2. Specifically, the model was:

$$Y_{ij} = \mu_j + g_i\gamma_j + (1 - g_i)\delta_{0i} + g_i\delta_{1i} + \epsilon_{ij}, \quad (3.5)$$

where $g \in \{0, 1\}$, $\delta_0 \sim N(0, \sigma_0^2)$, and $\delta_1 \sim N(0, \sigma_1^2)$. Otherwise, the notation and assumptions are the same as in Model 3.1.

The fixed effects portion of this model was estimated as in Equation 3.3. The covariance parameters are estimated as in Equation 2.24, letting $\hat{Y}_{ij} = \hat{\mu}_j + g_i \hat{\gamma}_j$:

$$\begin{aligned} \hat{\tau}^2 &= \frac{1}{N(V-1)} \sum_{i=1}^N \sum_{j=1}^V (Y_{ij} - \hat{Y}_{ij})^2 \\ &\quad - \frac{1}{NV(V-1)} \sum_{i=1}^N \sum_{j=1}^V \sum_{k=1}^V (Y_{ij} - \hat{Y}_{ij})(Y_{ik} - \hat{Y}_{ik}) \end{aligned} \quad (3.6)$$

$$\begin{aligned} \hat{\sigma}_0^2 &= \frac{1}{NV(V-1)} \left(\frac{VN_0 + (V-1)N_1}{VN_0} \sum_{i \in C} \sum_{j=1}^V \sum_{k=1}^V (Y_{ij} - \hat{Y}_{ij})(Y_{ik} - \hat{Y}_{ik}) \right. \\ &\quad \left. + \frac{1}{V} \sum_{i \in A} \sum_{j=1}^V \sum_{k=1}^V (Y_{ij} - \hat{Y}_{ij})(Y_{ik} - \hat{Y}_{ik}) - \sum_{i=1}^N \sum_{j=1}^V (Y_{ij} - \hat{Y}_{ij})^2 \right) \end{aligned} \quad (3.7)$$

$$\begin{aligned} \hat{\sigma}_1^2 &= \frac{1}{NV(V-1)} \left(\frac{VN_1 + (V-1)N_0}{VN_1} \sum_{i \in A} \sum_{j=1}^V \sum_{k=1}^V (Y_{ij} - \hat{Y}_{ij})(Y_{ik} - \hat{Y}_{ik}) \right. \\ &\quad \left. + \frac{1}{V} \sum_{i \in C} \sum_{j=1}^V \sum_{k=1}^V (Y_{ij} - \hat{Y}_{ij})(Y_{ik} - \hat{Y}_{ik}) - \sum_{i=1}^N \sum_{j=1}^V (Y_{ij} - \hat{Y}_{ij})^2 \right). \end{aligned} \quad (3.8)$$

The estimated covariance matrix of the observations \mathbf{Y} is:

$$\hat{\Sigma} = (\mathbf{I} - \text{diag}(\mathbf{g}))_{N \times N} \otimes (\hat{\sigma}_0^2 \mathbf{J}_{V \times V} + \hat{\tau}^2 \mathbf{I}_{V \times V}) + \text{diag}(\mathbf{g})_{N \times N} \otimes (\hat{\sigma}_1^2 \mathbf{J}_{V \times V} + \hat{\tau}^2 \mathbf{I}_{V \times V}),$$

where $\text{diag}(\mathbf{g})$ is the diagonal matrix whose $(i, i)^{\text{th}}$ entry is the group indicator g_i . Its inverse is given in closed-form by:

$$\begin{aligned} \hat{\Sigma}^{-1} &= \frac{1}{V\hat{\tau}^2\hat{\sigma}_0^2 + \hat{\tau}^4} (\mathbf{I} - \text{diag}(\mathbf{g}))_{N \times N} \otimes ((V\hat{\sigma}_0^2 + \hat{\tau}^2) \mathbf{I}_{V \times V} - \hat{\sigma}_0^2 \mathbf{J}_{V \times V}) \\ &\quad + \frac{1}{V\hat{\tau}^2\hat{\sigma}_1^2 + \hat{\tau}^4} \text{diag}(\mathbf{g})_{N \times N} \otimes ((V\hat{\sigma}_1^2 + \hat{\tau}^2) \mathbf{I}_{V \times V} - \hat{\sigma}_1^2 \mathbf{J}_{V \times V}). \end{aligned}$$

The covariance matrix of the fixed effects estimates, $\hat{\beta}$ can again be calculated as:

$$\text{Var}(\hat{\beta}) = (\mathbf{X}' \hat{\Sigma}^{-1} \mathbf{X})^{-1},$$

and t -statistic or p -value maps can be produced for the group difference parameters as before.

The maximized log-likelihood for this model is given by:

$$\begin{aligned}
-2 \log L(\mathbf{Y}; \hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\Sigma}}) &= -2 \log \left(\frac{1}{(2\pi)^{N/2} |\hat{\boldsymbol{\Sigma}}|^{1/2}} \exp \left(-\frac{1}{2} (\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}})' \hat{\boldsymbol{\Sigma}}^{-1} (\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}) \right) \right) \\
&= N \log(2\pi) + \log(|\hat{\boldsymbol{\Sigma}}|) + \left((\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}})' \hat{\boldsymbol{\Sigma}}^{-1} (\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}) \right) \quad (3.9) \\
&= N \log(2\pi) + N_0 \log(\tau^{2V} + V\sigma_0^2\tau^{2(V-1)}) + N_1 \log(\tau^{2V} + V\sigma_1^2\tau^{2(V-1)}) \\
&\quad + \left((\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}})' \hat{\boldsymbol{\Sigma}}^{-1} (\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}) \right).
\end{aligned}$$

3.2.2.5 LME With Combined Heterogeneous Variance The final approach involved fitting another LME analogous to that in Section 3.2.2.4, but allowing for group heterogeneity in residual variance in addition to heterogeneity in between-subjects variance. Closed-form estimation for such models was described in Section 2.4.3. Specifically, the model was:

$$Y_{ij} = \mu_j + g_i\gamma_j + (1 - g_i)(\delta_{0i} + \epsilon_{0ij}) + g_i(\delta_{1i} + \epsilon_{1ij}), \quad (3.10)$$

where $\delta_0 \sim N(0, \sigma_0^2)$, $\delta_1 \sim N(0, \sigma_1^2)$, $\epsilon_0 \sim N(0, \tau_0^2)$ and $\epsilon_1 \sim N(0, \tau_1^2)$. Otherwise, the notation and assumptions are the same as in Model 3.1.

The fixed effects portion of this model was estimated as in Equation 3.3. The covariance parameters are estimated as in Equation 2.27:

$$\hat{\sigma}_0^2 = \frac{1}{N_0 V(V-1)} \sum_{i \in C} \sum_{j \neq k}^V (Y_{ij} - \hat{Y}_{ij})(Y_{ik} - \hat{Y}_{ik}) \quad (3.11)$$

$$\hat{\tau}_0^2 = \frac{1}{N_0 V} \sum_{i \in C} \sum_{j=1}^V (Y_{ij} - \hat{Y}_{ij})^2 - \hat{\sigma}_0^2 \quad (3.12)$$

$$\hat{\sigma}_1^2 = \frac{1}{N_1 V(V-1)} \sum_{i \in A} \sum_{j \neq k}^V (Y_{ij} - \hat{Y}_{ij})(Y_{ik} - \hat{Y}_{ik}) \quad (3.13)$$

$$\hat{\tau}_1^2 = \frac{1}{N_1 V} \sum_{i \in A} \sum_{j=1}^V (Y_{ij} - \hat{Y}_{ij})^2 - \hat{\sigma}_1^2. \quad (3.14)$$

The estimated covariance matrix of the observations \mathbf{Y} is:

$$\hat{\boldsymbol{\Sigma}} = (\mathbf{I} - \text{diag}(\mathbf{g}))_{N \times N} \otimes (\hat{\sigma}_0^2 \mathbf{J}_{V \times V} + \hat{\tau}_0^2 \mathbf{I}_{V \times V}) + \text{diag}(\mathbf{g})_{N \times N} \otimes (\hat{\sigma}_1^2 \mathbf{J}_{V \times V} + \hat{\tau}_1^2 \mathbf{I}_{V \times V}).$$

Its inverse is given in closed-form by:

$$\begin{aligned}\hat{\Sigma}^{-1} &= \frac{1}{V\hat{\tau}_0^2\hat{\sigma}_0^2 + \hat{\tau}_0^4}(\mathbf{I} - \text{diag}(\mathbf{g}))_{N \times N} \otimes ((V\hat{\sigma}_0^2 + \hat{\tau}_0^2)\mathbf{I}_{V \times V} - \hat{\sigma}_0^2\mathbf{J}_{V \times V}) \\ &\quad + \frac{1}{V\hat{\tau}_1^2\hat{\sigma}_1^2 + \hat{\tau}_1^4}\text{diag}(\mathbf{g})_{N \times N} \otimes ((V\hat{\sigma}_1^2 + \hat{\tau}_1^2)\mathbf{I}_{V \times V} - \hat{\sigma}_1^2\mathbf{J}_{V \times V}).\end{aligned}$$

The covariance matrix of the fixed effects estimates, $\hat{\beta}$ can again be calculated as:

$$\text{Var}(\hat{\beta}) = (\mathbf{X}'\hat{\Sigma}^{-1}\mathbf{X})^{-1},$$

and t -statistic or p -value maps can be produced for the group difference parameters as before.

The maximized log-likelihood for this model is given by:

$$\begin{aligned}-2 \log L(\mathbf{Y}; \hat{\beta}, \hat{\Sigma}) &= -2 \log \left(\frac{1}{(2\pi)^{N/2} |\hat{\Sigma}|^{1/2}} \exp \left(-\frac{1}{2} (\mathbf{Y} - \mathbf{X}\hat{\beta})' \hat{\Sigma}^{-1} (\mathbf{Y} - \mathbf{X}\hat{\beta}) \right) \right) \\ &= N \log(2\pi) + N_0 \log(\tau_0^{2V} + V\sigma_0^2\tau_0^{2(V-1)}) + N_1 \log(\tau_1^{2V} + V\sigma_1^2\tau_1^{2(V-1)}) \\ &\quad + \left((\mathbf{Y} - \mathbf{X}\hat{\beta})' \hat{\Sigma}^{-1} (\mathbf{Y} - \mathbf{X}\hat{\beta}) \right).\end{aligned}\tag{3.15}$$

3.2.3 Results

Below, I describe the results of the FA analyses, region-by-region.

3.2.3.1 Region 1 Figure 3.7 displays the maps of t -statistics summarizing group difference at each voxel of the splenium region from each of the five analytic methods discussed in Section 3.2.2, in order of model complexity. Figure 3.8 displays the corresponding maps of p -values. Note that the regional mean model (Section 3.2.2.1) yields a single t -statistic and p -value. These have been converted to maps by assigning the value of the regional t -statistic or p -value to each voxel in the region.

The direction of the differences for the t -statistic maps are ASDs mean - control mean, so that blue blocks reflect voxels in which control participants had a higher mean FA than ASDs participants, and red blocks reflect voxels in which the ASDs mean FA was higher. Significant p -values ($p < .05$) are displayed in red on the p -value maps.

Model 1 yielded $t_{118} = -1.38$, $p = .17$, indicating no significant difference between the groups in mean FA averaged over the whole region. The overall patterns of t -statistics

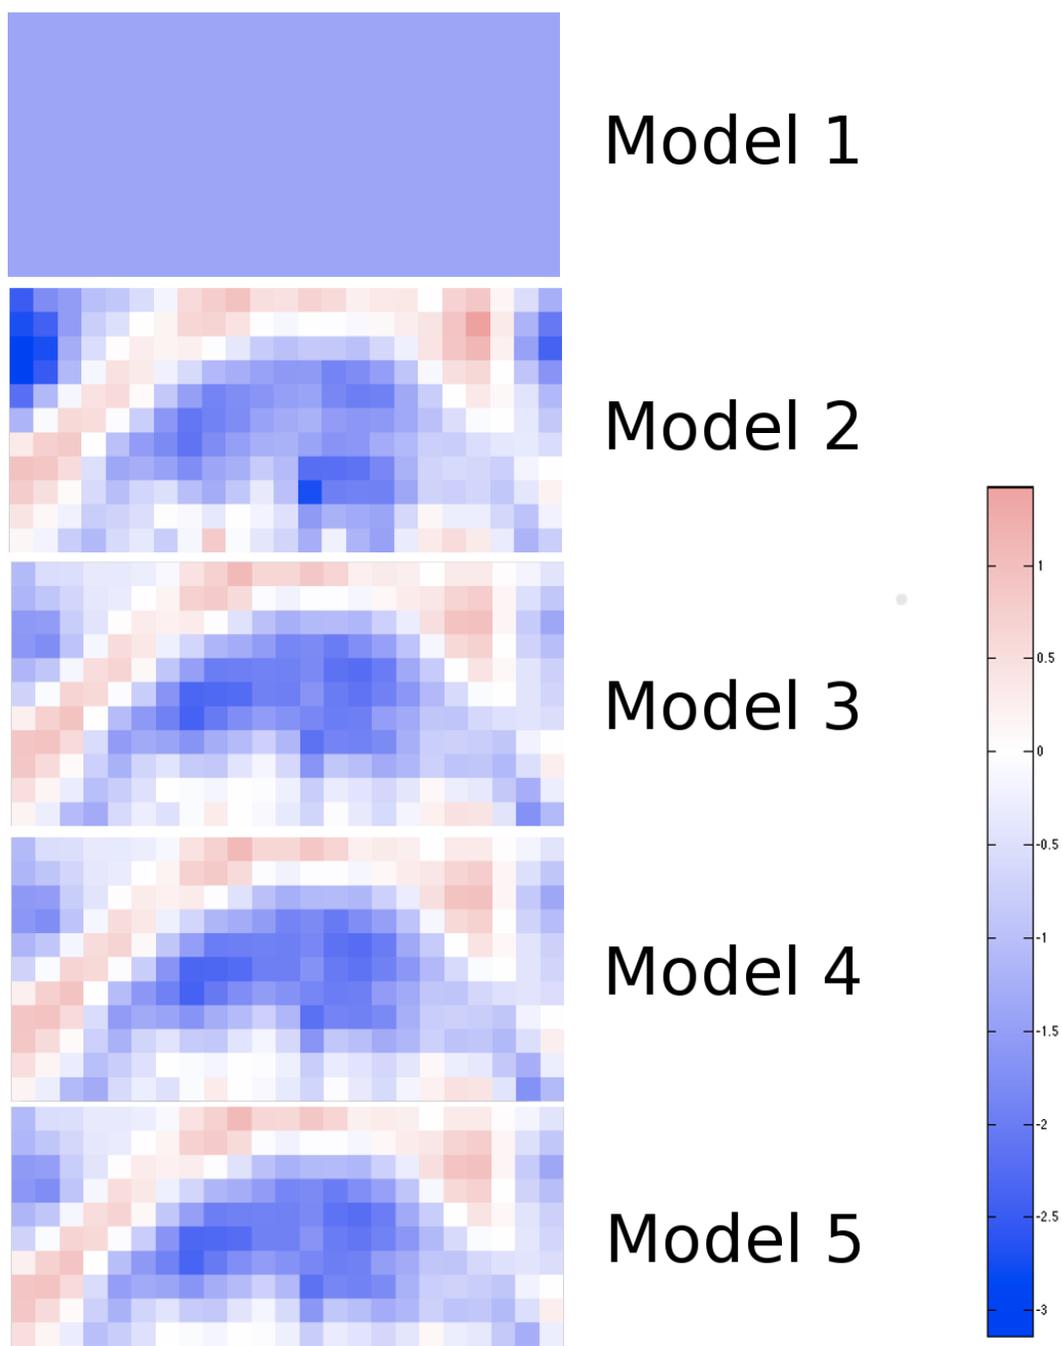


Figure 3.7: Group difference t-statistics for FA in Region 1

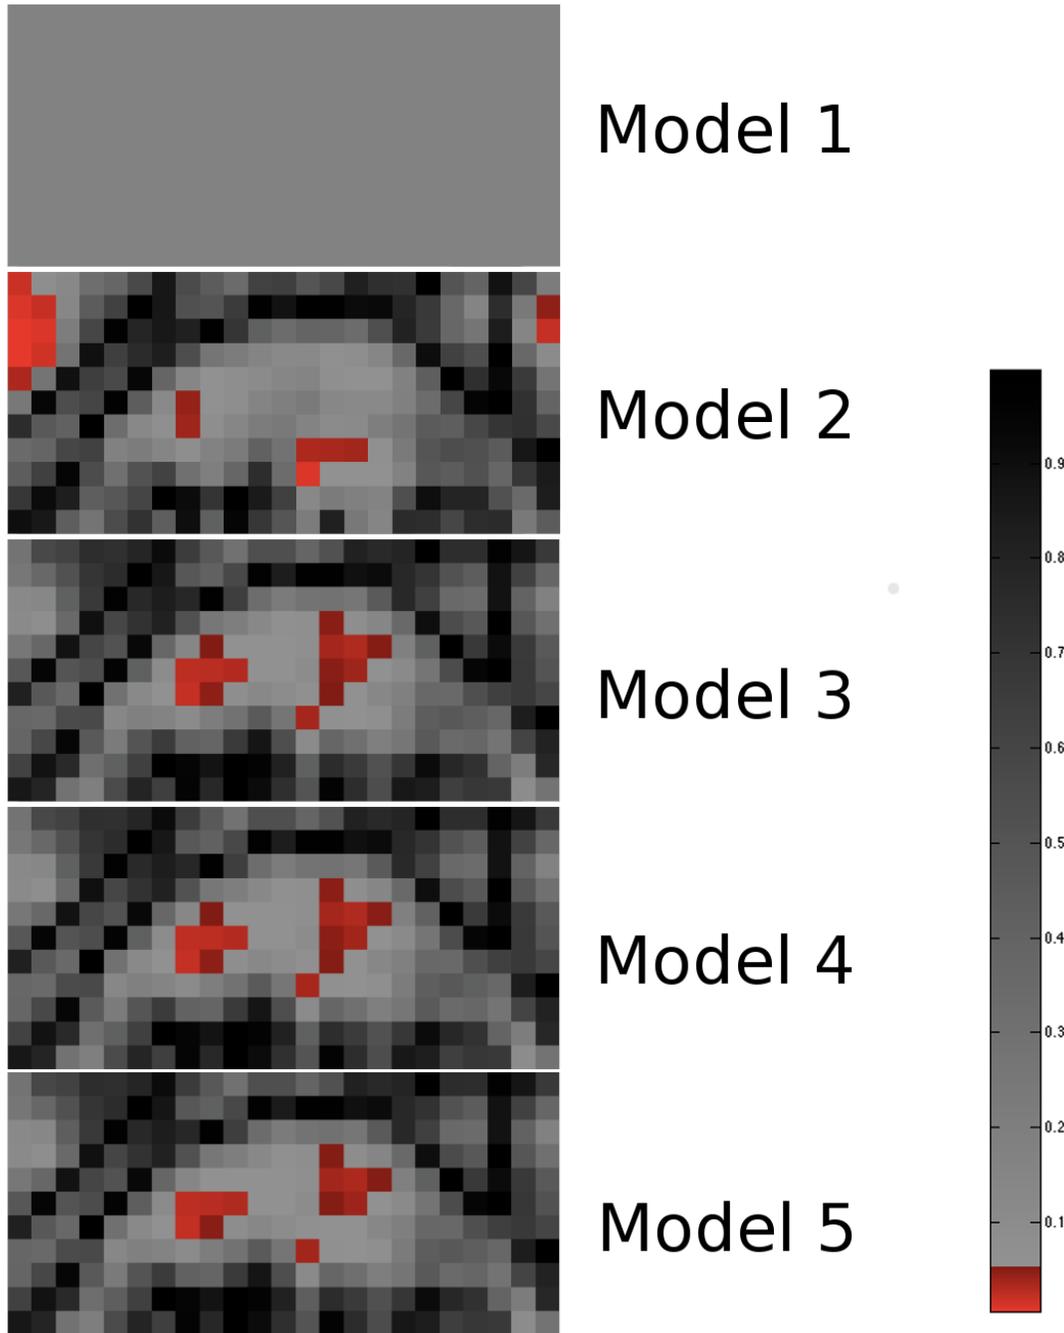


Figure 3.8: Group difference p-values for FA in Region 1

are comparable between Models 2 - 5, with mean FA for control participants higher than for ASDs participants in the main body of the splenium, and mean FA higher for ASDs participants in the anterior boundary of the splenium.

There are, however, some interesting differences among the models. The t map for Model 2 (voxelwise t-tests) is less smooth than those for the Models 3 - 5: Contrasts between neighboring voxels tend to be more abrupt for Model 2 than for the LME-based methods, especially in the lower part of the figure (posterior boundary of the splenium). And, when looking at the p-value maps, the voxels with significant group differences for Models 3 - 5 form a roughly bilaterally symmetric pattern in the densest part of the splenium. On the other hand, the areas of significant difference under Model 2 are less coherent. An echo of the bilateral mid-splenium difference is still visible, but the area of largest difference is at the upper-left portion of the figure, corresponding to the boundary between the splenium and adjacent structures.

The differences among the three LME-based approaches (Models 3 - 5) are much more subtle. The t maps are visually indistinguishable, and the p-value maps are also very similar, although a few voxels that show significant group difference under Models 3 and 4 are no longer significant under Model 5. According to likelihood ratio tests, Model 5 provides a significantly better fit than Model 3 ($\chi^2_2 = 16.91, p = .0002$) or Model 4 ($\chi^2_1 = 16.73, p < .0001$), whereas there is no evidence of a difference in fit between Models 3 and 4 ($\chi^2_1 = 0.18, p = .67$).

3.2.3.2 Region 2 The results for the region incorporating the genu are displayed in Figures 3.9 (t -statistic maps) and 3.10 (p-value maps). Model 1 shows a significant overall regional group difference in FA ($t_{118} = -2.48, p = .02$), with higher mean FA for control participants than for ASDs participants. Models 2 - 5 also show increased FA for control relative to ASDs participants, particularly in the left hemisphere (right side of the figures). Interestingly, the magnitude of the effect is substantially attenuated for the voxelwise t -test approach (Model 2) relative to the LME-based approaches (Models 3 - 5). This is particularly evident in the p-value map, where only three voxels of the genu (and one voxel anterior to the genu) show significant group differences, compared to approximately 20 voxels

with significant group differences for each of the LME-based models.

As with the Region 1 analyses, the differences among the LME-based approaches are quite subtle. Likelihood ratio tests indicate that incorporating group heterogeneity in variance is justified: $\chi_1^2 = 6.00, p = .01$ for Model 3 vs. Model 4, $\chi_2^2 = 50.16, p < .0001$ for Model 3 vs. Model 5, and $\chi_1^2 = 44.15, p < .0001$ for Model 4 vs. Model 5. Thus, we would have greatest confidence in the results for the model incorporating group heterogeneity in both between-subject and residual variance parameters, summarized in the bottom panels of Figures 3.9 and 3.10.

3.2.3.3 Region 3 Figure 3.11 displays the t statistics for the ventricular region as calculated by the five models under consideration, and Figure 3.12 shows the associated p -values. There was no overall regional difference in mean FA according to Model 1 ($t_{118} = -1.02, p = .31$). Models 2 – 5 show no group differences in the ventricles themselves, as is appropriate. The mean FA is higher in the control group than the ASDs group in the anterior portion of the figure, bordering on the genu. This difference is significant under the LME-based models, not under voxelwise t -tests. There is also a trend toward the mean FA being higher for the ASDs group than for controls in an anterior – posterior strip in the lower medial part of the figure. This strip may correspond to the septum pellucidum, which separates the anterior horns of the lateral ventricles.

There are no visually obvious differences among Models 3, 4 and 5, apart from a small number of voxels above or below threshold. Likelihood ratio tests indicate no difference between Models 3 and 4 ($\chi_1^2 = 0.42, p = .52$). Model 5, however, is superior to both ($\chi_2^2 = 14.78, p = .0006$ vs. Model 3 and $\chi_1^2 = 14.36, p = .0002$ vs. Model 4).

3.3 FULL TENSOR ANALYSES

The analyses in Section 3.2 demonstrated the application of closed-form LME based methods to the analysis of fractional anisotropy, which is the most widely-used of the scalar indices derived from diffusion tensor data. While this is an important demonstration, there is an

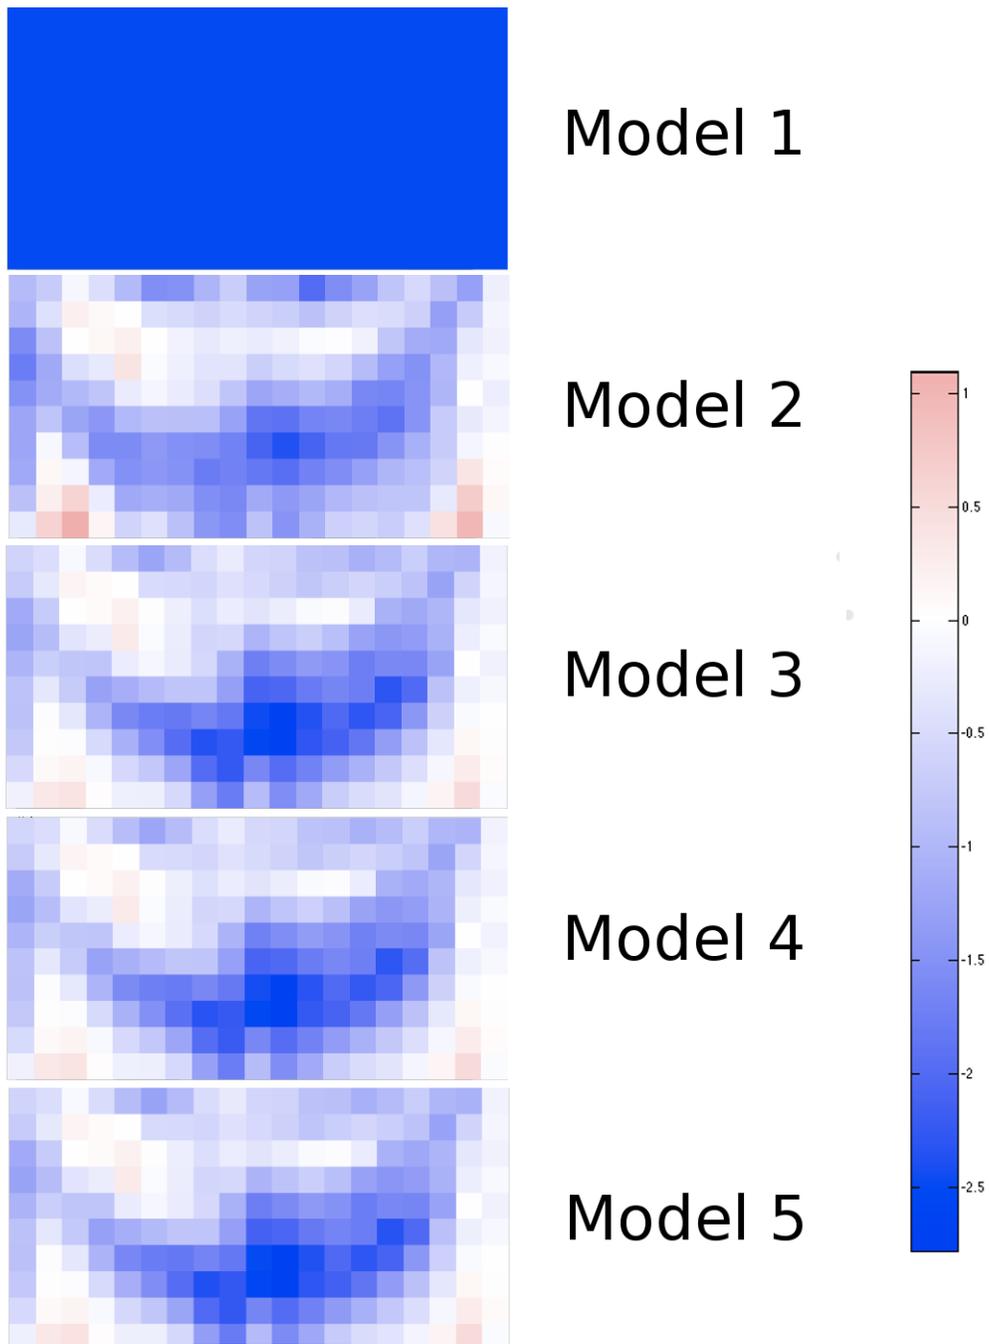


Figure 3.9: Group difference t-statistics for FA in Region 2

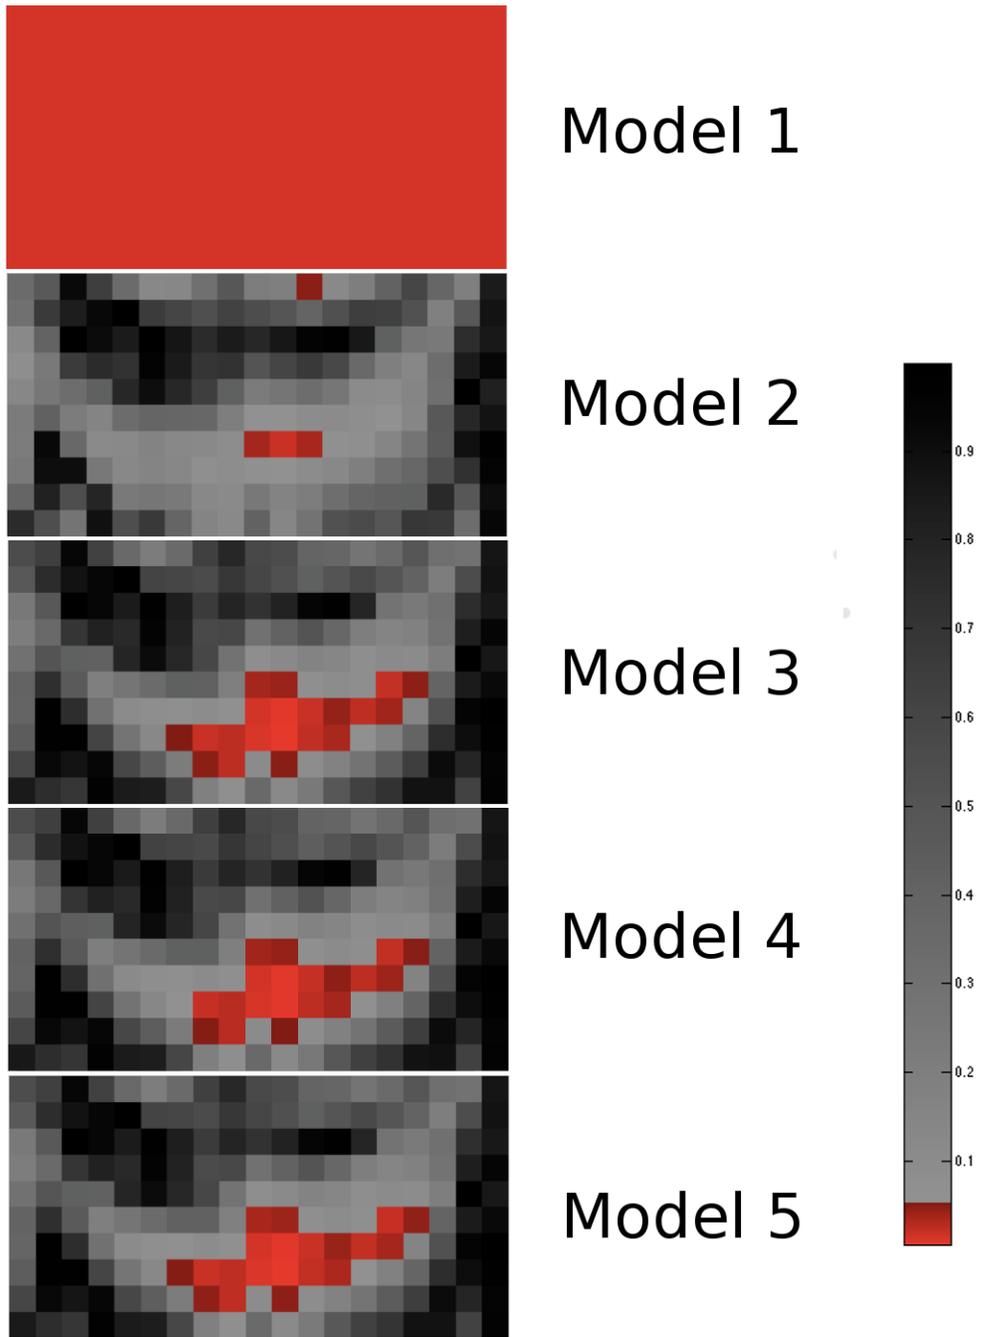


Figure 3.10: Group difference p-values for FA in Region 2

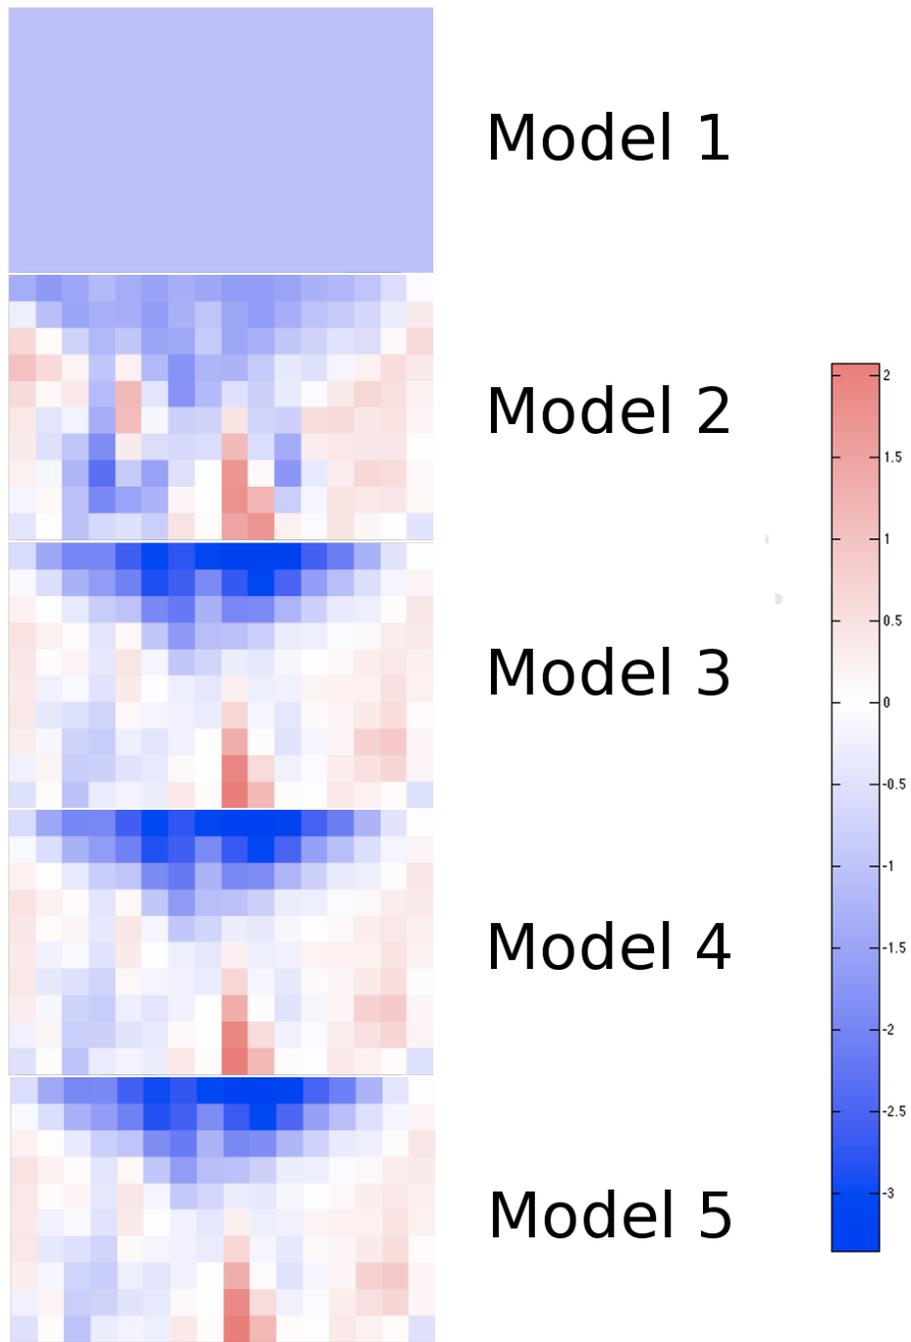


Figure 3.11: Group difference t-statistics for FA in Region 3

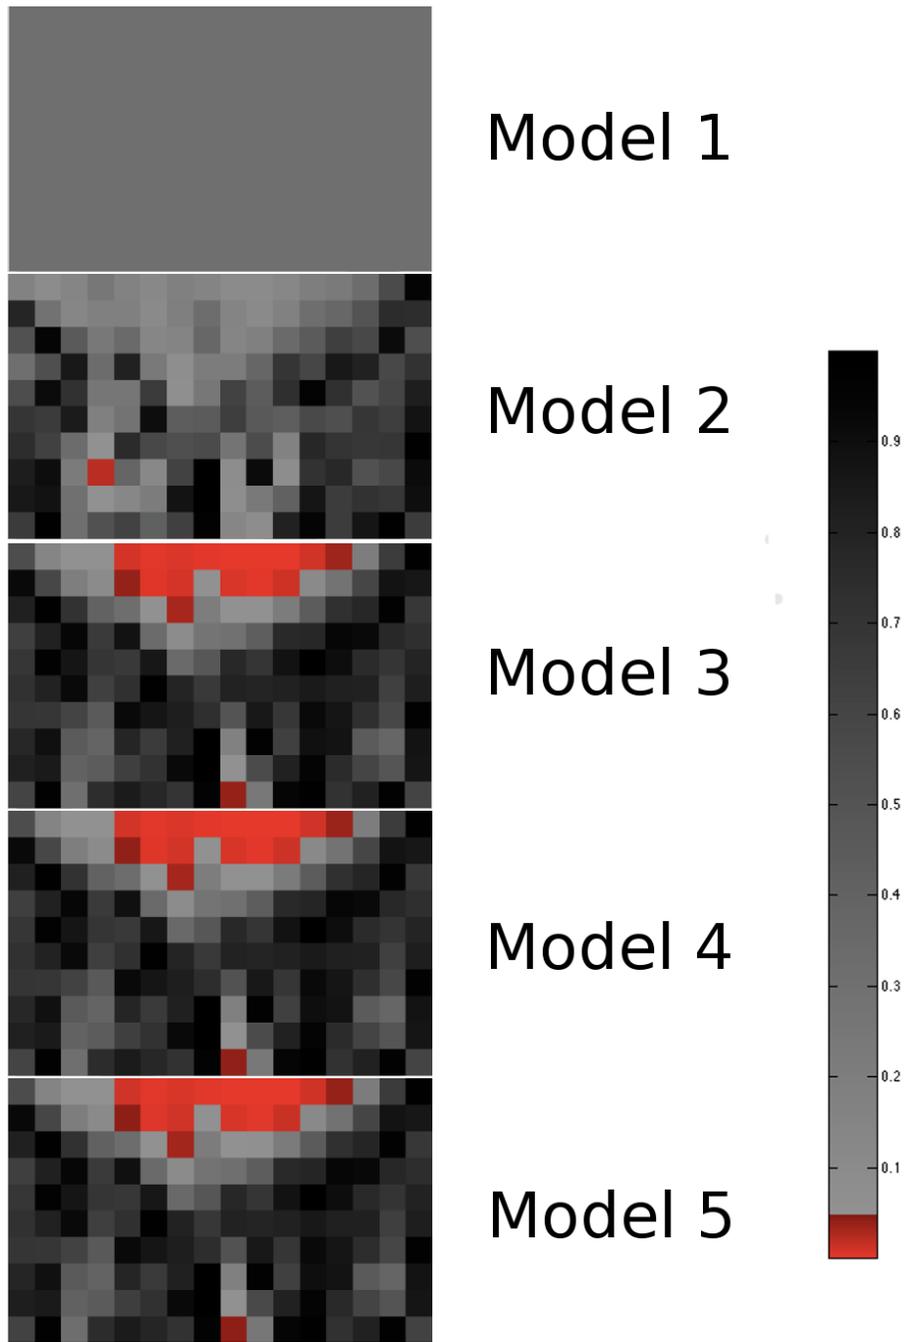


Figure 3.12: Group difference p-values for FA in Region 3

inherent limit to the sensitivity of any approach based on only one tensor-derived scalar index to detect white matter abnormalities. Diffusion tensors are six dimensional and any scalar reduction entails a loss of information, and there does not yet appear to be any principled method by which to identify which scalar index or indices will be sensitive to important group differences.

In this section, I will describe analyses I performed to compare the entire six-dimensional tensor data between groups. This approach is potentially sensitive to any white matter abnormality that can be captured by the tensor model.

The first step was to apply a decorrelating transformation to the tensor elements in order to justify analyzing them separately. I then applied five multivariate analytic approaches analogous to the univariate approaches in Section 3.2 to the decorrelated tensor data for each region described in Section 3.2.1, obtaining in each case regional maps of T^2 statistics and p-values.

3.3.1 Decorrelating Transformation

Because the tensor elements are not statistically independent, including all six together as dependent variables in a single LME model would require estimation of inter-element covariance terms, greatly increasing the complexity and the computational burden of the model. For that reason, I employed the transformation described in Section 2.1 to decorrelate the tensor elements. This allowed me to treat the transformed tensor elements as independent in each model, considerably simplifying the analysis.

Specifically, suppose a vectorized diffusion tensor is given by:

$$\mathbf{d} = [D_{xx} \ D_{yy} \ D_{zz} \ D_{xy} \ D_{xz} \ D_{yz}]'$$

Then the transformed tensor elements given by:

$$\tilde{\mathbf{d}} = \begin{bmatrix} \tilde{D}_{xx} \\ \tilde{D}_{yy} \\ \tilde{D}_{zz} \\ \tilde{D}_{xy} \\ \tilde{D}_{xz} \\ \tilde{D}_{yz} \end{bmatrix} = \begin{bmatrix} \frac{1}{\sqrt{3}} & \frac{1}{\sqrt{3}} & \frac{1}{\sqrt{3}} & 0 & 0 & 0 \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} & 0 & 0 & 0 & 0 \\ \frac{1}{2\sqrt{2}} & \frac{1}{2\sqrt{2}} & -\frac{1}{\sqrt{2}} & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \mathbf{d}$$

are independent under the assumption of rotational invariance of the gradient sampling scheme. The elements of $\tilde{\mathbf{d}}$ served as the dependent variables in the analyses described below.

3.3.2 Data Analyses

The data analyses for the full tensor comparisons were closely analogous to those for fractional anisotropy described in Section 3.2.2. The difference is that these are analyses with multiple dependent variables, and the methods used previously had to be adapted to the multivariate context. The five methods employed were:

1. A group comparison of regional multivariate means
2. A group comparison of voxelwise multivariate means
3. A series of LMEs with homogeneous between-subject variance and homogeneous residual variance
4. A series of LMEs with heterogeneous between-subject variance and homogeneous residual variance
5. A series of LMEs with heterogeneous between-subject variance and heterogeneous residual variance

3.3.2.1 Regional Multivariate Means The simplest analysis was a direct multivariate analogue of the region-wise t -test described in Section 3.2.2.1. To accommodate the multivariate nature of the transformed tensor data, I used a two-sample Hotelling's T^2 test. Specifically, let

$$\bar{\mathbf{d}}_i = [\bar{d}_{i1} \ \bar{d}_{i2} \ \bar{d}_{i3} \ \bar{d}_{i4} \ \bar{d}_{i5} \ \bar{d}_{i6}]'$$

denote the sample mean of the vector of transformed tensor elements for subject i in a single region, and let $\bar{\mathbf{d}}_C$ and $\bar{\mathbf{d}}_A$ denote the group means of these mean vectors for the control and ASDs groups, respectively. I then calculated the pooled variance of each transformed tensor element, d_l , as

$$\text{Var}(d_l) = \frac{1}{N_0 + N_1 - 2} \left(\sum_{i \in C} (\bar{d}_{li} - \bar{d}_{lC})^2 + \sum_{i \in A} (\bar{d}_{li} - \bar{d}_{lA})^2 \right),$$

and these were combined into a variance matrix as:

$$\mathbf{W} = \text{diag}(\text{Var}(d_1), \dots, \text{Var}(d_6)).$$

The two-sample T^2 statistic was then given by

$$t^2 = \frac{N_0 N_1}{N} (\bar{\mathbf{d}}_A - \bar{\mathbf{d}}_C)' \mathbf{W}^{-1} (\bar{\mathbf{d}}_A - \bar{\mathbf{d}}_C). \quad (3.16)$$

This quantity is proportional to a statistic with an F distribution, and so I compared

$$f = \frac{N - 7}{6(N - 2)} t^2$$

to an F -distribution on $(6, N - 7)$ degrees of freedom to obtain p-values.

3.3.2.2 Multivariate Voxelwise Means The second approach to whole-tensor comparisons was a direct multivariate analogue to the voxelwise scalar analyses introduced in Section 3.2.2.2. Rather than performing two-sample univariate t -tests at each voxel, I performed two-sample T^2 tests, using the calculations described above in Section 3.3.2.1, but substituting individual voxel values rather than regional means. This allowed me to create maps of T^2 statistics and corresponding p-values.

3.3.2.3 Multiple Linear Mixed Effects Models Approaches 3 - 5 in the list above were direct extensions of the models described in Sections 3.2.2.3 – 3.2.2.5. The novelty is that, since there are six distinct dependent variables, I fit a separate model for each transformed tensor element under each approach.

For simplicity, consider the homogeneous LME (approach 3). The basic model I used was that given in Equation 3.1, but I fit a separate model for each transformed tensor element d_l :

$$d_{ijl} = \mu_{jl} + g_i\gamma_{jl} + \delta_{il} + \epsilon_{ijl}, \quad (3.17)$$

where d_{ijl} is transformed tensor element l at voxel j for subject i , μ_{jl} is the overall mean of element l at voxel j , g_i is a group indicator as before, γ_{jl} an ASDs group deviation from the mean for element l at voxel j , $\delta_{il} \sim \text{MVN}(\mathbf{0}, \sigma_l^2)$ is a subject-specific random intercept for element l for subject i , and $\epsilon_{ijl} \sim \text{MVN}(\mathbf{0}, \tau_l^2)$ is a residual error term.

The estimation for each of these six models was conducted as in Section 3.2.2.3. This yielded an estimate of ASDs - control group difference, $\hat{\gamma}_{jl}$, for each voxel and each dependent variable, along with the variance of this estimate, $\text{Var}(\hat{\gamma}_{jl})$. In order to summarize the group tensor differences at each voxel, I combined these estimates into a single voxelwise T^2 statistic by setting:

$$\hat{\boldsymbol{\gamma}}_j = [\hat{\gamma}_{j1} \ \dots \ \hat{\gamma}_{j6}]',$$

and

$$\hat{\mathbf{\Gamma}}_j = \begin{bmatrix} \text{Var}(\hat{\gamma}_{j1}) & 0 & 0 & 0 & 0 & 0 \\ 0 & \text{Var}(\hat{\gamma}_{j2}) & 0 & 0 & 0 & 0 \\ 0 & 0 & \text{Var}(\hat{\gamma}_{j3}) & 0 & 0 & 0 \\ 0 & 0 & 0 & \text{Var}(\hat{\gamma}_{j4}) & 0 & 0 \\ 0 & 0 & 0 & 0 & \text{Var}(\hat{\gamma}_{j5}) & 0 \\ 0 & 0 & 0 & 0 & 0 & \text{Var}(\hat{\gamma}_{j6}) \end{bmatrix}.$$

I then formed a T^2 statistic for each voxel as:

$$t_j^2 = \hat{\boldsymbol{\gamma}}_j' \hat{\mathbf{\Gamma}}_j^{-1} \hat{\boldsymbol{\gamma}}_j$$

and compared

$$f = \frac{N - 7}{6(N - 2)} t^2$$

to an F -distribution with $(6, N-7)$ degrees of freedom to obtain p-value maps.

For the scalar index analyses in Section 3.2, I relied on likelihood ratio tests to determine whether heterogeneous variance models were justified by the data. One complication of using multiple univariate models to test for a difference in multivariate outcomes is that it is no longer possible to use a single likelihood ratio test to compare nested models. To work around this difficulty, I extended the likelihood ratio test to this context in the following way. Suppose we wish to compare Model A and Model B (with Model B nested within Model A) for a single transformed tensor element, d_l . We would form the likelihood ratio statistic:

$$\text{LR}_l = 2 \left(\log(L_A(d_l; \hat{\beta}, \hat{\Sigma})) - \log(L_B(d_l; \hat{\beta}, \hat{\Sigma})) \right),$$

which has an asymptotic χ^2 distribution on p degrees of freedom (the difference in the number of parameters between Models A and B). Noting that the sum of statistics with independent χ^2 distributions also has a χ^2 distribution with degrees of freedom equal to the sum of the degrees of freedom of the summands, I formed the statistic:

$$\text{LR}_{\mathbf{d}} = \sum_{l=1}^6 \text{LR}_l, \tag{3.18}$$

which has an asymptotic χ^2 distribution on $6p$ degrees of freedom. This statistic can be used to test the null hypothesis that Model A and Model B provide equivalent fits to the six dependent variables.

3.3.3 Results

The following sections summarize the results for the whole-tensor comparison, one region at a time.

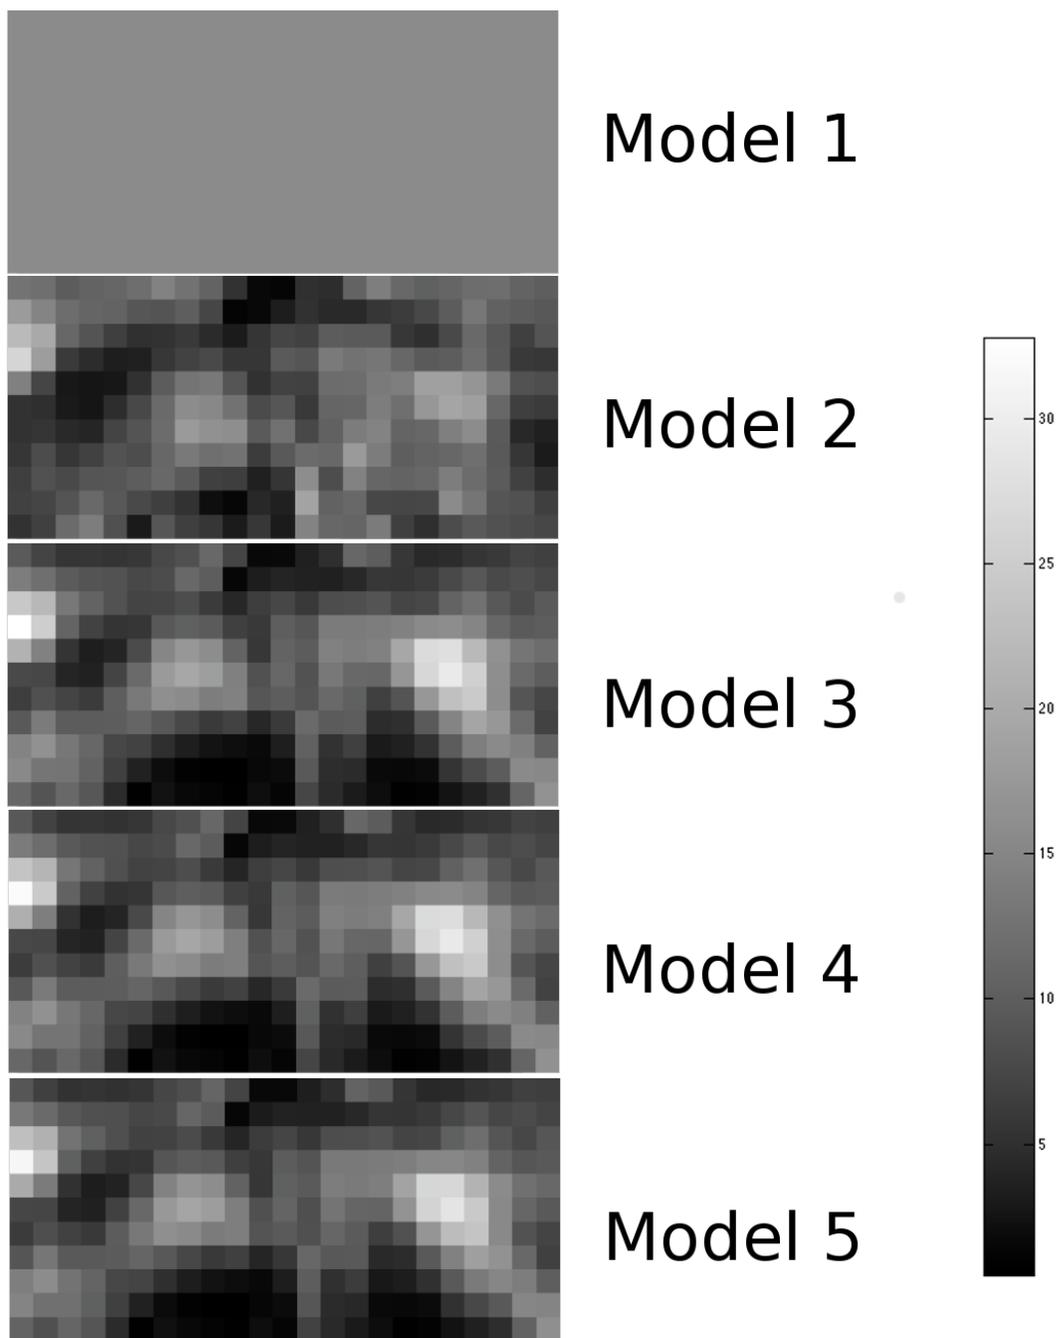


Figure 3.13: Group tensor difference T^2 statistics in Region 1

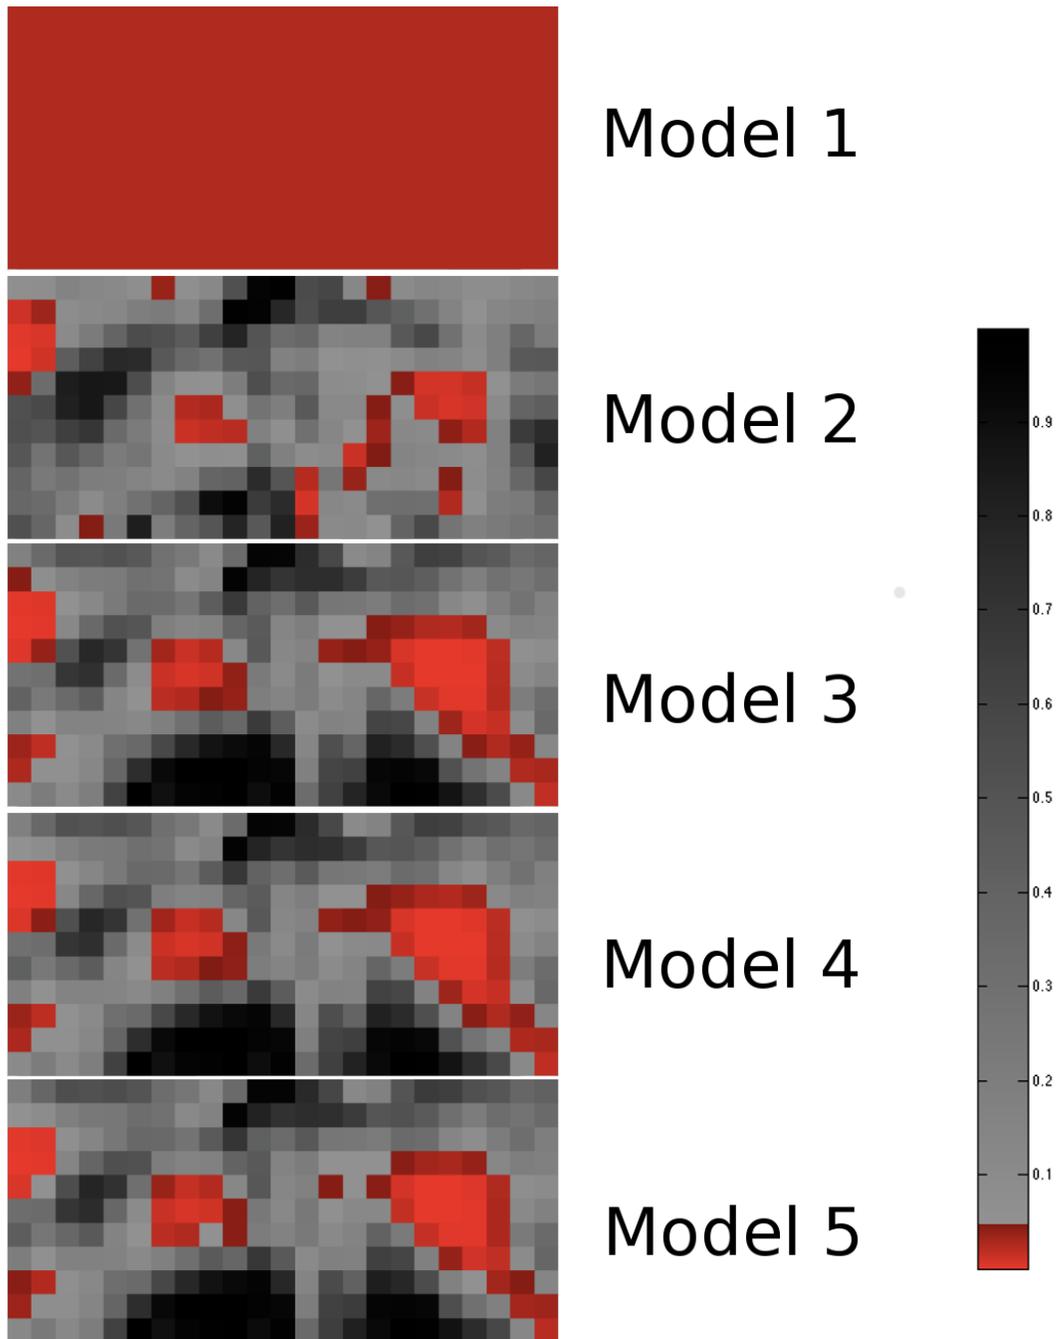


Figure 3.14: Group tensor difference p-values in Region 1

3.3.3.1 Region 1 The maps of T^2 statistics obtained from each of the models described in Section 3.3 for the splenium are displayed in Figure 3.13 and the associated p-values are shown in Figure 3.14.

Model 1 shows an overall regional difference in diffusion characteristics between ASDs and control participants ($f_{6,113} = 2.50, p = .03$). Model 2 shows a very scattered pattern of group difference, with very blurry boundaries between the dense white matter of the splenium and surrounding tissue. The p-value map shows scattered significant voxelwise differences in mean diffusion properties.

Models 3 – 5 reveal a much more cohesive pattern of group differences, with roughly bilaterally symmetric and largely spatially contiguous group differences in the splenium, as well as a right-hemisphere (left side of the figure) region anterior to the splenium. The differences among the LME models are visually subtle and are largely in the direction of attenuation of group difference effects with the incorporation of increasing heterogeneity, although not uniformly so.

The likelihood ratio test described in Equation 3.18 indicated significant improvement in fit with increased heterogeneity: $\chi_6^2 = 17.40, p = .008$ for Model 4 vs. Model 3, $\chi_{12}^2 = 280.58, p < .0001$ for Model 5 vs. Model 3, and $\chi_6^2 = 263.18, p < .0001$ for Model 5 vs. Model 4.

3.3.3.2 Region 2 Figure 3.15 displays the maps of T^2 statistics for each model for the genu region; associated p-values are shown in Figure 3.16.

Model 1 shows an overall significant regional difference in mean diffusion properties ($f_{6,113} = 4.78, p = .0002$). The T^2 map for Model 2 shows a somewhat confused pattern of group difference. The p-value map indicates that there are significant group differences in a region of the genu in the left hemisphere as well as two somewhat distinct clusters in the right hemisphere, including significant group difference in a posterior region that overlaps the anterior portion of the right lateral ventricle. Group differences in diffusion in the ventricles make very little sense, and these results are likely irreproducible.

Models 3 – 5 show a much more coherent pattern of group difference, with the greatest difference in the right hemisphere of the genu, but also greater difference in the left hemi-

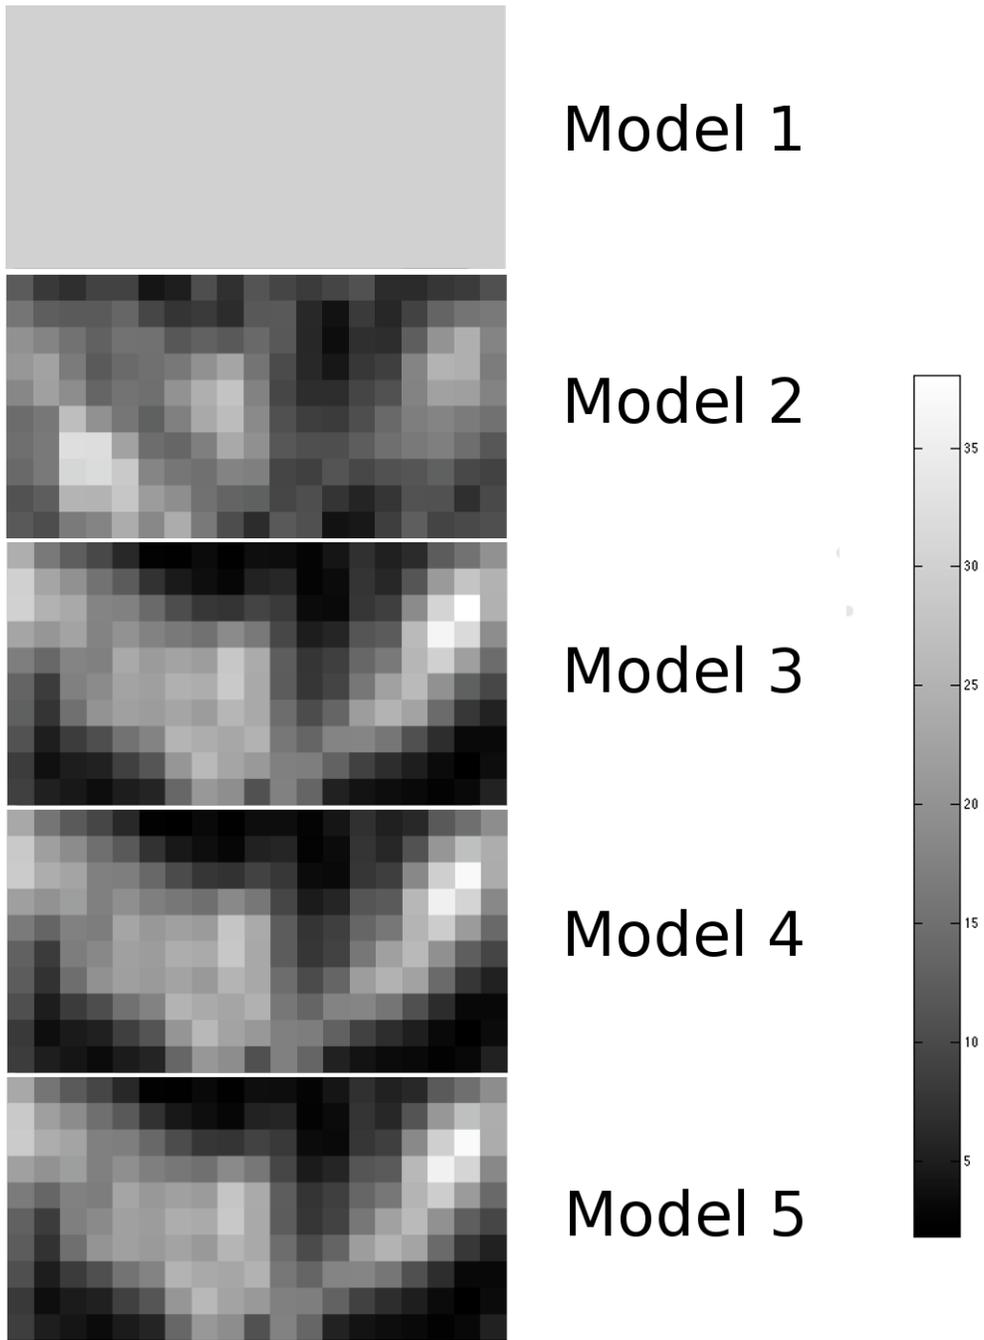


Figure 3.15: Group tensor difference T^2 statistics in Region 2

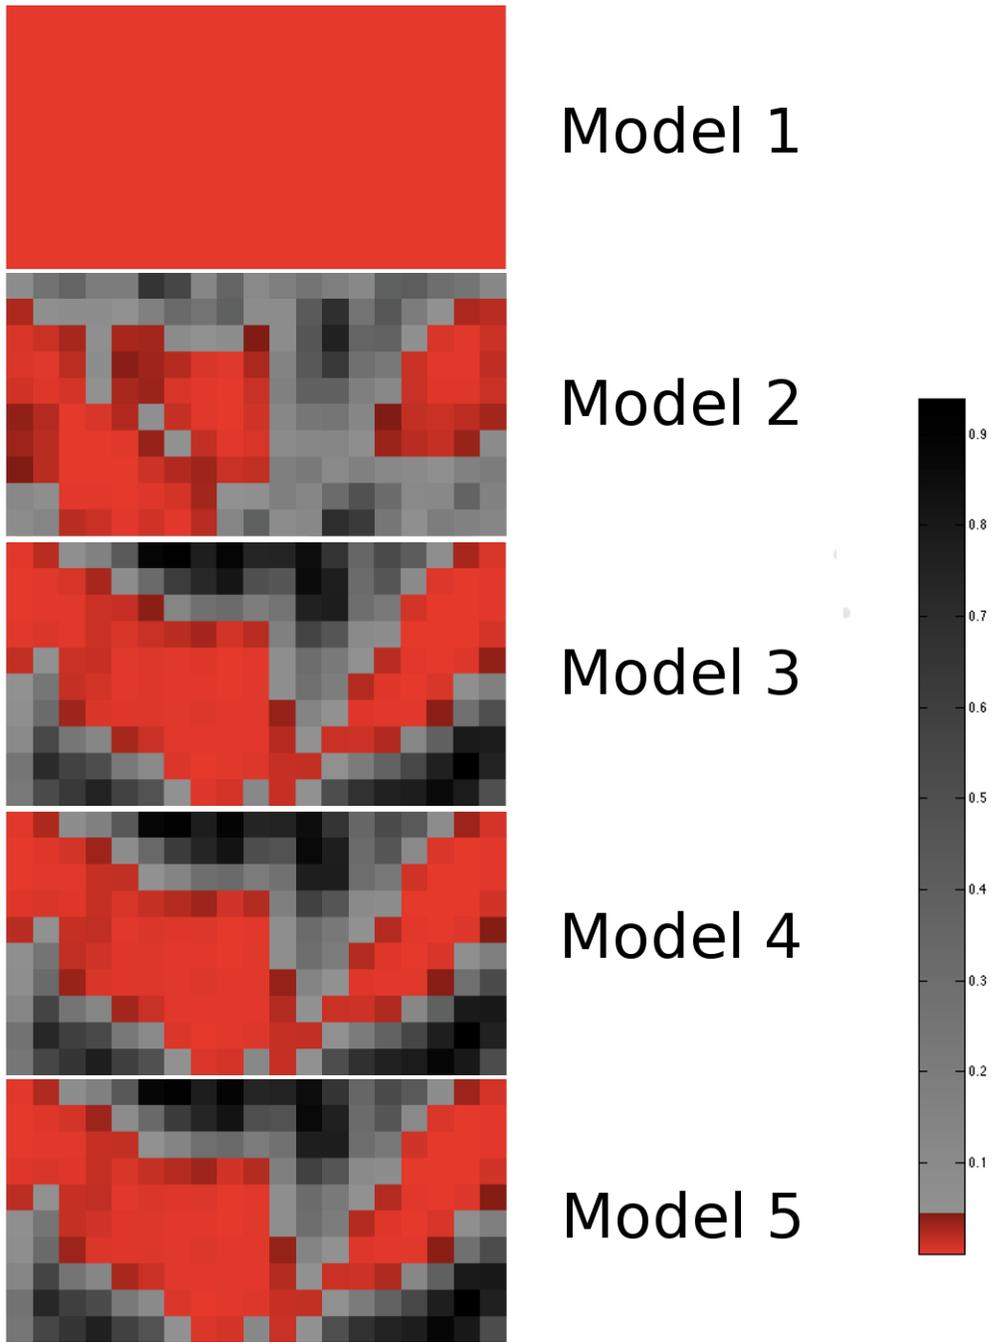


Figure 3.16: Group tensor difference p-values in Region 2

sphere than was shown by Model 2. There is no evidence of group difference in the lateral ventricles with the LME-based approaches. The generalized likelihood ratio test defined in Equation 3.18 shows no evidence that Model 4 is better than Model 3 ($\chi_6^2 = 6.22, p = .40$), but there is evidence that including heterogeneous residual variance provides a significantly better fit: $\chi_{12}^2 = 123.49, p < .0001$ for Model 3 vs. Model 5 and $\chi_6^2 = 117.27, p < .0001$ for Model 4 vs. Model 5.

3.3.3.3 Region 3 Figure 3.17 displays the maps of T^2 statistics for each model for the ventricular region, and Figure 3.18 contains the associated p-values.

According to Model 1, there is no overall regional difference in diffusion between ASDs and control participants ($f_{6,113} = 1.76, p = .11$). Model 2 shows a patchwork pattern of group differences in diffusion, including plausible significant bilateral differences in the anterior horn of the internal capsule, and some significant differences in the upper portion of the figure where the region intersects the posterior portion of the genu. However, there are also voxels that show significant differences in the right lateral ventricle (left side of the figure), and the ventricles overall are not well-defined in the figure, with many voxels trending toward significant differences.

This is not the case in Models 3 – 5, which yielded clearly defined patterns of group difference, with significant differences in the portions of the figure bordering on the genu and the internal capsule, and no suggestion of significant differences in diffusion in the lateral ventricles. There is also a medial strip of group difference that may correspond to the septum pellucidum.

As with the previous analyses, there is little visible difference among Models 3, 4 and 5. According to the generalized likelihood ratio test (Equation 3.18), there is again no significant difference between Models 3 and 4 ($\chi_6^2 = 1.80, p = .94$), but Model 5 provides a significantly better fit than Model 3 ($\chi_{12}^2 = 228.49, p < .0001$) or Model 4 ($\chi_6^2 = 226.68, p < .0001$).

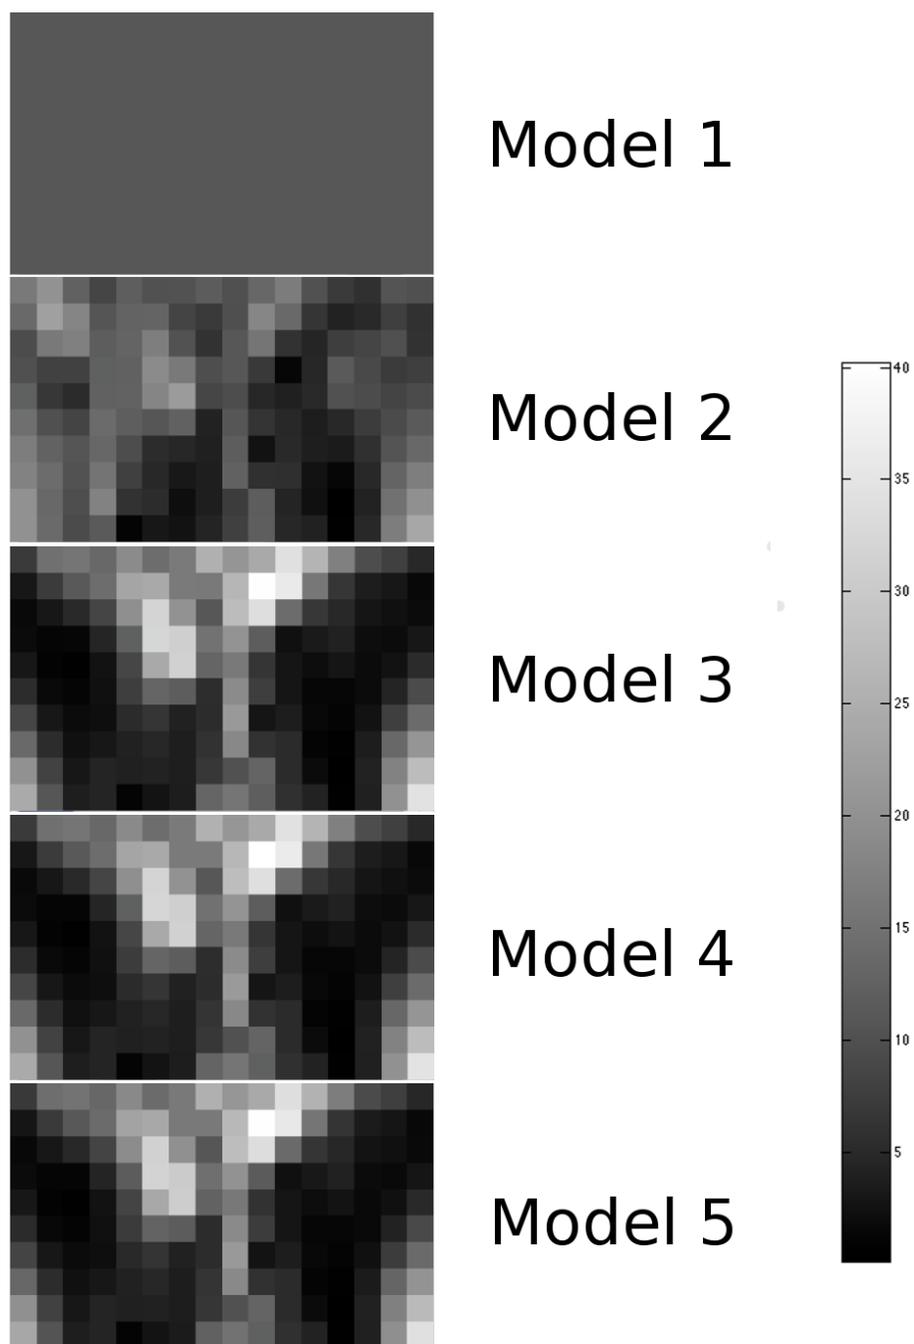


Figure 3.17: Group tensor difference T^2 statistics in Region 3

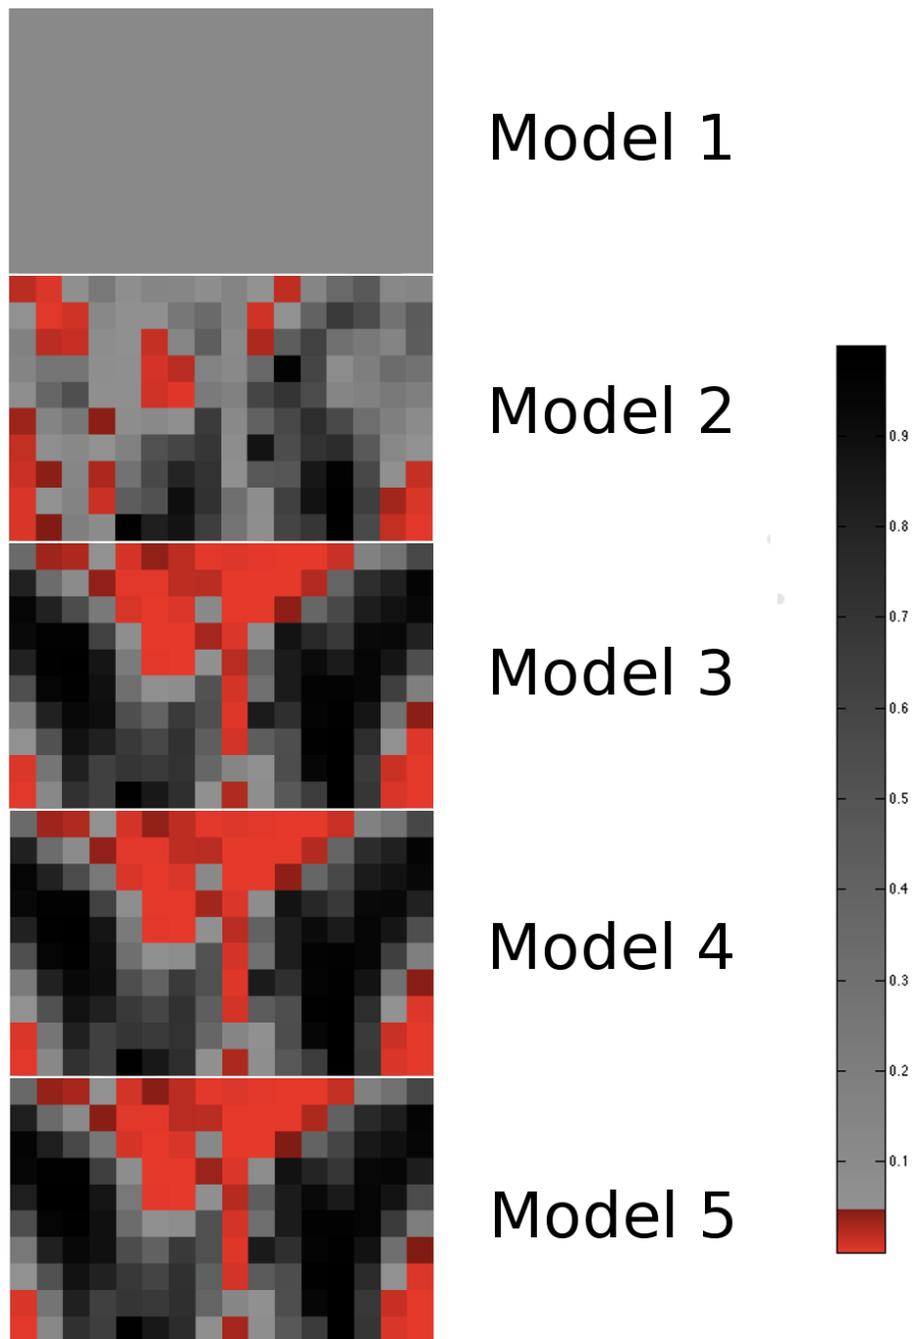


Figure 3.18: Group tensor difference p-values in Region 3

3.4 DISCUSSION

The results presented in this chapter demonstrate that LMEs were able to detect plausible and consistent group differences in white matter anatomy between a large sample of individuals with ASDs and comparable controls. In analyses involving the whole tensor information as the dependent variable, I detected significant differences in general diffusion properties bilaterally in the body of the splenium (Figure 3.14) and in the genu (Figure 3.16), while there was the predicted absence of group difference in the lateral ventricles (Figure 3.18). There was qualitative evidence of asymmetry in the patterns of group difference in the splenium and genu, which may prove to be an interesting area of future investigation.

In the analyses with FA as the dependent variable, I found significantly increased anisotropy in the control group relative to the participants with ASDs bilaterally in the splenium and the genu (Figures 3.8 and 3.10), and no evidence of differences in anisotropy in the lateral ventricles.

These findings are in broad agreement with those reported on the same sample by Alexander et al. (2007), who found decreased anisotropy in ASDs relative to control populations in both the genu and the splenium using an ROI-based analysis.

3.4.1 LMEs Compared To Traditional Approaches

Each set of LME-based analyses I performed was accompanied by parallel ROI-based (Model 1) and voxelwise (Model 2) analyses. The ROI-based analyses provide qualitatively different information than the LME-based analyses, so a direct comparison of the results isn't necessarily possible. The ROI-based approach provides a single overall indication of the presence of group difference over a whole region, while the LME-based approaches yield a distinct measure of group difference at each voxel in the region. In the FA analyses, Model 1 showed no overall significant difference between ASDs and control participants in the splenium or the ventricular area, while Models 3 – 5 revealed significant differences at the level of individual voxels in the densest white matter portions of the regions. And, in the genu, while Model 1 revealed an overall decrease in FA in ASDs relative to controls, Models 3 – 5 yielded a more

nuanced picture, with significant difference in anisotropy in the center of the genu, but less difference toward the exterior of the region. The comparison between Model 1 and Models 3 – 5 is much the same in the multivariate whole-tensor comparisons as in the FA analyses.

The Model 2 analyses provide the same sort of voxelwise quantitative information as Models 3 – 5, and it is here that we see the most interesting contrasts between traditional methods and LME-based approaches. In the FA analyses, in all three regions considered, the Model 2 analysis produced a less coherent picture of significant group difference. Models 3 – 5 showed consistent group differences in the densest white matter of each region. Model 2 showed virtually no areas of significant group difference in the genu, including the portion of the genu abutting the ventricular region (Region 3). In Region 1, Model 2 revealed a few voxels with significant group difference in the densest part of the splenium, but also showed surprising group differences in parts of the region outside of the splenium.

In the whole-tensor analyses, the LME-based approaches were more sensitive than the voxelwise analysis to overall group differences in diffusivity in the white matter tracts in Regions 1 and 2. In both cases, the voxels with significant group difference under Model 2 were a subset of those under Models 3 – 5, and showed less spatial coherence. In Region 3, Model 2 was not sensitive to group differences in the part of the region abutting the genu or in the septum pellucidum, and did show group differences in two voxels of the right lateral ventricle.

The results justify LME-based approaches as a useful alternative to traditional ROI or voxelwise analyses. Models 3 – 5 each showed greater and more cohesive areas of group difference than Model 2, both in univariate analyses of FA and multivariate analyses of the whole tensor. And each provides more nuanced information than Model 1. Statistically, the strength of LMEs relative to voxelwise t -tests in this context is that the LMEs borrow strength across the whole region to estimate variance components, while the voxelwise analyses require estimation of a distinct pooled variance at each voxel. Naturally, some of these pooled variance estimates will be too low, potentially leading to spurious findings of group difference, and some estimates will be too high, leading to voxelwise Type II errors.

3.4.2 Full Tensor Analyses Compared To FA Analyses

The standard approach to group comparisons of DTI data is to use ROI or voxelwise methods with one or more tensor-derived scalar indices as the dependent variable(s). I have shown the utility of LME-based approaches in this context by performing univariate analyses of FA. I have also performed a series of multivariate analyses comparing all six dimensions of diffusion information encoded in the tensors between groups. Multivariate and univariate approaches each have strengths and weaknesses, and it is worthwhile to discuss these in light of the results presented in this chapter.

We saw the main weakness of the univariate approach in Sections 2.1 and 3.3: any single tensor-derived scalar index necessarily entails a loss of information from the full six-dimensional diffusion tensor. Thus, it is entirely possible that genuine group differences could be overlooked by focusing attention on only one or a small number of tensor-derived scalar indices. The major strength of univariate analyses of chosen scalar indices is ease of interpretation. Several of the most popular tensor-derived indices (see Section 1.2.2.3) have simple interpretations in terms of the geometry of the underlying diffusion. Furthermore, the fact that analyses of FA and mean diffusivity in particular have become somewhat conventional is itself a recommendation for continuing to perform such analyses, in order to have a firm ground for comparing new results with previous literature.

The strengths and weaknesses of the multivariate approach are exactly complementary to those of scalar analyses. By analyzing all six transformed tensor components, it is possible to identify the location of group differences of any sort in diffusion properties captured by the tensor model, potentially including differences in anisotropy, diffusivity, shape and orientation of the diffusion ellipsoid. However, having identified an overall difference in multivariate means among the transformed tensor elements at a single voxel, the interpretation of this difference is non-trivial. While it would be a simple matter to follow up an overall significant multivariate comparison with six individual voxelwise comparisons of the transformed tensor elements, this would gain us little in terms of interpretability. As described in Section 2.1, the first transformed tensor element is proportional to mean diffusivity and is thus potentially directly interpretable to practitioners, but the remaining five transformed

elements have much more problematic interpretations. The second and third have a loose interpretation as anisotropy indices, but they are sensitive only to differences in amount of diffusion between cardinal directions of the scanner frame of reference. Thus, they are not rotationally invariant, meaning that a change of coordinate system would lead to a change in value of these indices. And the final three elements are the untransformed off-diagonal elements of the original tensor, and have no useful direct interpretation in terms of diffusion properties.

The tension between the strengths and weaknesses of univariate and multivariate approaches is evident in the results presented in this chapter. In each of the three regions, the voxels that showed significant group differences in anisotropy (univariate FA analyses) were a subset of those that showed significant multivariate group differences. That is, the multivariate approach was more sensitive to group differences in local diffusion properties. However, while we can directly interpret the FA results as indicating voxels where control participants had (usually) higher anisotropy than ASDs participants, and therefore likely more compact, well-organized and/or highly myelinated white matter fibers, we have no such simple interpretation for the voxels shown to have significant multivariate group differences. These observed differences may be due to anisotropy, but may just as well be due to diffusivity or tensor orientation. Nor is it meaningful to speak of the direction of difference in the multivariate analyses.

3.4.3 Homogeneous Compared To Heterogeneous Variance Models

One of the novelties of the data analytic approach in this chapter was the application of LME models with group heterogeneity in variance structure to DTI data. In the univariate analyses of FA as well as the multivariate analyses of the whole tensor data, I fit three different classes of LMEs. Model 3 assumed homogeneous variance between groups for both the within- and between-subject variance components, Model 4 assumed homogeneous variance for within-subject variance, but group heterogeneity in between-subjects variance, and Model 5 assumed group heterogeneity in both within- and between-subject variance structure.

The differences in the several p-value maps obtained from each of these models were visually quite subtle. In most cases, there was a slight attenuation of estimated group difference effects with increasing variance heterogeneity, with a small number of voxels crossing the threshold from significant to non-significant in the transition from Model 3 to Model 5.

Table 3.2: Likelihood ratio tests comparing FA analyses under Models 3–5

Comparison	Region 1			Region 2			Region 3		
	χ^2	df	p	χ^2	df	p	χ^2	df	p
Model 3 vs. Model 4	0.18	1	.67	6.00	1	.01	0.42	1	.52
Model 3 vs. Model 5	16.91	2	.0002	50.16	2	<.0001	14.78	2	.0006
Model 4 vs. Model 5	16.73	1	<.0001	4.15	1	<.0001	14.36	1	.0002

Table 3.3: Likelihood ratio tests comparing full tensor analyses under Models 3–5

Comparison	Region 1			Region 2			Region 3		
	χ^2	df	p	χ^2	df	p	χ^2	df	p
Model 3 vs. Model 4	17.40	6	.008	6.22	6	.40	1.80	6	.94
Model 3 vs. Model 5	280.58	12	<.0001	123.49	12	<.0001	228.49	12	<.0001
Model 4 vs. Model 5	263.18	6	<.0001	117.27	6	<.0001	226.68	6	<.0001

Tables 3.2 and 3.3 display the likelihood ratio comparisons between the three LME based models for the FA and full tensor analyses, respectively. In every region, for both univariate and multivariate analyses, Model 5 provided a significantly better fit to the data than either Models 4 or 3 according to likelihood ratio tests. Model 4 provided a significant improvement in fit to Model 3 in some cases and not in others. This suggests that incorporation of group heterogeneity in variance structure, and especially in residual variance, is justified, and that the results reported from Model 5 are the most reliable of those presented in this chapter.

3.4.4 Conclusions

Based on these results, I propose a hybrid two-stage approach to DTI group comparisons in future analyses. The first stage would be an LME-based multivariate analysis such as performed in Section 3.3. The goal of this stage would be to localize areas of group difference in overall diffusion properties. Once regions displaying overall group differences had been identified, the second-stage analysis would involve performing a series of univariate analyses of tensor-derived scalar indices within those regions in order to determine the nature of the group difference in an interpretable framework. This two-stage approach is analogous to the common statistical practice of comparing multivariate means with a MANOVA, and following significant omnibus MANOVA results with univariate ANOVA to more specifically isolate the group differences. One advantage of this approach over the analysis of scalar indices only is that the first-stage multivariate analysis provides some protection against inflation of Type I error due to multiple comparisons (in this case, comparisons of multiple scalar indices). The other major advantage is that, as we have seen in this chapter, the multivariate comparison may be sensitive to group differences that are not apparent in the given choice of tensor-derived scalars. These multivariate differences, while not directly interpretable on their own, may point the way for future investigation or for consideration of a greater variety of scalar indices in the regions involved.

4.0 GENERALIZED ICC IN NEUROIMAGING MODELS

My research involves the application of linear mixed effects models (LMEs) to the analysis of DTI datasets. In previous chapters, my focus was on using LMEs to perform group comparisons in order to address research questions involving the identification of white matter abnormalities associated with psychiatric disorders. In this chapter, I instead address the evaluation of reliability in DTI datasets.

The reliability investigations are connected with the previous chapters in that I investigate a generalized intraclass correlation coefficient (ICC) based on variance components estimated from LMEs. I present analytic results showing that the generalized ICC is robust to misspecification of the random effects structure of mixed models, and simulation results that demonstrate the generalized ICC is also robust to misspecification of the residual covariance structure. These properties make the generalized ICC an attractive choice for quantifying reliability in an exploratory model-building framework.

4.1 RELIABILITY

There are many different contexts in which quantifying reliability is important. However, the literature to date on reliability in DTI analyses has been fairly sparse, and focused on simple techniques. For instance, a few studies have assessed reliability in the sense of within-subject, between-scan reproducibility of a small number of scalar indices, using the coefficient of variation to quantify reproducibility[20, 36].

4.1.1 Local Spatial Reliability

In this chapter, I present an application of the generalized ICC for quantifying and comparing the local spatial reliability of tensor-derived scalar indices. As discussed in Section 1.2.2.3, investigators exploring white matter abnormalities with DTI typically use summary statistics calculated from one or more scalar indices (formally, scalar-valued tensor functionals) to quantify important aspects of white matter architecture. There are many such indices available, most of which fit into one of three broad classes: diffusivity indices (e.g. mean diffusivity, Frobenius norm), anisotropy indices (e.g. fractional anisotropy, relative anisotropy, volume ratio) and shape indices (e.g. Westin’s coordinates, tensor mode)[16, 31, 75, 95]. Although indices from these three categories are not independent (shape and anisotropy indices in particular are closely related), it is generally assumed that measures from each category reflect functionally distinct aspects of white matter neuroanatomy. For a given research question, it may be clear whether data analysis should focus on diffusivity, anisotropy, and/or shape. In many cases, it is appropriate to perform analyses involving more than one of these categories.

What is generally less clear is which representative of a given category should be used in a particular analysis. Fundamental properties of the indices can help inform this decision. For example, fractional anisotropy is more sensitive to variability at the low end of the anisotropy spectrum than is the volume ratio, and so might be more appropriate for analyzing cortical white matter microstructure and its abnormalities [45]. In practice, however, such properties seem to provide inadequate guidance for data analyses, and the choice of which scalar index to use within a given class appears to be driven more by external forces such as software availability and laboratory tradition than by intrinsic properties of the indices themselves or of the underlying research questions.

Empirical evidence of differences in reliability among scalar indices would provide additional practical guidance as to the choice of index for a given problem. In this context, I am interested in quantifying the local spatial reliability of scalar indices; that is, in quantifying how much variability in scalar indices is due to individual differences vs. how much is due to within-subject variability. Since within-subject variability in white matter characteris-

tics certainly swamps between-subject differences on the scale of the entire brain, I examine reliability in very localized regions (e.g., $3 \times 3 \times 3$ voxel).

A model-based approach to quantifying such reliability is potentially advantageous, in that it may allow us to calculate a reliability index in the presence of complicated fixed effects structures, multiple sources of between-subject variability, and spatially autocorrelated residuals. However, it is not immediately obvious how to generalize traditional reliability indices to a mixed effects model context, and how changes in the structure of such models would affect reliability estimates. I address these issues in this chapter. Although the focus here is on DTI, these reliability investigations are very generally applicable to other correlated data.

4.1.2 Other Applications

Although the emphasis in this chapter is on local spatial reliability, there are many other potential applications of a generalized approach to calculating reliability for neuroimaging data. These include assessment of test-retest reliability and inter-rater reliability. In the context of DTI, evaluation of test-retest reliability could be very useful in understanding the variance attached to individual scans of subjects, and might well be found to depend on equipment, gradient encoding scheme, and the population under consideration. Inter-rater reliability in this context could be used to quantify variability in scans due to different scanners, or different processing streams.

4.2 INTRACLASS CORRELATION COEFFICIENT BACKGROUND

The intraclass correlation coefficient (ICC) is a standard approach to inter-rater reliability and test-retest reliability (in the case of more than two timepoints)[66]. For instance, in the context of inter-rater reliability in which N judges provide ratings on n targets, the ICC represents the average within-judge correlation across targets.

4.2.1 ICC Definition

The ICC is typically defined in terms of the variance parameters in the one-way random effects ANOVA model as the ratio of the between-subject variance to the total variance. Suppose we have N subjects each with n observations. Then the model is:

$$\mathbf{Y}_s = \mu \mathbf{1}_n + \alpha_s \mathbf{1}_n + \boldsymbol{\epsilon}_s, \quad (4.1)$$

where \mathbf{Y}_s is the vector of observed responses for the s^{th} subject, μ is the population mean, α_s is a subject-level random intercept, and $\boldsymbol{\epsilon}_s$ is a vector of residual errors. It is typically assumed that $\alpha_s \sim N(0, \sigma_\alpha^2)$, $\boldsymbol{\epsilon}_s \sim N(\mathbf{0}, \tau^2 \mathbf{I}_n)$, and that α and $\boldsymbol{\epsilon}$ are independent. These assumptions induce a compound symmetric covariance structure on the observations: $\text{Var}(\mathbf{Y}_i) := \boldsymbol{\Sigma} = \tau^2 \mathbf{I}_n + \sigma_\alpha^2 \mathbf{J}_n$, where $\mathbf{J}_n = \mathbf{1}_n \mathbf{1}_n'$. In this context, the ICC, ρ , is defined as:

$$\rho = \frac{\sigma_\alpha^2}{\sigma_\alpha^2 + \tau^2}. \quad (4.2)$$

Lange, Jones and Pierpaoli (2004) used this model in DTI data analysis for the first time, leading to more efficient Empirical Bayes estimates of individual and population tensor fields[53].

4.2.2 ICC Estimation

There are several estimators of ρ . The classical ANOVA estimator is

$$\hat{\rho}_a = \frac{\text{BMS} - \text{WMS}}{\text{BMS} + (N - 1)\text{WMS}}, \quad (4.3)$$

where

$$\text{BMS} = \sum_{s=1}^N \frac{n}{N} (\bar{Y}_s - \bar{Y}_{..})^2$$

denotes the between-subject mean square,

$$\text{WMS} = \sum_{s=1}^N \sum_{t=1}^n \frac{1}{N(n-1)} (Y_{st} - \bar{Y}_s)^2$$

denotes the within-subject mean square, $\bar{Y}_s = \sum_{t=1}^n Y_{st}/n$ and $\bar{Y}_{..} = \sum_{s=1}^N \sum_{t=1}^n Y_{st}/(nN)$. $\hat{\rho}_a$ is probably the most widely described estimator (for instance, see Shrout and Fleiss

(1979)[86]). Equation 4.3 assumes a balanced design, but the estimator can be algebraically extended to accommodate unbalanced designs[27]. The ANOVA estimator has also been extended to the ANCOVA model with one covariate by Stanish and Taylor[88], but further generalizations in this direction are not straightforward.

Donner and Koval (1980)[28] derived the likelihood equation for ρ , which can be optimized numerically to obtain the maximum likelihood estimate,

$$\hat{\rho}_m = \min_{\rho} M(1 + \log \sigma^2 + \log 2\pi) + (M - N) \log(1 - \rho) + \sum_{s=1}^N \log(1 + (n_s - 1)\rho),$$

where $M = \sum_{s=1}^N n_s$. In the balanced data case ($n_s = n, \forall s$), $\hat{\rho}_a = \hat{\rho}_m$.

Another approach is to obtain ML or REML[35] estimates of the variance components $\hat{\sigma}_\alpha^2$ and $\hat{\tau}^2$ using standard techniques (i.e., $\hat{\sigma}_\alpha^2 = (\text{BMS} - \text{WMS})/n$ and $\hat{\tau}^2 = \text{WMS}$ for balanced cases, the EM algorithm for unbalanced and/or incomplete cases[25]) and construct a ‘variance components’ estimator, $\hat{\rho}_{vc} = \hat{\sigma}_\alpha^2 / (\hat{\sigma}_\alpha^2 + \hat{\tau}^2)$. One advantage of this approach is that it is applicable to models with arbitrarily more complicated fixed effects structures than that in Equation 4.1. Lange and Ryan (1989) corrected for “plug-in” effects of this type in LME models[55].

Although the ANOVA and variance components reliability estimators are sensible in many simple applications, none are designed or appropriate for more complex models involving correlated data, complex fixed effects structures, and covariance structures other than compound symmetry. Since models for neuroimaging data often exhibit these features, a more general approach is needed.

4.2.3 Generalized ICC

The alternative to the ANOVA-based formulation I develop is based on the proposal by Kistner and Muller (2004)[47], who defined a generalized intraclass correlation coefficient based on the covariance matrix of the observations, Σ :

$$\rho = \frac{[\mathbf{1}'_n \Sigma \mathbf{1}_n - \text{Tr}(\Sigma)]/n(n-1)}{\text{Tr}(\Sigma)/n}. \quad (4.4)$$

Equation 4.4 can be interpreted as the ratio of the average of the off-diagonal elements of Σ to the average of its diagonal elements or, in other words, the ratio of the average covariance among within-subject observations to the average variance. As an estimator of ρ , Kistner and Muller proposed replacing Σ in Equation 4.4 with its maximum likelihood estimate, namely the sample covariance matrix \mathbf{S} . They derived the exact distribution of $\hat{\rho}$ defined in this way.

Unlike the quantities described previously, the generalized ICC can be calculated for any covariance matrix Σ and thus is potentially applicable to more general random effects structures than that in Equation 4.1. In particular, $\hat{\rho}$ can be calculated from model-based estimates of $\hat{\Sigma}$ in mixed-effects models with arbitrary random effects and residual covariance structures. In this respect, the generalized ICC has much in common with the ideas of Generalizability Theory (G-Theory), which extends classical notions of test reliability to a much more general linear model-based context[85]. However, the quantities defined in G-Theory are dependent upon the use of relatively simple ANOVA-like models with independent random effects and i.i.d. residuals. On the other hand, the generalized ICC is applicable to any model-based framework in which a covariance matrix is estimable, including potentially non-linear models and models with non-normal errors.

Throughout the following sections, ρ will refer to the generalized ICC, Equation 4.4.

4.3 GENERALIZED ICC IN GROWTH CURVE MODELS

An immediate question is how an estimator $\hat{\rho}$ will behave in the context of mixed-effects models with more complicated random effects structures than that of the one-way random effects ANOVA (Equation 4.1). In this section, I consider this question in the context of growth curve models - an analytically tractable and important subset of mixed effects models, closely related to the examples considered in Chapter 2. In the following, I show that within the class of mixed-effects growth curve models with i.i.d. residuals, $\hat{\rho}$ is invariant to the number of random effects under ML estimation and approximately so under REML estimation. I also provide simulation results to demonstrate that $\hat{\rho}$ is robust to misspecification

of the residual covariance structure. These robustness properties indicate that $\hat{\rho}$ is a good candidate for quantifying reliability in the context of the models I have proposed for group comparisons of DTI data, such as those presented in Chapter 3.

4.3.1 Notation And Assumptions

Let \mathbf{Y}_s denote the vector of n observed responses for subject s , $s = 1, \dots, N$. Note that in the case of the DTI analyses performed in Chapter 3, \mathbf{Y}_s would be the vector of observed values of a scalar index over voxels in a region of interest. Let matrix $\mathbf{Z} = [\mathbf{z}_1 \ \mathbf{z}_2 \ \dots \ \mathbf{z}_p]$ be an $n \times p$ fixed effects design matrix with $1 \leq p < n$. Without loss of generality, assume that the columns of \mathbf{Z} form an orthonormal set of vectors in \mathbb{R}^n : $\mathbf{z}'_i \mathbf{z}_j = \delta_{ij}$, where δ_{ij} is Kronecker's δ , and assume that the first column of \mathbf{Z} corresponds to an intercept term: $\mathbf{z}_1 = (1/\sqrt{n})\mathbf{1}_n$. \mathbf{Z}_r will denote a random effects design matrix formed from the first r columns of \mathbf{Z} , $1 \leq r \leq p$. $\boldsymbol{\beta}$ will be a $p \times 1$ vector of fixed effects coefficients and \mathbf{b}_s an $r \times 1$ vector of random effects coefficients. I will assume $\mathbf{b} \sim N(\mathbf{0}, \mathbf{G})$ where $\mathbf{G}_{ij} = \sigma_{ij}$. The vector of residual errors, $\boldsymbol{\epsilon}_s$, is $p \times 1$, with $\boldsymbol{\epsilon} \sim N(\mathbf{0}, \tau^2 \mathbf{I}_n)$. With this notation, the model I am considering is

$$\mathbf{Y}_s = \mathbf{Z}\boldsymbol{\beta} + \mathbf{Z}_r \mathbf{b}_s + \boldsymbol{\epsilon}_s. \quad (4.5)$$

Note that $\text{span}(\mathbf{Z}_r) \subseteq \text{span}(\mathbf{Z})$, a necessary condition for growth curve models of this form [78]. To facilitate comparisons between models with varying numbers of random effects, I denote the variance components from the model with r random effects as $\sigma_{ij,r}$ and τ_r^2 , and the generalized ICC as ρ_r .

The variance of \mathbf{Y} for these models is given by $\boldsymbol{\Sigma} = \mathbf{Z}_r \mathbf{G} \mathbf{Z}'_r + \tau_r^2 \mathbf{I}_n$. $\hat{\boldsymbol{\Sigma}}$ is obtained by substituting $\hat{\mathbf{G}}$ for \mathbf{G} and $\hat{\tau}_r^2$ for τ_r^2 . I will denote the estimator of the generalized intraclass correlation coefficient obtained by substituting $\hat{\boldsymbol{\Sigma}}$ into Equation 4.4 by $\hat{\rho}_r$.

4.3.2 Preliminary Results

In the results below, I will use the following observation: for $j > 1$,

$$\sum_{i=1}^n z_{ji} = 0, \quad (4.6)$$

which follows from $\mathbf{z}_1 \perp \mathbf{z}_j$:

$$\begin{aligned} \mathbf{z}'_1 \mathbf{z}_j &= 0 \\ \mathbf{1}'_n \mathbf{z}_j &= 0 \\ \sum_{i=1}^n z_{ji} &= 0. \end{aligned}$$

With the notation and assumptions, we have the following lemmas needed for the main result:

Lemma 4.3.1.

$$\mathbf{1}'_n \mathbf{z}_i \mathbf{z}'_j \mathbf{1}_n = \begin{cases} n & \text{if } i = 1 = j \\ 0 & \text{otherwise} \end{cases}$$

Proof. There are two cases. Case 1: $i = j = 1$

$$\begin{aligned} \mathbf{1}'_n \mathbf{z}_1 \mathbf{z}'_1 \mathbf{1}_n &= \mathbf{1}'_n \frac{1}{n} \mathbf{J}_n \mathbf{1}_n \\ &= \frac{1}{n} n^2 \\ &= n. \end{aligned}$$

Case 2: $i > 1$ or $j > 1$. Assuming without loss of generality that $j > 1$, we have:

$$\begin{aligned} \mathbf{1}'_n \mathbf{z}_i \mathbf{z}'_j \mathbf{1}_n &= \sum_{k=1}^n \sum_{l=1}^n z_{ik} z_{jl} \\ &= \sum_{k=1}^n \left(z_{ik} \sum_{l=1}^n z_{jl} \right) \\ &= 0, \text{ by Equation 4.6.} \end{aligned}$$

□

Lemma 4.3.2.

$$\text{Tr}(\mathbf{z}_i \mathbf{z}_j') = \begin{cases} 0 & i \neq j \\ 1 & i = j \end{cases}$$

Proof. This follows from the orthonormality of the columns of \mathbf{Z} . □

Lemma 4.3.3.

$$\hat{\tau}_r^2 = \hat{\tau}_p^2 + \frac{1}{m} \sum_{j=r+1}^p \hat{\sigma}_{jj,p},$$

and for $1 \leq i \leq r$,

$$\hat{\sigma}_{ii,r} = \hat{\sigma}_{ii,p} - \frac{1}{m} \sum_{j=r+1}^p \hat{\sigma}_{jj,p},$$

$$\text{where } m = \begin{cases} n - r & \text{for ML estimation} \\ n - r + \frac{1}{N-1}(n - p) & \text{for REML estimation} \end{cases}$$

These results follow directly from Equations 19 and 20 in Lange and Laird (1989) [54].

4.3.3 Invariance Of The Generalized ICC To Number Of Random Effects

I am now prepared to state the main result of this section.

Theorem 4.3.4. *The estimator of ρ obtained by substituting the ML estimate of the within-subject covariance, $\hat{\Sigma}$, from a model of the class described in Equation 4.5 for Σ in Equation 4.4 is invariant to r , for $1 \leq r \leq p$.*

Proof. The proof of this result is given in Appendix A.1. □

This invariance of $\hat{\rho}$ to the number of random effects is specific to ML estimation of the model in Equation 4.5. Under REML estimation, $\hat{\rho}$ varies modestly with the number of random effects, attaining its minimum with the saturated random effects model, $r = p$. The following algebraic identity will be used several times in the proof of this result.

Lemma 4.3.5.

$$\left(n - r + \frac{1}{N-1}(n - p) \right)^{-1} = \frac{1}{(n - r)} - \frac{(n - p)}{(n - r)^2(N - 1) + (n - r)(n - p)}.$$

Proof.

$$\begin{aligned}
\frac{1}{n-r} - \frac{(n-p)}{(n-r)^2(N-1) + (n-r)(n-p)} &= \frac{1}{n-r} - \frac{\frac{(n-p)}{N-1}}{(n-r)^2 + \frac{(n-r)(n-p)}{N-1}} \\
&= \frac{1}{n-r} - \frac{\frac{n-p}{N-1}}{(n-r) \left[n-r + \frac{n-p}{N-1} \right]} \\
&= \frac{\left[n-r + \frac{n-p}{N-1} \right] - \frac{n-p}{N-1}}{(n-r) \left[n-r + \frac{n-p}{N-1} \right]} \\
&= \left(n-r + \frac{1}{N-1}(n-p) \right)^{-1}.
\end{aligned}$$

□

Theorem 4.3.6. *For any $1 \leq r < p$, $\hat{\rho}_r > \hat{\rho}_p$ where $\hat{\rho}$ is the estimator of ρ obtained by substituting the REML estimate of the within-subject covariance, $\hat{\Sigma}$, from a model of the class described in Equation 4.5 for Σ .*

Proof. The proof of this result is provided in Appendix A.2. □

Although $\hat{\rho}$ is not invariant to the number of random effects under REML estimation of $\hat{\Sigma}$ as it is under ML estimation, $\hat{\rho}_r - \hat{\rho}_p$ will be small under most realistic conditions. The numerator and denominator of $\hat{\rho}_r$ each differ from those of $\hat{\rho}_p$ by a multiple of $\sum_{j=r+1}^p \hat{\sigma}_{jj,p}$. In practice, the magnitudes of variance components $\hat{\sigma}_{jj,p}$ tend to decrease with increasing j , so the magnitude of this sum can be expected to be relatively small. Furthermore, the multiples of $\sum_{j=r+1}^p \hat{\sigma}_{jj,p}$ by which the numerators and denominators of $\hat{\rho}_r$ and $\hat{\rho}_p$ differ, $(n-p)((n-1)[(n-r)(N-1) + (n-p)])^{-1}$ and $(n-p)[(n-r)(N-1) + (n-p)]^{-1}$, respectively, decrease rapidly with N . Even at the modest value of $N = 10$, these terms have maxima of 0.026 (at $n = 3$, $p = 2$, $r = 1$) and 0.089 (at $n = 9$, $p = 2$, $r = 1$), respectively.

4.3.4 Autoregressive Error

I have shown that the generalized ICC is invariant to the number of random effects in a balanced growth model with i.i.d. errors under ML estimation and is modestly larger for a model with $r < p$ random effects than for the saturated random effects model with p random effects under REML estimation. An immediate question of interest is how ρ

behaves in models with possibly non-i.i.d. errors. This extension is particularly important the application of ρ to MRI data that is susceptible to noise from a variety of sources, leading to spatially autocorrelated errors.

It is apparent that ρ provides an intuitively appealing measure of intraclass correlation in one-way ANOVA designs with possibly non-i.i.d. errors. As an example, consider again the model in Equation 4.5, with the assumption that, instead of being independent, the errors stem from a first-order autoregressive (AR(1)) process. I use this particular example because it leads to a non-i.i.d. covariance structure with only two parameters that is a reasonable model in many situations with serially autocorrelated errors. Specifically, $\epsilon_s \sim N(\mathbf{0}, \Sigma_\epsilon)$, where the $(i, j)^{\text{th}}$ entry of Σ_ϵ is given by $\tau^2\theta^{|i-j|}$. Under this assumption, the variance of the observations takes the form of $\Sigma = \mathbf{Z}'_r \mathbf{G} \mathbf{Z}_r + \Sigma_\epsilon$, and

$$\rho_{\text{AR}(1)} = \frac{\hat{\sigma}_{11} - \frac{1}{n-1} \sum_{i=2}^r \hat{\sigma}_{ii} + n\tau^2 \frac{\sum_{j=1}^{n-1} w_j \theta^j}{\sum_{j=1}^{n-1} w_j}}{n\hat{\tau}^2 + \sum_{i=1}^r \hat{\sigma}_{ii}}, \quad (4.7)$$

with $w_j = (2(n-j))/(n-1)$. Thus, the numerator of Equation 4.7 incorporates a weighted average of the within-subject covariance terms, where the weights are given by the number of times each term appears in the covariance matrix. In practice, we can estimate $\rho_{\text{AR}(1)}$ by substituting ML or REML estimates of $\{\sigma_{ii}^2\}$, τ^2 and θ into Equation 4.7 to obtain $\hat{\rho}_{\text{AR}(1)}$.

Unfortunately, the generalized ICC is not invariant to changes in the residual covariance structure. This can be shown with simple empirical examples. Nor do there appear to be simple analytic relationships between the estimated variance components from the growth curve model with i.i.d. errors and the model with AR(1) errors that would make it possible to establish inequalities between $\hat{\rho}$ calculated from the two models.

4.3.5 Simulations

I conducted a series of simulations to evaluate and compare several different methods of estimating ρ for data generated according to the autoregressive model described in Section 4.3.4. The results of these simulations are given in Tables 4.1 – 4.3. Each dataset was simulated with $N = 100$, $n = 20$, $r = 1$ and $\sigma_{11} = 1$. The values of τ^2 are 0.33 for Table 4.1, 1.00 for Table 4.2 and 3.00 for Table 4.3. The value of θ varies by row for each table. The

columns of each table contain the mean percent bias and empirical 95% confidence intervals of the percent bias over 500 simulations of five different estimators of ρ :

- $\hat{\rho}_1$ is the estimator based on ML estimation under the assumption of i.i.d. residuals.
- $\hat{\rho}_2$ is the estimator based on ML estimation under the assumption of AR(1) residuals.
- $\hat{\rho}_3$ is the estimator based on the sample covariance matrix.
- $\hat{\rho}_4$ is the estimator based on REML estimation under the assumption of i.i.d. residuals.
- $\hat{\rho}_5$ is the estimator based on REML estimation under the assumption of AR(1) residuals.

Table 4.1: Mean percent bias and empirical 95% confidence intervals of percent bias of $\hat{\rho}$ for 500 simulations under an AR(1) process with $N = 100$, $n = 20$, $\sigma_{11} = 1$, and $\tau^2 = 0.33$.

θ	$\hat{\rho}_1$	$\hat{\rho}_2$	$\hat{\rho}_3$	$\hat{\rho}_4$	$\hat{\rho}_5$
0.00	*	*	*	*	*
0.25	-0.5 (-6.5, 5.0)	-0.5 (-6.5, 5.0)	-0.3 (-6.2, 5.2)	-0.3 (-6.2, 5.2)	-0.3 (-6.2, 5.2)
0.50	-0.7 (-6.5, 4.4)	-0.7 (-6.5, 4.4)	-0.5 (-6.2, 4.6)	-0.5 (-6.3, 4.6)	-0.5 (-6.3, 4.7)
0.75	-0.3 (-5.2, 4.2)	-0.4 (-5.4, 4.1)	-0.2 (-5.1, 4.4)	-0.2 (-5.0, 4.3)	-0.2 (-5.3, 4.2)

Table 4.2: Mean percent bias and empirical 95% confidence intervals of percent bias of $\hat{\rho}$ for 500 simulations under an AR(1) process with $N = 100$, $n = 20$, $\sigma_{11} = 1$, and $\tau^2 = 1.00$.

θ	$\hat{\rho}_1$	$\hat{\rho}_2$	$\hat{\rho}_3$	$\hat{\rho}_4$	$\hat{\rho}_5$
0.00	-1.2 (-14.2, 10.4)	-1.2 (-14.1, 10.4)	-0.7 (-13.8, 10.9)	-0.7 (-13.7, 10.9)	-0.7 (-13.6, 10.9)
0.25	-0.8 (-13.2, 10.7)	-0.8 (-13.4, 10.7)	-0.4 (-12.8, 11.1)	-0.4 (-12.7, 11.2)	-0.4 (-12.9, 11.1)
0.50	-1.1 (-12.4, 9.4)	-1.1 (-12.0, 9.4)	-0.6 (-11.9, 9.9)	-0.6 (-12.0, 9.8)	-0.7 (-11.6, 9.8)
0.75	-0.5 (-11.4, 9.5)	-0.5 (-11.0, 9.4)	-0.2 (-11.1, 9.8)	-0.2 (-11.1, 9.8)	-0.2 (-10.7, 9.7)

Table 4.3: Mean percent bias and empirical 95% confidence intervals of percent bias of $\hat{\rho}$ for 500 simulations under an AR(1) process with $N = 100$, $n = 20$, $\sigma_{11} = 1$, and $\tau^2 = 3.00$.

θ	$\hat{\rho}_1$	$\hat{\rho}_2$	$\hat{\rho}_3$	$\hat{\rho}_4$	$\hat{\rho}_5$
0.00	-1.3 (-22.6, 19.0)	-1.3 (-22.5, 19.0)	-0.5 (-21.8, 19.8)	-0.5 (-21.8, 19.9)	-0.5 (-21.8, 19.9)
0.25	-1.1 (-19.8, 18.6)	-1.1 (-19.7, 18.7)	-0.3 (-18.9, 19.6)	-0.3 (-19.1, 19.5)	-0.3 (-19.1, 19.5)
0.50	-1.4 (-20.4, 16.7)	-1.3 (-19.7, 16.2)	-0.7 (-19.7, 17.4)	-0.7 (-19.8, 17.5)	-0.7 (-19.1, 16.8)
0.75	-0.9 (-16.1, 13.0)	-0.8 (-14.0, 12.5)	-0.5 (-15.7, 13.7)	-0.4 (-15.5, 13.6)	-0.4 (-13.6, 12.9)

These simulations highlight several interesting results. First, all five estimators considered are slightly biased downward. Second, the two model-based estimators with ML

estimation ($\hat{\rho}_1$ and $\hat{\rho}_2$) have almost identical bias, and the sample covariance based estimator ($\hat{\rho}_3$) and the two model-based estimators with REML estimation ($\hat{\rho}_4$ and $\hat{\rho}_5$) have almost identical bias, which is less than that of $\hat{\rho}_1$ and $\hat{\rho}_2$. The variances of all the estimators are comparable, and increase with τ^2 and decrease slightly with θ . Interestingly, this decrease is slightly more marked for the AR(1) model-based estimators.

The results in this section suggest that the generalized ICC performs well under misspecification of the residual error structure, at least in the case of spatial autoregressive error. This is an additional robustness finding that supports the use of the generalized ICC in exploratory analyses of neuroimaging data. Spatially autocorrelated residuals are ubiquitous in these data, but the exact distribution is generally unknown, and may be difficult to model for computational reasons.

The model-based REML estimators performed approximately equivalently to the sample covariance-based estimator in these simulations. However, in data with more complicated fixed-effects structures than a simple global mean, it is much more sensible to model that mean and use model-based variance-covariance estimates as the basis for calculating the gICC than it would be to calculate gICC based on the simplistic model implicit in the sample covariance based estimator.

4.4 GENERALIZED ICC APPLICATION

In order to provide an empirical example of the use of the generalized ICC for neuroimaging data, I performed a series of data analyses designed to compare the local spatial reliability (Section 4.1.1) of different scalar indices. Specifically, I was interested in evaluating the reliability of three different measures of anisotropy: fractional anisotropy (FA, Equation 1.19), standardized relative anisotropy (sRA, Equation 1.23) and volume ratio (VR, Equation 1.24). All three of these indices are designed to measure the degree to which tissue at a given voxel is compact and coherently organized. While FA is the most widely used anisotropy index in the DTI literature, sRA has the elegant mathematical property of scaling linearly with the largest eigenvalue in tensors that reflect cylindrical symmetry, and VR has an

attractive direct geometrical interpretation as the ratio between the volume of the diffusion tensor and the volume of a sphere of equal total diffusivity.

The three indices each have a distinct sensitivity range to anisotropy: FA is most sensitive to differences among low levels of anisotropy, VR is most sensitive to differences among high levels of anisotropy, and sRA is in between[45]. It seems entirely possible that different indices may have different levels of spatial reliability depending on the region of the brain or type of tissue under consideration. However, these considerations have not generally been used to inform the choice of a particular index for a given analysis, nor have there been any empirical investigations of such regional differences in reliability to date.

4.4.1 Methods

I used the same data for this application as was used in Chapter 3. The sample consisted of 80 children and young adults with autism spectrum disorders, and 40 otherwise comparable controls. All of the participants in the sample were male. The DWI acquisition, preprocessing, diffusion tensor estimation and inter-subject registration were all performed as described in Section 3.1

I first created binary brain masks for each participant using the Brain Extraction Tool from the FSL software suite. I then multiplied the masks together to obtain an overall brain mask for the entire sample. I calculated the FA, sRA and VR at each voxel in the masked space for each participant.

I then fit a series of local LME models for each of the three anisotropy measures to create local reliability maps. The procedure was the same for all three indices; for simplicity, I will fix FA as the dependent variable in the following discussion. For each voxel v in the masked space, I extracted the FA values in the $3 \times 3 \times 3$ voxel neighborhood ($216mm^3$) centered at v for each subject. Then, if every voxel in the neighborhood was within the masked brain area, I fit the following model, closely analogous to the model described in Section 3.2.2.5:

$$FA_{ij} = \mu_j + g_i\gamma_j + (1 - g_i)(\delta_{0i} + \epsilon_{0ij}) + g_i(\delta_{1i} + \epsilon_{1ij}), \quad (4.8)$$

where FA_{ij} is the FA for subject i at voxel j , $j = 1, \dots, 27$, g_i is a group indicator ($g_i = 0$ if

subject i is in the control group, $g_i = 1$ if subject i is in the ASDs group), μ_j is an unknown fixed parameter reflecting the control mean at voxel j , γ_j is an unknown fixed parameter reflecting the ASDs group deviation from the control mean at voxel j , $\delta_0 \sim N(0, \sigma_0^2)$ and $\delta_1 \sim N(0, \sigma_1^2)$ are subject-specific random intercepts, and $\epsilon_0 \sim N(0, \tau_0^2)$ and $\epsilon_1 \sim N(0, \tau_1^2)$ are residual error terms. This model thus incorporates group heterogeneity in both the between-subject and the residual variance, which was justified in each analysis performed in Chapter 3.

Each of these models has closed-form likelihood estimation, as discussed extensively in Chapter 2, and, more specifically, Section 2.4.3. I estimated the model parameters using REML estimation (Section 2.3.6) obtaining in particular an estimated covariance matrix $\hat{\Sigma}$. This covariance matrix is block diagonal with blocks of two distinct types, one for control participants ($\hat{\Sigma}_0$) and one for participants with ASDs ($\hat{\Sigma}_1$). I then used these group-specific estimated covariance matrices to calculate group-specific generalized ICCs according to Equation 4.4, ρ_0 for the control group and ρ_1 for the ASDs group.

Note that this procedure was repeated for each anisotropy measure at every voxel in the interior of the masked space. Thus, the analysis resulted in multiple reliability maps, one for each group for each scalar index. I also calculated maps containing the ratios of the generalized ICC estimates obtained for the three scalar indices in order to compare the regional reliability between different indices.

4.4.2 Results

Figure 4.1 displays the generalized ICC maps for FA for control participants, while Figures 4.2 and 4.3 display the control ICC maps for sRA and VR, respectively. The maps for ASDs participants are contained in Figures 4.4 – 4.6.

In order to compare the local reliability between scalar indices and groups, it is more instructive to examine a single slice. Figure 4.7 displays the midaxial slice of the gICC maps for FA, sRA and VR for control and ASDs participants. In this figure, the reliability maps for FA and sRA are nearly indistinguishable. The reliability maps for VR are also very similar to those for FA and sRA, but there is some suggestion of generally decreased reliability,

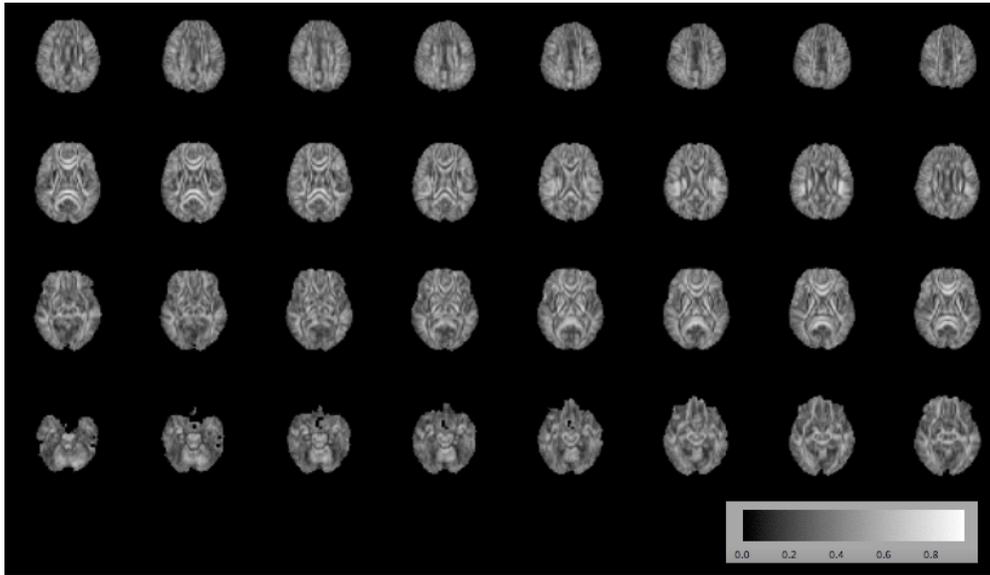


Figure 4.1: gICC maps for FA, control participants

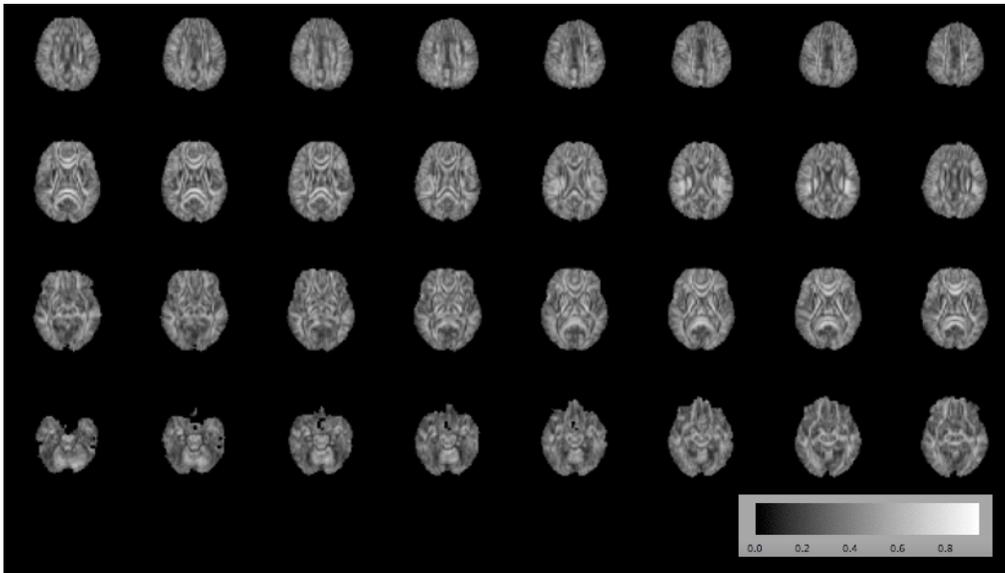


Figure 4.2: gICC maps for sRA, control participants

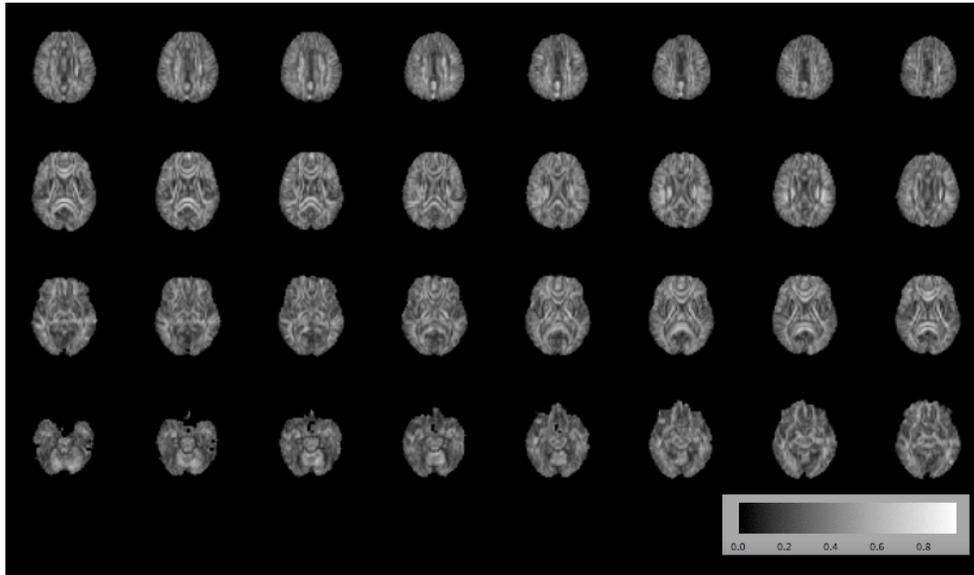


Figure 4.3: gICC maps for VR, control participants

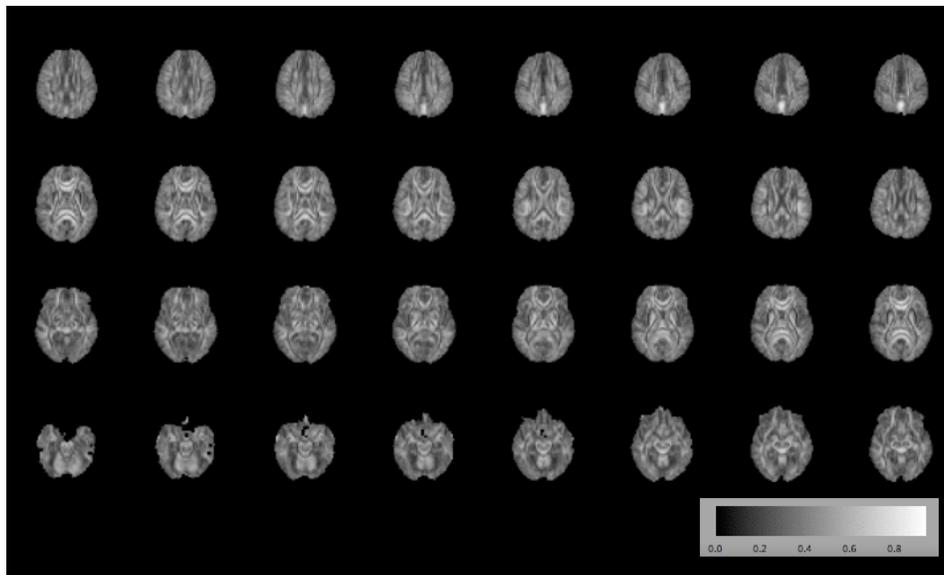


Figure 4.4: gICC maps for FA, ASDs participants

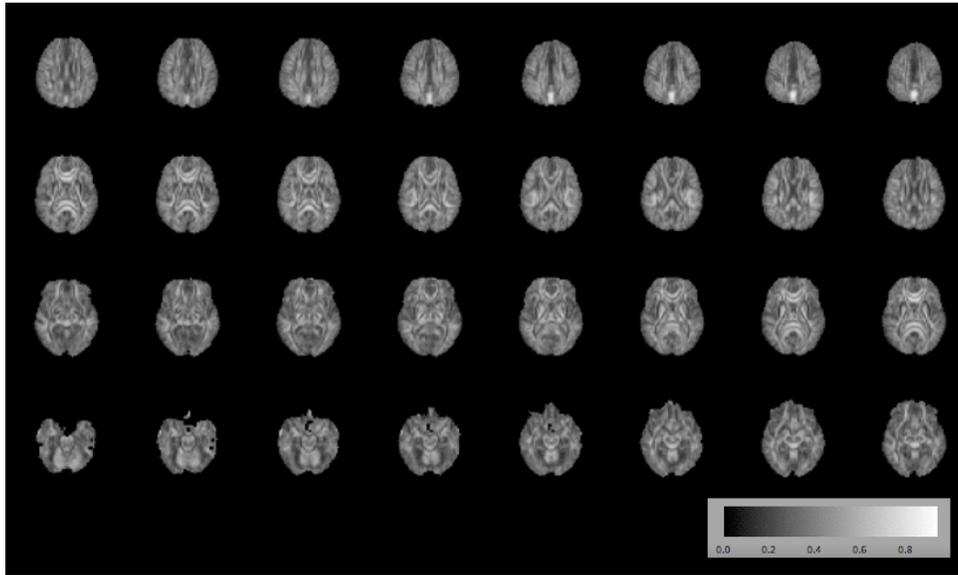


Figure 4.5: gICC maps for sRA, ASDs participants

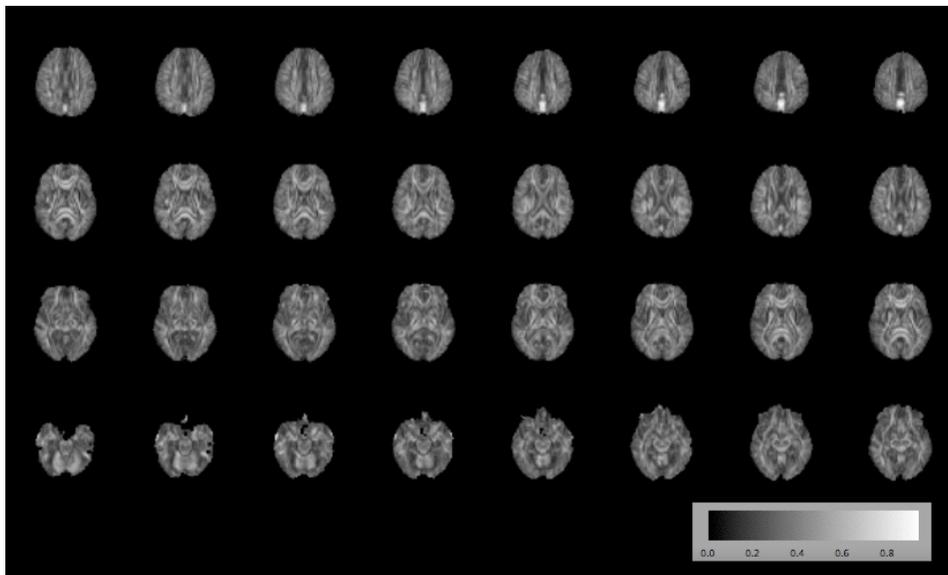


Figure 4.6: gICC maps for VR, ASDs participants

except in the frontal cortical portion of the figure. Table 4.4 presents summary statistics for reliability in the midaxial slice by group and anisotropy index. FA and sRA have very similar average reliability, as do controls and ASDs participants, although the reliability for ASDs participants is uniformly slightly lower than that for controls. However, it should be noted that one of the strengths of the approach in this chapter is that it allows us to obtain a more nuanced, local picture of reliability than is traditionally available through simple summary statistics.

In all six panels, there is a characteristic pattern of decreased relative reliability in the medial portions of major white matter tracts. These decreases appear as dark lines through the middle of otherwise high-reliability areas. This is due to the relative contributions of between- and within-subject variance to the total variance in the densest parts of major white matter structures compared to the periphery. I discuss this phenomenon in detail in Section 4.5.

Table 4.4: Mean and standard deviation of local $\hat{\rho}$ in the midaxial slice

	Control		ASDs	
Index	Mean	S.D.	Mean	S.D.
FA	0.413	0.175	0.407	0.175
sRA	0.408	0.176	0.402	0.172
VR	0.373	0.167	0.368	0.164

Figure 4.8 displays the pairwise percentage difference in gICC for the midaxial slice among FA, sRA and VR for the control group. The top pane in the figure shows the ratios of local reliability for sRA compared to FA. It appears from this figure that FA has higher reliability (light areas) in the densest callosal white matter regions, while sRA has higher reliability (dark areas) in the midbrain white matter structures lateral to the corpus callosum, apparently including the superior longitudinal fasciculus and the corticospinal tract. These areas show differences in reliability on the order of about 5%, indicating fairly subtle differences in reliability between FA and sRA.

The middle pane of Figure 4.8 shows the reliability of FA as a percentage of that for VR.

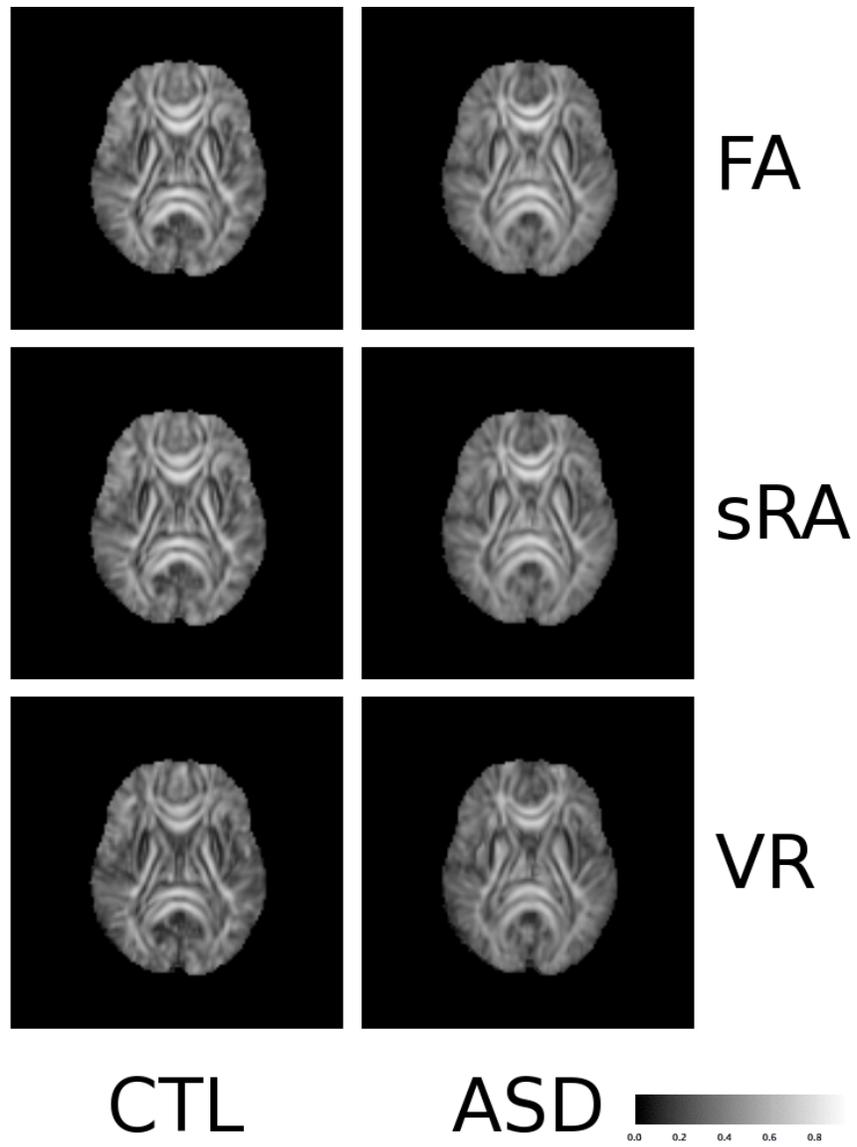


Figure 4.7: gICC maps comparison, midaxial slice

In this image, FA seems to have higher reliability (light areas) in most of the major white matter structures, while VR has higher reliability (dark areas) in some cortical areas and along the boundaries between different types of tissues. There is a particularly distinctive increase in reliability for VR relative to FA in the frontal cortex. The decreased reliability for FA relative to VR in the frontal region is on the order of 10 – 15%, while the increases in reliability for FA in the major white matter tracts are in the range of 15 – 20%. The areas that show the greatest increase in reliability for FA relative to VR are the fluid-filled ventricles, where the reliability of FA is as high as 60% higher than that of VR. However, as between-subject diffusivity differences in these regions are not apparently interpretable, it seems likely that this is an artifact of very low baseline reliability in this region for VR.

The comparison between VR and sRA reliability (bottom pane of the figure) is much the same as that between VR and FA, as would be expected due to the similarity of sRA and FA.

4.5 CONCLUSIONS

I have shown that the generalized ICC ρ is an intuitive and useful measure of reliability in mixed effects growth curve models. It is robust to misspecification of the random effects model under assumptions of balanced data and i.i.d. errors, in that it is invariant to the number of random effects under ML estimation and approximately so under REML estimation. Based on the simulation results in Section 4.3.5, it is also robust to misspecification of the residual covariance structure. These properties make ρ an attractive candidate to quantify reliability in an exploratory model-building framework in which random effects misspecification is entirely possible.

In an example with a sample of DTI data from an autism study, I have shown that the generalized ICC can be used to compare the reliability of various tensor-derived scalar indices. In my analysis of the local spatial reliability of anisotropy indices in a sample of 80 participants with ASDs and 40 controls, it appeared that FA and sRA had comparable reliability throughout the brain, while VR has substantially less reliability than either in

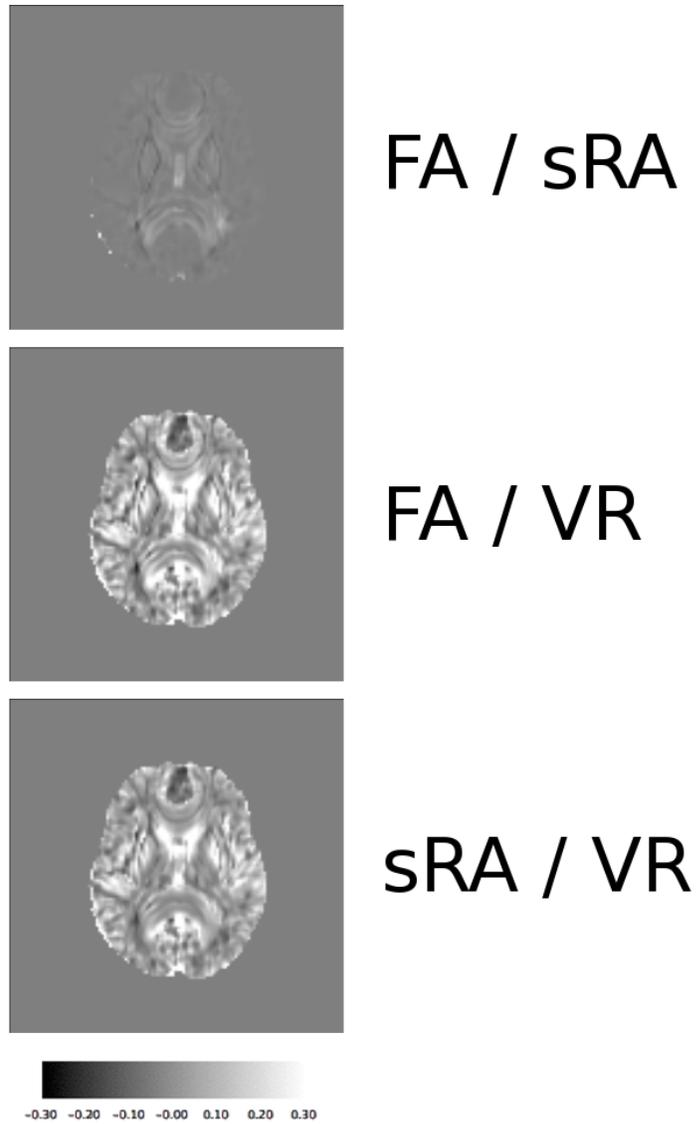


Figure 4.8: gICC ratio maps, midaxial slice

major white matter tracts. This would seem to indicate that FA or sRA would be statistically preferable to VR as an anisotropy index. However, in analyses involving substantial amounts of cortical white matter, especially in the frontal cortex, VR might prove to be a more reliable index than FA or sRA.

The reliability maps themselves contain an important lesson about DTI data analyses. Under the classical ICC model, the local spatial reliability at each voxel would represent the proportion of variance accounted for by between-subjects variation, and this is a useful framework in which to view the generalized ICC results as well. Thus, ρ can be taken as an index of the extent to which variability in a measurement at that voxel represents some true individual anatomical characteristic rather than noise. The reliability maps in Figure 4.7 show generally high reliability in major white matter tracts, which is to be expected for anisotropy indices.

However, the characteristic pattern of decreased reliability in the medial areas of these structures is perhaps more surprising. The explanation for this finding is that all scanned subjects tend to have comparably high anisotropy in these densest portions of the white matter structures. Therefore, the between-subject variance is lower here relative to the within-subject variance, which is more uniform throughout the region.

The implication is that between-subject differences and, by extension, between-group differences will be more difficult to find in the central portion of major white matter structures than differences outside of the dark stripes seen on Figure 4.7. That is, all else being equal, a DTI experiment that is adequately powered to detect group differences of a certain magnitude in the high reliability areas of the anisotropy maps will not necessarily be able to detect differences of the same magnitude in the center of the major white matter tracts. As Fleiss (1986) observed, if a sample size of n^* is required to have power $1 - \beta$ to detect an effect of size δ in the population for a variable measured without error, then a sample of size $n = n^*/\rho$ is required to have the same power to detect the same sized effect for a variable with reliability $\rho < 1$ [32].

It is important to remember that reliability is only one consideration in choosing a scalar index - the validity of the quantity for any specific analytic purpose must be established separately. Post-mortem, animal, phantom and simulation studies can all contribute to our

understanding of the validity of DTI-derived quantities, as can finding replicable results in clinical populations (convergent validity). However, since reliability is a prerequisite for validity, evaluations of reliability constitute an important first step, and one which has been largely absent from the literature to date. While the application in this chapter centered on local spatial reliability, the methods are more generally applicable to the quantification of test-retest or inter-rater (inter-scanner or inter-processing stream) reliability as well. Through the example in this chapter, we have seen that the reliability of anisotropy indices in the brain is not necessarily a simple matter, and could have profound implications for the design and interpretation of future DTI experiments.

APPENDIX

PROOFS OF THEOREMS IN CHAPTER 4

A.1 PROOF OF THEOREM 4.3.4

I will prove that $\hat{\rho}_r = \hat{\rho}_p$ for arbitrary r , $1 \leq r < p$.

First, expanding $\hat{\rho}_r$ yields:

$$\begin{aligned}
 \hat{\rho}_r &= \frac{[\mathbf{1}'_n \hat{\Sigma} \mathbf{1}_n - \text{Tr}(\hat{\Sigma})]/n(n-1)}{\text{Tr}(\hat{\Sigma})/n} \\
 &= \frac{\mathbf{1}'_n (\mathbf{Z}_r \hat{\mathbf{G}} \mathbf{Z}'_r + \hat{\tau}_r^2 \mathbf{I}_n) \mathbf{1}_n - \text{Tr}(\mathbf{Z}_r \hat{\mathbf{G}} \mathbf{Z}'_r + \hat{\tau}_r^2 \mathbf{I}_n)}{(n-1) \text{Tr}(\mathbf{Z}_r \hat{\mathbf{G}} \mathbf{Z}'_r + \hat{\tau}_r^2 \mathbf{I}_n)} \\
 &= \frac{\mathbf{1}'_n \mathbf{Z}_r \hat{\mathbf{G}} \mathbf{Z}'_r \mathbf{1}_n - \text{Tr}(\mathbf{Z}_r \hat{\mathbf{G}} \mathbf{Z}'_r) + \mathbf{1}'_n \hat{\tau}_r^2 \mathbf{I}_n \mathbf{1}_n - \text{Tr}(\hat{\tau}_r^2 \mathbf{I}_n)}{(n-1) \text{Tr}(\mathbf{Z}_r \hat{\mathbf{G}} \mathbf{Z}'_r + \hat{\tau}_r^2 \mathbf{I}_n)} \\
 &= \frac{\mathbf{1}'_n \mathbf{Z}_r \hat{\mathbf{G}} \mathbf{Z}'_r \mathbf{1}_n - \text{Tr}(\mathbf{Z}_r \hat{\mathbf{G}} \mathbf{Z}'_r)}{(n-1) \text{Tr}(\mathbf{Z}_r \hat{\mathbf{G}} \mathbf{Z}'_r + \hat{\tau}_r^2 \mathbf{I}_n)}, \text{ as } \mathbf{1}'_n \hat{\tau}_r^2 \mathbf{I}_n \mathbf{1}_n = \text{Tr}(\hat{\tau}_r^2 \mathbf{I}_n) \\
 &= \frac{\mathbf{1}'_n (\sum_{i,j=1}^r \hat{\sigma}_{ij} \mathbf{Z}_i \mathbf{Z}'_j) \mathbf{1}_n - \text{Tr}(\sum_{i,j=1}^r \hat{\sigma}_{ij} \mathbf{Z}_i \mathbf{Z}'_j)}{(n-1) (\text{Tr}(\sum_{i,j=1}^r \hat{\sigma}_{ij} \mathbf{Z}_i \mathbf{Z}'_j) + n \hat{\tau}_r^2)} \\
 &= \frac{n \hat{\sigma}_{11,r} - \sum_{i=1}^r \hat{\sigma}_{ii,r}}{(n-1) (n \hat{\tau}_r^2 + \sum_{i=1}^r \hat{\sigma}_{ii,r})}, \text{ by Lemmas 4.3.1 and 4.3.2} \\
 &= \frac{\hat{\sigma}_{11,r} - \frac{1}{n-1} \sum_{i=2}^r \hat{\sigma}_{ii,r}}{n \hat{\tau}_r^2 + \sum_{i=1}^r \hat{\sigma}_{ii,r}}
 \end{aligned}$$

Similarly,

$$\hat{\rho}_p = \frac{\hat{\sigma}_{11,p} - \frac{1}{n-1} \sum_{i=2}^p \hat{\sigma}_{ii,p}}{n \hat{\tau}_p^2 + \sum_{i=1}^p \hat{\sigma}_{ii,p}}.$$

To establish the equality of $\hat{\rho}_r$ and $\hat{\rho}_p$, I will separately establish the equality of the numerators and denominators, making repeated reference to Lemma 4.3.3. First, the denominators:

$$\begin{aligned}
n\hat{\tau}_r^2 + \sum_{i=1}^r \hat{\sigma}_{ii,r} &= n \left(\hat{\tau}_p^2 + \frac{1}{n-r} \sum_{j=r+1}^p \hat{\sigma}_{jj,p} \right) + \sum_{i=1}^r \left(\hat{\sigma}_{ii,p} - \frac{1}{n-r} \sum_{j=r+1}^p \hat{\sigma}_{jj,p} \right) \\
&= n\hat{\tau}_p^2 + \sum_{i=1}^r \hat{\sigma}_{ii,p} + \frac{n}{n-r} \sum_{j=r+1}^p \hat{\sigma}_{jj,p} - \frac{r}{n-r} \sum_{j=r+1}^p \hat{\sigma}_{jj,p} \\
&= n\hat{\tau}_p^2 + \sum_{i=1}^p \hat{\sigma}_{ii,p}
\end{aligned}$$

And the numerators:

$$\begin{aligned}
\hat{\sigma}_{11,r} - \frac{1}{n-1} \sum_{i=2}^r \hat{\sigma}_{ii,r} &= \left(\hat{\sigma}_{11,p} - \frac{1}{n-r} \sum_{i=r+1}^p \hat{\sigma}_{ii,p} \right) - \frac{1}{n-1} \left(\sum_{i=2}^r \hat{\sigma}_{ii,p} - \frac{r-1}{n-r} \sum_{j=r+1}^p \hat{\sigma}_{ii,p} \right) \\
&= \hat{\sigma}_{11,p} - \frac{1}{n-1} \sum_{i=2}^r \hat{\sigma}_{ii,p} + \left(\frac{r-1}{(n-1)(n-r)} - \frac{1}{n-r} \right) \sum_{j=r+1}^p \hat{\sigma}_{jj,p} \\
&= \hat{\sigma}_{11,p} - \frac{1}{n-1} \sum_{i=2}^r \hat{\sigma}_{ii,p} + \left(\frac{(r-1) - (n-1)}{(n-1)(n-r)} \right) \sum_{j=r+1}^p \hat{\sigma}_{jj,p} \\
&= \hat{\sigma}_{11,p} - \frac{1}{n-1} \sum_{i=2}^p \hat{\sigma}_{ii,p}
\end{aligned}$$

□

A.2 PROOF OF THEOREM 4.3.6

In the proof of Theorem 4.3.4, we observed that

$$\hat{\rho}_r = \frac{\hat{\sigma}_{11,r} - \frac{1}{n-1} \sum_{i=2}^r \hat{\sigma}_{ii,r}}{n\hat{\tau}_r^2 + \sum_{i=1}^r \hat{\sigma}_{ii,r}}$$

and

$$\hat{\rho}_p = \frac{\hat{\sigma}_{11,p} - \frac{1}{n-1} \sum_{i=2}^p \hat{\sigma}_{ii,p}}{n\hat{\tau}_p^2 + \sum_{i=1}^p \hat{\sigma}_{ii,p}}$$

To establish the inequality $\hat{\rho}_r > \hat{\rho}_p$, I will separately establish inequalities for the numerators and denominators of these expressions. Starting with the numerators, Lemmas 4.3.3 and 4.3.5 yield

$$\begin{aligned}
& \hat{\sigma}_{11,r} - \frac{1}{n-1} \sum_{i=2}^r \hat{\sigma}_{ii,r} = \\
& \hat{\sigma}_{11,p} - \left(n - r + \frac{1}{N-1}(n-p) \right)^{-1} \sum_{j=r+1}^p \hat{\sigma}_{jj,p} \\
& - \frac{1}{n-1} \sum_{i=2}^r \left[\hat{\sigma}_{ii,p} - \left(n - r + \frac{1}{N-1}(n-p) \right)^{-1} \sum_{j=r+1}^p \hat{\sigma}_{jj,p} \right] \\
& = \hat{\sigma}_{11,p} - \frac{1}{n-1} \sum_{i=2}^r \hat{\sigma}_{ii,p} - \left(n - r + \frac{1}{N-1}(n-p) \right)^{-1} \sum_{j=r+1}^p \hat{\sigma}_{jj,p} \\
& + \frac{r-1}{n-1} \left(n - r + \frac{1}{N-1}(n-p) \right)^{-1} \sum_{j=r+1}^p \hat{\sigma}_{jj,p} \\
& = \hat{\sigma}_{11,p} - \frac{1}{n-1} \sum_{i=2}^r \hat{\sigma}_{ii,p} \\
& + \left(\frac{r-1}{n-1} - 1 \right) \left(n - r + \frac{1}{N-1}(n-p) \right)^{-1} \sum_{j=r+1}^p \hat{\sigma}_{jj,p} \\
& = \hat{\sigma}_{11,p} - \frac{1}{n-1} \sum_{i=2}^r \hat{\sigma}_{ii,p} + \frac{r-n}{n-1} \left(\frac{1}{n-r} \right) \sum_{j=r+1}^p \hat{\sigma}_{jj,p} \\
& - \frac{r-n}{n-1} \left(\frac{(n-p)}{(n-r)^2(N-1) + (n-r)(n-p)} \right) \sum_{j=r+1}^p \hat{\sigma}_{jj,p} \\
& = \hat{\sigma}_{11,p} - \frac{1}{n-1} \sum_{i=2}^r \hat{\sigma}_{ii,p} - \frac{n-r}{(n-r)(n-1)} \sum_{j=r+1}^p \hat{\sigma}_{jj,p} \\
& + \frac{(n-r)(n-p)}{(n-1)[(n-r)^2(N-1) + (n-r)(n-p)]} \sum_{j=r+1}^p \hat{\sigma}_{jj,p} \\
& = \hat{\sigma}_{11,p} - \frac{1}{n-1} \sum_{i=2}^r \hat{\sigma}_{ii,p} \frac{n-p}{(n-1)[(n-r)(N-1) + (n-p)]} \sum_{j=r+1}^p \hat{\sigma}_{jj,p} \\
& > \hat{\sigma}_{11,p} - \frac{1}{n-1} \sum_{i=2}^r \hat{\sigma}_{ii,p}
\end{aligned}$$

Thus, the numerator of $\hat{\rho}_r$ is strictly greater than that of $\hat{\rho}_p$. Next, considering the denomi-

nators,

$$\begin{aligned}
n\hat{\tau}_r^2 + \sum_{i=1}^r \hat{\sigma}_{ii,r} &= \\
& n \left[\hat{\tau}_p^2 + \left(n - r + \frac{1}{N-1}(n-p) \right)^{-1} \sum_{j=r+1}^p \hat{\sigma}_{jj,p} \right] \\
& + \sum_{i=1}^r \left[\hat{\sigma}_{ii,p} - \left(n - r + \frac{1}{N-1}(n-p) \right)^{-1} \sum_{j=r+1}^p \hat{\sigma}_{jj,p} \right] \\
& = n\hat{\tau}_p^2 + \sum_{i=1}^r \hat{\sigma}_{ii,p} + (n-r) \left(n - r + \frac{1}{N-1}(n-p) \right)^{-1} \sum_{j=r+1}^p \hat{\sigma}_{jj,p} \\
& = n\hat{\tau}_p^2 + \sum_{i=1}^r \hat{\sigma}_{ii,p} + \frac{n-r}{n-r} \sum_{j=r+1}^p \hat{\sigma}_{jj,p} \\
& \quad - \frac{(n-r)(n-p)}{(n-r)^2(N-1) + (n-r)(n-p)} \sum_{j=r+1}^p \hat{\sigma}_{jj,p} \\
& = n\hat{\tau}_p^2 + \sum_{i=1}^r \hat{\sigma}_{ii,p} - \frac{n-p}{(n-r)(N-1) + n-p} \sum_{j=r+1}^p \hat{\sigma}_{jj,p} \\
& < n\hat{\tau}_p^2 + \sum_{i=1}^r \hat{\sigma}_{ii,p}
\end{aligned}$$

So the denominator of $\hat{\rho}_r$ is strictly less than that of $\hat{\rho}_p$. This inequality, together with the inequality between the numerators, yields the result: $\hat{\rho}_r > \hat{\rho}_p$. \square

BIBLIOGRAPHY

- [1] O. Abe, Y. Masutani, S. Aoki, H. Yamasue, H. Yamada, K. Kasai, H. Mori, N. Hayashi, T. Masumoto, and K. Ohtomo. Topography of the human corpus callosum using diffusion tensor tractography. *Journal of Computer Assisted Tomography*, 28(4):533–9, 2004.
- [2] A. L. Alexander, J. E. Lee, M. Lazar, R. Boudos, M. B. Dubray, T. R. Oakes, J. N. Miller, J. Lu, E. K. Jeong, W. M. McMahon, E. D. Bigler, and J. E. Lainhart. Diffusion tensor imaging of the corpus callosum in autism. *Neuroimage*, 34(1):61–73, 2007.
- [3] D. C. Alexander, C. Pierpaoli, P. J. Basser, and J. C. Gee. Spatial transformations of diffusion tensor magnetic resonance images. *IEEE Transactions on Medical Imaging*, 20(11):1131–9, 2001.
- [4] T. W. Anderson. Asymptotically efficient estimation of covariance matrices with linear structure. *The Annals of Statistics*, 1(1):135–141, 1973.
- [5] S. Aoki, N. K. Iwata, Y. Masutani, M. Yoshida, O. Abe, Y. Ugawa, T. Masumoto, H. Mori, N. Hayashi, H. Kabasawa, S. Kwak, S. Takahashi, S. Tsuji, and K. Ohtomo. Quantitative evaluation of the pyramidal tract segmented by diffusion tensor tractography: Feasibility study in patients with amyotrophic lateral sclerosis. *Radiation Medicine*, 23(3):195–9, 2005.
- [6] Y. Assaf, D. Ben-Bashat, J. Chapman, S. Peled, I. E. Biton, M. Kafri, Y. Segev, T. Hendler, A. D. Korczyn, M. Graif, and Y. Cohen. High b-value q-space analyzed diffusion-weighted MRI: Application to multiple sclerosis. *Magnetic Resonance in Medicine*, 47(1):115–26, 2002.
- [7] R. Bammer. Basic principles of diffusion-weighted imaging. *European Journal of Radiology*, 45(3):169–84, 2003.
- [8] N. Barnea-Goraly, H. Kwon, V. Menon, S. Eliez, L. Lotspeich, and A. L. Reiss. White matter structure in autism: Preliminary evidence from diffusion tensor imaging. *Biological Psychiatry*, 55(3):323–6, 2004.
- [9] T. R. Barrick and C. A. Clark. Singularities in diffusion tensor fields and their relevance in white matter fiber tractography. *Neuroimage*, 22(2):481–91, 2004.

- [10] T. R. Barrick, I. N. Lawes, C. E. Mackay, and C. A. Clark. White matter pathway asymmetry underlies functional lateralization. *Cerebral Cortex*, 17(3):591–8, 2007.
- [11] P. J. Basser and D. K. Jones. Diffusion-tensor MRI: Theory, experimental design and data analysis - a technical review. *NMR in Biomedicine*, 15(7-8):456–67, 2002.
- [12] P. J. Basser, J. Mattiello, and D. LeBihan. Estimation of the effective self-diffusion tensor from the NMR spin echo. *Journal of Magnetic Resonance, Series B*, 103(3):247–54, 1994.
- [13] P. J. Basser and S. Pajevic. Statistical artifacts in diffusion tensor MRI (DT-MRI) caused by background noise. *Magnetic Resonance in Medicine*, 44(1):41–50, 2000.
- [14] P. J. Basser and S. Pajevic. A normal distribution for tensor-valued random variables: Applications to diffusion tensor MRI. *IEEE Transactions on Medical Imaging*, 22(7):785–94, 2003.
- [15] P. J. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi. In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*, 44(4):625–32, 2000.
- [16] P. J. Basser and C. Pierpaoli. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *Journal of Magnetic Resonance, Series B*, 111(3):209–19, 1996.
- [17] C. Beaulieu. The basis of anisotropic water diffusion in the nervous system - a technical review. *NMR in Biomedicine*, 15(7-8):435–55, 2002.
- [18] C. Beaulieu, C. Plewes, L. A. Paulson, D. Roy, L. Snook, L. Concha, and L. Phillips. Imaging brain connectivity in children with diverse reading ability. *Neuroimage*, 25(4):1266–71, 2005.
- [19] D. Ben Bashat, V. Kronfeld-Duenias, D. A. Zachor, P. M. Ekstein, T. Hendler, R. Tarasch, A. Even, Y. Levy, and L. Ben Sira. Accelerated maturation of white matter in young children with autism: A high b-value DWI study. *Neuroimage*, 37(1):40–7, 2007.
- [20] D. Bonekamp, L. M. Nagae, M. Degaonkar, M. Matson, W. M. Abdalla, P. B. Barker, S. Mori, and A. Horska. Diffusion tensor imaging in children and adolescents: Reproducibility, hemispheric, and age-related differences. *Neuroimage*, 34(2):733–42, 2007.
- [21] O. Ciccarelli, T. E. Behrens, D. R. Altmann, R. W. Orrell, R. S. Howard, H. Johansen-Berg, D. H. Miller, P. M. Matthews, and A. J. Thompson. Probabilistic diffusion tractography: A potential tool to assess the rate of disease progression in amyotrophic lateral sclerosis. *Brain*, 129(Pt 7):1859–71, 2006.
- [22] L. Concha, C. Beaulieu, and D. W. Gross. Bilateral limbic diffusion abnormalities in unilateral temporal lobe epilepsy. *Annals of Neurology*, 57(2):188–96, 2005.

- [23] L. Concha, D. W. Gross, B. M. Wheatley, and C. Beaulieu. Diffusion tensor imaging of time-dependent axonal and myelin degradation after corpus callosotomy in epilepsy patients. *Neuroimage*, 32(3):1090–9, 2006.
- [24] I. Corouge, P. T. Fletcher, S. Joshi, S. Gouttard, and G. Gerig. Fiber tract-oriented statistics for quantitative diffusion tensor MRI analysis. *Medical Image Analysis*, 10(5):786–98, 2006.
- [25] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):1–38, 1977.
- [26] E. DiCicco-Bloom, C. Lord, L. Zwaigenbaum, E. Courchesne, S. R. Dager, C. Schmitz, R. T. Schultz, J. Crawley, and L. J. Young. The developmental neurobiology of autism spectrum disorder. *Journal of Neuroscience*, 26(26):6897–906, 2006.
- [27] A. Donner. A review of inference procedures for the intraclass correlation coefficient in the one-way random effects model. *International Statistical Review*, 54(1):67 – 82, 1986.
- [28] A. Donner and J. J. Koval. The estimation of intraclass correlation in the analysis of family data. *Biometrics*, 36(1):19 – 25, 1980.
- [29] P. Dubey, A. Fatemi, H. Huang, L. Nagee-Poetscher, S. Wakana, P. B. Barker, P. van Zijl, H. W. Moser, S. Mori, and G. V. Raymond. Diffusion tensor-based imaging reveals occult abnormalities in adrenomyeloneuropathy. *Annals of Neurology*, 58(5):758–66, 2005.
- [30] D. Ducreux, I. Huynh, P. Fillard, J. Renoux, M. C. Petit-Lacour, K. Marsot-Dupuch, and P. Lasjaunias. Brain MR diffusion tensor imaging and fibre tracking to differentiate between two diffuse axonal injuries. *Neuroradiology*, 47(8):604–8, 2005.
- [31] D. B. Ennis and G. Kindlmann. Orthogonal tensor invariants and the analysis of diffusion tensor magnetic resonance images. *Magnetic Resonance in Medicine*, 55(1):136–46, 2006.
- [32] J. L. Fleiss. *Design and Analysis of Clinical Experiments*. John Wiley and Sons, New York, NY, 1986.
- [33] C. J. Galban, S. Maderwald, K. Uffmann, and M. E. Ladd. A diffusion tensor imaging analysis of gender differences in water diffusivity within human skeletal muscle. *NMR in Biomedicine*, 18(8):489–98, 2005.
- [34] Y. Ge, M. Law, and R. I. Grossman. Applications of diffusion tensor MR imaging in multiple sclerosis. *Annals of the New York Academy of Sciences*, 1064:202–19, 2005.
- [35] D. A. Harville. Maximum likelihood approaches to variance component estimation and to related problems. *Journal of the American Statistical Association*, 72(358):320–338, 1977.

- [36] E. Heiervang, T. E. Behrens, C. E. Mackay, M. D. Robson, and H. Johansen-Berg. Between session reproducibility and between subject variability of diffusion MR and tractography measures. *Neuroimage*, 33(3):867–77, 2006.
- [37] K. Hoffman and R. Kunze. *Linear Algebra, Second Edition*. Prentice Hall, Englewood Cliffs, NJ, 1971.
- [38] E. W. Hsu, D. L. Buckley, J. D. Bui, S. J. Blackband, and J. R. Forder. Two-component diffusion tensor MRI of isolated perfused hearts. *Magnetic Resonance in Medicine*, 45(6):1039–45, 2001.
- [39] H. Huang, J. Zhang, P. C. van Zijl, and S. Mori. Analysis of noise effects on DTI-based tractography using the brute-force and multi-ROI approach. *Magnetic Resonance in Medicine*, 52(3):559–65, 2004.
- [40] D. K. Jones, L. D. Griffin, D. C. Alexander, M. Catani, M. A. Horsfield, R. Howard, and S. C. Williams. Spatial normalization and averaging of diffusion tensor MRI data sets. *Neuroimage*, 17(2):592–617, 2002.
- [41] D. K. Jones, A. Leemans, and G. C. Koay. On the design of statistically rotationally invariant sampling schemes for diffusion tensor MRI. *IEEE Transactions on Medical Imaging*, in press.
- [42] T. A. Keller, R. K. Kana, and M. A. Just. A developmental study of the structural integrity of white matter in autism. *Neuroreport*, 18(1):23–7, 2007.
- [43] M. Kim, I. Ronen, K. Ugurbil, and D. S. Kim. Spatial resolution dependence of DTI tractography in human occipito-callosal region. *Neuroimage*, 32(3):1243–9, 2006.
- [44] P. B. Kingsley. Introduction to diffusion tensor imaging mathematics: Part I. Tensors, rotations, and eigenvectors. *Concepts in Magnetic Resonance Part A*, 28A(2):101–122, 2006.
- [45] P. B. Kingsley. Introduction to diffusion tensor imaging mathematics: Part II. Anisotropy, diffusion-weighting factors, and gradient encoding schemes. *Concepts in Magnetic Resonance Part A*, 28A(2):123–154, 2006.
- [46] P. B. Kingsley. Introduction to diffusion tensor imaging mathematics: Part III. Tensor calculation, noise, simulations, and optimization. *Concepts in Magnetic Resonance Part A*, 28A(2):155–179, 2006.
- [47] E. O. Kistner and K. E. Muller. Exact distributions of intraclass correlation and Cronbach’s alpha with Gaussian data and general covariance. *Psychometrika*, 69(3):459–474, 2004.
- [48] M. Kubicki, C. F. Westin, S. E. Maier, M. Frumin, P. G. Nestor, D. F. Salisbury, R. Kikinis, F. A. Jolesz, R. W. McCarley, and M. E. Shenton. Uncinate fasciculus findings in

- schizophrenia: A magnetic resonance diffusion tensor imaging study. *American Journal of Psychiatry*, 159(5):813–20, 2002.
- [49] J. E. Lainhart. Advances in autism neuroimaging research for the clinician and geneticist. *American Journal of Medical Genetics Part C: Seminars in Medical Genetics*, 142(1):33–9, 2006.
- [50] N. M. Laird and J. H. Ware. Random-effects models for longitudinal data. *Biometrics*, 38:963 – 974, 1982.
- [51] N. Lange. Statistical approaches to human brain mapping by functional magnetic resonance imaging. *Statistics in Medicine*, 15:389–428, 1996.
- [52] N. Lange. Statistical procedures for functional MRI. In C. Moonen and P. A. Bandettini, editors, *Functional MRI*, pages 301–335. Springer-Verlag, New York, 1999.
- [53] N. Lange, D. Jones, and C. Pierpaoli. A closed-form method for improving inter-subject coherence in diffusion tensor magnetic resonance imaging. In *Proceedings of the 2004 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, Arlington, VA, April 2004. IEEE.
- [54] N. Lange and N. M. Laird. The effect of covariance structure on variance estimation in balanced growth-curve models with random parameters. *Journal of the American Statistical Association*, 84(405):241 – 247, 1989.
- [55] N. Lange and L. Ryan. Assessing normality in random effects models. *Annals of Statistics*, 17(2):624–642, 1989.
- [56] D. Le Bihan. Looking into the functional architecture of the brain with diffusion MRI. *Nature Reviews Neuroscience*, 4(6):469–80, 2003.
- [57] D. Le Bihan, J. F. Mangin, C. Poupon, C. A. Clark, S. Pappata, N. Molko, and H. Chabriat. Diffusion tensor imaging: Concepts and applications. *Journal of Magnetic Resonance Imaging*, 13(4):534–46, 2001.
- [58] J. E. Lee, E. D. Bigler, A. L. Alexander, M. Lazar, M. B. DuBray, M. K. Chung, M. Johnson, J. Morgan, J. N. Miller, W. M. McMahon, J. Lu, E. K. Jeong, and J. E. Lainhart. Diffusion tensor imaging of white matter in the superior temporal gyrus and temporal stem in autism. *Neuroscience Letters*, 424(2):127–32, 2007.
- [59] K. O. Lim and J. A. Helpert. Neuropsychiatric applications of DTI - a review. *NMR in Biomedicine*, 15(7-8):587–93, 2002.
- [60] X. Lin, C. R. Tench, P. S. Morgan, G. Niepel, and C. S. Constantinescu. ‘Importance sampling’ in MS: Use of diffusion tensor tractography to quantify pathology related to specific impairment. *Journal of the Neurological Sciences*, 237(1-2):13–9, 2005.

- [61] C. Liston, R. Watts, N. Tottenham, M. C. Davidson, S. Niogi, A. M. Ulug, and B. J. Casey. Frontostriatal microstructure modulates efficient recruitment of cognitive control. *Cerebral Cortex*, 16(4):553–60, 2006.
- [62] M. J. Lowe, C. Horenstein, J. G. Hirsch, R. A. Marrie, L. Stone, P. K. Bhattacharyya, A. Gass, and M. D. Phillips. Functional pathway-defined MRI diffusion measures reveal increased transverse diffusivity of water in multiple sclerosis. *Neuroimage*, 32(3):1127–33, 2006.
- [63] J. F. Mangin, C. Poupon, C. Clark, D. Le Bihan, and I. Bloch. Distortion correction and robust tensor estimation for MR diffusion imaging. *Medical Image Analysis*, 6(3):191–8, 2002.
- [64] *MATLAB R2007a*. Natick, MA, 2007.
- [65] J. Mattiello, P. J. Basser, and D. LeBihan. Analytical expressions for the b matrix in NMR diffusion imaging and spectroscopy. *Journal of Magnetic Resonance, Series A*, 108(2):131–141, 1994.
- [66] K. O. McGraw and S. P. Wong. Forming inferences about some intraclass correlation coefficients. *Psychological Methods*, 1(1):30 – 46, 1996.
- [67] P. McGraw, L. Liang, and J. M. Provenzale. Evaluation of normal age-related changes in anisotropy during infancy and childhood as shown by diffusion tensor imaging. *American Journal of Roentgenology*, 179(6):1515–22, 2002.
- [68] M. E. Moseley and K. Butts. Diffusion and perfusion. In D. D. Stark and J. Bradley, W. G., editors, *Magnetic Resonance Imaging*, volume 3, pages 1515 – 1538. Mosby, St. Louis, 3rd edition, 1999.
- [69] K. W. Muir, A. Buchan, R. von Kummer, J. Rother, and J. C. Baron. Imaging of acute stroke. *Lancet Neurology*, 5(9):755–68, 2006.
- [70] J. J. Neil, S. I. Shiran, R. C. McKinstry, G. L. Schefft, A. Z. Snyder, C. R. Almlie, E. Akbudak, J. A. Aronovitz, J. P. Miller, B. C. Lee, and T. E. Conturo. Normal brain in human newborns: Apparent diffusion coefficient and diffusion anisotropy measured by using diffusion tensor MR imaging. *Radiology*, 209(1):57–66, 1998.
- [71] C. J. Newschaffer, L. A. Croen, J. Daniels, E. Giarelli, J. K. Grether, S. E. Levy, D. S. Mandell, L. A. Miller, J. Pinto-Martin, J. Reaven, A. M. Reynolds, C. E. Rice, D. Schendel, and G. C. Windham. The epidemiology of autism spectrum disorders. *Annual Review of Public Health*, 28:235–58, 2007.
- [72] R. C. Oldfield. The assessment and analysis of handedness: The Edinburgh inventory. *Neuropsychologia*, 9:97–113, 1971.

- [73] N. G. Papadakis, C. D. Murrills, L. D. Hall, C. L. Huang, and T. Adrian Carpenter. Minimal gradient encoding for robust estimation of diffusion anisotropy. *Magnetic Resonance Imaging*, 18(6):671–9, 2000.
- [74] H. J. Park, M. Kubicki, M. E. Shenton, A. Guimond, R. W. McCarley, S. E. Maier, R. Kikinis, F. A. Jolesz, and C. F. Westin. Spatial normalization of diffusion tensor MRI using multiple channels. *Neuroimage*, 20(4):1995–2009, 2003.
- [75] C. Pierpaoli and P. J. Basser. Toward a quantitative assessment of diffusion anisotropy. *Magnetic Resonance in Medicine*, 36(6):893–906, 1996.
- [76] C. Pierpaoli, P. Jezzard, P. J. Basser, A. Barnett, and G. Di Chiro. Diffusion tensor MR imaging of the human brain. *Radiology*, 201(3):637–48, 1996.
- [77] H. W. Powell, G. J. Parker, D. C. Alexander, M. R. Symms, P. A. Boulby, C. A. Wheeler-Kingshott, G. J. Barker, U. Noppeney, M. J. Koepp, and J. S. Duncan. Hemispheric asymmetries in language-related pathways: A combined functional MRI and tractography study. *Neuroimage*, 32(1):388–99, 2006.
- [78] C. R. Rao. The theory of least squares when the parameters are stochastic and its application to the analysis of growth curves. *Biometrika*, 52(3/4):447–458, 1965.
- [79] M. Ries, R. A. Jones, F. Basseau, C. T. Moonen, and N. Grenier. Diffusion tensor MRI of the human kidney. *Journal of Magnetic Resonance Imaging*, 14(1):42–9, 2001.
- [80] G. K. Rohde, A. S. Barnett, P. J. Basser, and C. Pierpaoli. Estimating intensity variance due to noise in registered images: Applications to diffusion tensor MRI. *Neuroimage*, 26(3):673–84, 2005.
- [81] C. H. Salmond, D. K. Menon, D. A. Chatfield, G. B. Williams, A. Pena, B. J. Sahakian, and J. D. Pickard. Diffusion tensor imaging in chronic head injury survivors: Correlations with learning and memory indices. *Neuroimage*, 29(1):117–24, 2006.
- [82] C. Schmitz and P. Rezaie. The neuropathology of autism: Where do we stand? *Neuropathology and Applied Neurobiology*, 34(1):4–11, 2008.
- [83] T. Schonberg, P. Pianka, T. Hendler, O. Pasternak, and Y. Assaf. Characterization of displaced white matter by brain tumors using combined DTI and fMRI. *Neuroimage*, 30(4):1100–11, 2006.
- [84] S. R. Searle. *Linear Models*. John Wiley and Sons, New York, 1971.
- [85] R. J. Shavelson, N. M. Webb, and G. L. Rowley. Generalizability theory. *American Psychologist*, 44(6):922–932, 1989.
- [86] P. E. Shrout and J. L. Fleiss. Intraclass correlations: Uses in assessing rater reliability. *Psychological Bulletin*, 86(2):420–428, 1979.

- [87] C. H. Sotak. The role of diffusion tensor imaging in the evaluation of ischemic brain injury - a review. *NMR in Biomedicine*, 15(7-8):561–9, 2002.
- [88] W. M. Stanish and N. Taylor. Estimation of the intraclass correlation coefficient for the analysis of covariance model. *The American Statistician*, 37(3):221 – 224, 1983.
- [89] P. C. Sundgren, Q. Dong, D. Gomez-Hassan, S. K. Mukherji, P. Maly, and R. Welsh. Diffusion tensor imaging of the brain: Review of clinical applications. *Neuroradiology*, 46(5):339–50, 2004.
- [90] T. H. Szatrowski. Necessary and sufficient conditions for explicit solutions in the multivariate normal estimation problem for patterned means and covariances. *The Annals of Statistics*, 8(4):802 – 810, 1980.
- [91] P. R. Szeszko, B. A. Ardekani, M. Ashtari, A. K. Malhotra, D. G. Robinson, R. M. Bilder, and K. O. Lim. White matter abnormalities in obsessive-compulsive disorder: A diffusion tensor imaging study. *Archives of General Psychiatry*, 62(7):782–90, 2005.
- [92] L. Thivard, S. Lehericy, A. Krainik, C. Adam, D. Dormont, J. Chiras, M. Baulac, and S. Dupont. Diffusion tensor imaging in medial temporal lobe epilepsy with hippocampal sclerosis. *Neuroimage*, 28(3):682–90, 2005.
- [93] M. Walterfang, S. J. Wood, D. Velakoulis, D. Copolov, and C. Pantelis. Diseases of white matter and schizophrenia-like psychosis. *Australian and New Zealand Journal of Psychiatry*, 39(9):746–56, 2005.
- [94] C. F. Westin, M. Martin-Fernandez, C. Alberola-Lopez, J. Ruiz-Alzola, and H. Knutsson. Tensor field regularization using normalized convolution and Markov random fields in a Bayesian framework. In J. Weickert and H. Hagen, editors, *Visualization and Image Processing of Tensor Fields*, Mathematics and Visualization, pages 381–398, 464–467. Springer, 2006.
- [95] C. F. Westin, S. Peled, H. Gudbjartsson, R. Kikinis, and F. A. Jolesz. Geometrical diffusion measures for MRI from tensor basis analysis. In *Proceedings of the 5th Annual Meeting of the International Society for Magnetic Resonance in Medicine*, page 1742, Vancouver, BC, Canada, April 1997.
- [96] S. Xie, J. X. Xiao, Y. H. Wang, H. K. Wu, G. L. Gong, and X. X. Jiang. Evaluation of bilateral cingulum with tractography in patients with alzheimer’s disease. *Neuroreport*, 16(12):1275–8, 2005.
- [97] C. S. Yu, K. C. Li, Y. Xuan, X. M. Ji, and W. Qin. Diffusion tensor tractography in patients with cerebral tumors: A helpful technique for neurosurgical planning and postoperative assessment. *European Journal of Radiology*, 56(2):197–204, 2005.