

# **Three-Dimensional Morphometric Analysis of Brain Ventricles Using Statistical Shape Models**

A thesis submitted to the University of Manchester  
for the degree of Doctor of Philosophy  
in the Faculty of Medicine, Dentistry, Nursing and Pharmacy

2004

Kolawole .O. Babalola

Imaging Science and Biomedical Engineering

# Contents

<b>1</b>	<b>Introduction</b>	<b>20</b>
1.1	Schizophrenia . . . . .	21
1.2	The Lateral Ventricles . . . . .	22
1.2.1	Anatomical Description of the Brain Ventricles . . . . .	22
1.2.2	Function of the Brain Ventricles . . . . .	23
1.3	Ventricular Changes in Schizophrenia . . . . .	24
1.3.1	Volumetric Studies . . . . .	25
1.3.2	Morphometric Studies . . . . .	25
1.4	Aims and Motivation . . . . .	26
1.5	Overview of Thesis . . . . .	27
<b>2</b>	<b>Background</b>	<b>28</b>
2.1	Overview . . . . .	28
2.2	Magnetic Resonance Imaging . . . . .	29
2.3	Volumetric Images . . . . .	31
2.3.1	Image Resolution and the Partial Volume Effect . . . . .	32
2.3.2	Re-sampling and Interpolation . . . . .	35
2.3.3	Image Registration . . . . .	36



2.4	Image Preprocessing . . . . .	39
2.4.1	MR Intensity Non-Uniformity Correction . . . . .	40
2.4.2	Segmentation . . . . .	41
2.5	Shape Analysis in Medical Imaging . . . . .	44
2.5.1	Shape . . . . .	45
2.5.2	Shape Analysis . . . . .	45
2.5.3	A Selection of Methods for 3D Shape Analysis of Neuroanatomical Structures . . . . .	58
2.5.4	Normalisation for Brain/Head Shape . . . . .	65
2.6	Summary . . . . .	66
<b>3</b>	<b>Data and Image Preprocessing</b>	<b>68</b>
3.1	Overview . . . . .	68
3.2	Description of Data . . . . .	68
3.3	Application of MR Non-Uniformity Correction . . . . .	69
3.4	Segmentation of the Lateral Ventricles . . . . .	70
3.5	Image Combination to Improve Spatial Resolution . . . . .	73
3.5.1	Evaluation of Three Registration Packages . . . . .	76
3.5.2	Combination Method 1 - Image Combination by Registration and Interpolation . . . . .	81
3.5.3	Combination Method 2 - Image Combination by Solution of Linear Equations . . . . .	82
3.5.4	Experiments on Image Combination . . . . .	87
3.5.5	Selection of Method of Choice for Image Combination . . . . .	89
3.6	Obtaining Brain Size Parameters . . . . .	91

3.7	Summary . . . . .	91
<b>4</b>	<b>3D Point Distribution Models - Theory and Practice</b>	<b>95</b>
4.1	Overview . . . . .	95
4.2	Point Distribution Models - The Theory . . . . .	96
4.2.1	Mathematical Definition of 3D PDMs . . . . .	97
4.2.2	Steps to Create a Point Distribution Model . . . . .	100
4.3	A Review of Approaches to Construct 3D PDMs . . . . .	101
4.4	Quantitative Assessment of 3D PDMs . . . . .	108
4.4.1	Generalisation Ability . . . . .	109
4.4.2	Specificity . . . . .	110
4.4.3	Compactness . . . . .	111
4.4.4	A Distance Transform Metric for Comparing Two Surfaces . .	112
4.5	Our Approach to the Construction of a 3D PDM of the Lateral Ventricles	113
4.6	Curvature-Based Landmarks on the Ventricular Surface . . . . .	115
4.6.1	Curvature of 3D Surfaces . . . . .	117
4.6.2	Obtaining Curvature-based Characteristics of 3D Surfaces . .	125
4.6.3	The Use of Crest Points and Crest Lines as Curvature-Based Landmarks on the Ventricular Surface . . . . .	132
4.7	Generating Point-to-Point Correspondences . . . . .	134
4.7.1	A Closest Point Approach to Point Matching . . . . .	134
4.7.2	A Transportation Approach to Point Matching . . . . .	136
4.7.3	Quantitative Comparison of the Transportation and Closest Point Methods . . . . .	140
4.8	Dependence on Choice of Template . . . . .	159

4.8.1	Dependence of PDM Construction on Choice of Template without Re-projection . . . . .	163
4.8.2	Re-Projection as a Solution to Dependence on the Template . . . . .	165
4.8.3	Dependence of PDM Construction on Choice of Template, Using Re-projection . . . . .	167
4.9	Summary . . . . .	175
<b>5</b>	<b>Morphometric Analysis</b>	<b>176</b>
5.1	Overview . . . . .	176
5.2	Experimental Methods . . . . .	177
5.3	Using the Parameters of the PDM for Morphometric Analysis . . . . .	178
5.4	Comparisons we Wish to Make . . . . .	178
5.5	Shape (Group) Comparisons Using Linear Discriminant Analysis . . . . .	180
5.6	Asymmetry Comparisons Using Paired Linear Discriminant Analysis . . . . .	185
5.7	Results of Morphometric Analysis . . . . .	188
5.7.1	Shape Differences Between Schizophrenics and Controls . . . . .	188
5.7.2	Differences in Extent of Asymmetry within Schizophrenic and Control Groups . . . . .	188
5.8	Discussion . . . . .	199
5.8.1	Discussion of Results of Shape Comparisons . . . . .	199
5.8.2	Discussion of Asymmetry Results . . . . .	201
5.9	Summary . . . . .	204
<b>6</b>	<b>Conclusions</b>	<b>205</b>
6.1	Contributions of this Thesis . . . . .	205
6.2	Further Work . . . . .	207

6.3	Final Statement . . . . .	208
<b>A</b>	<b>Operations on the Triangulated Surfaces</b>	<b>209</b>
A.1	Outward Normals to a Triangulated Surface . . . . .	209
A.2	Intersection of a Triangulated Surface By a Line Segment . . . . .	210
A.3	Generating Voxel Representations of Triangulated Surfaces . . . . .	213
<b>B</b>	<b>Subjects Participating in the Study</b>	<b>215</b>
<b>C</b>	<b>Movies on Accompanying CD</b>	<b>233</b>

# List of Figures

1.1	A view of the brain showing the ventricular system (source [26]) . . .	23
1.2	A view of the brain showing the lateral ventricles with the four main parts labelled (source [26]) . . . . .	24
2.1	Schematic illustration of the use of nuclear magnetic resonance in the generation of signals for MRI . . . . .	30
2.2	Principal components of a complete magnetic resonance imaging system (source [132]) . . . . .	32
2.3	Sub-figures illustrating a 2D cross-section (a) and the discrete nature of images in 2D (b) and 3D (c) . . . . .	33
2.4	The three orthogonal orientations of MRI cross-sections through a volumetric 3D image of the head . . . . .	34
2.5	An example to illustrate shape differences using two different outlines of the hand . . . . .	46
2.6	Illustration of classification of landmarks to anatomical, mathematical and pseudo-landmarks . . . . .	55
3.1	Qualitative results of intensity non-uniformity correction . . . . .	70
3.2	Images of a coronal slice showing the 3D edges and the resulting segmented contours of the ventricle . . . . .	74
3.3	A schematic illustration of the 2D case for image combination by linear equations . . . . .	83

3.4	The simple 2D case to illustrate the smoothness term . . . . .	85
3.5	Qualitative results of the application of reconstruction methods to test images . . . . .	88
3.6	Qualitative results of the application of reconstruction methods to an image of one of the subjects of this study . . . . .	90
3.7	Rendering of a ventricular surface before and after image combination	92
3.8	A brain surface of one subject is shown with an ellipsoid fitted to it .	93
4.1	Outline of algorithm for alignment of shapes in a training set . . . . .	102
4.2	Binary contours of two slices of the ventricle (a) and (c) and their respective distance transforms . . . . .	114
4.3	Schematic diagram of the process of creating the 3D PDM of the ventricles . . . . .	116
4.4	A regular surface patch $S$ over a 2D plane $(u,v)$ . . . . .	118
4.5	Illustration of tangent and normal planes and the normal section through a surface . . . . .	120
4.6	Characterisation of local surface properties using principal curvatures	122
4.7	Representation of voxels as nodes in a rectangular array for the Marching Lines Algorithm . . . . .	129
4.8	Summary of the steps in using the Marching Lines Algorithm to extract crest lines . . . . .	130
4.9	Four ventricular surfaces colour-mapped according to mean curvature with crest lines (after filtering out noisy ones) superimposed on them	131
4.10	Illustration of the result of point-to-point matching . . . . .	135
4.11	Outline of algorithm for point-to-point matching . . . . .	137
4.12	Graphs illustrating quantitative comparison of point matching methods	143
4.13	Application of closest point and transportation methods to a pair of synthetic lines. Results before monotonicity and injectivity enforced .	146

4.14	Application of closest point and transportation methods to a pair of synthetic lines. Results after monotonicity and injectivity enforced . .	147
4.15	Plot of the index of symmetry for each ventricle for both the transportation and the closest point methods . . . . .	148
4.16	Point-to-point correspondences obtained on a pair of ventricles giving the worst measure of symmetry for the closest point method . . . . .	149
4.17	Point-to-point correspondences obtained on a pair of ventricles giving the best measure of symmetry for the closest point method . . . . .	150
4.18	Comparison of the Generalisation Ability of the PDMs constructed by the closest point and transportation methods . . . . .	152
4.19	Comparison of the Specificity of the PDMs constructed by the closest point and transportation methods . . . . .	154
4.20	Compactness of the PDMs created by the transportation and closest point methods . . . . .	155
4.21	Variance captured by the the PDMs created by the transportation and closest point methods . . . . .	155
4.22	A schematic diagram illustrating how the quantitative measure used by [53] is obtained. . . . .	161
4.23	Mean shapes of PDMs of the same training set obtained from four different templates . . . . .	164
4.24	Qualitative comparison of a member of the corresponding training set produced by four different choices of template with and without the re-projection step . . . . .	168
4.25	Mean shapes of PDMs of the same training set obtained from four different templates by three different methods . . . . .	170
4.26	First mode of variation of the PDMs of the same training set obtained from four different templates . . . . .	171
4.27	First mode of variation of the PDMs after projection of landmark points onto the surfaces of the actual ventricles in the training set . .	172

5.1	Illustration of an ideal case of separation of two groups (a) compared with the actual case for the first and second shape parameters of the male schizophrenics and male controls . . . . .	181
5.2	The importance of the direction of the discriminant vector in obtaining good separation is illustrated. In (a) the direction is chosen arbitrarily, whereas in (b) it is that which maximises Fisher's criterion . . . . .	183
5.3	Illustration of finding the discriminant vector maximising the separation between a group of pairs. The figure shows the vector obtained after application of the methods described in section 5.6 . . . . .	186
5.4	Results of Comparison 1 - Shape differences between all schizophrenic subjects and all control subjects . . . . .	189
5.5	Results of Comparison 2 - Shape differences between all male schizophrenic subjects and all male control subjects . . . . .	190
5.6	Results of Comparison 3 - Shape differences between all female schizophrenic subjects and all female control subjects . . . . .	191
5.7	Results of Comparison 4 - Scalar values for projections of left and right pairs of all schizophrenic and all control subjects . . . . .	193
5.8	Results of Comparison 4 - Colour-mapped ventricular surface of shape differences between means of left and right pairs of all schizophrenic and all control subjects after projection onto the discriminant vector .	194
5.9	Results of Comparison 5 - Scalar values for projections of left and right pairs of male schizophrenic and male control subjects onto the discriminant vector . . . . .	195
5.10	Results of Comparison 5 - Colour-mapped ventricular surface of shape differences between means of left and right pairs of male schizophrenic and male control subjects after projection onto the discriminant vector	196
5.11	Results of Comparison 6 - Scalar values for projections of left and right pairs of female schizophrenic and female control subjects onto the discriminant vector . . . . .	197
5.12	Results of Comparison 6 - Colour-mapped ventricular surface of shape differences between means of left and right pairs of female schizophrenic and female control subjects after projection onto the discriminant vector	198



A.1	A facet $\mathbf{F}$ with vertices $V_0$ , $V_1$ , and $V_2$ lying in the plane $\pi$ . The line segment from the point $P_0$ in the direction $\hat{\mathbf{n}}_v$ , intersects $\mathbf{F}$ at $P_x$ . . .	211
A.2	An illustration of using planes to intersect the ventricular surface in creating a voxel representation of the triangulated ventricular surface	214
B.1	Lateral ventricles of male schizophrenic subjects (1 to 5 of 30) . . . .	217
B.2	Lateral ventricles of male schizophrenic subjects (6 to 10 of 30) . . . .	218
B.3	Lateral ventricles of male schizophrenic subjects (11 to 15 of 30) . . .	219
B.4	Lateral ventricles of male schizophrenic subjects (16 to 20 of 30) . . .	220
B.5	Lateral ventricles of male schizophrenic subjects (21 to 25 of 30) . . .	221
B.6	Lateral ventricles of male schizophrenic subjects (26 to 30 of 30) . . .	222
B.7	Lateral ventricles of female schizophrenic subjects (1 to 5 of 9) . . . .	223
B.8	Lateral ventricles of female schizophrenic subjects (6 to 9 of 9) . . . .	224
B.9	Lateral ventricles of male control subjects (1 to 5 of 17) . . . . .	225
B.10	Lateral ventricles of male control subjects (6 to 9 of 17) . . . . .	226
B.11	Lateral ventricles of male control subjects (10 to 13 of 17) . . . . .	227
B.12	Lateral ventricles of male control subjects (14 to 17 of 17) . . . . .	228
B.13	Lateral ventricles of female control subjects (1 to 5 of 13) . . . . .	229
B.14	Lateral ventricles of female control subjects (6 to 9 of 13) . . . . .	230
B.15	Lateral ventricles of female control subjects (10 to 13 of 13) . . . . .	231

# List of Tables

3.1	Mean Translation and Rotation Parameters for AIR . . . . .	79
3.2	Mean Translation and Rotation Parameters for MPR . . . . .	80
3.3	Mean Translation and Rotation Parameters for FLIRT . . . . .	80
3.4	Summary of values of $d_m$ (mm) for AIR, MPR and FLIRT . . . . .	80
3.5	Average values of NCC (n=5) for comparison of reconstructed image with original . . . . .	89
4.1	Value of distance transform metric for comparison of surfaces of the corresponding training sets with their respective original surfaces . . .	147
4.2	Values of the distance transform metric for the comparison of mean shapes of PDMs . . . . .	169
4.3	Values of the distance transform metric for the comparison of mean shapes of PDMs after projection of normals and smoothing steps were added to the method of construction of the PDMs . . . . .	173
4.4	Means of the diagonal and off-diagonal terms of Tables 4.2 and 4.3 . .	173
5.1	Summary of results of shape comparisons . . . . .	192
5.2	Asymmetry results - pair-wise $t$ -tests of differences in left-right pairs.	192
5.3	Summary of the quantitative results of the asymmetry analysis . . . .	199
B.1	Age and sex data for schizophrenic subjects . . . . .	216

B.2	Age and sex data for control subjects . . . . .	232
-----	---	-----

## Abstract

It is widely accepted that schizophrenia is accompanied by an increase in the volume of the lateral ventricles of the brain. It is of interest to investigate if the volumetric changes are accompanied by specific localised shape changes. If these can be shown they will aid in understanding what is a complex disease. Studies based on volumetric measurements have been carried out to further characterise the observed differences. However, volume is not a powerful shape descriptor - as ventricles with the same volume can have quite different shapes. Additionally, because of the intrinsic natural variation of biological structures, statistical methods which allow the separation of variability due to disease from natural variability are of interest.

Statistical shape models (SSMs) are a tool from computer vision that capture the inherent variation in the shape of a specified class of object based on statistics learnt from a sample representative of the population of interest. Our choice of SSM for carrying out morphometric analysis is the point distribution model (PDM). This type of model has been widely used as a segmentation tool in 2D and application in 3D is a subject of active research. The aim of this thesis is to use a 3D PDM to quantify localised changes in the shape of the lateral ventricles associated with schizophrenia.

Starting with magnetic resonance (MR) images of 30 control subjects and 39 age and sex matched schizophrenics, we describe the pre-processing of the MR images to improve their quality prior to segmentation. A semi-automatic approach to the segmentation of the lateral ventricles was adopted. We give details of an automatic approach to the construction of the 3D PDM using crest points as curvature-based landmarks on the ventricles. We use the transportation algorithm to solve the problem of defining correspondences of crest points. Issues regarding the comparison of 3D PDMs are also discussed.

Applying discriminant analysis to the most important shape parameters obtained from the PDM, the means of the schizophrenic and control groups are significantly different ( $p < 10^{-12}$ ). The shape changes observed were localised to three regions : the temporal horn (its tip near the amygdala, and along its body near the parahippocampal fissure), the central part of the lateral ventricles around the corpus callosum, and the tip of the anterior horn in the region of the frontal lobe. The differences in the temporal region and anterior horns are in regions close to structures thought to be implicated in schizophrenia. Investigations of asymmetry between left and right ventricle pairs was also undertaken. This showed that a degree of asymmetry existed in both schizophrenics and controls, and that the nature of asymmetry was different in the two groups.

# Declaration

No portion of the work referred to in the thesis has been submitted in support of an application for another degree or qualification of this or any other university or other institute of learning.

# Copyright

1. Copyright in text of this thesis rests with the Author. Copies (by any process) either in full, or of extracts, may be made **only** in accordance with instructions given by the Author and lodged in the John Rylands University Library of Manchester. Details may be obtained from the Librarian. This page must form part of any such copies made. Further copies (by any process) of copies made in accordance with such instructions may not be made without permission (in writing) of the Author.
2. The ownership of any intellectual property rights which may be described in this thesis is vested in the University of Manchester, subject to any prior agreement to the contrary, and may not be made available for use by third parties without the written permission of the University, which will prescribe the terms and conditions of any such agreement.

Further information on the conditions under which disclosures and exploitation may take place is available from the Head of the Division of Imaging Science and Biomedical Engineering.

# Acknowledgements

Firstly I give thanks to God, because I believe it is only by divine intervention that I have managed to complete the production of this body of work.

I would also like to thank the following people for their help during my research:

My supervisor, Dr Jim Graham for being truly outstanding. He has displayed a great deal of patience, and he provided a crucial vote of confidence after the end of a disappointing first year. He provided ideas which solved some of the difficulties encountered, and yet did not prevent me from exploring my own attempts at solutions.

My wife Shola, for taking care of domestic matters and putting up with the various extremes I have gone to over the past four years. Our children Seun and Ayo for being a constant source of joy, and for helping me to put this PhD in perspective with the wider context of life. My mother, and my brothers and sisters (and their partners) for providing moral and financial support to both me and my family.

The members of ISBE for always being willing to help with mathematical and computing problems. It has been a pleasure to work in an atmosphere that combines a unique mix of individual flair with collective harmony.

Thanks also to Dr Lili Kopala and Dr Donna Lang of the University of British Colombia for providing the images used in this study. The Epidaure group of INRIA, France, for the Marching Lines Algorithm code, and Professor P. Klienschmidt of Passau University for the transportation algorithm code. Dr Amit Herwadkar, Professor Alan Jackson and Professor Bill Deakin, of the University of Manchester for helping with segmentation and interpreting the results.

# Dedication

To the memory of my father, Samuel Oladeji Babalola.



# Perface

The following publications have arisen from the research presented in this thesis:

- K.O. Babalola, J. Graham, L. Kopala, and R. Vandorpe. Using the Transporation Algorithm to Improve Point Correspondences in the Construction of 3D PDMs. *In Proceedings of the 6<sup>th</sup> Annual Conference on Medical Image Understanding and Analysis*, Portsmouth, UK, Pages 141-144, July 2002
- K.O. Babalola, J. Graham, L. Kopala, and R. Vandorpe. Investigation of Shape Changes in the Lateral Ventricles Associated with Schizophrenia : A Morphometric Study Using a Three-Dimensional Point Distribution Model. *In Proceedings of the 7<sup>th</sup> Annual Conference on Medical Image Understanding and Analysis*, Sheffield, UK, Pages 97-100, July 2003
- K.O. Babalola and J. Graham. Characterisation of Morphometric Changes in the Lateral Ventricles Due to Schizophrenia using 3 Dimensional Point Distribution Models. *In Proceedings of the 11<sup>th</sup> Annual Meeting of the International Society for Magnetic Resonance in Medicine (ISMRM)*, Toronto, July 2003
- K.O. Babalola, J. Graham, W. Honer, L. Kopala, D. Lang, and R. Vandorpe. An Investigation of Morphometric Changes in the Lateral Ventricles of Schizophrenic Subjects. *In Proceedings of the 6<sup>th</sup> Conference on Medical Image Computing and Computer Assisted Intervention*, LCNS 2879, pp521-529, Montreal, November 2003

# Chapter 1

## Introduction

The internal structure of the human body is an age old fascination of the human race. Some scientists have gone to macabre lengths such as body-snatching in the nineteenth century [131], in their thirst for more knowledge about it. The fascination with our internal structure still continues to the present day - consider for instance the controversial Bodyworlds exhibitions of Professor Gunther von Hagens [55].

However, seeing *into* the body is of prime importance in medicine. It aids diagnosis, helps in understanding the aetiology of diseases, and is now commonly used in surgical and radiotherapy planning. The discovery of X-rays by William Roentgen in 1895 heralded the birth of non-invasive diagnosis. There are now four main methods of non-invasively imaging the body - Computed Tomography (CT) and other X-ray based methods, Magnetic Resonance Imaging (MRI), Positron Emission Tomography (PET) and Ultrasound Scanning (US). Novel imaging methods using devices such as the Superconducting Quantum Interference Device (SQUID) are continually being reported.

The processing power of computers and the capacity for storage of digital data are also increasing rapidly. This has lead to a proliferation in the availability of high quality digital data. Research in medical image analysis and the application of computer

vision techniques to solving biomedical problems has also increased at a rapid pace. Unfortunately, the development of automated tools has not matched the needs posed by the proliferation of imaging data. Nevertheless, important new directions in the analysis and interpretation of medical images are being realised.

This thesis explores one of these new directions in which there is a lot of interest - the application of morphometric analysis to 3D surfaces of neuroanatomical structures. By morphometric analysis, we mean the measurement of shape changes or differences. Our goal is to investigate shape and asymmetry differences between the lateral ventricles in the brains of control and schizophrenic subjects. En route to achieving this we describe, in 3D, the construction of a Point Distribution Model (PDM) [32] of the lateral ventricles, and comprehensively demonstrate its application to answering questions posed by morphometric analysis.

## 1.1 Schizophrenia

Schizophrenia is a collective term for mental disorders in which affected individuals exhibit grossly abnormal behaviour. The broad range of its symptoms include reduction in cognitive ability, hallucinations, and emotional instability. The general population risk of being affected is 1%, although genetic predisposition can considerably increase this risk. Onset usually occurs before an individual is 30 years old, and about 85% of sufferers respond to drug therapy. However, given its relatively common occurrence, and the severity of its symptoms this represents a reduction in the quality of life of a significant number of people who cannot be treated for the disease. An even greater number have their lives affected because of the interactions of schizophrenics with their families and with the society.

Although a genetic predisposition to schizophrenia has been recognised, the aetiology of schizophrenia is still unknown. Schizophrenia is divided into two types (Type I and Type II) and a number of subclasses of each type exist depending on the symptoms

exhibited by a patient. However, categorisation of patients is difficult as they usually exhibit symptoms from more than one subclass. Investigations of size and shape changes of the brain associated with schizophrenia may help with categorisation.

There is an intrinsic degree of asymmetry between the left and right hemispheres of human brains. It is believed that these occur in the developmental stage before birth. However, several volumetric MRI studies e.g. [10] and postmortem investigations [49] have suggested that the degree of asymmetry in the brains of schizophrenic subjects is less than that in the normal population.


Several theories about the origin of the disease have been put forward, e.g. biochemical hypotheses implicating neurotransmitters such as dopamine and glutamate [27]. A hypothesis by Professor T J Crow [35] - often referred to as the *Crow Hypothesis*, is of particular interest to morphometric studies. Crow hypothesised that the asymmetry in the brain is an evolutionary trait associated with increased linguistic ability, and that the loss of asymmetry in schizophrenics accounts for some of the symptoms associated with the disease. Morphometric studies of specific structures in the brain can allow further investigation of this theory.

## 1.2 The Lateral Ventricles

### 1.2.1 Anatomical Description of the Brain Ventricles

The brain ventricles are extensive cavities within the brain containing cerebro-spinal fluid (CSF). There are four ventricles - the right and left lateral ventricles, the third ventricle and the fourth ventricle. The structure of these within the human brain is shown in Figure 1.1. The two lateral ventricles are the subjects of interest in this work. They are the largest of the ventricles, one in each hemisphere of the brain. Each lateral ventricle has a central part, and anterior, posterior and inferior horns as shown in Figure 1.2. The third ventricle connects directly with the lateral ventricles. The


third and fourth ventricles are connected by a narrow canal, the “cerebral aqueduct” or aqueduct of Sylvius. The fourth ventricle is contiguous with the central canal of the spinal cord.



**Figure 1.1:** A view of the brain showing the ventricular system (source [26])

### 1.2.2 Function of the Brain Ventricles

The brain ventricles produce and circulate CSF within the brain. Each ventricle contains a choroid plexus which produces CSF within it. The CSF performs the physical function of buffering the brain from shock due to internal or external forces and facilitates biochemical activity by draining excess fluids and acting as a medium for the transfer of substances between the blood and the nervous tissues. The lateral ventricles contain the largest amount of choroid plexus and produce the majority of the CSF. This flows into the third ventricle then the fourth ventricle where CSF produced by each of these augments that produced by the lateral ventricles. The CSF



**Figure 1.2:** A view of the brain showing the lateral ventricles with the four main parts labelled (source [26])

flows from the fourth ventricle into the subarachnoid space from which majority of it flows over the cerebral hemisphere, then over the spinal chord before being reabsorbed back into blood through the arachnoid villi and through the walls of capillaries of the central nervous system.

### 1.3 Ventricular Changes in Schizophrenia

Hemispheric asymmetry in the brain is due to differences in the distribution of brain tissue and differences in the size and shape of neuroanatomical structures in the left and right hemispheres. The brain is a jelly-like material composed mainly of gray matter (GM), white matter (WM), and CSF enclosed in the rigid casing of the skull. It therefore occupies a fixed volume, and changes giving rise to asymmetry will be reflected in the distribution and relative composition of GM, WM and CSF in the brain. As majority of the CSF is contained in the lateral ventricles, changes in their size and shape may indicate structural abnormalities in adjacent brain structures due to disease. Volumetric investigations of differences in the lateral ventricles due to schizophrenia have been carried out. More recently, a number of shape based studies

have also been carried out. A brief review follows.

### 1.3.1 Volumetric Studies

Schizophrenia is widely believed to be accompanied by an increase in size of the lateral ventricles. This was first reported by Johnstone et. al. [67] in a CT study of age-matched controls and 17 institutionalized schizophrenics. The nature of the volume changes have been investigated by many groups. For instance Chance et. al. [28] report selective enlargement of the temporal horns of the ventricles, and Puri et. al. [94] report significant differences in volumetric asymmetries of two groups of schizophrenics, one group exhibiting active syndromes and the other withdrawn syndromes. [90] gives a recent review of volumetric findings in schizophrenia.

### 1.3.2 Morphometric Studies

Buckley et al. [24] use 48 manually defined landmarks corresponding to curvature extrema on the surface of the ventricles of 20 schizophrenic patients and 20 control subjects to investigate shape differences. They considered the whole ventricular system and reported no overall shape differences between the entire patient group and the entire schizophrenic group. However, when only the males of both groups were considered, significant shape differences were identified in the proximal parts of the temporal horn of the lateral ventricles and in the foramen of Monro.

Narr et al. [87] obtained average maps of anatomical differences based on voxel values of the limbic structures and the lateral ventricles of 25 schizophrenic and 28 control subjects. Their analysis showed that significant shape differences occurred in the left lateral ventricles. In particular, there was enlargement of the superior and posterior horns. There were also noticeable differences in the part of the lateral ventricles in the vicinity of the caudate head.

The results of the morphometric findings of this thesis will be compared with the above results in chapter 5.

## 1.4 Aims and Motivation

The fact that ventricular enlargement occurs in schizophrenia is now well established. Further work is needed to quantify and characterize shape and asymmetry differences. However, as noted by Ballester [7] volume is not a powerful shape descriptor - as two objects with the same volume can have quite different shapes. Furthermore, volumetric methods usually give global differences in shape and it is not easy to obtain information on specific localised shape changes. Gerig et al. [54] showed that shape measures reveal new information in addition to size or volumetric differences, which might assist in the understanding of structural differences due to neuroanatomical diseases.

A number of groups have recently proposed methods of morphometric analysis applicable to neuroanatomical structures, such as [87] and [24] mentioned in section 1.3.2. We propose that PDMs offer advantages in carrying out this morphometric analysis because they allow separation of shape changes due to disease in the presence of natural variation, and characterise shape by a small number of “modes of shape variation”, providing a compact parameterisation.

Our aims are :

1. To construct a 3D PDM of the lateral ventricles of a group of schizophrenic and control subjects
2. To use the PDM to investigate difference in shape and asymmetry associated with schizophrenia in a quantitative manner
3. To demonstrate the use of PDMs to facilitate quantitative and qualitative investigations of specific localised shape differences



## 1.5 Overview of Thesis

The work described in this thesis begins from the point at which MRI images of the subjects were obtained. The work done was in three main parts - the processing of the MRI images to obtain the structure of interest (the lateral ventricles), the building of the PDM of the lateral ventricle, and the application of the PDM to morphometric analysis. Each chapter (except the conclusions) begins with an overview of its contents and ends with a summary of the main points discussed.

Chapter 2 gives a background to image processing and shape analysis. Chapter 3 presents the MRI data we used and describes the preprocessing we applied to obtain the lateral ventricles. Chapter 4 gives a comprehensive description of PDMs and describes our approach to the construction of the PDM of the lateral ventricles. Chapter 5 shows the PDM can be applied to morphometric analysis, and gives the results of its application to investigation of shape and asymmetry differences between the control and schizophrenic groups. Chapter 6 concludes the thesis by summarising the main contributions made by this work.

Three appendices are included at the end. The first gives further details of some techniques mentioned in the text. The second contains images of all the lateral ventricles used. The third gives details of the files on the CD accompanying this thesis. This CD contains additional images and movies relevant to the work presented.

# Chapter 2

## Background

### 2.1 Overview

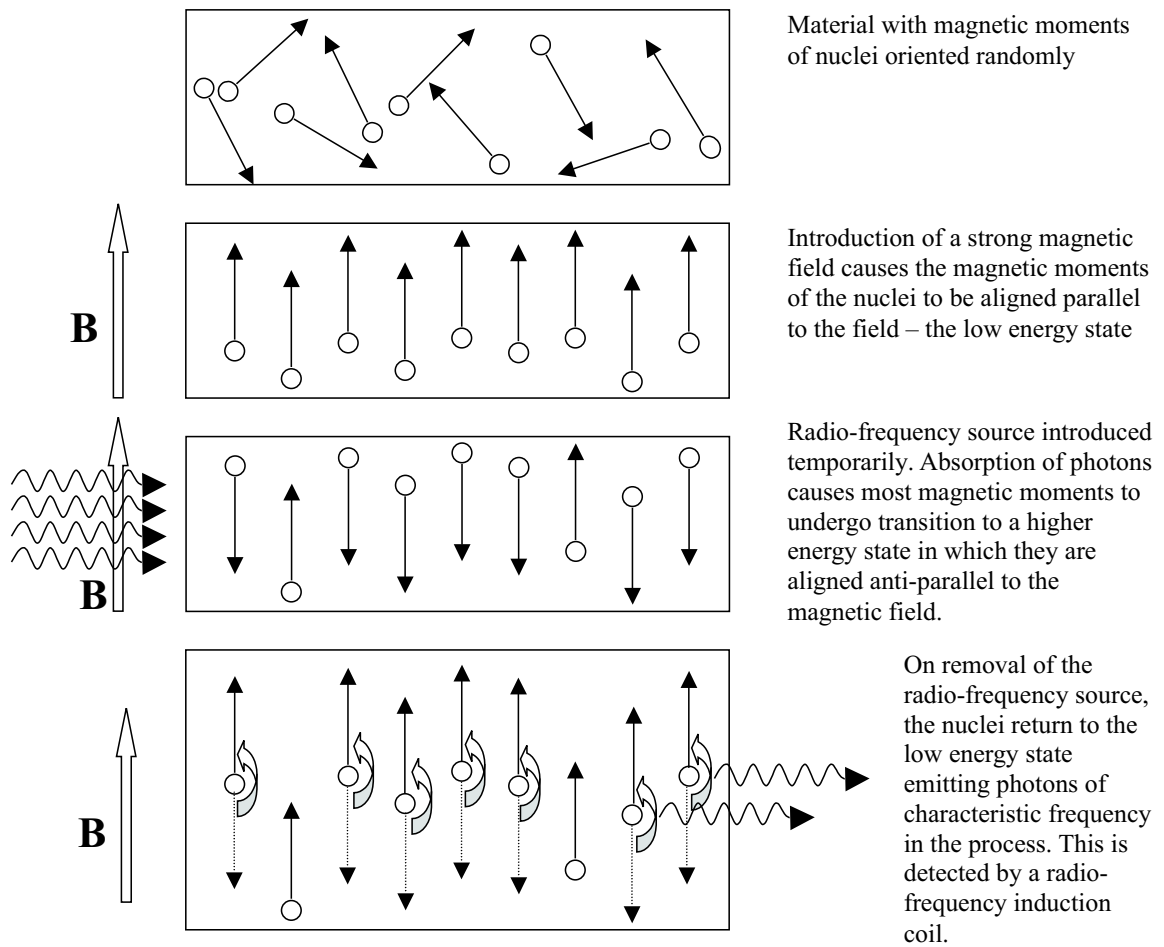
In this chapter we give background details relevant to the work involved in this thesis. MRI was used to obtain cross-sections of the brains of the subjects involved in the work we describe, and section 2.2 is a brief overview of the theory of MRI. The output of MR scans are arrays of values (gray levels) proportional to image contrast within the subject being imaged. In the 3D case these can be represented as 3D arrays referred to as volumetric images. Section 2.3 introduces volumetric images. Section 2.4 discusses the image analysis techniques used in the pre-processing of the images. It should be noted that the techniques are applicable to images obtained by other modalities. Section 2.5 discusses shape analysis in medical imaging, and introduces statistical shape models (SSMs) and PDMs. Section 2.6 is a summary of the main points of the chapter.

## 2.2 Magnetic Resonance Imaging

In this section we give only the briefest of descriptions of MRI. Standard texts such as [132] can be consulted for more detailed information. MRI is based on the magnetic properties of atoms by exploiting the principle of nuclear magnetic resonance (NMR) discovered in 1945 by Felix Bloch and Edward Purcell. NMR is a phenomenon in which the nuclei of atoms of a material placed in a strong magnetic field absorb radio-waves of certain frequencies.

A rigorous description of NMR requires the use of quantum mechanics, however, for our purposes a simplified classical description is sufficient. In classical physics, the nucleus of an atom is viewed as a concentrated spinning positive charge. Associated with this is a nuclear magnetic moment indicating the direction and strength of the field generated by the nucleus. In a non-magnetic material the direction of the nuclear magnetic moments of its atoms are oriented randomly, giving rise to no net magnetisation. The introduction of a strong magnetic field causes the nuclear magnetic moments to become aligned parallel to the direction of the magnetic field and they are said to be in a low-energy state. Introduction of electromagnetic radiation with frequency in the radio-wave range causes the nuclei to absorb photons of specific frequency and move to high-energy states in which their magnetic moments are aligned anti-parallel to the direction of the external magnetic field. Removal of the source of radio-frequency energy results in a proportion of the nuclei in the high energy state returning to the low energy state, and in the process emitting photons of their characteristic resonant frequency. These can be detected by a radio-frequency induction coil. Figure 2.1 summarises the details given in this paragraph.

Hydrogen atoms are more amenable to NMR than most others, and given that they are constituents of water and organic molecules they have a relatively high concentration in the body. The characteristic resonant frequency given out by each hydrogen atom is determined by the other types of atom it is bound to and its local intracellular environment. This results in MRI giving good anatomical and functional




**Figure 2.1:** Schematic illustration of the use of nuclear magnetic resonance in the generation of signals for MRI

contrast of cross-sections of the body. Further enhancements to the images can be obtained by using different protocols for the introduction and detection of the radio-frequency waves. Two of these are the so called  $T_1$ -weighted and  $T_2$ -weighted images.  $T_1$ -weighted images are acquired when the net magnetisation in the direction of the external field returns to 63% of its original value, whilst  $T_2$ -weighted images are acquired when the magnetisation in the direction transverse to the external field returns to 63% of its original value.

To obtain a cross-sectional image through the body in a particular plane, a gradient field is applied in a direction orthogonal to the plane, whilst the strong magnetic field causing NMR is applied parallel to the plane. The gradient field is used in conjunction with a readout gradient to prevent measurement of radio-frequency signals from nuclei not in the plane of interest. Standard tomographic (slice) techniques such as measurement of signals from different angles are applied to reconstruct point-by-point values in the cross-section of interest. Figure 2.2 is a schematic diagram of a complete MRI system. The data acquired by the MR scan is exported in digital format. Most modern scanners export data conforming to the DICOM standard for radiologic images [1].

## 2.3 Volumetric Images

As mentioned in the previous section, the data resulting from a MR scan are tomographic (slice by slice) cross-sections of the body. Each slice represents a 2D cross-section in which the image is composed of discrete picture elements (pixels - see Figure 2.3) of specified size and with a specified *gray level value* indicating its relative contrast within the image. It should be noted that these 2D slices actually represent cross-sections with finite thicknesses, although this is not evident when they are displayed. The cross-sections could be in axial, sagittal or coronal orientations as shown in Figure 2.4.




**Figure 2.2:** Principal components of a complete magnetic resonance imaging system (source [132])


Volumetric 3D images are sequentially ordered stacks of 2D slices which may have been acquired contiguously, or with a specified interval between consecutive slices (an inter-slice gap). The discrete unit of volumetric images are volume elements (voxels), which are basically the pixels of a 2D slice attributed an extra spatial dimension to account for the slice thickness. Figures 2.3(c) and 2.4 show a schematic diagram of voxels and cross-sections through a volumetric image. In the following subsections, we discuss some properties of volumetric images, and some operations that are routinely performed on them.

### 2.3.1 Image Resolution and the Partial Volume Effect


The spatial resolution of 3D images is determined by the size of the voxels. The stronger the external magnetic field the greater the spatial resolution can be. High resolution images are desirable because they reduce the probability of several different types of tissues contributing to the signal in a particular voxel - the *partial volume effect*. They also allow greater clarity in the images and reveal greater anatomical



(a) 2D coronal cross-section through the head




(b) 2D images are grids of discrete picture elements (pixels). The resolution of the image is determined by the size of the pixels



(c) 3D images are arrays of cuboid or cubic volume elements (voxels). These are formed from consecutive 2D slices. In-plane resolution is determined by the size of the pixels in the 2D slices, and out of plane resolution is determined by the slice thickness


**Figure 2.3:** Sub-figures illustrating a 2D cross-section (a) and the discrete nature of images in 2D (b) and 3D (c)



(a) Coronal Orientation



(b) Axial Orientation



(c) Sagittal Orientation

**Figure 2.4:** The three orthogonal orientations of MRI cross-sections through a volumetric 3D image of the head



detail. This is especially important when performing morphometry.

However, several practical factors have to be considered in determining the resolution of a 3D image. Firstly, there is an interplay between the slice thickness, the acquisition time and the signal-to-noise ratio (SNR) in the acquired images. The greater the slice thickness the greater the SNR, but the lower the out-of-plane resolution. The longer the acquisition time the greater the SNR. Because the body is a dynamic system the acquisition time has to be kept small. Therefore in many cases the slice thickness is increased to improve SNR.

Secondly, the size of the resulting images and the storage requirements are also considered when determining the resolution of images. The higher the resolution within a slice (in-plane resolution) and the smaller the slice thickness (out-of-plane resolution), the greater the number of voxels in the slice and therefore, the storage requirements increase. For these reasons many MR images are acquired with lower out-of-plane resolution, than in-plane resolution, leading to cuboid rather than cubic voxels - voxel anisotropy.

### **2.3.2 Re-sampling and Interpolation**

Re-sampling is simply redivision of the voxel grid of an image to give voxels of another size. Interpolation is necessary to ensure the distribution of gray level values in the re-sampled voxels is similar to that of the original image. Re-sampling and interpolation may be required to convert anisotropic voxels to isotropic voxels before application of specific image analysis techniques. They are used in re-slicing volumes to allow visualisation through them in specified directions. They are also required during image registration as discussed in section 2.3.3.

The most commonly used method for interpolation is trilinear interpolation. In this the value of a particular voxel is obtained by distance-weighted linear interpolation of the values of its neighbouring voxels. Trilinear interpolation is popular because

it produces relatively good results and is computationally efficient. A faster method is the nearest neighbour method. This simply substitutes the gray level value at a particular voxel of the re-sliced image with that of the nearest voxel in the original image. This method produces low quality results. However, interpolation using cubic spline functions gives smoother transitions between voxel values because of the continuity properties of splines. Interpolation using sinc functions addresses issues such as the sampling rate, and if implemented correctly has the potential to produce high quality re-sliced images. Unfortunately, sinc interpolation requires convolution of images with large kernels to produce accurate results. Although the computational requirements of this limits its application to 3D images, there is an active body of research to address this e.g. [114]. A comprehensive survey of image interpolation methods is available in [76].

### 2.3.3 Image Registration

Image registration is the process of bringing two images (in our case 3D volumetric images) into spatial correspondence. It involves applying geometric transformations to one of the images to optimise a metric indicating the similarity of the two images. Re-sampling and interpolation are needed to produce a registered image with appropriate voxel sizes and gray level values.

The three classes of geometric transformations and the main image similarity metrics of relevance to this thesis are described below. We are interested in registration of images of the same subject acquired with the same modality using voxel based similarity measures, and as such the discussion below is limited to these. [81] gives a comprehensive review of medical image registration. It should be noted that the term “registration” is sometimes used generically to refer to image registration as well as to the geometric transformation of point-sets. However, which is meant should be clear from the context.

#### Geometric Transformations

*Rigid-body transformations* are transformations in which global rotations and translations are carried out and in some cases includes isotropic scaling. In the 3D case there are 6 degrees of freedom (rotation about the  $x$ ,  $y$  and  $z$  axes, and translations in the  $x$ ,  $y$  and  $z$  directions). These transformations preserve all Euclidean distances and angles. If isotropic scaling is included there are 7 degrees of freedom. In this case distances between points change, but relative distances and the angles do not change.

*Affine transformations* give a greater degree of freedom than rigid registrations by allowing anisotropic scaling in the  $x$ ,  $y$  and  $z$  directions, and the application of *shears* in each of these directions. If anisotropic scaling is the only addition to the 6 degree of freedom rigid registration case then the affine registration has 9 degrees of freedom. However, if both anisotropic scaling and shearing are included it becomes a 12 degree of freedom affine transformation. Affine transformations do not preserve distances or angles, but do preserve the linearity of structures (i.e. straight lines remain straight). Details of rigid and affine transformations are given in [52].

*Non-rigid transformations* involve application of different transformations at different spatial locations of an image in a localised manner. This results in a greater number of degrees of freedom than rigid and affine transformations. A wide variety of methods exist for performing non-rigid transformations, and a comprehensive review is given by Toga [122]. In this thesis we implement a spline-based method outlined in chapter 14 of [122] to perform our non-rigid transformations.

### Image Similarity Metrics

The main voxel-based similarity metrics used in volumetric image registration are the *Sum of Squared Distances* (SSD), the *Normalised Cross Correlation* (NCC), the variance of intensity ratios (also called the *Woods Function*), and *Mutual Information*. Image similarity metrics are described in more detail in chapter 6 of [6].

The SSD is the simplest of the similarity metrics. It simply gives a measure of the global difference between the voxel values at corresponding spatial locations in two

images. However, it does not take into account differences in voxel values that may arise due to different scales being used in the two images or offsets in voxel values. In the registration of two images,  $A$  and  $B$ , the SSD is given by:

$$\frac{1}{N} \sum_{i=1}^N |A(i) - B(i)|^2 \quad (2.1)$$

Where  $N$  is the number of *overlapping* voxels in corresponding spatial locations in  $A$  and  $B$ .

The NCC is a relatively straightforward metric that is invariant to scale and offset differences. It gives the correlation between the values of all pairs of voxels at corresponding spatial locations in the two images being compared. This correlation is normalised by the geometric mean of the variance of voxel values in each image. The disadvantage of NCC is that it assumes that the voxel intensities in an image can be represented by a linear function which is not necessarily true. The NCC is given by:

$$\frac{\sum_i ((A(i) - \bar{A})(B(i) - \bar{B}))}{(\sum_i (A(i) - \bar{A})^2 \sum_i (B(i) - \bar{B})^2)^{\frac{1}{2}}} \quad (2.2)$$

The Woods function addresses non-linearity of intensities in the two images being compared. It requires that the reference image is partitioned into areas of similar intensity (isoets). The image being registered is then partitioned into regions using the boundaries of the isoets of the reference image. The variance within each region is calculated and the sum of the variance for all regions is normalised by the mean of their intensities to give the Woods function, i.e.

$$\sum_{a \in I_A} \left( \frac{n_A(a)}{N} \frac{\sigma_B(a)}{\mu_B(a)} \right) \quad (2.3)$$

Where

$I_A$  contains of all the isoets in image  $A$ , and  $a$  refers to a particular isonet.

$n_A(a)$  is the number of voxels in the  $a^{th}$  isonet

$\sigma_B(a)$  is the standard deviation of the voxel values in the region in image  $B$  corresponding to the  $a^{th}$  isoset in  $A$ , and  $\mu_B(a)$  is the mean of the voxel values in this region.

The Mutual Information cost function is a more general similarity metric. It is based on the concept of the entropy of an image. The entropy of an image is a measure of the dispersion of the probability distribution of the voxel values of the image. The measure of Mutual Information is given by the difference between the joint entropy of the two images being registered and their individual entropies. The mutual information is sometimes used in normalised form. The normalised mutual information (NMI) has several definitions - see [6], [81]. One of these is the joint entropy of the two images normalised by the sum of their individual entropies. i.e.

$$\frac{H(A) + H(B)}{H(A, B)} \quad (2.4)$$

Where the entropy of image  $A$  is given by

$$H(A) = - \sum_{i \in I_A} p(a) \log p(a) \quad (2.5)$$

$I_A$  being the range of intensities in image  $A$ , and  $p(a)$  the probability of the occurrence of the intensity in  $A$  - The proportion of voxels in  $A$  with this intensity. The entropy for  $B$  and the joint entropy,  $H(A, B)$ , are calculated in a similar manner.

## 2.4 Image Preprocessing

Raw images obtained from a MRI scanner usually require preprocessing before being amenable for their desired purpose. In the following subsections we discuss the correction of MR image intensity inhomogeneity, and image segmentation to obtain particular structures of interest.

### 2.4.1 MR Intensity Non-Uniformity Correction

Intensity non-uniformity is the smooth variation in intensity present in MR images. This can occur within a slice (intra-slice non-uniformity) or between slices (inter-slice non-uniformity). Intensity non-uniformity is an artifact of the techniques used in MR imaging. It is due to factors such as radio-frequency excitation field inhomogeneity, non-uniform reception coil sensitivity, eddy currents driven by field gradients as well as electrodynamic interactions with the subject being imaged.

Correction is needed because the effects may be severe enough to impair visual inspection of a scan. With newer MRI scanners the effects may be subtle enough not to be detected on visual inspection. However, they can affect automated image analysis techniques that assume homogeneity of intensity within each tissue type, for example data driven segmentation. Despite the need for correction of intensity inhomogeneity the problem is sometimes assumed to be inconsequential and not dealt with. This may be valid in cases where processing such as segmentation is done manually.

In correction of non-uniformity there have been several approaches. Dawant et. al. [40] propose a direct fit method. This requires the identification of specific points within the image to be corrected. The intensity of the image is then modelled to give a profile of the radio-frequency induced inhomogeneity. This is subtracted from the image to give a corrected version. This method has the disadvantage that it requires substantial user intervention. The authors proposed an automated version but the results were poorer than those of the direct method.

Wells et. al. [129] in implementing an automated method of segmentation correct for inter-slice and intra-slice non-uniformity use a priori knowledge about the tissues being imaged. They assume each pixel value in an image is given by a composite function of a gain field and a tissue field. If the gain field is known the tissue field can be determined and conversely the gain field can be obtained if the tissue field is known. This method requires a priori knowledge of the tissues being imaged and the

implementation is dependent on the types and distribution of the tissues.

Sled et. al. [103] describe an approach not requiring a model of the tissue classes present, in which intensity non-uniformity is modelled as a smooth multiplicative field. The measured signal giving rise to image intensity is assumed to be a sum of three components - the true signal, multiplicative noise, and additive Gaussian noise. Using the image intensities in a slice, an iterative algorithm is applied to estimate the multiplicative bias field and the distribution of true tissue intensities. Sub-sampling of the image is necessary to reduce the computational time for the algorithm.

Vorkuka et. al. [127] describe a technique to generate estimates of non-uniformity maps for correction of inter-slice and intra-slice non-uniformity. Their approach to non-uniformity correction was based solely on the assumption that the various sources of non-uniformity give rise to smooth variations in image intensity, and that these variations can be extracted and corrected for. The algorithm used is

1. Estimate the additive noise in the image
2. Estimate the normalised local intensity gradients in  $x$  and  $y$  directions in the slice
3. Estimate the smooth local derivatives using statistical averaging
4. Re-integrate the derivatives to determine a “relative” non-uniformity map
5. Computation of the “true” non-uniformity map  $G(x, y)$

### 2.4.2 Segmentation

Neuroanatomical segmentation is defined as “the extraction of a specific, precise and comprehensive 3D morphological description of the neuroanatomical structure of the subject’s brain that is obtained robustly and practically from volumetric data” [135]. Different neuroanatomical structures in the brain appear as varying combinations of

GM, WM and CSF depending on the types of cells they are made of. Segmentation at one extreme can be considered as the classification of voxels into one of these types based on their intensity values. However, in practice, segmentation requires a priori knowledge of characteristics of the structure of interest.

Several issues, some particular to MR imaging, result in segmentation not being a straightforward task. These include the structural complexity of the brain and neuroanatomical structures, the variability between brains of individuals, the partial volume effect, and the unavailability of a “Gold Standard” to allow quantitative assessment of the results of segmentation.

### **Manual, Automatic and Semi-automatic Segmentation**

Segmentation could be done manually, automatically or semi-automatically. In manual segmentation, an independent rater e.g. a radiologist examines an image and interactively delineates the boundaries of the structure of interest based on their knowledge and experience. Manual segmentation has disadvantages in that it is subjective and uses human resources that could be better applied elsewhere. Despite these disadvantages, it is still used because it is usually straightforward and produces more acceptable results than automatic methods.

Automatic segmentation offers the benefits of freeing valuable human resources and overcomes the problem of the requirement for segmentation being a bottleneck in clinical research. However, at present fully automated methods of segmentation are usually tailored for specific goals and are not used clinically. Semi-automatic segmentation combines manual and automatic methods and is widely used.

### **Approaches to Automatic Segmentation**

As with most of the information in this background chapter, segmentation is an entire topic in its own right and a full review is beyond the scope of this thesis. Chapter 5 of [105] provides a thorough overview of segmentation.

The most basic of automatic methods is Thresholding. This involves setting a thresh-



old on voxel values either using a priori knowledge of the expected voxel values within the structure of interest or using adaptative techniques to determine suitable thresholds. However, coming up with a good adaptative scheme is not trivial. Morphological operations such as opening and closing are usually required to give acceptable results. Another class of methods using gray level information are region based methods which attempt to segment an image by modelling their gray levels according to predefined functions. Statistical Methods also use gray level values. In the case of neuroanatomical segmentation the aim is usually to classify voxels as belonging to GM, WM or CSF by fitting them to probability distributions.

Edge-based methods rely on the assumption that the boundaries between different structures give rise to local differences in voxel values that can be detected. The algorithm of Canny [25] is very popular with these methods. However, the major problem faced is how to convert edge elements (edgels) into meaningful boundaries. Watershed algorithms provide a solution to the boundary conversion problem. Watershed algorithms transform an image into a disjointed set of catchment basins each of which can be imagined to be filled with water until it reaches its watershed. The image is segmented into regions in this manner. Watersheds can be applied to gray level images, but produce better results when applied to images in which the edges have been enhanced.

Model based methods come in a variety of forms. They include the use of shape models, which are discussed in section 2.5. They offer the advantages of allowing the incorporation of *a priori* anatomical information to the problem of segmentation. They also offer in many cases (such as snakes - [68]) an intuitive interface for manual interaction to modify boundaries, making them suitable for semi-automatic segmentation.

A growing area of research in segmentation is the use of anatomical atlases [119]. These are high resolution digital images of specified regions of the body for instance the brain, which have been labelled in detail. To segment an image of the same

anatomic region (for instance that of another brain), a non-rigid registration process is required. This is used to bring the two images into spatial correspondence. The labels from the atlas can then be transferred to the image to be segmented.

### **Accuracy of Segmentation**

The verification of the accuracy of the result of segmentation of images of structures within the human body cannot be carried out directly because there is no non-invasive method of determining their exact boundaries - no “Gold Standard”. Verification is therefore performed by a radiologist or other suitably qualified individual in a subjective manner. The variation from person to person on the same set of images should be small and statistical confidence in the assessment can be increased by using a number of independent raters. However, such assessments are time consuming, repetitive and laborious. An alternative approach, which provides scope for limited quantitative evaluation of automatic methods, is to use test images where the ground truth is known. Images such as those provided by the simulator of the Montreal Neurological Institute [30] can be used for this purpose.

## **2.5 Shape Analysis in Medical Imaging**

In this section we describe what is meant by shape and shape analysis, concentrating on 3D applications to neuroanatomical structures of the brain. We discuss methods used in 3D shape analysis, and explain why the statistical shape model approach was chosen. Section 2.5.1 describes shape from a computational viewpoint and illustrates this using hand outlines. Section 2.5.2 introduces shape analysis and the representations of shapes. Section 2.5.3 reviews methods used in 3D shape analysis of neuroanatomical structures, and section 2.5.4 discusses the effect of head shape on the analysis of the shape of neuroanatomical structures in the brain.

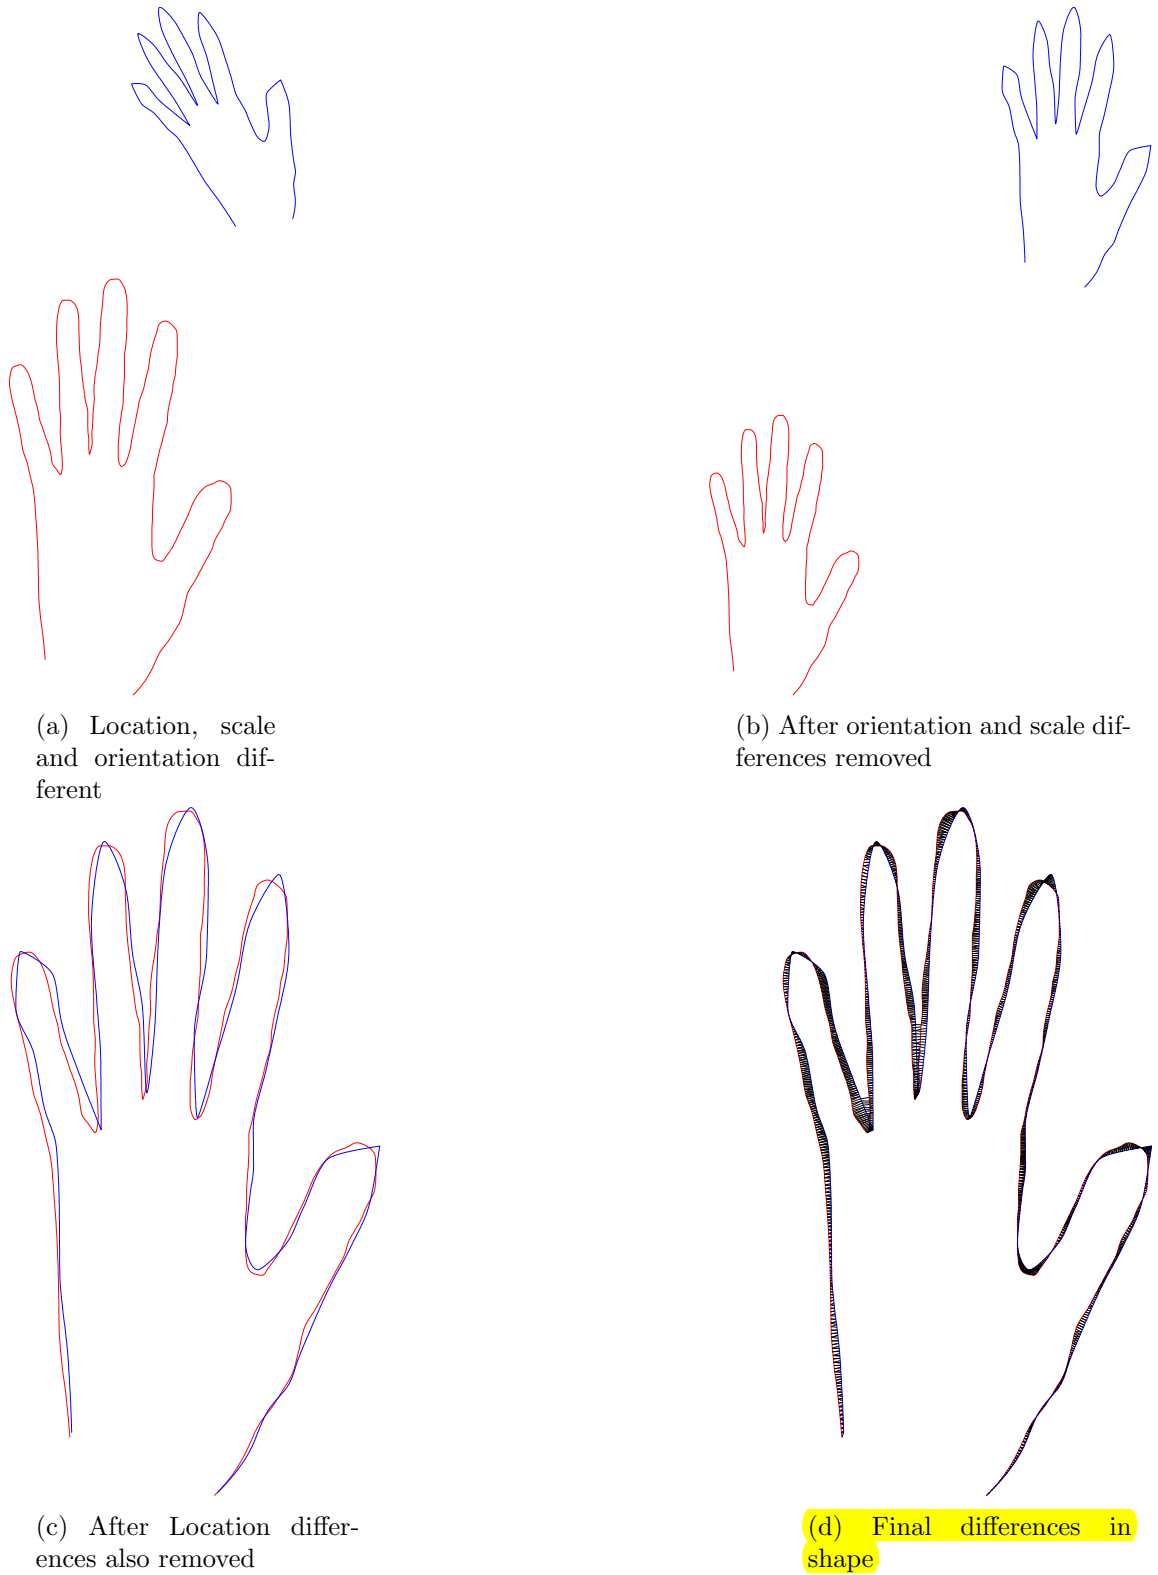
### 2.5.1 Shape

Objects of the same class e.g. hands, have a common physical form which is generally referred to as the shape of the class. The shape of an object comprises both its topology and its intrinsic geometry (relative sizes and orientations of its various sub-parts). We take our formal definition of shape as “All the geometric information remaining when translation, scale and rotational effects are filtered from an object” [44]. Hence, by definition, shape is invariant under Euclidean transformations of translation, isotropic scaling and rotation.

Figure 2.5 is used to illustrate what is meant by the formal definition given above. It shows two versions of the left hand of an individual (the original outline and a re-sampled and interpolated version with random noise added) drawn to scale at different locations and with different orientations (subfigure (a)). In this case location is removed by translating so that the centroids are coincident; orientation is removed by alignment of the principal axes; and scale is removed by approximating ellipsoids to the convex hulls of each and scaling the major axes of one of the hand outlines to be the same as that of the other. The remaining differences between the hands are shape differences.

### 2.5.2 Shape Analysis

In everyday situations the shape of individual members of a class are qualitatively described relative to an idealised notion of the shape of the class. In the same way shape analysis is concerned not with measurement of actual shape, but with investigating the extent of *differences* in the shape of two or more objects of the same or different classes, depending on the application. Shape analysis is used in many other fields for example biology [45], gait recognition [128], and character recognition [97]. The development of the application of multivariate statistical methods to shape analysis started with the independent work of Kendall [70], Bookstein [14], and Ziezold [138].



**Figure 2.5:** An example to illustrate shape differences using two different outlines of the hand

Shape analysis methods applied to medical imaging can be classified in a variety of ways. One such classification is into *statistical shape analysis methods* and *non-statistical shape analysis methods*. Non-statistical methods generally perform shape analysis on individual members of a sample of a population with respect to one member of the sample. Statistical methods on the other hand obtain a mean of the sample and perform shape analysis of individual members using the mean. In the case of structures within the human body, there is great natural variation in addition to variation due to the factor(s) of interest. The use of statistical methods allow more robust evaluation of shape differences to be made as they allow for natural variation.

To perform shape analysis the following have to be done: Physical representations of the shapes being analysed have to be obtained and transformed to numerical or structural representations in which variations due to Euclidean transformations have been removed. The removal of variations due to Euclidean transformations involves establishing *correspondence* (identification of homologous parts) between the shapes being analysed, and performing *registrations* (spatial transformations) to attain Euclidean invariance using the observed correspondences. From these Euclidean invariant representations, quantification of global and/or local shape differences have to be determined.

It is also desirable to obtain statistical inferences of the significance of any computed differences. Additionally, an intuitive presentation of the observed differences, although not an essential requirement, is very useful as it allows greater understanding of the observed differences, and determination of whether they make intuitive sense.

A number of shape representations are considered below, along with the ways of identifying correspondences and performing registration to obtain Euclidean invariance. There are a diverse variety of shape representations, but the review below is limited to 3D representations used in medical image analysis. Although it is convenient to consider the shape representations separately, in practical applications different representations may be combined. Examples of application to medical image analysis

are given in section 2.5.3.

### **Voxel-Based Representations**

Voxel-based representations are 3D arrays of values. At one extreme they could simply be volumetric image given as the output of magnetic resonance (MR) or another imaging modality over a specified section of the body. At the other extreme they could be binary values representing the surface or volume of specified structures or regions. In between these extremes they could also represent transformations of the output of an imaging modality e.g. distance transforms and probability maps.

In this section the first case, in which voxel values represent the output of an imaging modality, is of interest. The other two cases are usually used in the preliminary stages of the formation of other shape representations as described in the subsections below. The first case has the greatest information content as the other two cases are derived directly or indirectly from it. Correspondence can be measured using image similarity metrics as described in section 2.3.3. Euclidean invariance can be established by applying non-rigid registrations to minimise the chosen image similarity metric. Landmarks (see 2.5.2) can also be used to augment the non-rigid registrations. In the case of brain images a common method of obtaining correspondence is to transform the images into the standard Talairach coordinate space. This is a coordinate system within the brain defined by Talairach [110] which is widely used in neuroanatomy and available in digital format. Because of the method of application these representations are sometimes termed deformable anatomical templates (not to be confused with deformable shape templates discussed in section 2.5.2).

Advantages of this form of representation include the fact that it is derived directly from the output of the imaging modality and needs little or no manual intervention. Additionally, when used in shape analysis of brain data it allows the investigation of the shape of multiple structures within images. However, it has disadvantages in that the use of voxels involves applying complex non-rigid registrations to large

number of elements, the validity of the warps is not guaranteed, statistical analysis is complicated (due to the number of elements), and visualisation of the resulting shape differences is not straightforward. Voxel-based morphometry is a popular method that uses this representation. This is discussed in section 2.5.3

## **Surface Based Representations**

Structures within the body are three-dimensional, and their boundaries are therefore surfaces. The methods considered in this subsection rely solely on the use of the surfaces of the objects whose shapes are being analysed. This means that a segmentation step must have occurred, and this can be a disadvantage in some situations. However, surface based representations are more amenable to mathematical treatment than voxel based representations. Three categories are reviewed here.

### **Triangulated Meshes**

Triangulated meshes are approximations of a surface using vertices on the surface and a notion of the connectivity of the vertices to give polygonal faces. The marching cubes algorithm of Lorensen and Cline [78] and a Delaunay based method [13] are popular triangulation methods.

Triangulated surfaces are mainly used for visualisation purposes and as a starting point for other shape representation methods. However, correspondences can be sought by applying point based techniques such as the iterative closest point (ICP) method [9] and rigid or non-rigid transformations to minimise Euclidean distance between closest points. Unfortunately, correspondences sought in this manner are usually crude as they are not based on any physical properties of the surfaces, and are sensitive to outliers.

### **Surface Parameterisation**

Parameterisation gives a continuous representation of a 3D surface, with each surface point being uniquely represented by parameters specifying a point on a 2D manifold.

The parameterisation can be considered a mapping that must be one-to-one and “onto” (each point on the surface is mapped onto a point on the manifold) for which an inverse exists. Ideally, both the mapping and its inverse should be differentiable. It is also desirable that the magnitude of shape differences in parameter space be linearly related to the shape differences in Euclidean space.

We make the assumption that the surfaces to be parameterised represent whole structures in the body and are therefore closed surfaces. This assumption is quite valid in practice, and is important, as the problem of parameterisation of a closed surface is not as straightforward as that of the parameterisation of a surface patch because of the requirement here of one-to-one correspondence.

Brechtbühler [21] described a method of surface parameterisation using spherical harmonic functions defined on a unit sphere. Firstly, the surface to be parameterised is triangulated and mapped to a unit sphere. In Euclidean space the spatial relationship of vertices resulting from the triangulation can be viewed as a surface net or a graph. The mapping to parameter space is posed as an optimisation problem, with a cost function that minimises the distortion of the surface net whilst embedding it on the surface of a unit sphere. Secondly, the shape of the surface in parameter space is represented as the summation series of the inner product of the vertex points in parameter space and the spherical harmonic basis functions defined on the unit sphere. This gives the parameterised representation of the surface. It is assumed that higher frequency spherical harmonics contain noise, and a cutoff is selected so that harmonics with frequency above the cutoff are not included in the parameterised representation.

Location invariance is obtained by using object centred coordinates in Euclidean space e.g. by using the centroid as the origin of the coordinate system. Orientation invariance is obtained by rotations in parameter space to align the axis of the first spherical harmonic terms which are ellipsoids. This alignment is also assumed to establish correspondence. Scale invariance is implicitly obtained as all surfaces are



mapped to a unit sphere. Spherical harmonic shape descriptors, which are Euclidean invariant representations of the surfaces, are obtained from the parameterisations and are analysed to obtain differences in shapes.

Surface parameterisation methods allow continuous description of surfaces and arbitrary re-sampling, and the method of Brechbülher has been applied to a number of shape analysis problems e.g. [54]. However, current surface parameterisation methods have topological restrictions which limit their use. Such restrictions include the requirement that surfaces have no folds or self intersections. Additionally, there is not an intuitive relationship between the shape differences obtained from the shape descriptors in parameter space, and actual changes in Euclidean space.

3D Fourier transforms for surface parameterisation have also been described e.g. [136]. However, the only example of the use of 3D Fourier Descriptors in shape analysis [102] involved a modification of the spherical harmonic method of Brechbülher in which surfaces are mapped to spheres using an area preserving method, and shape descriptors were combined Fourier and spherical harmonic terms. Zhang [137] describes the parameterisation of 3D surfaces using harmonic maps. However, this is limited to disc like surfaces.

### **Semi-Parameterisation of Surfaces**

Spline functions [42] are popularly used to represent curves and boundaries in 2D. They are piece-wise polynomial functions defined over a domain of control or “knot” points that are continuously differentiable to a specified order. The control points are the non-parametric portion as they have to be specified in advance. Splines have a number of desirable properties - differentiable continuity (smoothness), local shape controllability and invariance under certain transformations. Cubic splines, Bezier curves and surfaces, and non-uniform rational b-splines (NURBS) are of particular interest in shape representation. Shape descriptors have been defined in 3D using NURBS [23]. Bezier functions have been used in visualisation of surfaces [29].

During the literature review, we have encountered no examples of 3D splines being

used directly as shape representations in shape analysis. Reasons for this may include the requirement that changes in the spline parameters be linearly related to changes in the shape of the actual surfaces being represented. Additionally, the creation of spline surfaces may require the discretisation of the surface into patches, or the manual specification of surface correspondences to allow registrations to remove location, scale and orientation differences. However, because of their smoothness and local control properties, spline based functions, for example thin plate splines [15], are widely used in non-rigid registrations of other shape representations.

### **Deformable Shape Templates**

Deformable shape templates start with a primitive surface (the template) which is deformed to fit the surface of interest. They are closely related to deformable models whose use in medical image analysis was introduced by Terzopoulos [113]. In the present case the primitives are closed surfaces or 3D materials of simple geometry. Arbitrary correspondences are defined based on properties of the surface to be represented, and the primitive is deformed to closely match the surface being represented. Quantification of shape differences is based on analysis of the deformations required to match the primitive to the different surfaces being represented. A number of primitives and methods of deformations have been described. Two of these (superquadrics and physically based templates) are considered. The common drawback of these methods are the need to specify material properties of the template such as elasticity or stiffness, and the arbitrary ways in which correspondences are generated. Additionally, most of these methods tend to be computationally intensive. It should be noted that some representations that fall into the voxel-based class of section 2.5.2 in this thesis have been termed deformable templates elsewhere e.g. [61]. However, they are more precisely deformable anatomical templates.

### **Superquadrics**

Superquadrics (such as ellipsoids) are generalisations of quadric surfaces. They are

parametric surfaces that can model a wide variety of complex shapes realistically with a compact number of parameters. They allow controllability and intuitive meaning of their deformations, and their parameters can be obtained robustly.

The use of superquadrics as primitives is described by [112]. Initial correspondences (as well as location, orientation and scale invariance) are obtained from the centroids and central moments of the surface to be represented. To obtain the parameters for the surface, stiffness parameters are assigned to the superquadric template, then global and local shape parameters are obtained by solving equations of motion of the model under the action of externally applied forces (from the surface to be represented). The global parameters contain scale, location and orientation terms, whilst the local parameters allow quantification of local shape differences.

### **Physically Based Templates**

In this case the primitive is taken as a solid material with physical properties that determine its deformation. Pentland and Scarloff [91] use an ellipsoidally shaped elastic material whose elastic properties are specified according to the surface to be represented. Nodes on the template are attached to the surface by virtual springs. The point of attachment of the virtual springs determines the correspondences of the surfaces to be represented. To obtain these, the centroid and principal axes of the template are aligned with those of the surface being represented. The projection of points on the surface onto the template determine the points of attachment of the virtual springs.

A constant load due to the surface to be represented causes forces to be exerted at the nodes of attachment on the template, and these nodes vibrate until the system reaches an equilibrium. The nature of the vibrations and the nodal displacements at equilibrium are determined using the finite element method (FEM). Modal analysis is performed on the vector of the nodal displacements to give the modes of vibration. Shape invariant descriptors are obtained from these by discarding the modes corresponding to the lowest frequencies, which are the modes of translation, rotation and

scaling.

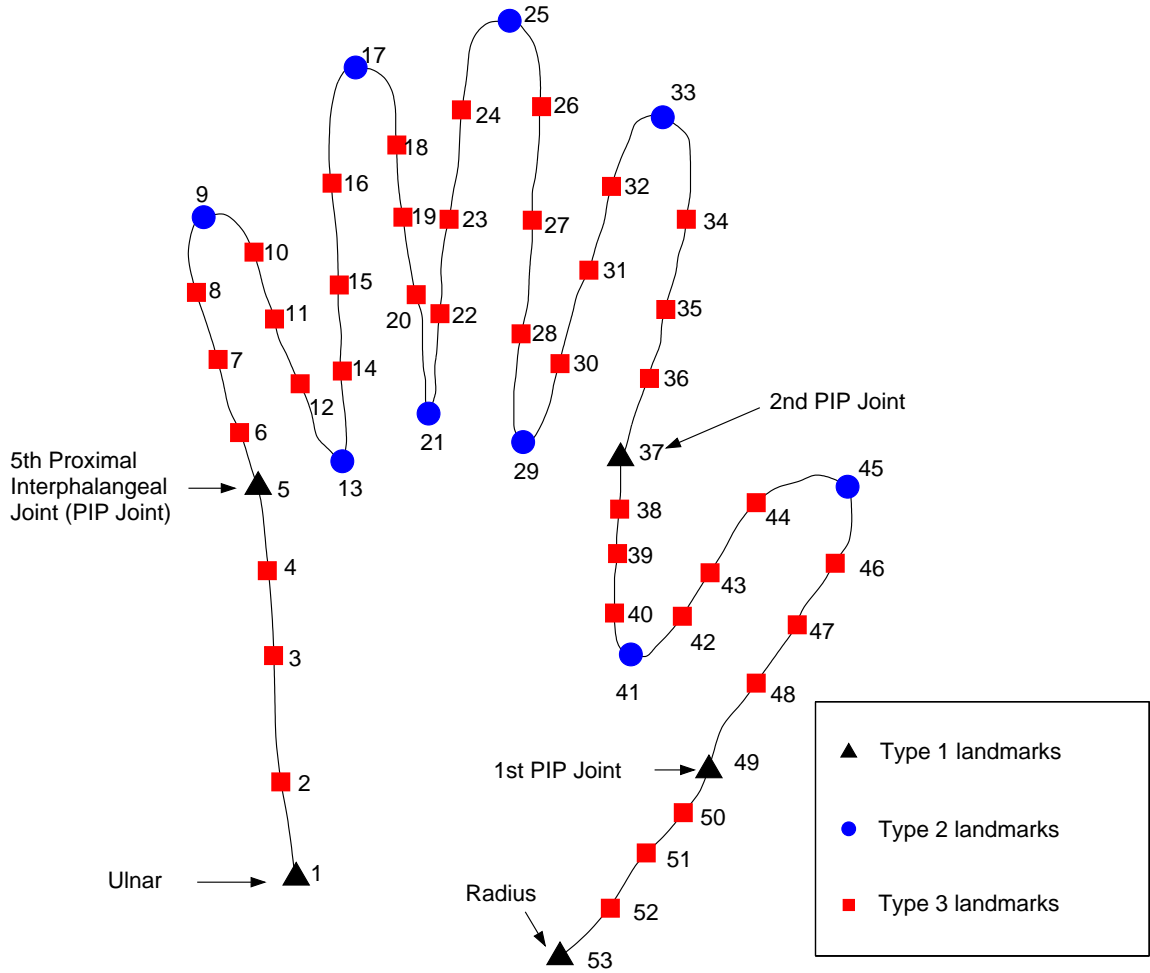
### Landmark (Point) Based Representations

Landmarks in the context of this thesis are points on a 3D surface or within a 3D volume chosen according to some predefined criteria. Landmarks have been categorised into three types according to Bookstein [16].

1. **Anatomical landmarks** which are points marking salient parts of an object with particular application dependent significance, e.g. the positions of the interphalangeal joints on an outline of a hand or some suitable point on the surface of a ventricle.
2. **Mathematical landmarks** which are points marking mathematical or geometrically meaningful features which are application independent, e.g. points of extreme curvature like the tips of the fingers.
3. **Pseudo-landmarks** whose positions can be interpolated from points of type 1 and type 2, e.g. placing a fixed number of points between each type 2 point on a boundary.

Figure 2.6 uses a 2D outline of the hand to illustrate examples of the three types of landmark points.

With landmark based representations correspondence is achieved by placing landmarks at homologous points of each of the shapes to be analysed. The same number of landmarks have to be placed on each of the shapes. There is no hard and fast rule setting the number of landmarks needed to represent a structure, however, they should be sufficient in number and placement to be truly representative of the shape of the original structure. Euclidean invariance is achieved by registrations of the shapes using the landmarks. For instance by applying Procrustes analysis [58] in



**Figure 2.6:** Illustration of classification of landmarks to anatomical, mathematical and pseudo-landmarks. The anatomical landmarks are the labelled anatomical points, the mathematical landmarks are the curvature extrema of the hand outline, and the pseudo-landmarks are sets of three equally spaced points between consecutive anatomical and/or mathematical landmarks (not drawn to scale in the figure)

3D using iterations of seven degree of freedom rigid body registrations to minimise a criterion based on the Euclidean distance between corresponding points.

The main advantage of landmark based representations is that there is a wealth of knowledge about their statistical analysis. They can be put into vector/matrix form and multivariate analysis techniques performed on them. The quantification and visualisation of shape differences is also straightforward as the differences can be represented in Euclidean space. The main disadvantage of landmarking methods is that it is impractical to manually place a large number of landmarks accurately on a boundary. This is particularly true of placing landmarks on a 3D surface. The numbers of landmarks needed to truly represent a 3D object is much larger than for 2D (usually starting at a few hundred landmarks). This is the “correspondence problem” and various methods of overcoming this are discussed in section 4.3.

### **Distance Transforms**

Given a binary image containing a region of interest or its boundary, a distance transform assigns a value to each image element giving its distance from the boundary of the region. A sign convention can be used to distinguish between the interior and exterior points. Borgefors [20] describes methods for obtaining distance transforms in 3D.

Golland et. al. [56] describe a method of obtaining Euclidean invariant shape descriptors of 3D binary data using distance transforms. The distance transform is inherently location invariant as it is object centred. Orientation and size invariance are obtained by applying rigid body transformations to align the centroids and moments of the objects to be analysed and to scale them to the same volume. The rigid body alignment is assumed to establish correspondence between the objects. The distance transforms of the aligned versions are recomputed and re-sampled to give feature vectors which are the shape descriptors.

This method of representation has advantages in that it is relatively straightforward and easy to compute, does not require manual intervention (except maybe for a segmentation step), and copes well with complex topology. However, the assumption that alignment of the centroid and moments of the distance transform establish correspondence is not justified, and may lead to invalid correspondences. Additionally the uniform sampling used to produce the shape descriptors may lead to insensitivity to some local shape differences.

### **Medial Axis Representations**

The terms medial axes and skeletons are sometimes used interchangeably in computer vision literature. Here we take them to mean the same thing. That is the locus of the centres of the maximal spheres touching the surface of the object being represented whilst being completely enclosed within the object. The set of points making the medial axis can therefore be thought of as nodes on a graph, hence this representation has also been called a graphical representation by skeletonisation (chapter 6 of [105]).

Medial axis transforms or skeletonisation are generic names used for operations that result in the medial axis being obtained from a binary version of an object or a sampling of points on its surface. Skeletonisation was introduced into computer graphics by Blum [11]. There are a large number of methods for performing medial axes transforms in 3D, some of which address the undesirable properties of this form of representation.

Invariant shape descriptors can be obtained from medial axes in a variety of ways. These commonly assume that the shapes being compared are similar, and share a common form from which the connectivity of the nodes of a generic skeleton (the branch topology) for the particular class of shapes can be defined. The generic skeleton is then fitted to each shape to be analysed, and the fitting process removes orientation and location dependencies. Reparameterisation of the medial axes removes scale differences. The invariant medial axis is then uniformly sampled to obtain

correspondence between the shapes being analysed. Radial distances to the object surface at specific angles at the sampled points are used to give shape descriptors.

This form of shape representation is intuitive and differences in the shape descriptors can be illustrated clearly using the skeletons or can be projected back onto the actual shapes. The major drawback of this method is that it is highly sensitive to noise in the boundary. Additionally, the need to specify the branching topology for particular classes of objects in advance reduces the desirability of this method.

Pizer et. al. [92] and Styner and Gerig [107] describe a method using 3D Voronoi graphs to obtain skeletons and automatically specify the branching topology. In this case skeletons were obtained for all shapes to be analysed, and the deformation of the skeletons was used as the shape descriptor. They term their medial axis description of shape “m-reps”. Golland et. al. [57] used fixed topology skeletons computed from 3D distance transforms of the surfaces. Saha and Chaudhuri [99] obtain skeletons using morphological operations, and classify these to specify the branching topology.

### **2.5.3 A Selection of Methods for 3D Shape Analysis of Neuroanatomical Structures**

Various aspects of shape analysis have been mentioned in the above survey. Here we give examples of application of 3D shape analysis to investigation of shape differences in neuroanatomical structures of the brain. The applications are in four categories, and in no way illustrate the full diversity of applications of shape analysis in computer vision, they are simply to illustrate the application of shape analysis techniques to the brain. A review of other methods applied in computer vision can be found in [77].



### Neuroanatomical Brain Atlases

Neuroanatomical brain atlases are based on voxel-representations and are a non-statistical method. The anatomical structures of interest in the gray-level image of the brain of a subject are manually or automatically labelled to give the neuroanatomical brain atlas e.g. [123]. Given volumetric images of the set of subjects to be analysed, the atlas is then deformed to each using high dimensional transformations. These transformations can have up to  $10^6$  parameters e.g. [120].

Csernansky et. al. [36] analysed a group of 15 matched pairs of schizophrenics and control subjects for shape differences in the hippocampus using MR images. Their neuroanatomical atlas in which the hippocampus was labelled was created using a control subject. Anatomical landmarks were placed in the MR images of the atlas and each subject to be analysed. These were used to obtain a coarse alignment of the atlas to each subject's brain. The coarse alignment was refined using a fluid transformation. A triangulated mesh was described on the surface of the hippocampus of the template, and this was carried along during the deformation. Vector fields were computed from the displacement of the vertices of the triangulated mesh for each subject. A covariance matrix was obtained using the vector fields to which principal component analysis (PCA) was applied. The first 15 principal components were used for shape analysis. The results showed abnormalities occurred in specific subregions of the head and body of the hippocampus between the control and schizophrenic groups. Furthermore, the shape analysis gave greater statistical significance between the groups than volumetric analysis.

### Probabilistic Brain Atlases

Probabilistic brain atlases use voxel-based representations and are a statistical shape method. They use the notion of an average brain of a normal population in the analysis of the brain of a particular subject for shape differences. They are termed

probabilistic because they assign a probability value to each anatomical position in the brain under investigation. The value gives the probability of the anatomical point being abnormal when compared to the location of corresponding points in a database of brains of normal subjects. Probabilistic brain atlases allow quantification of the full range of normal variation of anatomical structures, and help distinguish changes due to disease by providing confidence limits on normal variability.

Guimond et. al. [59] and workers at the Montreal National Institute [48] have generated average brains. However, the Laboratory on Neuro Imaging (LONI) group of the University of California, Los Angeles, is foremost in this field, and Thompson et. al. [119] from that group give a review of methods of constructing average brain atlases.

Thompson and Toga [118] of the LONI group describe a method for the creation of a probabilistic brain atlas. They demonstrate its use on a set of ultra high resolution MR images of 12 Alzheimer's patients and matched controls, and on a set of 6 cryosections of the brains of cadavers of a set of patients with metastatic tumours in the brain and matched controls. To establish orientation, scale and location invariance, all images are transformed to Talairach space using manually defined landmarks within their 3D volumes. As a first step to establishing correspondence, surfaces of several different anatomical structures in all the brains e.g. the cortical sulci, are segmented manually, and overlaid with triangulated meshes.

Each brain from the disease group is then matched to every single brain in the control group in turn to compute probability values. High dimensional warps are used to transform the Talairach space version of the disease brain to be exactly anatomically equivalent to the Talairach version of each of the control brains via a spherical mapping using correspondences obtained from the labelled anatomical structures. 3D probability density functions of the distribution of the mapping of the same anatomical point in each of the control brains is used to obtain a probability value of the location of that point in the diseased brain being abnormal. The probability values

obtained decrease in a monotonic manner with the probability that the location of an anatomic point is abnormal. These values were mapped via a logarithmic lookup table onto a standard colour range, and superimposed within the brain of a reference image for visualisation. For group comparisons, the root mean square of the probability values for the members of each group can be obtained.

The results of [118] were for demonstration of the method only. In the case of a subject with Alzheimer's, there were major differences in the extent of the lateral ventricles and the cortical boundaries. The cerebellum was smaller for the Alzheimer's brain, and there were differences in the cortical sulci. In the case of a subject with metastatic tumours, it was shown that there were significant structural abnormalities in structures in the immediate vicinity of the tumour, but structures in regions further away from the tumour did not exhibit significant structural abnormalities. Furthermore, they showed that when normal controls were analysed, the probability values were not significantly different from those in the database.

### **Voxel-Based Morphometry**

“At its simplest, voxel-based morphometry (VBM) involves a voxel-wise comparison of the local concentration of gray matter between two groups of subjects.” This is a quote from [4] which gives comprehensive details of VBM. The basis of the approach is similar to that of probabilistic brain atlases. However, in the case of VBM, non-rigid registrations are applied to remove global differences in shape, whereas local differences are retained, and these are quantified using the difference of concentration of gray matter (GM). A fundamental assumption is that differences in concentration of GM indicate local structural abnormalities in corresponding areas of brains in the same stereotactic space.

The two high resolution images (or the images from the two groups) being compared are spatially normalised to the same stereotactic space using a template. Ideally, the template should be an average brain based on a large number of images, however, at

the risk of introducing bias into the results, single brains can also be used. Spatial normalisation involves using 12 degree of freedom affine registrations, and linear combinations of smooth spatial basis functions to minimise the sum of squared difference between the image and the template. This is assumed to correct for global shape differences.

The spatially normalised images are segmented into GM, white matter (WM) and “other” tissue classes using cluster analysis. The GM distribution is smoothed with a gaussian kernel whose size depends on the nature of the differences being sought. This means the density of GM in each voxel depends on those of its surroundings, and makes the properties of the voxel values more like those in the statistical assumptions. The statistical analysis of the smoothed data uses a general linear model to identify regions of GM concentration significantly related to the effects being studied. It gives a framework in which different standard parametric tests e.g. t-tests and F-tests can be applied to hypothesis. The significance of any differences can be ascertained using the theory of gaussian random fields.

The validity of VBM has been questioned by Bookstein [17][18] who argued amongst other things that the spatial normalisation involved in VBM may make the statistics about voxelwise comparisons uninformative about group differences. Ashburner and Friston who pioneered VBM agreed that some findings deduced by using the method may pertain to systematic errors [5], but state that Bookstein’s argument about voxelwise tests do not apply to real data. Despite this ongoing debate the method is widely used in the neuroimaging community. The facts that it is automated and relatively straightforward undoubtedly contribute to this. Furthermore, it is more reproducible than manually defined volumetric regions of interests - the current gold standard.

Kubicki et al [73] use VBM in the analysis of brains from patients with first episode schizophrenia, and first episode affective disorder. The schizophrenic and affective disorder groups each comprised of 16 subjects. The control group comprised 18

matched subjects. VBM comparisons of the control and schizophrenic group showed significant differences in the left superior temporal gyrus (STG) of the schizophrenic group. Additional group differences were observed bilaterally in the anterior cingulate gyri and insula, and unilaterally in the parietal lobe. The use of a smaller smoothing kernel showed left sided hippocampal group differences. No significant differences were observed between the controls and the affective disorder group. However, localisation of comparisons to the same regions found defective in the schizophrenic group showed the same pattern of differences in the left STG and the left hippocampus of the affective group as with the schizophrenic group.

Davatzikos et al [37] describe a modification of VBM whereby the spatial normalisation is aided by manually defining a number of sulcal curves within the brain. Using simulated atrophy within the STG and the precentral gyrus, they demonstrate that this method performed substantially better than that of Ashburner et. al [4]. described above. However, this was at the cost of manual intervention.

### **Statistical Shape Models**

Statistical shape models (SSMs) are models that attempt to capture the inherent variation in the shape of a specified class of object based on statistics learnt from a training set (a sample from the population of interest). As the name suggests, they are a statistical method. In addition to shape analysis, they have been used in 2D for object tracking and segmentation. SSMs have an advantage over other methods in that they can allow the comparison of subgroups within a class

3D SSMs have been built using combinations of all the shape representations discussed in section 2.5.2. To construct 3D SSMs, 3D representations of the training set are needed. The members of the training set are normally obtained by segmentation of non-invasively acquired images of the structure of interest. The training set is supposed to be a truly representative sample of the population to be modelled, but obtaining adequate numbers in practice is usually not possible. Three criteria are

usually used in assessing how good a SSM is - Compactness, generalisation ability, and specificity. These are discussed in chapter 4.

*Point distribution models* (PDMs), were introduced by Cootes et al [32]. As the name suggests, a shape is represented by a cloud of [landmark] points. Given a training set, corresponding points on each member of the training set are obtained. The spatial coordinates of the corresponding points for each member of the training set are concatenated into a vector - the “shape vector”. Procrustes analysis or some other form of alignment and normalisation are applied to remove variations due to differences in size and orientation between members of the training set to give an aligned shape vector for each member of the training set. A mean shape and a covariance matrix based on deviations from the mean shape are then obtained from the aligned shape vectors. PDMs are of particular interest in this thesis and will be discussed more fully in chapter 4.

Styner and Gerig [107] present a framework for building SSMs combining surface representations based on the spherical harmonics of [21] and the medial axis of the structures to be modelled. Surface parameterisations are obtained for each member of the training set and these are brought into correspondence by normalising their first degree ellipsoids as in [69]. The continuous parameterisations of the surface are then finely sampled to give point representations of each member of the training set, from which medial axes are obtained by skeletonisation. The medial axes of all members of the training set are used to construct an “m-rep” model. Shape analysis is performed by deforming this to the surface parameterisation of each member of the training set. This method has been used to demonstrate shape differences in the ventricles of monozygotic and dizygotic twins with schizophrenia, and to investigate differences in the hippocampus of schizophrenic and control subjects.

Martin and Pentland [83] describe a finite element method using a physically based deformable template. The intercranial cavities (ICCs) of all members of the training set were segmented along with the structure(s) of interest in the brain. Correspon-

dence was achieved by rigid body registration to align the ICCs of all members of the training set. An average ICC is determined, and a FEM of this is created by assuming the ICC is a linear elastic material. Variations in the shape of the intercranial structures due to differences in head shape are normalised by deforming the FEM to the ICC of each member of the training set and warping the structures of interest accordingly. Modal analysis of the displacement vectors was used to allow shape based discrimination of Alzheimer's disease patients from normals, and schizophrenic patient from normals.

Golland et al [56] use a 3D distance transform as has been described in 2.5.2, and support vector analysis for shape discrimination between diseased and control anatomical structures. They demonstrated that discrimination of the hippocampus-amygdala complex between schizophrenic and control subjects was increased using their method, compared to using a volumetric technique.

#### **2.5.4 Normalisation for Brain/Head Shape**

In the special case of morphometric analysis of neuroanatomical structures in the brain, it is reasonable to assume that the size and shape of the brain is influenced by that of the head as the brain is a jelly like material. The size and shape of the neuro-anatomical structures in the brain are correlated with that of the brain and ultimately the head. It could be argued that in studying two groups, incorporating large enough samples in each could reduce the effects of size and shape between the groups. However, the numbers of subjects used in current trials is usually between ten and fifty per group and cannot be fully representative of variations in head size and shape.

In morphometric studies of this kind, it is important to eliminate as many variables as possible that may influence the size and shape to increase confidence that the residual differences are due to pathological effects of the disease in question. Normalisation of the internal structures of the brain for brain shape has been shown to result in

better shape discrimination. Martin et. al. [83] use a finite element method to accomplish this. They made approximations of the mechanical properties of the brain and accounted for these in their finite element model used to perform shape analysis. They showed that applying their normalisation increased the classification rate of structures into control and disease groups. Gerig et. al. [54] also normalise for brain size. They used the first mode of the spherical harmonic parameterisation of the ventricles to scale all shapes to a unit volume. They reported that normalisation resulted in improved shape differences.

The mechanical properties of the brain are not well known and the scaling factors applied by Gerig et. al. represent a volume parameter and therefore accounts for global size and not shape. We approach normalisation by obtaining an ellipsoidal approximation to the brain and using the three main axes of the ellipsoids to re-scale the ventricles of each brain relative to a chosen template in the three orthogonal directions. This is described in section 3.6.

## 2.6 Summary

This chapter has given details of the background necessary for the work done in this thesis. The principles behind MRI have been explained, as well as relevant properties of the images formed. The preprocessing necessary for our investigations have been described. A review of shape analysis was undertaken and as morphometric analysis is our goal we give a summary of the main points discussed in the section on shape analysis.

Section 2.5 discussed methods of shape analysis, and has given examples of practical application in 3D brain imaging emphasising the desirability of statistical methods. The three statistical methods discussed in section 2.5.3 each have their strengths and weaknesses.



VBM has advantages in that it can be fully automatic, it is relatively straightforward and it is applicable to multiple structures located throughout the brain. However, the validity of the principles behind its use are the subject of ongoing scientific debate. The fact that some structures in the brain e.g. some of the cortical sulci vary greatly or may be completely absent in some brains means the spatial registrations may not be accurate. Additionally, it does not directly measure shape differences of specific neuroanatomical structures.

Probabilistic brain atlases involve a great deal of interaction, and can compare one subject to a database of control subjects as well as compare disease and control groups. They are well suited for qualitative visual analysis of deviations of individual subjects from a group or deviations between two groups. However, the obtaining the high dimensional warps required for registration is computationally expensive. Furthermore, shape differences are expressed in parameters with a high dimensionality. They also suffer from the same problem as VBM in that the high dimensional warps may or may not make real differences.

SSMs also require manual intervention to various degrees. A segmentation step is usually required, and obtaining shape representations and correspondence in 3D may also require manual intervention. SSMs are usually applied to individual brain structures, which although at first sight may appear a disadvantage, allows the detailed characterisation of the particular structure. They allow the comparison of groups and yield compact parameters to which discriminant analysis can be applied.

# Chapter 3

## Data and Image Preprocessing

### 3.1 Overview

In this chapter we describe the MR data used in our work. We also describe the preprocessing steps carried out to obtain the surface representations of the lateral ventricles used in building the shape model. Section 3.2 describes the MR data, and section 3.3 the MR non-uniformity correction applied to this data. Segmentation of the lateral ventricles is the subject of section 3.4. Image combination was used to deal with problems caused by voxel anisotropy and inadequate spatial resolution, this is described in section 3.5. To facilitate better shape discrimination, the surface representations of the ventricles were normalised for differences in brain shape. Obtaining the parameters required for this is the subject of section 3.6.

### 3.2 Description of Data

Volumetric  $T_1$ -weighted and  $T_2$ -weighted MR scans of 39 schizophrenic subjects and 30 age and sex matched controls were used. The age range in the schizophrenic group was 14-48 years, and that for the control group was 13-45 years. For the

control subjects there were 13 females and 17 males, whilst the schizophrenic group was composed of 9 females and 30 males.

For each subject  $T_1$  MR scans had been acquired in the coronal orientation, and  $T_2$  scans in the sagittal, coronal and axial orientations using a 1.5 Tesla magnet. The acquisition in each orientation was independent (i.e. represented an entirely separate scan). The voxel dimensions for scans acquired in the coronal orientation was 0.78mm by 0.78mm in the plane of the slice. For the sagittal and axial orientations the in-plane size of voxels was 0.86mm by 0.86mm.

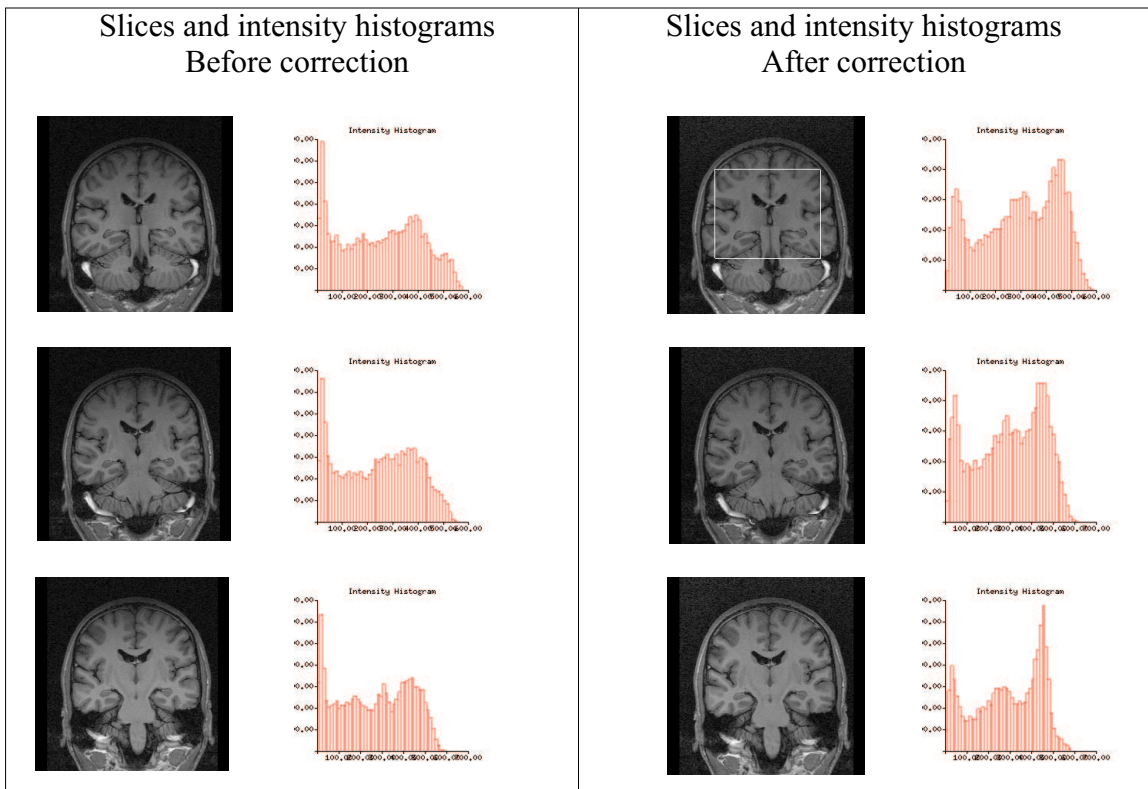
For all orientations there were  $256 \times 256$  voxels in each slice, with each slice having a thickness of 5mm. There was a 1mm intra-slice gap between consecutive slices. The images were originally in DICOM format but were converted to Analyze<sup>TM</sup> 7.5 format for viewing purposes. All subsequent processing took place using images in the Analyze format.

### 3.3 Application of MR Non-Uniformity Correction

All images were corrected for MR inhomogeneity using the method of Vorkuka et. al. [127] outlined in section 2.4.1. This method was used because it is a proven method and had been implemented in the Tina [121] environment by the authors. Figure 3.3 shows the results of non-uniformity correction for three consecutive slices from a  $T_1$ -weighted scan of one of the subjects in the study. The results of intensity histograms of the slices before and after application of non-uniformity correction are also shown next to each slice.

The intensity histograms show a reduction in low intensity values after non-uniformity correction. The profiles after correction show tri-modal peaks in all slices, which were not observed in the histograms before correction. The peaks correspond to CSF,

GM and WM intensity distributions respectively. The results show that although non-uniformity correction results in no significant visual improvement, the tri-modal nature of the histograms indicates an improved gray level distribution into tissue classes.



**Figure 3.1:** .

Qualitative results of intensity non-uniformity correction. The histograms of intensity values show distinct tri-modal peaks attributable to CSF, GM, and WM, after application of non-uniformity correction. Voxel intensities of the same region in all slices (see image on the right in row 1) were used in obtaining the histograms

### 3.4 Segmentation of the Lateral Ventricles

Our approach to segmentation was to use a 3D edge detector [85] to give edges which were manually linked to form closed contours in each slice. In determining our

approach to segmentation the following were considered:

1. **The Quality of Results.** As the purpose of segmentation was shape analysis of the resulting ventricular surface it was paramount that the ventricles were segmented as accurately as possible. The quality of segmentations was judged by two neuro-radiologists.
2. **Efficiency.** 69 Images had to be segmented, and the lateral ventricles spanned approximately 90 slices in each image. Therefore, a considerable amount of data had to be segmented. The method of choice had to make as much use as possible of the characteristics peculiar to the segmentation of the ventricles. The most prominent of these is the fact that the ventricles contain CSF, whilst the surrounding tissues contain GM and WM. Therefore in most slices there is usually good contrast between the ventricles and the surrounding tissues.
3. **Intuitive Interaction.** Even though there is good contrast between the ventricles and surrounding tissues in most slices, this is not the case in the temporal horns where the ventricles are thin. The partial volume effect is a problem at the temporal horns, and automated segmentations would require manual editing in this region. Additionally, the left and right ventricles may come into close contact in some slices. There is therefore the need for a significant amount of manual interaction. Intuitive methods of carrying this out help to reduce frustration and increase the accuracy of segmentation.
4. **Quick Implementation.** Developing methods of segmentation is an active research topic. However, it is not the subject of this thesis. Therefore we sought a method that met the criteria listed above, and was either readily available or could be implemented in a straightforward manner.

Before selecting the 3D edge detector approach a number of other methods were considered. Two readily available methods, the region-based method in Analyze and a statistical method available in the Tina environment were evaluated. Segmentation

using Analyze required substantial user interaction, and the results obtained still had to be modified manually in most cases.

The statistical method of Tina involved fitting Gaussian kernels to the distribution of gray levels for GM, WM and CSF (see figure 3.3). The partial volume effect was accounted for by using a linear approximation to the tails of the Gaussian distributions where two of the kernels overlapped (it is assumed that only two tissue types overlap in any one voxel). The results of the fit are used to obtain probabilities that a voxel is drawn from a population of CSF, GM, or WM - resulting in a probability image that can be thresholded. Segmentation using Tina gave reasonable results in most slices. However, it also required user interaction defining regions on which to base its probability estimations. It could not distinguish between left and right ventricles when they were close. Furthermore, it gave poor results for the temporal horns and at the extremities of the ventricles. The statistical method of Tina was not used because of these unfavourable features.

A watershed method [98] was implemented. This required the user to place seed points within the ventricles, so that their catchment basins could be identified. Again this method gave good results for most of the ventricle, but could not distinguish between right and left ventricles when they were in close proximity and did not work well for the temporal horns. Furthermore, in some instances the catchment basins overflowed. Although several modifications were made in attempts to solve this they did not completely eradicate the problem.

The edge-based method was selected because the edges are well defined for most slices of the ventricle. The problem of linking edgels was tackled manually as we did not want to spend time developing heuristics that needed further modification. Figure 3.2(a) shows the edgels superimposed on a slice of one of the subjects of the study and the resulting segmented contours. We were also able to present the interaction in an intuitive way using a user interface generated with Matlab (see Figure 3.2(b)). The 3D detector was used rather than a 2D one because the neuro-radiologists felt it


gave more accurate results.

The steps involved in segmentation were:

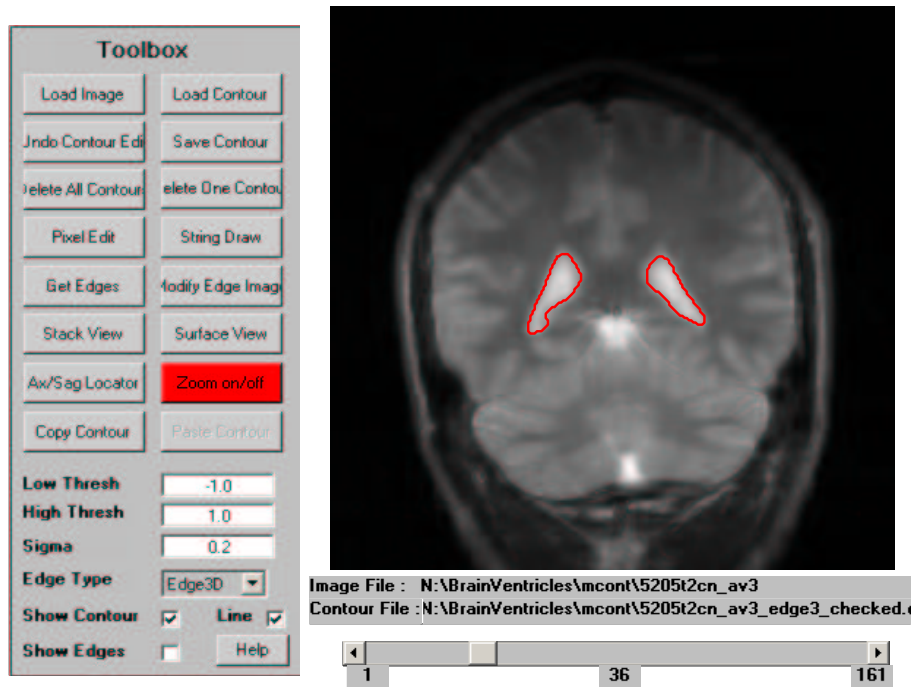
1. Load image into Matlab segmentation tool (see figure 3.2)
2. Obtain edges using 3D edge detector by clicking the “Get Edges” pushbutton in the toolbox on the left of figure 3.2.
3. The binary edge image is superimposed on the gray level image as shown in figure 3.2(a).
4. The “Modify Edge Im” pushbutton in the toolbox on the left of figure 3.2 allows the pixels of the binary edge image to be turned *on* or *off*. This is used to isolate the edgels around the ventricles to form closed contours. Figure 3.2(b) shows the contours for the slice displayed in figure 3.2(a).
5. The contours around the ventricles of each slice in the image are obtained in this manner.
6. The contours are then saved as list of x,y,z coordinates in a text file.

### 3.5 Image Combination to Improve Spatial Resolution

The voxels of the acquired images are anisotropic (see section 3.2), with the effective slice thickness of 6mm, leading to “steps” on the ventricular surface (see Figure 3.7). The reasons for acquiring images with anisotropic voxels were explained in section 2.3.1. However, from a shape analysis point of view we would like to represent the ventricular surface with as fine a resolution as possible (preferably 1mm or less). We would also like isotropic voxels to simplify the construction of the shape model by having the same distance units in the three orthogonal directions.



(a) 3D edges overlaid on a coronal slice



(b) Contours of segmented ventricle

**Figure 3.2:** Images of a coronal slice showing the 3D edges and the resulting segmented contours of the ventricle. The images are displayed within the GUI used for segmentation



A straightforward solution to both voxel anisotropy and inadequate out-of-plane resolution is re-slicing the images to give isotropic voxels of an adequate resolution using one of the interpolation methods discussed in section 2.3.3. Contour based interpolation [124] is an alternative method. This involves segmentation of the desired structure of interest from the low out-of-plane resolution image. Interpolation is performed on the segmented contours to approximate contours that would have been obtained had the out-of-plane resolution been higher. Contour interpolation methods have the advantage that they reduce the amount of segmentation required, as the low out-of-plane resolution image has thicker and hence fewer slices than one that would have resulted from conventional image interpolation.

However, inherent in the independently acquired  $T_2$  coronal, axial and sagittal views of a subject is sub-millimetre resolution in all three orthogonal directions. Our aim in this section was to produce high resolution isotropic voxel images of the brain of each subject using information from *all* three views. Our motivation being that conventional application of re-slicing and interpolation or contour based interpolation would only use information from one of the three views.

In the following subsections we describe how we arrived at our chosen approach to *image combination*. The layout is as follows. Image combination requires the accurate registration of the three views. **Registration is required because the views were acquired independently and there would have been a degree of movement in the patient between the acquisitions.** Section 3.5.1 describes the evaluation of three freely available registration packages to select an appropriate one for performing our registrations. In addition to straightforward interpolation, two approaches to image combination were investigated. The first (combination method 1), that which was selected as the method of choice, is described in section 3.5.2. The second approach (combination method 2) involved forming sets of linear equations from the images and solving for voxel values. This is described in section 3.5.3. Section 3.5.4 presents the results of application of both image combination methods, as well as conventional re-slicing and interpolation to synthetic test images and to the actual images used

in this study. The method of choice for image combination to apply in this study is discussed in section 3.5.5.

### 3.5.1 Evaluation of Three Registration Packages

#### The Registration Software

The three packages assessed were, version 3.08 of the Automated Image Registration (AIR) package of Woods et. al. [134]. The “MPR” registration package from the Computational Imaging Sciences Group at King’s College London (now called the VTK CISG registration toolkit) [106], and the FMRIB Linear Image Registration Tool (FLIRT) from the University of Oxford [66]. AIR uses the Woods function as its similarity metric, whilst MPR uses normalised mutual information, and FLIRT uses the correlation ratio.

We use the AIR convention when referring to the registration of a set of images. The image to which the others are registered is called the *standard* image, and the image being registered is the *reslice* image. For the registration of each reslice image to the standard image we obtain six parameters  $T_x$ ,  $T_y$  and  $T_z$  for translations and  $R_x$ ,  $R_y$  and  $R_z$  for rotations.

The aim of this evaluation was to investigate the robustness, accuracy and consistency of the registration software in order to select the best one for our purposes. It should be noted that the evaluation carried out here is limited to rigid-body registration of images of the same subject obtained by the same imaging modality (mono-modal intra-subject rigid-body registration). Other more comprehensive evaluations of multi-modal inter-subject registration have been carried out e.g. [51]

#### Gold Standards and Measure of Registration Accuracy

An MR simulator [30] was used to generate five T2 images of a phantom model at a resolution of  $1\text{mm} \times 1\text{mm} \times 1\text{mm}$ . These served as gold standards for the assessment of the accuracy of registration and were also used in quantitative evaluation of

image combination. The number of elements in the  $x$ ,  $y$  and  $z$  directions were  $181 \times 181 \times 216$ . The images were generated with 3% noise and 20% MR inhomogeneity. They were converted into “thickened” anisotropic voxel images in the coronal, axial and sagittal orientations, each with  $1\text{mm} \times 1\text{mm}$  in-plane resolution and 5mm slice thickness and 1mm slice gap, by averaging adjacent voxels in the appropriate direction.

A triangulated surface of the lateral ventricles of the phantom model was generated. The coordinates of the vertices,  $V$ , of this were stored for use in the computation of a similarity metric to assess the various registration packages.

An Euclidean metric  $d_m$  (Equation 3.1), was calculated using transformation parameters for a particular registration and the pre-stored vertices of the ventricular surface of the phantom. Artificial mis-registrations were introduced into the test images by rotating them about the coordinate axes. The same mis-registration was applied to  $V$  to give a new set of coordinate points  $V'$ . The transformation parameters returned by a registration package for the registration of the mis-registered and original images were applied to  $V'$  to give  $V''$ .  $d_m$  is the mean Euclidean distance between corresponding points in  $V$  and  $V''$ .

$$d_m = \frac{1}{N} \sum_{i=1}^N [(x_i - x''_i)^2 + (y_i - y''_i)^2 + (z_i - z''_i)^2]^{\frac{1}{2}} \quad (3.1)$$

where

$N$  = number of vertices on the ventricular surface

$x$ ,  $y$  and  $z$  are coordinates of the vertices on  $V$

$x''$ ,  $y''$  and  $z''$  are the coordinates of corresponding vertices of  $V''$

### Experiments Performed

Experiments were performed as follows. The axial version of each of the five thickened test images of resolution  $1\text{mm} \times 1\text{mm} \times 6\text{mm}$ , was used to generate both a standard image and a reslice image. The standard image was formed by re-slicing the thickened

image into  $1\text{mm} \times 1\text{mm} \times 1\text{mm}$  voxels using spline interpolation. The reslice image was formed by rotation of the thickened image by  $0^\circ$ ,  $5^\circ$  or  $10^\circ$  to give an image with resolution of  $1\text{mm} \times 1\text{mm} \times 6\text{mm}$ . The rotations were done about the  $x, y$  and  $z$  axes independently. Identical rotations were applied to the vertices of the ventricular surface,  $V$ , to give a new set of vertices  $V'$ . The reslice image was registered to the standard image by MPR, AIR and FLIRT, and the resulting transformation parameters tabulated. These were used to form transformation matrices applied to the respective  $V'$  to give  $V''$ . A value of  $d_m$  was then calculated using equation 3.1.

For MPR and AIR the above was repeated using images in the coronal orientation (FLIRT only works on images in the axial orientation).

### Results of Evaluation Experiments

Tables 3.1, 3.2 and 3.3 give average values of the translation parameters for each registration package. The values in each column represent the mean over the five test images. The first column states the initial rotation applied to give an artificial mis-registration. The values along each row give the rotation and translation parameters applied by the registration package to correct for the mis-registration. Table 3.4 summarises the values of  $d_m$  for each set of experiments.

In general the rotations applied by all the packages were correct to two decimal places. However, AIR mis-registered in three of five cases (highlighted in Table 3.1) when rotated by  $10^\circ$  about the  $x$ -axes in the axial orientation, and in all five cases when rotated by  $10^\circ$  in the coronal orientation.

The translation parameters returned differed for each method. This was the main contributor to differences in  $d_m$  between the methods. For AIR and MPR the maximum translations were 0.52mm in  $x$ , and 0.42mm in  $y$ . In the  $z$  direction the maximum translation was 2.32mm for AIR and 1.24mm for MPR. For FLIRT the maximum translations were 3.31mm in  $x$ , 2.95mm in  $y$  and 2.00mm in  $z$ . The greater translation in the  $z$  direction may be expected because of the lower resolution of the reslice images in this direction. This also explains the fact that even when no rotations are

Initial	Tx (mm)		Ty (mm)		Tz (mm)		Rx (degrees)		Ry (degrees)		Rz (degrees)		dm (mm)	
Rotation	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
Axial Orientation														
0	0.058	0.001	0.049	0.072	-2.324	0.005	0.020	0.027	0.078	0.002	0.019	0.002	2.327	0.006
Rx 5	-0.111	0.000	-0.290	0.001	-1.977	0.002	-5.007	0.002	-0.195	0.001	-0.015	0.000	2.000	0.002
Rx 10	-0.266	0.146	0.186	0.615	-2.228	0.300	-4.065†	5.396	-0.046	0.138	-0.012	0.000	5.825	3.528
Ry 5	-0.280	0.002	-0.172	0.002	-2.004	0.010	-0.072	0.003	-5.034	0.003	-0.009	0.001	2.034	0.009
Ry 10	-0.520	0.001	-0.176	0.001	-1.950	0.005	-0.074	0.001	-10.010	0.001	-0.006	0.000	2.030	0.004
Rz 5	0.001	0.000	-0.001	0.001	-2.320	0.001	0.003	0.001	0.001	0.000	-5.003	0.000	2.320	0.001
Rz 10	0.001	0.000	0.000	0.001	-2.319	0.001	0.003	0.001	0.001	0.000	-10.003	0.000	2.319	0.001
Coronal														
0	0.002	0.005	0.022	0.003	0.660	0.002	0.080	0.001	-0.068	0.097	-0.015	0.022	0.670	0.009
Rx 5	-0.167	0.001	0.007	0.002	0.924	0.003	-4.974	0.004	0.008	0.002	-0.016	0.001	0.937	0.003
Rx 10	-0.375	0.004	-0.420	0.002	0.355	0.002	-0.173†	0.001	-0.006	0.002	0.024	0.001	5.232	0.001
Ry 5	-0.004	0.001	-0.180	0.001	1.055	0.003	0.000	0.004	-5.003	0.002	-0.004	0.001	1.071	0.003
Ry 10	-0.090	0.179	-0.228	0.109	0.892	0.188	0.000	0.008	-7.999†	4.389	0.009	0.004	1.663	1.509
Rz 5	-0.007	0.001	0.017	0.001	0.689	0.000	0.002	0.000	-0.004	0.001	-5.020	0.000	0.690	0.000
Rz 10	-0.007	0.001	0.016	0.001	0.687	0.000	0.003	0.000	-0.002	0.001	-10.020	0.000	0.688	0.000

**Table 3.1:** Mean Translation and Rotation Parameters for AIR. The values along each row are the parameters applied to correct for the initial rotations. Rx5 indicates an initial 5° rotation about the  $x$ -axis and Rx10 an initial 10° rotation etc. The column headings show the parameters required to register the image with its unrotated version. †Highlighted values show where gross mis-registrations occurred

present, there is still significant translation in the  $z$  direction.

Table 3.4 shows MPR gave the lowest value of  $d_m$  for registrations in both the coronal and the axial orientations. The value of  $d_m$  for AIR was affected by the gross mis-registrations at 10° . The fact that gross mis-registrations occurred shows AIR is less robust than MPR and FLIRT. These mis-registrations were probably due to the optimisation method finding a local minimum of the voxel similarity metric cost function as opposed to the global minimum. The value of  $d_m$  for FLIRT was high because of the error resulting from the registration. Additionally the use of FLIRT would mean restriction to the axial orientation. On the basis of these results MPR was selected as the registration tool for the subsequent experiments.

Initial	Tx (mm)		Ty (mm)		Tz (mm)		Rx (degrees)		Ry (degrees)		Rz (degrees)		dm (mm)	
Rotation	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
Axial Orientation														
0	0.000	0.000	0.000	0.000	-1.000	0.000	0.000	0.000	0.000	0.000	-0.120	0.000	1.012	0.000
Rx 5	0.120	0.000	0.000	0.000	-1.210	0.000	-4.970	0.000	-0.120	0.000	0.000	0.000	1.215	0.000
Rx 10	0.120	0.000	0.000	0.000	-1.236	0.058	-9.960	0.000	-0.120	0.000	0.000	0.000	1.240	0.058
Ry 5	0.000	0.000	0.120	0.000	-1.236	0.058	0.048	0.066	-5.064	0.058	0.000	0.000	1.240	0.057
Ry 10	0.000	0.000	0.120	0.000	-1.210	0.000	0.000	0.000	-10.090	0.000	0.000	0.000	1.217	0.000
Rz 5	0.000	0.000	0.000	0.000	-0.880	0.000	0.000	0.000	0.000	0.000	-5.030	0.000	0.881	0.000
Rz 10	0.000	0.000	0.000	0.000	-0.880	0.000	0.000	0.000	0.000	0.000	-10.040	0.000	0.882	0.000
Coronal														
0	0.000	0.000	0.000	0.000	-1.000	0.000	0.120	0.000	-0.072	0.107	0.000	0.000	0.996	0.011
Rx 5	0.000	0.000	0.120	0.000	-0.960	0.000	-5.120	0.000	-0.120	0.000	0.000	0.000	0.959	0.000
Rx 10	0.000	0.000	0.120	0.000	-0.960	0.000	-10.160	0.000	-0.120	0.000	0.000	0.000	0.959	0.000
Ry 5	0.000	0.000	0.000	0.000	-0.910	0.000	0.000	0.000	-5.030	0.000	0.000	0.000	0.907	0.000
Ry 10	0.000	0.000	0.000	0.000	-0.988	0.071	0.000	0.000	-10.040	0.000	0.000	0.000	0.983	0.071
Rz 5	0.000	0.000	0.120	0.000	-1.000	0.000	0.120	0.000	0.000	0.000	-5.090	0.000	1.006	0.000
Rz 10	0.000	0.000	0.120	0.000	-1.000	0.000	0.120	0.000	0.000	0.000	-10.090	0.000	1.005	0.000

**Table 3.2:** Mean Translation and Rotation Parameters for MPR. See caption of 3.1 for description of row and column headings

Initial	Tx(mm)		Ty(mm)		Tz(mm)		Rx (degrees)		Ry (degrees)		Rz (degrees)		dm (mm)	
Rotation	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
Axial Orientation														
0	-0.036	0.004	-0.018	0.013	1.999	0.009	-0.032	0.013	-0.004	0.013	-0.044	0.003	1.999	0.008
Rx 5	-0.026	0.015	0.380	0.001	3.333	0.006	-5.005	0.006	-0.016	0.018	0.000	0.000	3.355	0.006
Rx 10	-0.026	0.028	0.895	0.003	4.722	0.013	-10.037	0.017	0.001	0.063	0.017	0.019	4.804	0.013
Ry 5	0.266	0.006	-0.031	0.004	0.266	0.002	0.000	0.001	-5.057	0.004	0.000	0.000	0.382	0.003
Ry 10	0.353	0.014	-0.033	0.017	-1.481	0.011	0.024	0.025	-10.074	0.005	-0.012	0.004	1.524	0.011
Rz 5	-1.579	0.008	-1.532	0.004	1.990	0.002	-0.021	0.003	-0.010	0.004	-4.971	0.014	2.968	0.006
Rz 10	-3.309	0.019	-2.948	0.017	1.987	0.004	-0.006	0.008	-0.004	0.012	-9.972	0.005	4.858	0.018

**Table 3.3:** Mean Translation and Rotation Parameters for FLIRT. See caption of 3.1 for description of row and column headings

Package	Orientation	No Rotation	Rx 5	Rx 10	Ry 5	Ry 10	Rz 5	Rz 10	Overall Mean	S.D.
AIR	Axial	2.327	2.000	5.825	2.034	2.030	2.320	2.319	2.693	1.389
AIR	Coronal	0.670	0.937	5.232	1.071	1.663	0.690	0.688	1.564	1.655
MPR	Axial	1.012	1.215	1.240	1.240	1.217	0.881	0.882	1.098	0.168
MPR	Coronal	0.996	0.959	0.959	0.907	0.983	1.006	1.005	0.974	0.035
FLIRT	Axial	1.999	3.355	4.804	0.382	1.524	2.968	4.858	2.841	1.668

**Table 3.4:** Summary of values of  $d_m$  (mm) for AIR, MPR and FLIRT. See caption of 3.1 for description of row and column headings

### 3.5.2 Combination Method 1 - Image Combination by Registration and Interpolation

In performing image combination by registration and interpolation, the coronal, axial, and sagittal images of a subject are registered and re-sliced to isotropic voxels of the same size. Overlapping regions in all three views are averaged to give the combined high resolution output with isotropic voxels. An outline of the experimental procedure for the ventricle study images is as follows.

For each subject:

1. The coronal image is selected as the standard image. Its original resolution is  $0.78\text{mm} \times 0.78\text{mm} \times 6.0\text{mm}$  in coronal space. It is resliced into  $0.78\text{mm} \times 0.78\text{mm} \times 0.78\text{mm}$  isotropic voxels using spline interpolation to give a high resolution standard image.
2. The axial image is taken as the current reslice image. Its original resolution is  $0.86\text{mm} \times 0.86\text{mm} \times 6.0\text{mm}$ . This is converted to the coronal orientation by a 90 anticlockwise rotation about the x - axis in axial space (see Figure 1). The reslice image is therefore the axial image in coronal orientation with a resolution of  $0.86\text{mm} \times 6.0\text{mm} \times 0.86\text{mm}$ .
3. The re-slice image is registered with the standard image of step 1. This registered image is resliced into  $0.78 \times 0.78 \times 0.78\text{mm}$  voxels using trilinear interpolation.
4. The sagittal image is taken as the current reslice image. Its original resolution is  $0.86 \times 0.86 \times 6.0\text{mm}$  in sagittal space. This is converted to the coronal orientation by a 90 rotation about the y axis in sagittal space. The current reslice image is therefore the sagittal image in the coronal orientation with resolution  $6.0 \times 0.86 \times 0.86\text{mm}$ .
5. Step 3 is repeated using the current reslice image.

We therefore have the axial and sagittal images in coronal space registered with the coronal image and having the same voxel dimensions as the coronal image. A combined image is formed from these three images by averaging overlapping regions.

Specifically:

- Overlapping regions in all three images are averaged.
- Overlapping regions in axial and coronal but not sagittal are averaged.
- Overlapping regions in sagittal and coronal but not axial are averaged.
- The coronal values are used to fill any voxels not already assigned a value.

The combined image therefore has isotropic voxels and occupies the same spatial extent as the coronal image but contains information from the coronal, axial, and sagittal views. This procedure was applied to the thickened versions of test images (resolutions in this case are  $1 \times 1 \times 6\text{mm}$ ). The reconstructed images were compared with the gold standards and the results given in section 3.5.4. The results of the application to the images of the ventricle study are also given in this section.

### 3.5.3 Combination Method 2 - Image Combination by Solution of Linear Equations

The coronal, axial and sagittal views are different representations of the same original image. A system of linear equations can be set up using the known values from these views to solve for a combined image of a given resolution. In an attempt to present the method clearly, a simple 2D case is used to illustrate the basic principle.

The following are assumed.

1. All three views are in spatial correspondence and have a common origin at the




top left. This was achieved by registration of the axial and sagittal views to the coronal view using MPR.

2. The value of each voxel/pixel is uniformly distributed within it
3. The coronal, axial, and sagittal image voxel values represent averages of the gray levels over the spatial extent they occupy

### The Simple 2D Case

In this case the combined image we want to form,  $\mathbf{W}$ , is a  $m \times n$  matrix with pixels of unit resolution. We have a matrix  $\mathbf{R}$ , which represents the sum (actually average) over  $d_r$  rows in our original  $m \times n$  image and a matrix  $\mathbf{C}$ , which is the sum over  $d_c$  columns in our original image - see Figure 3.3.



**Figure 3.3:** A schematic illustration of the 2D case for image combination by linear equations

In the case where  $m = n$ , and  $d_r = d_c = d$  (see Figure 3.3), we can form the following sums using the three assumptions made in the beginning of this subsection.

$$\begin{aligned} R_{ij} &= \sum_{k=d \times (i-1)+1}^{d \times i} W_{kj} \\ C_{ij} &= \sum_{k=d \times (j-1)+1}^{d \times j} W_{ik} \end{aligned} \quad (3.2)$$

Where  $R$ ,  $C$ , and  $W$  are elements in  $\mathbf{R}$ ,  $\mathbf{C}$ , and  $\mathbf{W}$  respectively. The first subscript of each element is the row index, and the second the column index.

The equations in 3.2 can be written in matrix form as

$$\mathbf{A}\mathbf{w} = \mathbf{b} \quad (3.3)$$

Where


$\mathbf{A}$  is a matrix whose elements refer to the indices of the elements in  $\mathbf{W}$

$\mathbf{w}$  is a vector representing the pixel values in the combined image (the values we seek to find)

$\mathbf{b}$  is a vector formed by the concatenation of the known pixel values from  $\mathbf{R}$  and  $\mathbf{C}$ .

Equation 3.3 is underdetermined in that the number of unknowns is greater than the number of equations (in the case when  $d = 2$  the number of unknowns and equations are the same, however, some of the equations will be linearly dependent). The solution given by solving Equation 3.3 for  $\mathbf{w}$  are constrained to sum to the respective values of  $\mathbf{R}$  and  $\mathbf{C}$ . A further constraint is introduced to impose a degree of smoothness on the image by minimising differences between the values of neighbouring elements. This reduces discontinuities in the resulting image. Figure 3.4 in conjunction with Equation 3.4 illustrates our reasoning.

To impose smoothness we want to minimise the differences between neighbours.



**Figure 3.4:** The simple 2D case to illustrate the smoothness term

Hence for the  $x$  (column) direction the term to be minimised is

$$\sum_{i=1}^m \sum_{j=1}^{n-1} (W_{i,j+1} - W_{i,j}) \quad (3.4)$$

A similar equation holds for the  $y$  direction.

With the elements of  $\mathbf{W}$  put in the form of a column vector  $\mathbf{w}$ , the smoothing terms can be represented in matrix form as  $|\mathbf{A}_x \mathbf{w}|^2$ , and  $|\mathbf{A}_y \mathbf{w}|^2$ , where  $\mathbf{A}_x$  and  $\mathbf{A}_y$  are matrices selecting the coefficients of  $\mathbf{w}_{ij}$  in Equation 3.4. Recalling that Equation 3.3 is underdetermined, its least squares solution is found by minimising  $|\mathbf{A}\mathbf{w} - \mathbf{b}|^2$ . Hence to include the smoothness constraint we have to minimise

$$E = |\mathbf{A}\mathbf{w} - \mathbf{b}|^2 + \alpha |\mathbf{A}_x \mathbf{w}|^2 + \beta |\mathbf{A}_y \mathbf{w}|^2 \quad (3.5)$$

Where  $|\mathbf{A}_x \mathbf{w}|^2$  represents the sum of absolute differences in neighbouring pixels on a row, and  $|\mathbf{A}_y \mathbf{w}|^2$  gives the equivalent for a column.  $\alpha$  is a factor that determines the smoothness in the  $x$  direction  $\beta$  is a factor that determines the smoothness in the  $y$  direction.

Differentiating 3.4 and setting the result to 0 (at a minimum) gives (see chapter 9

[80] for details)

$$(\mathbf{A}^T \mathbf{A} + \alpha \mathbf{A}_x^T \mathbf{A}_x + \beta \mathbf{A}_y^T \mathbf{A}_y) \mathbf{w} = \mathbf{A}^T \mathbf{b} \quad (3.6)$$

Generally we require  $\alpha = \beta$ . If  $\alpha = \beta = 0$  then Equation 3.4 reduces to Equation 3.3. The 3D case simply requires a term in the  $z$  direction to be included in Equation 3.4 and the addition of the known values in this direction (in a matrix we refer to as  $\mathbf{P}$ ) to  $\mathbf{b}$ . The 3D case is given by :

$$(\mathbf{A}^T \mathbf{A} + \alpha \mathbf{A}_x^T \mathbf{A}_x + \beta \mathbf{A}_y^T \mathbf{A}_y + \gamma \mathbf{A}_z^T \mathbf{A}_z) \mathbf{w} = \mathbf{A}^T \mathbf{b} \quad (3.7)$$

### Implementation Considerations

Obtaining a combined image by solving the linear equations requires a considerable amount of memory for storing the matrix elements, not to mention that required for the necessary computations. For the case of the test images, we want a combined image with  $180 \times 180 \times 216$  elements. This is generated from coronal axial and sagittal images each with  $(180 \times 180 \times 216)/6$  elements. The matrix  $\mathbf{A}$  for the test images has  $(6,998,400 \times 3,499,200)$  elements requiring 178 terabytes of storage. However, it only contains 20,995,200 non-zero elements and is therefore created as a sparse matrix requiring 160MB of storage.  $\mathbf{A}_x$ ,  $\mathbf{A}_y$  and  $\mathbf{A}_z$  are also generated as sparse matrices.

Applying Equation 3.4 to solve for the  $(180 \times 180 \times 216)$  test images did not finish after 1 week using 500MHz Pentium class processor running Matlab 5.3.1. An implementation whereby the matrices  $\mathbf{A}$ ,  $\mathbf{A}_x$ ,  $\mathbf{A}_y$  and  $\mathbf{A}_z$  were created for a  $6 \times 6 \times 6$  unit was therefore used. The inverse of  $\mathbf{A}$  and the transposes of  $\mathbf{A}_x$ ,  $\mathbf{A}_y$  and  $\mathbf{A}_z$  had to be calculated only once for the  $6 \times 6 \times 6$  unit. The whole image solution was obtained by consecutively applying these to successive values of  $\mathbf{R}$ ,  $\mathbf{C}$  and  $\mathbf{P}$ . The solution in this case took approximately 1 minute for each test image using equation 3.3 to give a minimum norm solution, and about 21 minutes for each test image using equation 3.4 to obtain a least squares solution.

### 3.5.4 Experiments on Image Combination

The two image combination methods and spline interpolation were applied to the thickened images of the phantom described in section 3.5.1 to reconstruct images of the same resolution as the original. This enabled a quantitative measure of the fidelity of the reconstructed images to the original versions by each method. We used the NCC measure (see Equation 2.2) for this. Qualitative comparisons of the reconstructed and original test images are also presented.

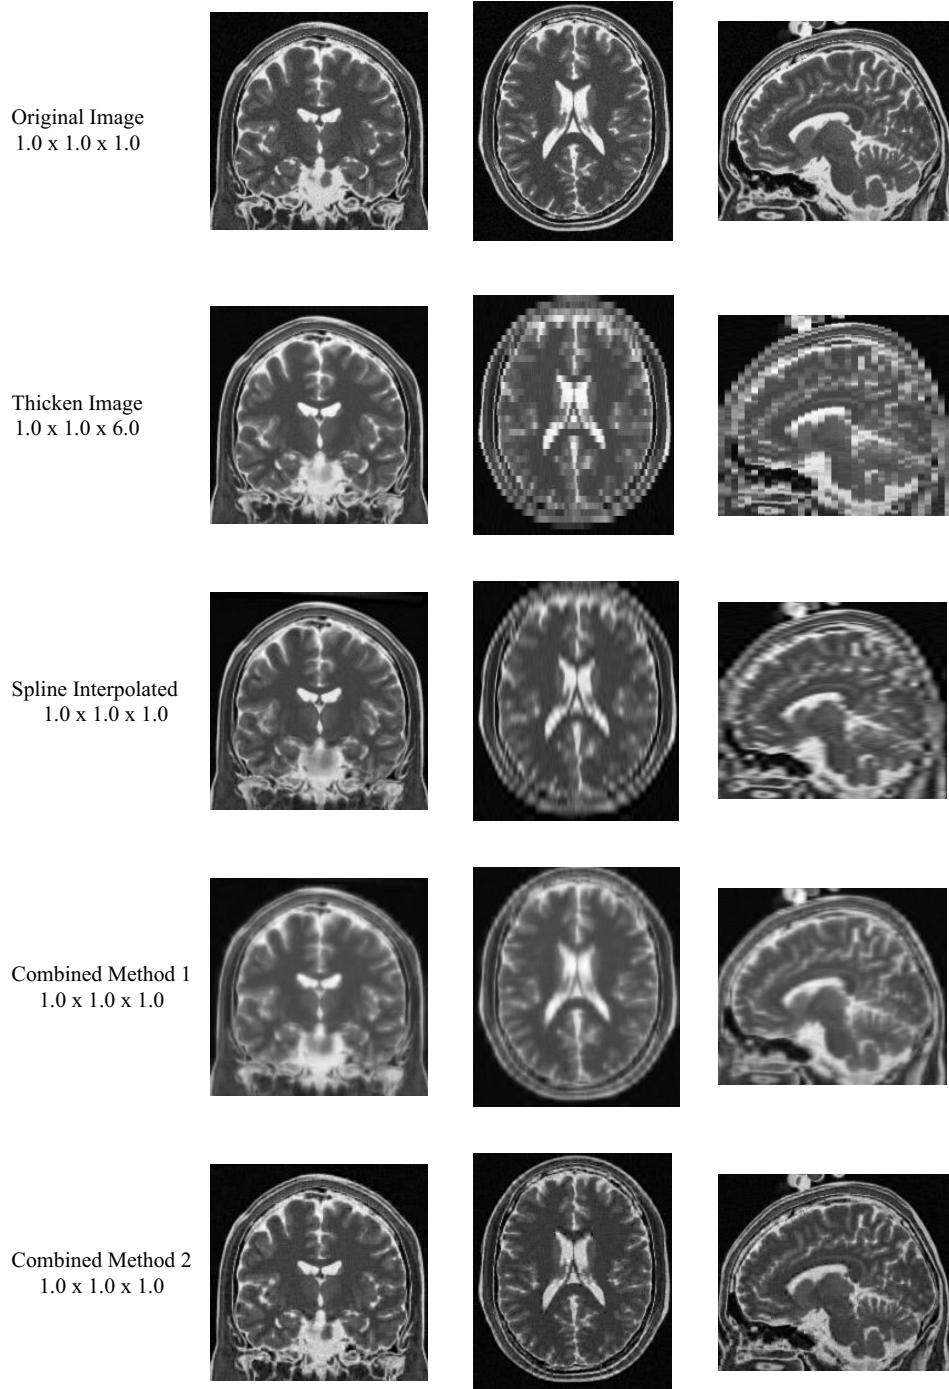
The combination methods and spline interpolation were also applied to images of this study to produce images with 0.78mm isotropic voxels. This resolution was chosen because it represented the lowest in-plane resolution in the acquired images. However, as there was no ground truth, only qualitative results are presented.

#### Results for Test Images

Table 3.5 gives the NCC values obtained for comparison of results of the two image combination methods and the spline interpolated image with the original image. Figure 3.5 shows a slice from the original image in coronal, axial, and sagittal orientations, and the corresponding slice from each of the methods. Both the qualitative and quantitative results show that combination method 2 using the linear equation approach gives consistently better results for all orientations. Combination method 1 using registration and interpolation gives more consistent results for all orientations than simply interpolating. This is shown in Figure 3.5. The coronal reconstruction using spline interpolation is sharper than that of the corresponding reconstruction by combination method 1. However, the sagittal and axial reconstructions of combination method 1 look better than those resulting from spline interpolation.

#### Results for Images for this Study

Figure 3.6 shows results of coronal, sagittal and axial views for the images acquired for this study. As in the results for the test images, combination method 2 gives more consistently sharper images. However, on close inspection of the results of combina-



**Figure 3.5:** Qualitative results of the application of reconstruction methods to test images

Reconstruction Method	NCC	Standard Deviation
Spline Interpolation	0.8839	0.0001
Combination Method 1	0.8856	0.0011
Combination Method 2	0.9773	0.0000

**Table 3.5:** Average values of NCC (n=5) for comparison of reconstructed image with original

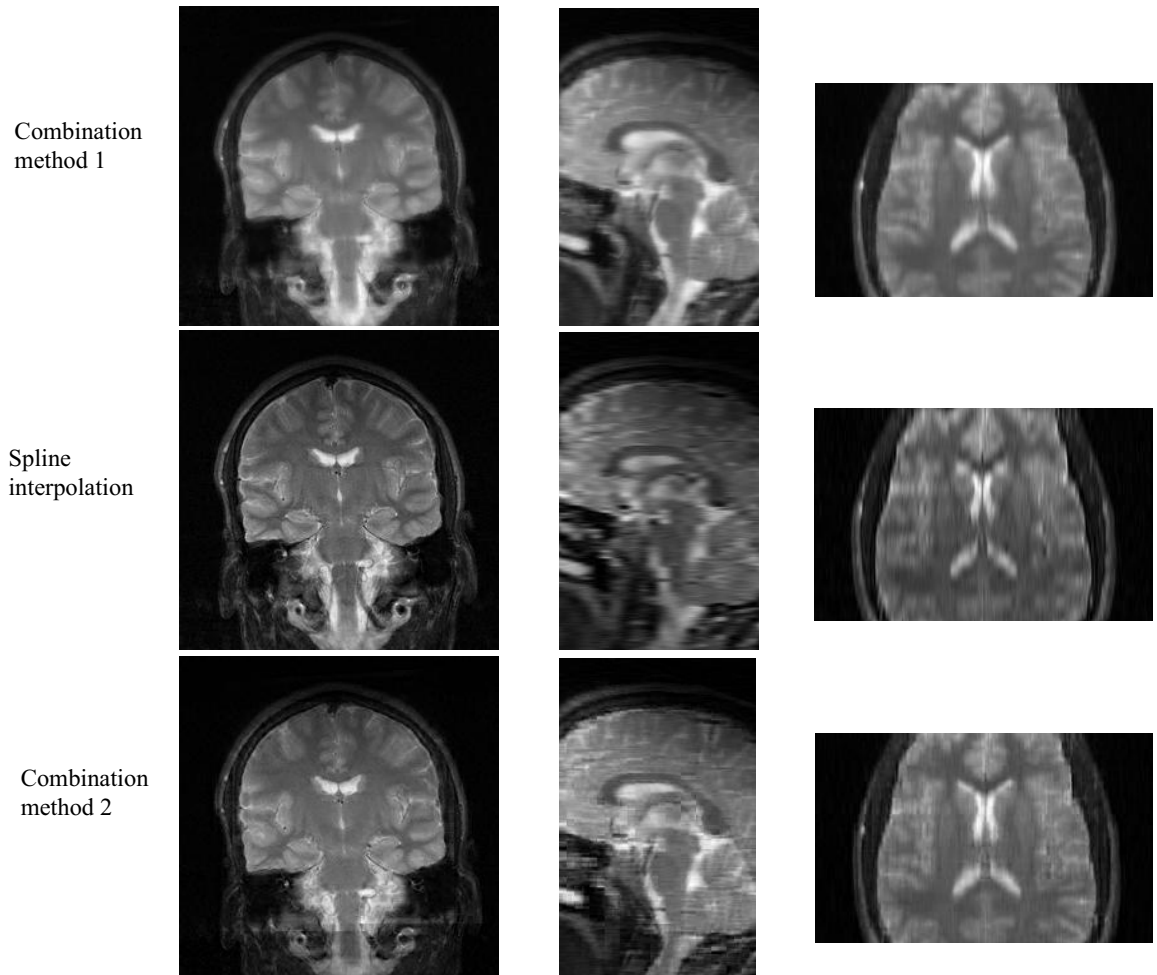
tion method 2 for the actual images of this study there were “blocky discontinuities”. The spline interpolation gives a sharp image for the coronal orientation but fuzzy images in the sagittal and axial orientations. The results for combination method 2 appeared slightly blurred in all the three orientations, but more so in the sagittal and axial orientations.

### 3.5.5 Selection of Method of Choice for Image Combination

The quantitative results of the test images indicated that combination method 2 performed best of all followed by combination method 1, then the spline interpolation. This order was also evident on qualitative inspection of both the test images and the images acquired for this study.

However, the neuroradiologists assisting with the segmentation expressed a preference for the images obtained by combination method 2. They felt those of combination method 1, although sharper were not “natural” because of the blocky discontinuities arising because of the computational constraints. Furthermore, they felt that the quality of the images produced by combination method 1 were adequate for the purposes of verifying the segmentation of the lateral ventricles.

The “blockiness” of combination method 1 was most probably due to the modified approach to the solution of the linear equations. It may be possible to solve the linear equation for a complete image, given significant computational resources, but that route was not followed in this study.



**Figure 3.6:** Qualitative results of the application of reconstruction methods to an image of one of the subjects of this study



The surface of the images used in the construction of the PDM of the lateral ventricle were segmented from combined images of the  $T_2$  axial, sagittal and coronal views of each subject using combination method 2 to give an image of 0.78mm isotropic voxels. Figure 3.7 shows the surface of one ventricle obtained from a coronal view of resolution  $0.78mm \times 0.78mm \times 6mm$  compared with one obtained from the combined image of the same subject.

## 3.6 Obtaining Brain Size Parameters

The need for normalisation for brain shape was discussed in section 2.5.4. For each subject, brain size parameters were obtained as follows. Skull stripping was performed on each MR image using the Brain Extraction Tool of the FMRIB group at Oxford [104], and ellipsoids were fitted to the resulting brains. The lengths of the three principal axes of the ellipsoids were stored as the brain size parameters. The ventricular surfaces were aligned to a canonical coordinate system using their centroids and the three principal axes obtained from the distribution of the coordinates of their surface points. The brain size parameters were then used to scale each ventricle independently in the three orthogonal directions for normalisation for brain size with respect to the brain size of an arbitrarily chosen template brain. Figure 3.8 shows the surface of the brain of one subject and an ellipsoid fitted to it.

## 3.7 Summary

In this section we have given details of the data used. We have shown that the application of non-uniformity correction allows better classification of voxels into the three primary classes of brain tissue and therefore results in improved images. A semi-automatic segmentation process was adopted which allows segmentation of individual left and right ventricles. To improve spatial resolution, independently acquired coro-




(a) Ventricular surface obtained from segmentation of the coronal view of a subject. The image resolution was  $0.78mm \times 0.78mm \times 6mm$



(b) Ventricular surface obtained after performing image combination to give a resolution of  $0.78mm \times 0.78mm \times 0.78mm$  using combination method 1

**Figure 3.7:** Rendering of a ventricular surface before and after image combination



**Figure 3.8:** A brain surface of one subject is shown with an ellipsoid fitted to it. The length of the three main axis of the ellipsoid give the scaling parameters to be applied in the coronal, sagittal and axial directions

nal sagittal and axial views of each subject were combined to give isotropic voxels with nominal sub-millimetre resolution. Brain size parameters were obtained by fitting ellipsoids to the brains of each subject. These will be used in normalisation for brain size during shape analysis.

# Chapter 4

## 3D Point Distribution Models - Theory and Practice

### 4.1 Overview

This chapter centres on the construction of the 3D PDM used for morphometric analysis of the lateral ventricles. We base our approach to construction of the PDM on the framework suggested by Subsol et. al. [108] in which correspondence is established using curvature-based landmarks automatically extracted from the surface of segmented ventricles.

We begin by giving a more in-depth introduction to PDMs (compared to that in chapter 2) in section 4.2 then we review approaches that have been taken to the construction of 3D PDMs in section 4.3. Section 4.4 discusses how to compare PDMs in a quantitative manner, and in this section we introduce a distance transform approach which we believe is more appropriate than current methods.

Section 4.5 gives an overview of our approach to the construction of the 3D PDM of the lateral ventricles. This involves creating point-to-point matches of the curvature-

based landmarks and selection of one ventricle to initiate the model building process. In section 4.6 we describe surface curvature and how curvature-based landmarks on the ventricular surface are obtained. Point-to-point matching is the subject of section 4.7. In this section we describe how point matching can be formulated as a transportation problem. Section 4.8 investigates the dependence of our approach on the ventricle chosen to initiate the model building process. The chapter ends with a summary of the main points discussed.

## 4.2 Point Distribution Models - The Theory

PDMs were briefly mentioned in section 2.5.3. Here a more in-depth description is given as well as the framework by which they are created. The description in this section is based on the seminal work of Cootes et al. [33].

The aim of a PDM is to provide a compact model of the variability within corresponding points of aligned versions of the training set. For the present purposes, each example of the training set is a point in a  $3n$ -dimensional space, where  $n$  is the number of landmark points, with each landmark point having  $x$ ,  $y$ , and  $z$  cartesian coordinates. A set of  $N$  example shapes gives a cloud of  $N$  points in this  $3n$ -dimensional space. The region of space defined by the cloud is called the *Allowable Shape Domain* i.e. any point in this region will give a set of landmarks whose shape is similar to that of the class of shapes in the training set.

The approach of [33] was to attempt to model the shape of this cloud in a reduced dimensional space (number of dimensions less than  $3n$ ) in order to capture the relationships between the positions of the individual landmark points. It is assumed that the shape of the cloud is that of a  $3n$ -dimensional hyperellipsoid and that it can be defined by calculating its centre and its major axes. The centre is taken as the mean shape of the training set and orthogonal basis vectors (the major axes) are calculated using Principal Component Analysis (PCA). This allows a reduction in

dimensionality by describing the shape in terms of a set of basis vectors, which are the axes of the ellipsoid representing most of the variability. Any point (shape) in the reduced dimensional space can be reached by using the mean shape (origin) and a linear combination of the basis vectors, and a PDM encapsulates this.

### 4.2.1 Mathematical Definition of 3D PDMs

Given a training set composed of  $N$  members, with each member having  $n$  landmark points in 3D cartesian space, a *shape vector* is formed for each member of the training set by concatenation of the  $x$ ,  $y$  and  $z$  coordinates of its landmark points. For the  $i^{th}$  member of the training set this is :

$$\mathbf{x}_i = [x_{i0}, x_{i1}, \dots, x_{in-1}, y_{i0}, y_{i1}, \dots, y_{in-1}, z_{i0}, z_{i1}, \dots, z_{in-1}]^T \quad (4.1)$$

To attain invariance to Euclidean transformations, differences due to orientation, scaling and position are removed by an alignment step described in section 4.2.2. The shape vectors referred to in the rest of this section are the Euclidean invariant aligned shape vectors. From the aligned shape vectors of all members of the training set, a mean shape vector can be obtained :

$$\bar{\mathbf{x}} = \frac{1}{N} \sum_{i=1}^N \mathbf{x}_i \quad (4.2)$$

and principal component analysis can be carried out by obtaining the eigenvalues and the eigenvectors of the covariance matrix of the aligned shape vectors:

$$\mathbf{S} = \frac{1}{N} \sum_{i=1}^N (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})^T \quad (4.3)$$

Eigen-decomposition of the sample covariance matrix  $\mathbf{S}$ , gives:

$$\mathbf{S}\mathbf{p}_k = \lambda_k \mathbf{p}_k \quad (4.4)$$

where  $\lambda_k$  is the eigenvalue corresponding to the  $k^{th}$  eigenvector  $\mathbf{p}_k$ , with  $k$  ranging from 1 to  $3n$ . If the number of examples is much less than the number of landmark points ( $N \ll 3n$ ), which is generally the case for 3D PDMs, the maximum number of degrees of freedom of the model is  $N - 1$  which is also the maximum number of non-zero eigenvalues the  $3n$  by  $3n$  covariance matrix may have. Therefore, it is not necessary to calculate all the  $3n$  eigenvectors, and [33] gives an appendix showing an efficient method for the derivation of the eigenvectors of the non-zero eigenvalues in this case (where  $k$  ranges from 1 to  $N - 1$ ).

The dimensionality of the shape space is reduced to  $t$  dimensions from  $k$  dimensions by selecting  $t$  so that the  $3n$ -dimensional hyperellipsoid is approximated by a  $t$ -dimensional hyperellipsoid.  $t$  is a positive integer less than  $k$  and represents a cutoff so that eigenvalues for  $\lambda_k$  ( $k > t$ ) are ignored.  $t$  is chosen so that a large proportion of the total variance,  $\lambda_T = \sum_{k=1}^{3n} \lambda_k$ , is encompassed by  $\sum_{k=1}^t \lambda_k$  where  $\lambda_k \geq \lambda_{k+1}$  (i.e. the eigenvalues are arranged in descending order).

$$\sum_{k=1}^t \lambda_k \geq f_v \lambda_T \quad (4.5)$$

where  $f_v$  specifies the fraction of the total variation retained.

Each eigenvector describes a *mode of variation* of the PDM. Eigenvectors corresponding to the largest values of  $\lambda_k$  describe the longest axes of the ellipsoidal approximation to the shape space, and give the most significant modes of variation. Any shape in the training set and indeed any allowable shape in the class can be reached by taking the mean shape and a linear combination of the eigenvectors :

$$\mathbf{x} = \bar{\mathbf{x}} + \mathbf{P}\mathbf{b} \quad (4.6)$$



The above equation is essentially the mathematical definition of a PDM.  $\mathbf{P}$  is a matrix whose columns are the first  $t$  (in descending order) eigenvectors of the covariance matrix i.e.  $\mathbf{P} = [\mathbf{p}_1 \ \mathbf{p}_2 \ \mathbf{p}_3 \ \cdots \ \mathbf{p}_t]$ , and  $\mathbf{b}$  is a vector of weights.

The values of the elements of  $\mathbf{b}$  can be constrained in two ways. Firstly, the variance of the eigenvalue corresponding to a given element in  $\mathbf{b}$  can be used. Assuming a normal distribution, a given percentage of the population will be expected to lie within a range of standard deviation units - e.g. 95% within 2 sd's giving  $-2\sqrt{\lambda_k} \leq b_k \leq 2\sqrt{\lambda_k}$ . This constrains shape vectors to be within a hypercuboid in shape space, centred on the mean vector.

Secondly, the Mahalanobis distance ( $D_m^2$ ) can be used. The Mahalanobis distance is a distance metric in which the Euclidean distance is weighted by the inverse of the sample covariance matrix.

$$D_m^2 = (\bar{\mathbf{x}} - \mathbf{x})^T \mathbf{S}^{-1} (\bar{\mathbf{x}} - \mathbf{x}) \quad (4.7)$$

where  $\bar{\mathbf{x}}$  is the mean shape vector,  $\mathbf{S}$  is the sample covariance matrix, and  $\mathbf{x}$  is the shape vector of the shape instance under consideration. A given value of  $D_m^2$  specifies a hyperellipsoid centred about the mean shape, thus confining instances generated by the PDM to a specified region of shape space. The elements of  $\mathbf{b}$  are chosen so that the Mahalanobis distance from the mean is less than a pre-selected value  $D_{max}^2$ .

$$D_m^2 = \sum_{k=1}^t \left( \frac{b_k^2}{\lambda_k} \right) \leq D_{max}^2 \quad (4.8)$$

$D_{max}^2$  is determined from the distribution of the Mahalanobis distances of members of the training set.

In the above, there has been an implicit assumption that the position of each shape instance about the mean can be obtained by a linear combination of the eigenvectors of the covariance matrix. In certain cases the correlations between landmark points may not be linear and the hypercuboid/hyperellipsoid constraints do not always result in

meaningful (“legal”) shapes. However, as long as the position of individual points do not vary very much, the linear assumption is acceptable. Non-linear PDMs [126] are being researched in attempts to address cases where linear PDMs are inappropriate.

### 4.2.2 Steps to Create a Point Distribution Model

Given a training set, the first step in the creation of a PDM is labelling each member of the training set with the same number of homologous landmarks to establish correspondence. It is important that the points on each example in the training set are labelled as accurately as possible so that variation in position of homologous landmarks between examples is not due to erroneous labelling. Davies et al. [39] have shown that certain properties of SSMs depend critically on finding accurately corresponding landmark points. The labelling issue and correspondence were introduced in section 2.5.2. Section 4.3 below reviews various methods that have been used to establish correspondence in the construction of 3D PDMs and sections 4.6 to 4.8 detail the approach taken here to automatically establish correspondence.

Once labelling is completed, the resulting shape vectors formed by the concatenation of the coordinates of the labelled points (as specified by equation 4.1) have to be aligned to remove differences due to location, scale and orientation. To remove differences due to location, an object-centred coordinate system is obtained using the centroid of each shape to translate the coordinates of its landmark points accordingly. e.g. for the  $x$  coordinate of each landmark point of the first member of the training set, the object-centred coordinate of the  $i^{th}$  landmark point,  $x'_{0i}$  is :

$$x'_{0i} = x_{0i} - \frac{1}{n} \sum_{i=1}^{n-1} x_{0i} \quad (4.9)$$

and the object-centred  $y$  and  $z$  coordinates are obtained similarly.

To remove scale and orientation differences, Generalised Procrustes Analysis [58] is carried out within the iterative framework suggested in [31]. Generalised Procrustes

Analysis is a widely used method in shape analysis for the least squares alignment of two or more shapes represented as matrices of coordinates of corresponding landmark points. It optimally matches each landmark configuration in the sample to an average configuration determined from all the shapes in the sample. This is achieved by finding translation, rotation and isotropic scaling factors for each shape that minimises a criterion based on the sum of squared differences of corresponding points across the whole training set.

The iterative framework of [31] includes weighting factors for each landmark point based on the variance of the position of the landmark over the training set, and penalising those with greatly varying positions. The intention of this is to give greater significance to points that tend to be more “stable”. Weighting is applied using a vector of weights,  $\mathbf{w}$ , of length  $n$ . Each element of  $\mathbf{w}$  is the weight applied to the corresponding landmark. The weight for the  $k^{th}$  landmark point,  $w_k$  is given by

$$w_k = \left( \sum_{l=0}^{n-1} V_{kl} \right)^{-1} \quad (4.10)$$

where  $V_{kl}$  is the variance of the Euclidean distance between landmark  $k$  and landmark  $l$  over all shapes in the training set. The iterative procedure applied during the construction of the 3D PDMs of the ventricles is outlined in figure 4.1. The rigid registrations referred to in the outline have 7 degrees of freedom - translation in the three orthogonal directions, rotation about the three axes, and isotropic scaling. The aligned shape vectors obtained are used to calculate the PDM parameters described in section 4.2.1

### 4.3 A Review of Approaches to Construct 3D PDMs

A key step in building a PDM is establishing correspondence between points on training images. In section 2.5.2 we mentioned that establishing correspondence in 3D was not straightforward. The manual approach to landmarking usually used in 2D is not

Given a training set of labelled shapes:

1. Transform the coordinates of the landmark points of each shape to obtain object-centred coordinates using equation 4.9.
2. Calculate weight vector using equation 4.10.
3. Align each shape to the first shape in the training set by finding the rigid body transformation that minimises the sum of the weighted Euclidean distance between the corresponding landmark points of the first shape and the current shape.
4. Repeat the following :
  - (a) Calculate a current mean using equation 4.2.
  - (b) Align the current mean shape to the first shape in the training set (as in 3 above).
  - (c) Align all shapes in the training set to the current mean (as in 3 above).
5. Until there is convergence of Euclidean distance between corresponding points.

**Figure 4.1:** Outline of algorithm for alignment of shapes in a training set

practical for the 3D case because of the large number of points needed to truly represent shape and the practical difficulty of viewing and comparing surfaces. A number of approaches have been developed to address the issue of landmark generation in 3D and the related issue of automation of landmark generation in general. This section reviews methods that have been applied to solve the correspondence problem in the building of 3D PDMs. A common feature of all the methods reviewed (except those of Hill et. al. [63] and Andresen et. al. [3]) is that the 3D PDM is constructed from a representation of the surfaces of the members of the training set in which correspondence has been established. That is to say the landmarks are not placed on the actual surfaces of the members of the training set but on an approximated surface derived from the point locations. We refer to the training set used for the construction of the 3D PDM as the *corresponding training set*. We will address issues arising from this approximation in section 4.8.

The first practical extension of PDMs from 2D to 3D was by Hill et. al. [63]. They

used 2D contour data from contiguous slices of a 3D volume as the basis for generation of their 3D PDM of the lateral ventricles (minus the temporal horns, modelled as three sub-parts), the caudate nucleus, a section of the brain stem and a section of the subcutaneous fat under the skin of the head. They assumed that the structures being modelled had simple topology which did not vary from slice to slice. Contour data were landmarked interactively with the same number of landmarks on contours in each slice. Within each slice the landmarks followed an ordered sequence which was used to impose correspondence between slices, and curves were drawn through the locus of points of corresponding landmarks on each slice. These curves in 3D space were then subdivided and discretely sampled to give the points used to construct the 3D PDM. The major drawbacks of this method are that it is limited to cylindrically shaped structures, it involves extensive manual interaction in placing landmarks on each slice, and the correspondences obtained are not based on anatomical features. Furthermore, the requirement that the same number of landmarks is placed on contours in each slice is impractical, as the size of contours usually taper towards the ends of structures.

Hill et. al. [64] introduced a novel framework for automatic landmark identification to establish correspondence over training sets of the shapes represented by their boundaries in 2D. The main contribution of their work was the description of a *pair-wise corresponder* to allow the construction of a mean shape from the members of the training set in a hierarchical manner. Brett and Taylor [22] extended this work to describe a method of automated 3D PDM construction. Triangulated polyhedral surface representations of each member of the training set were generated from 2D contours. The pair-wise corresponder was used to construct a binary tree with a mean shape at the root and the members of the training set at the leaf nodes. At the leaf node level of the tree, the pair-wise corresponder generated sparse triangulated approximations of pairs of members of the training set and established correspondences between their sparse approximations. A dense triangulation using the pair of corresponding points was then obtained by recursively inserting vertices between the vertices of the sparse triangulation. This was repeated at each level until the

root level was reached. The dense triangulation at the root level is the mean shape, and the vertices of this are the landmark points of the mean shape. These landmark points are propagated down the tree to the members of the training set on the leaf nodes generating the labelled training set. A drawback of this method is that the selection of landmark points is dependent on the triangulation of the surface, and only loosely coupled to the shape of the class of objects. Furthermore, the use of the pair-wise corresponder requires the evaluation of all possible pairings in the training set, which may become prohibitive as the size of the training set increases.

Gerig et. al. [54] created a 3D PDM of the lateral ventricles using the SPHARM parameterisation of Brechbühler [21]. The ventricular surface of each member of the training set was mapped onto a unit sphere from which spherical harmonic functions were obtained. Correspondence was imposed by aligning the ventricles using the major axes of the ellipsoids given by the first term of their harmonic functions. The continuous parameterisations of the surface are then finely sampled to give point representations of each member of the training set. However, the physical relationships of the corresponding points obtained by this method are unclear. Davies [38] has shown that the properties of PDMs created in this manner are inferior (according to criteria for quantitative assessment of PDMs described in section 4.4) to those of PDMs created by optimisation of a minimum description length cost function. Gerig et. al. have subsequently modified their method to incorporate a medial axis transform. This was described in section 2.5.2.

Davies [38] describes a method for automatically placing landmarks on 2D and 3D PDMs by optimisation of a minimum description length (MDL) objective function. The members of the training set are segmented contours of the structure of interest, and have to be parameterised and mapped to a unit sphere. Landmark points are placed on the parameterised representations in a multi-resolution manner by minimisation of a MDL based objective function. The final landmarks are then inversely mapped back unto the spatial domain of the members of the training set. This method has advantages in that it is relatively general as it can potentially be applied

to any class of shapes as long as they can be mapped to a sphere. Additionally, it considers all the members of the training set simultaneously unlike other methods e.g. [22], [79] that match the members of the training set in a pair-wise manner. However, the parameterisation of 3D shapes is not trivial and the effectiveness of the model is determined by how well the parameterisation method works. This, combined with the fact that the complexity of the shape of some structures in the body do not allow straightforward mapping of their surfaces unto a sphere, currently limit the use of this method in 3D.

Lorenz and Kranhnsstöver [79] use a template based method in their construction of 3D PDMs. They create a landmarked template of the class of objects to be modelled, and use this to coat a surface representation of each member of the training set. The template was acquired by Delaunay triangulation [13] of the binary voxel representation of a randomly selected member of the training set to give vertex points and triangle faces. The vertex points were taken as the landmark points of the template. The template is registered to each member of the training set using a small number of manually defined landmarks which have to be placed interactively on the surface of each member of the training set. The method was applied to the construction of a 3D PDM of lumbar vertebrae. This method is not fully automatic as it involves user interaction in placing landmarks. Accurate placement of landmarks on 3D surfaces is difficult to achieve and the choice of template may influence the properties of the resulting model.

Shen et al. [100] define point correspondences using an “Adaptative Focus Deformable Model (AFDM)”. This uses both a triangulated representation of the surface of the structure to be modelled generated from segmented contours, as well as image data from the actual 3D volumetric MR images. Attribute vectors are created for each vertex of the triangulation of each member of the training set. The attribute vectors reflect the geometric structure of the model in a hierarchical manner from local to global level. Correspondences are based on moving the vertices of the triangulations to minimise an energy term created from both the triangulations (using the attribute

vectors of the vertices) and the image data (designed to move the vertices towards boundary structures). This method of establishing correspondence attempts to use both segmented data and actual image data - the distinction being that binary segmented data only provides geometric information, whereas image data allows gray level values of a neighbourhood about two points to be compared. Disadvantages of this method include the requirement of manual intervention in some cases to “pull” the triangulated surface towards image boundaries, and the fact that the initial correspondences are defined by the triangulation of segmented data.

Frangi et al. [53] use a distance transform method to obtain a template of a training set of segmented binary voxel images. Landmark propagation from the template to each member of the training set is achieved by “volumetric elastic registration”. One member of the training set is randomly selected as the initial template. Each member of the training set is then registered with the template using a nine degree of freedom quasi-affine registration to minimise a similarity metric based on normalised mutual information. Euclidean distance transforms of each registered image are obtained using a convention of negative values within the subject of interest and positive values outside it. An average image is computed by obtaining the mean of the distance transformed images. This is converted to a binary representation by thresholding at the zero-level set. This average image is taken as the new template. A triangulated surface is obtained by applying the marching cubes algorithm [78] to its binary image and the vertices of the triangles are used as the landmark points of the template. To propagate the landmarks of the template onto each member of the training set, the binary image of the fiducial member of the training set is warped onto the average image using a volumetric elastic registration process. The inverse transformation is applied to the landmarks of the template to map them onto the surface of the fiducial image. This gives a set of corresponding points from which a 3D PDM can be built. The method was demonstrated by building 3D PDMs of the head of the radius (a bone of the arm) and the caudate nucleus of the brain.

Lamecker et al. [75] describe a “patchification” method to tackling the correspon-



dence problem. They use manually placed landmarks to divide the surface of each member of the training set into the same number of surface patches. Each patch represents similar regions, and is topologically equivalent to a disk. Each patch is mapped onto a disk, and correspondence between equivalent patches of a pair from the training set are obtained by solving a sparse linear system of equations to obtain the mapping that minimises local scaling and shearing. They applied this method to the construction of a 3D shape model of the liver. The shape model was intended to be applied to segmentation of the liver from abdominal MR images. An obvious disadvantage of this method is the need for manual intervention to specify the landmarks defining the surface patches. Furthermore, the fact that the surface is divided in this manner may result in non-optimal correspondences, as the correspondences obtained are directly dependent on the surface subdivision. The authors cite an advantage of the method as the fact that it is applicable to shapes of any topology, and state that the intention is to apply automatic techniques to the surface sub-division stage. However, the application of automatic techniques will result in restriction to surfaces on which salient landmarks can be reliably detected automatically.

Subsol et al. [108] describe a framework for automatically constructing 3D morphometric atlases which they proposed could be used as a basis for the construction of 3D PDMs. Their method was based on the identification of curvature extremum points on the surfaces of the members of the training set. A sparse correspondence between each member of the training set and a template is obtained in a pair-wise manner, and this is used to define a dense correspondence for warping a predefined set of landmark points from the surface of the template onto the surface of each member of the training set. Andersen et. al. [3] have used this framework as the basis of their method of construction of a 3D PDM of the human mandible. This framework was also used as the basis of the construction of the 3D PDM of the lateral ventricles used for the shape analysis presented in this thesis. Section 4.6.3 reviews the work of Andersen et. al. and sections 4.6 and 4.7 give further details of the framework suggested by Subsol et. al.

## 4.4 Quantitative Assessment of 3D PDMs

A number of methods by which 3D PDMs have been created were reviewed in the previous section. However, there is a dearth of studies in which different methods of 3D PDM construction are quantitatively compared. Furthermore most of these methods use an approximation to the surfaces in the original training set (the corresponding training set) as opposed to the actual surfaces of the members of the training set in constructing PDMs. We encountered no studies in which the fidelity of the surfaces of the corresponding training set to those of the original training set were investigated quantitatively. Lastly, most of the methods reviewed involved the selection of one member of the training set to initiate the model building process. Again there has been a dearth of quantitative investigation of the effect the initial choice has on the final PDM.

In this section three measures used for quantitatively assessing 2D PDMs - *Generalisation Ability*, *Specificity*, and *Compactness* are described. To our knowledge the only quantitative comparison of two methods of 3D PDM construction was undertaken by Davies [38]. This compared a PDM of the hippocampus produced by the MDL approach with one obtained from the same training set using the SPHARM approach of Gerig et. al. [54] and showed that the MDL approach produced PDMs with better properties.

The comparisons were with respect to the corresponding training set (approximations to the surfaces of the actual training set in which correspondence has been established). However, we argue that they should be with respect to the actual training set. We propose the use of 3D distance transforms in a novel manner to give a metric showing the similarity of a triangulated surface with the surface of a binary volumetric image. This metric can be used to allow comparisons of PDMs based on the surfaces of the original training set (if they are available as binary volumetric images). It also allows the corresponding training set to be compared with the original training set, and investigation of the dependence of the resulting PDM on the

member of the training set selected to initiate the PDM building process.

#### 4.4.1 Generalisation Ability

A PDM models a class of shapes using a sample of the class assumed to be representative of the class. Cootes [31] describes the generalisation ability of the resulting PDM as a measure of the ability of the model to exhibit all the variation expected in the class of shapes being modelled.

To obtain a quantitative measure of the generalisation ability, Cootes proposed a set of leave-one-out experiments. For a training set composed of  $N$  members, one member is left out in the construction of the model. The model is then fitted to the left out example, and the error of the fit is used as a quantitative measure of the generalisation ability. The leave-one-out procedure is repeated for each member of the training set. To further characterise the properties of the PDM, the model fitting stage can be performed as a function of the number of modes of variation in the model. To do this, during the model fitting stage, the number of modes used is increased from the first (most significant mode of variation) mode to the last (least significant mode of variation) mode in a cumulative manner. For two PDMs constructed in different ways, this allows a mean error-of-fit and the associated error of the mean to be calculated as a function of the number of modes used in the fitting. This gives a measure of generalisation that is independent of compactness (see section 4.4.3). The PDM with the lower error-of-fit is said to have better generalisation ability.

The traditional method used in obtaining the error of the fit is the sum of squared differences between corresponding points of the model fit and the member of the training set that the fit was applied to. For 2D PDMs where the corresponding training set is an exact representation of the actual shape, but on which corresponding points have been placed, this can be considered an adequate way to obtain the error-of-fit. However, for 3D PDMs where the corresponding training set is usually an approximation of the actual surfaces, AND this can be different for two different

methods, it makes more sense to obtain the error-of-fit from the actual surfaces of the training images, which is the same for both methods. Doing this ensures that the error-of-fit is not dependent on the fidelity to which the corresponding training set represents the original surfaces. However, it also presents the problem of how to obtain the measure as there are no correspondences between the generated instance and the actual surface. In measuring generalisation ability we evaluate both metrics: the point-to-point distance between a generated instance and the annotated points on the training surface and the distance transform between the generated instance and the original surface.

In performing the leave-one-out experiments, one alignment step including all the members of the training set is performed. When building a model with an example left out, re-alignment is not performed. This is because the alignment step for all 138 ventricles takes a considerable amount of time, and we can expect that the omission of one example would not result in significant changes to the overall alignment. For the  $i^{th}$  left out example,  $\mathbf{x}_i$ , the weight vector  $\mathbf{b}_i$  is given by rearranging equation 4.6:

$$\mathbf{b}_i = \mathbf{P}^T(\mathbf{x}_i - \bar{\mathbf{x}}) \quad (4.11)$$

The model fit to  $\mathbf{x}_i$  is obtained by using  $\mathbf{b}_i$  in equation 4.6.

#### 4.4.2 Specificity

The specificity of a PDM is a measure of its ability to generate “valid” shapes. As mentioned in section 4.2, the PDM is a linear approximation of the allowable shape domain. The parameters of the PDM are learnt from the training set, and the elements of the  $\mathbf{b}$ -vector are weights which restrict the generation of new shapes to those in the allowable shape domain. If the allowable shape domain is not modelled accurately, some of the shapes generated by the PDM may be “illegal” for the class of shapes being modelled and the resulting PDM is not very specific.

The specificity of a model can be qualitatively assessed by viewing randomly generated shapes and forming a subjective opinion of their validity. However, for objective assessments, quantitative measures are needed. These can be obtained in two ways. One is to generate a number of random examples and obtain a measure of their similarity to the shapes in the training set. The other is to assume the shapes in the training set are samples from a probability density function. Then, given any shape, we can calculate the probability that it is a plausible example of the class of shapes described by the training set.

Examples of practical application of the second method were not encountered in the literature reviewed. This may be because it is not easy to estimate the probability distribution function of the class of shapes. Davies [38] used the first method and defined the measure of similarity between two shapes as the sum of squared differences between corresponding points on the two shapes. However, as explained in section 4.4.1, this may not be adequate in the present case. Therefore, in addition to using the sum of squared differences, we have used the distance transform metric to obtain a measure of the similarity of the generated random shape to the actual surfaces of the training images.

The generation of random shapes is carried out by randomly generating values for the elements of the  $\mathbf{b}$  vector, using a multivariate gaussian distribution. For each  $\mathbf{b}$  vector, the shape generation is carried out as a function of the number of modes in a cumulative manner as in section 4.4.1.

### 4.4.3 Compactness

The amount of variation encompassed by a PDM is captured in the eigenvalues of the eigenvectors of the covariance matrix  $\mathbf{S}$  of equation 4.3. The compactness of a PDM is a relative measure based on the amount of variation captured by the model. For two PDMs of the same training set, a more compact model represents the same amount of variation with a smaller number of modes of variation than a less compact

one. The compactness is measured in a cumulative manner by summing the values of the eigenvalues of each mode of variation and normalising by the overall sum of all eigenvalues, i.e. for the  $i^{th}$  mode of variation the compactness  $C(i)$  is given by

$$C(i) = \frac{\sum_{j=1}^i \lambda_j}{\lambda_T} \quad (4.12)$$

Where  $\lambda_T$  is the total variance (sum of all eigenvalues)

#### 4.4.4 A Distance Transform Metric for Comparing Two Surfaces

In this section we describe the use of a 3D distance transform to facilitate the comparison of a triangulated surface with those of a surface represented by a binary voxel image. We assume that the surfaces being compared are surfaces of objects of the same class and therefore have similar shape and topology.

Borgefors [20] defines a distance transform as “an operation that converts a binary picture, consisting of feature and non-feature elements, to a picture where each element has a value that approximates the distance to the nearest feature element”. In the present case, the “binary picture” is the 3D array of voxels of the segmented ventricles, with the feature elements being the voxels that are on the surface of the ventricles. A *D-Euclidean* metric [20] is used to compute values for the transform. This assigns positive values to each voxel in proportion to an approximation of its lowest Euclidean distance from the surface voxels (the values of the distance transform of the surface voxels are zero). Figure 4.2 shows two slices through the binary voxel image of one ventricle, and the resulting distance transforms of the contours.


For our purposes, one surface is a triangulated surface made of vertices and faces, whilst the other is a 3D binary array of voxels. The triangulated surface is embedded in the distance transform of the binary voxel image. For each vertex of the triangu-

lated surface, the value of the distance transform of the voxel in which it is embedded is obtained. The mean of these values is used as an objective function for a rigid registration. The *fminsearch* function of Matlab is used to obtain the 6 degree-of-freedom rigid registration giving the minimum value of the objective function. This minimum value is taken as the final index of the similarity of the triangulated surface and the voxel image. The rigid registrations are needed because the vertices of the triangulated surface are in an object-centred coordinate system relative to the centroid of the vertices. The coordinates of the binary voxel image on the other hand are relative to the bottom left of the voxel array and are dependent on the position of the patient's head during the MR scan.

## 4.5 Our Approach to the Construction of a 3D PDM of the Lateral Ventricles

In this section we describe our approach to construction of the 3D PDM of the lateral ventricles, and introduce some terminology that will be used frequently in the subsequent text. Correspondence amongst members of the training set is established in a pair-wise manner. One ventricle ( $\mathbf{v}_t$ ) was used as a *template* and its surface represented as vertices and vertex faces defined by triangular triplets of the vertices. The initial triangulation of the binary voxel image of  $\mathbf{v}_t$  produced about 10,000 vertices, but for computational reasons these were decimated by an order of magnitude to give 1,291 vertices. We use the term *fiducial* to refer to a particular member of the training set on which an operation is being performed.

Curvature based landmarks (crest lines and crest points described in section 4.6) were obtained for each ventricle. The crest lines of the template were matched to those of the remaining 137 ventricles  $\mathbf{v}_i \in \{\mathbf{v}_1, \dots, \mathbf{v}_{137}\}$  in a pair-wise manner. The matching was performed in both directions i.e  $\mathbf{v}_t \rightarrow \mathbf{v}_i$  and  $\mathbf{v}_i \rightarrow \mathbf{v}_t$ . The method of point-to-point matching is the subject of section 4.7.



**Figure 4.2:** Binary contours of two slices of the ventricle (a) and (c) and their respective distance transforms



The symmetric subset of matched points gives a sparse correspondence which is used to obtain coefficients defining a final spline based warp allowing transformation of the vertex points of  $\mathbf{v}_t$  onto the surface of each  $\mathbf{v}_i$ , giving the corresponding training set. The ventricles  $\mathbf{v}_i$ , to which the template is deformed, is referred to as the *target*. The landmark positions on the corresponding training set were re-projected onto triangulated surfaces of the binary voxel image of each member of the training set, and smoothed as described in section 4.8.

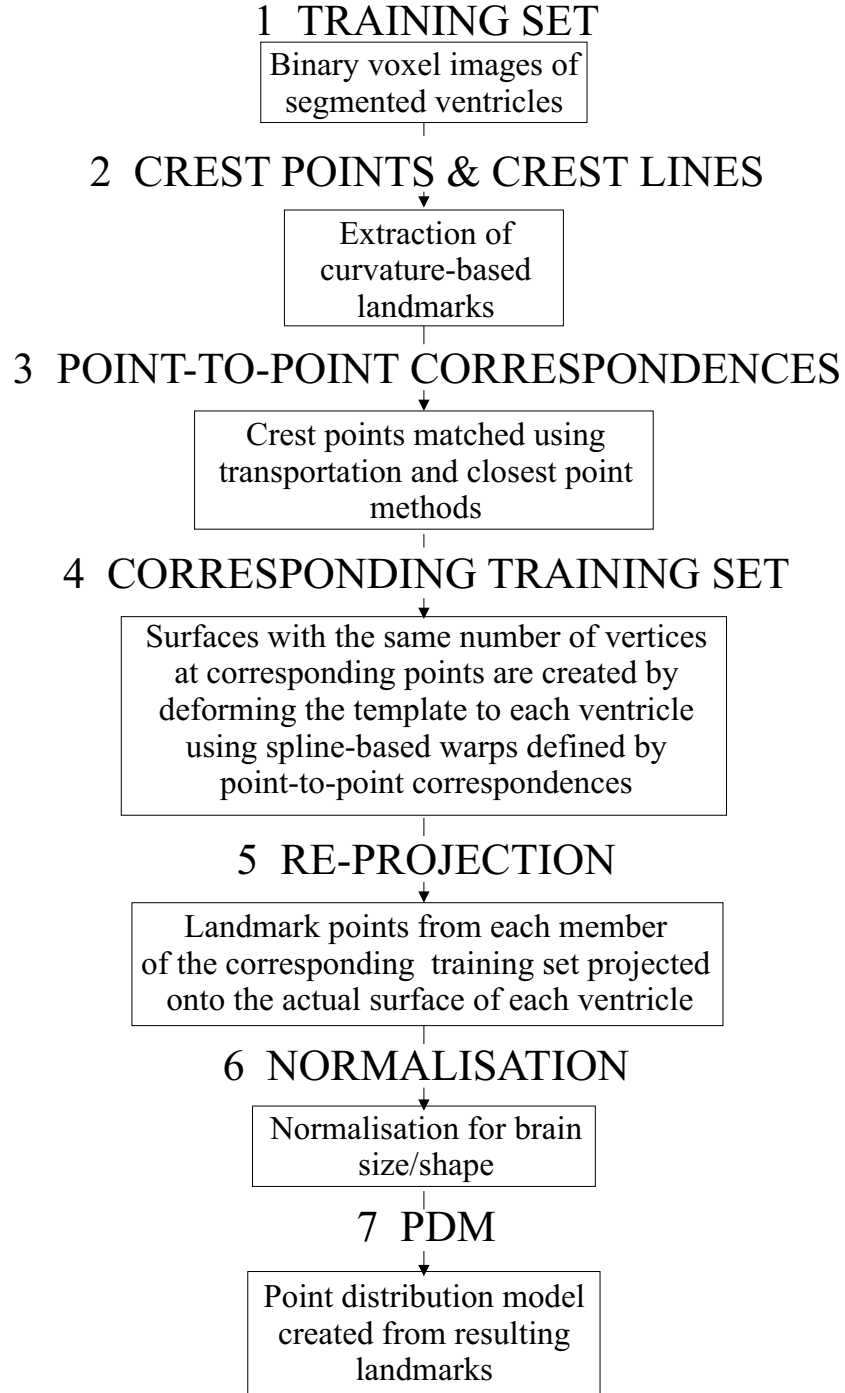
The brain size parameters that had been obtained as described in section 3.6 were used to normalise the brain of each member of the training set with respect to that of the template in order to reduce effects due to different brain shapes and sizes. The PDM was built in the usual manner (section 4.5) using the resulting landmark points.

The schematic diagram of figure 4.3 summarises the key stages in the construction of the PDM. The curvature-based landmarks on the surface of the ventricles are the subject of section 4.6. Section 4.7 discusses establishing a sparse correspondence across the training set using these curvature-based landmarks. Section 4.8 investigates the consequence of the need to select one ventricle as a template in our pair-wise approach on the resulting PDM. The key question being asked is whether the resulting PDM is dependent on the ventricle chosen as a template.

## 4.6 Curvature-Based Landmarks on the Ventricular Surface

Curvature is an important geometrical property of surfaces, and as boundaries of segmented structures in 3D images are surfaces it is not surprising that surface curvature is of importance here.

Section 4.6.1 gives a brief overview of surface curvature. Detailed mathematical descriptions are available in differential geometry texts e.g. [43], and [65] (this gives



**Figure 4.3:** Schematic diagram of the process of creating the 3D PDM of the ventricles. Quantitative analysis are carried out at steps 3 and 5.

a less mathematical but more practical treatment). Definitions of some curvature related terms are given including that of the *normal curvature* of a surface. The existence of maximal and minimal curvatures are stated and characterisation of surfaces based on these are illustrated. Finally, curvature extrema are discussed, and *crest points* and *crest lines*, which are used as curvature-based landmarks on the surface of the lateral ventricles, are defined.

Section 4.6.2 discusses different ways to obtain crest points and crest lines. The Marching Lines algorithm that was used as the method of choice in this thesis is described here. The use of crest points and crest lines as curvature-based landmarks is discussed in section 4.6.3. Examples of previous applications of crest points and crest lines as curvature based landmarks are reviewed, and their appropriateness in the case of the lateral ventricles is illustrated.

### 4.6.1 Curvature of 3D Surfaces


#### Basic Framework and Definitions

Consider a *regular* surface patch,  $S$ , above a plane in 2D,  $U$ , as shown in figure 4.4. Let  $U$  be a subset of two-dimensional real space  $\mathbb{R}^2$  parameterised by  $u$  and  $v$  ( $0 \leq u, v \leq 1$ ). Let  $\mathbf{m}$  be a mapping which takes points in  $U$  onto the 3D manifold  $S$ , i.e.  $\mathbf{m}(u, v) : U \mapsto S; S \subset \mathbb{R}^3$ .

For a curve defined in  $U$  as  $\mathbf{s}'(t) = (u(t), v(t))$  there exists its mapping on  $S$  given by

$$\mathbf{s}(t) = \mathbf{m}(u(t), v(t)) \quad (4.13)$$

The partial derivative  $\frac{\partial \mathbf{s}}{\partial u}$  gives a tangent vector to the curves on the surface for which  $v = \text{constant}$ , and is denoted  $\mathbf{s}_u$ . A tangent vector  $\mathbf{s}_v$  can also be defined in a similar manner.  $\mathbf{s}_u$  and  $\mathbf{s}_v$  are sometimes called the *basic vectors* of the surface.



**Figure 4.4:** A regular surface patch  $S$  over a 2D plane  $(u, v)$ . The coordinate lines show the curves on the surface for which  $u$  and  $v$  are constant. These define a net of curves on the surface. The basic vectors are also depicted on the surface.

### Tangents to Curves and Surfaces

The tangent to  $\mathbf{s}(t)$  at any point is obtained from its derivative :

$$\dot{\mathbf{s}}(t) = \dot{u}\mathbf{s}_u + \dot{v}\mathbf{s}_v \quad (4.14)$$

Where  $\dot{u} = \frac{du}{dt}$  and  $\dot{v} = \frac{dv}{dt}$ .

Now  $\dot{\mathbf{s}}(t)$  is a tangent vector to  $\mathbf{s}(t)$  and hence a tangent vector to  $S$ . Therefore, for any point  $\mathbf{p}$  on  $\mathbf{s}$  the tangent to  $\mathbf{s}$  at  $\mathbf{p}$  is also a tangent to  $S$ . Conversely, for any tangent  $\mathbf{t}_\theta^p$  to any point  $\mathbf{p}$  on  $S$ , there exists a curve through  $\mathbf{p}$  such that the tangent to the curve at  $\mathbf{p}$  is parallel to  $\mathbf{t}_\theta^p$ .

There are an infinite number of directions in which a curve along the surface can pass through  $\mathbf{p}$ . Each direction,  $\theta$ , has an associated tangent vector,  $\mathbf{t}_\theta^p$ , and all the possible tangent vectors lie in a plane called the tangent plane (see figure 4.5). There is however only one normal vector at  $\mathbf{p}$ , and this can be specified by its direction (the

unit normal,  $\mathbf{n}$ ) :

$$\mathbf{n} = \frac{\mathbf{s}_u \times \mathbf{s}_v}{|\mathbf{s}_u \times \mathbf{s}_v|} \quad (4.15)$$

### Curvature of a 3D Curve

For curves, curvature is a scalar denoting the “bendiness” of the curve at a point (i.e. the inverse of the radius of curvature of the circle that best fits the curve around the point). It is proportional to the rate of change of the tangent along the curve. For any curve  $\mathbf{s}(t)$  on  $S$ , its curvature at any point  $\mathbf{p}$  is given by the standard formula:

$$\kappa = \frac{|\ddot{\mathbf{s}}|}{\nu} \quad (4.16)$$


where  $\ddot{\mathbf{s}}$  is the derivative of  $\frac{\dot{\mathbf{s}}}{|\dot{\mathbf{s}}|}$ , the unit tangent to the curve at that point, and  $\nu$  is the rate of change of arc length of the curve (the speed of the curve).

### Surface Curvature

The curvature of a surface is a more complicated matter, with several different curvatures being defined. The notion of curvature of a space curve is used to define the *normal curvature* of a surface at a point  $\mathbf{p}$ . Given a tangent at  $\mathbf{p}$  in the direction  $\theta$ ,  $\mathbf{t}_\theta^p$ , the normal curvature is the curvature at  $\mathbf{p}$  of the curve formed by the intersection of the plane  $\Pi$  containing both the unit normal at  $\mathbf{p}$  and  $\mathbf{t}_\theta^p$  with the surface  $S$ . The curve formed by the intersection of the normal plane  $\Pi$  with the surface is called the *normal section* of the surface at  $\mathbf{p}$  in the direction  $\theta$  (see figure 4.5). The value of the normal curvature is given by:

$$\kappa(\mathbf{t}_\theta^p) = \frac{L\dot{u}^2 + 2M\dot{u}\dot{v} + N\dot{v}^2}{E\dot{u}^2 + 2F\dot{u}\dot{v} + G\dot{v}^2} \quad (4.17)$$

where  $E = \mathbf{s}_u \cdot \mathbf{s}_u$ ,  $F = \mathbf{s}_u \cdot \mathbf{s}_v$ ,  $G = \mathbf{s}_v \cdot \mathbf{s}_v$ ,  $L = \mathbf{s}_{uu} \cdot \mathbf{n}$ ,  $M = \mathbf{s}_{uv} \cdot \mathbf{n}$ , and  $N = \mathbf{s}_{vv} \cdot \mathbf{n}$ , are standard abbreviations used in differential geometry.  $\mathbf{s}_{uu}$ ,  $\mathbf{s}_{uv}$ , and  $\mathbf{s}_{vv}$  are the respective second order partial derivatives of the basic vectors. The numerator of equation 4.17 is called the *first fundamental form* and the denominator the *second*



**Figure 4.5:** Illustration of tangent and normal planes and the normal section through a surface

*fundamental form* of  $S$ .

The normal curvature is also a function of direction, therefore as the direction of  $\mathbf{t}_\theta^p$  changes (i.e. the angle  $\theta$  changes), the value of the normal curvature in that direction will in general change. According to a theorem by Euler (expressed as Euler's formula), there exists a maximum and a minimum value of normal curvature at every point on a surface (except *umbilic points* where the curvature in all directions is the same, such as on a sphere). The maximum and minimum curvatures are called the principal maximum and principal minimum curvatures,  $\kappa_{max}$  and  $\kappa_{min}$ . The directions of  $\mathbf{t}_\theta^p$  for which these extreme values are reached,  $\mathbf{t}_{max}^p$  and  $\mathbf{t}_{min}^p$ , are the principal directions at that point, and are orthogonal to each other.

Two further curvatures, of importance in characterising a surface are defined from the principal curvatures. The *mean curvature*,  $H$ , is the arithmetic mean of the principal curvatures, and the *Gaussian curvature*,  $K$  is their product (the square of their geometric mean) :

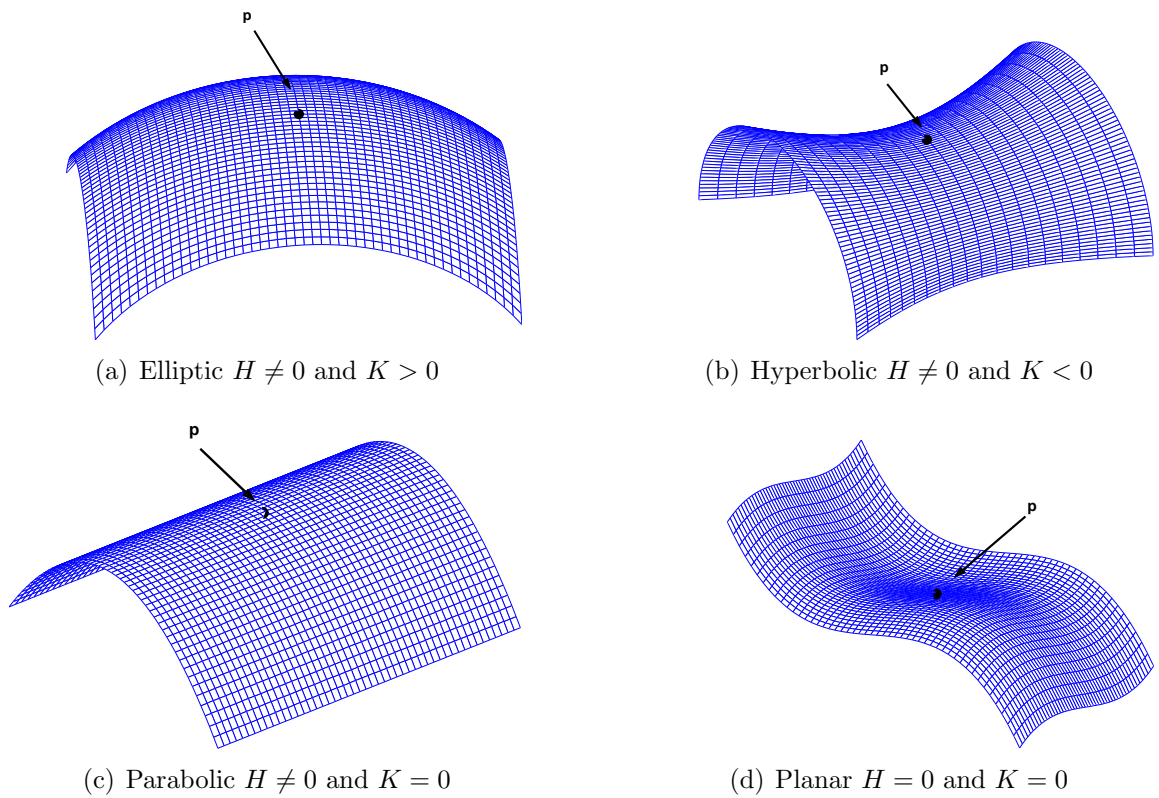
$$H = \frac{\kappa_{max} + \kappa_{min}}{2} \quad (4.18)$$

$$K = \kappa_{max}\kappa_{min} \quad (4.19)$$

### Characterisation of Surfaces using Curvature

The local area about a point on the surface can be characterised using the curvatures defined above. One such characterisation is as follows (See figure 4.6) :

1. **Elliptic Point** ( $H \neq 0$  and  $K > 0$ ). At these points both  $\kappa_{max}$  and  $\kappa_{min}$  are non-zero and have the same sign. The normal sections of both principal curvatures have the same curvature profile, and the local surface around it is ellipsoidal. If  $\kappa_{max} = \kappa_{min}$  the local surface becomes spherical and this special case of an elliptic point is called an *umbilic point*.
2. **Hyperbolic Point** ( $H \neq 0$  and  $K < 0$ ). Here both  $\kappa_{max}$  and  $\kappa_{min}$  are non-zero with opposite signs. The normal sections of both principal curvatures have



**Figure 4.6:** Characterisation of local surface properties using the principal curvatures. The shape of the surface around the point  $\mathbf{p}$  can be deduced from the values of the mean curvature ( $H$ ) and the gaussian curvature ( $K$ ) at the point



opposite curvature profiles, and the local surface around such a point is saddle shaped.

3. **Parabolic Point** ( $H \neq 0$  and  $K = 0$ ). In this case one of the principal curvatures is zero while the other one is non-zero. The curvature of the normal section of the zero magnitude principal curvature is therefore linear, and the local shape of the surface is that of a parabolic cylinder (a *ridge* or *trough*).
4. **Flat or Planar Point** ( $H = K = 0$ ,  $\kappa_{max} = \kappa_{min} = 0$ ) The normal sections corresponding to both principal curvatures are linear, and the local surface is flat.

The above classification of points based on the principal curvatures are also summarised by Dupin's indicatrix [65] (page 77).

### Crest Points and Crest Lines

The use of curvature in identifying characteristic points and regions on surfaces is of importance in this thesis. The beauty of using these is that the points described are intrinsic to the shape of the surface and are insensitive to orientation and choice of coordinate system. We are specifically interested in curvature extremum points and the loci of these points.

Firstly, the *lines of curvature* should be defined. These are curves on a surface whose tangents are always in the direction of the principal curvature. The equations for these lines are :

$$(MG - NF)dv^2 - (NE - LG)dudv + (FL - ME)dv^2 = 0 \quad (4.20)$$

where  $E$ ,  $F$ , etc are as defined in Equation 4.17

Along a line of curvature the value of the principal curvature changes. It is the loci of the extrema of these curves that are of interest here. Their equations are obtained from equation 4.20 by including the curvature extremum condition  $d\kappa = \kappa_u du + \kappa_v dv = 0$ ,

where  $\kappa_u$  and  $\kappa_v$  are the partial derivatives of the normal curvature with respect to the parameters  $u$  and  $v$ , giving :

$$(MG - NF)\kappa_v^2 - (NE - LG)\kappa_u\kappa_v + (FL - ME)\kappa_u^2 = 0 \quad (4.21)$$

The loci of the set of points satisfying equation 4.21 on a surface have been called *principal curvature extremum curves* in [65] and *creases* in [47] which are further classified into the loci of *weak* and *strong ridge points*. Here we take ridge points to mean points where equation 4.21 is satisfied, and a *ridge* or *ridge line* to be the loci of such points. *Crest points* are a sub-set of ridge points corresponding to the points where lines of principal maximal curvature have maximum values (in this case the maximal curvatures denote maximum of the absolute values of the principal curvatures) , i.e.

$$e_1(\mathbf{p}) = \nabla(\kappa_{max}(\mathbf{p}) \cdot \mathbf{t}_{max}^p) = 0 \quad (4.22)$$

subject to the following conditions

$$\text{If } \kappa_{max}(\mathbf{p}) < 0, \quad \nabla(e_1(\mathbf{p}) \cdot \mathbf{t}_{max}^p) > 0 \quad (4.23)$$

$$\text{If } \kappa_{max}(\mathbf{p}) > 0, \quad \nabla(e_1(\mathbf{p}) \cdot \mathbf{t}_{max}^p) < 0 \quad (4.24)$$

The quantity  $e_1$  is called the first extremality and is discussed further in section 4.6.2. The conditions 4.23 and 4.24 limit the points classified as crest points to only the maximum extremum points of lines of maximum curvature.

Other characteristic curves and points on surfaces such as sub-parabolic lines, silhouette curves, and gradient extremum points and curves can be defined, but these are not of interest here.

### 4.6.2 Obtaining Curvature-based Characteristics of 3D Surfaces

Three approaches can be taken to obtain curvature based characteristics on a surface. If the analytic equations of the surface in question are known, the curvature values and the required characteristics can be determined directly. However, in practice these are not usually available for the surfaces of structures of interest in medical image analysis. The common approach is to locally fit surface patches to the surface, from which parameters for the calculation of the curvatures and curvature extrema are obtained. A different approach suggested by Thirion et al. [117] is based on the extremality function of a surface which they obtain directly from the 3D voxel values. These three approaches are discussed further.

#### The Analytic Approach

If the analytic equations of a surface are known in parametric form, analytic expressions for the mean and gaussian curvature can be obtained, as well as expressions for calculating the ridge points and crest points using the appropriate equations from section 4.6.1. Whilst in theoretical examples the analytical approach can be used, it is not widely used in practical situations (in the literature reviewed no examples of use of the analytical approach were encountered). Reasons for lack of practical use include the fact that in practical situations it is often impossible to obtain analytic expressions for the surfaces of interest because of their complex shape. Even in cases where analytic expressions can be obtained, the parameterisation of these is complicated by the fact that most surfaces of interest in medical image analysis are closed surfaces. Obtaining one to one mappings from a plane to closed surfaces is usually impossible - although alternative parameterisations such as cylindrical and spherical ones can be used to achieve this.

### Local Surface Fitting

An alternative approach which circumvents the need for obtaining parameterised analytic equations of the surface in question is to fit surface patches locally around regions of the surface. This approach is widely used in practice with the main difference in methods being in the way surface patches are fitted and the way the problem of accounting for variations in the surface due to noise is handled.

Kwak et al. [74] used a fourth-order least squares B-spline surface-fitting method. The method was used to perform curvature analysis on the articular cartilage of the human patellofemoral joint (the joint between the kneecap (patella) and thigh bone (femur)). 3D surface points were obtained by stereophotogrammetry. As mentioned in section 2.5.2 splines provide smooth continuity between points and surfaces. The use of the fourth order spline was necessary to ensure continuous derivatives up to the third order were obtainable, which are needed for the calculation of curvature extrema. The principal curvatures were calculated from the B-spline parameters and maps of lines of curvature were obtained and these were used to locate ridges on the surface.

Shi et al. [101] perform motion tracking of the left ventricle (LV) of the heart using an algorithm based on surface shape properties. The surface of the LV was triangulated at various time points during the cardiac cycle. Curvature calculations were carried out at the vertices of the triangulated surface by fitting a local biquadratic surface patch to each vertex. Smoothing to account for distortions in the surface due to noise was achieved by including different orders of neighbouring points when fitting the surface patch to a fiducial point. The first order neighbours are vertices of triangles defining faces directly connected to the fiducial vertex. The second order neighbours are vertices directly connected to the first order neighbours and so on. The greater the order of neighbourings the coarser the curvature estimation. A multi-scale approach was adopted with the curvature estimations being refined by decreasing the order of neighbours used. The curvatures were calculated directly from the parameters of the quadratic surface patch.

Kent et al. [71] use *kriging* as their method of smoothing in the determination of ridge curves on surfaces. Kriging is based on the assumption that there is an underlying process that gives rise to errors in the data and that this can be modelled as a linear combination of the data. Ridge points on lines of curvature were identified using a modified version of equation 4.21. Ridge curves were formed by applying a zero tracing algorithm to identifying the ridge points on triangulated representations of the surfaces of interest. The method was illustrated by applying it to detect ridge curves on 3D laser range data of the human head.

The main advantage of the local surface fitting methods is that their implementation to determine curvature is straightforward. However, despite the smoothing measures taken the fitting of local surface patches may still be affected by noise. An additional disadvantage is that the continuity of the ridge curves is not guaranteed across the boundaries of surface patches.

### **The Marching Lines Algorithm**

The Marching Lines algorithm [117] developed by researchers at INRIA circumvents the need for local surface fitting by obtaining the curvature extrema directly from the voxel values of 3D images.


Iso-surfaces can be defined on 3D images as the surface boundary between regions whose intensity is higher than or equal to a given threshold, and regions whose intensity is lower than the threshold. The Marching Lines algorithm describes a method for computing with sub-voxel accuracy, the 3D curves formed at the intersection of two iso-surfaces of two voxel images.

For the purposes of the algorithm the gray levels of individual voxels are assumed to be concentrated at their bottom left corner. The voxel array is regarded as a rectangular lattice of cubic cells with the voxel values defined at the nodes (see figure 4.7). Voxels where the iso-surfaces intersect can be determined by examining the respective values in both images. The locations along the voxel edges where the


intersections occur are found by linear interpolation of predetermined possible cases. The extracted curves are shown to be continuous.

The application of the Marching Lines algorithm to the extraction of crest points is made possible using the extremality criterion defined in equation 4.22. In [117] previous work by Monga and Benayoun [84] allowing the calculation of curvature from first and second derivatives of image values was extended. In particular it was shown that the principal curvatures and associated principal directions of an iso-surface defined on a 3D voxel image can be computed directly from the first and second derivatives of image voxel values. These in turn allow the calculation of the extremality values at each voxel. The Marching Lines algorithm can then be applied to the intersection of the iso-surface of the zero level set of the extremality of the image and a suitable threshold value of the actual voxel values. The curves at the intersection of these two iso-surfaces are the crest lines on the iso-surface of the image.


For binary segmented data convolution with a gaussian kernel of specified width is used to create an image with varying voxel values (between 0 and 1). This serves the dual role of allowing the calculation of derivative information at voxels and acting as a smoothing factor - the greater the width of the kernel the smaller the change between neighbouring voxel values, hence the greater the smoothing. For the binary segmentations of the lateral ventricles used in this study, the width of the kernel was 1.5 voxel units, and the threshold for determining the image iso-surface was 0.5. A summary of the application of the Marching Lines Algorithm to the extraction of crest lines is given in figure 4.8. Figure 4.9 shows the crest lines obtained from four of the ventricles used in the study.



(a) Image as an array of voxels



(b) Image as a rectangular lattice with voxel values concentrated at nodes



 (c) A  $2 \times 2 \times 2$  neighbourhood of voxels forms a cubic cell in the lattice. The nodes and the edges of the cell are labelled

**Figure 4.7:** The Marching Lines Algorithm works by examining voxel values in a  $2 \times 2 \times 2$  neighbourhood. Subfigure (a) shows a 3D image as an array of voxels. (b) shows the voxel array as a rectangular lattice and (c) shows a neighbourhood of the voxel array as a cell in the lattice.

1. The input to the algorithm is the binary segmented voxel data. *On* voxels have value 1 and are within the lateral ventricles. *off* voxels have value 0 and are not ventricle.
2. Two images are derived from the input.
  - One is a smoothed version of the input data. This is obtained by convolution of the input data with a gaussian kernel to give an image with values between 0 and 1. This is converted to a binary image by thresholding at the 0.5 level set. This image is referred to as the *isosurface* image
  - The other is referred to as the *extremality* image (see equation 4.22). This is obtained by repeatedly convolving the input image with derivative of gaussian filters to obtain derivatives to the required order.
3. The crest lines are extracted.
  - (a) *March* from cell to cell in the isosurface image (see figure 4.7(c) for description of what a cell is)
  - (b) For each cell, if all the nodes do NOT have the same value then the isosurface passes through this cell. Otherwise, the isosurface does not pass through this cell - *march* on to next cell.
  - (c) Check the signs of the values at the corresponding nodes in the extremality image. If the signs at all eight nodes are NOT the same the isosurface and the zero level-set surface of the extremality function may intersect. Otherwise, no intersection - *march* on to next cell.
  - (d) Use the values of the nodes of the cell in the extremality image and the isosurface image to calculate where the isosurface intersects the zero level-set surface of the extremality function. See [115] and [117] for details.
  - (e) Follow the crest line to the next cell and repeat from 3(b).

**Figure 4.8:** Summary of the steps in using the Marching Lines Algorithm to extract crest lines





**Figure 4.9:** Four ventricular surfaces colour-mapped according to mean curvature with crest lines (after filtering out noisy ones) superimposed on them

### 4.6.3 The Use of Crest Points and Crest Lines as Curvature-Based Landmarks on the Ventricular Surface

Here justification for the use of crest points and crest lines as curvature-based landmarks in the construction of 3D PDMs of the lateral ventricle is presented. Curvature has already been stated to be an important property of surfaces. In section 4.6.1 it was shown that the mean and Gaussain curvatures of a point allow the characterisation of the local shape of the surface around that point. Koenderick in his treatise on 3D shape [72] stated (page 319) that “the *local* shape on a surface depends only on the values of the principal curvatures”. However, he also warned that the mean and Gaussian curvatures together do not completely determine the *global* shape of a surface (page 253) “Two surfaces may possess equal Gaussian and average curvatures at corresponding points but still look quite different! Curvature isn’t all there is to surfaces”.

The use of crest points here pays heed to Koenderick’s warning. Crest lines associated with crest points are robust in the sense that if the surface on which they are defined is deformed they deform accordingly [86]. Furthermore, the crest points are specifically used to establish correspondence and are not used to measure shape differences. It is the analysis of the coordinates of corresponding points that is used in the investigation of shape differences. Ridge curves (of which crest lines are a subset) found repeatedly on biological surfaces have been argued to be equivalent to anatomically meaningful landmarks [19],[115]. Additionally, Thirion and Gourdon [116] have shown direct correspondence between crest lines and anatomical features in medical images.

Andresen et. al. [3] use crest lines to obtain correspondences when building a 3D PDM of the mandibles of six subjects with Apert’s syndrome to model and predict mandible growth. The crest points and crest lines on the surface of the mandible of each subject were obtained using the matching lines algorithm. One mandible was chosen as a template and its surface was triangulated. Correspondence between the training set was established in a similar manner to that described in section 4.5. How-

ever, Andresen et. al. used a closest point method (we use a transportation-based method described in section 4.7) in establishing point-to-point correspondences, and a *geometry-constrained diffusion* algorithm to re-project landmark points onto the surface of each member of the training set after deforming the template to approximate the shape of that member.

Of greater importance here is the fact that crest points have been utilised in the creation of a 3D SSM of the ventricles of the brain - albeit by a manual method. Buckley et al. [24] and Dean et al. [41] manually subdivided the surface of the entire ventricular system into surface patches. The surface subdivision was based on geodesic lines between salient crest points on the surface. The forty-eight crest points used were known to be reliably detectable on ventricular surfaces. They were manually specified by clicking a mouse on the ventricular surfaces that had been colour-mapped according to magnitude of curvature. Average positions of the landmarks over the sets of ventricles analysed were created using thin plate spline warping. This method was applied to investigate shape differences between the lateral ventricles of schizophrenic and control subjects. In this work crest lines were shown to be reliable consistent landmarks on the ventricular surface. This proposition is backed up by figure 4.9 and qualitative examination of the crest lines of all the ventricles used in the study. These show that even though the shape of the ventricles varies quite markedly, the crest lines on their surface seem consistent. Crest lines have also been used in creating average morphometric atlases of the human skull [108].

A further reason for using the curvature-based approach is the fact that the application of other methods of automatically constructing 3D PDMs to the construction of PDMs of the lateral ventricles often results in failure. In the approach of Davies et al. [38] the method of reparameterisation used cannot be applied to the whole of the lateral ventricles because its highly curved shape cannot be projected easily onto a sphere. The method of Brett et al. [22] is also inapplicable to the whole of the lateral ventricle as their method is restricted to disc-like surfaces. Finally, the use of crest points and crest lines to define correspondence as we have done here directly

exploits a physically significant structure of the ventricles. This is in contrast to other automatic methods that rely on arbitrary triangulations or other reparameterisations.


## 4.7 Generating Point-to-Point Correspondences

Having determined the crest points and crest lines of the set of ventricles in our training set, they are used to define a sparse correspondence across the training set in a pair-wise manner. In this section we summarise the automatic method for establishing correspondence between a pair of ventricles  $A$  and  $B$ . We compare two methods of point-to-point matching. One is based on matching locally close points as suggested by [108] and in the other we present the point matching problem as a transportation problem and apply an efficient algorithm to solve it.

Section 4.7.1 gives details of our implementation of the closest point method suggested in [108]. Section 4.7.2 describes how the point matching is posed as a transportation problem, and the application of an efficient algorithm to solve it. Section 4.7.3 focuses on the quantitative comparison of the transportation and closest point methods. The comparison is done at three levels. Firstly, the point-to-point correspondences produced by the two matching methods are compared. Secondly, the fidelity of the corresponding training sets generated by the two methods to the actual surfaces are compared and lastly, the PDMs generated by both methods are compared.

### 4.7.1 A Closest Point Approach to Point Matching

The point matching method described by Subsol et. al. [108] is an adaptation of the Iterative Closest Point (ICP) algorithm of Besl and McKay [9]. As a preliminary step, the origin of the coordinate system for the crest points in each ventricle being matched is translated to the centroid of the crest points. In matching the pair of ventricles  $A$  and  $B$ , the matching can be done in two “directions”,  $A \rightarrow B$  and  $B \rightarrow A$ . As



**Figure 4.10:** Illustration of the result of point-to-point matching. (a) Point-to-point matching between lower and upper lines using closest points results in cross matches (heavy dotted lines) and multiple matches (light dotted lines). (b) removal of multiple and cross matches results in a single correspondence for each point. Some correspondences may be missed after this post-processing

will be discussed in section 4.7.3, the results of matching in both directions are not generally the same. Here we consider matching in the direction  $A \rightarrow B$ .

The matching is an iterative procedure. At each iteration, point-to-point correspondences are obtained, then registrations are applied to  $A$  to make it more similar to  $B$ . The final set of corresponding points are used to obtain parameters for a spline-based warp allowing a set of predetermined vertices on the surface of  $A$  to be mapped to corresponding points on the surface of  $B$ .

In the object centred coordinates of  $A$  and  $B$ , initial point-to-point correspondences are obtained by finding for each crest point of  $A$  the closest crest point in  $B$  in terms of Euclidean distance. This initial set of correspondences may result in multiple matches (where more than one point in  $A$  is matched with the same point in  $B$ ), and cross-matches where the lines joining matched points “cross-over” as shown in the example of figure 4.10.

To obtain a valid set of correspondences, each crest point in  $A$  is constrained to have only one matching point in  $B$ . For each crest line in  $A$  the matching crest line in  $B$  is found. The crest lines are used to impose an ordering on the crest points to prevent cross-matches. These requirements are encapsulated in two constraints. An *injectivity* constraint that each crest point on  $A$  can be matched to at most one

crest point on  $B$ , and vice versa. And a *monotonicity* constraint that maintains the same order of matched crest points along crest lines  $A$  and  $B$ . We used a modified version of the heuristic algorithm proposed in [108] to enforce these constraints. This is shown in figure 4.11

The main modification is that we do not restrict each crest line in  $A$  to be matched to only one crest line in  $B$ . Instead, we allow *partial line matching* where a contiguous part of any line in  $A$  can be matched to a contiguous part of any line in  $B$  and vice versa. We found this necessary because our approach to filtering “spurious” crest lines differed from that of [108]. They obtained crest lines using a low and a high width Gaussian smoothing kernel, and then filtered the spurious crest lines obtained with the less smoothed low width kernel using those obtained with the more smoothed high width kernel. However, application of that method to our images resulted in some salient crest lines not being obtained. We therefore used only a low width smoothing kernel, and applied a threshold based on absolute value of curvature to filter out the spurious crest lines as suggested in [117].

For the first 10 iterations of the point-to-point matching, rigid body transformations are performed on  $A$  to minimise the Euclidean distance between pairs of matched points given by the point matching algorithm. For the next 10 iterations 12 d.o.f affine transformations are applied. For the last 10 iterations spline-based warps are used. The spline-based warps were implemented as described in chapter 14 of [122].

### 4.7.2 A Transportation Approach to Point Matching

The previous section described the closest point method of obtaining crest point correspondences on the surfaces of a pair of ventricles. This method of finding correspondences is not necessarily optimal. Matching points with closest Euclidean distance can result in rejection of matches which are good, but not locally “best”. This may result in solutions which are not globally optimal across all points. In this section, we demonstrate the presentation of point matching as a transportation problem, and

Given crest points and crest lines for a pair of ventricles  $A$  and  $B$  to perform point-to-point matching of crest points in  $A$  with crest points in  $B$  ( $A \rightarrow B$ ):

1. For each crest point in  $A$  find the closest crest point in  $B$
2. For each crest line in  $A$ 
  - (a) Enforce monotonicity and injectivity constraints for its point-to-point matches (as described in [108])
  - (b) Find full or *partial matching lines* in  $B$  using the voting mechanism described in [108]
  - (c) Obtain the closest point matches for the crest points in each matched section of the crest line in  $A$  and  $B$  anew and enforce the monotonicity and injectivity constraints
3. On each iteration perform rigid, affine or spline registrations as appropriate
4. The resulting matches after the final spline registration are the output of the point matching process

**Figure 4.11:** Outline of algorithm for point-to-point matching

discuss the application of an efficient method of obtaining solutions.

### Point-to-Point Matching as a Transportation Problem

A transportation problem is a class of linear programming problem in operational research [60]. It derives its name from solving the problem of distributing goods from a set of sources with given production to a set of outlets with known demands along paths of known costs. The requirement is to provide the overall lowest-cost distribution of all goods produced at the sources to the outlets, subject to the constraints that the total production at all sources equals the total demand at all outlets, and the quantity transported between source and outlet is non-negative. The transportation paradigm has been applied in some medical image analysis problems, for example chromosome matching [125] and mammography [12]. Belongie et al [8] have applied the transportation algorithm to point-to-point matching and generating correspondences. However, we arrived at our approach independently, using a different algorithm to solve the resulting *assignment problem* (see section 4.7.2). Furthermore our application is in 3D and can be used in 2D whereas the method of Belongie et

al was applied to 2D problems. The *Softassign* Procrustes Matching Algorithm [95] also includes an *assignment* step in establishing point-to-point correspondences.

To illustrate how point matching is cast as a transportation problem, consider the matching of crest points of two ventricles, ventricle  $A$  with  $m$  crest points, and ventricle  $B$  with  $n$  crest points. We assume for the moment that the number of crest points in both are equal (i.e.  $m = n$ ). The crest points from  $A$  are taken as sources each with unit production and those of  $B$  as outlets each with unit demand. We seek to transport a quantity (create a match) from each source in  $A$  to each outlet in  $B$ . The cost of each transportation (match) is simply the Euclidean distance between  $A$  and  $B$ . We seek to minimise this cost over all matches. This gives the advantage of permitting matches that are not locally closest, potentially allowing more matches to be made. Point matching can therefore be written as the following transportation problem: Given

$$\sum_{i=1}^m a_i = \sum_{j=1}^n b_j \quad (4.25)$$

where  $\forall i, j; a_i$  (availability at source) =  $b_j$  (demand at outlet) = 1.

If  $x_{ij}$  is the quantity transported from source  $i$  to outlet  $j$  (For the purposes of point matching, the quantity transported is always 1 when a match is created, or 0 otherwise).

Find  $x_{ij} \geq 0$  which satisfy

$$\sum_{j=1}^n x_{ij} = a_i (= 1), i = 1, \dots, m \quad (4.26)$$

and

$$\sum_{i=1}^m x_{ij} = b_j (= 1), j = 1, \dots, n \quad (4.27)$$

and which minimise

$$z = \sum_i^m \sum_j^n D_{ij} x_{ij} \quad (4.28)$$



Where  $\mathbf{D}$  is a  $m \times n$  distance matrix with element  $D_{ij}$  of the  $i$ th row and  $j$ th column given by,

$$100 \times d_{ij} \tag{4.29}$$

$d_{ij}$  being the square of Euclidean distance between crest point  $i$  on  $A$  and crest point  $j$  on  $B$ .

Although the number of crest points on any two ventricles are of the same order, in general they are not equal (i.e. for any two ventricles  $A$  and  $B$ , in general,  $m \neq n$ ). In such cases, dummy sources or dummy outlets are introduced to allow expression in the form of equations 4.25 to 4.28. If  $n > m$ , we introduce  $n - m$  dummy sources, each with unit availability, and amend equations 4.25 to 4.28 as necessary. On the other hand, if  $m > n$ , a dummy outlet with demand equal to the difference between  $m$  and  $n$  is introduced. The costs for transporting to or from dummy sources and outlets are set to an arbitrarily high value i.e. transportation to/from dummy sources and outlets capture the worst matches. The matched points are the  $(i, j)$  pairs for which  $x_{ij} \neq 0$ , and  $i$  or  $j$  is not a dummy source or outlet. In section 4.7.1 we noted the need to apply injectivity and monotonicity post-processing. Due to the constraints stated in equations 4.26 and 4.27, the transportation method intrinsically enforces the injectivity criterion (but not the monotonicity one).

### **An Efficient Solution to the Transportation Problem**

The classical method for solving transportation problems is the *Stepping Stone algorithm* [60]. However, the fact that the quantity transported in the present case is always of unit value makes it similar to (but not quite) a special case of transportation problems termed *assignment problems*. This special case has unit demand at each outlet and unit availability at each source. The formulation of the point-to-point matching problem given above is not quite an assignment problem, because in the case where a dummy outlet is used, its demand is equal to the excess in the number of sources, and not necessarily a unit demand. The classical method of solving the assignment problem is the Hungarian method [96]. However, within the operational

research field there are a variety of more efficient solutions to the assignment problem. One such solution is that of Achatz et al. [2]. Their algorithm allows the assignment of  $m$  unit source nodes to  $n$  outlets in time  $O(mn^2)$ , where  $m$  is the number of sources, and  $n$  is the number of outlets - allowing its application to the formulation of the point matching problem given above.

As the number of sources for the point matching problem are of the order  $10^3$  (relatively large for transportation problems), the computational efficiency and memory requirements of the algorithm used to find a solution are important. The NAG-C implementation of the Stepping Stone Algorithm did not produce an optimal solution for point matching of 1000 sources and 1000 outlets after 24hrs. However, an implementation of the algorithm of [2] produced an optimal solution for the same problem in about 5 seconds. Both functions had been coded in C and were called using the Mex functionality of Matlab 5.3.1 running on a standard PC with a 500Mhz Pentium processor and 256 MB RAM. The running time of the Achatz et al. algorithm is comparable to that of the closest point method. To match 1000 pairs of points it takes 2.5 seconds with the closest point method. The transportation method takes 4.5 seconds for the matching, and an additional 2 seconds are needed for calculation of the cost function.

### 4.7.3 Quantitative Comparison of the Transportation and Closest Point Methods

Using the training set of all 138 ventricles (69 right ventricles of control and schizophrenic subjects, and their left counterparts reflected to give the same pose as the right ones), two PDMs were constructed as outlined in section 4.5 without the re-projection and normalisation steps. For one PDM, the closest point approach was used to match crest points, whilst in the other the transportation method was used.

With reference to figure 4.3, quantitative comparison of both methods is carried out

at steps 3, 4 and 7. At step 3, the results of the point-to-point correspondences of both methods are compared. The four criteria used for this comparison are detailed in the subsection below. In step 4, the corresponding training sets produced by both methods are compared using the 3D distance transform metric described in section 4.4.4. Lastly, the resulting PDMs are compared. The criteria of section 4.4 are used for this comparison. In all results in the following subsections (except where stated), the data plotted are the mean of the values for all 138 ventricles. In all cases the measure of error used is the standard error of the mean obtained in the usual manner (i.e.  $\frac{\sigma}{\sqrt{N}}$ , where  $\sigma$  is the sample standard deviation, and  $N$  is the number of samples).

### **Comparing the Point-to-Point Correspondences Produced by the Transportation Method and the Closest Point Method**

The point-to-point correspondences obtained for the crest points are used to drive a spline-based warp. The spline-based warp is an approximating function that maps each 3D coordinate of the template onto the fiducial ventricle. The greater the number of accurately determined corresponding points, the better the approximation function. As the process of obtaining the point-to-point correspondences is iterative, it is also desirable to ascertain the level of convergence that had been attained at the end of the matching process.

In this section the point-to-point correspondences obtained by the transportation method and the closest point method are quantitatively compared. The four measures used for this comparison are: Number of point matches at each iteration, Number of matches changed at each iteration, Mean distance between matched points, and the symmetry of the matched points. Each measure is described more fully below, and their results for the present experiment shown.

#### **Number of Matched Points**

During point matching, the point-to-point correspondences are stored in a vector. The index of the elements of the vector represent points on one ventricle, and the


value of the element is the corresponding point on the other ventricle. At each iteration of the point matching process, the number of point-to-point matches after monotonicity and injectivity constraints have been enforced is measured. This is achieved by simply obtaining the number of non-zero elements in the vector defining the point-to-point correspondences.

Figure 4.12(a) shows plots of the mean number of matched points at each iteration for both methods. The transportation approach gives almost twice the number of point matches than the closest point method on each iteration.


### Changes in Matched Points Between Iterations

Using the vector defining point-to-point correspondences, the number of changes between matched points at each iteration can be obtained by comparing the elements of the vector for consecutive iterations. This gives a measure of the stability of the matching method as the iterations progress. As the template becomes more similar to the ventricle to which it is being matched, the number of changes in the matched points should decrease.


Figure 4.12(b) shows the mean percentage of changes in matched points from one iteration to the next. The plots for both methods show the same pattern - an initially high proportion of points changing correspondences, which reduces in a somewhat exponential manner as the number of iterations for a particular transformation type increases. When the transformation type changes, there is a discontinuity (also evident to a lesser extent in figure 4.12(a) for the number of points and figure 4.12(c) for the mean distance between matched points). The proportion of points in which the match changes between iterations, stabilises more quickly using the transportation method than the closest point method during rigid and affine transformations. However, the rate of reduction in point changes is similar during spline transformations.



(a) Total number of matched points



(b) % number of changes in matched points



(c) Mean distance between matched points

**Figure 4.12:** Evolution of the number of matched points after injectivity and monotonicity enforced (a), the percentage number of changes in matched points (b), and the mean distance between matched points (c), with iterations for the transportation method (blue) and the closest point method (red) during the point-to-point matching process. The discontinuities in the curves (most obvious in (b)) occur where the registration changes from rigid to affine and then to spline-warping.

### Mean Distance Between Corresponding Points

At each iteration, the mean Euclidean distance between matched points can be obtained in a straightforward manner using the correspondences between points at that iteration, and the coordinates of the points of the ventricles being matched. As the crest lines of the template are deformed to become more similar to the ventricle it is being matched to, it would be expected that the distance between the matched points decreases. The mean distance between corresponding points serves as a measure of the convergence of the iterations.

Figure 4.12(c) shows that the mean distance between matched points evolves in the same pattern for both methods. However, its value is smaller for the closest point method at each iteration. This is, of course, to be expected, since the transportation method does not exclude matches that are not locally shortest.

### Symmetry of Matches

Although the matching algorithms for both matching methods do not require matches to be symmetric (that ventricle  $A$  matched to ventricle  $B$ ,  $A \rightarrow B$ , gives the same result as  $B$  matched to  $A$ ,  $B \rightarrow A$ ), we would have greater confidence in symmetric correspondences. This was briefly touched upon in sections 4.7.1 and 4.7.2.

To quantitatively measure the symmetry of matches, an index of symmetry,  $I_s$ , is calculated using matched points lying on symmetrically matched partial-lines. Partial-lines were described in section 4.7.1, and symmetrically matched partial-lines are overlapping partial-lines from matches  $A \rightarrow B$  and  $B \rightarrow A$ .  $I_s$  is the ratio of the number of point-to-point matches in overlapping partial lines for  $A \rightarrow B$  and  $B \rightarrow A$ , to the total number of point-to-point matches in both directions, i.e.

$$I_s = \frac{\text{number of matches on symmetrically matched partial - lines}}{(\text{number of matches in } A \rightarrow B + \text{number of matches in } B \rightarrow A)} \quad (4.30)$$


For the transportation method, if the number of crest points on both  $A$  and  $B$  are

equal (i.e.  $m = n$ ), then the  $x_{ij}$  given as a solution to the transportation problem are the same for both  $A \rightarrow B$  and  $B \rightarrow A$ . This is because firstly, the constraints given in equations 4.25 to 4.27 are not changed if  $a$  and  $b$  are interchanged. And secondly, for  $m = n$ ,  $\mathbf{D}$  is a square matrix, and  $D_{ij}^{A \rightarrow B} = D_{ji}^{B \rightarrow A} \forall i, j$  (i.e.  $\mathbf{D}^{A \rightarrow B} = (\mathbf{D}^{B \rightarrow A})^T$ ). Hence, the constraints and the cost function are not changed for this case, and the optimal solution is symmetric.

Figures 4.13 and 4.14 illustrate the fact that the transportation method is more symmetric than the closest point method using a pair of synthesised lines. Points were generated along each line and matched with both methods. The initial results of both methods have to be subjected to the injectivity and monotonicity constraints. In the case where  $m = n$  as for the synthetic lines, a symmetric match is expected for the transportation method. Although the initial matches are symmetric, after enforcing the monotonicity constraints (the transportation solution ensures injectivity) some matches may be deleted. Therefore, the output for the transportation method may not be symmetric even though the optimal solution was symmetric. However, as the transportation method gives a globally optimal solution, we would nevertheless expect the matches produced to be more symmetric than those of the closest point method.

Notice in subfigures (c) and (d) of Figure 4.13 that there is one point on each line that does not match “naturally” to the other line. The transportation method forces a match between these points. Despite the high cost for this particular match, the global match cost is minimised by the intuitively correct matches of the remaining points. The “incorrect” match is rejected by the application of the monotonicity constraint.

Figure 4.15 shows a line plot of the mean index of symmetry for the members of the evaluation set matched to the chosen target for both methods. Along the  $x$ -axis each data point represents a particular ventricle, there is no continuity between data points - the use of lines is purely for ease of visualisation. The mean index of symmetry



**Figure 4.13:** Initial results (before enforcement of monotonicity and injectivity constraints) for the closest point method in both directions (a) and (b). (c) and (d) are the equivalent results for the matches returned by the transportation method before enforcement of monotonicity and injectivity constraints. The initial results of the transportation method are always symmetric for the case where the number of points in  $A$  and  $B$  are the same, whereas those of the closest point method are not generally symmetric


for the transportation method was  $0.79 \pm 0.05$  (range 0.63-0.91), whilst that of the closest point method was  $0.70 \pm 0.11$  (range 0.26-0.89).

Figures 4.16(a) and (b) show the matches in both directions for the case where the closest point index of symmetry was lowest ( $I_s = 0.26$ ), and the corresponding set of matches for the transportation method ( $I_s = 0.78$ ) are shown in figures 4.16 (c) and (d). Figures 4.17 (a) and (b) show the matches in both directions for the case where the closest point index of symmetry was highest ( $I_s = 0.89$ ) and figures 4.17 (c) and (d) are the corresponding set of matches for the transportation method ( $I_s = 0.85$ ).

### Comparing the Corresponding Training Set Produced by the Transportation Method and the Closest Point Method

Each member of the corresponding training set is composed of the same number of vertices as the template. The similarity of each to its original version was determined by finding the rigid registrations that minimised the distance transform based objective function. The mean values are given in table 4.1. The results show that on average the transportation method gives a corresponding training set that approximates the original training set better than the corresponding training set of the





**Figure 4.14:** (a) and (b) are the results of the closest point method after enforcement of monotonicity and injectivity constraints. (c) and (d) are the results after enforcement of the constraints for the transportation method (note that the transportation method gives an injective map for the case where the number of points on  $A$  and  $B$  are the same). Again the transportation method is symmetric whereas the closest point method is not.

closest point method. However, the difference is not very significant ( $p = 0.0059$ ).

	Transportation	Closest Point
Mean value of Distance Transform metric	$1.01 \pm 0.16$	$1.08 \pm 0.28$
Range	0.71 - 1.74	0.69 - 2.60


**Table 4.1:** Value of distance transform metric for comparison of surfaces of the corresponding training sets produced by transportation and closest point methods with their respective original surfaces. The values shown are in voxel units (1 voxel = 0.78mm). A  $t$ -test showed that the difference between both groups was not very significant ( $p = 0.0059$ ).

### Comparison of Point Distribution Models Produced by the Transportation Method and by the Closest Point Method


The PDMs created by the two methods were compared using the criteria described in section 4.4. The results are presented below.

#### Generalisation Ability

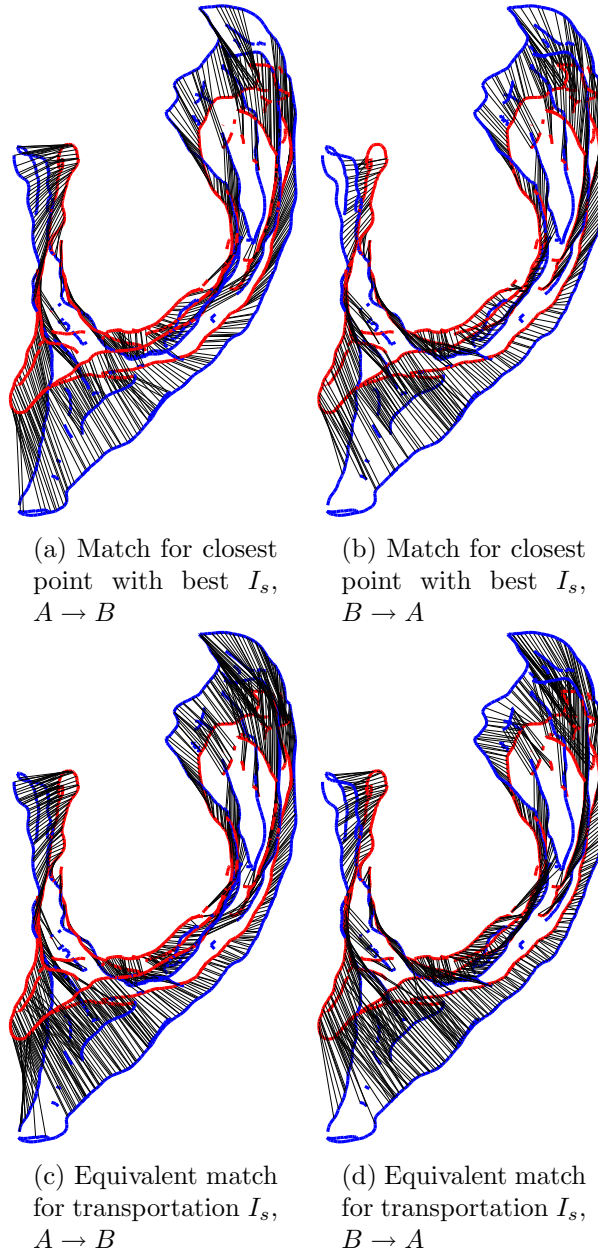
In section 4.4.1 we stated that the error-of-fit to measure the generalisation ability of a model can be obtained in two ways. One method (the *de facto* one) uses the



**Figure 4.15:** Plot of the index of symmetry for each ventricle for both the transportation and the closest point methods. The plots show that the closest point method is generally less symmetric, and its index of symmetry is more variable



**Figure 4.16:** The figure above shows point-to-point correspondences (black lines) obtained between crest points of a pair of ventricles giving the worst measure of symmetry for the closest point method ( $I_s = 0.26$ ). The red lines are the crest lines of one ventricle and the blue ones the crest lines of the other. (a) shows the matches in one direction for the closest point method, and (b) shows the matching in the other direction. (c) and (d) are the equivalent results for the same pair of ventricles using the transportation method. The symmetry value for the transportation method was 0.78, and the figure shows that the matches for the transportation method are more symmetric and evenly distributed.



**Figure 4.17:** The figure above shows point-to-point correspondences (black lines) obtained between crest points of a pair of ventricles giving the best measure of symmetry for the closest point method ( $I_s = 0.89$ ). The red lines are the crest lines of one ventricle and the blue ones the crest lines of the other. (a) shows the matches in one direction for the closest point method, and (b) shows the matching in the other direction. (c) and (d) are the equivalent results for the same pair of ventricles using the transportation method. The symmetry value for the transportation method was 0.85, and the figure shows that qualitatively the matches returned are similar.


approximations to the training set in which correspondence has been established (the corresponding training set), whilst the other (a novel product of the present work) uses the binary segmentations of the training set. Results for both methods are presented in this subsection and the subsection on specificity.

Figures 4.18(a) and (b) show plots of the generalisation ability of the PDMs constructed using the closest point and the transportation methods. Subfigure 4.18(a) uses the difference in Euclidean distance between the model fit and the shape vector of the left out example from the corresponding training set. Subfigure 4.18(b) uses the distance transform metric obtained from the model fit and the binary voxel-based image of the left out example.


The shapes of the plots in both cases are the same, with the error-of-fit decreasing as the number of modes used in generating the model fit increases. The comparison with respect to the corresponding training set suggests that the closest point method gives a model with greater generalisation ability than the transportation method. However, the comparison with respect to the binary voxel data of the left out example suggests that the model created by the transportation method gives better generalisation ability. It is also interesting to observe that the generalisation error-of-fit measured with respect to the corresponding training set decreases to zero as the number of modes increases, whereas the minimum of the distance transform metric is 1.3 voxel units. We believe these results are important as they suggest that the *de facto* method of measuring generalisation ability, which is a direct extension of that of the 2D case, is not appropriate in the 3D case.

### Specificity

25 random shapes were generated to investigate the specificity of the PDMs created by the transportation and closest point methods. The specificity measure and the generation of the random shapes were described section 4.4.2. As in the case of generalisation ability the specificity is measured both with respect to the corresponding training set and with respect to the binary voxel data of the training set.



(a) Generalisation ability with respect to the corresponding training set



(b) Generalisation ability with respect to actual surfaces of training set using the distance transform metric

**Figure 4.18:** Comparison of the Generalisation Ability of the PDMs constructed by the closest point and transportation methods. The comparison is done as a function of the number of modes. Subfigure (a) shows the generalisation ability measured as the Euclidean distance between corresponding points of the generated instance and its surface approximation in the corresponding training set. Subfigure (b) shows the generalisation ability measured by comparison of the instance generated by the model to its binary segmentation using the distance transform metric.


Figure 4.19(a) shows the specificity as a function of the number of modes for the two PDMs as measured by comparison of landmark points on each random shape with the landmark points on the members of the corresponding training set. Figure 4.19(b) shows the specificity measured by comparison of the surface of each random shape with the most similar surface of the binary representations of the ventricles using the distance transform metric. Each data point in the graphs is the mean of the values for all 25 random shapes, and the error is given as the standard error of the mean.

The shapes of the plots for both measures of specificity are similar rising in an exponential manner. For both measures, corresponding values of the error-of-fit for the PDM produced by the closest point method are lower than those of the PDM produced by the transportation method. The difference in error-of-fit for both methods is very significant for comparison using the corresponding training set than comparison using the distance transform metric. This suggests that the PDM produced by the closest point method is more specific than that produced by the transportation method.


### Compactness

Section 4.4.3 discussed how PDMs are compared on the basis of compactness. Figure 4.20 shows the compactness of the PDMs created using both point matching methods. In both cases the amount of the total variation encompassed by a given number of modes is similar. For both PDMs over 95% of the total variation was contained in the first 30 modes.

Figure 4.21 shows the amount of variation encoded by both PDMs as a function of the number of modes. The difference between Figures 4.20 and 4.21 is that the data points in the former have been normalised by the total variance ( $\lambda_T$  of Equation 4.12) whilst the later is simply a cumulative sum of the eigenvalues of the model. For each mode, the PDM created by the transportation method had over twice as much variation as that created by the closest point method.




(a) Specificity measured with respect to the corresponding training set




(b) Specificity measured with respect to actual surfaces of training set using the distance transform metric

**Figure 4.19:** Comparison of the Specificity of the PDMs constructed by the closest point and transportation methods. The comparison is done as a function of the number of modes. Subfigure (a) shows the specificity measured as the Euclidean distance between corresponding points of the randomly generated instances and those of the most similar member from the corresponding training set. Subfigure (b) shows the specificity measured in a similar manner using the distance metric to compare the randomly generated instance to the binary segmentations of the members of the training set.





**Figure 4.20:** Compactness of the PDMs created by the transportation and closest point methods as a function of number of modes. This is a plot of the cumulative variance normalised by the total variance ( $\lambda_T$  of Equation 4.12)



**Figure 4.21:** Variance captured by the the PDMs created by the transportation and closest point methods as a function of number of modes. This is simply a plot of the cumulative variance WITHOUT normalised by the total variance

## Discussion of Results of Comparisons

### *Point-to-Point Correspondences*

The results of the quantitative comparison of point-to-point correspondences produced by both matching methods suggest that the transportation method is better than the closest point method. However, figure 4.15 showed that in 18 out of 137 cases the symmetry of the closest point method was better than that of the transportation method. This may occur in cases where the shapes of the ventricles being compared are very similar. However, when the ventricles are very different (as quite a number of them are) the transportation method gives better results. In particular, the fact that the average number of point correspondences returned by the transportation method is almost twice that returned by the closest point method, and that the symmetry of the former is better in 119 out of 137 cases attest to the overall superiority of the transportation method to the closest point method.

### *Corresponding Training Set*

The translation of the better correspondences obtained by the transportation method to giving better spline-based warps is demonstrated by the results of the comparison of the corresponding training sets produced by both methods (Table 4.1). In this the similarity of the surface obtained by deforming the template to approximate the surface of each member of the training set was measured using the distance transform metric.

The mean value of the distance transform metric over the whole training set was lower for the transportation method. However, the difference was not very significant ( $p = 0.0059$ ). Furthermore, the minimum value of the distance transform metric for the closest point method was lower than that of the transportation method (0.69 cf. 0.71). We believe this is because when two ventricles of similar shape are being matched, the closest point and transportation methods give comparable results. However, when the shapes of the ventricles being matched are quite different, the transportation method out-performs the closest point method. This is reflected in the fact that the range of the values of the distance transform metric is smaller for the transportation

method (0.71-1.74) than for the closest point method (0.69-2.60). Comparison of the values of the distance transform metric for individual ventricles showed that for 82 ventricles out of the 138, the transportation method gave a better approximation to the original surface than the closest point method.

The fact that the corresponding training sets produced by both methods are different also suggests that its use as the basis of comparing the resulting PDMs, as is the convention in the 2D case, is not adequate in the 3D case.

### *PDMs*

The interpretation of the results obtained for the comparison of the PDMs produced by both methods requires some care. Firstly, the results of the compactness (Figure 4.20) and the amount of variation encoded by each PDM (Figure 4.21) are considered. The compactness results are similar for both PDMs. However, the PDM constructed by the closest point method may have a slightly greater degree of compactness than the transportation PDM at lower modes of variation.

The actual amount of variation encoded by the two PDMs is significantly different. The transportation PDM encodes over twice as much variation as that of the closest point method. According to Davies [38] this can be taken to mean that the transportation method is less compact. However, based on the results discussed in the previous subsection (those of the corresponding training set), we believe that the corresponding training set used to build the PDM constructed by the transportation method is more representative of the original training set than that used to build the PDM of the closest point method. To conclude, the compactness results suggest that the closest point and transportation PDMs have similar compactness. However, the transportation PDM encompasses more of the variability actually present in the data.

Both the generalisation ability and the specificity were measured in the conventional manner with respect to the corresponding training set and then in a novel manner with respect to the binary segmentations from MR data using the 3D distance trans-

form metric. Measurement of the generalisation ability in the conventional manner suggested that the closest point method gave a PDM that had better generalisation ability than that of the transportation method (Figure 4.18(a)). However, using the distance transform metric, the results were reversed (Figure 4.18(b)). If our assertion in section 4.4.1, that the comparison relative to the actual ventricular surfaces, as opposed to comparison with surfaces of the corresponding training set (which has been shown to be different for the closest point and transportation methods) is correct, then it can be concluded that the PDM created by the transportation method has better generalisation ability on the basis of the results obtained by the distance transform metric.

The measures of specificity with respect to the corresponding training set and with respect to the actual ventricular surfaces using the distance transform metric showed a similar trend. In both cases the PDM constructed by the closest point method had greater specificity. However, the significance of the difference in specificity was much greater when measured with respect to the corresponding training set. The trend shown in the plots of figure 4.19 are counter-intuitive. This is because it may be expected that as the number of modes used to construct a random example increases the accuracy of the surface produced should increase, and the error-of-fit with the nearest member from the training set should decrease. However, the shape of the graphs show that as the number of modes increases the error-of-fit also increases - i.e. the specificity decreases! This makes sense when it is considered from the point of view of the allowed shape domain. Each mode of variation is an axis of the hyperellipsoid approximating this domain. As the number of modes of variation are increased, the volume of this hyperellipsoid also increases. Each mode of variation therefore increases the distance between the shapes in this domain - the volume they are contained in is expanding and carrying the shapes along.

#### *Importance of Distance Transform metric*

The importance of these results is the demonstration that the evaluation of the generalisation ability by the conventional method and by the distance transform method

give conflicting results. We believe that in the comparison of PDMs the generalisation ability and specificity should be compared with respect to the actual members of the training set and not to their representations in the corresponding training set. This is justified by the results of the generalisation ability comparisons. However, in the case of specificity, the distance transform method is not ideal as the shape of the randomly generated instances are different from that of the members of the training set, and there is no guarantee of implicit correspondence.

## 4.8 Dependence on Choice of Template

The PDM building process requires the selection of a template. Landmark points are placed on the template by a process of triangulation and decimation. For the creation of the PDM we need the same number of corresponding points on each member of the training set. This is achieved by using crest points and crest lines to obtain a mapping of curvature based landmark points between the template and each member of the training set in a pair-wise fashion. The mapping is used to obtain coefficients which drive a spline-based warp of the template to place its landmark points at corresponding locations on each member of the training set. The deformation gives an approximation of each member of the training set, resulting in what we termed the *corresponding training set*.

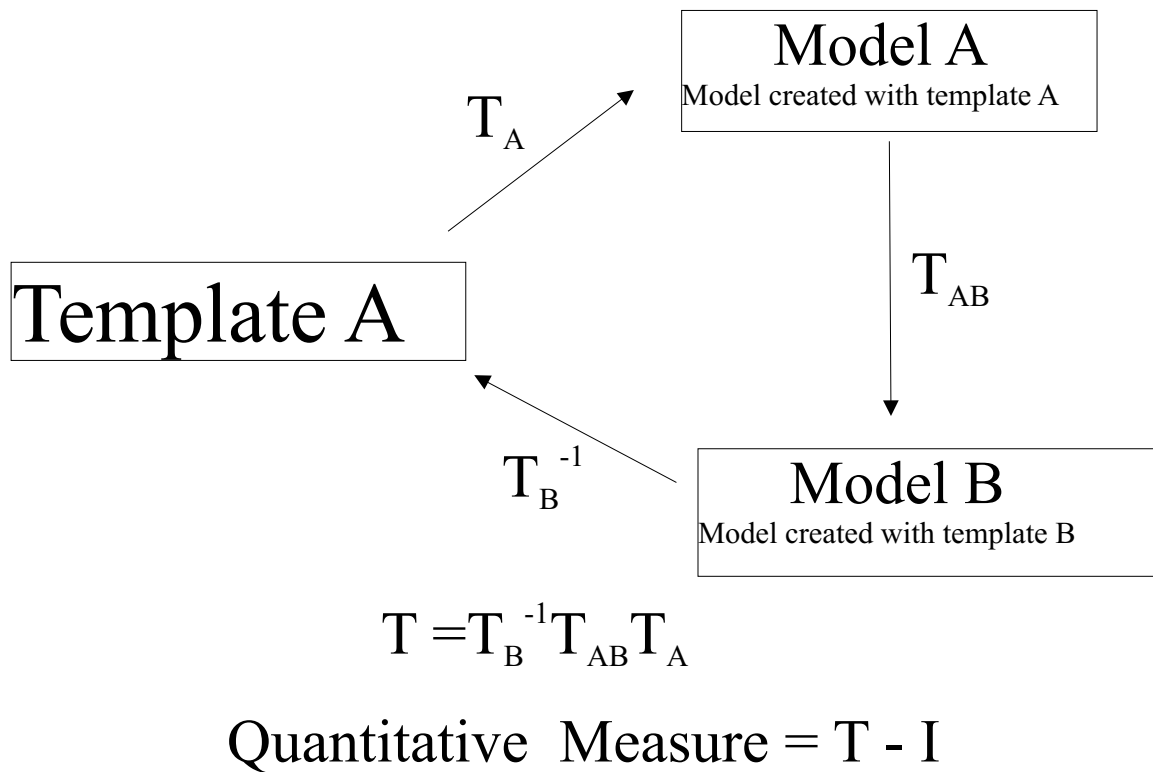
How is the template chosen? There is an argument that it should be chosen randomly. However, visualisation of the surfaces of all the ventricles used in this study (see appendix B) shows that there is great variation in the shape of the members of the training set, some being quite irregular even within the control group. The template was therefore chosen from a subset of the training set that were adjudged to be most representative of the entire training set.

There are several approaches to answer the key question posed in the overview to this chapter - whether the choice of template affects the resulting PDM. These must

involve comparison of PDMs built from the same training set using different members of the training set as the template. However, the manner in which the resulting PDMs are compared is not trivial. Quantitative methods based on the PDM equation (equation 4.6) may be devised. These may use the mean shape ( $\bar{\mathbf{x}}$ ), the principal components ( $\mathbf{P}$ ), the weights ( $\mathbf{b}$ ), or combinations of these. There is also the option of qualitative assessment of the mean shape and the modes of variation by visual inspection. Quantitative assessment is more desirable. However, with the present method there is a complication in that PDMs created using different templates do not generally have the same number of landmark points - as the number of landmark points on the template is determined by its triangulation. It is therefore not straightforward to compare the mean shapes, the principal components and the weight vectors as these are not constrained to have the same number of dimensions. Indeed this may explain why there is a dearth of quantitative comparisons of the effect of templates in methods of construction of PDMs that require a template to be chosen from the training set. In the literature reviewed, of five methods requiring pair-wise matching or choosing a template (i.e. [3],[22],[53],[75],[79]), only one ([53]) included a quantitative assessment of the dependence of the resulting model on the choice of template.

The quantitative assessment of the dependence of the final model on the choice of template in [53] was made in the following manner. Two models were constructed using different members of the training set (say  $A$  and  $B$ ) as templates. A volumetric elastic registration process was the core of their method (more details of their method are in section 4.3). The registration is applied to obtain the parameters of a 9 degree-of-freedom quasi-affine transformation for registering one template (say  $A$ ) to each of the two models and for registering the model created by template  $A$  to the model created by template  $B$ . These allow transformations for registering template  $A$  to itself in a cyclic manner (see figure 4.22) to be applied. The resulting composite transformation matrix would be expected to be the identity matrix. Therefore, the difference of its elements from those of corresponding elements of the identity matrix

is taken as a quantitative measure of the dependence of the models on the choice of template. The results obtained showed that the models did depend on the choice of template. However, no correction for this fact was introduced as the authors felt the dependance did not affect the ability of the model to perform segmentation. Figure 4.22 schematically illustrates how the quantitative measure was derived.



**Figure 4.22:** A schematic diagram illustrating how the quantitative measure used by [53] is obtained.

The use of the transformation parameters of the volumetric elastic registration to measure the dependence of the model on the choice of template is not ideal. This is because the volumetric elastic registration is itself used in the construction of the model. If a dependence on the choice of template does exist, it is very likely that the volumetric elastic registration process is a contributing factor. Therefore, the quantitative measure detailed in figure 4.22 reflects both the inaccuracies of the

registration process as well as the dependence of the model on the template.

Given a training set, the mean shape is expected to be independent of the ventricle chosen as the template. Therefore, if several PDMs are obtained using different members of the training set as templates, their mean shapes would be expected to be approximately equal. This is analogous to the expectation that the average of a set of numbers is independent on the order of the individual numbers. To quantitatively assess the dependence of our model generation scheme on the choice of template, we propose to use the distance transform metric described in section 4.4.4 to measure the similarity of the mean shapes of models generated with different choices of template.

The use of the distance transform metric requires a distance transform image of a binary voxel representation, and a triangulated surface. However, we only have triangulated representations of the mean shapes we intend to compare. We therefore need a binary voxel representation of each mean shape. We generated binary voxel representations (with 0.78mm isotropic voxels) of the surfaces of the mean shapes by obtaining contours in planes representing contiguous slices through the surface in coronal orientation (details are given in appendix A.3). The distance transform of the binary voxel representation can then be taken, and distance transform metrics can be obtained.

A scatter matrix (see Tables 4.2 and 4.3) is constructed whereby the triangulated surface of each mean shape is embedded in the distance transform of every other mean shape (including its own). The minimum values of the distance transforms obtained after optimisation of rigid registrations as described in section 4.4.4 is used as the measure of similarity. In the case of the results reported below, 100 initialisations were used when obtaining the distance transform metric. The values recorded in this chapter are means of the lowest three values. The standard deviations of the lowest three values were used to estimate the error of measurement. If there is no dependence on the choice of template, the scatter matrix would be expected to be a symmetric matrix with elements close to 0.



Additionally, qualitative assessments of the dependence of PDMs on choice of template can be carried out by rendering the mean shapes due to different templates and visualising the variations of the most significant modes of each PDM.

### 4.8.1 Dependence of PDM Construction on Choice of Template without Re-projection

Here we investigate the dependence of the current method of PDM construction (using the transportation method of matching points) on the choice of template without applying the re-projection step (step 5 of Figure 4.3) .

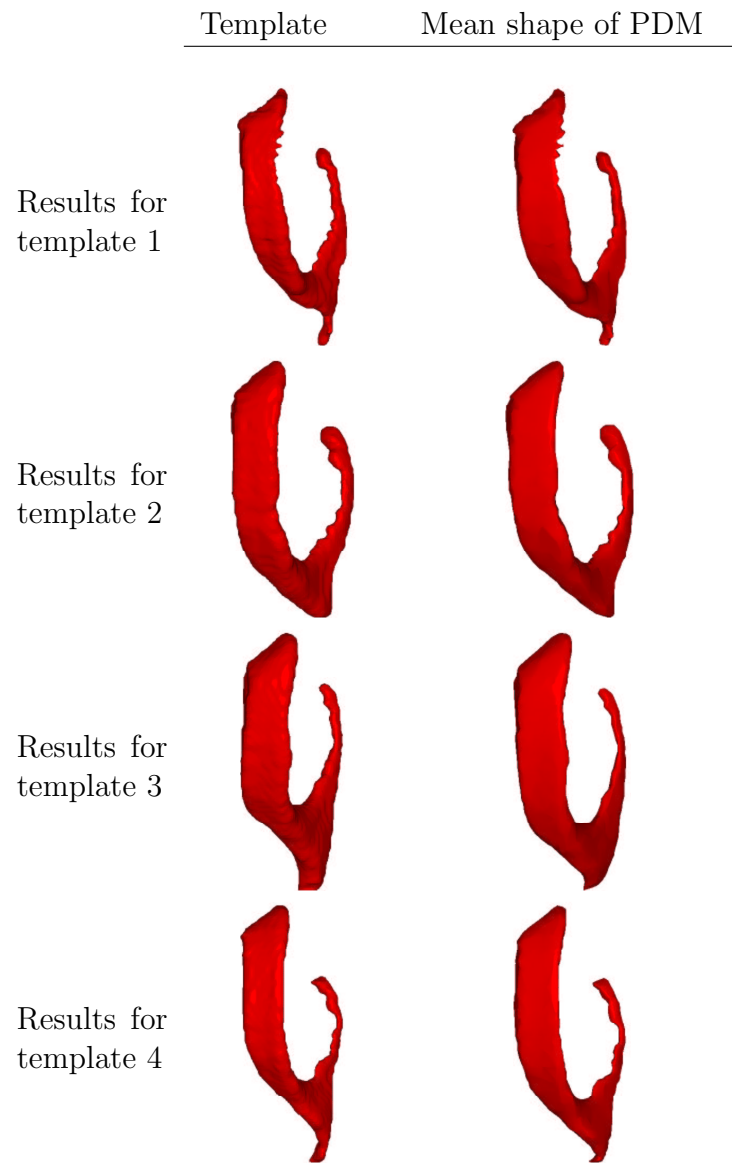
#### Experiments

The segmented ventricles of the 17 male control subjects were used. The left ventricles of each subject were reflected to give the same pose as the right ones giving a training set of 34 ventricles. Four ventricles, three chosen at random and the ventricle adjudged to be most representative of the training set were used as templates. Four PDMs were constructed in the prescribed manner, each using one of the chosen templates. The similarity of the mean shapes of each of the PDMs were compared using the distance transform metric.

#### Results

The qualitative results are shown in Figure 4.23. These show each of the templates (first column of each row) and the corresponding mean shape obtained using the current method (second column of each row). The fact that each mean shape retains features of the ventricle used as a template is obvious in each of the four cases.

Figure 4.26 shows the variation of the most significant mode of each PDM between three standard deviation units of the value at the mean. In all cases (looking from left to right along each row) the model seems to “elongate”. However, the different parts are affected to varying degrees. The width of the anterior and temporal horns



**Figure 4.23:** Mean shapes of PDMs of the same training set obtained from four different templates (down the rows)

decrease, and the temporal horns seem to move closer to the central part of the ventricle. However, in the PDMs constructed using the first and fourth templates, the elongation of the occipital horn is more pronounced than in the PDMs of the second and third templates.

Table 4.2 gives the quantitative results of the comparison of the differences in the mean shapes. The lowest values of the distance transform metric are in the diagonal of the table representing the embedding of a triangulated surface of the mean of a particular PDM within the distance transform of the voxel representation of the same mean. The values of the distance transform metric along the diagonal ranged from 0.079 to 0.095 with a mean of 0.088 ( $\pm 0.007$ ). The off-diagonal values ranged from 0.777 to 1.581 with a mean of 1.029 ( $\pm 0.245$ ). These values are in voxel units - each voxel unit being equivalent to 0.78mm. These results show that the mean shapes of PDMs of the different templates are different - as the distance transform metric of the diagonal entries of the table are an order of magnitude less than those of the other entries. However, the extent to which the mean shapes of the different models depend on the choice of template cannot be measured from this. These results are revisited for a relative comparison in the next section.

### 4.8.2 Re-Projection as a Solution to Dependence on the Template

Figure 4.25 shows that for the current method of construction of PDMs, if the re-projection step is not included the resulting model depends on the choice of template. Three possible reasons for this are put forward in this section. Two of these are unlikely, and the re-projection step was included to address the third reason.

The first is that the number and placement of landmark points might be different for different templates. The second is that the accuracy with which correspondence is established differs with the member of the training set chosen as the template. The

third is that in deforming the template to a target, the template may retain some of its features and not fully deform to the target.

The first reason may exist, but if it does it should not be significant, as the triangulation and the decimation aim to preserve the topological structure of the shapes they are applied to. Furthermore, it is the variation in position of the corresponding landmarks on each ventricle that is being measured, and not the number of landmarks. Therefore, as long as the landmarks fully represent the shape being modelled, and are placed in corresponding locations on each member of the training set their number and placement on the template should not matter. The second reason may also exist in that using a ventricle that has a highly irregular shape may result in worse pair-wise crest line correspondence with more members of the training set than a ventricle with a more regular shape.

However, the third reason seems more the culprit. Visualisation of the corresponding training set of the experiment in section 4.8.1 showed that although each template was deformed to approximate the target, the deformed template still retained some of its own characteristics. Row 2 of Figure 4.24(b) shows an example from the corresponding training set used to construct the PDM of the lateral ventricles of the controls which illustrates this. We therefore propose to improve the approximation of each member of the training set by the template. This can be achieved by re-projecting the landmark points of the deformed template approximating each ventricle (the corresponding training set) back onto the surface of the actual ventricle. For each landmark on the approximating ventricle, the re-projection is along the normal to the surface of the approximating ventricle. The closest vertex on the intersection of the normal with the actual ventricle is taken as the better approximation of the landmark. The premise is that the deformed template is a reasonable approximation of the target ventricle. Therefore, if the deformed template and the target it is approximating are aligned, the projection of the vertices of the deformed template along the normal to the vertices should intersect the target at a corresponding point. The landmarks on each member of the corresponding training set can therefore be

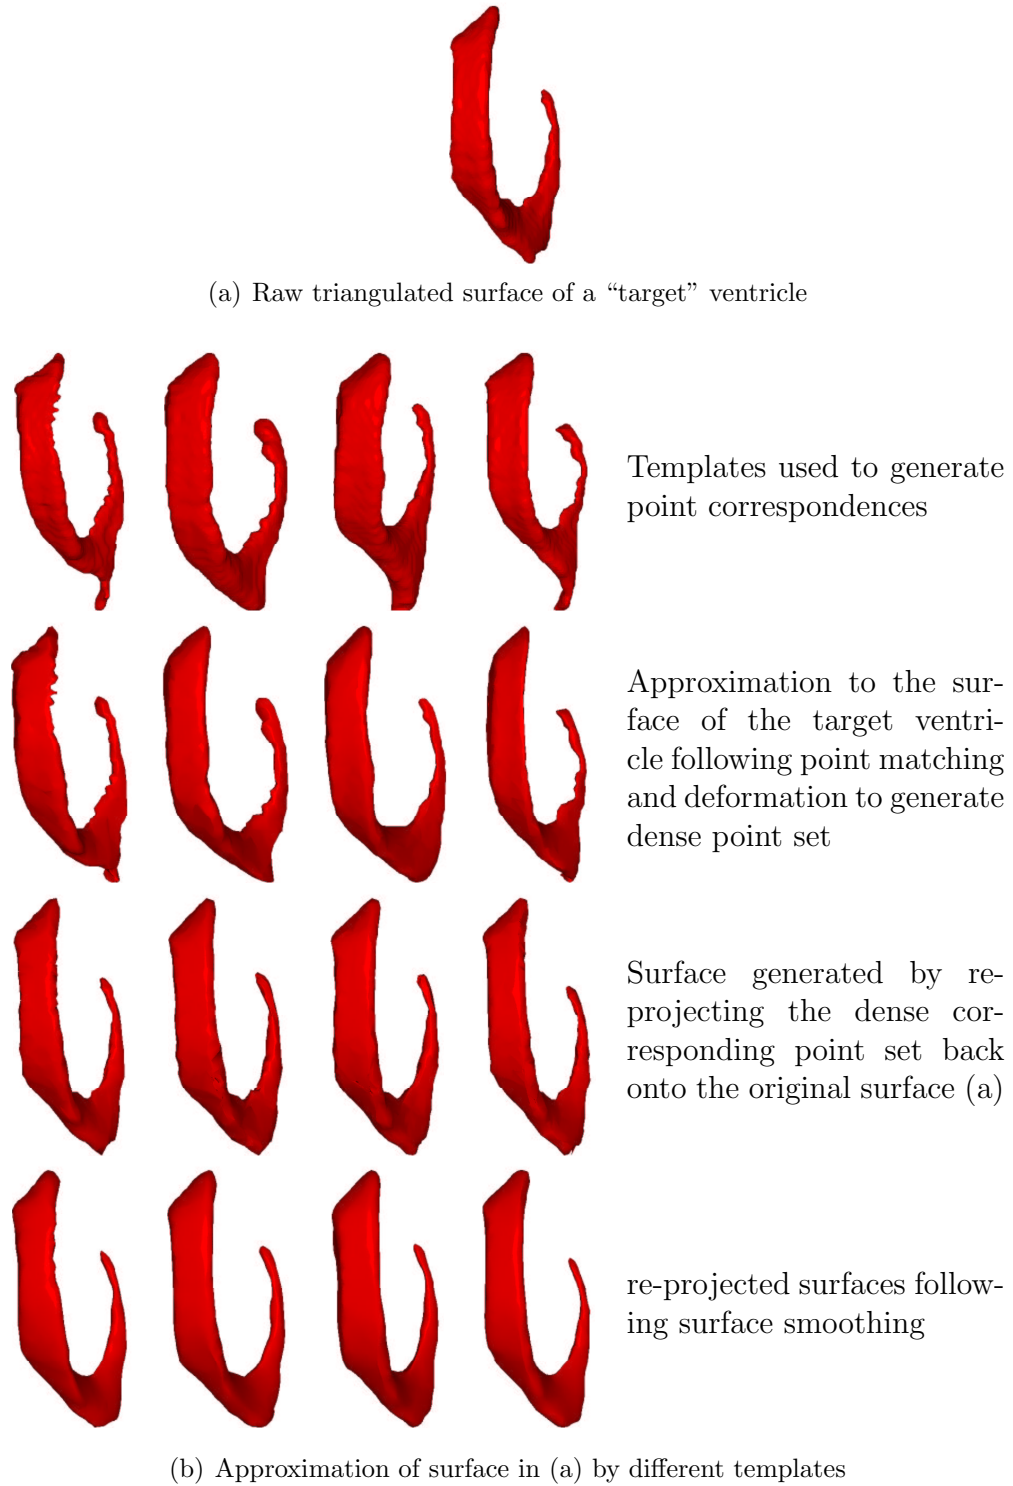
transferred onto the surfaces of the actual ventricles in this manner.

Two steps are needed to implement the surface projection described above. Firstly, we need to calculate the outward normals for the deformed template surface. Secondly, we need to calculate the intersection of these normals with the raw (triangulated) target surface. Appendix A give details of the two steps in this process: calculating the surface normals (appendix A.1) and calculating the intersection of these normals with the target surface (appendix A.2).

The smoothing method of Taubin [111] was applied to the resulting surfaces purely for visualisation purposes to reduce the “jaggedness” due to the discrete nature of the surface projection. The effect of this step on the positions of the vertices of the triangulated surface was negligible. The mean difference in Euclidean distance between vertices before and after smoothing was applied was  $0.87 \pm 0.59\text{mm}$ . Figure 4.24 shows an example of a member of the corresponding training set obtained before and after applying the projection of vertices and smoothing. These results are typical of the entire training set, and support the quantitative results that re-projecting decreases the dependence of the PDM on the choice of template (see section 4.8.3). However, these results suggest that although the influence of the template on the corresponding training set is greatly reduced by re-projection, this influence is not eliminated.

### 4.8.3 Dependence of PDM Construction on Choice of Template, Using Re-projection

The experiments described in section 4.8.1 were repeated here. The only difference was that prior to construction of the PDM, the vertices of each member of the corresponding training set were projected onto the actual surfaces as described in section 4.8.2.



**Figure 4.24:** Qualitative comparison of a member of the corresponding training set produced by four different choices of template with and without the re-projection step. Subfigure (a) shows the original triangulated surface of the “target”. In subfigure (b), the 1<sup>st</sup> row shows the four templates. Subsequent rows show the surfaces produced by deformation of each template to the target with and without re-projection.

	Surface 1	Surface 2	Surface 3	Surface 4
Dist Trans 1	0.0910 (0.0023)	0.9576 (0.0062)	0.9468 (0.0202)	0.7768 (0.0093)
Dist Trans 2	1.5807 (0.0142)	0.0947 (0.0036)	0.9432 (0.0073)	1.1902 (0.0158)
Dist Trans 3	1.3746 (0.0386)	0.7906 (0.0225)	0.0792 (0.0010)	1.1052 (0.0225)
Dist Trans 4	0.9614 (0.0169)	0.8374 (0.0029)	0.8819 (0.0219)	0.0868 (0.0004)

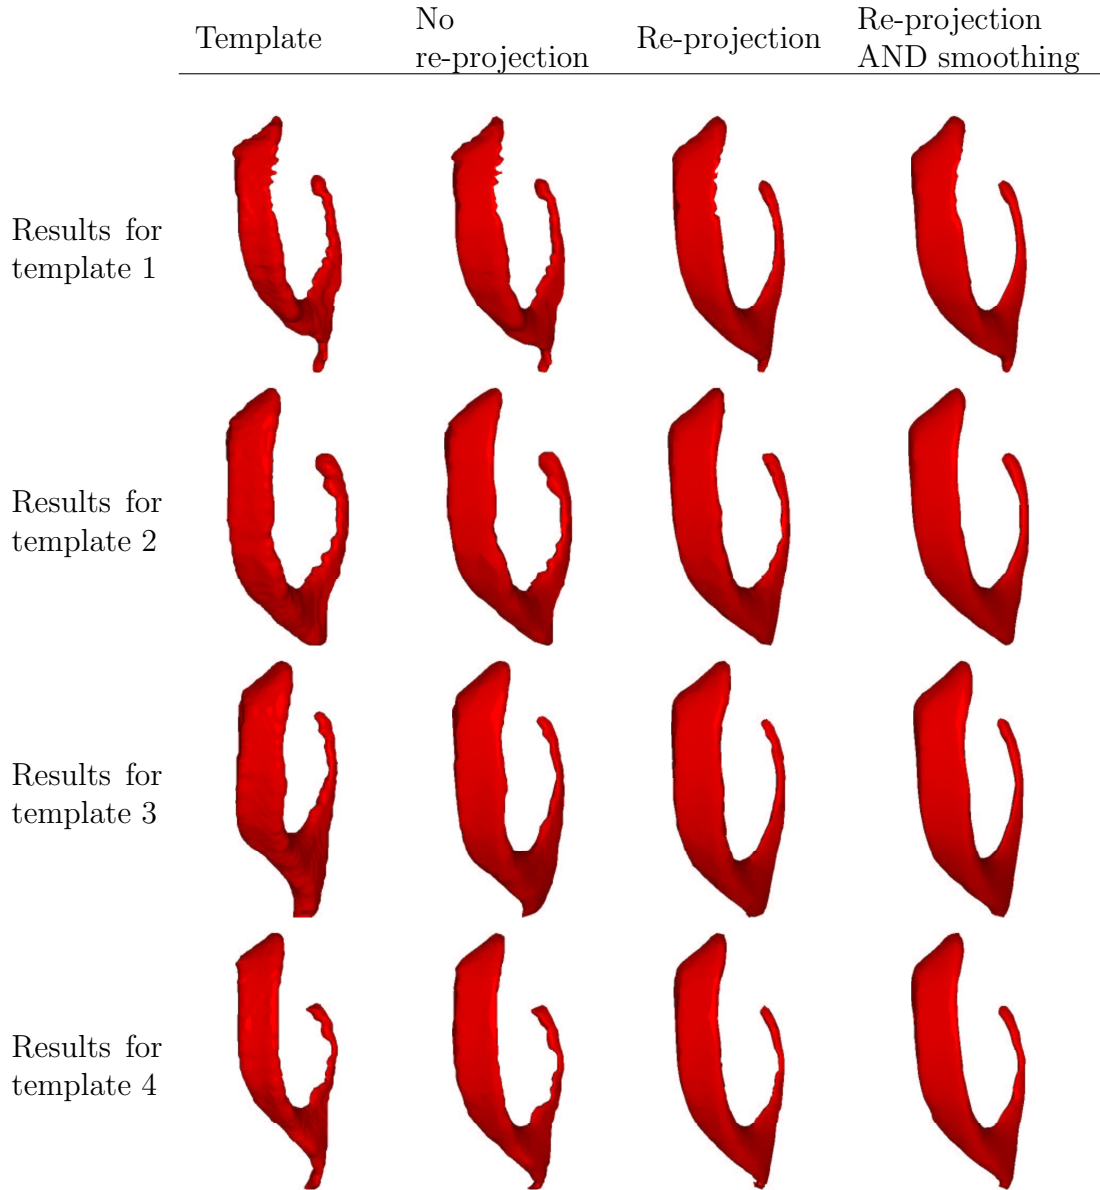
**Table 4.2:** Values of the distance transform metric for the comparison of mean shapes of PDMs. There are four PDMs of the same training set. Each PDM was created using a different member of the training set as the template. The values along each row represent the values of the metric for the triangulated surface of each mean shape embedded in the distance transform of the mean shape of the particular PDM. The values in parenthesis give the error estimate of the distance transform metric.

## Results

Qualitative results are shown in figures 4.25. The first, third and fourth columns of each row are relevant to this section. These show each of the templates (first column of each row), the corresponding mean shape after inclusion of the projection step (third column of each row), and the mean shape when smoothing was applied to each member of the corresponding training set after the projection (fourth column). Comparison of the third and fourth columns supports the assertion that the smoothing merely increases the aesthetic quality of the surfaces. The images of the mean shapes in the third and fourth columns show that they are much more similar to each other, and less similar to the ventricle used as the template than corresponding mean shapes in column 2.

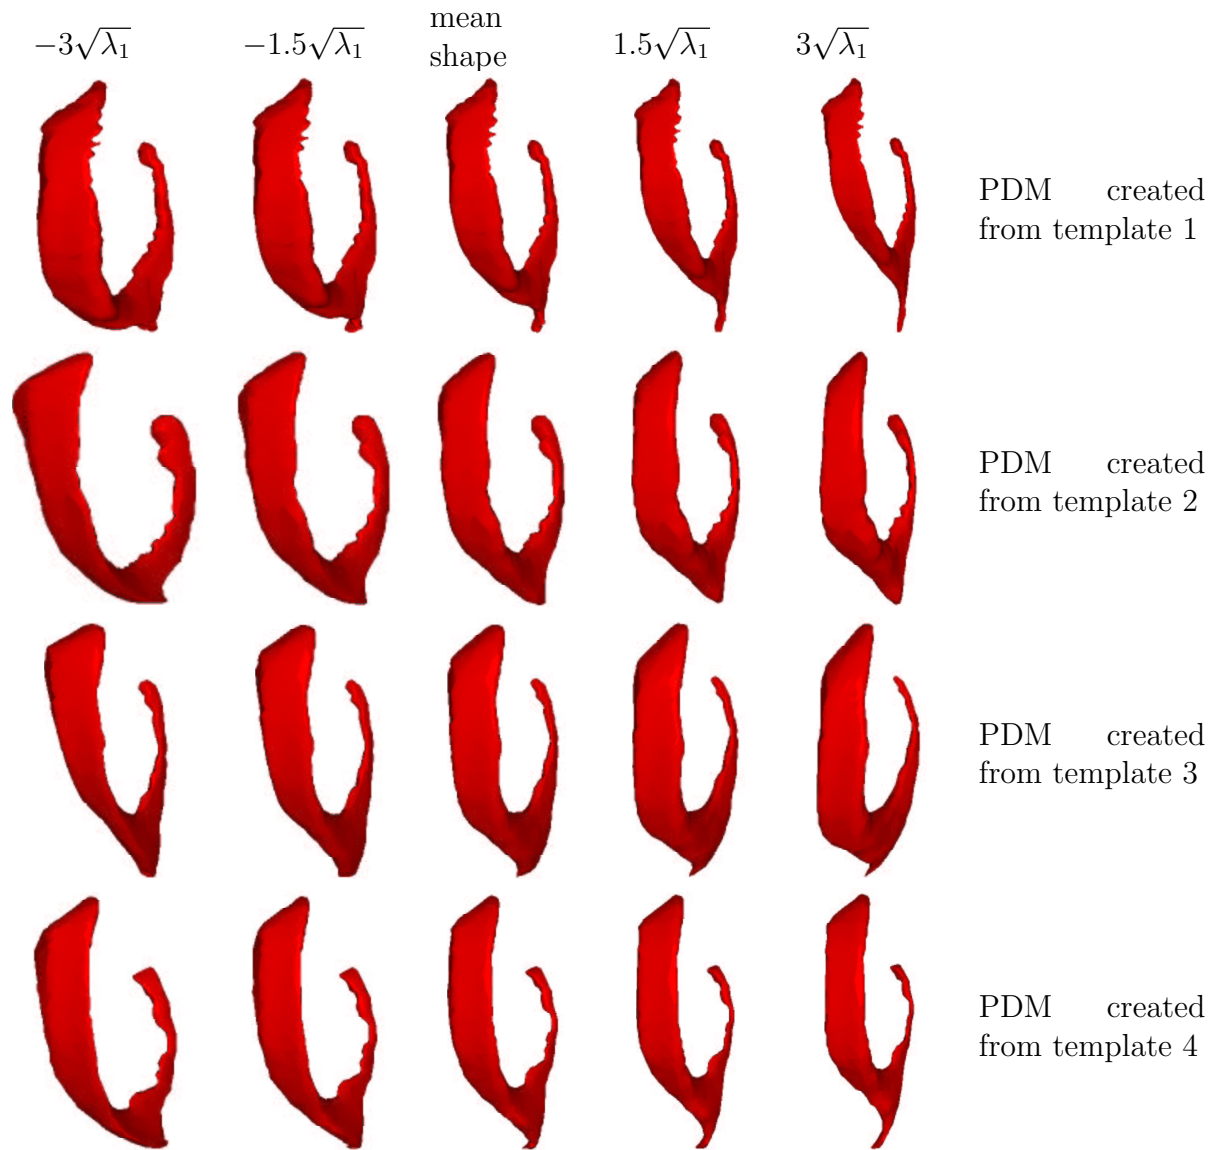
Figure 4.27 shows the variation of the most significant mode of each PDM out to three standard deviation units on either side of the mean. In all cases (looking from left to right along each row) the model seems to “elongate”. The modes of variation seem to be similar for all four PDMs, although the physical instances shown are different.

Table 4.3 gives the quantitative results of the comparison of the differences in the mean shapes. As with table 4.2 the lowest values of the distance transform metric are in the diagonal entries. These ranged from 0.079 to 0.095 with a mean of  $0.074 \pm 0.011$ . The off-diagonal values ranged from 0.266 to 1.208 with a mean of  $0.642 \pm 0.259$ .

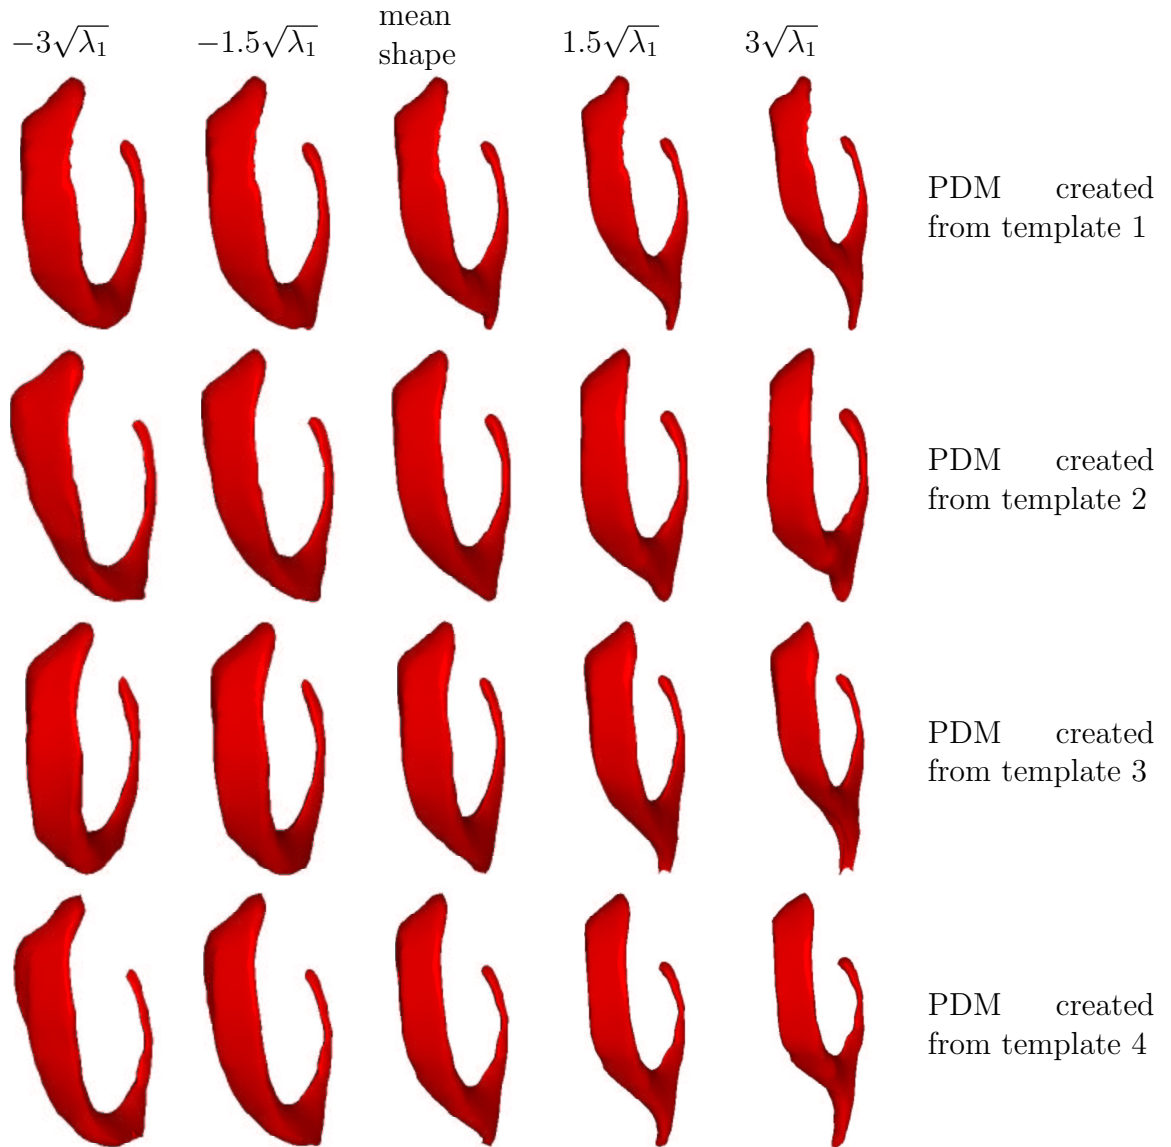


**Figure 4.25:** Mean shapes of PDMs of the same training set obtained from four different templates (down the rows) by three different methods (along the columns). Column 1 shows the surfaces of the templates, column 2 the mean shapes of the PDM before the re-projection, column 3 shows the surfaces of the mean shapes after the re-projection step, and column 4 shows the surfaces of the mean shapes after the re-projection step is followed by smoothing





**Figure 4.26:** First mode of variation of the PDMs of the same training set obtained from four different templates. The columns of each row are the modes of variation for -3 s.d. units (column 1), -1.5 s.d. units (column 2) the mean shape - 0 s.d. units (column 3), 1.5 s.d. units (column 4) and 3 s.d. units (column 5).



**Figure 4.27:** First mode of variation of the PDMs of the same training set as for figure 4.26 obtained from four different templates after the projection of landmark points of the corresponding training set onto the surfaces of the actual ventricles in the training set. The columns of each row are the modes of variation for -3 s.d. units (column 1), -1.5 s.d. units (column 2) the mean shape - 0 s.d. units (column 3), 1.5 s.d. units (column 4) and 3 s.d. units (column 5).

	Surface 1	Surface 2	Surface 3	Surface 4
Dist Trans 1	0.0683 (0.0067)	0.7041 (0.0054)	0.6618 (0.0170)	0.4525 (0.0135)
Dist Trans 2	1.2081 (0.0029)	0.0736 (0.0061)	0.3610 (0.0164)	0.6787 (0.0058)
Dist Trans 3	0.9923 (0.0182)	0.2663 (0.0200)	0.0896 (0.0043)	0.6628 (0.0230)
Dist Trans 4	0.6657 (0.0089)	0.4862 (0.0114)	0.5668 (0.1005)	0.0655 (0.0018)

**Table 4.3:** Values of the distance transform metric for the comparison of mean shapes of PDMs after projection of normals and smoothing steps were added to the method of construction of the PDMs. There are four PDMs of the same training set. Each PDM was created using a different member of the training set as the template. The values along each row represent the values of the metric for the triangulated surface of each mean shape embedded in the distance transform of the mean shape of the particular PDM. The values in parenthesis give the error estimate of the distance transform metric.

	Mean of diagonal terms	Mean of off-diagonal terms
Before re-projection	$0.075 \pm 0.005$	$1.029 \pm 0.245$
After re-projection	$0.074 \pm 0.011$	$0.642 \pm 0.259$
significance of difference	$p = 0.896$	$p = 0.001$

**Table 4.4:** Means of the diagonal and off-diagonal terms of Tables 4.2 and 4.3. The last row shows the significance of the difference of the means using a *t-test*.

## Discussion

The results of section 4.8.1 can be compared qualitatively with those obtained in this section. Firstly, the mean shapes of the different PDMs show that projection of the landmarks of the deformed template onto the surface of the target significantly reduces the similarity of the mean shape to the template. However, there might be a residual similarity to the target. This may explain why, although the mean shapes are more similar to each other for the case where the landmarks were projected onto the surface of the target, some noticeable difference still exist.

Secondly, the variation of the most significant mode between 3 standard deviation units about the mean shows similar trends for both the PDMs constructed before and after the surface projection step was introduced. However, in the case after the introduction of the surface projection step, the modes of variation were markedly more similar for all the PDMs. The most salient part of this was the elongation of the occipital horn. This result suggests that although the mean shapes may differ,

the underlying differences between the shapes are still being captured regardless of the choice of template. This result is important because our intention is to use the PDM for shape analysis. Therefore the relative difference between the members of the training set are what we are interested in.

The suggestion by the qualitative results, that the surface projection reduced the dependence of the resulting PDM on the choice of template, is backed up by the quantitative results of the distance transform metric given in tables 4.2 and 4.3. As suggested in section 4.8.2 these tables would have been expected to be symmetric, with the diagonal entries being zero. The diagonal entries in both cases are significantly smaller than the other entries. However, the tables do not seem to be symmetric. This suggests that the embedding of the triangulated surface of mean shape  $A$  into the distance transform of the binary voxel representation of mean shape  $B$  does not give the same value for the distance transform metric as the embedding of surface  $B$  in the distance transform of  $A$ . This observation is most likely due to the discretisation process necessary for the conversion of the triangulated surface into the voxel representation.

The mean values of the diagonal entries of Tables 4.2 and 4.3 are given in Table 4.4. The mean values of the non-diagonal entries are also given in the table. This shows that the inclusion of the surface projection step does not cause a significant change in the mean values of the diagonal entries. This is to be expected as the diagonal entries are simply the comparisons of the triangulated and voxelised representations of the same mean shape. However, the inclusion of the surface projection does cause a significant reduction in the distance transform metric of the non-diagonal entries. This is quantitative evidence to back up the qualitative observation that surface projection reduces the dependence of the PDM on the choice of template.

## 4.9 Summary

We have described the mathematical details of PDMs and issues pertaining to their construction in 3D. A method of PDM construction based on curvature and involving the transportation algorithm was outlined and used in the construction of the PDM of the lateral ventricles. The characteristics of the PDM were investigated quantitatively and qualitatively, and measures to reduce the dependence of the PDM on the choice of template were implemented. Appendix B shows the surfaces of all the ventricles and the corresponding surfaces used in construction of the PDM of all 138 (left and right) ventricles of all the subjects. The modes of variation of the first three modes of this PDM can be seen on the accompanying CD.

# Chapter 5

## Morphometric Analysis

### 5.1 Overview

This chapter gives the results of the morphometric analysis of the ventricles of the schizophrenic and control subjects. We investigate both shape differences between the schizophrenic and control groups, as well as differences in the extent of asymmetry of left and right ventricles in both groups. These comparisons are the main aim of the thesis, and the culmination of work described in previous chapters.

In chapter 1 the motivation for morphometric analysis, as well as related research was discussed. Chapter 2 described shape analysis and explained why the PDM approach was chosen. Description of what a PDM is, and the method of construction of the 3D PDM of the lateral ventricle was given in chapter 4. Parameters obtained from the PDM are used in the morphometric analysis performed in this chapter. The morphometric analysis is based on linear discriminant analysis (LDA). The PDM is also used in characterising the differences in shape inferred from the LDA by allowing instantiation of intermediate surface representations showing progression of differences from the schizophrenic extreme to the control extreme.

The chapter is organised as follows. Section 5.2 states the experimental method. Section 5.3 describes how the parameters obtained from the PDM are useful for morphometric analysis. The morphometric analysis involves investigating shape differences between a *pair of groups* (control and schizophrenics) and differences between degree of asymmetry within a *group of pairs* (left and right ventricles of the same subject). The investigations we wish to make are detailed in section 5.4. Section 5.5 describes the basic linear discriminant analysis used for the investigation of shape differences, and section 5.6 the *paired-linear discriminant analysis* used in the case of the asymmetry investigations. In section 5.7 we show that significant shape differences are observed between control and schizophrenic groups, and that these are adjacent to neuroanatomical regions that have known association with schizophrenia. However, we also note that because the discriminants are determined from the training set, the  $p$ -values quoted are optimistically biased. Results of the asymmetry analysis are also presented in this section. The results are discussed in section 5.8, and a summary of the main points of the chapter is given in section 5.9.

## 5.2 Experimental Methods

The PDM used in the subsequent analysis was constructed as described in section 4.5 using segmented ventricles of all subjects. The left ventricles of each subject were reflected to give the same pose as the right ventricles, giving an evaluation set of 138 ventricles - 60 from the control group (26 female, 34 male) and 78 from the age and sex matched schizophrenic group (18 female, 60 male).

### 5.3 Using the Parameters of the PDM for Morphometric Analysis

The parameters of the PDM are the mean shape,  $\bar{\mathbf{x}}$ , the principal components  $\mathbf{P}$ , and the weights vector,  $\mathbf{b}$ . The concept of *parameter space* or *shape space* was described in section 4.2. The mean shape can be taken as the origin of the shape space, and each principal component as defining an orthogonal direction from the mean shape in this shape space. The elements of  $\mathbf{b}$  are orthogonal shape features in the shape space. The dimensionality of the shape space is determined by the number of principal components kept. In the present case, the first 30 principal components were used. These explained over 95% of the variance captured by the PDM.

Each member of the training set is a point within this 30-dimensional space, the coordinates of which are obtained by inversion of equation 4.6. For the  $i^{th}$  member of the training set, this is given by :

$$\mathbf{b}_i = \mathbf{P}^T(\bar{\mathbf{x}} - \mathbf{x}_i) \quad (5.1)$$

The reduction of dimensionality (from 3873 to 30 in the present case) was one of the reasons for using the PDM approach. For each member of the training set we therefore have a 30-dimensional parametric representation obtained from the PDM. These are used for the analysis detailed in the rest of this section.

### 5.4 Comparisons we Wish to Make

We seek to identify differences in shape between the groups:

- Are there significant differences in shape between the ventricles of the control and schizophrenic groups?



- If such differences exist, how do they relate to brain anatomy and, in particular, regions believed to be associated with schizophrenia?
- Is there any difference between male and female subgroups in this respect?

We also wish to investigate whether differences exist in lateral asymmetry between groups. The degree of left-right asymmetry of the ventricular shape between the control and schizophrenic subjects is compared for the pooled male and female groups, and for each subgroup.

The six comparisons we carry out are listed explicitly below:


1. Shape differences between all schizophrenic subjects and all control subjects.
2. Shape differences between all *male* schizophrenic subjects and all *male* control subjects.
3. Shape differences between all *female* schizophrenic subjects and all *female* control subjects.
4. Difference in extent of asymmetry between left-right pairs of all control subjects compared with extent of asymmetry between left-right pairs of all schizophrenic subjects
5. Difference in extent of asymmetry between left-right pairs of all *male* control subjects compared with extent of asymmetry between left-right pairs of all *male* schizophrenic subjects
6. Difference in extent of asymmetry between left-right pairs of all *female* control subjects compared with extent of asymmetry between left-right pairs of all *female* schizophrenic subjects

## 5.5 Shape (Group) Comparisons Using Linear Discriminant Analysis


Linear discriminant analysis (LDA) forms the core of our investigations into shape and asymmetry differences. Brief details relevant to this particular application are given in this section and in section 5.6. Texts such as [46] (chapter 4) and [62] (chapter 4) give fuller details of LDA.

The shape comparisons we wish to carry out are essentially comparisons between two groups located in a 30-dimensional space. We wish to investigate if the two groups can be distinguished in this space, and describe the difference(s) between the groups. An ideal situation would be each group being a multivariate gaussian distribution centred about a mean with known standard deviations and no overlap in each dimension. Figure 5.1(a) illustrates this in the 2D case. However, in practice a degree of overlap will exist between the groups. Figure 5.1(b) illustrates this by plotting the values of the first and second elements of the parameter vectors of the male schizophrenic and male control groups. Furthermore the fact that the shape space is multidimensional means a hyperplane would be needed for separation of the groups.

The use of LDA allows simplification of the problem of separation of the two groups by reducing it to a one-dimensional problem. LDA involves finding a *discriminant vector* in shape space which gives a direction along which the difference between the groups is most marked. Shape differences can be quantified by projecting the individual shape vectors of each ventricle onto the discriminant vector to provide a scalar value representing each ventricle. The extent of the difference between the groups is mirrored by the difference between the means of the scalar values of the members of each group. Furthermore, the nature of the shape differences between the groups can be visualised by reconstructing shapes along this vector between the schizophrenic extreme and the control extreme. Specific differences correspond to



(a) Ideal case



(b) Real case

**Figure 5.1:** Illustration of an ideal case of separation of two groups (a) compared with the actual case for the first and second shape parameters of the male schizophrenics and male controls

locations on the surface where differences are observed between corresponding points of the reconstructed shapes.

The importance of the direction of the discriminant vector in obtaining optimum separation between the two groups is illustrated in figure 5.2. This shows the projection of the two non-overlapping gaussian distributions of figure 5.1(a) onto discriminant vectors in two directions. In one direction chosen arbitrarily (subfigure a), the separation is not good as there is an overlap of the projections of the two groups. The other direction is obtained by applying the steps discussed below to maximise Fisher's criterion. The separation in this case (subfigure b) is better as there is no overlap between the projections of the two groups.

To obtain the direction of the discriminant vector in the present case, we use Fisher's criterion [50]. According to Fisher, the direction of optimum separation is that in which the difference between the class means is maximised relative to the standard deviations of the classes in that particular direction. This is encapsulated in Fisher's criterion, which for our purposes can be written:

$$J_F = \frac{\text{distance between sample means}}{\text{standard deviation within sample means}} = \frac{\mathbf{w}^T \bar{\mathbf{b}}_c - \mathbf{w}^T \bar{\mathbf{b}}_s}{\sqrt{\mathbf{w}^T \mathbf{S} \mathbf{w}}} \quad (5.2)$$


where  $\mathbf{w}$  is a vector in the direction that  $(\mathbf{w}^T \bar{\mathbf{b}}_c - \mathbf{w}^T \bar{\mathbf{b}}_s)$  is maximised relative to the standard deviation in that direction.

$\bar{\mathbf{b}}_c$  is the mean of the control group


$\bar{\mathbf{b}}_s$  is the mean of the schizophrenic group

$\mathbf{S}$  is the total within class scatter, which gives a measure of the total within group variance of the schizophrenic and control groups

Differentiating equation 5.2 and setting the result to 0 allows the direction of the discriminant vector,  $\hat{\mathbf{w}}$  that maximises  $J_F$  to be determined. This is given by



(a) Bad separation - direction of discriminant vector  $w$  chosen arbitrarily



(b) Good separation - direction of discriminant vector  $w$  given by Fisher's discriminant analysis

**Figure 5.2:** The importance of the direction of the discriminant vector in obtaining good separation is illustrated. In (a) the direction is chosen arbitrarily, whereas in (b) it is that which maximises Fisher's criterion

$$\hat{\mathbf{w}} = \alpha \mathbf{S}^{-1}(\bar{\mathbf{b}}_c - \bar{\mathbf{b}}_s) \quad (5.3)$$

$\alpha$  is set to normalise the result so that the unit vector  $\hat{\mathbf{w}}$  giving the direction of  $\mathbf{w}$  can be obtained.

In the results presented, the definition of the total within class scatter matrix according to [46] was used :

$$\mathbf{S} = \mathbf{S}_c + \mathbf{S}_s \quad (5.4)$$

Where

$\mathbf{S}_c$  and  $\mathbf{S}_s$  are the unnormalised covariance matrices of the schizophrenic and control groups. i.e. for the control group,  $\mathbf{S}_c = \sum_{i \in s} (\mathbf{b}_i - \bar{\mathbf{b}}_c)(\mathbf{b}_i - \bar{\mathbf{b}}_c)^T$

The scalar values facilitate the investigation of shape differences in three ways. Firstly, the significance of the separation between the groups can be obtained by applying a  $t$ -test to the scalar values for the schizophrenic and control groups.

Secondly, localisation of the magnitude of the shape differences on the ventricular surface is given by the Euclidean distance between corresponding points of instantiations of ventricular surfaces corresponding to the means of the scalar values of both groups on the discriminant vector. The localisation of shape differences can be presented intuitively by colour-coding the ventricular surface using this difference. It should be noted that the differences are expressed in arbitrary *normalised shape units* due to the fact that in the creation of the PDM, the members of the training set were normalised for brain size, and subjected to transformations to make them Euclidean invariant.

Thirdly, the shape change involved in moving from the schizophrenic extreme to the control extreme along the discriminant vector can be animated (see accompanying CD). The extremes are taken as three standard deviation units from the means along

the discriminant vector. Instantiation of ventricle surfaces at intervals on  $\mathbf{w}$  between the schizophrenic and control extremes can be made by substituting the position vector of sampled points along  $\mathbf{w}$  into the PDM equation (equation 4.6) i.e.


$$\mathbf{x} = \bar{\mathbf{x}} + \mathbf{P}\lambda\hat{\mathbf{w}} \quad (5.5)$$

Where  $\lambda$  is the scalar value giving the position on  $\mathbf{w}$  at which the surface instantiation occurs. The first and third methods above are similar to the analysis applied to the hippocampi of schizophrenic and control subjects by Davies [38].

## 5.6 Asymmetry Comparisons Using Paired Linear Discriminant Analysis

In this section we describe the use of LDA in the analysis of asymmetry within the schizophrenic and the control groups. Whereas, in section 5.5 we sought to maximise the separation between the schizophrenic and control groups, here we seek to maximise the separation between left and right ventricles of the same subject. It is the difference between the scalar values of left and right counterparts projected onto the discriminant vector that we are interested in. It should be noted that the left ventricles had been reflected to give the same pose as the right ventricles. The analysis presented here follows a method suggested by Poxton et al. [93], used in the analysis of asymmetry in 2D hippocampal slices. The premise is that there is an underlying trend in the separation of left and right pairs within shape space. The hypothetical example of figure 5.3 illustrates this.

For the control group, the discriminant vector we seek is that which maximises separation of all left and right pairs in shape space. The same goes for the schizophrenic group, hence for each of comparisons 4 to 6 we obtain one discriminant vector for the control group and one for the schizophrenic group. To obtain the direction of



**Figure 5.3:** Illustration of finding the discriminant vector maximising the separation between a group of pairs. The figure shows the vector obtained after application of the methods described in section 5.6

the discriminant vector, Poxton et al. substituted the differences of the means in the numerator of Fisher's criterion (equation 5.2) with the mean difference between left and right pairs. The scatter matrix,  $\mathbf{S}$ , was replaced with a *paired covariance matrix*,  $\mathbf{S}_p$ .

$$J_P = \frac{\text{mean difference between left and right pairs}}{\text{standard deviation of difference}} = \frac{\mathbf{w}_p^T \bar{\mathbf{d}}}{\sqrt{\mathbf{w}_p^T \mathbf{S}_p \mathbf{w}_p}} \quad (5.6)$$

where  $\mathbf{w}_p$  is a vector in the direction that  $\mathbf{w}_p^T \bar{\mathbf{d}}$  is maximised relative to the standard deviation in that direction.

$\bar{\mathbf{d}} = \frac{1}{n} \sum_{i=1}^n \mathbf{d}_i$  - with  $\mathbf{d}_i$  being the difference in the parameter vectors of the  $i^{th}$  left and right pair.

The measure of the standard deviation of the difference between the left and right



pairs is taken as the *paired covariance matrix*  $\mathbf{S}_p$  which is simply the covariance matrix of the differences between the left and right pairs :

$$\mathbf{S}_p = \frac{1}{n-1} \sum_{i=1}^n (\mathbf{d}_i - \bar{\mathbf{d}})(\mathbf{d}_i - \bar{\mathbf{d}})^T \quad (5.7)$$

The direction of the discriminant vector in this case is

$$\hat{\mathbf{w}}_p = \alpha \mathbf{S}_p^{-1} \bar{\mathbf{d}} \quad (5.8)$$

We assess comparisons 4 to 6 in four ways. The first is by using a pair-wise  $t$ -test to investigate whether the scalar values of the projections of the left and right ventricles onto the discriminant vector for each group are significantly different. This indicates whether a significant degree of asymmetry exists within each group. Secondly, as in section 5.5 we show the localisation of shape differences between left and right using the difference between corresponding points on instantiations of ventricular surfaces of the left and right means on the discriminant vector.

The above two assessments are similar to those carried out in [93]. However, we go further by doing two more assessments. Firstly, we perform a  $t$ -test on the differences in the scalar values of left-right pairs in the schizophrenic and control groups. This indicates whether there is a significant difference in the magnitude of left-right asymmetry in the two groups. Secondly we find the scalar product of the unit vectors in the directions giving best discrimination for the control and schizophrenic group. This indicates whether the shape differences between the left and right pairs are of the same nature in schizophrenics and controls. If the magnitude of the scalar product is close to 1 the discriminant vectors are co-directional and represent shape changes of the same nature. On the other hand, if the magnitude of the scalar product is close to 0, it indicates that the discriminant vectors are orthogonal, representing different kinds of shape change.

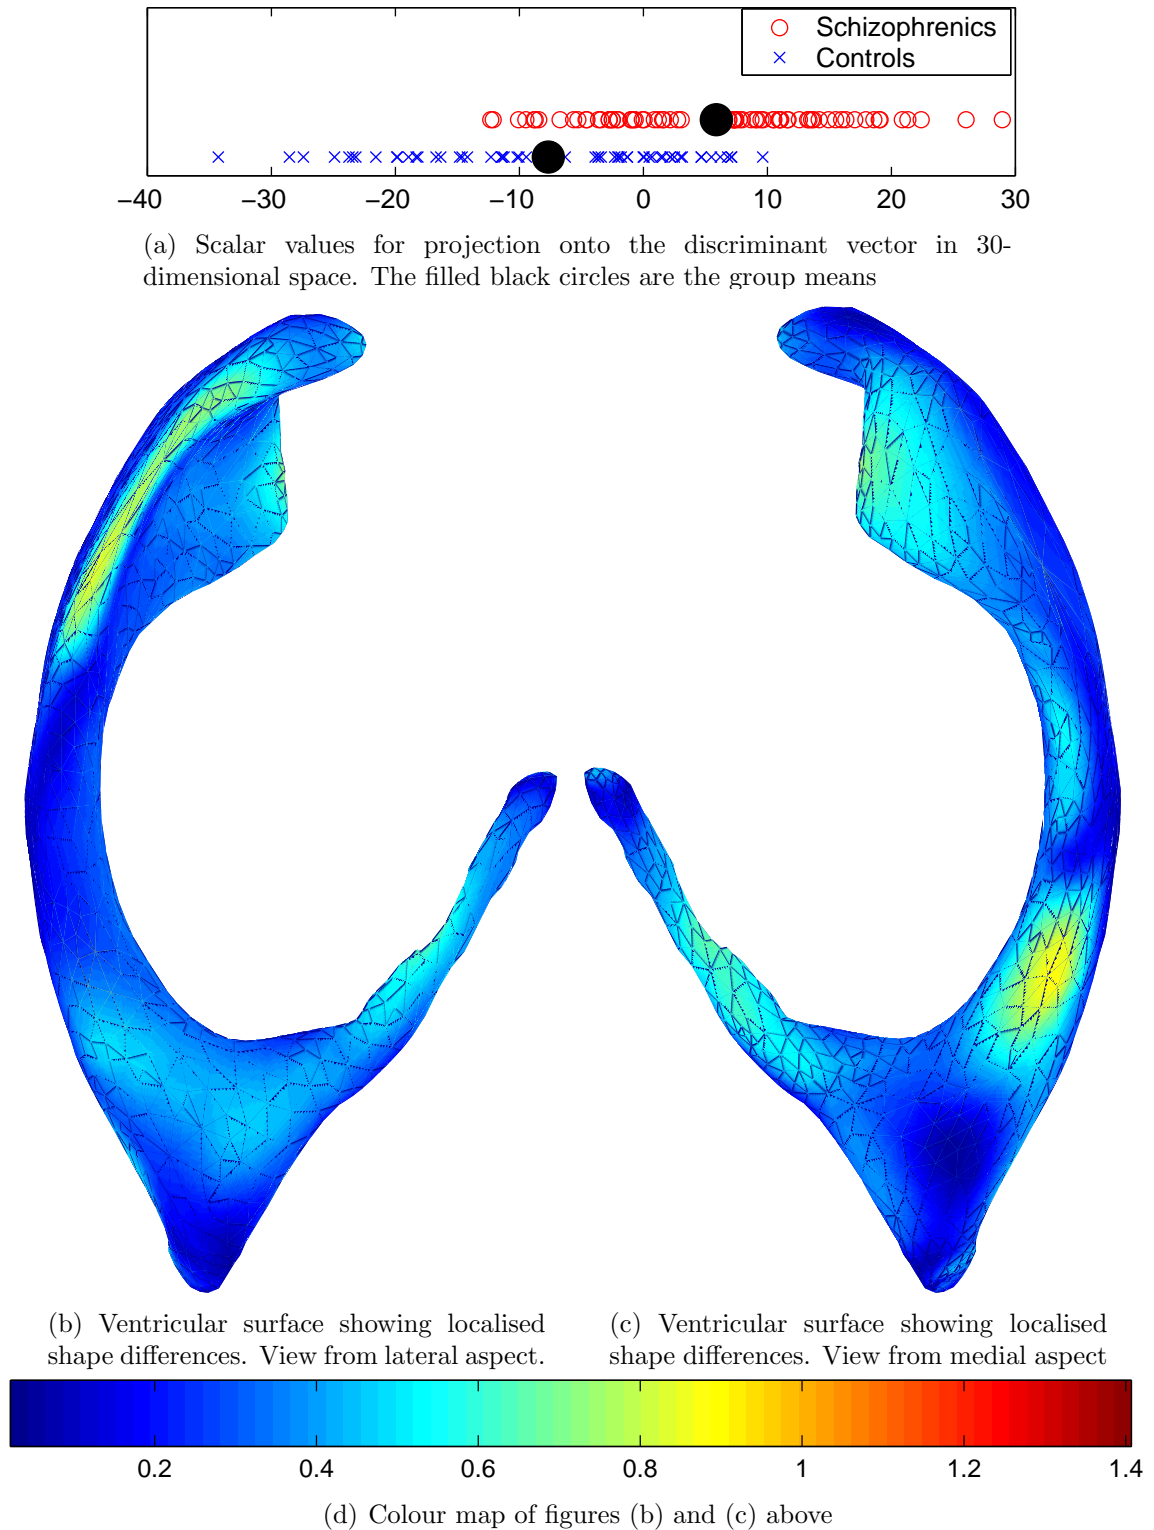
## 5.7 Results of Morphometric Analysis

### 5.7.1 Shape Differences Between Schizophrenics and Controls

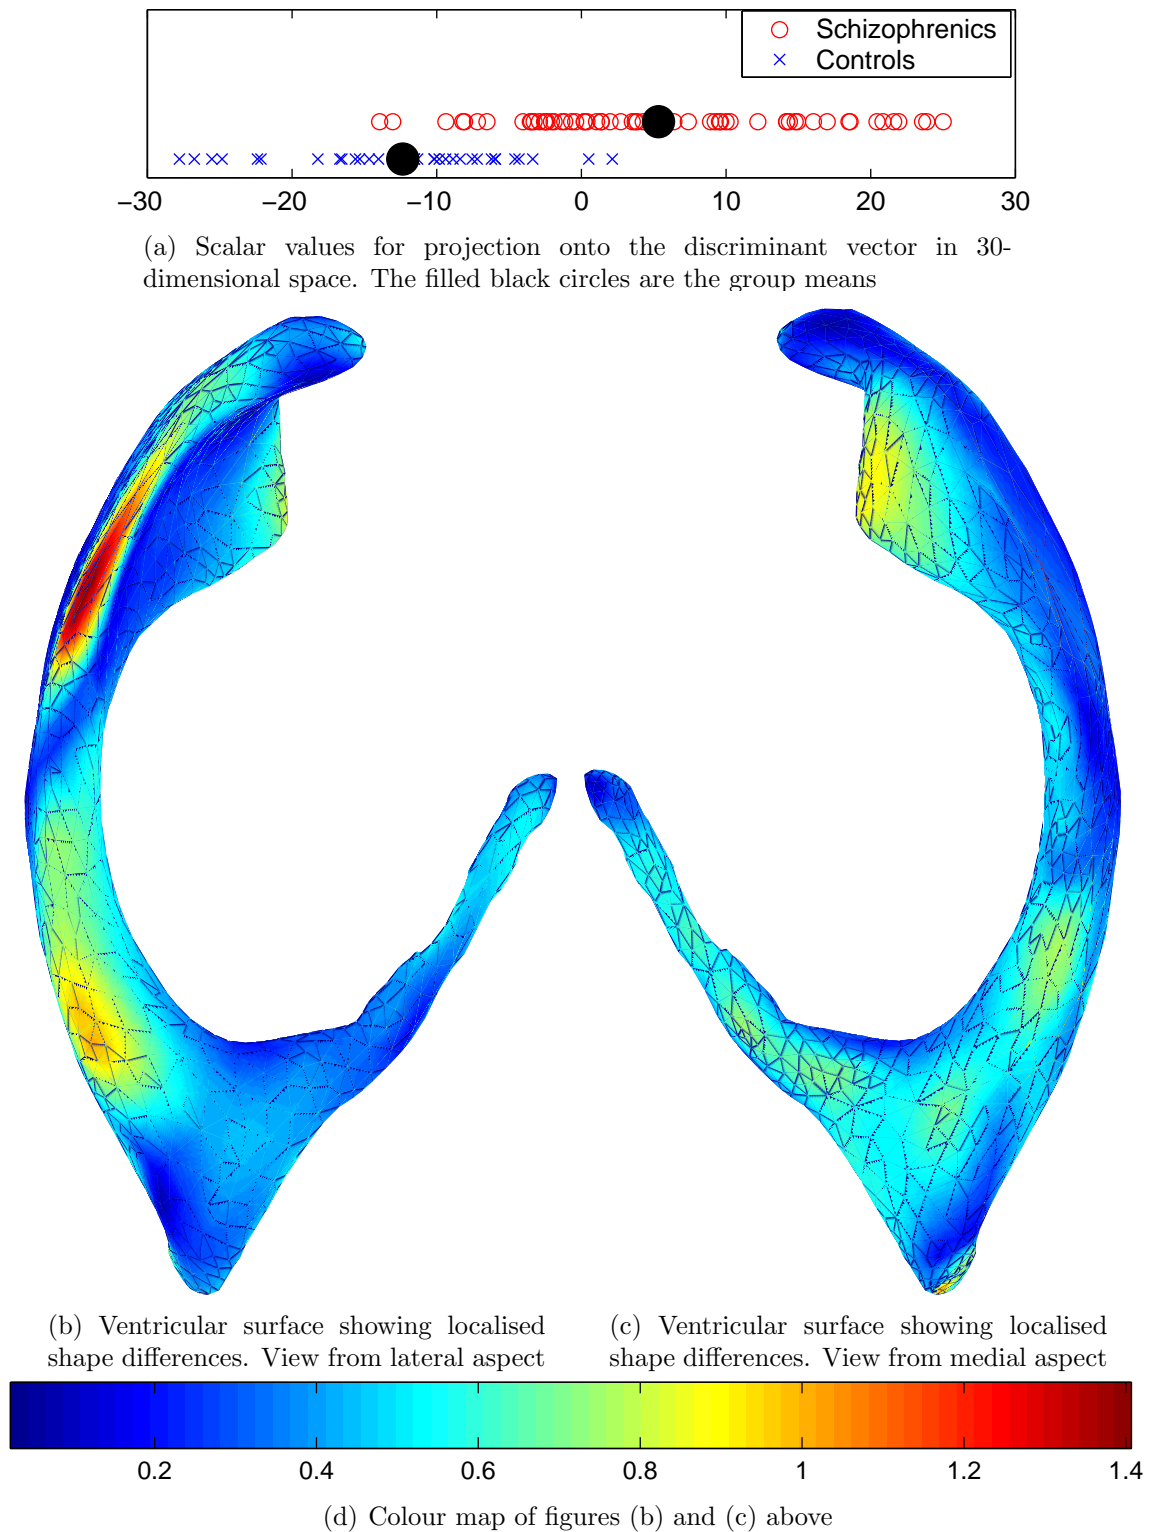
Figures 5.4 to 5.6 show the results for each of the three shape comparisons. The scalar values of the projections of the ventricles onto the discriminant vector in the 30-dimensional shape space are shown in (a) of each figure. The difference between ventricles at the schizophrenic mean and the control mean of the projections onto the discriminant vector are shown in (b) and (c). In each case sub-figure (b) is the view of the ventricle from the lateral aspect - i.e. the midline of the brain. Sub-figure (c) is the view from the medial aspect - i.e. from the side of the head. Figure (d) in each case shows the colour-map of the normalised shape distance. The same colour-map is used in the results of comparisons 1 to 3 as it is the relative magnitudes of the shape differences that interest us. AVI files of animations showing the change in the shape of the ventricular surface between the schizophrenic and control extremes along the discriminant vector are included on the accompanying CD.

### 5.7.2 Differences in Extent of Asymmetry within Schizophrenic and Control Groups

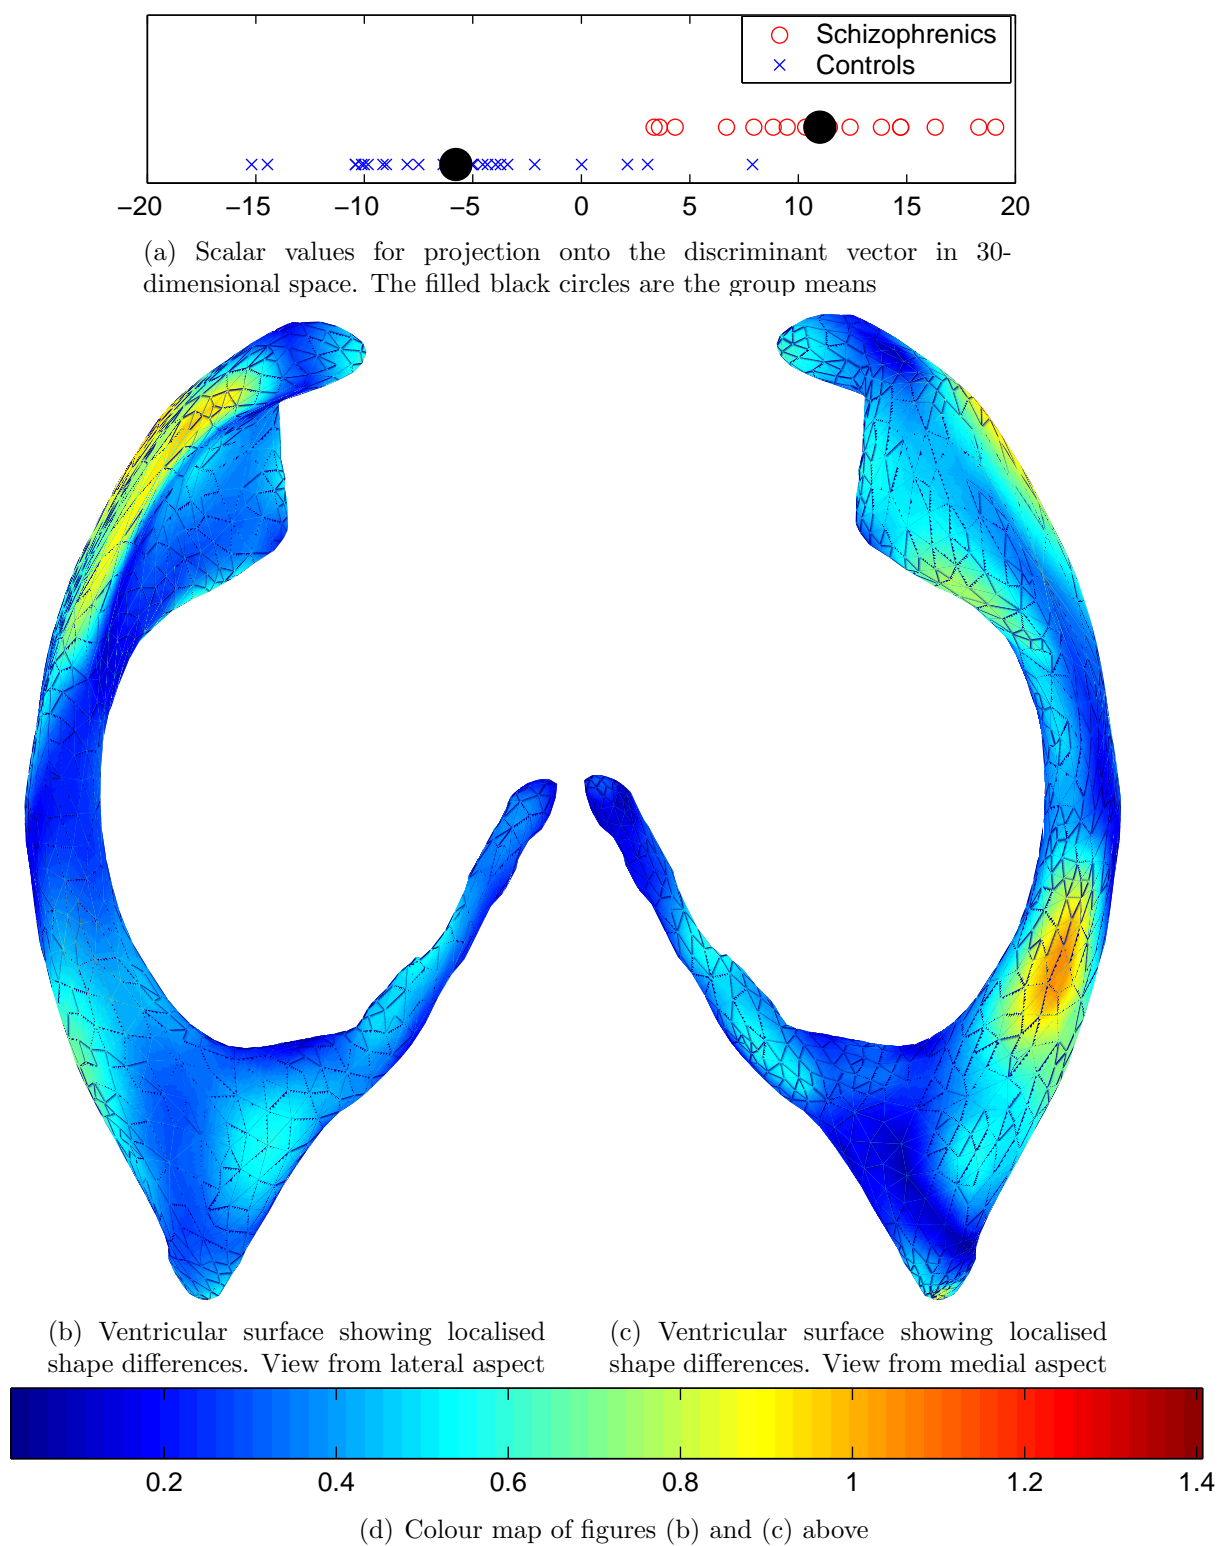
Figures 5.7, 5.9 and 5.11 show the results of the paired LDA for the three comparisons in this section. In each case sub-figure (a) shows two plots. The first is that of the scalar values of the left-right pairs of the schizophrenic group projected onto the discriminant vector. The second is a similar plot for the control group. Sub-figure (b) is also a plot of the projections onto the discriminant vector. However, in this case the intention is to highlight the relationship between the scalar values of left-right pairs on the discriminant vector. The vertical distance between the pairs was added simply to aid the clarity of presentation.



**Figure 5.4:** Results of Comparison 1 - Shape differences between all schizophrenic subjects and all control subjects



**Figure 5.5:** Results of Comparison 2 - Shape differences between all male schizophrenic subjects and all male control subjects



**Figure 5.6:** Results of Comparison 3 - Shape differences between all female schizophrenic subjects and all female control subjects



	Schizophrenic mean	Control Mean	$\overline{D}_{shape}$	$p_{shape}^{\dagger}$
Comparison 1	5.89 (9.41)	-7.66 (10.62)	13.54	$7.35 \times 10^{-13}$
Comparison 2	5.34 (9.94)	-12.34 (7.54)	17.68	$2.87 \times 10^{-14}$
Comparison 3	11.00 (4.72)	-5.78 (5.23)	16.78	$8.30 \times 10^{-14}$

**Table 5.1:** Summary of results of shape comparisons. The mean values for the scalar result of projection of members of the schizophrenic and control groups onto the discriminant vector is given for comparison 1 (both males and females), comparison 2 (males only) and comparison 3 (females only).  $\overline{D}_{schiz}$  is the mean difference between the scalar projections of the schizophrenic group and the control group (standard deviations in parenthesis), and  $p_{shape}$  is the probability associated a t-test between the two groups.  $^{\dagger}$  The  $p$ -values are optimistically biased as stated in section 5.1

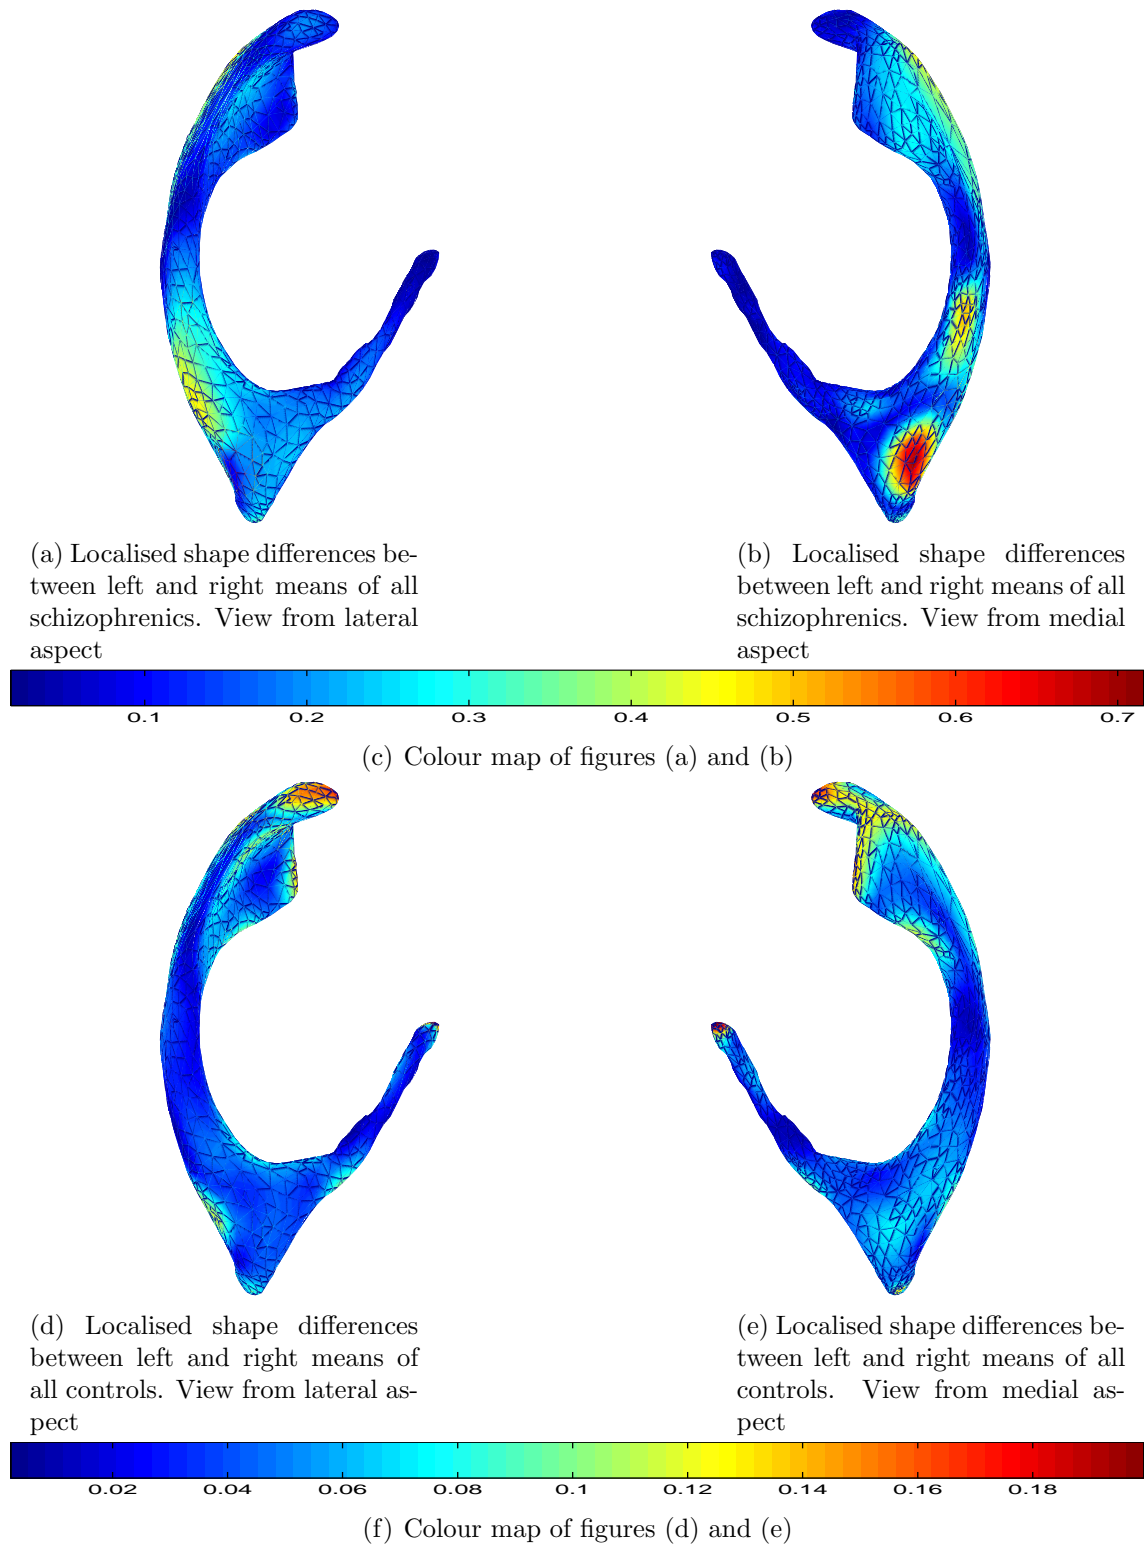
	$\overline{D}_{schiz}$	$p_{schiz}$	$\overline{D}_{cont}$	$p_{cont}^{\dagger}$
Comparison 4	7.94 (2.38)	$2.16 \times 10^{-22}$	2.49 (0.33)	$3.55 \times 10^{-27}$
Comparison 5	3.59 (1.07)	$1.69 \times 10^{-17}$	19.42 (11.23)	$2.39 \times 10^{-6}$
Comparison 6	23.15 (19.65)	$1.61 \times 10^{-02}$	32.45(32.03)	$1.13 \times 10^{-03}$

**Table 5.2:** Pair-wise  $t$ -tests of differences in left-right pairs.  $\overline{D}_{schiz}$  is the mean difference between the scalar projections of the left-right pairs within the schizophrenic group (standard deviations in parenthesis), and  $p_{schiz}$  is the probability associated with the pair-wise t-test of the scalar values.  $\overline{D}_{cont}$  and  $p_{cont}$  are corresponding values for the control groups.  $^{\dagger}$  The  $p$ -values are optimistically biased as stated in section 5.1

Figures 5.8, 5.10 and 5.12 show the localisation of shape asymmetry between left-right pairs. Sub-figures (a) and (b) are views for the schizophrenic group from the lateral aspect and the medial aspect respectively. Sub-figure (c) is the colour-map for (a) and (b). Sub-figures (d), (e) and (f) give similar results for the control group. The same colour-map is not used in all the figures because the magnitude of the shape differences varies significantly for the different sub-groups. Table 5.2 shows the values for the pair-wise  $t$ -test of the scalar values of the left and right ventricles for comparisons 4, 5 and 6. Comparison 4 examines differences in extent of asymmetry between left-right pairs of all control subjects and left-right pairs of all schizophrenic subjects. Comparison 5 does the same for the male subgroup, and comparison 6 for the female subgroup. Table 5.3 summarises the quantitative comparison of differences in extent of asymmetry between the schizophrenic and control subjects in each of the groupings.




**Figure 5.7:** Results of Comparison 4 - Scalar values for projections of left and right pairs of all schizophrenic and all control subjects




**Figure 5.8:** Results of Comparison 4 - Colour-mapped ventricular surface of shape differences between means of left and right pairs of all schizophrenic and all control subjects after projection onto the discriminant vector



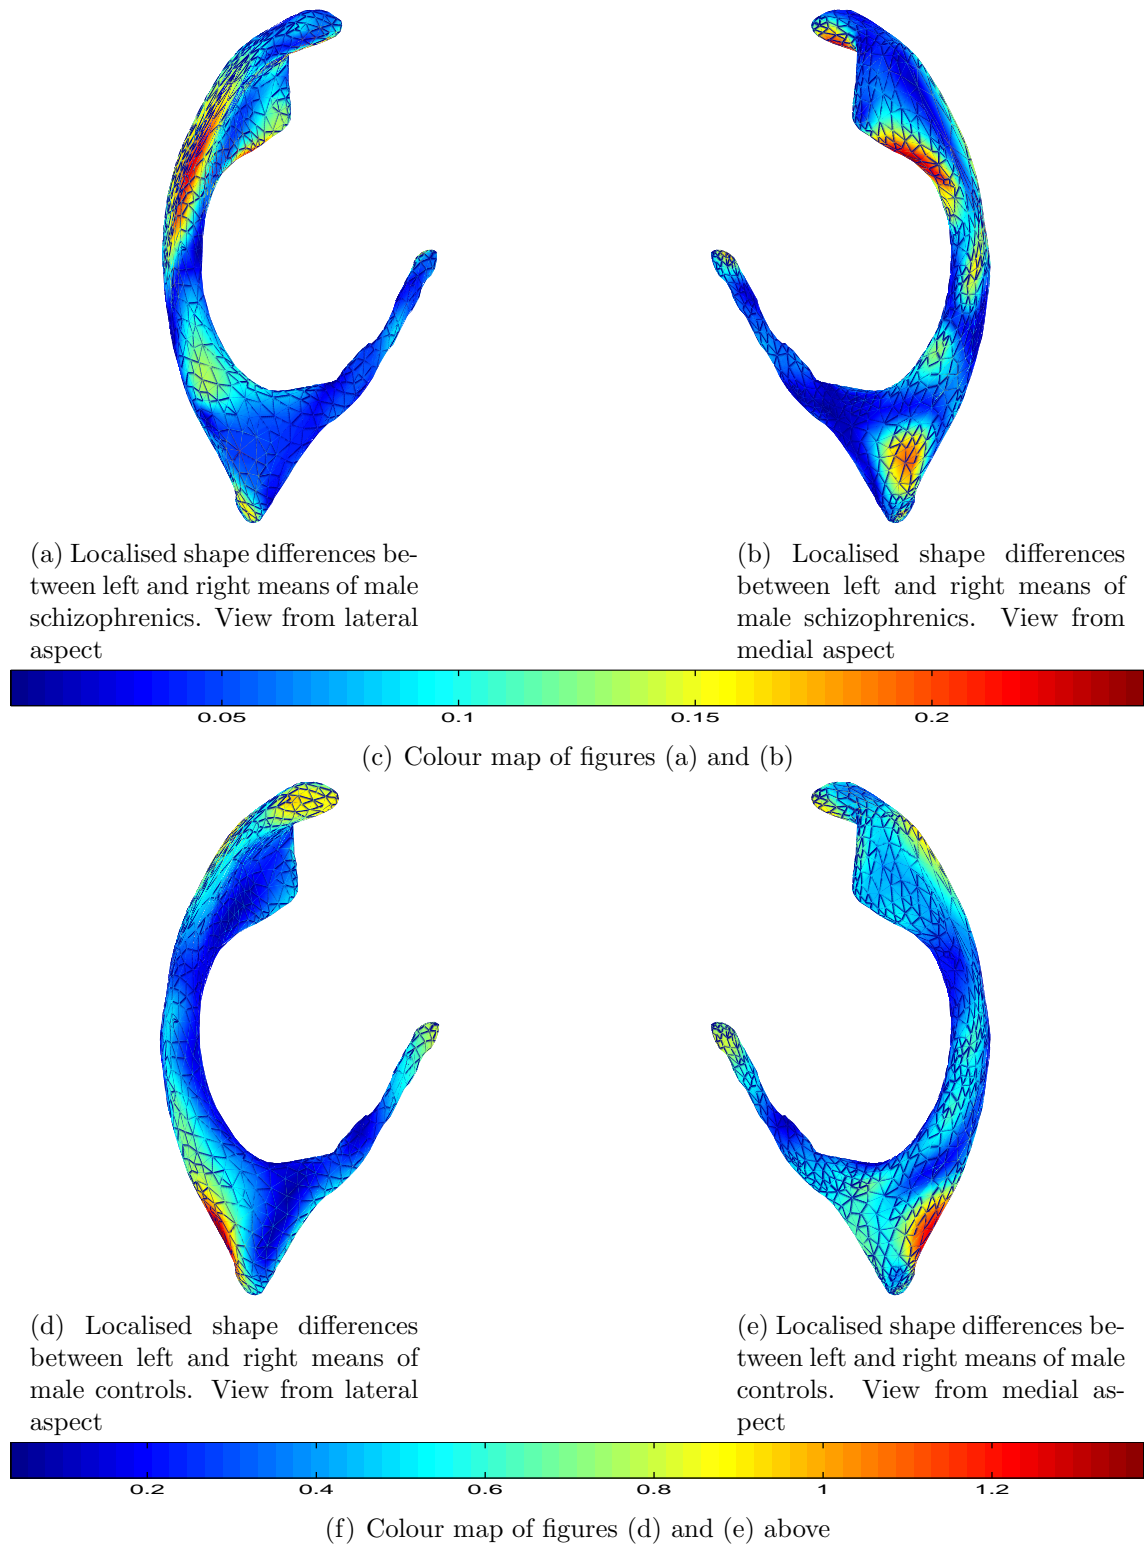


(a) Scalar values for projection onto the discriminant vector.




(b) Scalar values for projections showing relationship of values of left and right counterparts


**Figure 5.9:** Results of Comparison 5 - Scalar values for projections of left and right pairs of male schizophrenic and male control subjects onto the discriminant vector



**Figure 5.10:** Results of Comparison 5 - Colour-mapped ventricular surface of shape differences between means of left and right pairs of male schizophrenic and male control subjects after projection onto the discriminant vector

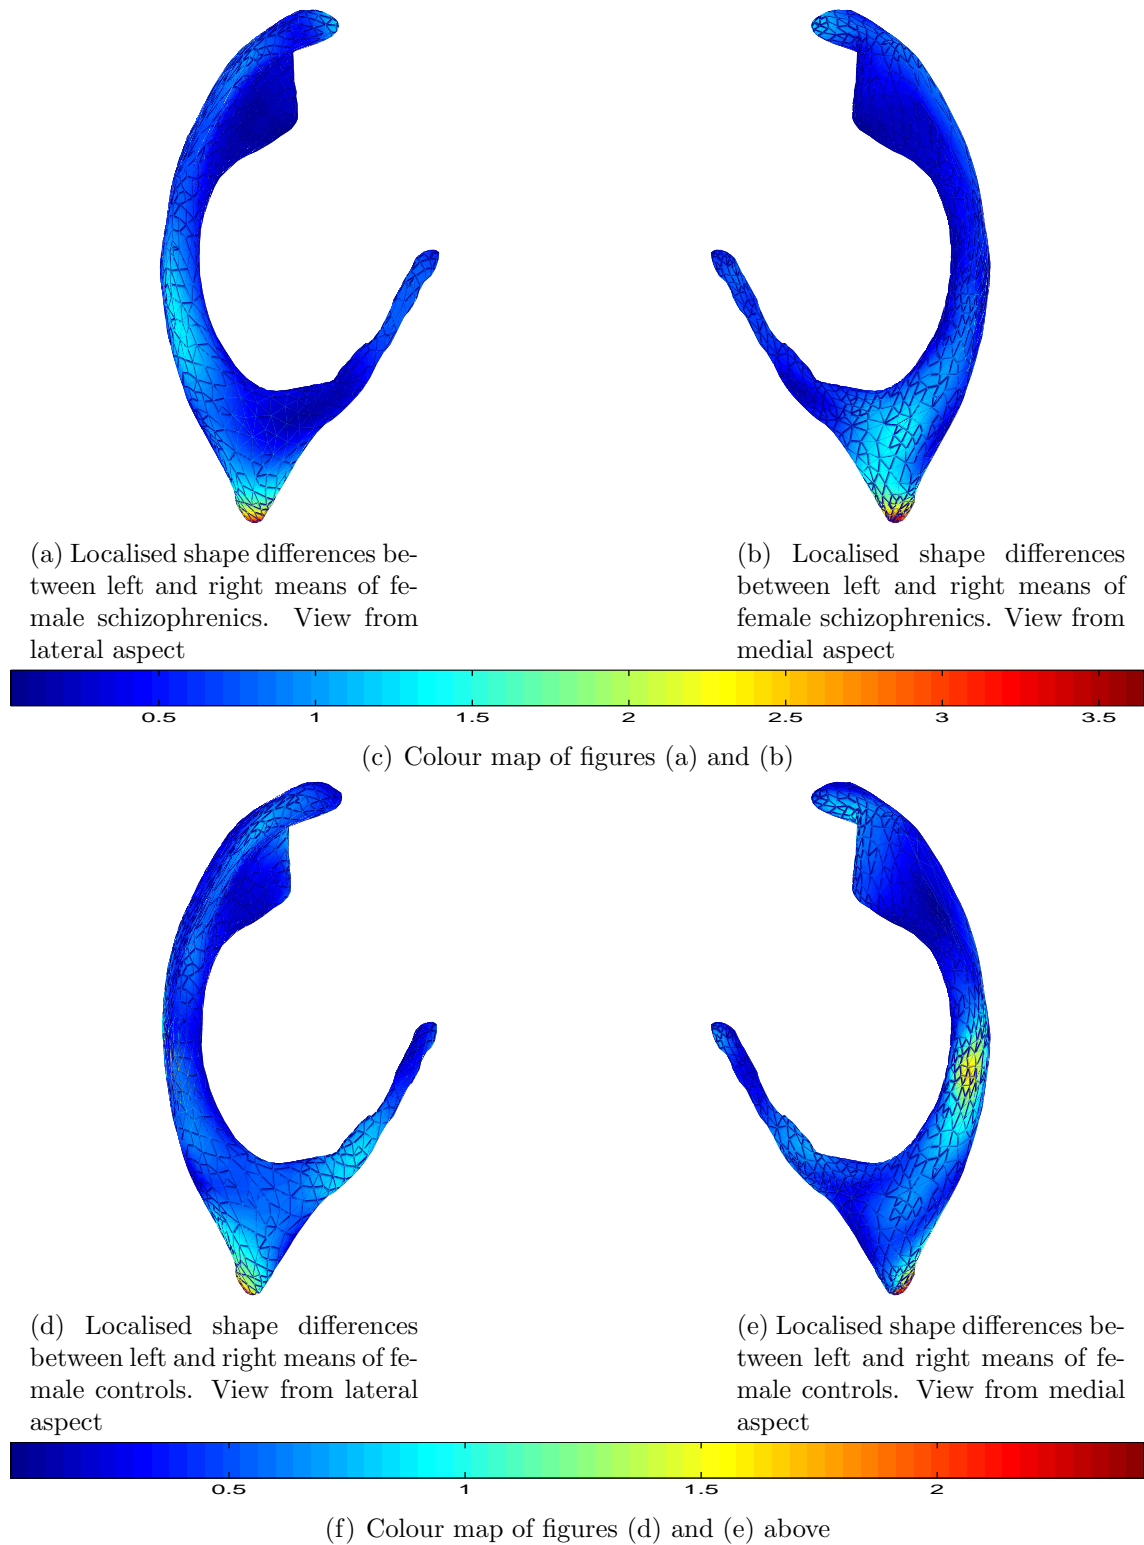


(a) Scalar values for projection onto the discriminant vector.



(b) Scalar values for projections showing relationship of values of left and right counterparts

**Figure 5.11:** Results of Comparison 6 - Scalar values for projections of left and right pairs of female schizophrenic and female control subjects onto the discriminant vector



**Figure 5.12:** Results of Comparison 6 - Colour-mapped ventricular surface of shape differences between means of left and right pairs of female schizophrenic and female control subjects after projection onto the discriminant vector

	$\overline{D}_{schiz}$	$\overline{D}_{cont}$	$p^\dagger$	Dot Product  (angle)
Comparison 4	7.94 (2.38)	2.49 (0.33)	$4.5 \times 10^{-19}$	0.18 ( $80^0$ )
Comparison 5	3.59 (1.07)	19.42 (11.23)	$8.7 \times 10^{-10}$	0.26 ( $105^0$ )
Comparison 6	23.15 (19.65)	32.45(32.03)	0.41	0.43 ( $66^0$ )

**Table 5.3:** Summary of the quantitative results of the asymmetry analysis.  $\overline{D}_{schiz}$  is the mean difference between the scalar projections of the left-right pairs within the schizophrenic group (standard deviations in parenthesis).  $\overline{D}_{cont}$  is the mean for the control group.  $p$  is the value of the  $t$ -test comparing  $D_{schiz}$  and  $D_{cont}$ . The last column contains the magnitude of the dot product of the unit vectors in the direction of the schizophrenic and control discriminant vectors, and the angle between them is given in brackets.  $^\dagger$  The  $p$ -values are optimistically biased as stated in section 5.1

## 5.8 Discussion

### 5.8.1 Discussion of Results of Shape Comparisons

Comparison 1 shows that significant shape differences exist between the schizophrenic and control groups. The means of the scalar values of projection of ventricles from each group onto the discriminant vector are significantly different ( $p < 0.0005$  by two tailed t-test). Significant shape differences also exist between the subgroups of schizophrenic males and control males, and schizophrenic females and control females with  $p < 0.0005$  in both cases. For all three comparisons the localisation of the shape differences was presented using colour-mapped ventricular surfaces.

For comparison 1 the shape differences were greatest at the lateral aspect of the anterior horn near the caudate nucleus, the region of the medial aspect of the anterior horn adjacent to the corpus callosum, the medial aspect of the lower part of the central body of the ventricle, and parts of the inferior (temporal) horn.

For comparison 2, the greatest shape differences were observed in the anterior horn and in the temporal horn as in comparison 1. However, the magnitude of the shape differences in the anterior horn were greater when only the male subgroups of schizophrenics and controls were compared. Shape differences were also present in

the medial aspect of the lower part of the central body of the ventricle, however their magnitude was much less than that of comparison 1. There were shape differences in the lateral aspect of the lower part of central body of the ventricle which were not observed in comparison 1. The shape differences between schizophrenics and controls in the female subgroup (comparison 3) were more similar to those of the combined male and female (comparison 1), than to those of the male subgroup. Shape differences were present in the anterior and temporal horns, and also in the medial aspect of the lower part of the central body of the ventricle. The differences in the lateral aspect of the anterior horn were over a larger local area in the case of the female subgroup than the combined male and female group and the male subgroup. There was no noticeable difference in the lateral aspect of the lower part of the central body as observed in the male subgroup.

Table 5.1 summarizes the differences in the scalar values of the projections onto the discriminant vector for comparison 1 (both males and females), comparison 2 (males only), and comparison 3 (females only). values of the differences in the means of the projections onto the discriminant vector for the schizophrenic and control members in each of the comparisons is of the same order (between 13 and 17). The significance of the differences in the means is high for all three comparisons ( $p \ll 10^{-12}$ ). However, the maximum difference was observed in the male only group (comparison 2), followed by the female only group (comparison 3). Narr et. al. [87] have also noted that shape differences between male schizophrenics and male controls were greater than those between female schizophrenics and female controls. Buckley et. al. reported that no significant shape differences were observed when comparing the lateral ventricles of a combined group of male and female schizophrenics. However, within the male subgroup significant shape differences were observed in the temporal horn.

The differences in the anterior horn and the temporal horn are in regions thought to be associated with schizophrenia. The corpus callosum connects the two hemispheres of the brain and a number of MRI studies e.g. [133] have reported abnormalities in this structure in the brains of schizophrenics.

Schizophrenia frequently involves cognitive and behavioural impairment, and these are often associated with frontal lobe damage [89]. The anterior horn extends into the frontal lobe of the brain. Some functional studies of the brains of schizophrenic subjects using positron emission tomography (PET) and functional magnetic resonance imaging (fMRI) have shown abnormalities in the prefrontal cortex, which occupies the anterior portions of the frontal lobe. Some structural MRI studies have also reported abnormalities in the Frontal lobe. However, there have also been studies reporting no significant differences in this area between schizophrenic and control subjects. A review of studies of the involvement of the frontal lobe in schizophrenia is given in [130] .

The shape differences observed in the temporal horn are also interesting. These were observed in the region around the hippocampus. Changes in the function of the hippocampus are thought to be important in the symptoms and aetiology of schizophrenia.

The shape differences in the lower part of the central body of the ventricle could not be associated with published schizophrenic findings. We also did not observe significant shape differences at the tip of the temporal horn, which is in the region of the amygdala. The amygdala is closely associated with the hippocampus, and also of importance in the pathology of schizophrenia. Shape differences in the frontal horn were reported by Narr et. al. [87] in their morphometric analysis of schizophrenia. However, they also reported significant shape differences in the posterior horn which we have not observed.

### 5.8.2 Discussion of Asymmetry Results

The results of the asymmetry comparisons show that there is a significant degree of asymmetry in left and right ventricle pairs within both the schizophrenic and control groups. For comparisons 4 and 5 (all subjects and males subgroup respectively), pair-wise  $t$ -tests of the scalar values of the projections onto the discriminant vector (Table

5.2) showed  $p \ll 0.0005$ . For the the female schizophrenic subgroup (comparison 6)  $p = 0.016$  and for the female control subgroup,  $p = 0.0011$ . Sub-figure (b) of figures 5.7, 5.9 and 5.11 shows that in almost all cases the scalar value of the projection of the right ventricle onto the discriminant vector was greater than that of its left counterpart. This supports the assumption depicted in the hypothetical example of figure 5.3, that there is an underlying trend in the separation of left and right pairs in the shape space.

Table 5.3 shows that the magnitude of the asymmetry within the schizophrenic group was significantly different to that within the control group for comparisons 4 and 5 ( $p \ll 0.0005$ ). The pattern of asymmetry for the control group was also different from that of the schizophrenic group for comparisons 4 and 5. This was reflected quantitatively by the magnitude of the dot products of the discriminant values. The fact that the patterns of asymmetry are different is also illustrated qualitatively in Figures 5.8 and 5.10. The asymmetry in the projections of the means of the left-right pairs of all schizophrenic subjects on the discriminant vector (5.8 (a) and (b)) was greatest in the lower part of the main body of the ventricle and in the anterior horn. In the case of all control subjects (5.8(d) and (e)), asymmetry was greatest at the tip of the temporal horn and the tip of the anterior horn. The maximum of the magnitude of the asymmetry in the schizophrenic case was over three times as much as that in the control case (0.72 vs 0.20 normalised shape units).

When only the male subgroup was considered (comparison 5), the extent of asymmetry in the control group was greater than that in the schizophrenic group - a reversal of the observation in comparison 4. However, as in comparison 4, the localisation of asymmetry was different in the control and schizophrenic groups. For the schizophrenic group asymmetry was greatest in the lateral aspect of the anterior horn in a similar region to that where shape differences were observed between the controls and schizophrenic males. Asymmetry was also present in lower part of the main body of the ventricle, the medial aspect of the anterior horn, the occipital horn, and the tip of the temporal horn. For the male control group the greatest asymmetry



was observed around the tip of the anterior horn, in the lower part of the central body of the ventricle, and in the temporal horn.

Table 5.3 shows that the extent of the asymmetry observed in the female subgroup (comparison 6) was not significantly different between the control and schizophrenic groups ( $p = 0.41$ ). The localisation of the asymmetry was also similar for the controls and the schizophrenic females. This was mainly in the occipital horn, although in the control group there was also some asymmetry in the middle part of the medial aspect of the main body of the ventricle. This similarity is also shown quantitatively in Table 5.3, as the magnitude of the scalar product for this sub-group (0.43) is greater than for the combined male and female group, and the male only sub-group. Figures 5.11 shows that there is much higher variability between the left and right pairs of the female subgroup, particularly within the schizophrenic subjects. This is most probably due to the small number of subjects in these groups (18 female schizophrenics and 26 female controls, compared with 60 male schizophrenics and 34 male controls).

In conclusion, the asymmetry results indicate that there a difference in the extent of asymmetry of the shape of the ventricles for schizophrenic subjects and for control subjects. Our results showed that when the pooled groups of males and females are compared (comparison 4) the magnitude of asymmetry is greater in schizophrenics than in controls, and that the nature of the asymmetry in both groups is different. However, comparison of the male subgroup showed that the magnitude of the asymmetry in controls was greater than that in schizophrenics, and again the nature of asymmetry in both groups was different. The comparison of the female subgroup showed that there was no significant difference in the magnitude of asymmetry between schizophrenics and controls. Furthermore the nature of the observed asymmetry was similar for schizophrenics and controls. However, the results for this subgroup may have been affected by the small sample size.

Crow's hypothesis was mentioned in section 1.1. This proposed that brains of schizophrenic

subjects are more symmetric than those of control subjects. The results of comparison 5 (male subgroup) support Crow's hypothesis as the magnitude of the asymmetry is greater in the male schizophrenics than in the male controls. On the other hand the results of comparison 4 (combined male and female group) do not support Crow's hypothesis as they indicate that the magnitude of asymmetry of the schizophrenics is greater than that of the controls. However, the most interesting aspect of our results are the suggestion that the nature of asymmetry in the schizophrenic group is different from that in the control group.

## 5.9 Summary

In this section we have outlined the construction of the PDM used to perform morphometric analysis on the lateral ventricles of the schizophrenic and control subjects. We explained the application of LDA to investigate the differences between the groups. The results showed that there were significant shape differences in areas of the lateral ventricles adjacent to structures in the brain implicated in schizophrenia. There was also a significant degree of asymmetry between the left and right ventricles within both the control and the schizophrenic groups. The pattern of asymmetry was similar for controls and schizophrenics in the female sub-group. However, in the male sub-group, and in the combined male and female group the pattern of asymmetry was different for controls and schizophrenics.

# Chapter 6

## Conclusions

### 6.1 Contributions of this Thesis

We have described a method of construction of 3D PDMs of the lateral ventricles using salient curvature-based physical landmarks. These were the crest lines of the ventricles obtained using the marching lines algorithm. The use of crest lines in constructing 3D PDMs was proposed by Subsol [108], and Andresen et. al. [3] have applied this to the construction of a 3D PDM of the human mandible.

However, our method differs from that of [3] in that we have used the transportation algorithm to address the problem of obtaining point-to-point correspondences. We have shown that this method gives better correspondences, and that the resulting PDMs capture the variation in the training set better than those obtained by correspondence using the closest points. The transportation approach we described can be applied to other point matching problems.

The application of PDMs to 3D problems is relatively new and their construction usually involves automatic location of homologous points on structurally significant parts of surface representations of each member of the class of object being modelled.

This is less clearly true in the MDL approach of Davies et. al. [39] since they never relate the optimal points to structural features on the surface. In the case of the lateral ventricles, 3D PDMs of parts of the lateral ventricle have been constructed [22][38]. We are only aware of one group (Gerig et. al. [54]) that have constructed a comprehensive 3D PDM of this structure. However, in that case the corresponding points were simply distributed uniformly across the surface (taking no account of structural features). Furthermore, the entire ventricle surface was not modelled, the examples shown in [54] do not include the temporal horns and the posterior horns. Therefore, the PDM described here is the first true PDM of the entirety of the lateral ventricles and our method of construction has explicitly incorporated normalisation for brain shape, as opposed to other methods that apply global normalisation based on brain size.

We have performed extensive quantitative analysis of our PDM. We introduced the use of a 3D distance transform method for the quantitative comparison of the similarity of triangulated surfaces with voxel images. We have applied this to the comparison of the *corresponding training set* with the original training set to quantitatively measure the fidelity with which the original surfaces are represented in the building of the PDM. This comparison has not been reported for other methods of building 3D PDMs. Additionally, we quantitatively investigated the dependence of our model on the choice of template, and took steps to reduce this.

The approach we adopted to shape analysis gave both quantitative and qualitative indications of shape and asymmetry differences between the schizophrenic and control groups being compared. Our shape analysis showed that the ventricular shape of the schizophrenic and control groups are significantly different ( $p < 10^{-12}$ ). The shape differences observed were localised to four main regions. Three of these are close to structures thought to be implicated in schizophrenia: the temporal horn (its tip near the amygdala, and along its body near the parahippocampal fissure), the central part of the lateral ventricles around the caudate nucleus, and the tip of the anterior horn in the region of the frontal lobe and near the corpus callosum.

Investigations of asymmetry between left and right ventricle pairs showed that a degree of asymmetry in shape existed in both schizophrenics and controls. The extent of asymmetry in the combined group of male and female subjects was greater in schizophrenics than in controls. However, when the male subgroup was examined the degree of asymmetry in controls was greater - which would be expected according to the “Crow hypothesis”. For the female subgroup the degree of asymmetry was not significantly different for schizophrenics and controls. This may have been because of the small number of subjects within this subgroup. Our analysis also showed that the nature of the asymmetry in the combined male and female group and in the male only subgroup was different for schizophrenics and controls. This is interesting as it may suggest that it is not the degree of asymmetry alone that changes with schizophrenia, but the nature of the actual asymmetry. The PDM model can be used to investigate this further as it allows instantiation of surfaces that gives insight into the physical nature of the differences. These results contribute to the sparse (but growing) results of differences in the lateral ventricles of schizophrenic and control subjects provided by 3D morphometric techniques.

## 6.2 Further Work

The dependence of the model on the choice of template needs to be investigated further. It is possible that the use of a better method of non-rigid registration in the step involving deformation of the template to each member of the training set might reduce the effect. One possible solution may arise from work being carried out at ISBE on clamped-plate splines [82], being extended to the case of approximating spline functions.

The crest-line approach to the construction of 3D PDMs is application specific in that it can only be applied in cases where consistent salient curvature landmarks exist e.g the ventricles and the skull. There are three other prominent methods of constructing SSMs - the MDL based approach of Davies et.al [39], the combined SPHARM and

medial transform approach of the Styner and Gerig [107], and the high-dimensionality transformation method of Thompson and Toga used in [87]. Comparison of these methods against each other, and indeed methods of comparison of different SSMs in general are desirable. This will allow the effectiveness of different methods to be compared - it may be that application-specific methods such as the present one are better than general methods in certain instances, or the converse might be the case. The form of this comparison may be in applying different methods to the same data set to see if similar results are produced.

PDMs can also be used in model based segmentation. As indicated in chapter 2, segmentation is usually a bottleneck in the analysis of medical images, and methods of automating segmentation are desirable. The 3D PDM of the ventricle could be used in the construction of a 3D active shape model (ASM) which can be applied to segmentation.

### 6.3 Final Statement

3D Shape based morphology is an active area of research. We have presented a principled method of quantitative shape analysis. In the case of schizophrenia we do not claim to have a method that allows automatic categorisation into control and schizophrenic groups, but one which will aid the understanding and monitoring of schizophrenia, and other neurological diseases.

# Appendix A

## Operations on the Triangulated Surfaces

Dealing with the 3D surfaces of the ventricles presented some problems that are faced in the computer graphics field, and for which solutions are available. This appendix gives details of the operations carried out in chapter ... where results from the computer graphics field that were used. Texts such as [34][88] give further details. However, the solutions used here were based on notes obtained from [109].

### A.1 Outward Normals to a Triangulated Surface

Triangulation of a surface involves having a list of coordinate points on the surface - the vertices, and a list of triplets of vertices defining triangular facets which tessellate the surface.

The vertices can be stored in a  $M \times 3$  matrix  $\mathbf{V}$ . The number of rows of the matrix,  $M$ , gives the number of vertices, and there are three columns because we are dealing with 3D space. Each row in the matrix represents the triplet of real numbers  $V$  giving the 3D coordinates of a particular vertex.

The facets are represented by a  $N \times 3$  Matrix,  $\mathbf{F}$ . Each row of this matrix stores three indices to the rows in  $\mathbf{V}$  containing the coordinates of the vertices defining a facet. We assume (quite rightly in practice) that the order of vertices in each triplet defining a facet is the same - i.e. applying the *right hand rule* to the vectors formed by the vertex edges gives consistent results for all facets. We therefore have a sense of outward and inwards for vectors on the surface.

The unit normal to a facet is given by the cross-products

$$\hat{\mathbf{n}}_f = \frac{\sum_{i=1}^2 (V_{i+1} - V_i) \otimes (V_{i+2} - V_{i+1})}{\|\sum_{i=1}^2 (V_{i+1} - V_i) \otimes (V_{i+2} - V_{i+1})\|} \quad (\text{A.1})$$

where  $i$  is the column in  $\mathbf{F}$  from which the index to a vertex is obtained. Note that when  $i = 2$ , the column referred to by  $i + 2$  is the first column.

An estimation to the unit normal to a vertex is obtained from the unit vectors of the neighbouring facets which share the vertex :

$$\hat{\mathbf{n}}_v = \frac{\sum_{\text{neighbours}} \hat{\mathbf{n}}_f}{\|\sum_{\text{neighbours}} \hat{\mathbf{n}}_f\|} \quad (\text{A.2})$$


## A.2 Intersection of a Triangulated Surface By a Line Segment

One of the steps in the construction of PDMs described in this thesis required projection of vertices from one surface onto another surface. Each vertex was projected in the direction of its unit normal vector obtained using Equation A.2.

The problem can be stated as follows : Given a vertex point  $P_0$ , and a direction  $\hat{\mathbf{n}}_v$  (the unit normal of the vertex), we wish to determine the point of intersection of a line segment through  $P_0$  in the direction  $\hat{\mathbf{n}}_v$ , with the triangular facet  $\mathbf{F}$  in a plane  $\pi$ .  $\mathbf{F}$  has vertices  $V_0$ ,  $V_1$ , and  $V_2$ , and unit normal vector  $\hat{\mathbf{n}}_f$ . Figure A.1 illustrates



this.



**Figure A.1:** A facet  $\mathbf{F}$  with vertices  $V_0$ ,  $V_1$ , and  $V_2$  lying in the plane  $\pi$ . The line segment from the point  $P_0$  in the direction  $\hat{\mathbf{n}}_v$ , intersects  $\mathbf{F}$  at  $P_x$ .

Firstly, we need to find the point  $P_x$ , at which the line segment from  $P_0$  in the direction  $\mathbf{n}_f$  intersects the plane  $\pi$ . This is obtained from the parametric equation of the line segment :

$$P_i = P_0 + \lambda \hat{\mathbf{n}}_v \quad (\text{A.3})$$

Where

$$\lambda = -\frac{\hat{\mathbf{n}}_f \cdot \mathbf{w}}{\hat{\mathbf{n}}_f \cdot \hat{\mathbf{n}}_v}$$

with  $\mathbf{w}$  being the vector in the plane  $\pi$  joining  $P_x$  and  $V_0$ , i.e.  $\mathbf{w} = (P_x - V_0)$ .

Secondly, we check if the intersection point,  $P_x$ , is in the triangle defined by the facet,  $\mathbf{F}$ . The parametric equation of  $P_x$  in the plane is :

$$V(s, t) = V_0 + s\mathbf{u} + t\mathbf{v} \quad (\text{A.4})$$

where  $s, t \in \mathbb{R}$ ,  $\mathbf{u} = (V_1 - V_0)$ , and  $\mathbf{v} = (V_2 - V_0)$  as shown in Figure A.1.

$P_x$  is in  $\mathbf{F}$  only if ALL of the follow three conditions are true,

$$\begin{aligned} s &\geq 0 \\ t &\geq 0 \\ s + t &\leq 1 \end{aligned} \tag{A.5}$$

[109] derive an expression allowing  $s$  and  $t$  to be computed using only scalar products :

$$s = \frac{(\mathbf{u} \cdot \mathbf{v})(\mathbf{w} \cdot \mathbf{v}) - (\mathbf{v} \cdot \mathbf{v})(\mathbf{w} \cdot \mathbf{u})}{(\mathbf{u} \cdot \mathbf{v})^2 - (\mathbf{u} \cdot \mathbf{u})(\mathbf{v} \cdot \mathbf{v})} \quad t = \frac{(\mathbf{u} \cdot \mathbf{v})(\mathbf{w} \cdot \mathbf{u}) - (\mathbf{u} \cdot \mathbf{u})(\mathbf{w} \cdot \mathbf{v})}{(\mathbf{u} \cdot \mathbf{v})^2 - (\mathbf{u} \cdot \mathbf{u})(\mathbf{v} \cdot \mathbf{v})} \tag{A.6}$$

The algorithm for the projection of the vertices of one surface  $\mathbf{S}_1$  onto another triangulated surface,  $\mathbf{S}_2$ , is :

1. For each vertex in  $\mathbf{S}_1$ 
  - (a) Find parametric equation of line segment in direction of vertex normal using Equation A.3
  - (b) Find position of intersection of the line segment with planes of ALL facets of  $\mathbf{S}_2$  by solving for  $\lambda$  in each case
  - (c) Find the facets for which the intersection points satisfy Equation A.5
2. Repeat above steps in the direction  $-\hat{\mathbf{n}}_v$  (as we do not know if the vertex is on the inside or the outside of surface  $\mathbf{S}_2$ )

If there is more than one candidate point on surface  $\mathbf{S}_2$  for a vertex on  $\mathbf{S}_1$ , we select the closest point whose facet normal is in the same direction as  $\hat{\mathbf{n}}_v$ . On the other hand, if no candidate points were returned - which happened in less than 5% of vertices, the vertex is projected onto the closest vertex on surface  $\mathbf{S}_2$ .

The algorithm outlined above seems exhaustive. However, a *vectorised* implementation in Matlab was used. This took less than two minutes for each pair of surfaces.

### A.3 Generating Voxel Representations of Triangulated Surfaces

To generate voxel representations of a triangulated surface, planes representing contiguous slices are used to intersect the surface at uniform intervals (see Figure A.2). The coordinates of the points at which the surface intersects the plane are obtained as a finely sampled contour using the algorithm outlined below. These coordinates are rounded to the nearest integer to get the indices of the surface voxels for that plane (slice).

#### Algorithm for voxelisation of a surface

1. Translate surface so all vertices have positive coordinate values
2. In the present case unit normal to plane is always in  $Z$ -direction, i.e.  $(0,0,1)$
3. Repeat for each plane,  $i := \text{floor}(\text{minimum } Z - 1) \text{ To } \text{ceil}(\text{maximum } Z) + 1$  In Unit integer steps
  - (a) for each facet determine if it intersects plane (by comparing  $Z$  coordinates of its three vertices with  $i$ )
  - (b) for each intersecting facet find the pair of points on the facet edges,  $P_1$  and  $P_2$  at which intersection with the plane occurs using the steps outlined in **A.2** above
  - (c) linearly interpolate between  $P_1$  and  $P_2$  to obtain a finely sampled contour



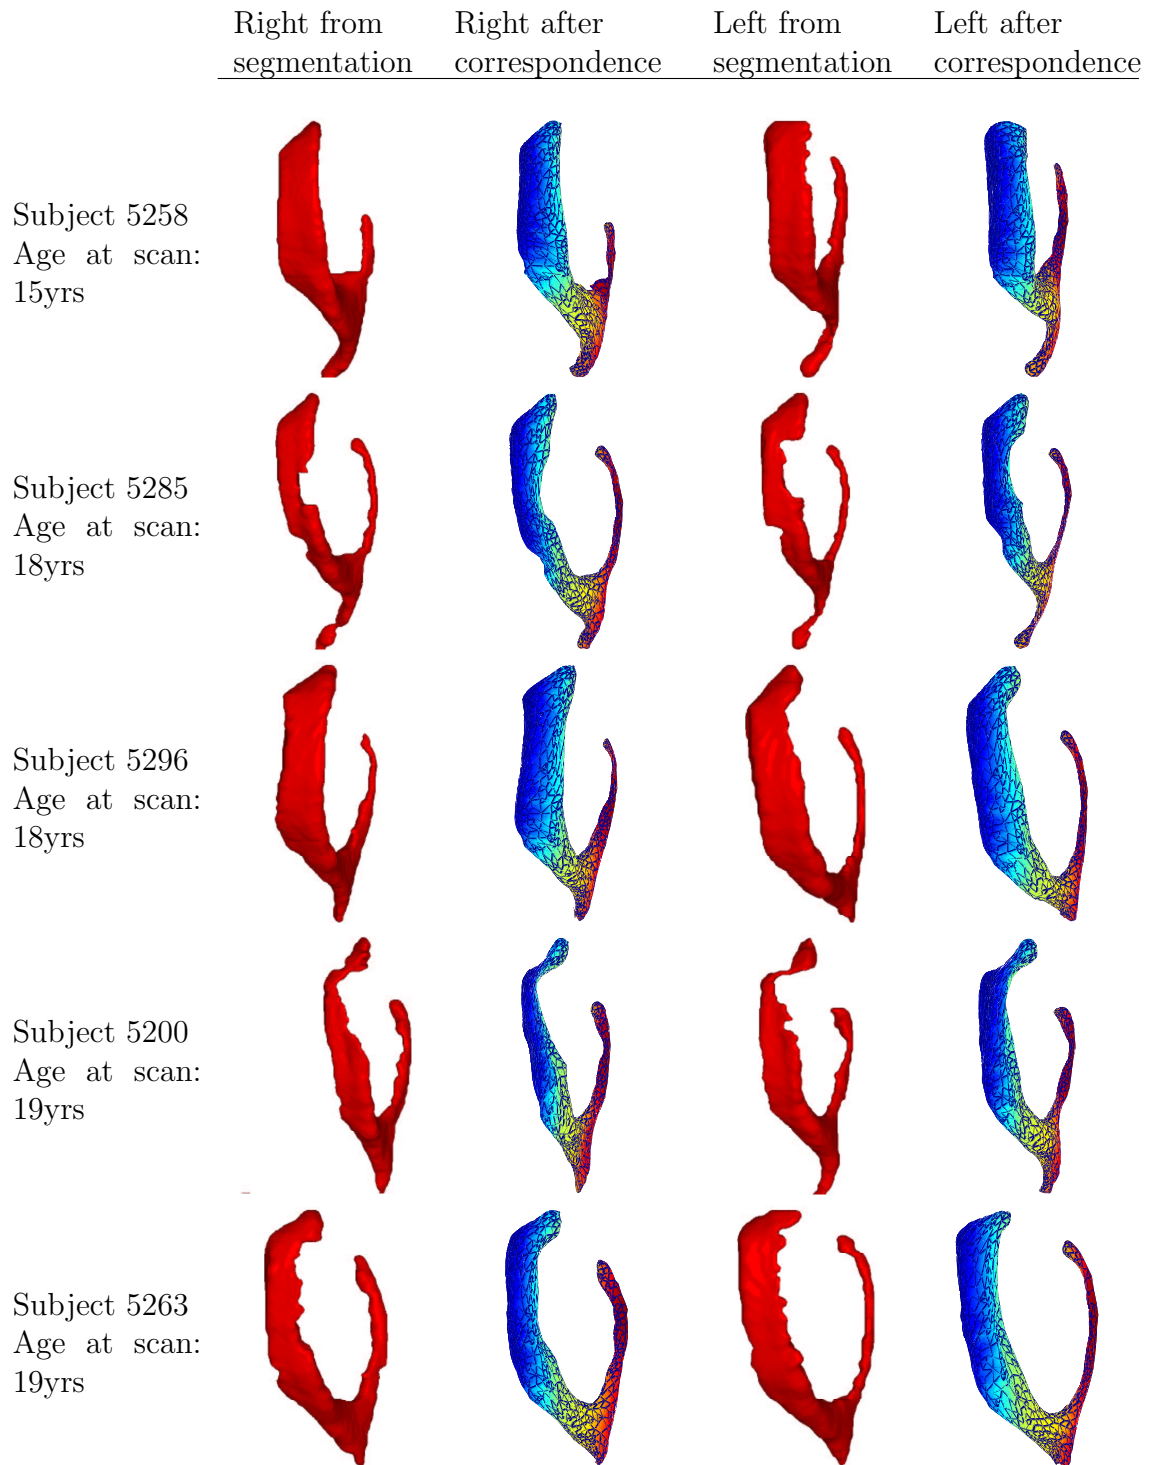
**Figure A.2:** An illustration of using planes to intersect the ventricular surface in creating a voxel representation of the triangulated ventricular surface

## Appendix B

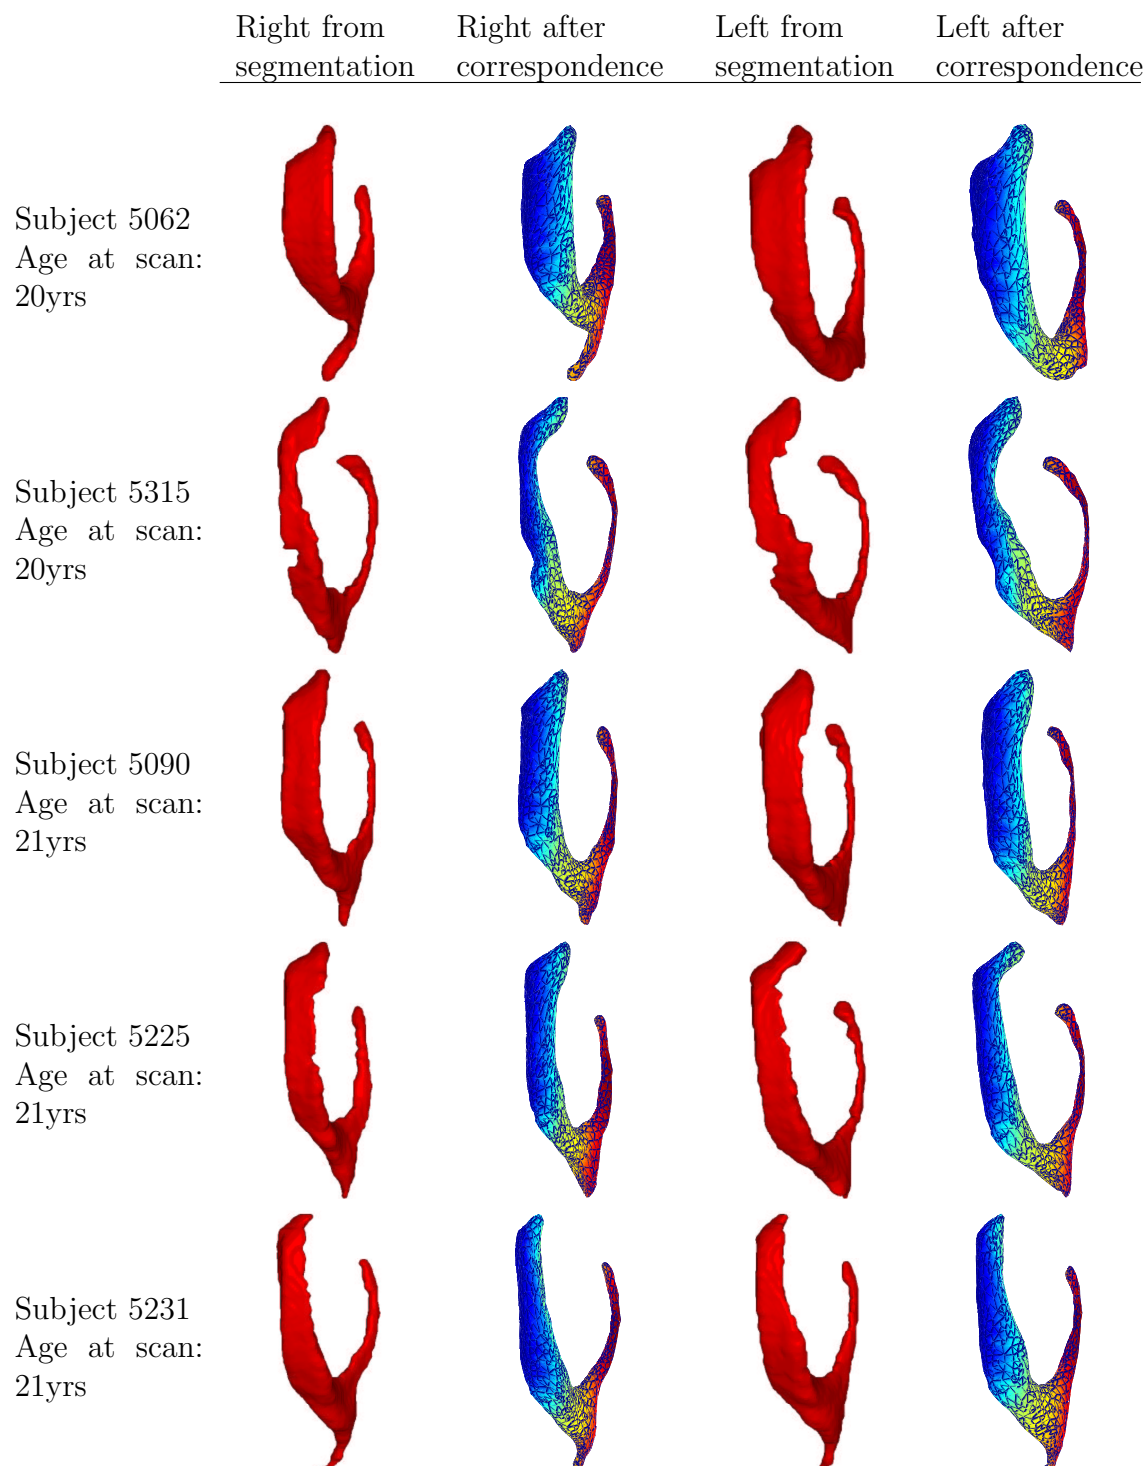
### Subjects Participating in the Study

Male Schizophrenics		Female Schizophrenics	
Subject Number	Age	Subject Number	Age
5258	15	5131	14
5285	18	5179	17
5296	18	5203	18
5200	19	5314	18
5263	19	5019	19
5062	20	5237	23
5315	20	5265	23
5090	21	5091	30
5225	21	5116	31
5231	21		
5266	21		
5274	21		
5002	22		
5087	22		
5160	22		
5077	23		
5054	24		
5280	24		
5148	25		
5256	26		
5025	27		
5197	27		
5283	27		
5284	27		
5141	28		
5276	28		
5233	30		
5171	33		
5140	36		
5175	48		

**Table B.1:** Age and sex data for schizophrenic subjects

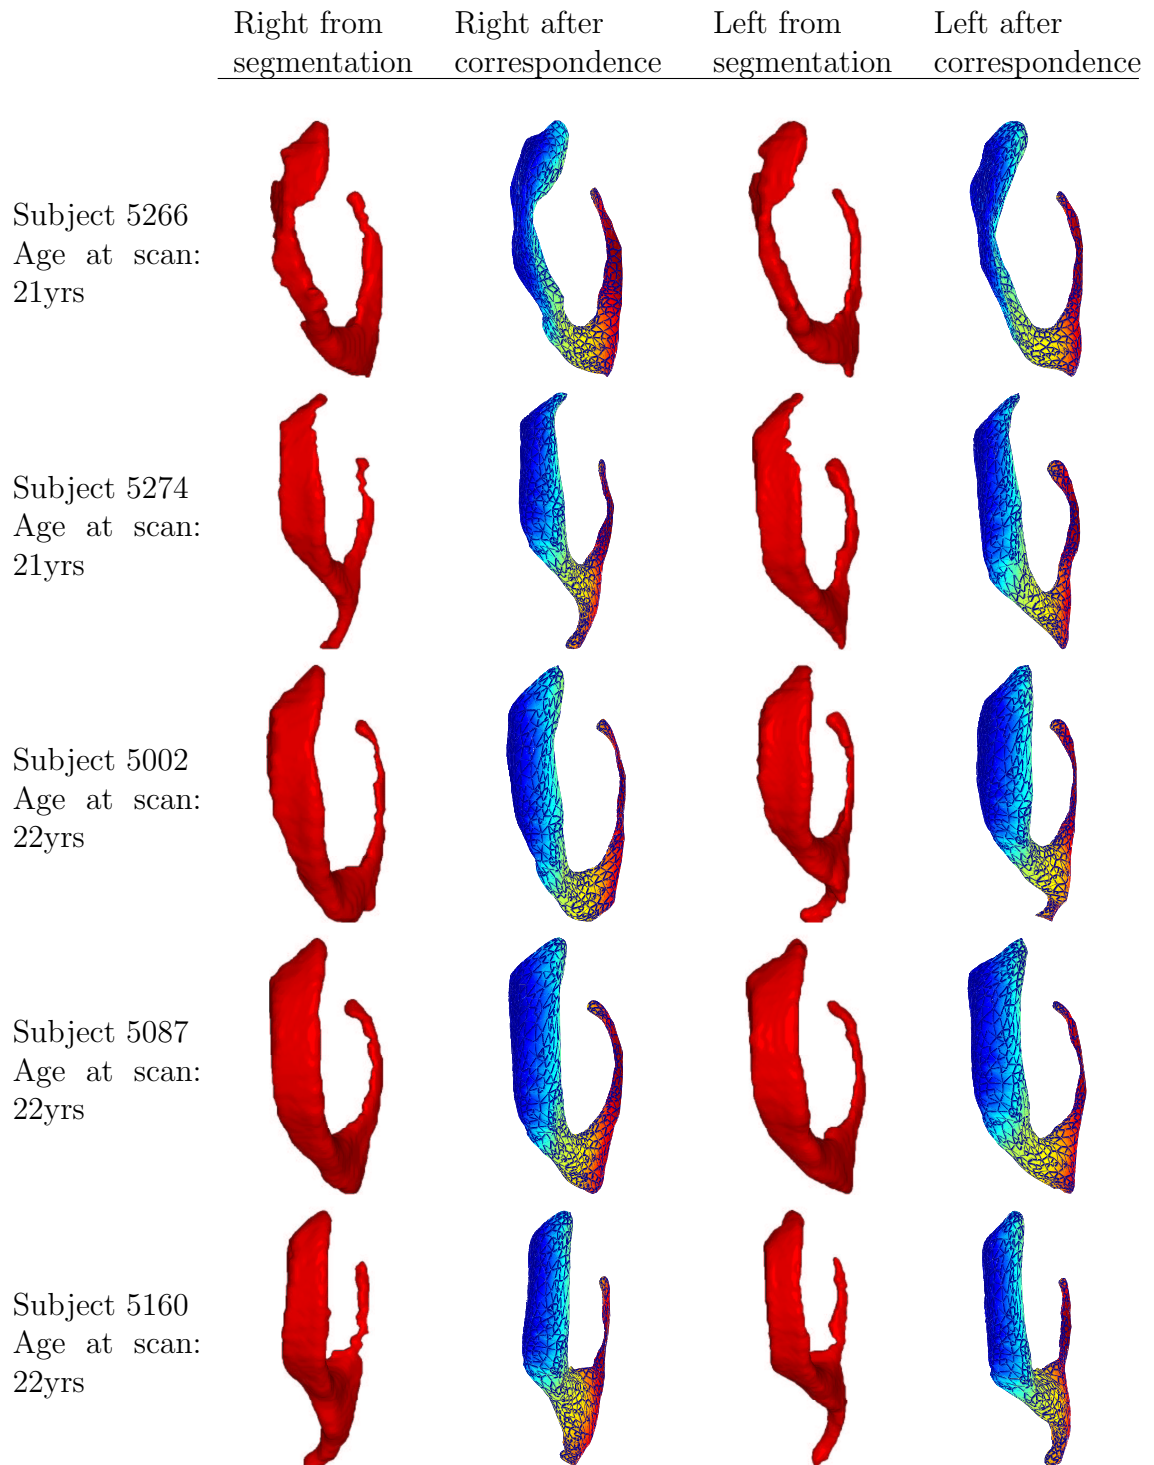


**Figure B.1:** Lateral ventricles of male schizophrenic subjects (1 to 5 of 30). The 1<sup>st</sup> column shows the surface of the segmented right ventricle of the specified subject. The 2<sup>nd</sup> column is the surface of the member of the training set used in construction of the PDM. The surface has been colour-coded such that corresponding parts of the surface have the same colour. The 3<sup>rd</sup> and 4<sup>th</sup> columns give respective versions for the left ventricle (reflected to give same pose as right) of the subject.

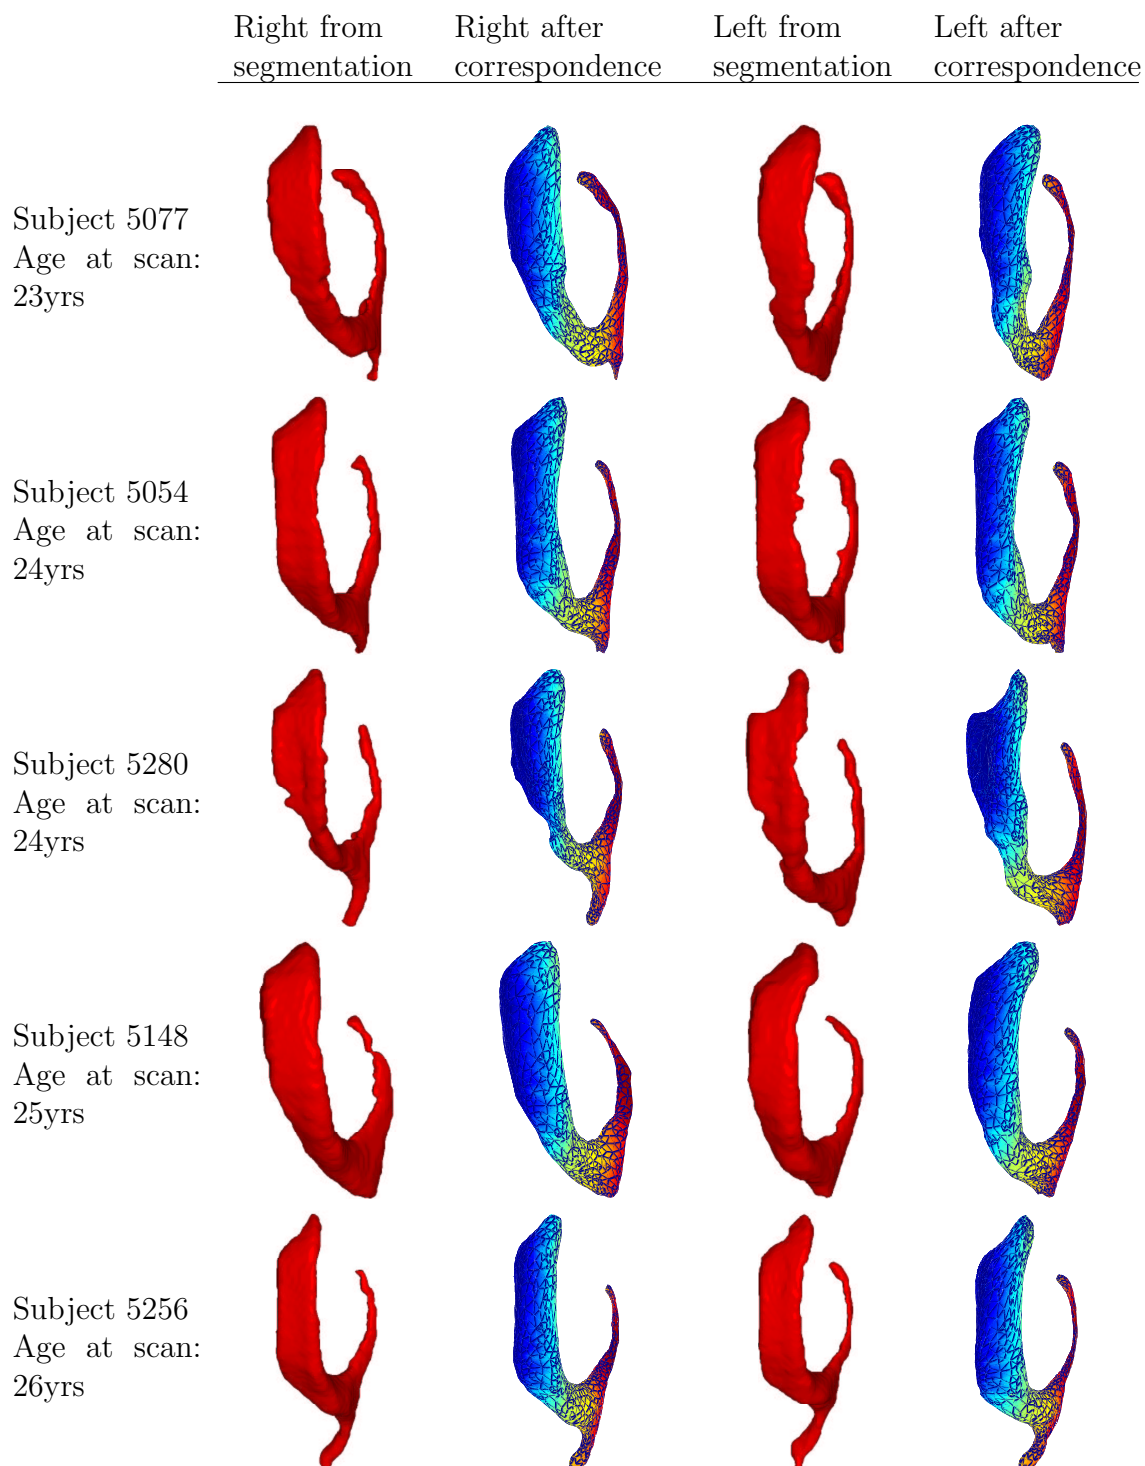


**Figure B.2:** Lateral ventricles of male schizophrenic subjects (6 to 10 of 30). The 1<sup>st</sup> column shows the surface of the segmented right ventricle of the specified subject. The 2<sup>nd</sup> column is the surface of the member of the training set used in construction of the PDM. The surface has been colour-coded such that corresponding parts of the surface have the same colour. The 3<sup>rd</sup> and 4<sup>th</sup> columns give respective versions for the left ventricle (reflected to give same pose as right) of the subject.

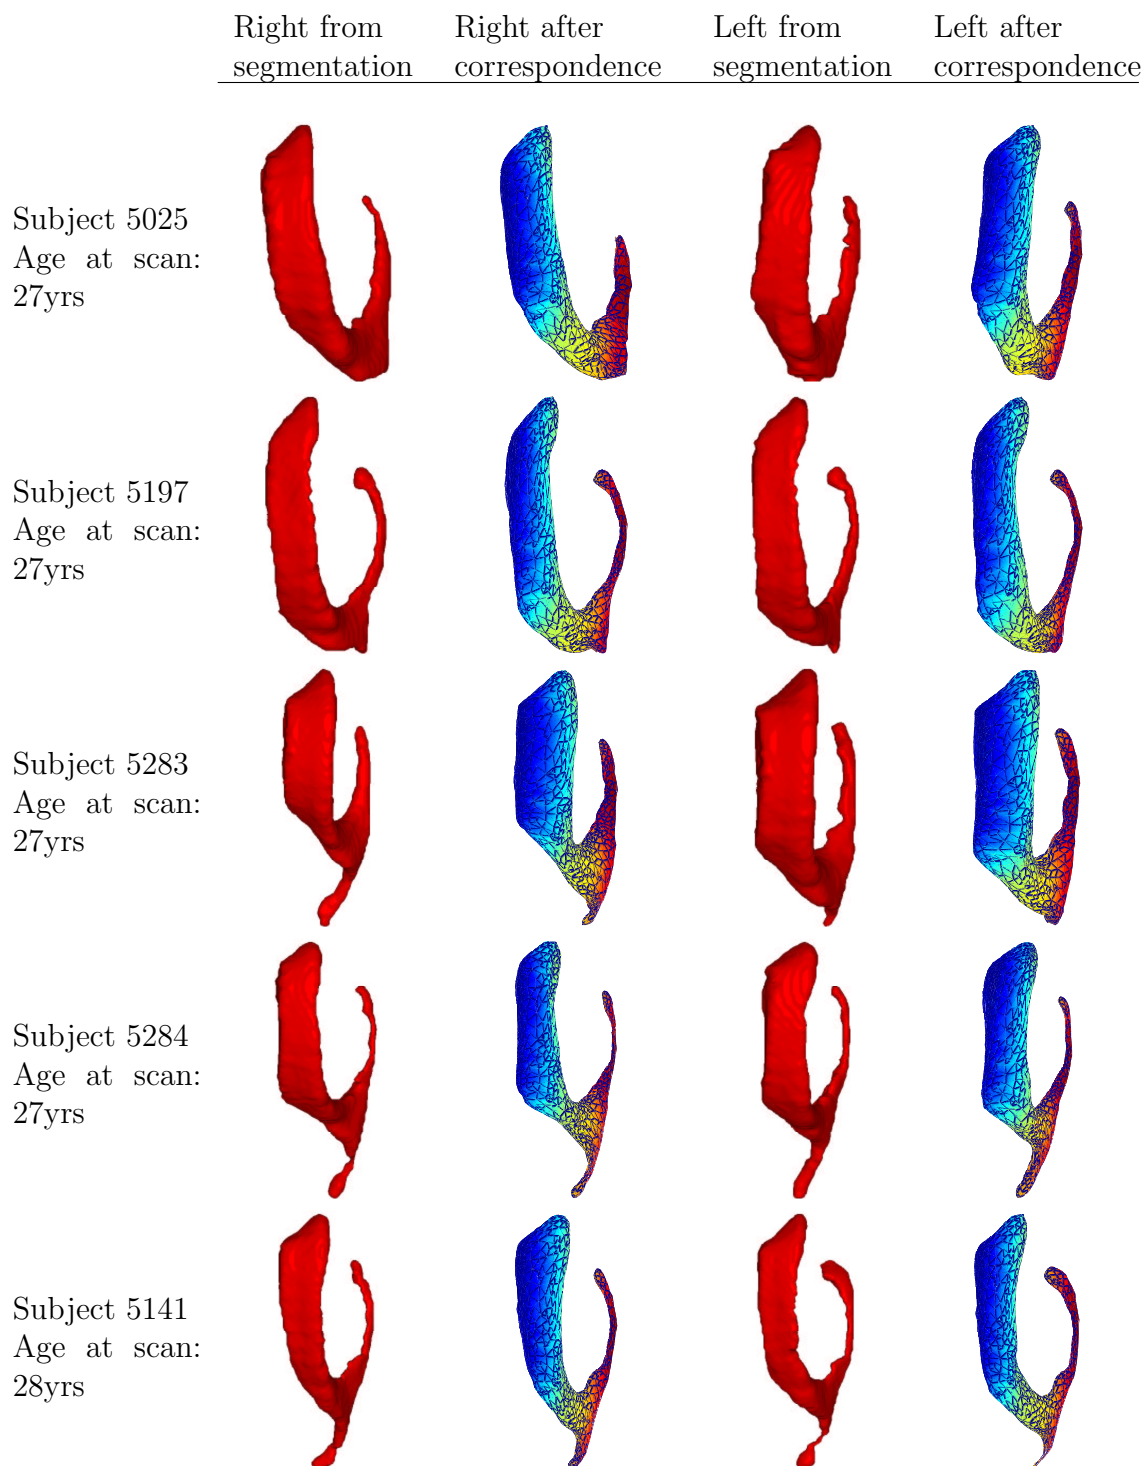




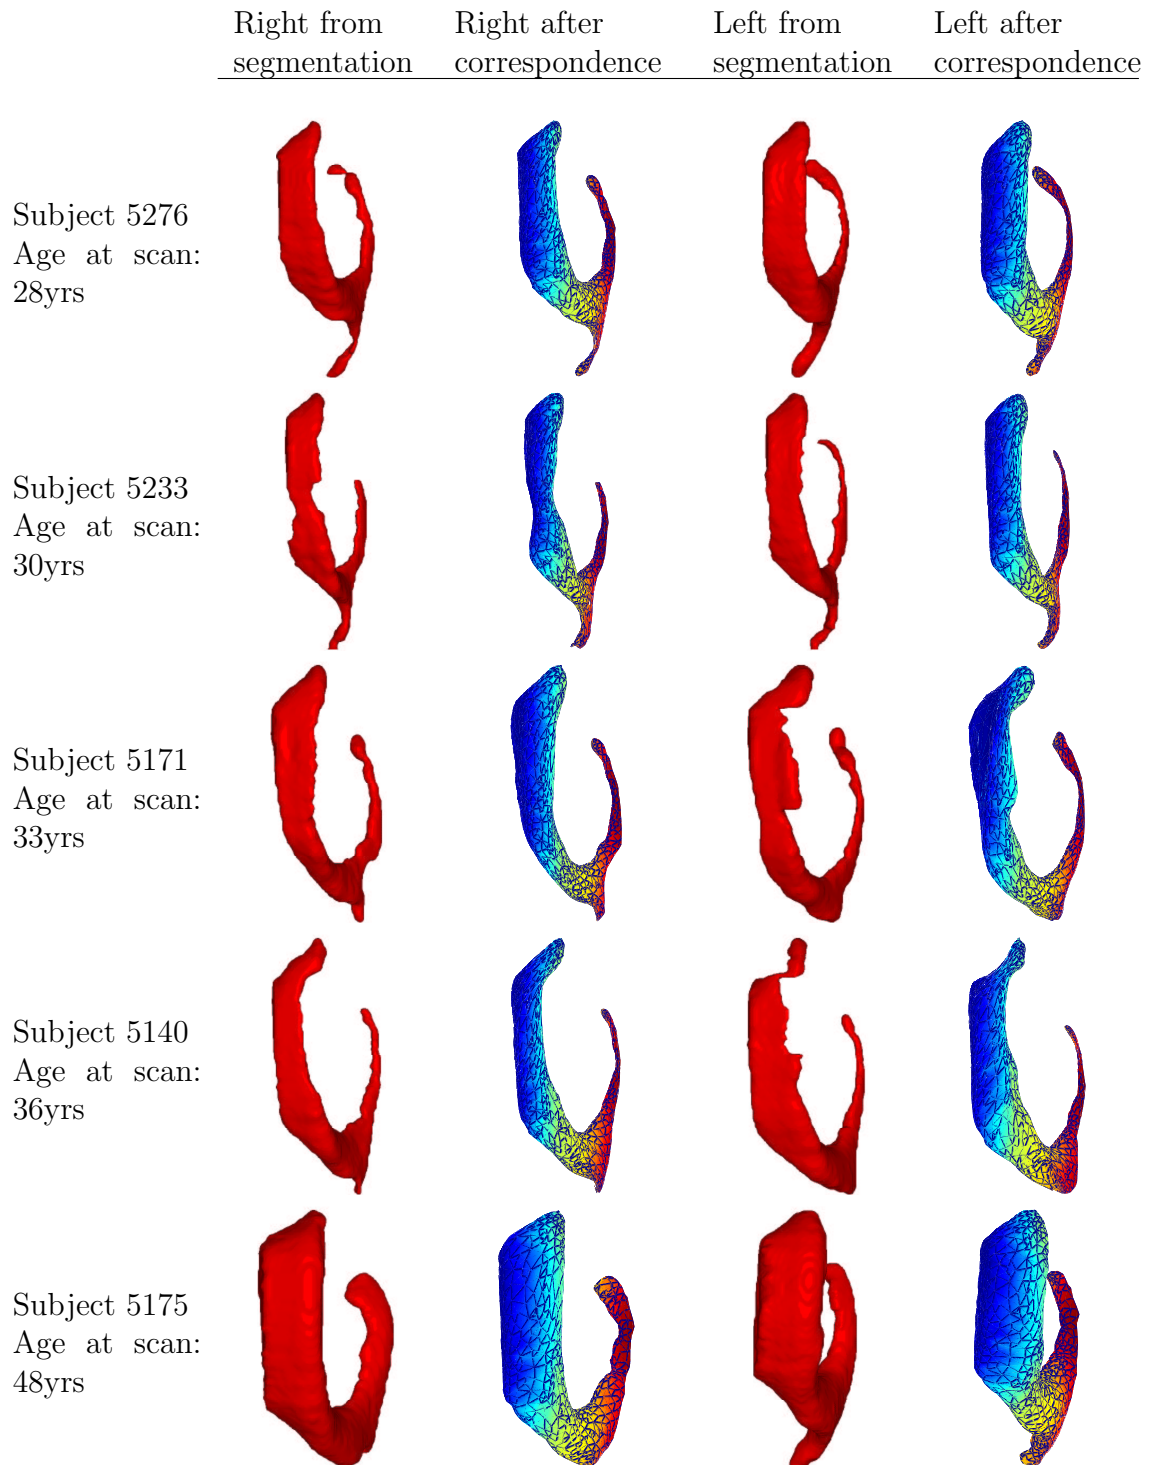
**Figure B.3:** Lateral ventricles of male schizophrenic subjects (11 to 15 of 30). The 1<sup>st</sup> column shows the surface of the segmented right ventricle of the specified subject. The 2<sup>nd</sup> column is the surface of the member of the training set used in construction of the PDM. The surface has been colour-coded such that corresponding parts of the surface have the same colour. The 3<sup>rd</sup> and 4<sup>th</sup> columns give respective versions for the left ventricle (reflected to give same pose as right) of the subject.



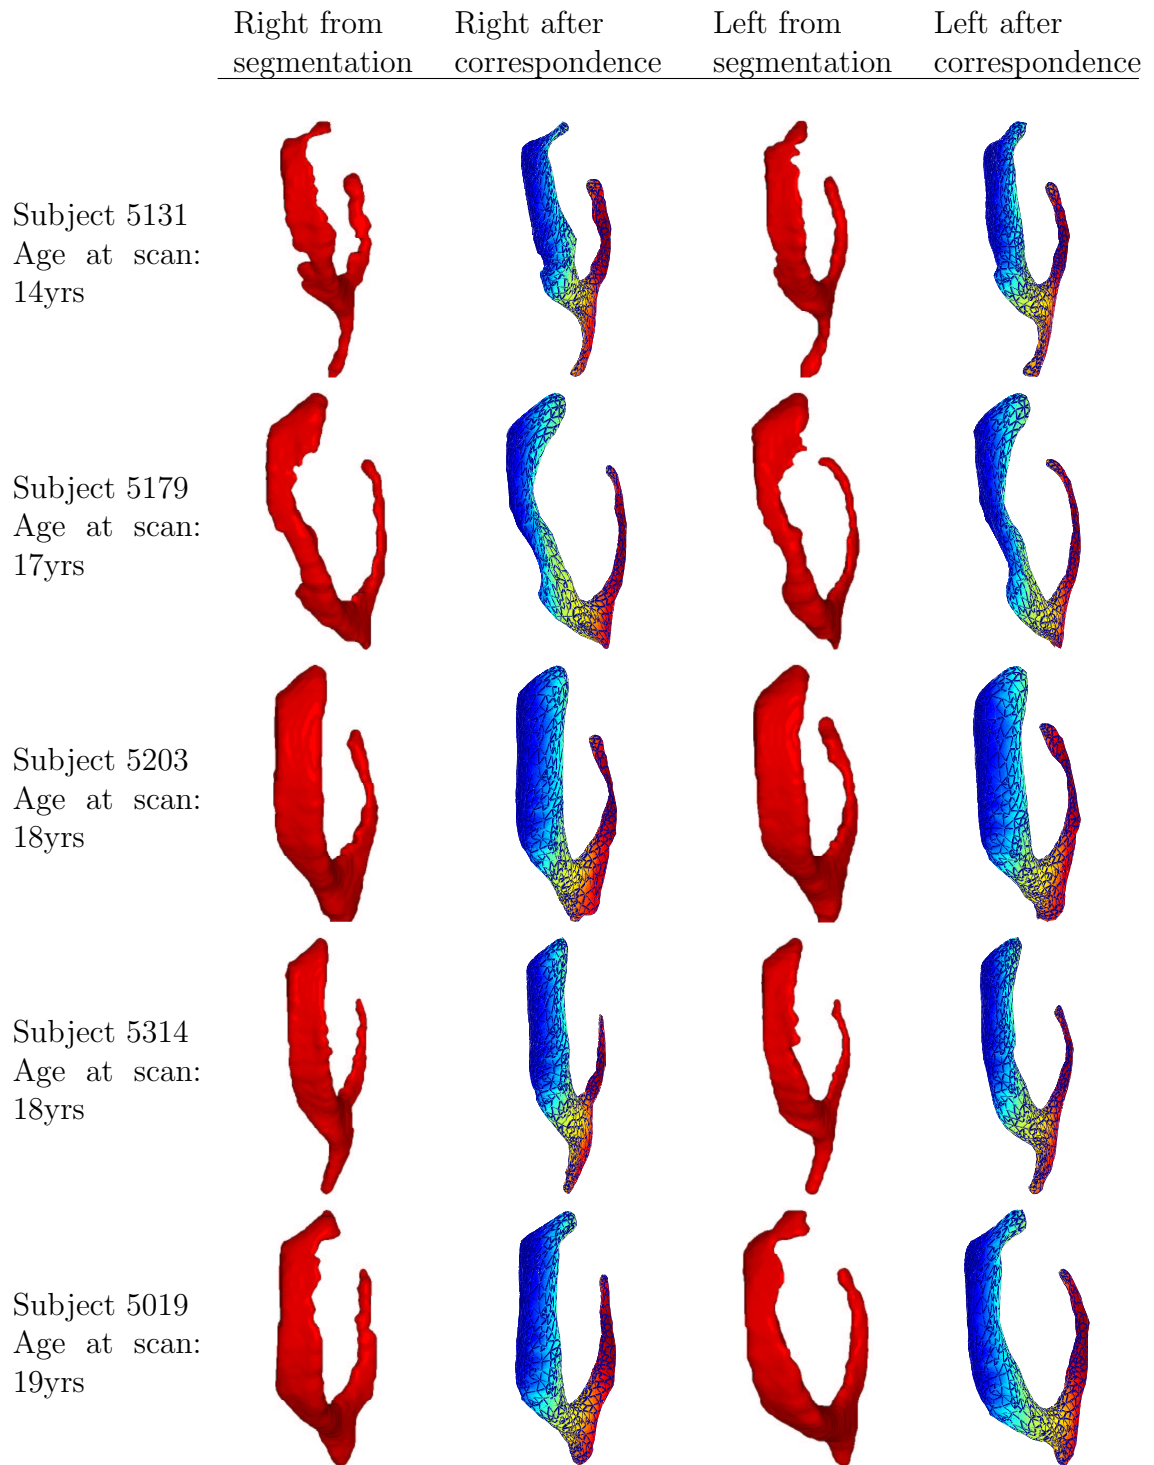
**Figure B.4:** Lateral ventricles of male schizophrenic subjects (16 to 20 of 30). The 1<sup>st</sup> column shows the surface of the segmented right ventricle of the specified subject. The 2<sup>nd</sup> column is the surface of the member of the training set used in construction of the PDM. The surface has been colour-coded such that corresponding parts of the surface have the same colour. The 3<sup>rd</sup> and 4<sup>th</sup> columns give respective versions for the left ventricle (reflected to give same pose as right) of the subject.



**Figure B.5:** Lateral ventricles of male schizophrenic subjects (21 to 25 of 30). The 1<sup>st</sup> column shows the surface of the segmented right ventricle of the specified subject. The 2<sup>nd</sup> column is the surface of the member of the training set used in construction of the PDM. The surface has been colour-coded such that corresponding parts of the surface have the same colour. The 3<sup>rd</sup> and 4<sup>th</sup> columns give respective versions for the left ventricle (reflected to give same pose as right) of the subject.

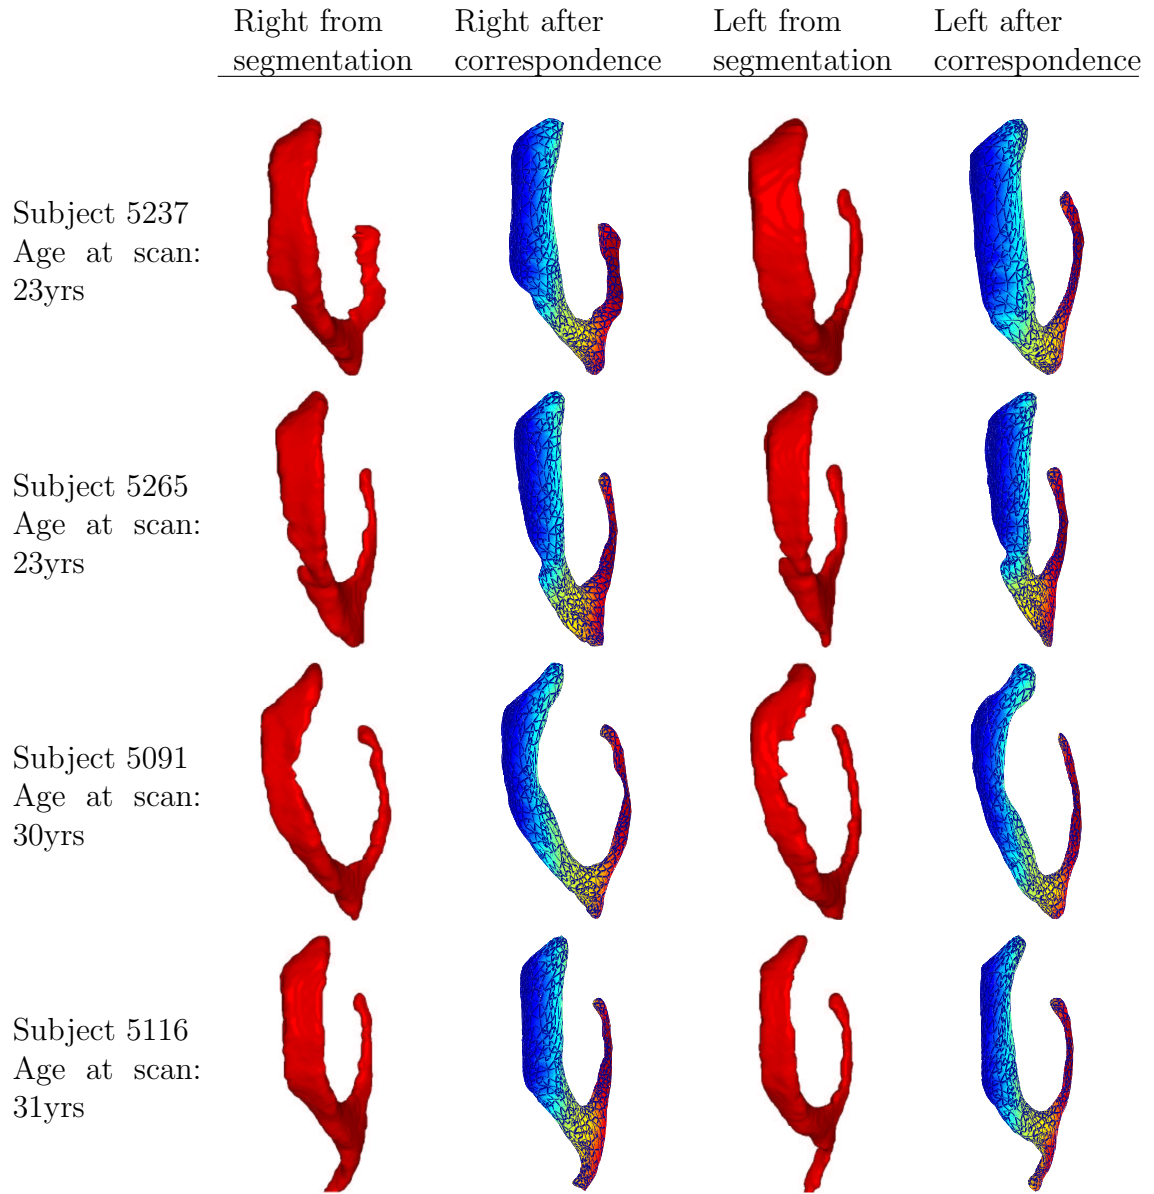


**Figure B.6:** Lateral ventricles of male schizophrenic subjects (26 to 30 of 30). The 1<sup>st</sup> column shows the surface of the segmented right ventricle of the specified subject. The 2<sup>nd</sup> column is the surface of the member of the training set used in construction of the PDM. The surface has been colour-coded such that corresponding parts of the surface have the same colour. The 3<sup>rd</sup> and 4<sup>th</sup> columns give respective versions for the left ventricle (reflected to give same pose as right) of the subject.

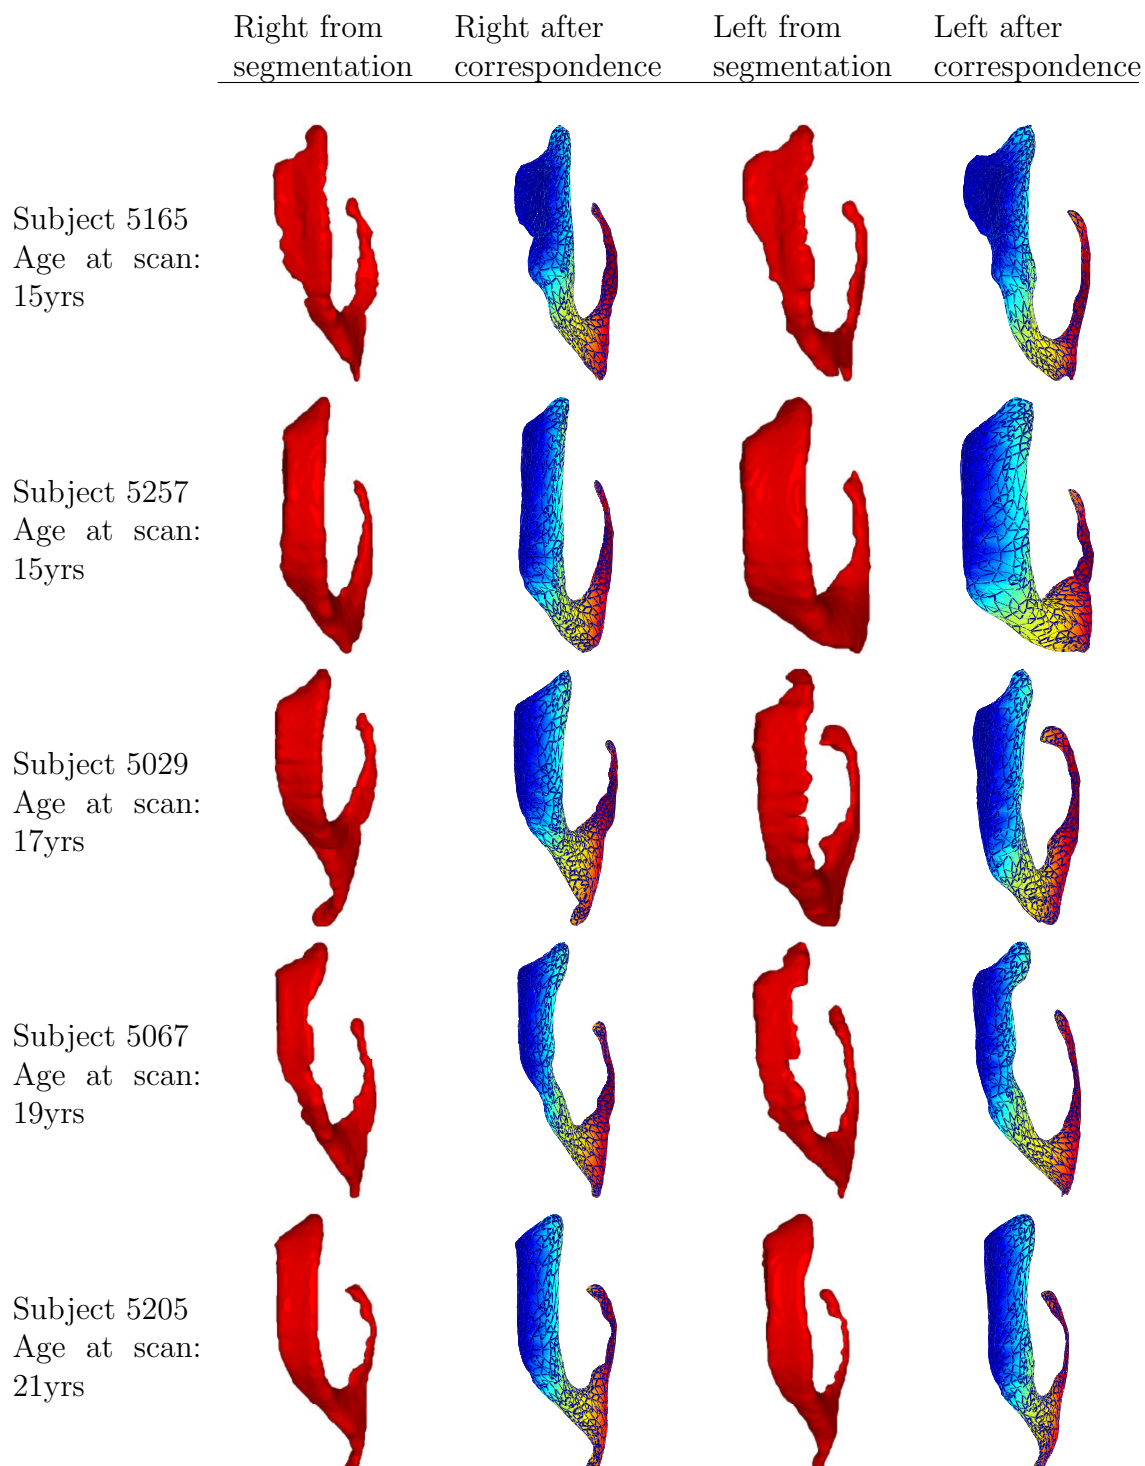


**Figure B.7:** Lateral ventricles of female schizophrenic subjects (1 to 5 of 9). The 1<sup>st</sup> column shows the surface of the segmented right ventricle of the specified subject. The 2<sup>nd</sup> column is the surface of the member of the training set used in construction of the PDM. The surface has been colour-coded such that corresponding parts of the surface have the same colour. The 3<sup>rd</sup> and 4<sup>th</sup> columns give respective versions for the left ventricle (reflected to give same pose as right) of the subject.

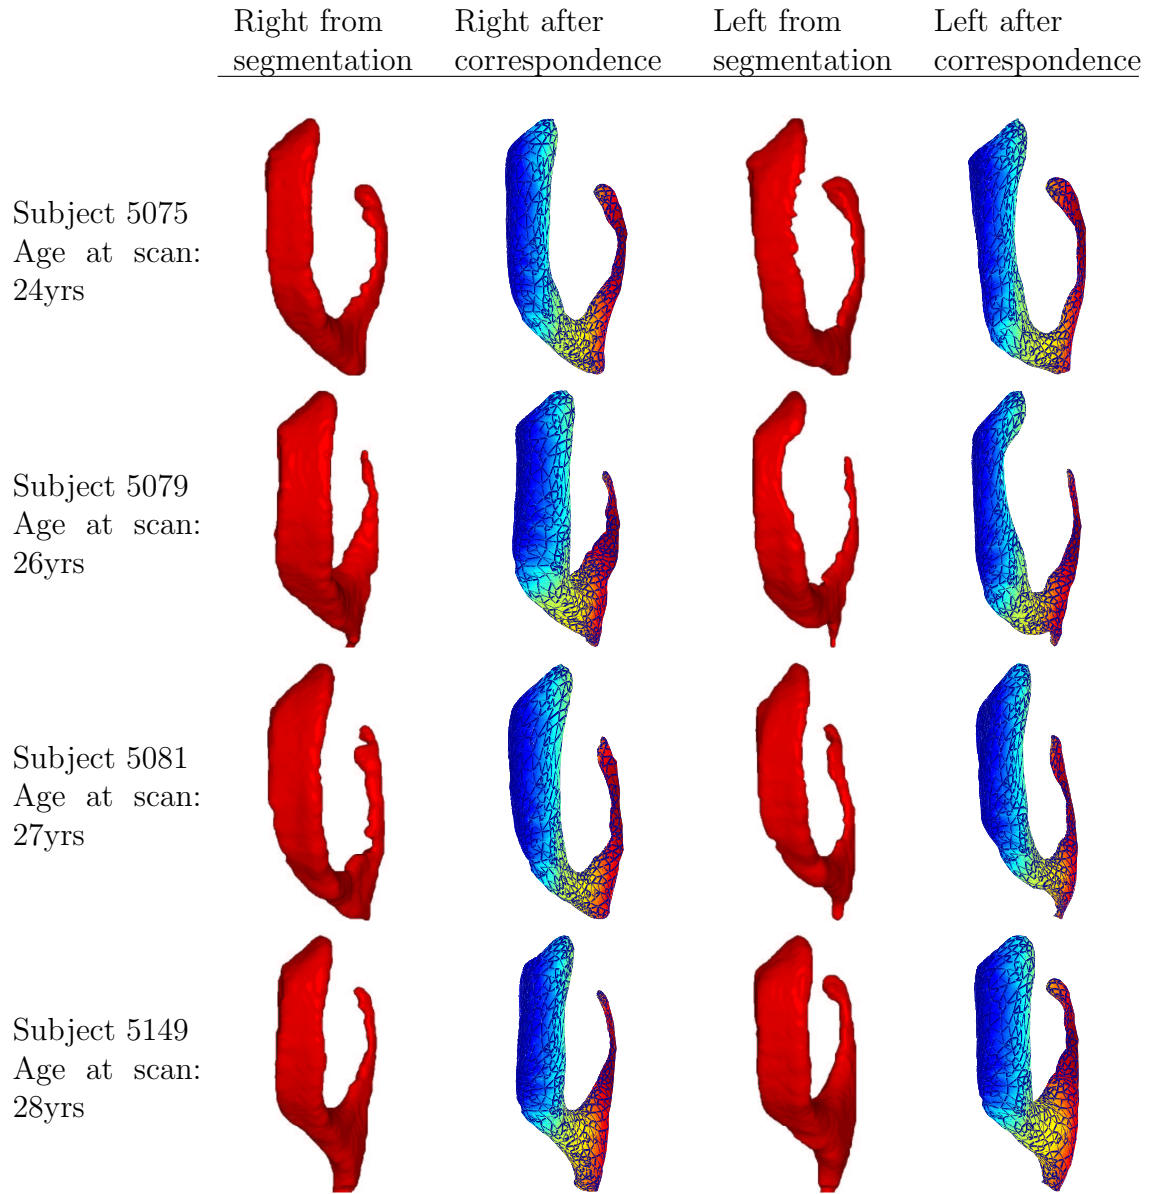




**Figure B.8:** Lateral ventricles of female schizophrenic subjects (6 to 9 of 9). The 1<sup>st</sup> column shows the surface of the segmented right ventricle of the specified subject. The 2<sup>nd</sup> column is the surface of the member of the training set used in construction of the PDM. The surface has been colour-coded such that corresponding parts of the surface have the same colour. The 3<sup>rd</sup> and 4<sup>th</sup> columns give respective versions for the left ventricle (reflected to give same pose as right) of the subject.

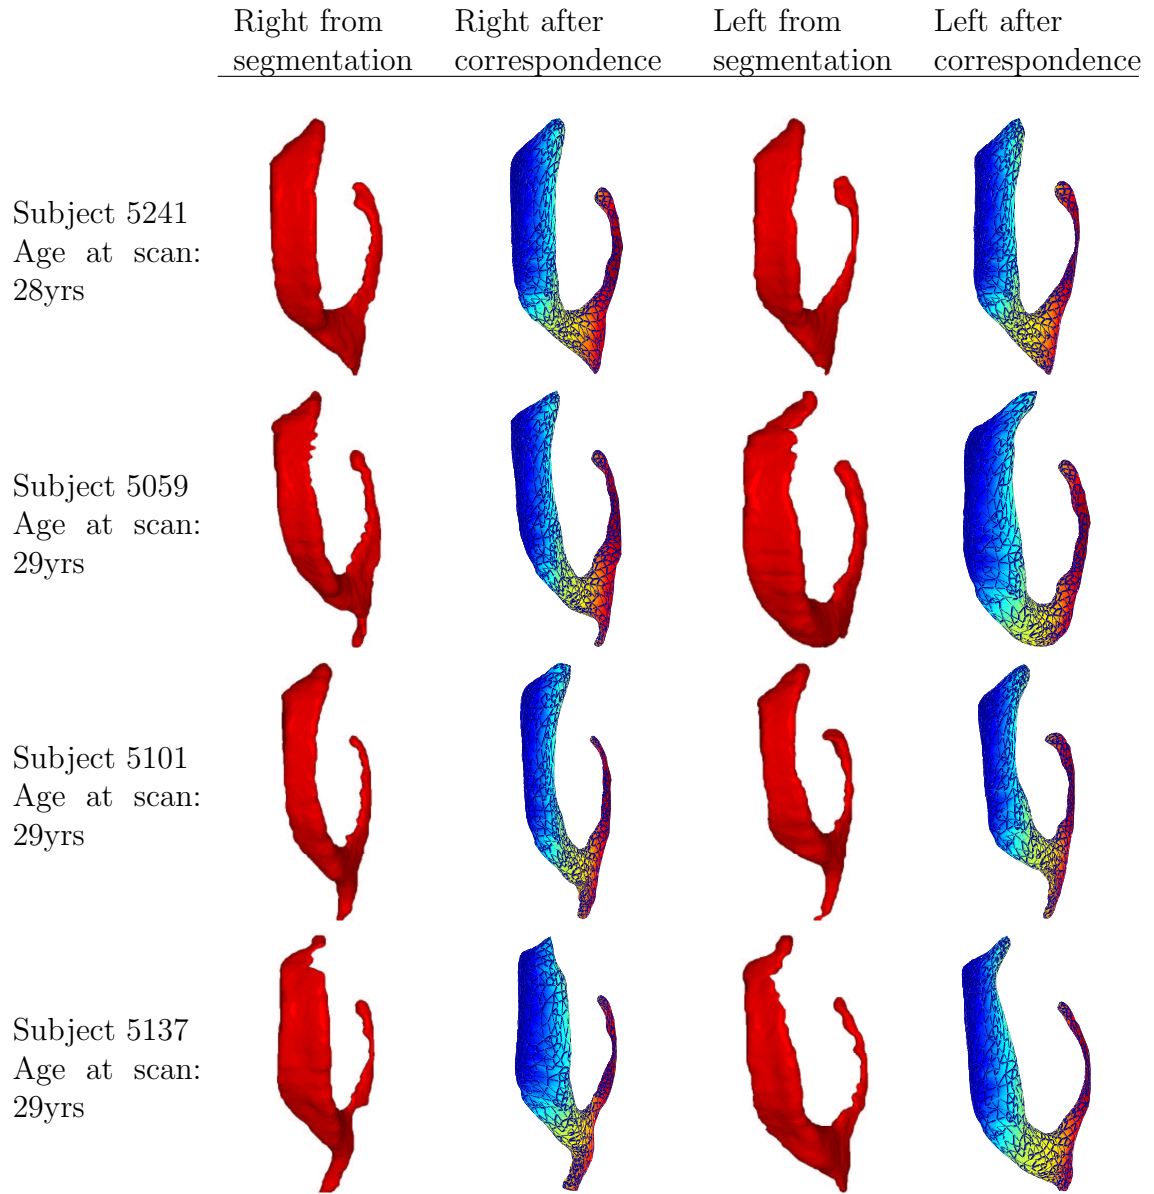


**Figure B.9:** Lateral ventricles of male control subjects (1 to 5 of 17). The 1<sup>st</sup> column shows the surface of the segmented right ventricle of the specified subject. The 2<sup>nd</sup> column is the surface of the member of the training set used in construction of the PDM. The surface has been colour-coded such that corresponding parts of the surface have the same colour. The 3<sup>rd</sup> and 4<sup>th</sup> columns give respective versions for the left ventricle (reflected to give same pose as right) of the subject.

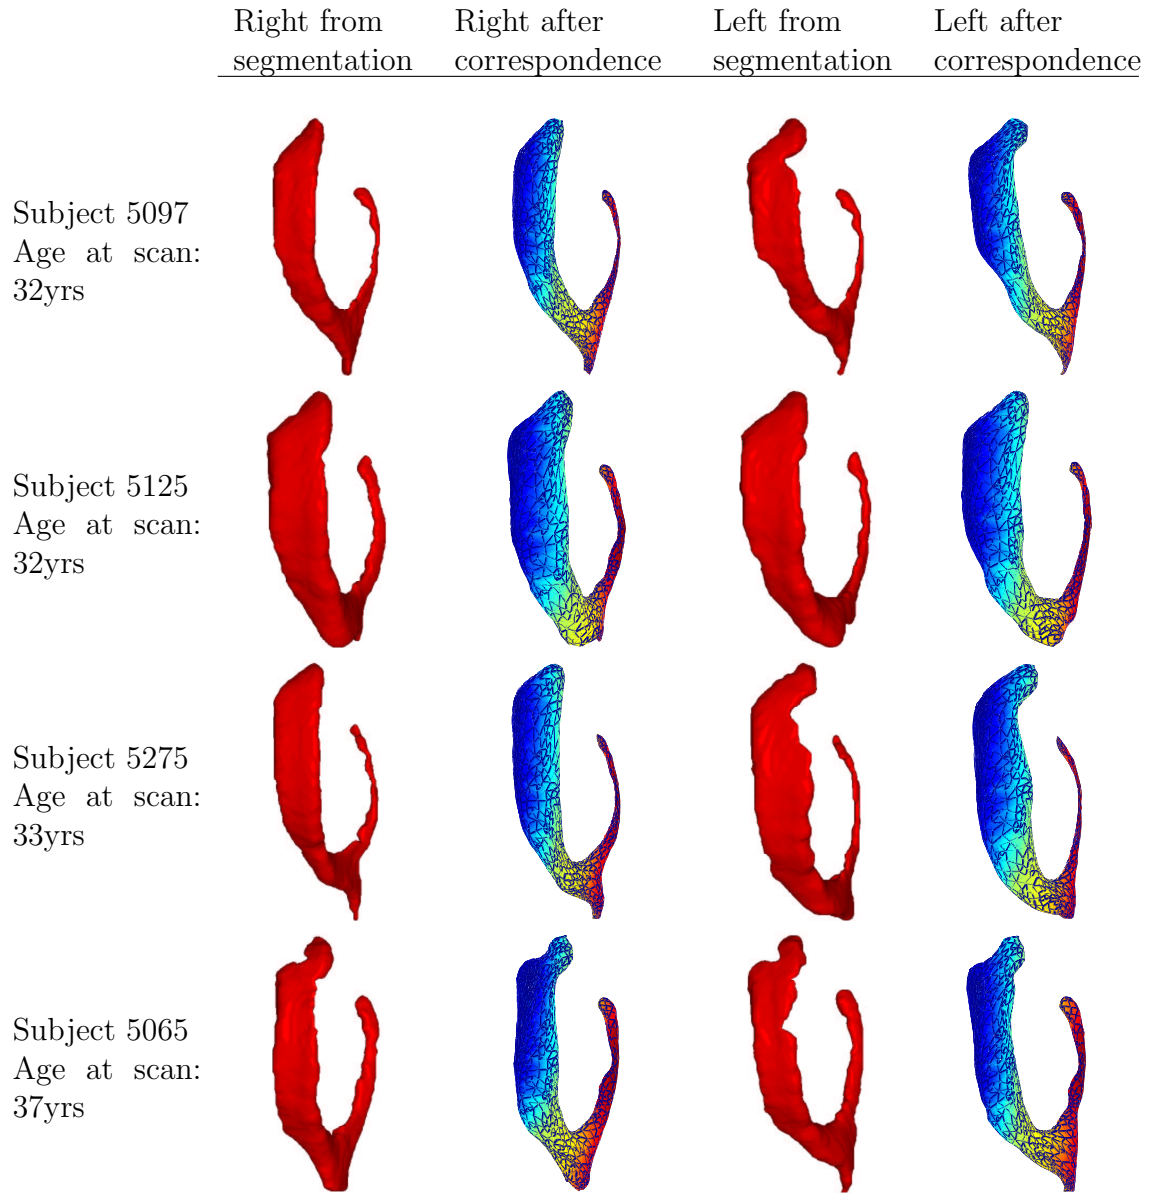


**Figure B.10:** Lateral ventricles of male control subjects (6 to 9 of 17). The 1<sup>st</sup> column shows the surface of the segmented right ventricle of the specified subject. The 2<sup>nd</sup> column is the surface of the member of the training set used in construction of the PDM. The surface has been colour-coded such that corresponding parts of the surface have the same colour. The 3<sup>rd</sup> and 4<sup>th</sup> columns give respective versions for the left ventricle (reflected to give same pose as right) of the subject.

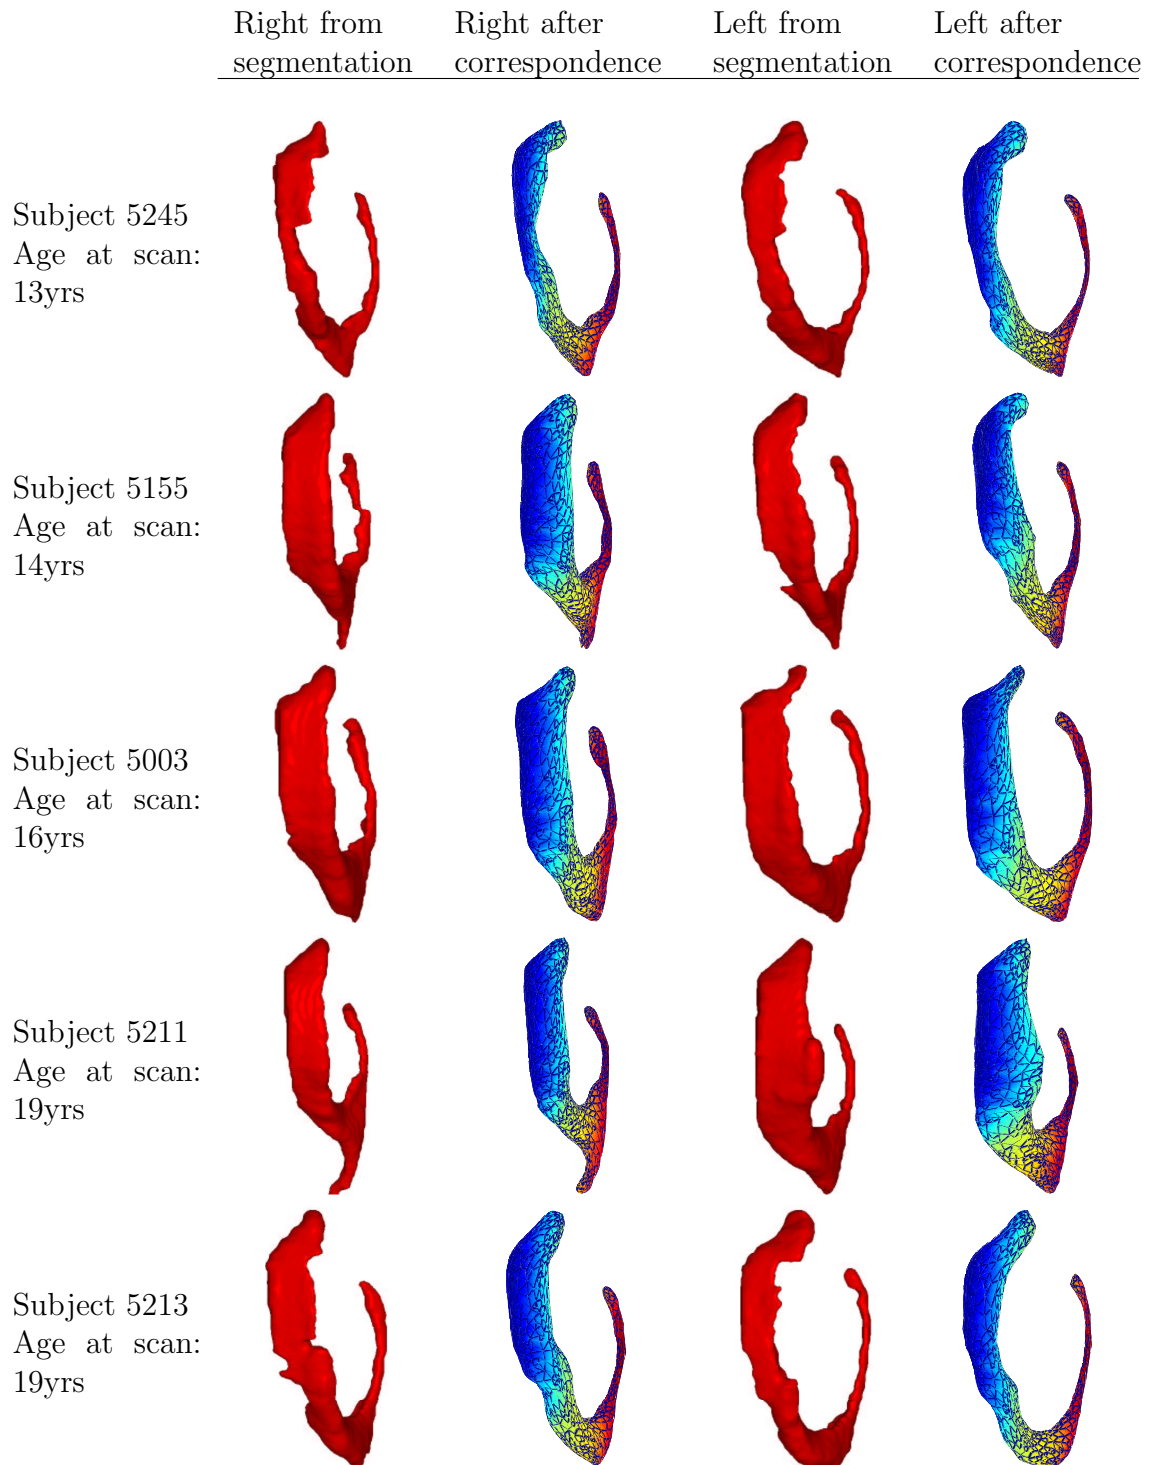




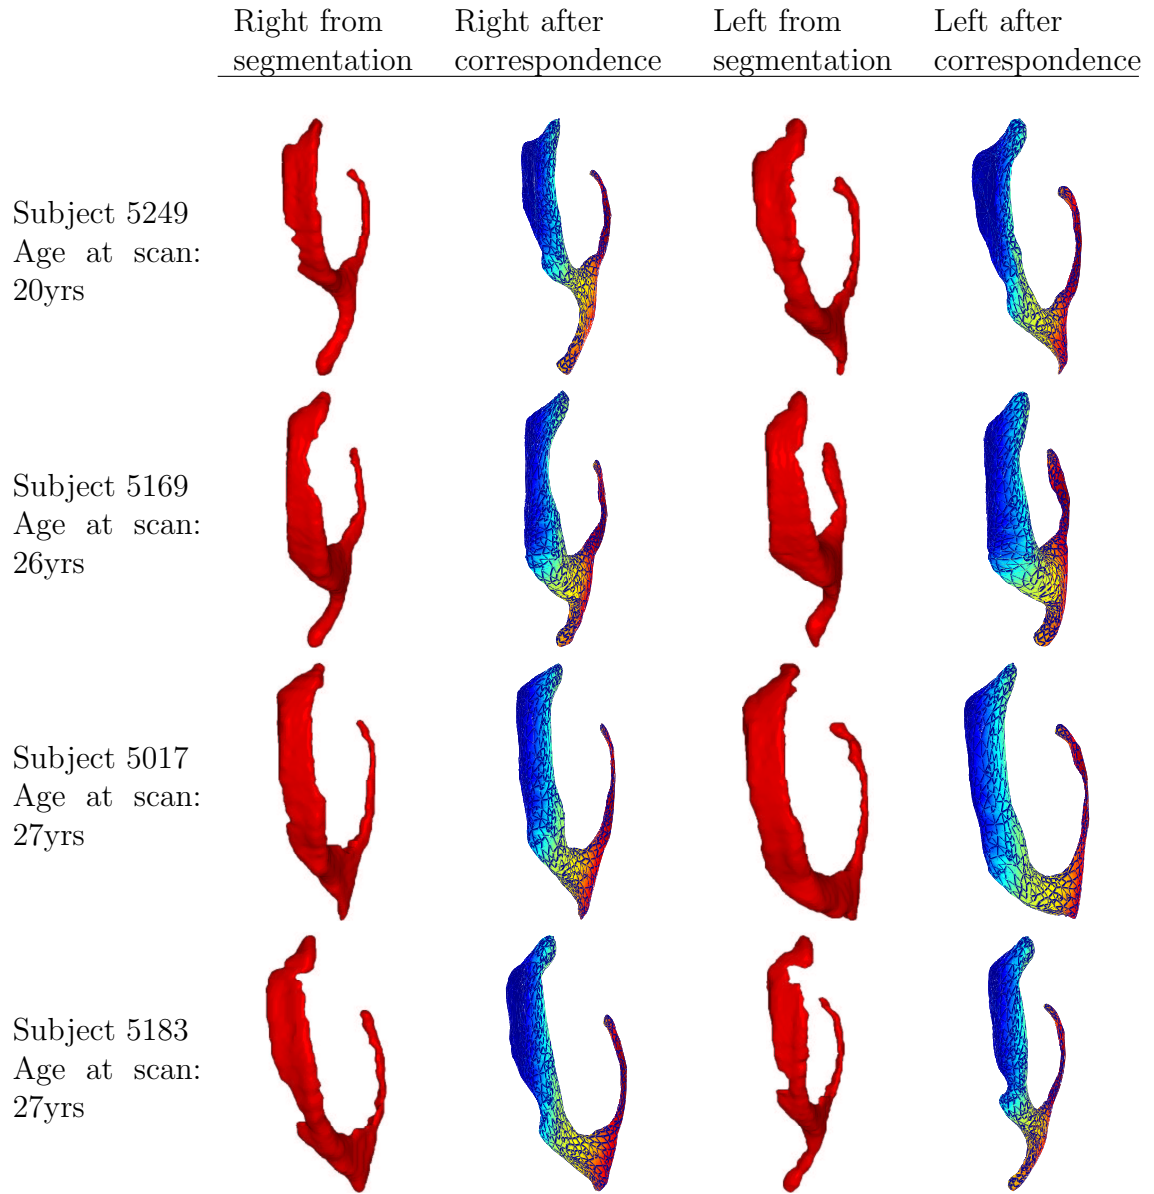
**Figure B.11:** Lateral ventricles of male control subjects (10 to 13 of 17). The 1<sup>st</sup> column shows the surface of the segmented right ventricle of the specified subject. The 2<sup>nd</sup> column is the surface of the member of the training set used in construction of the PDM. The surface has been colour-coded such that corresponding parts of the surface have the same colour. The 3<sup>rd</sup> and 4<sup>th</sup> columns give respective versions for the left ventricle (reflected to give same pose as right) of the subject.



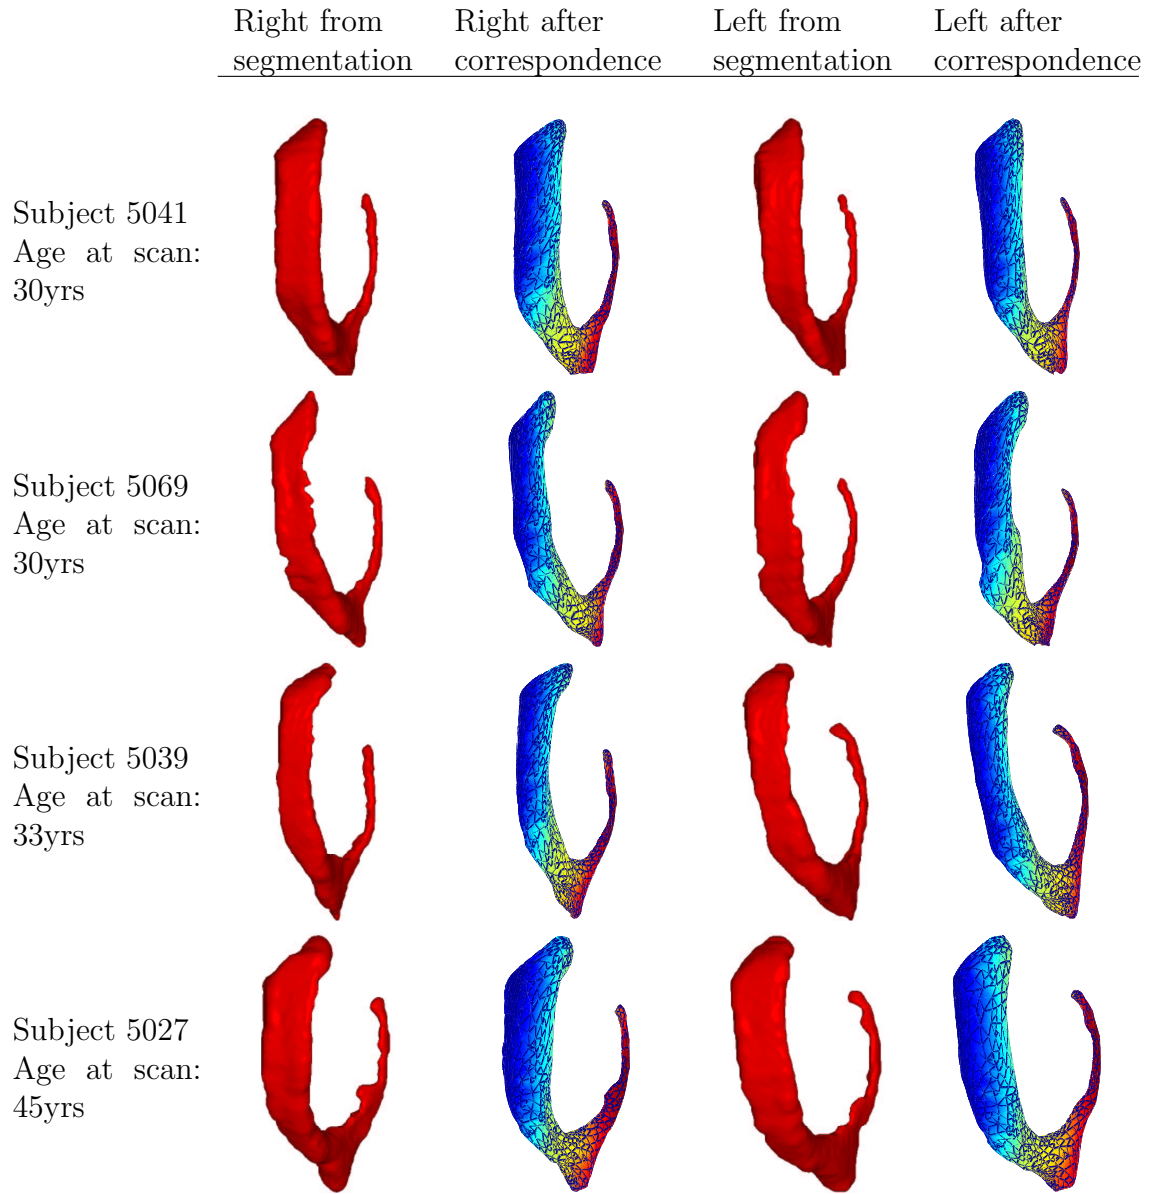
**Figure B.12:** Lateral ventricles of male control subjects (14 to 17 of 17). The 1<sup>st</sup> column shows the surface of the segmented right ventricle of the specified subject. The 2<sup>nd</sup> column is the surface of the member of the training set used in construction of the PDM. The surface has been colour-coded such that corresponding parts of the surface have the same colour. The 3<sup>rd</sup> and 4<sup>th</sup> columns give respective versions for the left ventricle (reflected to give same pose as right) of the subject.



**Figure B.13:** Lateral ventricles of female control subjects (1 to 5 of 13). The 1<sup>st</sup> column shows the surface of the segmented right ventricle of the specified subject. The 2<sup>nd</sup> column is the surface of the member of the training set used in construction of the PDM. The surface has been colour-coded such that corresponding parts of the surface have the same colour. The 3<sup>rd</sup> and 4<sup>th</sup> columns give respective versions for the left ventricle (reflected to give same pose as right) of the subject.



**Figure B.14:** Lateral ventricles of female control subjects (6 to 9 of 13). The 1<sup>st</sup> column shows the surface of the segmented right ventricle of the specified subject. The 2<sup>nd</sup> column is the surface of the member of the training set used in construction of the PDM. The surface has been colour-coded such that corresponding parts of the surface have the same colour. The 3<sup>rd</sup> and 4<sup>th</sup> columns give respective versions for the left ventricle (reflected to give same pose as right) of the subject.



**Figure B.15:** Lateral ventricles of female control subjects (10 to 13 of 13). The 1<sup>st</sup> column shows the surface of the segmented right ventricle of the specified subject. The 2<sup>nd</sup> column is the surface of the member of the training set used in construction of the PDM. The surface has been colour-coded such that corresponding parts of the surface have the same colour. The 3<sup>rd</sup> and 4<sup>th</sup> columns give respective versions for the left ventricle (reflected to give same pose as right) of the subject.

Male Controls		Female Controls	
Subject Number	Age	Subject Number	Age
5165	15	5245	13
5257	15	5155	14
5029	17	5003	16
5067	19	5211	19
5205	21	5213	19
5075	24	5249	20
5079	26	5169	26
5081	27	5017	27
5149	28	5183	27
5241	28	5041	30
5059	29	5069	30
5101	29	5039	33
5137	29	5027	45
5097	32		
5125	32		
5275	33		
5065	37		

**Table B.2:** Age and sex data for control subjects

# Appendix C

## Movies on Accompanying CD

There are fourteen movies on the accompanying CD. They can be viewed by running the Powerpoint presentation - ThesisMovies.ppt. Alternatively, the individual avi files can be viewed using a program such as Realplayer. Contents of the individual avi files are as follows

1. **03-CoronalViewOfSegmentedVentricle.avi** - Coronal View Showing Contours Resulting from Segmentation of a Ventricle
2. **03-RotatingSegmentedVentricle.avi** - Surface of a Segmented Ventricle(rotating)
3. **04-PdmFirstModeOfVariation.avi** - First mode of variation of the PDM of all 138 ventricles (between -2 standard deviations ( $-2\sqrt{\lambda_1}$ ) and +2 standard deviations ( $+2\sqrt{\lambda_1}$ ) of the mean shape)
4. **04-PdmSecondModeOfVariation.avi** - Second mode of variation of the PDM of all 138 ventricles (between -2 standard deviations and +2 standard deviations of the mean shape)
5. **04-PdmThirdModeOfVariation.avi** - Third mode of variation of the PDM of all 138 ventricles (between -2 standard deviations and +2 standard deviations of the mean shape)

6. **05-Comparison1\_ShapeDiffAll.avi** - Shape comparisons of all schizophrenic subjects vs all control subjects. The movie shows movement along the discriminant vector from the schizophrenic mean + 2 standard deviations to the control mean + 2 standard deviations.
7. **05-Comparison2\_ShapeDiffMale.avi** - Shape comparisons of male schizophrenic subjects vs male control subjects. The movie shows movement along the discriminant vector from the schizophrenic mean + 2 standard deviations to the control mean + 2 standard deviations.
8. **05-Comparison3\_ShapeDiffFemale.avi** - Shape comparisons of female schizophrenic subjects vs female control subjects. The movie shows movement along the discriminant vector from the schizophrenic mean + 2 standard deviations to the control mean + 2 standard deviations.
9. **05-Comparison4.1AllSchizAsymm.avi** - Left-Right asymmetry comparisons of all schizophrenic subjects. The movie shows movement along the discriminant vector from the right mean + 2 standard deviations to the left mean + 2 standard deviations.
10. **05-Comparison4.2AllControlAsymm.avi** - Left-Right asymmetry comparisons of all control subjects. The movie shows movement along the discriminant vector from the right mean + 2 standard deviations to the left mean + 2 standard deviations.
11. **05-Comparison5.1MaleSchizAsymm.avi** - Left-Right asymmetry comparisons of male schizophrenic subjects. The movie shows movement along the discriminant vector from the right mean + 2 standard deviations to the left mean + 2 standard deviations.
12. **05-Comparison5.2MaleControlAsymm.avi** - Left-Right asymmetry comparisons of male control subjects. The movie shows movement along the discriminant vector from the right mean + 2 standard deviations to the left mean + 2 standard deviations.



13. **05-Comparison6\_1femaleSchizAsymm.avi** - Left-Right asymmetry comparisons of female schizophrenic subjects. The movie shows movement along the discriminant vector from the right mean + 2 standard deviations to the left mean + 2 standard deviations.
14. **05-Comparison6\_2femaleControlAsymm.avi** - Left-Right asymmetry comparisons of female control subjects. The movie shows movement along the discriminant vector from the right mean + 2 standard deviations to the left mean + 2 standard deviations.

# Bibliography

- [1] *Digital Imaging and Communications in Medicine (DICOM)-PART 1: Introduction & Overview*. Nema Standards Publication PS3.1, 1993.
- [2] H. Achatz, P. Kleinschmidt, and K. Paparrizos. A dual forest algorithm for the assignment problem. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 4:1–11, 1991.
- [3] P. R. Andresen, F. L. Bookstein, K. Conradsen, B. K. Ersbøll, J. L. Marsh, and S. Kreiborg. Surface-bounded growth modeling applied to human mandibles. *IEEE Transactions on Medical Imaging*, 19(11):1053–1063, 2000.
- [4] J. Ashburner and K. J. Friston. Voxel-based morphometry - the methods. *NeuroImage*, 11:805–821, 2000.
- [5] J. Ashburner and K. J. Friston. Why voxel-based morphometry should be used. *NeuroImage*, 14(6):1238–1243, 2001.
- [6] R. Baldock and J. Graham. *Image Processing and Analysis. A Practical Approach*. Oxford University Press, 2000.
- [7] M. A. G. Ballester. *Morphometric Analysis of Brain Structures in MRI*. PhD thesis, Department of Engineering Science, University of Oxford, 1999.
- [8] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape context. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(4):509–522, April 2002.
- [9] P. Besl and N. McKay. A method for registration of 3D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(4), 1991.
- [10] R. M. Bilder, H. Wu, B. Bogerts, M. Ashtari, D. Robinson, M. Woemer, J. A. Lieberman, and G. Degreef. Cerebral volume asymmetries in schizophrenia and mood disorders: A quantitative magnetic resonance imaging study. *International Journal of Psychophysiology*, 34:197–205, 1999.

- [11] H. Blum. Biological shape and visual science (part 1). *Journal of Theoretical Biology*, 38:205–287, 1973.
- [12] M. Board and S. Astley. A novel method of evaluating feature detection algorithms in medical images. In *Proceedings of the 6<sup>th</sup> Conference on Medical Image Understanding and Analysis*, pages 69–72, Portsmouth, UK, 2002.
- [13] J. Boissonnat and G. Geiger. Three dimensional reconstruction of complex shapes based on the Delaunay triangulation. In R. S. Acharya and D. B. Goldgof, editors, *Biomedical Image Processing and Biomedical Visualization*, volume 1905, pages 964–975. SPIE, 1993.
- [14] F. L. Bookstein. The measurement of biological shape and shape change. *Lecture Notes on Biomathematics*, 24, 1978.
- [15] F. L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(6):567–585, 1989.
- [16] F. L. Bookstein. *Morphometric Tools for Landmark Data*. Cambridge University Press, London/New York, 1991.
- [17] F. L. Bookstein. Registration error and functional image analysis. In K. V. Mardia and R. G. Aykroyd, editors, *Functional and spatial data analysis*, pages 31–38. Leeds University Press, 2001.
- [18] F. L. Bookstein. “Voxel-Based Morphometry” should not be used with imperfectly registered images. *NeuroImage*, 14(6):1454–1462, 2001.
- [19] F. L. Bookstein and C. B. Cutting. A proposal for the apprehension of curving craniofacial form in three dimensions. In K. Vig and A. Burdi, editors, *Craniofacial Morphogenesis and Dysmorphogenesis*, pages 127–140. Center of Human Growth, University of Michigan, Ann Arbor, 1988.
- [20] G. Borgefors. Distance transforms in arbitrary dimensions. *Computer Vision, Graphics and Image Processing*, 27:321–345, 1984.
- [21] C. Brechbühler, G. Gerig, and O. Kubler. Parameterization of closed surfaces for 3D shape description. *Computer Vision and Image Understanding*, 61(2):154–170, 1995.
- [22] A. D. Brett and C. J. Taylor. A method of automated landmark generation for automated 3D PDM construction. *Image and Vision Computing*, 18:739–748, 2000.

- [23] D. Brujic, I. Ainsworth, M. Ristic, and V. Brujic. Efficient shape description using NURBS. In *Proceedings of the 4<sup>th</sup> International Workshop on Visual Form*, pages 643–653, Capri, Italy, 2001.
- [24] P. F. Buckley, D. Dean, F. L. Bookstein, L. Friedman, D. Kwon, J. S. Lewin, J. Kamath, and C. Lys. Three-dimensional magnetic resonance-based morphometrics and ventricular dysmorphology in schizophrenia. *Biological Psychiatry*, 45:62–67, 1999.
- [25] J. F. Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(1):679–698, 1986.
- [26] S. H. Cardoso. Brain ventricles - our internal lake. <http://www.epub.org.br/cm/n02/fundamentos/ventriculos.i.htm>.
- [27] M. Carlsson and A. Carlsson. Schizophrenia: a subcortical neurotransmitter imbalance syndrome? *Schizophrenia Bulletin*, 16:425–432, 1990.
- [28] S. A. Chance, M. M. Esiri, and T. J. Crow. Ventricular enlargement in schizophrenia: A primary change in the temporal lobe. *Schizophrenia Research*, 62:123–131, 2003.
- [29] J. Chou, S. Chen, G. S. Sudakoff, K. R. Hoffmann, C. Chen, and A. H. Dachman. Image fusion for visualization of hepatic vasculature and tumors. In *Proceedings of the SPIE International Symposium on Medical Imaging*, volume SPIE Proceedings Volume 2434, pages 157–163, San Diego, CA, USA, 1995.
- [30] D. L. Collins, A. P. Zijdenbos, V. Kollokian, J. G. Sled, N. J. Kabani, C. J. Holmes, and A. C. Evans. Design and construction of a realistic digital brain phantom. *IEEE Transactions on Medical Imaging*, 17(3):463–468, 1998. MR image generator available online at <http://www.bic.mni.mcgill.ca/brainweb>.
- [31] T. F. Cootes. Statistical models of appearance for computer vision. Technical report, Imaging Science and Biomedical Engineering, University of Manchester, February 2000. Available from <http://www.isbe.man.ac.uk/~bim/>.
- [32] T. F. Cootes, A. Hill, C. J. Taylor, and J. Haslam. The use of Active Shape Models for locating structure in medical images. *Image Processing and Machine Intelligence*, pages 33–47, June 1993.
- [33] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active Shape Models - their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, Jan. 1995.

- [34] D. Coxeter. *Introduction to Geometry*. John Wiley and Sons, 2<sup>nd</sup> edition, 1989.
- [35] T. J. Crow. Schizophrenia as failure of hemispheric dominance for language. *Trends in Neuroscience*, 20:339–343, 1997.
- [36] J. G. Csernansky, S. Joshi, L. Wang, J. W. Haller, M. Gado, J. Philip Miller, U. Grenander, and M. I. Miller. Hippocampal morphometry in schizophrenia by high dimensional brain mapping. *Proceedings of the National Academy of Sciences, USA*, 95:11406–11411, September 1998.
- [37] C. Davatzikos, A. Genc, D. Xu, and S. M. Resnick. Voxel-based morphometry using the RAVENS maps: Methods and validation using simulated longitudinal atrophy. *NeuroImage*, 14:1361–1369, 2001.
- [38] R. Davies. *Learning Shape: Optimal Models for Analysing Natural Variability*. PhD thesis, Division of Imaging Science and Biomedical Engineering, University of Manchester, 2002.
- [39] R. H. Davies, C. J. Twining, T. F. Cootes, J. C. Waterton, and C. J. Taylor. 3D statistical shape models using direct optimisation of description length. In *In: Proceedings of the 7th European Conference on Computer Vision, LCNS, Vol 2350(3)*, pages 3–20, 2002.
- [40] B. M. Dawant, A. P. Zijdenbos, and R. A. Margolin. Correction of intensity variations in mr images for computer aided tissue classification. *IEEE Transactions in Medical Imaging*, 12(4):770–781, 1993.
- [41] D. Dean, P. F. Buckley, F. L. Bookstein, L. Friedman, D. Kwon, J. S. Lewin, J. Kamath, and C. Lys. Three dimensional MR-based morphometric comparison of schizophrenic and normal cerebral ventricles. In *Visualisation in Biomedical Computing, Lecture Notes in Computer Science*, volume 1131, pages 363–372, 1996.
- [42] P. Dierckx. *Curve and surface fitting with splines*. Oxford University Press, 1993.
- [43] M. P. Do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice–Hall, College Division, 1976.
- [44] I. L. Dryden and K. V. Mardia. *Statistical Shape Analysis*. John Wiley and Sons, 1998.
- [45] L. C. Duarte, M. L. R, F. Jos Von Zuben, and F. Furtado Dos Reis. Variation in mandible shape in *thrichomys apereoides* (mammalia:rodentia): geometric analysis of a complex morphological structure. *Systematic Biology*, 49(3):563–578, 2000.

- [46] R. O. Duda and P. E. Hart. *Pattern Classification and Scene Analysis*. John Wiley and Sons Ltd, 1973.
- [47] D. Eberly, R. Gardner, B. Morse, S. Pizer, and C. Scharlach. Ridges for image analysis. *Journal of Mathematical Imaging and Vision*, 4:351–371, 1994.
- [48] A. C. Evans, D. L. Collins, S. R. Mills, E. D. Brown, R. L. Kelly, and T. M. Peters. 3D statistical neuroanatomical models from 305 MRI volumes. In *Proceedings of the IEEE-Nuclear Science Symposium and Medical Imaging Conference*, pages 1813–1817, 1993.
- [49] P. Falkai, B. Bogerts, B. Greve, U. Pfeiffer, B. Machus, B. Folsch-Reetz, C. Majtenyi, and I. Ovary. Loss of sylvian fissure asymmetry in schizophrenia: A quantitative post mortem study. *Schizophrenia Research*, 7:23–32, 1992.
- [50] R. A. Fisher. The use of multiple measures in taxonomic problems. *Annals of Eugenics*, 7:179–188, 1936.
- [51] J. M. Fitzpatrick, J. B. West, and C. R. Maurer. Predicting error in rigid-body, point-based registration. *IEEE Transactions on Medical Imaging*, 17:694–702, 1998.
- [52] J. D. Foley, A. van Dam, S. K. Feiner, and J. F. Hughes. *Computer Graphics - Principles and Practice*. Addison-Wesley, 2 edition, 1990.
- [53] A. F. Frangi, D. Rueckert, J. A. Schnabel, and N. W. J. Automatic 3D ASM construction via atlas-based landmarking and volumetric elastic registration. *Lecture Notes in Computer Science (IPMI 2001 Proceedings)*, 2082:78–91, 2001.
- [54] G. Gerig, M. Styner, D. Jones, D. Weinberger, and J. Lieberman. Shape analysis of brain ventricles using SPHARM. In *Proceedings of IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, pages 171–178, 2001.
- [55] F. Gibbons. Galleries tap rich vein of public’s lust for blood. In *The Gaurdian*. 9 January, 2003.
- [56] P. Golland, E. L. Grimson, M. E. Shenton, and R. Kikinis. Small sample size learning for shape analysis of anatomical structures. *Lecture Notes in Computer Science (MICCAI 2000 Proceedings)*, 1935:305–314, 2000.
- [57] P. Golland, W. E. L. Grimson, M. E. Shenton, and R. Kikinis. Deformation analysis for shape based classification. In *Proceedings of the 16<sup>th</sup> International Conference on Information Processing in Medical Imaging, LNCS 1613*, pages 517–530, 1999.
- [58] J. C. Gower. Generalized procrustes analysis. *Psychometrika*, 40:33–51, 1975.

- [59] A. Guimond, J. Muenier, and J. Thirion. Average brain models. *Computer Vision and Image Understanding*, 77:192–210, 2000.
- [60] G. Hadley. *Linear Programming*. Addison-Wesley, Reading, Massachusetts, 1975.
- [61] W. J. Haller, A. Bannerjee, G. E. Christensen, M. Gado, S. Joshi, M. I. Miller, Y. L. Sheline, M. W. Vannier, and J. G. Csernansky. Three-Dimensional hippocampal MR morphometry with high-dimensional transformation of a neuroanatomic atlas. *Radiology*, 202:504–510, 1997.
- [62] D. J. Hand. *Discrimination and Classification*. John Wiley and Sons Ltd, 1981.
- [63] A. Hill, T. A., and T. C. J. Model-based interpretation of 3D medical images. In *Proceedings of the 4<sup>th</sup> British Machine Vision Conference*, pages 339–348, Guildford, UK, 1993.
- [64] A. Hill, T. C. J., and B. A. D. A framework for automatic landmark identification using a new method of nonrigid correspondence. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(3):241–251, 2000.
- [65] M. Hosaka. *Modeling of Curves and Surfaces in CAD/CAM*. Springer-Verlag, 1992.
- [66] M. Jenkinson and S. Smith. A global optimisation method for robust affine registration of brain images. *Medical Image Analysis*, 5(2):143–156, 2001. Software available for download at <http://www.fmrib.ox.ac.uk/fsl/flirt/index.html>.
- [67] E. C. Johnstone, T. J. Crow, C. D. Frith, J. Husband, and L. Kreel. Cerebral ventricular size and cognitive impairment in chronic schizophrenia. *Lancet*, 7992(2):924–6, 1976.
- [68] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1988.
- [69] A. Kelemen, G. Szekely, and G. Gerig. Elastic model-based segmentation of 3-D neuroradiological data sets. *IEEE Transactions on Medical Imaging*, 18(10):828–839, 1999.
- [70] D. G. Kendall. The diffusion of shape. *Advances in Applied Probability*, 9:428–430, 1977.
- [71] J. T. Kent, K. V. Mardia, and W. J. M. Ridge curves and shape analysis. In *Proceedings of the 7th British Machine Vision Conference*, pages 43–52, 1996.
- [72] J. J. Koenderink. *Solid Shape*. The MIT Press, 1990.

- [73] M. Kubicki, M. E. Shenton, D. F. Salisbury, Y. Hirayasu, K. Kasai, R. Kinikis, F. A. Jolesz, and R. W. McCarley. Voxel-based morphometric analysis of gray matter in first episode schizophrenia. *NeuroImage*, 17:1711–1719, 2002.
- [74] S. D. Kwak, W. W. Colman, and G. A. Ateshian. Anatomy of the human patellofemoral joint articular cartilage: Surface curvature analysis. *Journal of Orthopaedic Research*, 15:468–472, 1997.
- [75] H. Lamecker, T. Lange, and M. Seebass. A statistical shape model for the liver. In *Proceedings of the 5<sup>th</sup> Conference on Medical Image Computing and Computer-Assisted Intervention, Lecture Notes in Computer Science*, volume 2489, pages 412–427, 2002.
- [76] T. M. Lehmann, C. Gönner, and K. Spitzer. Survey: Interpolation methods in medical image processing. *IEEE Transactions on Medical Imaging*, 18(11):1049–1075, 1999.
- [77] S. Loncaric. A survey of shape analysis techniques. *Pattern Recognition*, 31(8):983–1001, 1998.
- [78] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. *Computer Graphics*, 21(4):163–169, 1987.
- [79] C. Lorenz and N. Kranhnsstöver. Generation of point-based 3D statistical shape models for anatomical objects. *Computer Vision and Image Understanding*, 77:175–191, 2000.
- [80] J. R. Magnus and H. Neudecker. *Matrix Differential Calculus with Applications in Statistics and Econometrics*. Wiley, 1999.
- [81] J. B. A. Maintz and M. A. Viergever. A survey of medical image registration. *Medical Image Analysis*, 2(1):1–36, 1998.
- [82] S. Marsland and C. Twining. Clamped-plate splines and the optimal flow of diffeomorphisms. In *Proceedings of the 21st Leeds Annual Statistics Workshop*, Leeds, UK, 2002.
- [83] J. Martin, A. Pentland, S. Sclaroff, and R. Kikinis. Characterization of neuropathological shape deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(2):97–112, 1998.
- [84] O. Monga and S. Benayoun. Using partial derivatives of 3d images to extract typical surface features. *Computer Vision and Image Understanding*, 61(2):171–189, 1995.



- [85] O. Monga, R. Deriche, G. Malandain, and J. Cocquerez. Recursive filtering and edge tracking: two primary tools for 3D edge detection. *Image and Vision Computing*, 9(4):203–214, 1991.
- [86] R. Morris. The sub-parabolic lines of a surface. In G. Mullineux, editor, *In Mathematics of Surfaces VI - Proceedings of Sixth IMA Conference on the Mathematics of Surfaces*, pages 79–102. Clarendon Press, Oxford, 1994.
- [87] K. L. Narr, P. M. Thompson, T. Sharma, J. Moussai, R. Blanton, B. Anvar, A. Edris, R. Krupp, J. Rayman, M. Khaledy, and A. W. Toga. Three-dimensional mapping of tempro-limbic regions and the lateral ventricles in schizophrenia: Gender effects. *Biological Psychiatry*, 50:84–97, 2001.
- [88] J. O’Rourke. *Computational Geometry in C*. Cambridge University Press, 2<sup>nd</sup> edition, 1998.
- [89] C. Pantelis, F. Z. Barber, T. R. Barnes, H. E. Nelson, A. M. Owen, and T. W. Robbins. Comparison of set-shifting ability in patients with chronic schizophrenia and frontal lobe damage. *Schizophrenia Research*, 37:251–70, 1999.
- [90] G. D. Pearlson and L. Marsh. Structural brain imaging in schizophrenia: A selective review. *Biological Psychiatry*, 46(5):627–649, 1999.
- [91] A. P. Pentland and S. Sclaroff. Closed-form solutions for physically based shape modelling and recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7):715–729, 1991.
- [92] S. M. Pizer, P. T. Fletcher, A. Thall, M. Styner, G. Gerig, and S. Joshi. Object models in multiscale intrinsic coordinates via m-reps. *Image and Vision Computing*, 21(1):5–15, 2003.
- [93] D. Poxton, J. Graham, and J. F. W. Deakin. Detecting asymmetries in hippocampal shape and receptor distribution using statistical appearance models and linear discriminant analysis. In *Proceedings of the Ninth British Machine Vision Conference*, pages 525–534, Southampton, UK, 1998.
- [94] B. K. Puri, A. J. Richardson, A. Oatridge, J. V. Hajnal, and N. Saeed. Cerebral ventricular a symmetry in schizophrenia: A high resolution 3D magnetic resonance imaging study. *International Journal of Psychophysiology*, 34:207–211, 1999.
- [95] A. Rangarajan, H. Chui, and B. F. L. The softassign procrustes matching algorithm. In J. Duncan and G. Gindi, editors, *Information Processing in Medical Imaging (IPMI’97), Lecture Notes in Computer Science LCNS 905*, pages 29–42, 1997.

- [96] A. Ravindran, D. T. Philips, and J. J. Solberg. *Operations Research, Principles and Practice*. John Wiley and Sons, 2<sup>nd</sup> edition, 1987.
- [97] J. Rocha and T. Pavlidis. A shape analysis model with applications to a character recognition system. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(4), 1994.
- [98] D. Rutovitz. Expanding picture components to natural density boundaries by propagation methods, the notions of fall-set and fall-distance. In *Proceedings of the 4<sup>th</sup> International Joint Conference on Pattern Recognition*, pages 657–664, Kyoto, Japan, 1978.
- [99] P. K. Saha and B. B. Chaudhuri. 3D digital topology under binary transformation with application. *Computer Vision and Image Understanding*, 63(3):418–429, 1996.
- [100] D. Shen, H. H. Edward, and D. C. An adaptive-focus statistical shape model for segmentation and shape modeling of 3-D brain structures. *IEEE Transactions on Medical Imaging*, 20(4):257–270, April 2001.
- [101] P. Shi, G. Robinson, and J. Duncan. Myocardial motion and function assessment using 4D images. In *In Proceedings of SPIE Conference on Visualisation in Biomedical Computing*, volume 2359, pages 148–159, 1994.
- [102] J. Sijbers, T. Ceulemans, and D. Van Dyck. Algorithm for the computation of 3D fourier descriptors. In *Proceedings of the 16<sup>th</sup> International Conference on Pattern Recognition*, Quebec, Canada, 2002.
- [103] J. G. Sled, A. P. Zijdenbos, and A. C. Evans. A non-parametric method for automatic correction of intensity non-uniformity in MRI data. *IEEE Transactions in Medical Imaging*, 17:87–97, 1998.
- [104] S. Smith. Fast robust automated brain extraction. *Human Brain Mapping*, 17(3):143–155, 2002. Software available for download at <http://www.fmrib.ox.ac.uk/fsl/bet/index.html>.
- [105] M. Sonka, V. Hlavac, and R. Boyle. *Image Processing, Analysis and Machine Vision*. Chapman and Hall, London, 2nd edition, 1999.
- [106] C. Studholme, D. L. G. Hill, and D. J. Hawkes. An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognition*, 32:71–86, 1999. Software available for download from <http://www.imageregistration.com>.
- [107] M. Styner and G. Gerig. Three-dimensional medial shape representation incorporating object variability. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, pages 651–656, 2001.

- [108] G. Subsol, J. Thirion, and A. N. A scheme for automatically building three-dimensional morphometric anatomical atlases: Application to a skull atlas. *Medical Image Analysis*, 2(1):37–60, 1998.
- [109] D. Sunday. <http://www.geometryalgorithms.com>.
- [110] J. Talairach and P. Tournoux. *Co-planar Stereotaxic Atlas of the Human Brain: 3-Dimensional Proportional System - an Approach to Cerebral Imaging*. Thieme Medical Publishers, New York, NY, 1988.
- [111] G. Taubin. A signal processing approach to fair surface design. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 351–358. ACM Press, 1995.
- [112] D. Terzopolous and D. Metaxas. Dynamic models with local and global deformations: Deformable Superquadrics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7):703–714, 1991.
- [113] D. Terzopoulos, A. Witkin, and M. Kass. Symmetry-seeking models for 3-D object reconstruction. *International Journal of Computer Vision*, 1(3):211–221, 1987.
- [114] N. A. Thacker, A. Jackson, and D. Moriarty. Improved quality of re-sliced MR images using re-normalized sinc interpolation. *Journal of Magnetic Resonance Imaging*, 10:582–588, 1999.
- [115] J. P. Thirion. The extremal mesh and the understanding of 3D surfaces. Technical Report RR-2149, INRIA, 1993.
- [116] J. P. Thirion and A. Gourdon. Computing the differential characteristics of isointensity surfaces. *Computer Vision and Image Understanding*, 61(2):190–202, 1995.
- [117] J. P. Thirion and A. Gourdon. The 3D marching lines algorithm. *Graphical Models and Image Processing*, 58(6):503–509, 1996.
- [118] P. Thompson and A. W. Toga. A surface based technique for warping three-dimensional images of the brain. *IEEE Transactions on Medical Imaging*, 15(4):402–417, 1996.
- [119] P. M. Thompson, M. S. Mega, K. L. Narr, E. R. Sowell, R. E. Blanton, and A. W. Toga. Brain image analysis and atlas construction. In I. Bankman, editor, *Chapter 17, Handbook of Medical Image Processing and Analysis*. Academic Press, 2000.

- [120] P. M. Thompson and A. W. Toga. Detection, visualisation and animation of abnormal anatomic structure with a deformable probabilistic brain atlas based on random vector field transformations. *Medical Image Analysis*, 1(4):271–294, 1996/7.
- [121] Tina. An open source image analysis environment. <http://www.tina-vision.net/index.php>.
- [122] A. W. Toga. *Brain Warping*. Academic Press, San Diego, California, USA, 1999.
- [123] A. W. Toga, T. P. M., and P. B. A. Modeling morphometric changes of the brain during development. In R. W. Thatcher, G. R. Lyon, T. Rumsey, and N. Krasnegor, editors, *Chapter in: Developmental Neuroimaging: Mapping the Development of Brain and Behavior*. Academic Press, New York, 1996.
- [124] G. M. Treece, R. W. Prager, A. H. Gee, and L. Berman. Surface interpolation from sparse cross-sections using region correspondences. *IEEE Transactions on Medical Imaging*, 19(11):1106–1114, 2000.
- [125] M. Tso, P. Kleinschmidt, I. Mitterreiter, and J. Graham. An efficient transportation algorithm for automatic chromosome karyotyping. *Pattern Recognition Letters*, 12:117–126, 1991.
- [126] C. J. Twinning and C. J. Taylor. Kernel principal component analysis and the construction of non-linear active shape models. In T. F. Cootes and C. J. Taylor, editors, *Proceedings of the 12<sup>th</sup> British Machine Vision Conference*, volume 1, pages 23–32, September 2001.
- [127] E. A. Vorkuka, N. A. Thacker, and J. A. A fast model independent method for automatic correction of intensity nonuniformity in MRI data. *Journal of Magnetic Resonance Imaging*, 10(4):550–562, 1999.
- [128] L. Wang, H. Ning, W. Hu, and T. Tan. Gait recognition based on procrustes shape analysis. In *Proceedings of the 9<sup>th</sup> IEEE International Conference on Image Processing*, Rochester, New York, 2002.
- [129] W. M. Wells, W. E. L. Grimson, R. Kilinis, and F. A. Jolesz. Adaptive segmentation of MRI data. *IEEE Transactions in Medical Imaging*, 15(4):429–442, 1996.
- [130] C. G. Wible, M. E. Shenton, H. Hokama, R. Kikinis, F. A. Jolesz, D. Metcalf, and R. W. McCarley. Prefrontal cortex and schizophrenia. A quantitative magnetic resonance imaging study. *Archives of General Psychiatry*, 52:279–288, 1995.

- [131] G. Williams. *The Age of Miracles: Medicine and Surgery in the Nineteenth Century*. Academy Chicago, 1987.
- [132] A. B. Wolbarst. *Physics of Radiology*. Prentice-Hall, 1993.
- [133] P. W. Woodruff, G. D. Pearlson, M. J. Geer, P. E. Barta, and H. D. Calcite. A computerised magnetic resonance imaging study of corpus callosum morphology in schizophrenia. *Psychological Medicine*, 23:45–56, 1993.
- [134] R. P. Woods, S. T. Grafton, C. J. Holmes, S. R. Cherry, and J. C. Mazziotta. Automated image registration: I. general methods and intrasubject, intramodality validation. *Journal of Computer Assisted Tomography*, 22:139–152, 1998. Software available for download at <http://bishopw.loni.ucla.edu/>.
- [135] A. J. Worth, N. Makris, V. S. Caviness, and D. N. Kennedy. Neuroanatomical segmentation in mri: Technological objectives. *International Journal of Pattern Recognition and Artificial Intelligence*, 11(8):1161–1187, 1997.
- [136] M. F. Wu and H. T. Shen. Representation of 3D surfaces by two-variable fourier descriptors. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(8):858–863, 1998.
- [137] D. Zhang. *Harmonic Shape Images: A 3D Free-form Surface Representation and Its Applications in Surface Matching*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, November 1999.
- [138] H. Ziezold. On expected figures and a strong law of large numbers for random elements in quasi-metric spaces. In *Transactions of the 7<sup>th</sup> Prague Conference on Information Theory, Statistical Decision Functions, Random Processes and of the 1974 European Meeting of Statisticians*, volume Volume A, pages 591–602, Prague, 1974.