



Department of Engineering Science, University of Oxford

Morphometric Analysis of Brain Structures in MRI

Miguel Ángel González Ballester

(Green College)

Under the supervision of:

Prof. J. Michael Brady

Prof. Andrew P. Zisserman

A thesis submitted to the Department of Engineering Science, University of Oxford, in partial fulfilment of the requirements for the degree of Doctor of Philosophy

Michaelmas Term 1999

Morphometric Analysis of Brain Structures in MRI

A thesis submitted to the Department of Engineering Science, University of Oxford, in partial fulfilment of the requirements for the degree of Doctor of Philosophy – Michaelmas Term 1999

Miguel Ángel González Ballester (Green College)

Under the supervision of Prof. J. Michael Brady and Prof. Andrew P. Zisserman

Abstract

Medical computer vision is a novel research discipline based on the application of computer vision methods to data sets acquired via medical imaging techniques. This work focuses on magnetic resonance imaging (MRI) data sets, particularly in studies of schizophrenia and multiple sclerosis. Research on these diseases is challenged by the lack of appropriate morphometric tools to accurately quantify lesion growth, assess the effectiveness of a drug treatment, or investigate anatomical information believed to be evidence of schizophrenia. Thus, most hypotheses involving these conditions remain unproven.

This thesis contributes towards the development of such morphometric techniques. A framework combining several tools is established, allowing for compensation of bias fields, boundary detection by modelling partial volume effects (PVE), and a combined statistical and geometrical segmentation method. Most importantly, it also allows for the computation of confidence bounds in the location of the object being segmented by bounding PVE voxels. Bounds obtained in such fashion encompass a significant percentage of the volume of the object (typically 20-60%).

A statistical model of the intensities contained in PVE voxels is used to provide insight into the contents of PVE voxels and further narrow confidence bounds. This not only permits a reduction by an order of magnitude in the width of the confidence intervals, but also establishes a statistical mechanism to obtain probability distributions on shape descriptors (e.g. volume), instead of just a raw magnitude or a set of confidence bounds. A challenging clinical study is performed using these tools: to investigate differences in asymmetry of the temporal horns in schizophrenia. This study is of high clinical relevance. The results show that our tools are sufficiently accurate for studies of this kind, thus providing clinicians, for the first time, with the means to corroborate unproven hypotheses or reliably assess patient evolution.

Acknowledgements

*A mi familia: Antonio, Agustina, Carlos,
José Antonio, María José y Ágata.*

I have had the great honour of being supervised by two of the brightest and most exceptional persons I have ever met. Mike and Andrew were always accessible and helpful, with a never-ending stock of brilliant ideas, good advice, movie reviews, and great drinking stories! All my gratitude and admiration is for them, who made this thesis possible.

Several people have directly contributed to the work described in this thesis. In particular, I am very grateful to Dr. Hervé Delingette and Dr. Régis Guillemaud for sharing their code and providing very useful advice and discussion. Also, I would like to thank Dr. Patrick Marais for taking time off his thesis to help me on my first steps in this field, completely novel to me at the time. My gratitude to Prof. Alan Colchester and Dr. Fernando Bello for sharing their MRI phantom data set for validation of our technique. The data sets used for our clinical study are part of the BIOMORPH data pool, and were acquired at the State University of New York. Many thanks to Dr. Lynn Delisi for sharing them. I would also like to thank Dr. Paul Hayton for allowing me to use his GUI for the bias correction program.

Medical background and motivation was provided by the MRC Schizophrenia Research Group at the Radcliffe Infirmary, headed by Prof. Tim Crow. They helped driving my research towards clinically useful applications. Special thanks to Steven Chance and Dr. Rebecca Craven for their help in uncountable times.

During the three years of this DPhil I have had the pleasure to meet some of the leading researchers in the field. Many of them have provided very useful ideas, which have influenced in several ways the realisation of this thesis. Amongst them, I would like to thank Prof. Guido Gerig, Prof. Nicholas Ayache, Prof. James Duncan, Dr. Jacques Feldmar, Dr. Stephen Smith and Dr. Xavier Pennec. Also, thanks to all BIOMORPH partners for very useful discussion and enjoyable meetings.

My labmates at the Medical Vision Lab have been my family for the last three years, and I would like to thank them for their friendship. A couple of books like this one could be filled with reasons why I am grateful to each one of them and how they helped me with my DPhil. Instead, I will just name them: Alison, Ralph, Annette, Patrick, Seb, Dale, Gary, Paul H., Miguel, Maud, Andrew, Jacques, Gerardo, Kostas, Margaret, Albert, Gabriel, Guofang, Yongyue, Yasuyo, Djamal, Jerome, Chris, Robert, Paul G., Michael, Chus, Marius and Xujiong. Looking back at this list of names in the future will always bring me great memories from my times in Oxford.

I have met many wonderful people and have been blessed with their friendship. I would like to thank them for their support and companionship during these three years. Some of them were already mentioned in the list above, which goes to prove that the atmosphere at the lab has been fantastic. For the rest, even before I write it, I am sure that I am going to forget to mention quite a few people. To them, I would ask them to understand that my brain is recovering from being fried in front of a computer for three years, and that the fact that their names did not come into my mind at the moment I was writing this is to blame to all the consequent confused neurons in my brain. My love and gratitude to Blanca, Jorge, Alex Mira, Alex Ribés, Guillermo, Veit, Stefan, Jasmine, Xaro, Marian, Alexandra, Criel, Pilar, Ralph, Hans, Darren, Bernhard, Ana, Gustavo, Ismael, Izumi and Marta. Last but not least, thanks to Prof. Vicent Cervera, Dr. Mar Marcos, and Prof. Ángel P. del Pobil, from the Universitat Jaume I of Castellón.

Table of contents

1. INTRODUCTION	1
1.1 Computer vision meets medicine	1
1.2 Need for morphometric tools for neurological studies	3
1.2.1 Brain morphometry	3
1.2.2 Schizophrenia	4
1.2.3 Multiple sclerosis	7
1.3 Objectives of this thesis	9
1.4 Overview	10
2. BACKGROUND	12
2.1 Introduction	12
2.2 Magnetic resonance imaging (MRI)	12
2.2.1 Basics	12
2.2.2 Image formation	13
2.2.3 Relaxation times and tissue characterisation	15
2.3 Processing steps in medical computer vision	17
2.4 Image enhancement: application to MRI	19
2.4.1 Inhomogeneity correction	19
2.4.2 Compensation for partial volume effects	22
2.5 Segmentation	24
2.5.1 Voxel classification methods	25
2.5.2 Shape-based methods	25
2.5.3 Mixed methods	27
2.6 Shape representation	27
2.6.1 Voxel-based representations	28
2.6.2 Landmark-based methods	29
2.6.3 Parametric models	29
2.6.4 Meshes	30
2.6.5 Other shape models	31
2.7 Registration	32
2.8 Shape description	32
2.9 Summary	34
3. A FRAMEWORK FOR MORPHOMETRIC STUDIES OF MRI INCLUDING CONFIDENCE BOUNDS	35
3.1 Introduction	35
3.2 Bias correction and voxel-based segmentation: the EM algorithm	37
3.2.1 Basic formulation	37
3.2.2 Improvements and considerations	39
3.3 Detection of PVE voxels	40
3.4 Shape-based segmentation: the simplex mesh	41
3.5 Global segmentation scheme	44
3.6 Modelling complex biological shape	46
3.7 Estimation of confidence bounds	48
3.8 Shape description: volume computation	49
3.9 Validation issues	51

3.9.1	Validation with synthetic data sets	52
3.9.2	Validation with MRI phantoms	53
3.9.3	Validation with clinical data: repeatability	53
3.10	Validation results	54
3.10.1	Results for synthetic data sets	54
3.10.2	MRI phantoms	56
3.10.3	Results using clinical data	58
3.11	Flexibility of the shape model: representing the cortical surface	60
3.12	Discussion	62
4. ESTIMATION OF THE PARTIAL VOLUME EFFECT		63
4.1	Introduction	63
4.2	Mixture model	63
4.3	Example: Gaussian distributions	67
4.4	Confidence bounds	69
4.5	Propagation of uncertainty	70
4.5.1	Conservative bounds	70
4.5.2	Monte Carlo	71
4.6	Validation with synthetic images	72
4.7	Discussion and practical considerations	75
4.8	Sensitivity to tissue parameters: modelling uncertainty in parameter estimation	75
4.9	Use of contextual information to determine the contents of PVE voxels	77
4.10	Construction of PVE masks from inner and outer surface bounds	79
4.11	Experiments on MRI phantoms	82
4.12	Generation of PVE and intensity images from surface models	84
4.13	Conclusions	85
5. CLINICAL STUDY: SYMMETRY DIFFERENCES OF THE TEMPORAL HORNS IN SCHIZOPHRENIA		87
5.1	Introduction	87
5.2	Anatomical background	88
5.3	Related clinical studies	90
5.4	Problem statement	91
5.5	Data used in the study	92
5.6	Methods	93
5.6.1	Bias field correction	93
5.6.2	Determination of contents of PVE voxels	94
5.6.3	Identification of regions of interest	95
5.6.4	Voxel-based volume bounds	98
5.6.5	PVE analysis: conservative bounds and Monte Carlo simulation	98
5.7	Results	99
5.8	Analysis of results	102
5.9	Clinical conclusions	105
5.10	Discussion	107
6. CONCLUSIONS AND FUTURE WORK		109
6.1	Contributions of this thesis	109
6.2	Future work	110
6.2.1	Improved tissue models	110
6.2.2	Priors for PVE prediction	111

6.2.3	Use of anatomical knowledge	111
6.2.4	Surface fitting based on PVE constraints	112
6.2.5	Developing the concept of “statistical shape”	113
6.2.6	Anatomy-based global parameterisation	115
6.2.7	New shape descriptors for morphometric studies	116
6.2.8	Study of sampling limitations	116
6.3	Concluding remarks	117
APPENDIX A. TRIANGULAR GREGORY-BÉZIER PATCHES		118
A.1	Definition	118
A.2	G¹-continuity constraints	119
A.3	Determination of control points situated along the boundaries	121
A.4	Determination of the inner control points	123
APPENDIX B. ALGORITHMS		126
B.1	Oct-tree recursive subdivision algorithm	126
B.2	Quad-tree recursive subdivision algorithm	127
APPENDIX C. CLINICAL STUDY: TISSUE PARAMETERS		128
APPENDIX D. ANATOMY-BASED GLOBAL PARAMETERISATION		130
D.1	Introduction	130
D.2	Correspondence between global parameters and local patch coordinates	132
D.3	Re-parameterisation based on anatomical landmarks	133
D.4	Extension to volumetric objects	133
BIBLIOGRAPHY		135

Introduction

1.1 Computer vision meets medicine

The times when computer science was a purely endogamic discipline, dedicated solely to its own maturation, are long gone. The reader will find proof of this in every aspect of daily life. Computers have become an essential part of our society, not only in highly industrialised countries, but throughout the world. It is therefore not surprising that new disciplines are emerging from the application of computer engineering concepts to all other sciences and arts. The synergy created by such association is proving more and more successful every day, working in a symbiotic fashion towards the growth of both computer science and the whole body of *host* sciences. The area of study of this thesis falls into this multidisciplinary category, and in particular this work contributes towards what is known as *medical computer vision*.

The human, social, and economic relevance of medicine is evident, so it is obvious that an effort towards the application of computers in a medical setting is fully justified. In a sense, computers can provide more accuracy and speed than humans, thus providing the hand that drives the scalpel to the perfect location, the memory that remembers immensely larger amounts of data and retrieves them at faster speed, and the eye that can see closer into the human body. This last field, vision, is the subject of the present thesis.

The advent of medical imaging techniques such as X-rays, ultrasound, or magnetic resonance imaging (MRI) – to mention a few – provided the physician with tools to *see through* the human body. These developments unleashed a whole spectrum of new diagnostic and treatment possibilities. The power to inspect inside the human body *in-vivo* meant a considerable reduction on the number of biopsies, and offered the potential to investigate pathologies that were not fully understood to date.

The concept of *computer vision* was established in the 1960's as a discipline aimed at providing computers with an *understanding* of visual data. This embodies a large set of operations, which can be roughly classified in a bottom-up multilevel fashion as follows. Image enhancement is applied to process the image to remove noise or enhance contrast. Low-level vision aims at detecting elementary patterns in the image, such as lines or corners. High-level vision is dedicated to the analysis of such elementary patterns in order to reconstruct the global structure of the image and identify the objects present in it. Finally, at more abstract levels, artificial intelligence can be employed to provide further understanding of the information contained in the image and, if required, take a line of action (maybe based on prior knowledge). After a slow start, computer vision has taken off. Early developments were necessarily basic, both in the nature of the techniques employed and the scope of application. However, once these basic foundations were laid the subject matured at a fast pace, and applications of computer vision to most common daily activities are either already present or under way in the very near future.

Medical computer vision can be defined as the application of computer vision processes to medical data sets acquired using any medical imaging technique. The added value obtained from such association aims not only at reducing the time needed by the clinician to examine patient scans, but also at providing useful additional information to help inform a diagnosis. The nature of this complementary information can, for example, consist of highly accurate measurements of the anatomical structure or lesion of interest. Thus, it is not the aim of medical computer vision to *replace* the clinician, but rather to *complement* his/her role.

This thesis contributes towards this end by establishing new computer vision tools for the analysis of medical images. In particular, it focuses on the problem of detecting and measuring organs or lesions accurately and with explicit and tight bounds on the error. Although the techniques developed here have wider application, we will pay particular attention to neurological studies based on magnetic resonance images of the brain. Further details about such applications are given next, and the rest of the chapter establishes the objectives of this thesis and provides an outline of its contents.

1.2 Need for morphometric tools for neurological studies

1.2.1 Brain morphometry

In order to introduce the work developed in this thesis, it is first necessary to take a closer look at the target clinical applications. An engineering approach must, of necessity, start by clearly stating some objectives and identifying requisites and difficulties. The present work is located in the area of *brain morphometry*. The word morphometry means “measurement of shape”. This concept has two important implications. First, the concept of *shape* must be defined in a quantitative manner suitable for measurement. Shape descriptors will depend on the application at hand, but typical magnitudes such as volume, area or length can be employed as general shape descriptors. Second, the concept of *measurement* suggests that any given morphometric method must pay particular attention to issues such as accuracy, precision, and error analysis.

In particular, the techniques developed in this thesis aim at providing morphometric tools for neurological studies of the brain. The particular anatomy and sensitivity of the human brain is a perfect example of how medical imaging techniques have helped to better understand the human body. Prior to non-invasive techniques such as MRI or MRS (Magnetic Resonance Spectroscopy), diagnosis of neurological disease was extremely difficult, and life-threatening biopsies were frequently carried out. Because of its good contrast of soft tissues, MRI has helped to understand brain anatomy. Imaging evidence of a multitude of pathologies has been reported (Rodriguez, 1996), proving the effectiveness of MRI as a valuable diagnostic technique.

However, the analysis of certain clinical conditions remains difficult, mainly due to the fact that the anatomical evidence is very small relative to the imaging resolution. Advanced morphometric tools are then required to analyse these data and provide accurate measurements. Existing tools are not well suited to such a task, and new morphometric techniques for brain studies must be developed. This is the main objective of this thesis.

It should be noted that the social and economic relevance of this research is paramount. Diseases such as schizophrenia and multiple sclerosis, for example, affect millions of people world-wide, and necessitate enormous amounts of money for hospital beds, medical assistance, and drug

treatment. The critical lack of morphometric techniques was recognised by the European Union, which in 1996 established the research project BIOMORPH, under the framework of the BIOMED2 programme. This project involved collaborators from Switzerland (ETH Zurich), France (INRIA Sophia-Antipolis), Belgium (KU Leuven), and Great Britain (University of Kent at Canterbury and University of Oxford). The work described in this thesis was developed as a contribution to the BIOMORPH project, aimed at the establishment of new morphometric techniques for brain studies. Although the techniques are of a more general use, primary focus was set on applications to schizophrenia and multiple sclerosis. The particularities of these two diseases are introduced next. Unless otherwise stated, the main clinical reference for the next two sections is (Souhami and Moxham, 1994).

1.2.2 Schizophrenia

Just under 1% of the population develop schizophrenia at some point in their lives, a statistic that appears to hold true for all cultures and countries (Souhami and Moxham, 1994; Harrison, 1998). In the USA, 25% of all hospital beds are occupied by schizophrenic patients, and the total direct and indirect economic costs are approximately \$33 billion (Harrison, 1998). Regarding the UK, until the move towards community care began, about half of all hospital beds were in mental hospitals, and the majority of these were occupied by patients with a schizophrenic illness (Souhami and Moxham, 1994).

Schizophrenia usually starts in young adulthood, and has a profound effect both on mental processes and social interaction. Mental disturbances (see Figure 1) can be expressed as a disruption of the normal progression of logical thought (“thought disorder”). Auditory hallucinations are frequent, and they usually involve the voice of a “third person” commenting on the patient’s actions. Delusions about an external agent interfering with the patient’s thoughts by inserting, blocking, or transmitting them to other people are also common, as well as delusional perceptions, in which an apparently irrelevant fact acquires special and often frightening significance to the patient. Socially, patients tend to withdraw from social contact and spend many hours in isolation, and may give up their jobs and shun their friends.

- Delusional perception
- Auditory hallucinations
 - Audible thoughts (thought echo)
 - Voices arguing or discussing (“third person”)
 - Voices commenting on the patient’s actions
- Thought disorder
 - Thought withdrawal
 - Thought insertion
 - Thought broadcasting
- Passivity experiences: delusions of control
 - “Made” feelings (includes somatic hallucinations)
 - “Made” actions
 - Somatic passivity (body invaded from outside)

Figure 1 Mental disturbances characteristic to schizophrenia. This classification is known as Schneider’s first-rank symptoms.

About 30% of patients only have one episode, whereas 15% remain severely disabled and will still be in hospital a year later. Chronically mentally ill patients are severely handicapped. Their prospects of employment are poor, they may have little or no family support, and their capacity to care for and occupy themselves may be very limited.

The cause of the illness is unclear, and it has even been suggested that schizophrenia is not a single entity, but a group of related conditions. Both genetic and environmental factors seem to be important in the aetiology of schizophrenia. This is corroborated by studies showing that there is higher probability of being affected by the illness if relatives suffer the condition, and higher predisposition to a second relapse in situations where the patient lives in a hostile family environment.

Anatomical evidence of schizophrenia, and in particular its relation to brain asymmetry, has been reported in several studies (Johnstone et al., 1976; Crow, 1990; Bilder et al., 1994; DeLisi et al., 1995). The prevailing hypothesis (Crow et al., 1989; Crow, 1995) relates anatomical and genetic findings, proposing that schizophrenia results from a failure to develop the normal lateral (left-right) asymmetry which resulted from evolutionary processes of functional specialisation of brain structures. Normal brains show larger left occipital and temporal lobes, corresponding to language

areas, and this asymmetry is reversed in the frontal lobe (Figure 2). Results have been reported for cortical volumes and internal structures, showing a reduction of this asymmetry in schizophrenic brains.




Figure 2 Normal brain asymmetry is reduced in schizophrenia. This is illustrated on a top view of the brain. Human brains present larger left volumes in the posterior (occipital) area, corresponding to the zones of language specialisation, while the opposite holds in the anterior areas of the brain. Reduction or even reversal of this asymmetry has been reported in several studies.

Unfortunately, in many of the published studies the measurement error is at least as large as the effect being reported, so the ideas of lack of asymmetry sketched above remain hypotheses. Techniques used for these clinical studies include water displacement and measurement of lengths in (not necessarily aligned) slices of computer tomography (CT) or MRI. Additionally, the number of subjects is usually very limited, and some studies worked on *post-mortem* brains, which tend to deform considerably when extracted from the skull, so it is difficult to reliably measure their asymmetry. The techniques developed in this thesis aim at providing accurate error-bounded results for *in-vivo* studies. These techniques can be deployed to prove or refute current hypotheses about schizophrenia, as shown in the clinical study reported in Chapter 5.

1.2.3 *Multiple sclerosis*

Neurons, or nerve cells, are excitable cells which are specialised for the reception, integration, transformation and onward transmission of coded information. Neurons consist of a cell body or *soma*, from which branching structures emerge. These structures may be immediate connections to the soma, called *dendrites*, or conducts away from the cell body or *axons*. Neurons communicate with each other through *synapses*, which tend to be axodendritic or, in some instances, axosomatic. Coarsely, the brain structure can be thought of as an interconnection switchboard where the active elements (the somae and dendritic trees) are located on the outer surface (the cortex), and the interior consists of axons interconnecting these somae in complex patterns. Axons are laminated with a lipo-protein sheath of *myelin*. These sheaths are essentially concerned with the functional efficiency of the nerve fibres during impulse conduction, in terms of an optimal conduction velocity for a given diameter, and minimal energy requirements. They impart a pinkish-white opalescence to the fresh tissue – hence the name *white matter* (Williams and Warwick, 1980). Multiple sclerosis (MS) is a demyelinating disease, meaning that its primary pathological process is the destruction of the myelin laminae around the axons, without further axonal degeneration (Souhami and Moxham, 1994).

MS is an illness of unknown cause, in which discrete areas of demyelination develop at many sites of the brain and spinal cord. Lesions, or *plaques*, develop in different sites at different times, usually with some capacity for regeneration and restoration of function. This leads to the characteristic relapsing and remitting history in many patients; in others, a slowly progressive deficit occurs. Since plaques may appear at any site in central white matter, the clinical manifestations of MS are extremely variable. Lesions affecting the optic nerves cause a dimming of vision and pain around the eye; involvement of the cervical spinal cord causes motor, sensory, bladder, bowel and sexual disturbances. Vertigo is also common. Depression in long-standing MS is frequent and requires vigorous treatment. Widespread cerebral hemisphere demyelination leads to intellectual impairment. In about 20% of patients there is no significant disability after 5 years, and the average life expectancy overall from onset of symptoms is 20-30 years. In about 5% of patients, the disease is rapidly progressive and fatal within 5 years.

MS presents a strong geographical variation, being essentially a disease of temperate climates. It is rare in childhood and uncommon in early adolescence. There is an increasing incidence with age, peaking at about 30 years, and it is uncommon over the age of 50. Slightly more women than men are affected by MS. It is the most common progressive disabling neurological disease affecting young adults in the UK. In the USA 350,000 people suffer from MS, and with the exception of trauma it is the most frequent cause of neurologic disability (Harrison, 1998). Clinical drug trials are being carried out and mobilise enormous amounts of economic and human resources. To date, there was no specific treatment for MS. However, a drug called β -InterferonTM is in the last stages of trial and is very likely to obtain approval in the very near future.

MRI shows good contrast of MS lesions¹ (see Figure 3), but the size of the plaques is usually very small relative to the image resolution. Accurate detection and quantification of lesion size is crucial for the study of patient evolution and the analysis of the effect of tested drug treatments, since it reduces the subjectivity in their diagnosis. Therefore, the development of advanced morphometric techniques is urgently needed, and could lead to substantial savings in time, money and lives.




Figure 3 Axial slice of a P_D MRI of the brain of a patient affected by multiple sclerosis. Plaques have been marked with circles. Note that MS lesions appear as bright spots, whereas white matter shows dark intensities. This is due to the loss of myelin in the plaques.

¹ P_D and FLAIR are sequences commonly used for imaging of MS lesions. It is also common to use contrast agents, such as GadoliniumTM. MRS is also heavily used for MS studies.

1.3 *Objectives of this thesis*

The previous section sketched the clinical applications of the techniques to be developed, and posed some of the problems inherent in these clinical requirements. We now provide a more formal statement of the problem to be solved in this work. In a nutshell, this thesis is aimed at providing computer tools for morphometric studies, with a special focus on schizophrenia and MS. The key requirement is that these tools must provide sufficient accuracy in order to prove or disprove hypotheses about the aetiology of schizophrenia and allow for a close follow-up of MS patients.

This general problem can be subdivided into *segmentation*, or detection of interesting features such as MS plaques or anatomical structures of interest for the study of schizophrenia, and *measurement* or characterisation of their shape. Although the field of medical computer vision is young, there is a relatively abundant literature concerning segmentation. On the other hand, the problem of accuracy, precision, and error analysis of measurements has had much less attention, and the same holds for the establishment of appropriate shape descriptors for features of interest in particular clinical studies, such as asymmetry in the case of schizophrenia.

The lack of error analysis schemes for measurements obtained from medical computer vision techniques is especially remarkable. In a field where accuracy is paramount, most methods provide a raw magnitude (such as volume or area) and no notion of the error contained in the measurement. Apparatus such as confidence bounds should be incorporated, and are one of the key elements in the tools presented in the following chapters.

Another important issue is validation. The nature of the data we are dealing with poses the problem of lack of ground truth. Validation on post-mortem MR scans – where ground truth measurements can be made – is possible but prone to error, since brains taken out of the skull deform considerably and the chemicals used to fixate them affect the imaging process. Therefore, validation schemes should be developed. This issue is another focus of especial attention in this thesis.

Although this research has a marked technical inclination, it should be taken into account that any engineering work is motivated and aimed at a practical application. The employment of our work to medical purposes will be demonstrated by performing a clinical study about the anatomical evidence of schizophrenia.

1.4 Overview

Chapter 2 provides background information about MRI and medical computer vision. The principles of magnetic resonance and image formation are reviewed in section 2.2. The remainder of the chapter sketches the processing steps commonly followed in medical computer vision and provides literature review for them.

Chapter 3 describes an algorithmic framework for morphometric studies from MRI data. This framework includes correction for bias fields, detection of salient features based on identification of partial volume effects (PVE), segmentation, shape modelling, and shape description (exemplified by volume computation). The segmentation method is novel in that it combines statistical classification and geometry-based segmentation through active surfaces. Special attention is paid to establishing confidence bounds on the segmentation and subsequent measurements. This is achieved by creating a dual segmentation bounding PVE voxels. Thus, the result of the process is an *inner* and an *outer* surface bounding the real location of the surface of the object being segmented. An extensive validation study is performed, and the crucial problem of lack of ground truth measurements for validation is addressed. The particular choice of shape model, consisting of a G^1 -continuous surface based on triangular patches, provides a flexible means for representing biological shape, and this is demonstrated by segmenting the cortical surface from a data set. Additionally, the continuity inherent in such surface model constrains the shape of the surface and further narrows the confidence interval between inner and outer segmentations.

Such interval, delimiting PVE voxels, commonly accounts for a large portion of the volume of the object being segmented (typically 20-60%). Chapter 4 takes on the task of decoding the information contained in such PVE voxels in order to further narrow confidence bounds. A

statistical model of PVE voxels is created and an estimation technique established. Robustness is ensured by explicitly modelling the uncertainty on the parameters of the PVE model. Auxiliary methods to detect PVE voxels and determine their composition are defined. The statistical nature of our technique provides means to propagate the uncertainty on the estimation of the composition of individual PVE voxels, in order to create statistical distributions of shape descriptors (e.g. volume). Results are provided showing that the PVE estimation framework described in this chapter allows for a reduction of confidence intervals by more than an order of magnitude with respect to the ones obtained in the previous chapter (typical values are around 1-2%).

The usefulness of the techniques described in this thesis is demonstrated by a challenging clinical study, presented in Chapter 5. This study focuses in the temporal horns, a very thin structure part of the lateral ventricles (section 5.2 provides anatomical background). The importance of such structure in schizophrenia has been pointed in several studies. However, the lack of resolution of current morphometric techniques has not allowed for conclusive results to be established, due to the very small size of the temporal horns with respect to the voxel size. We show that the use of our techniques makes this problem, intractable to date, feasible (confidence bounds on volume measurements are around 6%, as opposed to 200% for voxel-based methods). This study is not only a validation for our techniques, but also an important clinical contribution to schizophrenia research.

Finally, Chapter 6 provides conclusions and ideas for future work. Several appendices are also included for clarity, and provide the necessary detail on particular issues.

Background

2.1 Introduction

This chapter provides some necessary background information. In particular, section 2.2 deals with MRI, describing the basics of its functionality and how images are generated, and the rest of the chapter consists of a literature review of the state of the art in medical computer vision. Section 2.3 lists the typical computer vision processes applied to medical data sets, and the following sections provide more detail for each of these processing steps.

2.2 Magnetic resonance imaging (MRI)

2.2.1 Basics

The magnetic properties of the nuclei of atoms commonly present in human tissues, such as ^1H , ^{13}C , or ^{19}F constitute the basic principle of MRI. In particular, these nuclei (^1H , commonly called a proton, is the most abundant and hence the one most frequently used for imaging purposes), when embedded in a magnetic field, can absorb the energy of radio-waves of characteristic frequencies. The exact frequency, known as the Larmor frequency, depends on the type of nucleus, the field strength, and the physical and chemical environment of the nucleus, and is given by the Larmor equation (Rinck, 1993). The absorption of the energy of such radio-waves causes a switch from the initial low-energy state (parallel to the magnetic field) to a higher, anti-parallel, energy state. This is the basic phenomenon utilised in MR imaging. The proportion of high and low energy spins in a sample determines its net magnetisation.

The patient is placed inside a static, stable and highly homogeneous magnetic field \mathbf{B}_0 . A radio-frequency (RF) transmitter is used to generate short pulses of radio-waves of an appropriate frequency to excite the nuclei of interest. These nuclei absorb the energy provided by the radio-

waves and liberate it when the transmitter is turned off. The released energy is captured by a receiver coil and measured, reflecting the proton density (P_D) inside each voxel. Localisation of each voxel is achieved by modifying the field \mathbf{B}_0 with the addition of gradient fields that contribute to \mathbf{B}_0 linearly in one direction. Three gradients are used to encode all three spatial dimensions inside the volume. Thus, the resulting magnetic field varies spatially, and by modifying the magnetic field at each point of the volume, the frequencies absorbed by each particular location inside it also vary linearly. This allows for the origin of the signals to be identified, and makes it possible to determine the net magnetisation \mathbf{M}_0 at each spatial location. Images can then be generated, as explained in the following section.

2.2.2 *Image formation*

A *slice* through the imaged volume is selected by applying a gradient field in the z direction, i.e. parallel to the main field \mathbf{B}_0 . The contribution of this gradient produces a linear increase of the field present in different locations of the patient (Figure 4). Only certain field strengths are excited by the subsequent RF pulse, as stated by the Larmor equation. The strength and duration of the application of the gradient field determines the excited area within the sample, i.e. the slice thickness. Changes in the frequency of the gradient field bias it in the z direction, so allowing the selection of the desired slice.




Figure 4 Slice selection. A gradient field is used to modify linearly the main magnetic field \mathbf{B}_0 . Only a slice through the patient will be excited, in virtue of the Larmor equation.

Phase is used to encode the y direction inside each slice. A short time after a magnetic field has been applied to a sample, the spins start dephasing (cf. definition of T_2 in section 2.2.3). In the presence of a gradient magnetic field, this dephasing is dependent on the spatial location of the spins. The phase angle induced by a gradient depends both on the duration and the amplitude of the gradient. Figure 5 shows how amplitude (and polarity) affects phase, assuming constant duration of the application of the gradient.




Figure 5 Effect of a gradient on the phase of the spins. Assuming constant duration of the application of the gradient, the phase angle (left) can be controlled by varying the polarity or amplitude of the gradient (right).

The *2D Fourier transform method* is used to locate voxels in a slice. First, a gradient in the y direction is switched on and turned off after a certain time, so modifying the phase of the sample at different locations. Then, a gradient in the x direction is used to complete the definition of the voxel by modifying the resonance frequency at different locations along the x axis. Thus, phase encodes y and frequency encodes x. Since the y gradient must be applied as a preparation step before the x gradient is used, the experiment must be repeated n times, where n is the desired resolution in the y direction (typical protocols use 256 acquisitions per slice). The amplitude of the signal for each location in the sampled volume is stored, and this representation is called *K-space*. One line of K-space is filled for each activation of the y gradient (Figure 6).

The K-space representation does not correspond directly to the image. In fact, it contains frequency information, in a fashion similar to the Fourier transform of an image. Positions near the centre of the matrix correspond to low frequencies, and points near the corners show high frequency information. A 2D inverse Fourier transform is used to convert this representation to an intensity image.




Figure 6 K-space. One line is filled during each activation of the y gradient.

2.2.3 Relaxation times and tissue characterisation

Two characteristic times, T_1 and T_2 , enable tissue type to be identified and give complementary information to that obtained from proton density. After a 90° RF pulse the net magnetisation of the sample is shifted by 90° , placing it on the x,y-plane (perpendicular to the direction of \mathbf{B}_0). When this pulse is turned off, the sample will progressively relax to the former stable position, that is, its magnetisation will be parallel to \mathbf{B}_0 . This recovery follows an exponential rate $\mathbf{M} = \mathbf{M}_0 e^{-t/T_1}$, where \mathbf{M}_0 is the initial magnetisation and T_1 is the *spin-lattice relaxation time*. Therefore, T_1 is the time required for the system to recover 63% of its equilibrium value after it has been exposed to a 90° pulse. The *spin-spin relaxation time*, T_2 , is related to the phenomenon known as transverse relaxation. After a spin has been excited by a RF pulse, it initially behaves in a

coherent fashion, i.e., all spins precess² in phase around the direction of the external field and add up to a net magnetisation \mathbf{M}_0 . However, as time passes, the net magnetisation starts to decrease due to a dephasing of the spins (see Figure 7). This is partially caused by small differences in the static magnetic fields at different locations of the samples, which determine slightly different Larmor frequencies (Rinck, 1993). T_2 characterises this decay in magnetisation due to dephasing of excited spins.




Figure 7 The contribution of all the spins precessing around the external magnetic field \mathbf{B}_0 produces a net magnetisation \mathbf{M}_0 . When a 90° RF pulse is applied, this net magnetisation is tipped onto the x,y-plane. Dephasing of the spins results in a quick decrease of the net magnetisation in the x,y-plane, characterised by T_2 . Such decay follows the equation: $\mathbf{M} = \mathbf{M}_0 e^{-t/T_2}$.

A number of acquisition protocols have been designed in order to be able to measure T_1 and T_2 for individual voxels. They consist of series of 90° and 180° pulses that excite the sample and permit monitoring of the decay in magnetisation due to return to equilibrium state (T_1) or dephasing of spins (T_2). Partial Saturation-Saturation Recovery and Inversion Recovery sequences are used to measure T_1 , while the Spin Echo Pulse Sequence can be employed to obtain T_2 . Cf. (Rinck, 1993) or (Webb, 1988) for a detailed explanation of these and other pulse sequences.

² When nuclei are embedded in a magnetic field, they align parallel to the direction of the field and exhibit a complex motion known as *precession*. It consists of a rotational movement around their own axis combined with a rotation around the axis of the field, in a similar fashion as a spinning top or ‘wobble’ (Rinck, 1993).

By modifying the parameters of these sequences (mainly times and strengths of pulses) contrast can be altered to enhance the delineation of particular tissues. In particular, there are three basic types of images: proton density (P_D), T_1 -weighted, and T_2 -weighted³. The information they convey is particular of certain tissue types (Rodriguez, 1996) and, to some extent, complementary. Figure 8 shows a P_D and a T_2 -weighted image of the same brain.




Figure 8 P_D and T_2 -weighted images of the same brain. P_D images give good white matter / grey matter contrast, but CSF shows intensities very similar to grey matter (notice the complicated delineation of the ventricles, due to the proximity of inner nuclei). T_2 shows excellent CSF contrast, but poor white matter / grey matter contrast.

2.3 *Processing steps in medical computer vision*

The term computer vision applies to the set of computer processes applied to an image⁴ in order to extract high-level information. These processes can be classified, based on the level of understanding they provide, as follows (see Figure 9):

- **Image enhancement.** Correction of image artefacts or noise, or transformation of the image in order to enhance certain elements.

³ The terms T_1 - and T_2 -weighted refer to images whose intensities are related to the T_1 and T_2 times, respectively, of the tissues.

⁴ The term *image*, in the present context, refers both to 2D and 3D data.

- **Feature detection and segmentation.** This is the first abstraction step in the image processing pipeline. Structural elements of interest are detected (*feature detection*), and objects present in the image are separated from the background and located (*segmentation*).
- **Shape modelling.** The segmented objects are represented in a convenient form. This step is closely related to segmentation, and frequently, segmentation methods take advantage of a particular shape model, as is the case for *deformable models* (cf. section 2.5.2).
- **Registration.** This step can be applied either at the image level (after image enhancement), or after obtaining a convenient shape representation. It brings into correspondence two images (or shape models) in order for them to be compared.
- **Shape description.** High-level information about the objects detected in the image is obtained. This information may be qualitative or quantitative. Clearly, medical morphometric studies should provide accurate quantitative information.

The following sections provide a review of the medical computer vision literature related to each of the steps described above.




Figure 9 Illustration of the processing steps performed in medical computer vision. After removing noise from the image (image enhancement), the object present in the image is identified and extracted (segmentation). An appropriate shape model is obtained, in this case a circle of centre c and radius r . Shape descriptors such as diameter and area can then be obtained. The process of image registration is also illustrated. Image B is registered to image A in order to identify the location in image A corresponding to the cross in image B.

2.4 *Image enhancement: application to MRI*

In an ideal world, based on the physics of MRI, voxels containing a particular tissue type should have nearly identical intensities, thus yielding piecewise constant images. Unfortunately, this is not always the case, and a set of image artefacts is present in MRI acquisitions, affecting both intensity values and the geometry of the structure being imaged. In this section we describe two of the most common problems: bias fields and partial volume effects, and sketch various methods reported in the literature to correct for them. Further discussion of these and other MRI artefacts, such as geometric distortions, ghosting, or motion-induced noise can be found in texts dedicated to the physics of medical imaging (Webb, 1988; Rinck, 1993).

2.4.1 *Inhomogeneity correction*

Spatial inhomogeneities, also known as “bias fields”, affect MR images by altering the intensity measured from homogeneous tissue across the image with a multiplicative, smoothly varying spatial contribution. The result is a correlation of the intensity present at a certain voxel and its spatial location, showing brighter intensities in certain locations of the image, with smooth transitions to darker regions. The effect of a bias field is highly disruptive when using surface coils, affecting intensities with a contribution of up to 60% of the maximum image intensity (Hayton, 1998). Acquisitions using body coils, such as the ones commonly performed for brain studies, create more uniform magnetic fields, thus reducing intensity variations to the range of 10%-20% of image amplitudes (Sled et al., 1998). Although in some cases this spatial variation can be barely noticeable by visual inspection, its effect on intensity-based segmentation techniques is significant, and this has motivated the effort of many research groups towards correcting for it. (Wicks et al., 1993) identify four sources of bias fields, namely: inhomogeneities of the \mathbf{B}_1 field emitted by the transmitting RF coil, non-uniformity of the received \mathbf{B}_1 field due to the sensitivity of the receiving RF coil, the effect of the receiver filter, which is an analogue filter in the receiver chain used to limit bandwidth and hence prevent aliasing in the frequency encode direction, and uncompensated gradient eddy currents.

Bias fields affect the image acquisition process as a loss of high frequencies, which correspond to locations furthest away from the centre of the K-space representation of the image (cf. section 2.2.2). Thus, the effect of the bias field can be modelled as a low-pass filter in K-space:

$$K_{real} = K_{ideal} * L,$$

where L is a low-pass filter and $*$ is the convolution operator. After inverse Fourier transform to generate the intensity image, this convolution becomes a multiplication:

$$I_{real} = F^{-1}(K_{real}) = F^{-1}(K_{ideal} * L) = F^{-1}(K_{ideal}) \cdot F^{-1}(L) = I_{ideal} \cdot B,$$

where B is the bias field observed in the image.

Most approaches to correct for bias fields rely on the assumption of spatial smoothness. (Listerud et al., 1989) claim that variations within the bias field are negligible over a distance of a fourth of the field of view (FOV), thus confining the spatial frequencies of the bias field to a circle of diameter $4/\text{FOV}$ in the K-space representation of the image. Filtering approaches isolate this area of low frequency and identify it to the bias field. Instead of working in frequency space, most methods apply low-pass filters to its intensity counterpart, then divide this blurred version into the original image (Axel et al., 1987). (Listerud et al., 1989) make use of an averaging filter, which is equivalent to applying Gaussian smoothing to the image, to estimate the bias field. (Axel et al., 1987) note that this sort of non-linear processing produced by dividing the image by a blurred version of itself is equivalent to homomorphic filtering. It is commonly recognised that this type of approach introduces edge effects, and only gives good results when applied to modalities with low contrast, such as P_D (Dawant et al., 1993; Wicks et al., 1993). Additionally, removing all low frequencies in the image also implies removing useful structural information. Other early techniques for inhomogeneity correction isolate the bias field by scanning a uniform phantom and building a model of it, which can then be removed from the patient data. (Axel et al., 1987) use the image of a uniform phantom and divide the patient data by it. (Wicks et al., 1993) use an exponential decay function model for the bias, while (Tincher et al., 1993) construct a polynomial representation of its variation across the image. A crucial assumption in these approaches is that of independence of the bias field with respect to placement of the patient in the coil, and negligible temporal and patient-dependent variations. These assumptions do not hold in most cases, and only acquiring a phantom for every patient would circumvent them (Dawant et al.,

1993; Gilles et al., 1996). This does not mean that phantoms are not useful. (Hayton, 1998) uses a phantom to remove a large proportion of the bias field (up to 90%) generated using surface coils in breast MR, and then uses complementary techniques to remove the remaining, patient-dependent contribution.

In a similar vein to (Tincher et al., 1993), several other techniques model the bias field as a polynomial surface fitted to the global intensity variations of the image, but do not use a phantom to obtain this surface. Instead, they fit the polynomial directly to homogeneous regions of the patient data. Examples can be found using thin-plate splines (Dawant et al., 1993), B-splines (Gilles et al., 1996), and Legendre polynomials (Brechtbuhler et al., 1996). (Hayton, 1998) uses the graduated non-convexity algorithm (Blake and Zisserman, 1987) to fit a weak membrane, allowing discontinuities at the transitions between tissues. These spline-based methods have proven successful for correcting intrascan inhomogeneities, but frequently require a slice-by-slice processing in order to correct interscan intensity inhomogeneities (Wells et al., 1996).

(Wells et al., 1996) propose a completely different approach to estimate of the bias field, also allowing for correction of interscan artefacts. Their technique is based on a statistical model of the imaging process, the construction of a Bayesian estimator for the bias field, and the use of the EM algorithm for iteratively interleaving the estimation of this bias field and the segmentation of the image into tissue classes. A further improvement to the method, based on the introduction of a uniform distribution to cater for tissues not explicitly modelled, such as air, bone and fat, or tissues that have a high variance in intensity, is described in (Guillemaud and Brady, 1997). This method is used in our work, and will be described in detail in Chapter 3. Several authors comment on the need for a piecewise constant structure of the tissues in the image in order for the method to give satisfactory results. This is a problem for breast MR studies (Gilles et al., 1996; Hayton, 1998), but the piecewise constant assumption is reasonable for brain tissues. Likewise, (Brechtbuhler et al., 1996; Hayton, 1998) claim that the method is quite sensitive to the parameters of the tissue classes, which must be specified by the user. (Guillemaud and Brady, 1997) provide some insight into the problem of automatically estimating the parameters for the tissue classes,

and (Zhang et al., 1999) describe a method for iteratively updating the estimates of such parameters within the EM framework.

Several methods to bring prior knowledge into the bias correction procedure have been devised. (Aylward et al., 1994a; Aylward et al., 1994b) use an augmented feature space including spatial location to characterise tissue properties. Their approach is quite interesting, and the results reported are excellent, but it requires a considerable amount of user interaction to obtain results of such quality. Several researchers (Rajapakse et al., 1997; Kapur et al., 1998; Zhang et al., 1999) model the expected piecewise nature of the segmentation using Markov Random Fields (MRF) to specify priors on neighbouring pixels. (Kapur et al., 1998) also builds a prior of relative geometry between tissues to model relations between anatomical structures. (Van Leemput et al., 1998) use a digital brain atlas with prior probabilities on tissue location in a standardised Talairach space. While all these approaches go in the right direction towards specifying prior information, it remains a difficult task to leverage the contribution of the prior to the final result. An over-weighting of the prior with respect to the data would tend to suppress detail and create artificial results (particularly affecting partial volume effects; cf. next section), while weighting the prior too low would effectively not contribute to the segmentation and be a computational burden. Finally, in the different field of aneurysm detection from angiography (MRA) data, (Wilson and Noble, to appear) circumvent the problems associated with bias fields by estimating different tissue parameters for different regions of the image in an oct-tree fashion.

2.4.2 *Compensation for partial volume effects*

MRI is an inherently volumetric imaging modality, in the sense that voxels are not infinitesimally small samples of 3D space, but whole 3D regions. Our brief review of the concepts of MR image formation (section 2.2.2) shows that the intensity corresponding to an individual voxel is a function of the signal returned by the whole contents of the 3D region isolated by the gradient fields in order to define the voxel⁵. This means that if several types of tissue are present in the voxel, its intensity will be an average of the contribution of all the tissues. This phenomenon is

⁵ Typical voxel sizes in present acquisition protocols are around 1mm^3 .

known as the *partial volume effect* (PVE). In what follows, voxels containing only one tissue type are referred to as *pure* voxels, whereas voxels containing more than one tissue will be called PVE voxels.

PVE especially affects segmentation methods based on assigning tissue type tags to voxels (see section 2.5.1). Some of the methods described next consider PVE as a corrupting factor and try to *correct* for it. On the other hand, the work described in this thesis is based on using PVE as an important image feature, and the focus is on *estimating* it in order to obtain sub-voxel accuracy.

(Roll et al., 1994) use an intensity thresholding segmentation approach to assign a tissue type tag to each voxel and establish an optimal segmentation threshold and a correction factor to take into account PVE. The particular application is MS lesion quantification. Although complicated by the fact that they take into account all voxels in the image, instead of only PVE voxels, their approach is very simple, and concludes that the optimal threshold is halfway between lesion and background intensities.

Several authors have proposed approaches for estimating PVE by solving a linear system using the information provided by several image modalities. (Soltanian Zadeh et al., 1993) develop a method for creating images with intensities proportional to the quantity of a certain tissue in each voxel, while maximising signal-to-noise ratio. The method is solid and a numerically stable solving procedure is detailed. However, it requires having at least the same number of (perfectly registered) image modalities as the number of tissues that we are interested in segmenting. (Thacker et al., 1998) set up a linear system with the intensities from two image modalities in order to segment grey matter, white matter, and CSF from neurological data sets, and provide a method for estimating the expected accuracy of the results. (Choi et al., 1991) argue that methods that combine different modalities linearly are prone to be very affected by noise, and they propose a Markov random field (MRF) prior as a way to smooth out the results. This introduces the typical problems of MRFs, i.e. determining the correct parameters and the weighting between the prior and the data, as well as huge computational cost. It must be noted that several imaging modalities

are not always readily available, and in such cases a single-channel PVE estimation method is necessary.

Another approach is to model the statistical distribution of pure tissue and mixture voxels and to fit these distributions to the image. (Santago and Gage, 1993) use a Gaussian model for the tissue classes based on a mean intensity and a common noise variance for all the tissues, and employ a uniform distribution to describe the behaviour of the proportion of each tissue in PVE voxels. They describe two methods for fitting these classes to the data, namely Bayes quantification and parameter quantification, and develop error measures for both techniques. The results are rather disappointing, however, since fitting only Gaussian tissue classes gives better results than explicitly modelling a distribution for PVE voxels. (Laidlaw et al., 1998) use the same tissue model assuming a single common variance for all tissues, and provide a sophisticated means for fitting the distributions to the histograms of the whole image and of single voxels, while constraining neighbour continuity. No quantitative comparison with other methods is provided, and visual comparison is performed with methods that do not give appropriate results for the example object. Therefore, it is difficult to draw conclusions about the performance of the method. Unfortunately, this technique is extremely time-consuming, even though it was implemented on parallel hardware.

(Marais, 1999) proposes a different PVE estimation strategy by using a model of the width of each tissue in profiles traversing the brain surface. His work was motivated by the need to deal with sparse, low resolution data sets, and the assumptions taken to model anatomical structure are too simplistic for its use on high resolution MRI.

2.5 Segmentation

Segmentation, even when restricted to the particular domain of medical three-dimensional images, is a very broad subject, which has received a great deal of attention in many applications. An exhaustive review of segmentation methods is outside the scope of this thesis, and the discussion below is limited to the main segmentation paradigms and gives some examples of related work. Segmentation and shape representation are closely related problems, so the reader will find a

significant overlap in the references provided in this and the following sections. The classification follows the taxonomy of (Marais, 1999) and (Jacob, 1999), based on the distinction between statistical voxel classification methods and techniques that incorporate some notion of expected shape.

2.5.1 *Voxel classification methods*

Voxel classification methods are based on establishing statistical intensity models for the different objects in the scene. Traditionally, thresholding or region growing techniques (Gonzalez and Woods, 1992) are applied in order to segment the data, but more complex models require more subtle classification schemes. Examples of this approach in the medical imaging literature can be found in (Dawant et al., 1993; Tincher et al., 1993; Maes et al., 1995; Wells et al., 1996; Guillemaud and Brady, 1997). Classification is conceived of as a local voxel-based process, and independence between neighbouring voxels is often assumed. Notable exceptions to this are models based on Markov random fields (Rajapakse et al., 1997; Kapur et al., 1998; Zhang et al., 1999), which include the notion of local structure by modelling neighbour interactions. Other attempts to bring a notion of shape into these approaches have been based on building probabilistic atlases for the location of anatomical structures, and therefore tissues (Kamber et al., 1992; Van Leemput et al., 1998). Although these last techniques aim at constraining global shape, their effect is still restricted to voxel level, and hence they are unable to enforce large-scale shape constraints.

2.5.2 *Shape-based methods*

The robustness of the segmentation process can often be improved by adding *a priori* topological and/or geometrical knowledge. *Deformable models* employ a shape representation of the surface (contour, in 2D) incorporating constraints about the expected topology and/or geometry, and expose it to a dynamical system which attempts to reach equilibrium between internal (shape constraints) and external (data-driven) forces. The following discussion will focus on the segmentation process itself, while shape representation techniques will be listed in the following section.

The seminal work of (Kass et al., 1987) describes a framework for active contours (snakes) based on 2D splines. They define a potential field combining information from the whole image. The internal force is provided by the spline bending energy, while the external force includes intensity and gradient information to drive the spline towards edges in the image. (Scott, 1987) proposes a very interesting contemporary approach using Fourier snakes and penalising high-order harmonics to ensure smoothness in the final result. Variations to the external force also abound, such as the one introduced by (Ivins and Porrill, 1994), which defines a region energy based on texture. All these techniques rely on global optimisation over the whole image, which is not time-effective. (Blake et al., 1993) modify the snake framework using fast local searches perpendicular to the contour. A parameter controls the search scale, and a statistical framework provides robustness to outliers.

Examples of the use of deformable models in medical imaging abound (Terzopoulos et al., 1987; Lipson et al., 1989; Delingette, 1994; Stoddart et al., 1994; Bardinet et al., 1995). A more complete review of such methods can be found in (McInerney and Terzopoulos, 1996). Methods differ mostly in the shape model used to represent the segmented structure, rather than the segmentation framework itself or the nature of the forces used to drive the fitting process.

Internal forces can be used to provide geometrical continuity and topological consistency, but the range of constraints that can be modelled by means of these forces is still limited. Using them, it is difficult to represent the concept of *normal shape* and *normal variation*. (Cootes et al., 1994) propose computing the mean shape from landmark positions on a set of manually segmented data sets, and using principal component analysis (PCA) on the covariance matrix in order to build a model for the normal shape variation of the structure at hand. Shape can then be constrained using a combination of the mean shape and the largest eigenvectors of this matrix (the principal *modes of variation*). This framework was originally proposed for landmarks in 2D images, but progress has recently been achieved to incorporate it into the internal force of deformable models (Szekely et al., 1996; Montagnat and Delingette, 1997; Jacob, 1999; Marais, 1999).

2.5.3 *Mixed methods*

Voxel classification methods may provide sophisticated statistical models of image intensities, but lack globality and tend to produce noisy or incomplete segmentations. On the other hand, shape-based methods efficiently constrain global shape, but only a very limited amount of the target data, namely the one located near the boundary of the object, is used. These two approaches have classically been considered alternative, and very few efforts have been done to integrate the local (voxel-based) and global (shape-based) information they provide.

As mentioned above, some voxel-classification schemes exist that incorporate more global information by means of Markov random fields and probabilistic atlases, but the effect is restricted to voxel level or small neighbourhoods. Region-driven deformable models (Ivins and Porrill, 1994; Ronfard, 1994; Zhu and Yuille, 1995; Brady et al., to appear) perform region growing while constraining the outer boundary of the segmentation using a snake-like approach that can incorporate shape constraints.

The work described in this thesis combines both statistical and shape information to drive the fitting process of a deformable model. It is also noticeable that very few segmentation methods provide a framework for quantifying the uncertainty of the resulting segmentation. This issue, of vital importance to medical applications, is also addressed in this thesis. A full description is provided in Chapter 3.

2.6 *Shape representation*

The complexity of shapes encountered in biological structures, and in particular the ones found in the brain, seriously challenges existing shape modelling techniques. Several attempts have been made to build models to encompass biological shape, and this section provides a review of such techniques. Let us first consider the characteristics that a shape representation for biological structures should have:

- **Topological flexibility.** Organs such as hippocampi and the brain cortex have spherical topology, whereas the ventricles, for example, have holes.
- **Geometrical flexibility,** for representing highly irregular and convoluted shapes. The geometrical complexity of the brain cortex is well known, not only for its convoluted shape,

but also for its intrinsic geometry, which makes it impossible to “flatten” its surface without introducing cuts or local stretches (Griffin, 1994).

- **Structural simplicity.** The shape model of choice is bound to be used as a deformable model, so in order to reduce computational cost, it is important that the model be simple in structure.
- **Dense resolution.** The model has to be able to closely approximate very complex shape, so sparse shape models are not appropriate. Additionally, geometrical continuity allows anatomical smoothness to be naturally modelled.

It should be noted that some of these constraints are somehow contradictory, so we need to establish a trade-off between them. High flexibility (topological or geometrical) and structural simplicity are usually at odds with dense resolution and accuracy.

The classification that is provided next is based on the structural complexity of the shape representation. In particular, the following paradigms are described: voxel-based representations, landmark-based methods, parametric models, representations based on meshes, and other techniques, such as symmetry-based representations and level sets.

2.6.1 Voxel-based representations

The results of voxel-classification segmentation methods are usually stored as a set of voxels in a mask, or labels for each voxel indicating the group it belongs to. If no post-processing is made to generate a more sophisticated representation, this is the final shape model for the segmented object. Most of the methods mentioned in section 2.5.1 fall into this category (Griffin et al., 1994; Maes et al., 1995; Kapur et al., 1996). Extensions to this paradigm may include more information about the tissue each voxel belongs to, such as a set of labels for each voxel (Tiede et al., 1996). Typically, these representations require large amounts of data and handling them is highly complex and time-consuming, thus rendering them unsuitable for applications that require flexibility. On the other hand, the high number of voxels involved in a segmentation ensures dense resolution (limited to the voxel size).

2.6.2 *Landmark-based methods*

At the other end of the spectrum of amount of represented information, landmark-based methods store very little data, corresponding to the locations of a set of (anatomically meaningful or otherwise chosen) landmarks. Landmarks can be points (Cootes et al., 1994; Bookstein, 1996; Amit, 1997), curves (Eberly et al., 1994; Subsol et al., 1995; Lohmann et al., 1997), or even surfaces embedded in 3D volumes (Collins et al., 1996). A more in-depth discussion on the choice of landmarks and a description of a particular way to analyse them to establish shape comparisons can be found in (Bookstein, 1992). These representations provide flexibility and a remarkable structural simplicity, but are extremely limited in the resolution of the representation. In fact, their usefulness is limited to studies where the location of the landmarks is the only important issue. Interpolation techniques can be used (Bookstein, 1996), but the limited number of landmarks makes it impossible to obtain true dense resolution.

2.6.3 *Parametric models*

Parametric surfaces, such as superquadrics (Lipson et al., 1989; Terzopoulos and Metaxas, 1991) provide a compact shape representation that is structurally simple. However, the range of shapes that such analytical models can represent in practice is rather limited. Extensions using free-form deformations (Bardinet et al., 1995), fractal-based parameter spaces (Vemuri and Radisavljevic, 1994), blending schemes to smoothly combine several surfaces (DeCarlo and Metaxas, 1995), and solid textures (Chover, 1996; Gonzalez Ballester, 1996) have been devised in order to enhance the representational power of such schemes. However, these extensions tend to increase enormously the number of parameters of the model, thus negatively affecting the structural load of the parametric model.

Fourier curves (Scott, 1987) and surfaces (Szekely et al., 1996) provide a convenient continuum between scales of detail, by decomposing the surface (curve in 2D) into *harmonics*. Although the formulation is mathematically nice, the number of harmonics required to model surfaces of high complexity, such as the brain cortex, render these methods ineffective.

Parametric models have no topological flexibility. The topology of the object must be known beforehand, since it is built into the analytical definition of the shape. Although none of the methods described above provides a completely satisfactory solution for complex biological shapes, it should be noted that having a parameter space to index the shape is very desirable, since it has several advantages for further shape description, as will be discussed in section 2.8.

2.6.4 Meshes

Discrete meshes are composed of a set of nodes (typically 3D points) and a connectivity function, which determines the connections between the nodes. Because of their regularity and the possibility to tessellate spherical topologies, triangulations have been used extensively as shape models. The *marching cubes* algorithm (Lorensen and Cline, 1987; Tiede et al., 1990), creates a triangulation from raw three-dimensional data by computing the location of an isosurface at a given intensity value. If the location of the nodes is known, a Delaunay triangulation (Fang and Piegl, 1995) can be built, in which the generated triangles are related to the specified nodes in that the closest node to all points in a given triangle is the same one. Triangulations generated directly from raw data tend to have an enormous number of nodes, making them very costly to manipulate. However, they remain effective for rendering purposes, since dedicated hardware is generally available.

Meshes can also be used as deformable models. In this case, the connections between nodes act as springs that enforce shape constraints between a node and its neighbours. Such is the case for active surfaces like the *simplex mesh* (Delingette, 1994), *slime* (Stoddart et al., 1994; Stoddart and Baker, 1998), or the ones described in (Nastar, 1994). The main advantage of such methods is their flexibility, both in the geometric and topological sense. A simple change in the connectivity function can alter the topology of the object. They are also structurally simple and easy to manipulate. Additionally, most of these techniques include some simplification and refinement processes, by which nodes can be dynamically removed or added, respectively, depending on the desired level of flexibility and/or dense resolution. Such is the case of the simplex mesh (Delingette, 1994), which also incorporates a very sophisticated segmentation and topological manipulation framework. Furthermore, simplex meshes are especially easy to manipulate, since

they are regular in their structure, which is dual to the one of the triangulation, i.e. each node is connected to exactly three other nodes.

The obvious drawback of discrete meshes is that of lack of density in their resolution. Interpolant patches can be used to this end. (Stoddart et al., 1994) uses S-patches (Loop, 1992) to create a G^1 -continuous surface interpolating the nodes of the slime mesh. (Schmitt et al., 1991) describe an interpolating procedure based on triangular Gregory-Bézier patches, which ensures G^1 continuity without the need to enforce conditions on the corners of the triangles, thus considerably reducing the computational cost. Alternative interpolants that can be used to create a spatially continuous mesh can be found in (Barnhill, 1974; Hoppe et al., 1994; Krishnamurthy and Levoy, 1996; Hong and Terzopoulos, 1997). The shape model used in this thesis incorporates triangular Gregory-Bézier patches (Schmitt et al., 1991) to the simplex mesh framework (Delingette, 1994), thus providing G^1 -continuity to the mesh. A more in-depth discussion of the structure and properties of this shape model is given in Chapter 3.

2.6.5 *Other shape models*

Exploiting the idea of *symmetry* in biological shapes has been one of the principal paradigms in shape modelling. The seminal works of (Blum 1973) and (Brady and Asada, 1984) are examples of early use of this idea in computer vision. Symmetry-based shape models are usually based on computing a symmetry axis (surface, in 3D) and a function that specifies the distance to the contour (or outer surface) for every point in the symmetry axis. (Terzopoulos et al., 1987) introduced symmetry-seeking deformable models. Symmetry-based representations applied to medical data can be found in (Gauch and Pizer, 1993; Burbeck and Pizer, 1995; Naf et al., 1996). It is worth noting that the extension of 2D symmetry axes to 3D symmetry surfaces is not trivial, and the resulting model may be very complex in structure (effectively another surface, or set of surfaces, representing the original surface), thus reducing the usefulness of these approaches for representing 3D structures.

Another shape representation paradigm is that of *level sets* (Whitaker, 1995; Xiaolan et al., 1998). In this case, the surface is modelled implicitly as the solution to a set of partial differential equations,

and no explicit mathematical representation is given. This means that another representation technique (voxel-based, usually) must be used to manipulate the resulting model. These methods provide a nice framework for segmentation, but lack structural simplicity and usually require highly time-consuming segmentation algorithms.

2.7 Registration

Although this thesis is not directly related to registration issues, it is convenient to sketch the main paradigms and introduce some concepts that may appear in what follows. An in-depth review of present (from 1993) registration techniques can be found in (Maintz and Viergever, 1998). Techniques prior to 1993 are reviewed in (Maurer and Fitzpatrick, 1993; van den Elsen et al., 1993).

Several taxonomies are reported in (Maintz and Viergever, 1998). The one most frequently used deals with the nature of the transformation used to map the two images (or shape models). *Rigid* registration performs a 6 degree of freedom (translation and rotation) transformation, while *affine* schemes also allow for scaling and shearing effects. Local techniques based on vector field mappings, commonly known as *non-rigid* registration, have the potential to model more complex deformations.

2.8 Shape description

Classical quantitative shape descriptors, such as length, area and volume can be obtained from most of the shape representations described in section 2.6. In particular, measures provided by voxel-based models are computed by simple addition. Volume, for example, is computed by adding up all the volumes of the voxels belonging to the object. It must be noted that a considerable error is introduced due to partial volume effects (cf. section 2.4.2), especially when dealing with complex convoluted biological structures with large surface area. Parametric models are specially suited to obtaining accurate quantitative measures, since they are defined by continuous analytical functions. Area and volume can be computed using concepts of differential geometry (Boehm and Prautzsch, 1994) and results from vector calculus, such as Gauss'

divergence theorem (Boas, 1983). These methods can also be applied to parametric interpolants fitted onto meshes (cf. Chapter 3 for an example of this). For planar triangulations, simpler expressions for volume and area can be obtained (see section 3.8).

Landmark-based methods do not provide full coverage of the shape of the object, but rather concentrate on specific meaningful locations. Therefore, global measures such as volume or area cannot be computed. Interesting measures focus on the variation of the location of these landmarks, either relative to each other, between different time points of a same patient, or in inter-patient studies. A detailed discussion of these techniques can be found in (Bookstein, 1992). Point Distribution Models (PDM), introduced by (Cootes et al., 1994), also provide a natural framework for studying intra- and inter-patient variation. Their landmark-based segmentation scheme, based on principal component analysis to constrain the allowed variation of shape, results in a set of coefficients for the eigenvectors corresponding to the principal modes of variation. Studying these coefficients may allow for the detection of abnormal variation with respect to the training set. This method has been applied to neurological and cardiac studies (Cootes et al., 1994), as well as orthopaedics (Solloway et al., 1996).

Differential characteristics of the contour or surface of the object of interest can be used to describe its shape based on concepts of differential geometry (mainly curvature). The most well known examples of use of this type of techniques are based on finding *ridges* or *crest-lines*. These curves are the loci of points of maximum curvature. They can be found using standard differential geometry on a parametric surface (Boehm and Prautzsch, 1994) or approximating differential properties directly from image data (Monga and Benayoun, 1995; Thirion and Gourdon, 1995). Crest-lines can be used as landmarks and compared to an atlas of normal variation (Subsol et al., 1995).

The analysis of vector fields resulting from non-rigid registration schemes is another approach for the study of shape differences. These vector fields establish a voxel-wise mapping between two images. A filtering process and a subsequent statistical analysis may then be applied in order to detect areas of significant differences. Examples of such techniques can be found applied to

multiple sclerosis time-evolution studies (Thirion and Calmon, 1997) and the assessment of anatomical evidence of schizophrenia (Prima et al., 1998). A similar approach is taken in (Thompson and Toga, 1997) to study anatomical variation on parameterised brain models.

2.9 *Summary*

Background information on the basics of MRI and medical computer vision has been provided in this chapter. The typical processing steps intervening in medical computer vision have been presented and literature reviews given for each of them.

A Framework for Morphometric Studies of MRI Including Confidence Bounds

3.1 Introduction

This chapter describes an algorithmic framework for the morphometric analysis of MRI scans. Several issues are addressed, corresponding to different stages of medical computer vision described in 2.3, namely: image enhancement, feature detection, segmentation, shape modelling, and shape description.

Image enhancement is performed to correct for magnetic field inhomogeneities, also known as bias fields (section 2.4.1). This is achieved using the EM algorithm to interleave the estimation of the bias field and the assignment of tissue tags to voxels in order to perform voxel-based segmentation (Wells et al., 1996; Guillemaud and Brady, 1997). The method is described in detail in section 3.2. Feature detection is accomplished by identifying voxels located at the boundaries of tissues in the brain by detecting PVE voxels. This method is based on the results of the EM algorithm and is outlined in section 3.3.

A mixed segmentation procedure combining the characteristics of voxel-based and shape-based methods is used. The two main elements of the segmentation framework are the EM algorithm and the simplex mesh (Delingette, 1994). Simplex meshes are active shape models, and will be described in section 3.4. Next, section 3.5 shows how both methods are combined to provide more robust segmentations.

The shape model used in this work satisfies most of the requirements stated in 2.6. In particular, it is given by a mesh of triangular Gregory-Bézier patches (Schmitt et al., 1991) providing G^1 -continuity. A detailed description is offered in section 3.6.

Confidence bounds are a key feature of this framework. It has been noted in the previous chapters that morphometric tools currently fail to provide an estimate of the error incurred during the measurement process. The principal obstacle to obtaining high accuracy on measurements performed in MRI is the resolution of the images and, as a direct consequence, the partial volume effect. In the present framework, PVE voxels are considered “uncertain” and a double segmentation is produced to generate *inner* and *outer* bounds on the location of the surface of the object being segmented. This is further considered in section 3.7.

Section 3.8 provides formulae for the computation of volume measurements from the simplex mesh and the tGB mesh. This general shape descriptor offers good means of comparison with the results of other methods. However, more sophisticated shape descriptors should be developed to study issues such as brain asymmetry.

Validation issues, in particular the need for ground truth to compare the results with are discussed in section 3.9. Extensive validation results and examples of the use of the method to clinical data are shown in section 3.10. Section 3.11 reports results on the application of the method to the segmentation of the cortex, proving that the shape model is flexible and useful for modelling complex biological shapes. Finally, section 3.12 discusses these results and concludes the chapter.

The contents of this chapter are based on the previously published articles “Combined statistical and geometrical 3D segmentation and measurement of brain structures” (Gonzalez Ballester et al., 1998a), “Measurement of brain structures based on statistical and geometrical 3D segmentation” (Gonzalez Ballester et al., 1998b), and “Segmentation and measurement of brain structures in MRI including confidence bounds” (Gonzalez Ballester et al., to appear).

3.2 Bias correction and voxel-based segmentation: the EM algorithm

3.2.1 Basic formulation

The first of the elements in our framework is based on the work of (Wells et al., 1996) and (Guillemaud and Brady, 1997) on bias correction. These methods estimate and correct for the bias field and produce a set of probability maps that will prove useful for the segmentation process.

The *expectation-maximisation* (EM) algorithm interleaves the estimation of two coupled distributions (Dempster, 1977). This is used in (Wells et al., 1996) to estimate simultaneously the (multiplicative) bias field and the probabilities of each voxel belonging to each tissue class. As a first step, the image is log-transformed so the multiplicative field becomes an addition. Let $Y = \{Y_i\}_{i=1..n}$ be the log-transformed set of voxels in the MR volume⁶. Each voxel has an additive contribution β_i from the bias field. Let $\beta = \{\beta_i\}_{i=1..n}$ denote the contributions for all voxels in the image. Following the model presented in section 2.4.1 (and ignoring additive noise), the effect of this logarithmic transform is as follows:

$$I_{real} = I_{ideal} \cdot B \Rightarrow \ln I_{real} = \ln I_{ideal} + \ln B = Y + \beta$$

The range of image intensities corresponding to a certain tissue type is modelled as a Gaussian distribution with small variance around a mean intensity value, and this model is extrapolated to the log-transformed intensities (further consideration of this assumption is given in the next section). Thus, the distribution on log-intensities of voxels containing (only) tissue j is assumed to be Gaussian with mean log-intensity μ_j and standard deviation ϕ_j . The probability that a voxel i , which contains only tissue j (represented by $\Gamma_j = 1$), and which is affected by a bias field contribution β_i , has log-intensity value Y_i is then:

$$p(Y_i | \Gamma_j, \beta_i) = G_{\phi_j}(Y_i - \mu_j - \beta_i)$$

where:

$$G_{\phi_j}(x) = \left(2\pi\phi_j^2\right)^{-\frac{1}{2}} \exp\left(-\frac{x^2}{2\phi_j^2}\right)$$

⁶ For simplicity, only one index i is used to locate a voxel in the volume, instead of the three indexes (x, y, z) that are commonly used.

The formulation also allows for the extension to multiple modalities, where a number m of perfectly registered images are available for the same data (for example, $m = 2$ for double-echo MRI acquisitions). In this case, the log-intensity model is assumed to be an m -dimensional Gaussian distribution:

$$G_{\phi_j}(\mathbf{x}) = (2\pi)^{-\frac{m}{2}} |\phi_j|^{-\frac{1}{2}} \exp\left(-\frac{1}{2} \mathbf{x}^T \phi_j^{-1} \mathbf{x}\right)$$

Likewise, in order to model the smoothness of the bias field, an n -dimensional Gaussian prior probability density with zero mean is used:

$$p(\beta) = G_{\phi_\beta}(\beta)$$

where:

$$G_{\phi_\beta}(\mathbf{x}) = (2\pi)^{-\frac{n}{2}} |\phi_\beta|^{-\frac{1}{2}} \exp\left(-\frac{1}{2} \mathbf{x}^T \phi_\beta^{-1} \mathbf{x}\right)$$

Statistical independence of the voxel intensities is assumed, so the expression for the probability of an image given the bias field is:

$$p(Y|\beta) = \prod_i p(Y_i|\beta_i) = \prod_i \left[\sum_j p(Y_i|\Gamma_j, \beta_i) p(\Gamma_j) \right]$$

Using Bayes' rule, an expression for the posterior probability of the bias field given the observed intensity data is obtained:

$$p(\beta|Y) = p(Y|\beta) \frac{p(\beta)}{p(Y)}$$

where $p(Y)$ is considered an unimportant normalising constant.

A zero-gradient condition on the logarithm of this posterior probability is used to estimate its maximum. Two coupled formulae are obtained, one estimating the probability of a voxel i belonging to tissue class j ($W_{i,j}$); the other computing the estimated bias field β . Both expressions are inter-dependent, and the EM algorithm is used to interleave their estimation iteratively. In particular, $W_{i,j}$ is computed as:

$$W_{i,j} = \frac{p(Y_i|\Gamma_j, \beta_i) p(\Gamma_j)}{\sum_j p(Y_i|\Gamma_j, \beta_i) p(\Gamma_j)} \quad (3.1)$$

β is expressed as a function of the mean residual R and mean inverse covariance φ^{-1} over all tissue classes:

$$R_i = \sum_j W_{i,j} \frac{Y_i - \mu_j}{\varphi_j}$$

$$\varphi_{i,k}^{-1} = \begin{cases} \sum_j W_{i,j} \varphi_j^{-1} & \text{if } i = k \\ 0 & \text{otherwise} \end{cases}$$

An additional operator H is defined as a function of the mean inverse covariance and the variance of the bias field:

$$H = [\varphi^{-1} + \varphi_\beta^{-1}]^{-1}$$

This operator is usually implemented as a low-pass filter. Finally, the bias field is estimated as:

$$\beta = HR \quad (3.2)$$

The coupled expressions (3.1) and (3.2) are alternatively evaluated, as dictated by the EM algorithm. Convergence is guaranteed and a good estimate is usually obtained after 4 or 5 iterations. Figure 10 shows an example of the application of this technique and the improvement achieved in the classification by removing the bias field.




Figure 10 Bias correction process: A) original image; B) image corrected for bias field; C) and D) segmentations corresponding to A) and B), respectively, obtained by assigning to each voxel the tissue with largest probability. Notice the improved performance of the segmentation after bias correction.

3.2.2 Improvements and considerations

The method developed by (Wells et al., 1996) requires that every voxel in the image be assigned to one of the Gaussian distributions corresponding to each tissue class. (Guillemaud and Brady, 1997) argue that in most images there are a number of voxels that do not belong to the modelled

classes. Modelling the set of such voxels as a Gaussian gives very large value for the standard deviation, and this leads to numerical instability. An additional class *other* with uniform probability distribution is introduced to model such voxels.

The algorithm requires the user to provide the number of tissues to be explicitly modelled, as well as their parameters μ_j and ϕ_j . We estimate these values interactively, by allowing the user to select a set of points in the tissue of interest and determining the statistics of the sample. This method has proven reasonable for our purposes, but more refined techniques, like that described in (Zhang et al., 1999), which introduces an update of the tissue parameters into the E-M framework could be used.

A note of warning should be made about the Gaussian model for log-intensities of the image in (Wells et al., 1996). The model most commonly used in MRI literature is based on Gaussian distributions to model the intensities corresponding to a certain tissue type. The same class of model is assumed to fit log-intensities, for practical reasons⁷. This simplification should have no effect on the largest part of the intensity range of the image. However, intensities close to zero (such as CSF in T₁-weighted images) suffer a distortion in range when log-transformed, so the Gaussian model for log-intensities may not provide the same results as a Gaussian on intensities.

3.3 *Detection of PVE voxels*

We now take on the task of defining an appropriate feature to detect the boundary of structures of interest. A careful consideration of the elements present in the imaging process leads to the conclusion that partial volume effects (PVE) are the best choice for a feature for segmentation. In fact, we are interested in detecting boundaries between tissues, but these are not accurately localised in MR images due to the discretisation of the images into voxels. PVE voxels contain more than one tissue, thus indicating that boundaries between tissues are present in them.

⁷ Personal communication with Sandy Wells.

In our framework, PVE voxels are detected by analysing the output of the EM-based statistical method, which consists of: an estimation of the bias field corrupting the MR (3D) image, a corrected version of the image after removing the bias field, and a set of probability maps for each tissue class. The usual approach is to obtain a voxel-based segmentation of the image into tissue classes by assigning to each voxel the tissue class with maximum probability (see Figure 10):

$$t_i = \arg \max_j W_{i,j}$$

We contend that a more careful treatment of the probability maps can provide additional useful information. In particular, for certain voxels the value of the maximum probability t_i is not very large, indicating that the voxel does not accord well with any of the tissue classes (including the uniform class *other*). This is particularly the case for PVE voxels. That is, a low value in the maximum of the probability maps is often a good indicator of boundariness. Thus, we detect PVE voxels by setting a threshold *thrs* on the minimum value allowed for this maximum probability:

$$T_i = \begin{cases} t_i & \text{if } W_{i,j} > \textit{thrs} \\ PVE & \text{otherwise} \end{cases}$$

Typically, a value of *thrs* in the range of 90-95% probability is used. It should be noted, however, that the particular choice of threshold depends on the shape of probability distributions for the tissues. In any case, the method is not very sensitive to small variations on the value of the threshold. This means that, in practice, a visual inspection of the results suffices to assess the adequacy of the threshold. Figure 11 shows an example of the use of the method.

3.4 *Shape-based segmentation: the simplex mesh*

Our segmentation framework is based on combining the EM voxel-based segmentation method described in section 3.2 with a shape-based model. This section describes the characteristics of the shape-based deformable model, and the following section shows how both methods are combined to provide more robust segmentations.




Figure 11 A) Original slice (coronal orientation); B) corrected slice; C) estimated bias field; D) segmentation using a tissue model for white matter and grey matter, plus a class with a uniform probability distribution for CSF, air, and other tissues. Voxels deemed to be places where the PVE is significant are coloured white ($thrs=0.95$).

Deformable models were introduced in section 2.5.2. In our framework we employ *simplex meshes* (Delingette, 1994). The fundamental property of a simplex mesh is that all its nodes have the same number of connections to other nodes. Formally, a k -simplex mesh M of \mathfrak{R}^3 is defined as a pair $\{V(M), N(M)\}$, where $V(M)$ is a set of vertices and $N(M)$ is a connectivity function between these vertices. A k -simplex mesh has $(k+1)$ -connectivity, i.e. each vertex is connected to exactly $(k+1)$ other vertices. A set of complementary conditions ensures the mathematical correctness of the construction of the mesh, by not allowing loops and guaranteeing the existence of a path connecting any two nodes (Delingette, 1994). 2-simplex meshes are the appropriate choice to model surfaces, and will be used in this work.

The topology of a simplex mesh is defined by its connectivity function $N(M)$. A set of operations is defined in order to provide tools to alter the mesh topology. These operations are based on the

addition and removal of nodes and on changes to the connectivity between nodes. An important property of simplex meshes is that they are dual to triangulations, i.e. a triangulation with the same topology can always be built from a simplex mesh. The concept of a contour as a connected subset of vertices and elements of $N(M)$ is also introduced in order to allow for better topological control.

Topology is independent of the embedding space, whereas geometry is highly related to the dimension of the space in which the mesh is located. The geometry of a simplex mesh is uniquely determined by its set of vertices $V(M)$ and their corresponding *simplex angles* and *metric parameters*. These last two sets of variables define the position of a vertex relative to its three (in the case of 2-simplex meshes) neighbours. Expressions for the mean and Gaussian curvatures are also derived in (Delingette, 1994), and prove useful for enforcing certain shape constraints to the mesh, as mentioned below.

The mesh is initialised and then exposed to a set of forces to make it lock on to the target data. Both intrinsic (shape) and extrinsic (fit to data) forces are introduced in the model, which is subjected to Newtonian dynamics:

$$m \frac{d^2 \mathbf{P}_i}{dt^2} = -\gamma \frac{d\mathbf{P}_i}{dt} + \mathbf{F}_{\text{int}} + \mathbf{F}_{\text{ext}}$$

where m is the mass of a vertex, \mathbf{P}_i is the position of the i^{th} node at time t , and γ is the damping factor. This is discretised assuming $t_i = t_0 + i\Delta t$ and using finite differences:

$$\mathbf{P}_i^{t+1} = (1 - \gamma) \cdot (\mathbf{P}_i^t - \mathbf{P}_i^{t-1}) + \mathbf{F}_{\text{int}} + \mathbf{F}_{\text{ext}}$$

where unit mass of the vertices is assumed and \mathbf{F}_{int} and \mathbf{F}_{ext} are computed at time t .

Internal forces determine the response of a physically-based model to external constraints. The response of elastic bodies may be derived by minimisation of an elastic energy. This elastic behaviour is implemented by \mathbf{F}_{int} . Rather than minimising a global elastic energy combining the positions of all nodes, the internal force is expressed in terms of simplex angles and metric parameters, so shape control is conceived as a local process of a vertex relative to its neighbours. Following this model, different types of constraints can be implemented, such as continuity of the

normal, surface orientation, mean curvature and shape constraint. This last constraint allows for the specification of a template or “rest” shape.

External forces are used to fit the mesh to three-dimensional data (in our case, the brain surface or some organ of interest contained in the MR volume). The method used in the implementation of the simplex mesh is based on distances from each point to the closest data point, and follows the *iterated closest point* approach (Besl and McKay, 1992). In particular, at each iteration and for every vertex \mathbf{P}_i the closest data point $\mathbf{M}_{Cl(i)}$ is searched for, and the expression for the force is:

$$\mathbf{F}_{\text{ext}} = \beta_i G \left(\frac{\|\overrightarrow{\mathbf{P}_i \mathbf{M}_{Cl(i)}}\|}{D} \right) \cdot (\overrightarrow{\mathbf{P}_i \mathbf{M}_{Cl(i)}} \cdot \mathbf{N}_i) \mathbf{N}_i$$

where \mathbf{N}_i is the normal vector at \mathbf{P}_i and $G(x)$ is the stiffness function, which has a constant value of 1 in the interval $x \in [0,1]$ and decreases exponentially from that point on. D determines the range of influence of \mathbf{F}_{ext} , and can be expressed as a number of voxels around \mathbf{P}_i . The term *data point* refers to some feature derived from the image data. In our case, it will be PVE voxels, as described in the following section. See Figure 12 for an example of the fitting process.

Additional topological control is obtained by two extensions to the normal behaviour of the mesh. First, a procedure has been implemented that adapts the mesh by moving vertices towards areas of high curvature, where more information is needed to represent the data. As a complement to this procedure, a method exists whereby the mesh is refined, i.e. the number of vertices is increased, in these areas (Delingette, 1994).

3.5 Global segmentation scheme

Voxel-classification and shape-based segmentation methods have traditionally been regarded as opposites, since voxel classification is a local process, whereas shape-based segmentation deals with global shape. However, some attempts have been made to reconcile these two paradigms (cf. 2.5.3). The method described in this chapter combines characteristics of both statistical voxel classification and shape-driven segmentation.




Figure 12 Simplex mesh fitting process: A) Original mesh (sphere) superimposed on a 2D slice of the MR volume; B) the mesh is locking on to the boundary; C) refined mesh; D) one slice of the target MRI data, showing a cross-section of the object to be segmented at top (cf. section 3.10.2 for details), and E) intersection of the resulting mesh with the slice.

The essence of the method consists of using the probability maps $W_{i,j}$ obtained from the application of the EM algorithm to create a pre-segmentation by classifying voxels into tissue classes, and detecting PVE voxels as explained in section 3.3. The simplex mesh uses PVE as an edge feature, and searches for PVE voxels to position its nodes. In-built shape constraints of the simplex mesh ensure smooth segmentations, effectively removing the noise in the pre-

segmentation obtained from the EM algorithm and providing robustness to outliers. Figure 13 illustrates this procedure. From the results of the simplex mesh program, a surface model is built and confidence bounds on the location of the object are computed. Next, shape descriptors are derived. These elements are described in the following sections.




Figure 13 Global segmentation scheme.

3.6 *Modelling complex biological shape*

The simplex mesh is a powerful tool for three-dimensional segmentation, since it has great topological flexibility; but much of this flexibility derives from the simplicity of the shape model, consisting of a set of connected 3D points. The representation of biological shape, and in particular the enormously complex shape of the brain surface, requires a more sophisticated representation technique. In addition, the applications of interest demand high accuracy in the location of the boundary of an object, and this is at odds with the simplicity of the simplex mesh.

Our approach is to construct a G^1 -continuous surface interpolating the positions and normals of the nodes of the simplex mesh. A triangulation is first derived from the mesh by adding a node at the centroid of each of its polygons and updating its position to the nearest data point by means of a local search, in the same way that nodes are updated during the fitting process⁸. Then, a set of triangular Gregory-Bézier (tGB) patches (Schmitt et al., 1991) are interpolated to the triangulation. TGB patches are expressed, in barycentric coordinates, as follows:

⁸ Alternatively, a triangulation can be built by joining the centroids of all the polygons in the simplex mesh.

$$\begin{aligned} \mathbf{GB}(u, v, w) = & u^3 \mathbf{P}_0 + v^3 \mathbf{P}_1 + w^3 \mathbf{P}_2 + 12u^2 v w \mathbf{P}_{211} + 12u v^2 w \mathbf{P}_{121} + 12u v w^2 \mathbf{P}_{112} + \\ & 3u^2 v(1-w) \mathbf{P}_{01} + 3u v^2(1-w) \mathbf{P}_{02} + 3v^2(1-u) w \mathbf{P}_{11} + \\ & 3(1-u) v w^2 \mathbf{P}_{12} + 3u(1-v) w^2 \mathbf{P}_{21} + 3u^2(1-v) w \mathbf{P}_{22} \end{aligned}$$

where $0 \leq u, v, w \leq 1$, $u + v + w = 1$ and:

$$\mathbf{P}_{211} = \frac{w \mathbf{P}_{211}^v + v \mathbf{P}_{211}^w}{w + v}, \quad \mathbf{P}_{121} = \frac{u \mathbf{P}_{121}^w + w \mathbf{P}_{121}^u}{u + w}, \quad \mathbf{P}_{112} = \frac{v \mathbf{P}_{112}^u + u \mathbf{P}_{112}^v}{v + u}$$




Figure 14 Triangular Gregory-Bézier (tGB) patch, defined by 15 control points. A G^1 -continuous mesh of tGB patches is used to interpolate the nodes of the simplex mesh and their corresponding normals.

Figure 14 shows a tGB patch, defined by its 15 control points. An important property of tGB patches is that the expressions for the first derivatives with respect to u , v , and w do not share any inner control point. In practice, this means that G^1 continuity can be guaranteed simply by constraining the connection along the boundary between two adjacent patches, avoiding the tedious process of considering continuity at corners (cf. Appendix A). G^1 continuity across patches is ensured by enforcing a coplanarity constraint between the two radial first derivative vectors

$$\left. \frac{\partial \mathbf{GB}_R(u_R, v_R, w_R)}{\partial \theta_{u_R}} \right|_{u_R=0} \quad \text{and} \quad \left. \frac{\partial \mathbf{GB}_L(u_L, v_L, w_L)}{\partial \theta_{u_L}} \right|_{u_L=0}$$

(where θ_{u_R} and θ_{u_L} take values in $\{u, v, w\}$, depending on the orientation of the patch), and the first order derivative vector $\Gamma^{(1)}(v)$ of the common boundary (see Figure 15):

$$\left(\frac{\partial GB_R(u_R, v_R, w_R)}{\partial \theta_{u_R}} \Big|_{u_R=0} \times \frac{\partial GB_L(u_L, v_L, w_L)}{\partial \theta_{u_L}} \Big|_{u_L=0} \right) \cdot \Gamma^{(1)}(v) = 0, \quad 0 \leq v \leq 1$$




Figure 15 Enforcing G^1 continuity in the boundary of neighbouring tGB patches.

3.7 Estimation of confidence bounds

Confidence bounds are one of the most important elements of our framework, and to estimate them we rely on the partial volume effect. PVE voxels can typically be seen as a band of variable width delineating the interface between tissues (see Figure 11). We interpret this width, corresponding to the transition from *pure* voxels of tissue A to *pure* voxels of tissue B , as the confidence interval we seek for our segmentation.

The mesh is used to obtain two surfaces representing the *inner* and *outer* estimates on the location of the true surface of the object. The exact location of these surfaces depends on the width of the PVE band. In particular, during each update step of the mesh, local searches in the direction of the normal of each node of the mesh are performed. When these searches hit PVE voxels, they produce upper and lower estimates by computing the intersection of the profile with the closest and farthest PVE voxel boundary. By employing this process on all the vertices of the mesh, the two bounding surfaces are built (see Figure 16). The *real* location of the surface is contained between these two bounding surfaces.




Figure 16 Outer (black) and inner (white) meshes segmenting the lateral ventricle of a patient. A) mesh superimposed on a coronal slice of the data set; B) close-up showing that the inner bound is contained inside the outer bound; C) intersection of both meshes with a slice. The separation between outer and inner bounds is quite small, so it is only reflected by a slight thickening of the white line in C).

3.8 *Shape description: volume computation*

Simple shape descriptors such as volume and area have very general application to anatomical studies. Formulae to compute the volume enclosed by a simplex mesh and a mesh of tGB patches are now given. It is recognised, however, that their descriptive power is very limited, and more

sophisticated shape descriptors need to be built to characterise biological shape. In particular, this is the case for the description of symmetry differences in schizophrenia.

Volume computations on the simplex mesh are obtained by first computing the centroid of each polygon and joining it to the neighbouring nodes, so creating a triangulation $T = \{\mathbf{P}_1^i, \mathbf{P}_2^i, \mathbf{P}_3^i\}_{i=1..n}$. Then, the volume is determined by the following expression, which is a simplification of Gauss' divergence theorem (Boas, 1983) when applied to triangulated surfaces:

$$V = \frac{1}{6} \sum_{i=1}^n \langle \mathbf{P}_1^i, \mathbf{P}_2^i, \mathbf{P}_3^i \rangle$$

where $\mathbf{P}_j^i \in \mathfrak{R}^3$ and $\langle \rangle$ denotes the scalar triple product operator. It is worth noting that computing the volume of a simplex mesh based on a triangulation of the surface poses a problem when it comes to establishing bounds on the surface, since undulations of the surface can “overflow” the triangulation. We circumvent this problem by displacing the position of the centre of each triangle to the nearest voxel boundary and updating the position of the nodes of the mesh accordingly. Figure 17 shows a 2D representation of this process. Grey circles represent nodes of the triangulation, and black circles are the data points found for the outer and inner bounds, which correspond to boundaries between voxels of different tissue or farthest boundaries of PVE voxels. A one pixel wide PVE band is shown in the 2D scheme. The real surface of the object is represented by the curve \mathbf{r} , while the outer and inner meshes are represented by \mathbf{o} and \mathbf{i} , respectively. The flat nature of the triangulation produces an underestimation of the outer mesh. To solve this problem, the position of the mid-point \mathbf{m} of the line \mathbf{o} in Figure 17 (centroid of the triangle, in 3D) is updated to the nearest data point, \mathbf{m}' , and the vertices of the outer mesh are recomputed⁹.

In the case of tGB patches, the volume is determined using the divergence theorem in its full form and the expression for the first fundamental form of a surface, found in classical differential geometry (Boehm and Prautzsch, 1994). Generally, the volume enclosed by a parametric surface $\mathbf{S}(u,v)$ is:

⁹ It should be noted that although this approach is highly conservative, “bound overflow” problems may still be present if the density of the mesh is much lower than the one required to encompass the curvature of the surface.

$$V = \int_{\Sigma} \mathbf{S}(u,v) \cdot \mathbf{n}(u,v) d\sigma$$




Figure 17 2D illustration of the construction of the outer estimate on surface location using the simplex mesh. A) the real boundary \mathbf{r} overflows the outer bound \mathbf{o} ; B) In order to avoid this, the mid-point \mathbf{m} is updated to the nearest data point, \mathbf{m}' , and a parallel surface is then built (see text for details).

where Σ is the surface enclosing the object and $\mathbf{n}(u,v)$ is the surface normal at the point (u,v) . The surface element $d\sigma$ is computed using the usual Jacobian expression from differential geometry:

$$d\sigma = \sqrt{EG - F^2} dudv,$$

where:

$$E = \left(\frac{\partial \mathbf{S}}{\partial u} \right)^2, \quad G = \left(\frac{\partial \mathbf{S}}{\partial v} \right)^2, \quad F = \frac{\partial \mathbf{S}}{\partial u} \cdot \frac{\partial \mathbf{S}}{\partial v}$$

Therefore, the expression for the volume enclosed by the surface $\mathbf{S}(u,v)$ is:

$$V = \iint \mathbf{S}(u,v) \cdot \mathbf{n}(u,v) \sqrt{EG - F^2} dudv$$

Unfortunately, the tGB formula is rational, not polynomial. This means that numerical integration procedures are required to compute the formula above, and this could incorporate round-off errors in the computation. A polynomial interpolant would circumvent this problem.

3.9 Validation issues

MR imaging in studies of neurological diseases such as schizophrenia and MS takes place *in vivo*. It is therefore difficult to define ground truth with which to compare the results obtained from a

segmentation method on clinical data. Several approaches to counter this problem are described next. First, we generate synthetic data sets simulating MR scans by applying a recursive subdivision procedure. This allows us to apply our segmentation and measurement method to objects of known geometry. Another approach consists of using an MR acquisition of a phantom, in this case developed for the study of MS lesions, with known volume. Finally, the measurement precision on clinical MRI is tested by scanning twice, in rapid succession, a patient whose head is oriented differently in the MRI machine (about 20° separation) and comparing the results (which we assume should be the same). Details of these ground truths are now given.

3.9.1 Validation with synthetic data sets

We first simulate MR acquisitions of objects of known geometry, for which an analytic formula is available (e.g. an ellipsoid), by means of a recursive subdivision procedure. At each voxel location, it is determined whether the 8 corners of the voxel are inside the simulated object by evaluating its formula. If this is the case, the voxel is assigned intensity value I_{in} , whereas value I_{out} is given to voxels whose 8 corners are all outside the object. The remaining voxels correspond to boundaries and their intensity values should simulate the partial volume effect. In order to determine the proportion α of the voxel that is inside the object, a recursive subdivision procedure is initiated by dividing the voxel by its centre into 8 smaller cubes. Testing for inclusion continues in the manner described above until all sub-voxels are assigned a value or a recursion limit is reached. Then, intensity value

$$I = \alpha I_{in} + (1 - \alpha) I_{out}$$

is assigned to the voxel. The recursion limit is chosen so that the contribution to the final value of α in the limit is smaller than a certain small value ϵ (typically 1×10^{-10}). Tissue-dependent Gaussian noise can also be applied by modelling I_{in} and I_{out} as Gaussian distributions.

This process effectively generates PVE. The primary interest of this validation technique is its flexibility, since different protocols can be simulated by acting on the intensity values to simulate image modalities (T_1 , T_2 , P_D , etc.), and different voxel sizes and inter-slice gaps can be used to simulate different spatial resolutions. Also, the fact that virtually any shape can be used for the

simulated object provides sufficient flexibility to study the effect of size (relative to voxel size), concavities, texture and global shape on the tested segmentation technique.

3.9.2 Validation with MRI phantoms

A phantom consisting of a group of shapes made from paraffin wax and embedded in an agarose gel is used for the second experiment. By measuring the density of the wax, the true volume can be derived from their weight to within a confidence interval of 2% (Roll et al., 1994). The phantoms were developed to simulate the size and shape of MS lesions, which usually are very small relative to the resolution of the MR acquisition. Figure 18 shows a slice through the MRI phantom (1 mm³ voxels, no gap) and a rendering of the shapes embedded in it. This phantom was generously provided by Dr. Fernando Bello and Prof. Alan Colchester of the NeuroMedIA group of the University of Kent at Canterbury.




Figure 18 One slice through and a rendering of the MRI phantom. Object “cube”, for which results are reported below, is indicated.

3.9.3 Validation with clinical data: repeatability

Ground truth can be generated using *in vivo* clinical data by testing repeatability. A volunteer was scanned twice, in quick succession, the second time with his head rotated through about 20-30° with respect to the first (Figure 19). The data sets are T₁-weighted and each consists of 124 slices of 256×256 voxels of size 0.781251×0.78125×1.7 mm³ (TE=9000ms, TR=24000ms).

3.10 Validation results

The measurements shown below are upper and lower bounds on the volume of the segmented object, corresponding to the volumes of the outer and inner surface estimates, respectively. Several methods are compared: the basic simplex mesh; the simplex mesh after applying the refinement process to lock closer to the data; the simplex mesh with tGB patches; and voxel-counting measurements obtained using an advanced hierarchical thresholding tool (Griffin et al., 1994; Colchester et al., 1996).




Figure 19 Slice 56 of the MRI acquisitions used for the study: straight (left) and rotated(right). Note the significant rotation of the right image.

For the results reported for the voxel-based thresholding tool, upper and lower estimates for the segmentation are obtained by setting two different thresholds. The results shown below were visually validated by an expert in the use of the tool as good segmentations of the object. Volume measurements are obtained by simple voxel counting.

3.10.1 Results for synthetic data sets

In our experiments, we simulated several acquisition protocols in use at our laboratory. Volume dimensions, voxel sizes, and inter-slice gaps are detailed in Table 1. For each of these protocols, several objects are used, ranging from spheres and ellipsoids to smooth objects with concavities or local textures simulating brain gyration.

Protocol	Dimensions	Voxel size (mm)	Gap (mm)
1. SUNY	256×256×20	0.9375×0.9375×5	2
2. Oxford	240×240×30	0.9375×0.9375×5	0
3. High res.	256×256×140	1×1×1	0
4. Med. res.	256×256×70	1×1×2	0
5. Low res.	256×256×35	1×1×4	0

Table 1 Simulated acquisition protocols. Protocols 1, 2, and 3 correspond to real clinical practices applied on data sets available in our data pool, whereas protocols 4 and 5 were included to estimate the effect of voxel size on volume measurements.

Volume measurements are obtained by applying the simplex mesh, simplex mesh with tGB patches, and voxel counting after thresholding. The real volume of the object is computed using its analytic formula. Results are reported next for one of the simulated objects (Figure 20). Volume measurements for the different protocols using the different tested methods are shown in Table 2 and Figure 21.




Figure 20 Slices 8, 11, 14, and 16 of simulation of the test object for which results are reported, using protocol SUNY. Notice the remarkable blurring in top and bottom slices due to PVE.

	SUNY	Oxford	Low res.	Med. res.	High res.
sm upper	220999	221591	218653	214162	212171
voxel upper	199410	200566	197520	190968	188200
tGB upper	191130	193318	191575	191220	190156
real volume	178447	178447	178447	178447	178447
tGB lower	153555	156952	160265	168950	172774
voxel lower	151668	152473	155200	162312	166008
sm lower	146130	149576	152505	160609	164326

Table 2 Volume measurements obtained for the synthetic test object. The middle row shows the real volume of the object. (**sm** = simplex mesh).

The first conclusion that can be drawn from these results is that the confidence bounds are correct in the sense that they bound the real volume. Also, the width of the bounds decreases as

the resolution of the protocol increases. The use of a smooth mesh of tGB patches improves significantly the performance of the simplex mesh. The tGB results are also better (except for the highest resolution protocol, in this example) than those obtained by voxel count based on thresholding. As we will see, these results hold consistently throughout our experiments.




Figure 21 Volume measurements for the test object. The horizontal axis reflects the different simulated protocols, from lowest to highest resolution. The vertical axis shows the volume.

3.10.2 MRI phantoms

We present the (typical) results for the phantom object *cube*, shown in Figure 18. Simplex mesh segmentation, tGB fitting, and simplex mesh refinement volume measures are obtained as described above. Voxel-based volumes are derived from intensity thresholding and posterior voxel count. Tables 3 and 4, and Figures 22 and 23 show these results.

Method	Volume (mm ³)
Lower simplex mesh	776
Lower refined sm	797
Lower voxel-based	821
Lower sm+tGB	823
Real volume	1028
Upper sm+tGB	1114
Upper voxel-based	1152
Upper refined sm	1228
Upper simplex mesh	1238

Table 3 Volume measurements for object cube of the MRI phantom.

Method	Width (%)
simplex mesh	44.9
refined sm	41.9
voxel-based	32.2
sm+tGB	28.3

Table 4 Width of the confidence interval relative to the size of the object ($\frac{V_{upper} - V_{lower}}{V_{real}} \times 100$).




Figure 22 Volume measurements, in mm³, for cube (sm = simplex mesh, U=upper, L=lower). The real volume is plotted as a straight line.




Figure 23 Width of the confidence interval as a percentage of the real volume. From left to right, simplex mesh, refined simplex mesh, voxel measurements by thresholding, and tGB patches fit over simplex mesh.

The results show that the best volume bound estimates are obtained by using a mesh of tGB patches interpolating the simplex mesh. This gives better volume estimates than those obtained using voxel-based methods, while also providing a continuous surface segmenting the object, which can be used for further shape description and processing. It is worth noting that although the refinement method of the simplex mesh improves the volume measurement significantly, the use of a continuous surface results in a better estimate. This is partly due to the fact that simplex mesh nodes are updated to avoid “bound overflow” problems, while this is not the case in the tGB mesh. It could be argued, however, that provided that the density of the mesh is extremely high, the need for such update would be removed and the results of the refined mesh would tend to the ones obtained by using continuous tGB patches. Results similar to the ones reported here were obtained when applying the method to the other shapes in the MRI phantom.

3.10.3 Results using clinical data

The data set is first bias-corrected assuming only one tissue encompassing white matter and grey matter, plus a uniform class modelling the rest of tissues plus CSF and air. Probability maps for the different tissues are generated, and a pre-segmentation step labels voxels with a probability smaller than 95% of belonging to one of the tissues as PVE voxels (see Figure 11). The left lateral ventricle of the patient is segmented on both straight and rotated data sets and volumes are compared. Simplex meshes are fitted to the data using the information derived from the maps to guide it and the refinement process is used to obtain better fit to the target data. Two meshes are fitted to obtain inner and outer bounds on the location of the surface. The numbers of vertices for the fitted refined meshes are 1538 (771 polygons), and 1558 for the rotated set (781 polygons). A set of tGB patches for each mesh is then built. For comparison, voxel-based segmentations are performed and validated by an expert. Volume measurements are shown in Table 5 and Table 6.

Method	Straight	Rotated
Upper simplex mesh	12192	11100
Upper voxel-based	10090	9770
Upper sm+tGB	9454	9574
Lower sm+tGB	7835	7998
Lower simplex mesh	7373	7483
Lower voxel-based	7314	7162

Table 5 Measured volume (in mm³) of the left ventricle (straight and rotated positions). Simplex mesh measurements were obtained after using the refinement process.

Method	Straight	Rotated
simplex mesh	4819	3617
voxel-based	2776	2608
sm+tGB	1619	1576

Table 6 Width of confidence interval (U-L).

Figure 24 shows the upper and lower bounds on the volume of the ventricle, and Figure 25 shows the width of the confidence interval. There are two main points to note about Figure 25. First, the confidence interval is significantly smaller for the tGB model, and second, the interval is almost invariant to the patient head rotation.




Figure 24 Volume estimates for the left ventricle (mm³) for straight and rotated configurations. The vertical bars show the difference between the upper and lower volume estimates.




Figure 25 Width of the confidence interval in the straight and rotated configurations. The best results are obtained using a mesh of tGB patches interpolating the simplex mesh. Note that the results for the tGB patches are almost invariant to patient head rotation.

3.11 Flexibility of the shape model: representing the cortical surface

Finally, we investigate the flexibility of our shape model by segmenting the cortical surface of a patient from a clinical MR data set (T_1 -weighted, 124 slices, voxel size $0.9375 \times 0.9375 \times 1.2 \text{ mm}^3$). The cortex is known for its highly convoluted shape, which makes it a very demanding surface for shape modelling. The volume is first bias corrected and a prior segmentation with a tissue model encompassing both grey matter and white matter, and another model with uniform probability density function to cater for other tissues is created (cf. Figure 11). An ellipsoidal simplex mesh with low number of nodes is first located roughly near the brain and the fitting process is started, driven by the probability maps computed during the bias correction step. Additionally, the refinement process of the mesh is enabled so more nodes are added in areas where more detail is required to fit closely to the target data. Figure 26 illustrates the process and shows visual results.




Figure 26 Mesh fitting process, starting from a mesh with the shape of an ellipsoid and low number of nodes. Refinement is applied to the mesh in order to add nodes in the areas where more detail is required. Different stages on the refinement-fitting process are shown in the first two rows, together with the intersection of the model with one of the slices of the MRI. Enlarged front and top views of the final model are shown in the last row.

3.12 Discussion

This chapter has presented a method for carrying out morphometric studies of structures embedded in MRI data. The method itself is novel in the sense of combining statistical voxel classification and shape-based segmentation. Nevertheless, the most important point in the technique is the establishment of confidence intervals in every measurement by bounding PVE. This is a crucial issue in morphometry that has too often been ignored. We described above several techniques for segmentation and compared their relative performance both in synthetic and clinical data, concluding that the smooth surface model based on the combination of the simplex mesh and triangular Gregory-Bézier patches provides the best results, i.e. gives narrower confidence bounds while still ensuring that the real surface is contained between them.

Regarding the computational cost of this technique, by far the most expensive part is the EM algorithm. When applied to data sets of size $256 \times 256 \times 124$, it may take up to 1.5 hours. This task can be done as a batch process and requires no user intervention other than the prior estimation of tissue parameters. The other components of the method are fast. In particular, mesh fitting takes approximately 4 minutes, including refinement of the mesh, and computation of tGB parameters is performed in less than 4 seconds. These times correspond to the ventricle example reported in section 3.10.3, on a relatively slow machine (SGI INDY R4400, 200MHz, 64 Mb RAM).

Our approach to bounding the real surface of the object to be segmented has so far been based on establishing conservative outer and inner bounds. This allows for an estimation of the error incurred in the segmentation process. It should be noted, however, that the percentage of the volume encompassed by these bounds is considerably large. In the following chapter we proceed to study the contents of the PVE area between the bounds in order to further narrow the width of the confidence intervals.

Estimation of the Partial Volume Effect

4.1 Introduction

The quest for accuracy, initiated in the previous chapter, is seriously challenged by the *partial volume effect* (PVE). The reported results show that the effect of discrete sampling in boundary locations is crucial, affecting volume measurements with an error in the range of 20-60%¹⁰. In this chapter we consider the task of decoding the information contained in PVE voxels in order to reconstruct the surface with sub-voxel accuracy and provide narrower confidence bounds on the measurements.

A framework is developed to model the partial volume effect and build a statistical distribution for it. Then, the uncertainties inherent in each PVE voxel are propagated to construct a distribution on measurements derived from the data (e.g. volume, area, etc.). This framework is novel in that it is strongly statistical. Thus, the end product of the process is not a single value for the shape descriptor of choice – volume will be used for the examples – but a statistical distribution from which we can obtain its probability density function, mode, mean, etc. In addition, confidence bounds to a certain probability value may be computed.

4.2 Mixture model

For simplicity, to date our model only considers PVE voxels containing two different types of tissue. This assumption does not limit the number of tissues modelled for the whole image, since the two tissues present in PVE voxels in separate locations of the image may be different. Cases in which more than two tissue types are present in a voxel are rare in brain studies, since PVE

¹⁰ See also (Niessen, 1997) for a study of the effect of PVE using synthetic brain phantoms. Percentages of volume encompassed by PVE voxels coincide with our estimations.

usually appears in the boundary between two structures. Although the model for PVE voxels developed in what follows can be extended to include such cases quite straightforwardly, the practical improvement expected from such an extension is questionable, and it would introduce instability and computational overhead.

The distribution of intensities generated from samples of a tissue type j is assumed to follow a certain distribution $p_j(I)$. No special assumptions are made regarding this distribution, which can be modelled as an *off the shelf* analytical statistical model or be built based on histograms from a training set. In the next section, the model is developed for the special case of Gaussian distributions.

PVE arises from the presence of more than one tissue in a voxel (cf. section 2.4.2), which in accordance with the physics of MRI, generates intensity values linearly dependent on the proportion of each tissue in the voxel¹¹ (Rinck, 1993). Following our assumption of only two tissue types present in PVE voxels, the intensity of a voxel is determined by the following expression:

$$I^i = \alpha^i I_1^i + (1 - \alpha^i) I_2^i \quad (4.1)$$

where i indexes all PVE voxels, and I_j^i is drawn from p_j , the distribution for tissue j , $j = \{1, 2\}$. In this context, α^i denotes the proportion of tissue 1 in the voxel i , $\alpha^i \in [0, 1]$.

The distributions for individual tissue types can be combined in order to build a mixture model, $p(I^i | \alpha^i)$. This distribution expresses the likelihood of a particular intensity value in a PVE voxel, given the proportion of each tissue in it. The expression for this mixture model can be obtained by analytical operations (cf. next section) or by means of Monte Carlo simulation.

¹¹ This holds for most common imaging modalities. It should be noted, however, that inversion recovery sequences may present problems in cases where one of the tissues present in the PVE voxel has negative magnetisation, while the other tissue has positive magnetisation. In such a case, PVE voxels can have lower intensities than the tissues present in them. This is an uncommon case, and although theoretically possible, we did not find it in any of the images in our data base.




Figure 27 A distribution for the intensities of PVE voxels, given the proportion of the two tissues present in it, is built based on the two tissue distributions, p_1 and p_2 .

In order to estimate PVE, we are interested in obtaining $p(\alpha^i | I^i)$, i.e. the statistical distribution of the proportion α of each tissue given the intensity in a PVE voxel i . Having this distribution, it is possible to compute its mode, which corresponds to the most likely value of α^i , as well as confidence intervals (section 4.4). Using Bayes' theorem, the expression for $p(\alpha^i | I^i)$ is as follows:

$$p(\alpha^i | I^i) = \frac{p(I^i | \alpha^i) p(\alpha^i)}{p(I^i)} \quad (4.2)$$

The term $p(I^i)$ is a normalising constant, and it is computed as:

$$p(I^i) = \int p(I^i | \alpha^i) p(\alpha^i) d\alpha \quad (4.3)$$

$p(\alpha^i)$ is the prior on α , and expresses the probability of a voxel i having a proportion α of tissue 1 (and therefore a proportion $1-\alpha$ of tissue 2). In this work we propose that, if we consider only PVE voxels, this prior can be modelled as a uniform distribution in the range $[0..1]$.

$$p(\alpha^i) = 1, \forall i \quad (4.4)$$

Note that this assumption only holds for the set of PVE voxels. If all voxels in the image were to be considered, a prior giving more probability to the values 0 and 1 would be more appropriate, since *pure* (non-PVE, i.e. $\alpha=0$ or $\alpha=1$) voxels are in general more common than PVE voxels (see Figure 28). Examples of such distributions can be found in (Roll et al., 1994). The problem of using a uniform distribution in the presence of many pure voxels is that small variations from 0 or 1 in the α 's of such voxels produce considerable errors, due to the high number of pure voxels in an image. It should be noted, however, that choosing the correct "U"-shaped prior is not trivial, and a bad choice can influence the results for the estimated α 's. A prior with wide tails, which

tends to “round off” too many values close to 0 and 1 will underestimate PVE, while a prior with narrow tails will not cancel the error due to small variations from 0 or 1 in pure voxels.




Figure 28 Distributions for $p(\alpha^i)$. **A)** Uniform distribution, corresponding to a set of PVE voxels; **B)** “U”-shaped distribution, corresponding to an image containing pure tissue voxels.

Let us now consider the assumptions of independence that can be made within the framework described above. We consider that $p(I^i|\alpha^i)$ is independent across voxels. A reminder of the intuitive meaning of the distribution will suffice to convince that this is a reasonable assumption. $p(I^i|\alpha^i)$ represents the probability of an intensity value in a voxel, given the proportion of each tissue. This only depends on the distributions of the tissues, p_1 and p_2 . The value of α is given, and there is no reason why neighbouring voxels should influence this distribution. Let us now consider $p(I^i)$, the probability of a given intensity in a voxel. This distribution is a direct function of $p(I^i|\alpha^i)$ and $p(\alpha^i)$ as stated in eq. 4.3. In addition, the effect of $p(I^i)$ in the computation of $p(\alpha^i|I^i)$ is irrelevant, since it acts as a normalising constant for the distribution.

The assumption of independence for $p(\alpha^i)$ is more tricky. If pure voxels are taken into account a neighbourhood relation should be incorporated, since voxels whose neighbours are pure tissue have a higher likelihood of being pure, due to the piecewise nature of the tissues. Markov random fields could be employed to that effect, but this is left for future work. In our case, however, only PVE voxels are considered, which opens the possibility for assuming independence. The variability of shapes passing through voxels is infinite (unless some assumptions are introduced to

limit their local geometry). Hence, no reliable prediction can be made to infer α^i from the values of α at neighbouring voxels. Independence is thus a reasonable choice. Further consideration to this argument is given in section 6.2.2.

4.3 Example: Gaussian distributions

We exemplify the framework described above for the case of a Gaussian. Here, the tissues present in the image are modelled as Gaussian distributions of small variance, i.e. the intensities corresponding to a given tissue j are assumed to be close to a mean value μ_j , with a standard deviation σ_j around that value:

$$p_j(I) = G(\mu_j, \sigma_j) = \frac{1}{\sqrt{2\pi}\sigma_j} \text{Exp}\left(-\frac{(I - \mu_j)^2}{2\sigma_j^2}\right) \quad (4.5)$$

This model is the one most commonly used in MRI studies¹². It could be argued that, assuming that the variance around a given mean value for a tissue is very small, a Gaussian would suffice to describe the distribution in intensities, even if the *real* distribution of the intensities obtained from the MRI were not perfectly symmetric.

Continuing with the case of a Gaussian, the expression for the mixture model can be computed analytically. Given a value of α , and assuming the mixture model in eq. 4.1, the distribution for the combined intensities follows a linear combination of two Gaussians, which is itself a Gaussian:

$$p(I^i | \alpha^i) = G(\alpha^i \mu_1 + (1 - \alpha^i) \mu_2, \sqrt{\alpha^{i2} \sigma_1^2 + (1 - \alpha^i)^2 \sigma_2^2}) \quad (4.6)$$

Given the parameters of the distributions for the two tissues present in a PVE voxel, we can plot $p(I^i | \alpha^i)$ for $\alpha^i \in [0..1]$. Figure 29 shows this.

Bayes' theorem can now be used to compute $p(\alpha^i | I^i)$ (eq. 4.2). As mentioned above, $p(\alpha^i)$ is assumed uniform in the interval $[0..1]$, so $p(\alpha^i) = 1$. The resulting formula has the same form as eq.4.6 (except for the normalising constant $p(I^i)$). However, this time it is not a Gaussian, since it is no longer a function of the intensity I , but of the proportion α . This function is plotted in

¹² Note that (Wells et al., 1996; Guillemaud and Brady, 1997) use a Gaussian distribution on the logarithm of the intensities. As discussed in section 3.2.2, this choice is not well suited for dark intensities.

Figure 30. It is simply a rotation of Figure 29, to reflect the change of variable . A few cross-sections at chosen values of I^i are shown in Figure 31. Note that the function is not symmetric.




Figure 29. $p(I^i|\alpha^i)$ plotted for all values of $\alpha^i \in [0..1]$. Each cross-section for a given value of α is a Gaussian following eq.4.6. In this example, $\mu_1 = 100$, $\sigma_1 = 10$, $\mu_2 = 200$, and $\sigma_2 = 40$.




Figure 30 $p(\alpha^i|I^i)$ plotted for all values of $I^i \in [0..255]$.




Figure 31 Cross-sections of $p(\alpha^i | I^i)$ at values $I = 100 (\mu_1)$, 120, 140, 160, 180 and 200 (μ_2).

4.4 Confidence bounds

The information contained in the individual distributions $p(\alpha^i | I^i)$ can be summarised by computing the most likely value of α^i and upper and lower bounds at particular confidence levels. It was shown above that $p(\alpha^i | I^i)$ is not necessarily symmetric, so the mean value α_{mean} is not, in general, the value of highest probability. Assuming that the distribution has only one maximum (the mode, α_{mode}), a simple search starting from the mean value is sufficient to compute this value. α_{mode} represents the most likely value of α given the intensity of the voxel and the distributions for the tissues present in it.

Similarly, confidence intervals on α^i can be set by constraining the area below the probability density function (pdf) of $p(\alpha^i | I^i)$. Upper and lower bounds are found by locating the values α_{upper} and α_{lower} , respectively, which encompass a certain area from the mode α_{mode} . This area is the lateral confidence interval (c_{lat}) for α .

$$\int_{\alpha_{lower}^i}^{\alpha_{mode}^i} p(\alpha^i | I^i)(\alpha) d\alpha = \int_{\alpha_{mode}^i}^{\alpha_{upper}^i} p(\alpha^i | I^i)(\alpha) d\alpha = c_{lat}$$

For example, an area of $c_{lat} = 0.45$ to each side of the mode enforces a 90% confidence interval, i.e. there is a 90% chance that the value of α^i given I^i is between the bounds computed in such a way. This point is illustrated in Figure 32.




Figure 32 Confidence bounds on $p(\alpha^i | I^i)$. α_{upper} and α_{lower} are set so as to make the areas A_1 and A_2 equal to c_{lat} .

4.5 Propagation of uncertainty

The framework developed above is local to each PVE voxel. In order to compute global magnitudes, a method must be devised to combine the statistical information contained in the individual distributions and create a distribution for global descriptors:

$$p(\alpha^i | I^i), \forall i \rightarrow p(V)$$

Common examples of the global descriptor V are volume and area, or any other shape descriptor that can be inferred from the local structure implied by the set $\{\alpha^i\}_{V^i}$.

4.5.1 Conservative bounds

The first approach towards propagating local information makes use of the confidence intervals developed in section 4.4. In particular, upper and lower bounds on the global magnitude V can be set by employing the upper and lower estimates of α for every PVE voxel. This obviously produces an overestimation of the confidence intervals (Pennec, 1996). The bounds computed in such a way will be referenced to as *conservative bounds*.

In the case of volume, the framework is straightforward, since V can be expressed as:

$$V = v_{\text{voxel}} \left(n_{\text{pure}} + \sum_{i \in \text{PVE}} \alpha^i \right)$$

where:

v_{voxel} is the volume of a voxel.

n_{pure} is the number of pure voxels.

i indexes all PVE voxels containing tissue of the object being measured.

Upper and lower estimates of the volume can then be computed as:

$$V_{\text{upper}} = v_{\text{voxel}} \left(n_{\text{pure}} + \sum_{i \in \text{PVE}} \alpha_{\text{upper}}^i \right)$$

$$V_{\text{lower}} = v_{\text{voxel}} \left(n_{\text{pure}} + \sum_{i \in \text{PVE}} \alpha_{\text{lower}}^i \right)$$

This method has two drawbacks. On the one hand, it produces a considerable overestimation of the confidence bounds. This means that using local bounds for each voxel at a confidence level c_{lat} will, in general, produce bounds for the global magnitude V of much greater confidence. This implies an unnecessary loss of resolution in the results reported by this method. On the other hand, establishing confidence bounds is a step back in the process of establishing a solid statistical model for the PVE. In fact, what we are after is a statistical distribution for V , not just two confidence bounds.

4.5.2 Monte Carlo

A simple way to construct this statistical distribution for V is to use Monte Carlo methods. The idea behind this scheme is to sample values from the individual distributions for each PVE voxel. Taking a sample from each PVE voxel, it is possible to compute a sample of V . By repeating this process a large number of times, the distribution for V is simulated and can thus be reconstructed. The process is spelled out in detail in Figure 33. In order to generate samples from the distributions, a rejection method was used (Leon-Garcia, 1994).

```

Algorithm MONTE-CARLO
For n=1 To number of Monte Carlo iterations Do
  volume[n]=0
  For i=1 To number of PVE voxels Do
     $\hat{\alpha}^i$  = sample from  $p(\alpha^i|I^i)$ 

    volume[n] = volume[n] +  $\hat{\alpha}^i$ 
  End For
  volume[n] = volume[n] * voxel volume
End For
pdf = normalised histogram of volume[]
End Algorithm

```

Figure 33 Algorithm for Monte Carlo propagation of uncertainty.

Monte Carlo methods are often computationally expensive, to the point of rendering them ineffective in practical applications. In our case, however, experiments carried out for computing volumes (cf. following section) were performed in satisfactory times (typically in the range of 2-5 minutes).

4.6 *Validation with synthetic images*

A synthetic data set was created using the program described in section 3.9.1. This phantom data set contains a sphere of radius 6.5 in a voxel grid of 20×20×20 voxels of dimensions 1×1×1 each. The dimensions of the sphere were chosen as to approximate the dimensions of the objects contained in the MRI phantom used for validation in section 3.10.2. Only two tissues (inside and outside the sphere) are represented. They are modelled as Gaussian distributions with the following parameters: $\mu_{in} = 200$, $\sigma_{in} = 2.5$, $\mu_{out} = 100$, $\sigma_{out} = 2$. The contents of PVE voxels are determined by means an oct-tree recursive subdivision process, as described in 3.9.1.

The ground truth volume of the sphere, computed analytically, is $V_{ground_theory} = 1150.3$. However, round-off errors in the oct-tree subdivision and the discretisation to create the image may amount to some error in the volume of the sphere represented in the image. In order to obtain a more accurate ground truth volume, a data set representing the same sphere but with no

noise in the tissue distributions (i.e. always using the means) was generated. In such image, the proportion of each tissue in a PVE voxel may be computed as follows:

$$\alpha^i = \frac{\mu_{in} - I^i}{\mu_{in} - \mu_{out}}$$

where I^i is the intensity at voxel i . The volume thus computed amounts to $V_{ground} = 1145.7$, which is the value that will be used as ground truth. The error with respect to the theoretical value is:

$$\frac{V_{ground_theory} - V_{ground}}{V_{ground_theory}} \times 100 = 0.46\%$$

PVE voxels are isolated by building a mask of values different than μ_{in} and μ_{out} in the noiseless image. The volume encompassed by these PVE voxels is 656, which establishes an area of uncertainty based on voxels of 57.3% of V_{ground} . We then proceed to create inner and outer surface bounds as described in Chapter 3. The width of these confidence bounds is computed as:

$$\frac{(V_{outer} - V_{inner}) \times 100}{V_{ground}}$$

The additional smoothness of surface representations is bound to reduce the uncertainty area, when compared to the one based on voxels. In particular, the width of the confidence interval on volume using simplex meshes for this phantom data set is 36.70% of V_{ground} , and this width reduces to 23.30% when using G^1 -continuous tGB patches.

We run the PVE estimation framework on the noisy image, as described in the previous sections of this chapter. The mixture model is built based on the ground truth parameters for the tissue distributions. Using the mode values α_{mode} of $p(\alpha^i | I^i)$ for every PVE voxel, the estimated volume is $V_{mode} = 1146.3$. The error with respect to V_{ground} is 0.05%.

Conservative bounds are established as described in 4.5.1. Several confidence thresholds were used, and results are reported in Table 7. Note that these confidence thresholds do not reflect the confidence on the value of the volume, but on local voxel-based estimates. This means that these confidence bounds are very conservative. In fact, the real volume is comfortably included between the bounds even in cases when a low threshold was used. Experiments performed on real MRI

data suggest that a value of 80% confidence ($c_{lat} = 0.4$) is reasonable, and a more conservative estimation is obtained when using a threshold of 90% confidence. Increasing this threshold to 99% is exaggerated, and the width of the confidence bounds in this case is unreasonably high. The values in Table 7 should be compared to the width of the confidence bounds using the simplex mesh (36.70%) and tGB patches (23.30%). There is a reduction by an order in magnitude.

Confidence threshold	Lower bound (%)	Upper bound (%)	Width %
80%	1131.6 (-1.23%)	1159.3 (+1.19%)	2.42%
90%	1127.6 (-1.58%)	1162.3 (+1.45%)	3.03%
95%	1121.3 (-2.13%)	1164.3 (+1.62%)	3.75%
99%	1012.4 (-11.63%)	1166.8 (+1.84%)	13.47%

Table 7 Conservative confidence bounds on volume for the synthetic phantom. Upper and lower bounds are computed using the values of α_{upper} and α_{lower} , respectively, as described in section 4.5.1. Percentages are computed with respect to $V_{ground} = 1145.7$.

Finally, a Monte Carlo process is used to generate 10000 samples of volume estimates computed from random samples of the distributions of $p(\alpha^i | I^i)$ for every PVE voxel i , as described in 4.5.2. The sample mean was $\mu_{MC} = 1145.5$ (-0.02% error with respect to V_{ground}), and the standard deviation was $\sigma_{MC} = 0.4559$. Confidence bounds based on this values are reported in Table 8. A further reduction by an order of magnitude in the width of the confidence intervals was obtained, the ground truth volume still being contained between the bounds.

Confidence	Lower bound (%)	Upper bound (%)	Width %
$2\sigma_{MC}$	1144.6 (-0.10%)	1146.6 (+0.08%)	0.18%
$3\sigma_{MC}$	1144.1 (-0.14%)	1146.9 (+0.10%)	0.24%

Table 8 Results of the Monte Carlo simulation. Confidence intervals were computed as a function of the standard deviation computed from 10000 samples.

It should be emphasised that the result of this Monte Carlo simulation is not a set of bounds, but a statistical distribution for the volume, from which bounds and other useful information can be derived.

4.7 *Discussion and practical considerations*

The results reported above show an impressive improvement by two orders of magnitude in the resolution of the measurements obtained from MRI data. It should be remarked, however, that in this synthetic example we made use of three pieces of information that are not likely to be available for real data:

- The exact parameters for the distribution of the tissues, $p_j(I)$.
- The composition of PVE voxels (trivial in the synthetic example, which contains only two tissues).
- The location of PVE voxels of interest.

These issues are described in more detail in the following three sections, and solutions are offered. In the following section, sensitivity of the method to slight deviations from the correct tissue parameters is analysed, and a solution based on explicitly modelling the uncertainty on the estimation of such parameters is proposed. Next, a method to enhance the PVE maps described in section 3.3 by including information about the composition of PVE voxels is described in 4.9. This section also describes how contextual knowledge may be used to obtain good voxel-based segmentations including PVE information. Finally, a method for generating PVE masks from the inner and outer surfaces developed in Chapter 3 is described in section 4.10.

4.8 *Sensitivity to tissue parameters: modelling uncertainty in parameter estimation*

The method as it stands is quite sensitive to the tissue parameters. In fact, the method assumes that the tissue distributions $p_j(I)$ are a perfect model. Thus, in the synthetic case presented in the previous section we obtained very narrow bounds around the ground truth volume. However, such extremely narrow bounds also mean that small errors on the means of the tissue distributions will shift the distribution for the volume, yielding bounds that do not contain the ground truth. For example, if in the synthetic phantom we change μ_{in} from 200 to 201 and run the Monte Carlo simulation, $3\sigma_{MC}$ bounds become $[1141.3, 1144.1]$, which do not contain $V_{ground} = 1145.7$. Conservative bounds are not so sensitive to the parameters, and although they slightly shift when incorrect parameters are used, V_{ground} is consistently between the bounds.

For our experiments with MRI data reported in this and the following chapters, tissue distributions are built from samples obtained using a manual “picking” tool. Tissue parameters estimated in such a way may vary in their accuracy and precision. One possible solution is to create more accurate tissue models based on the physics of MRI. We contend that a careful analysis of acquisition parameters and chemical properties of particular tissue types could yield a satisfactory model. More insight into this issue is provided in section 6.2.1.

Another solution is to sacrifice some of the resolution of the method, i.e. “widen” the pdf on the volume¹³ $p(V)$ in favour of accuracy. One possible approach is to model tissue parameters as statistical variables, in order to reflect the uncertainty on the estimation of the parameters. Thus, means and variances can have a most likely value and some standard deviation around it. This is the approach we take. Tissue parameters are estimated from the data set as described above, i.e. obtaining samples of the intensities of tissues by manual selection. However, this process is repeated several times, each yielding an estimate of the tissue parameters (e.g. the mean and standard deviation, for Gaussian models) of the distribution on intensities of the tissue. Next, the mean and standard deviation of such parameters are computed. Therefore, this model assumes that the distribution of the noise in the estimation of the parameters is Gaussian.

Monte Carlo simulation is performed in order to generate the final tissue model. At each iteration of the Monte Carlo process, samples from the distributions of the tissue parameters are drawn. This produces an instance of the intensity distribution for the tissue. A sample from such distribution is obtained and stored, and the process is repeated. After a number of iterations, the histogram of the samples is analysed to build the final intensity model for the tissue (Figure 34). Effectively, the process sketched above is a non-parametric generation of distributions.

Since the framework described in this chapter allows for any type of tissue distribution to be used, a mixture model is then constructed from the tissue distributions computed above and the remaining steps in the PVE estimation framework are as described in the previous sections.

¹³ Volume is used as an illustrative example; it should be kept in mind that V may be any shape descriptor.




Figure 34 Uncertainty on the estimation of tissue parameters is explicitly modelled by representing tissue parameters as statistical distributions. A Monte Carlo simulation is then used to build the final tissue model.

4.9 Use of contextual information to determine the contents of PVE voxels

Commonly, brain scans contain more than two tissue types. It is therefore necessary to determine which two tissues are contained in particular PVE voxels. This is a vital requirement, if we are to build a mixture model to estimate the contents of PVE voxels. Section 3.3 described how the probability maps obtained from the bias correction process can be used to identify PVE voxels. Namely, such voxels are detected by setting a threshold on the minimum value of the maximum probability across tissues. Therefore, if a particular voxel has a probability value greater than the threshold for a particular tissue, the voxel is identified as containing only that tissue type, and otherwise it is marked as PVE. An example is shown in Figure 35.




Figure 35 Voxel-based segmentation obtained as described in section 3.3. A) Original image; B) segmentation (cerebrospinal fluid: black, white matter: light grey; grey matter: dark grey; PVE: white).

In the following, it will be shown that the use of contextual information noticeably improves the results of such voxel-based classifications. Two intuitive ideas can be used to identify the contents of voxels marked as PVE, by looking at their 26 neighbours:

- 1) If a voxel is marked as PVE but all non-PVE neighbours have been given the same tissue tag t , the voxel has a large probability of having been misclassified as PVE. Tissue tag t is then assigned to it. In order to increase the robustness of this criterion, a minimum number of non-PVE neighbours is established (5, in our experiments).
- 2) Otherwise, the two tissues present in the PVE voxel are determined as the two most frequent tissue tags in its neighbourhood.

A further consideration may be taken in order to improve further the results. Let us focus on the T_1 -weighted MRI slice shown in Figure 35. Cerebrospinal fluid (CSF) shows dark intensities, white matter (WM) has bright intensities, and the range of intensities of grey matter (GM) is in between. It is obvious that the intensities of PVE voxels containing CSF and WM could fall into the distribution of GM, thus being wrongly classified. This is illustrated in Figure 36.




Figure 36 The distribution of GM is in the range of intensities of PVE voxels containing CSF and WM. Thus, such PVE voxels may be wrongly classified as GM.

The following criterion is used to avoid such misclassification:

- 3) If a voxel is classified as GM but some of its neighbours are CSF, then it must be a PVE voxel. If the most prevalent tissue across its neighbours is WM, it is classified as PVE(CSF,WM); if it is GM, it is identified as PVE(CSF,GM). This rule works for three tissue

types, which proves sufficient for most practical applications. If more tissues need to be modelled, additional criteria should be established.

Figure 37 shows the improvement on the voxel-based classification achieved using these three criteria. Note that noise in central areas of white matter has been considerably reduced, and the contents of PVE voxels determined. Additionally, voxels around the ventricles, which are PVE(CSF,WM) and were wrongly classified as GM in Figure 35 are now correctly identified.




Figure 37 Improvement on the voxel-based classification by using contextual information. Compare with Figure 35. CSF: black, WM: light grey, GM: dark grey, PVE(WM,GM): white, PVE(CSF,GM): red, PVE(CSF,WM): green.

4.10 Construction of PVE masks from inner and outer surface bounds

One of the assumptions taken in our PVE prediction model is that the prior on proportions of each tissue in a PVE voxel, $p(\alpha^i)$, is uniform. This assumption is valid for the set of PVE voxels, since the variability in shapes passing through them is infinite. If pure voxels are introduced, the prior should model the fact that voxels whose neighbours are pure have a high probability of being pure, due to the piecewise constant nature of the intensities of tissues corresponding to human organs (cf. section 4.2).

In order to isolate the PVE area around the structure of interest, we can make use of the surface bounds created in Chapter 3. These surfaces are constructed to encompass the region of uncertainty in the segmentation due to PVE. Once these inner and outer surfaces are constructed, the problem is reduced to creating a mask with the same voxel structure as the original image,

highlighting the region encompassed by the two bounds. Both masks are then combined and simple region growing techniques are used to include voxels contained between the bounds (combining the inside of the inner bound and the outside of the outer bound). We now describe how to build the masks. The technique described here is closely related to the algorithm used in section 3.9.1 to create synthetic images for phantom validation. In that case, an analytical formula for the object was available, and an oct-tree recursive subdivision process was used by testing the corners of the (sub-)voxels for inclusion in the object. The main difference with respect to the present case rests on the fact that the definition of the object was a volumetric one, whereas the bounds to be used here are surfaces.

The method used previously for generating synthetic MRI phantoms is based on recursively subdividing the voxel grid into smaller, more detailed sub-voxels. The criterion used for deciding whether to subdivide or not a given voxel is based on testing the eight corners of the (sub-)voxel for inclusion in the object. Subdivision is applied to those (sub-)voxels which had some corners inside and some outside. The algorithm is detailed in Appendix B, section B.1.

The two-dimensional nature of the surface bounds, defined as a set of tGB patches, makes it quite difficult to test for inclusion of a 3D point inside the object enclosed by the surface. Fortunately, for the application in mind in this section, we only need to locate and mark the voxels the surface passes through. To this effect, we can take advantage of the parametric definition of the surface. In general, any surface can be defined by a parametric mapping of the form:

$$\mathbf{S}(u, v) = (x(u, v), y(u, v), z(u, v)) \quad u \in D_u, v \in D_v \quad (4.7)$$

where D_u and D_v are the domains for the parameters u and v , respectively. Without loss of generality, in the following we will assume that $D_u = D_v = [0,1]$.

For the particular case of tGB patches, a prior step must be taken in order to establish a map \mathbf{b} from surface coordinates in the domain $[0,1] \times [0,1]$ to barycentric coordinates, used in the definition of tGB patches¹⁴ (see Figure 38):

¹⁴ Note that a simpler approach, consisting on tessellating the triangle formed by the barycentric domain, could be employed, avoiding the need for the map \mathbf{b} . The description above is of a more general nature, valid not only for barycentric definitions of the surface.

$$\begin{array}{lcl}
[0,1] \times [0,1] & \rightarrow & \{[0,1] \times [0,1] \times [0,1] \mid u' + v' + w' = 1\} \rightarrow \mathbb{R}^3 \\
(u, v) & \rightarrow & \mathbf{b}(u, v) = (u', v', w') \rightarrow \mathbf{tGB}(u', v', w') = (x, y, z)
\end{array}$$

This prior mapping is given by the formula:

$$\mathbf{b}(u, v) = (u, v \cdot (1 - u), 1 - u - v)$$




Figure 38 Mapping from surface coordinates (u, v) to three-dimensional points in a tGB patch. First, surface coordinates must be mapped to barycentric coordinates, in order to index the formula for the patch.

Let us then consider the more general problem of creating a mask for a surface defined parametrically in the form of equation 4.7, for any voxel grid. A possible solution would be to densely sample the surface guaranteeing sufficient resolution so as to not miss any voxel. However, it is difficult to establish a sampling step in the parametric space of the surface with these characteristics, unless an extremely conservative – and computationally inefficient – approach is taken.

A recursive subdivision technique is used. Instead of acting on the three-dimensional space of the surface and the voxel grid, this time the subdivision is performed on the two-dimensional parameter space. Therefore, the algorithm used is a *quad-tree* recursive subdivision. The approach consists on dividing the parametric space into squares and computing the location of the 3D points corresponding to the corners of each parametric square. The stopping criterion for the recursive subdivision of these squares is a minimum distance between the four 3D corner points. This minimum distance must, of course, be related to the voxel size. Typically, conservative values around 10^{-8} times the smallest side of a voxel are used, in order to not miss any voxel. The algorithm is illustrated in Figure 39, and detailed in section B.2.




Figure 39 Illustration of the quad-tree subdivision process. The rectangle pr in parametric space maps, through \mathbf{tGBob} , to the surface sub-patch r in 3D space. A voxel is shown, together with its intersection with the tGB patch. Since r does not fall completely inside the voxel, it is subdivided into 4 smaller rectangles, by subdividing pr . The rectangle r_1 falls completely inside the voxel, so the voxel will be marked. The remaining patches will continue the subdivision process to guarantee that no voxels are missed. A threshold area for the sub-patch is used as a recursion limit.

It should be noted that by using information about the possible location of the boundary to isolate PVE voxels we also reduce the effect of voxels incorrectly classified as PVE. For example, voxels which are pure but whose intensities fall in the tails of the tissue distribution may be classified as PVE, thus biasing the results. The number of such voxels is very small, and their incidence is spread over the whole image. Therefore, the number of such voxels occurring at the boundary (i.e. the area isolated by the mask) will be extremely small.

4.11 Experiments on MRI phantoms

The PVE prediction framework is now tested on a real MRI data set. This data set was used as a validation phantom in section 3.10.2. Then, bias field correction was performed, and inner and outer surface bounds were fit. Volume confidence bounds were reported using voxel-based measurements, simplex mesh, and tGB patches. These results are summarised in Table 9. The ground truth interval is provided for reference in the last row.

PVE voxels are identified by rendering the inner and outer surfaces into the voxel grid, as described in the previous section. Voxels contained between the bounds are included into the

PVE mask (Figure 40). The PVE mixture model is built based on tissue parameters computed from sampling a slice of the data set several times. As argued in section 4.8, this allows to explicitly model the uncertainty of the estimation of tissue parameters inherent to the sampling technique. Conservative bounds to a 90% confidence level ($c_{lat} = 0.45$) are established, a Monte Carlo simulation is used to generate 10000 samples of the volume, and statistics are derived from such samples. The mean is 1029.36, and the standard deviation is 1.47. $\pm 3\sigma$ confidence bounds are shown in Table 9.




Figure 40 A) A slice of the phantom data set; B) corresponding slice from the PVE mask generated from the outer and inner surfaces computed in section 3.10.2.

Method	Confidence interval	Width %
Simplex mesh	[899.6, 1264.7]	35.5 %
Voxel	[821.0, 1152.0]	32.2 %
tGB mesh	[954.8, 1166.9]	20.6 %
PVE conservative (90% bounds)	[980.9, 1112.0]	12.7 %
PVE Monte Carlo	$1029.4 \pm 3\sigma = [1024.6, 1033.4]$	0.9 %
Ground truth	$1028 \pm 2\% = [1007, 1049]$	2.0 %

Table 9 Summary of the results obtained for the MRI phantom “cube”, described in section 3.9.2. Percentages are computed with respect to the mean value of the ground truth confidence interval, 1028.

All the estimated bounds are correct, in the sense that they contain the range provided as ground truth. Nevertheless, the confidence interval computed by means of Monte Carlo simulation is smaller than the ground truth. This interval is contained inside the ground truth bounds, and the mean is almost identical to the ground truth mean, which hints at a sensible result. Additionally, $\pm 3\sigma$ comfortably contain the ground truth mean. However, it is not possible to determine whether the Monte Carlo bounds are correct or not, since the resolution of the ground truth interval provided by the manufacturer of the phantom is not sufficient.

4.12 Generation of PVE and intensity images from surface models

Although the algorithm designed in section 4.10 solves the problem at hand, i.e. the need for isolating PVE voxels, the technique can be extended to provide more useful information. In particular, a similar approach based on a quad-tree recursive subdivision of the parameter space can be used to compute the proportion of each voxel to each side of the surface. If we assume that this surface delimits the boundary between two different tissues, then this value corresponds to the proportion of inside (α_S) and outside ($1-\alpha_S$) tissue in the voxel. In other words, the method described next computes estimates of PVE for an object defined by its bounding surface, for any given voxel grid.

If the definition of the object at hand were volumetric, such as the case for the generation of synthetic phantoms, then an oct-tree algorithm such as the one used to generate synthetic images would suffice. However, the description of the object is a parametric surface. Several strategies may be adopted to solve this problem. A method based on the intensive use of computational geometry to compute the intersections of the surface with the voxels and then integrate their volumes can be employed. However, this method is very costly, for several reasons. First, the number of planes delimiting voxels – and hence the number of intersections to be computed – is enormous. Second, the analytical formula for tGB is not a simple, fast, polynomial. Finally, the integration process – to be performed in every PVE voxel – is extremely time-consuming. These factors render this approach impractical.

Our algorithm is based on that described in section 4.10 for creating PVE masks. In fact, the general idea consists of sub-sampling each PVE voxel into smaller sub-voxels, and then applying the same algorithm. This creates a mask in the sub-sampling resolution, marking the location of the surface passing through a PVE voxel. Then, a region growing algorithm is employed to fill the “inside” part of the voxel. See Figure 41 for an illustration of the method. It could be argued that this approach may be extremely memory consuming, due to the need for sub-sampling PVE voxels. However, two considerations should be taken into account. First, the sub-sampling process is not required for every voxel in the image, only for PVE voxels. Second, the storage space can be reduced considerably by storing the sub-voxels as a *height map*, in a similar vein to a

Z-buffer (Figure 41). The direction of this *height* is given by the normal of the first surface point hitting the voxel.




Figure 41 Illustration of the sub-sampling method used to create PVE images of objects defined as parametric surfaces (tGB patches, in our case). Each PVE voxel is sub-sampled and the intersection of the surface passing through it is computed. A region growing algorithm is used to “fill in” the inner part of the voxel (determined by the normal). A 2D implementation of the process, using height maps as Z-buffers, considerably reduces memory requirements.

Thus, PVE can be generated from the intersection of the surface and a voxel grid. Assuming that this surface delimits the boundary between two tissues, and if a model for the intensities of the tissues is at hand, intensity images can be generated. This means that MR images of any modality may be simulated.

4.13 Conclusions

This chapter has presented a framework for modelling and estimation of the partial volume effect. The mixture model allows for sub-voxel resolution in measurements obtained from MRI data. Such mixture model is propagated to build a distribution on global shape descriptors (e.g. volume), or estimate confidence bounds. Results on synthetic data sets show a significant reduction in the width of the confidence intervals, with respect to the results shown in Chapter 3.

Three key issues for the correct use of the mixture model in real MRI data sets have been tackled, namely: the sensitivity to the tissue prior distributions on the estimated distribution for the global shape descriptor, and the need to isolate PVE voxels and determine their contents. Even when taking a conservative approach, the width of confidence bounds is on the range of 1-2%, which means an improvement by more than an order of magnitude with respect to the bounds computed in Chapter 3. This opens the possibility to perform clinical studies on small structures where PVE voxels contain a high percentage of the volume of the structure. A clinical study of this kind is described in the following chapter, illustrating the use of our techniques.

Finally, section 4.12 presented a methodology to create intensity images from surface models, thus closing the link between voxels and continuous surface representations. The importance of this technique rests on the fact that it is a first step towards obtaining a truly sub-voxel segmentation framework. In particular, surface fitting may be enhanced with a criterion to minimise the difference between PVE estimates derived from voxel intensities, α^i , and PVE values determined by the intersection of the surface with the voxel grid, α_S^i . This point is described in more detail in section 6.2.4.

Clinical Study: Symmetry Differences of the Temporal Horns in Schizophrenia

5.1 Introduction

The tools developed in the previous chapter allow for greater resolution in measurements obtained from MRI. It is the purpose of this chapter to determine whether this additional resolution is sufficient to support or refute hypotheses relating to schizophrenia, and in particular its relation to brain asymmetry. The study described here is not only a clinical validation of our method, but an important clinical contribution towards understanding the aetiology of schizophrenia.

Asymmetry anomalies of several anatomical structures in schizophrenia have been reported, but findings so far have been handicapped by the lack of resolution of existing morphometric methods. This leads to a highly controversial literature, with findings that claim to corroborate or deny correlation between schizophrenia and cerebral asymmetry. In choosing an anatomical structure for our study, two criteria were adopted. First, clinical literature must exist relating the structure to schizophrenia. Second, measurement of the structure must be difficult due to its size relative to the voxel size, thus leading to inconclusive published results¹⁵. The *temporal horns* satisfy both criteria, and are the object of this study.

Section 5.2 provides anatomical background describing the temporal horns. Section 5.3 lists the existing literature regarding the importance of temporal horns in schizophrenia. A formal problem statement is given in section 5.4, and the data used in the study is described in 5.5. Methods are

¹⁵ In fact, all structures reported in the clinical literature comply with this condition.

described in 5.6. Sections 5.7 and 5.8 present and analyse the results, and a clinical interpretation is provided in section 5.9. Finally, conclusions and further discussion are provided.

5.2 *Anatomical background*

The human brain has an extremely complex structure. It is primarily composed of *grey matter* and *white matter*. Grey matter is formed by the nuclei of the neurons, and it is located in the outer surface of the brain (*cortex*), as well as in several internal *nuclei*. The brain's interior is mainly composed of white matter, which corresponds to the *axons*, establishing connections between neurons. Through evolution, the human brain has grown in complexity, both in number of neurons and, especially, connections between them. This process demanded an increase in cortical area, to cater for more active nuclei, but also required proximity between nuclei, so connections could be established. An optimal configuration was achieved by means of which the brain cortex became an extremely convoluted surface of high intrinsic geometry (Griffin, 1994). The cortex presents a large number of folds (*sulci*), which delimit areas called *gyri*. The location of sulci and gyri is far from random, and they determine the structural and functional organisation of the individual brain.

Globally, brains are structured into two roughly symmetric parts called *hemispheres*, connected at particular central locations in the brain, the most important of which is the *corpus callosum*¹⁶. Each hemisphere is in turn organised into *lobes*, defined by the location of specific sulci (Figure 42). Brains are not perfectly symmetric. As a result of evolution, several regions of the brain grew larger than their counterparts in the opposite hemisphere. Such is the case, for example, of the left posterior areas (temporal and occipital lobes) of the brain, which perform language-related tasks. The functional importance of cerebral asymmetry is well recognised, and literature exists relating anomalies in such asymmetry to schizophrenia (cf. section 1.2.2).

The brain obtains structural support via the *cerebrospinal fluid* (CSF), in which it is embedded. Several internal chambers, called *ventricles*, also contain CSF (Figure 43). The two *lateral ventricles* are

¹⁶ The importance of these central areas stems from the fact that all connections between left and right hemispheres must pass through them.

irregular cavities situated in the lower and medial parts of the cerebral hemispheres. They are almost separated from each other, but each communicates with the *third ventricle* and indirectly with each other through the *interventricular foramen*. Each lateral ventricle consists of a central body and three *horns*, anterior, posterior and inferior (Figure 44). Inferior horns are located inside the temporal lobe, and are also commonly known as *temporal horns*.




Figure 42 The brain can be partitioned into lobes, defined by the location of specific sulci. Lobes are coloured in the figure: frontal (blue), parietal (green), occipital (yellow), and temporal (violet). The cerebellum (pink), brain stem (magenta), and eyes (white) are also shown. Source: Digital Anatomist Program, University of Washington (<http://www.axess.com/users/ir/atlas>).



Figure 43 Ventricles are internal chambers containing cerebrospinal fluid (CSF).




Figure 44 Structure of the ventricles. There are two lateral ventricles, each one in a cerebral hemisphere, and the medially located third and fourth ventricles. The two lateral ventricles are divided into a main body and three horns. Our study focuses in the inferior, or temporal, horns. It should be noted that the size of these horns increases considerably with age, and the figure corresponds to a post-mortem brain. Temporal horns are commonly much thinner. Source: Digital Anatomist Program, University of Washington (<http://www.axess.com/users/ir/atlas>).

5.3 *Related clinical studies*

Ventricular shape differences between schizophrenics and controls have been reported in several studies. Most of them report enlarged lateral ventricles, and some report differences in asymmetry. Some studies have hinted at the importance of studying subdivisions of the lateral ventricles, and in particular the temporal horns. A detailed discussion of these studies falls out of the scope of this thesis. Reviews of these and other anatomical alterations reported in schizophrenia can be found in (Lawrie and Abukmeil, 1998; McCarley et al., 1999).

5.4 Problem statement

Clinical studies such as those mentioned in the previous section are based on creating binary voxel masks to isolate the temporal horns, and no attempt is made to estimate PVE and obtain sub-voxel accuracy. Figure 45 shows magnified slices through the medial part of the temporal horns. It should be apparent that PVE accounts for most of the volume of the temporal horn. In fact, the middle slice shows no clear trace of the temporal horn, since its diameter, when traversing that slice, is smaller than the voxel width. Thus, it is very difficult to represent the temporal horns using a voxel-based model. Figure 46 shows a rendering of one of such segmentations. Notice that the temporal horns are disconnected from the main bodies of the lateral ventricles.



Figure 45 Three MRI slices (numbers 25, 29, and 47, from back to front) showing the temporal horns. Notice that they traverse very few voxels. The middle slice, in particular, presents a very unclear delineation of them, remarking the need for sub-voxel resolution.



Figure 46 Rendering of a voxel-based segmentation of the ventricles. Note that the temporal horns are unconnected. It is very difficult to perform a binary classification, due to the extremely small size of the temporal horns with respect to the voxel size.

The purpose of this study is to accurately segment the temporal horns from MRI acquisitions of a set of schizophrenic patients and a set of normal controls. A statistical analysis will then be performed to assess whether there exist significant differences in left-right asymmetry of the temporal horns. The statistical apparatus provided by the PVE estimation framework developed in the previous chapter allows for the computation of “significance” measures.

5.5 *Data used in the study*

High resolution coronal T₁-weighted MRI data sets of 8 schizophrenic patients (SC) and 8 normal controls (NC) are analysed. The data sets consist of 124 slices (slice thickness 1.5 mm), each consisting of 256×256 voxels of dimensions 0.9375×0.9375 mm². Therefore, the field of view (FOV) is 240×240×186 mm³. A few slices through one of the data sets are shown in Figure 45.

These scans were acquired at the State University of New York (SUNY) in Stonybrook, NY, USA, and are part of the BIOMORPH data pool. The particular details of each data set are listed in Table 10.

Data set	Age	Sex	Handedness¹⁷
006_SC	36	M	R
007_SC	31	F	L
008_SC	29	F	R
010_SC	33	F	R
024_SC	28	M	R
025_SC	31	M	R
033_SC	22	F	R
039_SC	26	F	L
104_NC	30	F	R
105_NC	30	M	R
106_NC	42	M	R
107_NC	36	F	L
112_NC	29	F	R
114_NC	34	M	R
117_NC	31	M	R
118_NC	38	F	R

Table 10 Details of the subjects used in the study. Schizophrenic patients are labelled SC (all subjects in this study are chronic patients), and normal controls are indicated with the letters NC. M: male, F: female, L: left-handed, R: right-handed.

¹⁷ Handedness has an influence in lateral brain specialisation, and therefore in symmetry.

5.6 Methods

This section describes the processing steps performed in order to assess anatomical differences of the temporal horns in the data sets included in the study. Our analysis focuses on volume as a shape descriptor. More complex shape descriptors could be derived following the same processing scheme. Figure 47 illustrates the processing pipeline.




Figure 47 Processing steps followed in this study. Details are provided in the following subsections.

5.6.1 Bias field correction

Correction for bias fields is performed using the EM algorithm, as described in section 3.2. In order to run the algorithm, estimates for the tissue parameters must be obtained. The algorithm assumes that intensities for a given tissue follow a Gaussian distribution on the logarithm of the intensity. Intensity samples are taken from a slice of the data set. Sampling is restricted to a relatively small region in the image, to avoid intensity differences due to the bias field. This process is performed for the tissues present in the image, namely white matter (WM), grey matter (GM) and cerebrospinal fluid (CSF). Samples for each tissue type are log-transformed and the sample mean and standard deviation are computed. For completeness, these are listed in Appendix C, Table 18.

The EM algorithm is run allowing 4 iterations of the expectation-maximisation framework. Typical computation times are around 10 minutes. The output of the algorithm consists of an estimate of the bias field, a bias-corrected version of the image (on which we will work from now on), and three probability maps containing, for each voxel, the probability of it being CSF, GM and WM (Figure 48).




Figure 48 Output of the bias correction process on slice 13 of data set Main006_SC. A) Original image; B) estimated bias field; C) corrected image; D) CSF probability map; E) GM probability map; F) WM probability map.

5.6.2 Determination of contents of PVE voxels

Analysis of the probability maps allows for a classification of the voxels in the image, as previously described in section 3.3. The tissue whose probability value is highest at a given voxel is considered. If such a probability, as indicated in its probability map, is higher than a certain threshold (0.6 was used in the present experiment), it means that the intensity of the voxel adheres closely to the intensity distribution for the tissue. The voxel is then marked as a *pure* voxel containing only that tissue. Otherwise, the voxel is marked as PVE (Figure 49).



Figure 49 Initial classification into pure voxels and PVE voxels. CSF is black and PVE is coloured white. White matter is marked with a brighter intensity than grey matter.

Contextual information is used to further improve the classification and determine the composition of PVE voxels. Three basic ideas are employed, based on the analysis of the 26-voxel neighbourhood of PVE and GM voxels (cf. section 4.9). The use of contextual knowledge provides a cleaner segmentation. Additionally, it re-classifies PVE voxels, based on their composition, as PVE(GM,WM), PVE(CSF,GM), or PVE(CSF,WM). This is illustrated in Figure 50.




Figure 50 Improved voxel-based classification obtained using contextual information. CSF: black, GM: dark grey, WM: light grey, PVE(GM,WM): white, PVE(CSF,GM): red, PVE(CSF,WM): green.

5.6.3 Identification of regions of interest

The voxel-based segmentations obtained by combining the probability maps provide a high-quality decomposition of the image into tissue types, including the composition of PVE voxels. After this has been obtained, it is necessary to select the region of interest, i.e. those voxels that correspond to the structure of interest. This process is performed using a semiautomatic region

growing technique on the voxel segmentations, keeping an eye on the intensity images for visual feedback. The tool used for performing the segmentations is part of the preoperative planning software VISLAN (Colchester et al., 1996).

Seeds are selected manually and the thresholds of the region growing process are set so as to include voxels marked as CSF (pure voxels inside the temporal horns), PVE(CSF,GM) and PVE(CSF,WM). In most cases, the voxel-based classification provides natural boundaries for the structures being segmented. Otherwise, the software allows for manual editing of the results to stop the region growing process at some user-defined boundaries. Additionally, and due to the particular anatomical complexity of the temporal horns, some central areas of the horns may present problems. The width of the temporal horns in such areas may be extremely small with respect to the voxel size, occupying less than one voxel (Figure 51). Based on the analysis of the intensities of these voxels, they tend to be classified as pure GM or PVE(GM,WM). Contextual information may not be used to provide a better classification, since the temporal horn only traverses one voxel. Such voxels are segmented manually using the manual editing capabilities of the segmentation tool, and re-defined as PVE(CSF,WM). It should be noted that some heuristics could be designed to improve the detection of such cases in the central areas of the temporal horns. However, purely intensity-based heuristics, though effective in the affected voxels, render the voxel-based segmentation process much less robust, and worsen the results in other areas of the brain. Clearly, the incorporation of additional anatomical knowledge could help. It should be kept in mind that these cases are not very frequent, and for the most part of the volume of the temporal horns basic seeding and region growing on the voxel segmentations suffices.




Figure 51 The definition of the region of interest is complicated in some central areas of the temporal horns, where their thickness is minimal. Slice A) is located in such a location. Probability maps for CSF, GM and WM are shown in C), D), and E), respectively. Notice that very low CSF probability is assigned in the areas of the temporal horns, while voxels located in these areas present high GM probability values. This results in such voxels being classified as GM in the final segmentation (B).

The anterior boundary of the temporal horn is naturally defined by its anatomy (Figure 44). On the other hand, the anatomical definition of the posterior boundary of the temporal horns is quite vague. This is a common problem when working with biological structures of non-spherical topology, where a “cutting” criterion is required. This is very important for morphometric studies, and should be addressed by anatomists in order to define unambiguous separations for such structures as hemispheres, lobes, and sub-structures such as temporal horns. Consultations with neuroanatomists lead to the decision of using the first slice before the trigone (anterior tip of the atrium, Figure 44) appears when moving from front to back. Unfortunately, this definition is not independent of the orientation of the patient in the scanner. Several alternative criteria were considered, mainly based on fitting a cutting plane based on the location of anatomical landmarks in the ventricles. However, ventricles are very smooth structures and it is difficult to unambiguously define such landmarks. The uncertainty on locating such landmarks from voxel images where PVE is predominant would render the estimation of the cutting plane ineffective.

The reader may wonder why the scheme described in Chapter 3, based on the use of the simplex mesh and tGB patches, was not employed for segmenting the temporal horns. This technique was proven successful in segmenting ventricles (cf. 3.9.3) and even the cortical surface (section 3.11). Experiments were performed in order to segment the temporal horns, but results were not favourable. This is due to the fact that the technique described in Chapter 3 is guided by purely voxel-based information, and no sub-voxel capabilities are incorporated. This means that the mesh has problems segmenting extremely thin central areas of the temporal horn, where it traverses a portion of a voxel. The mesh may be upgraded to take into account sub-voxel information, as suggested in section 4.13. More detail about this possibility is provided in section 6.2.4.

5.6.4 *Voxel-based volume bounds*

Once the region of interest has been defined, voxel-based volume estimates may be obtained simply by multiplying the number of voxels times the voxel volume ($0.9375 \times 0.9375 \times 1.5 = 1.31836$). A lower bound on the volume is obtained by counting only voxels classified as CSF, and an upper bound is computed by including PVE voxels. Results are listed in section 5.7, Table 11.

For comparison, voxel-based segmentations performed by an expert neuroanatomist using the region growing tool on the intensity images are available. The result is a binary mask indicating which voxels are classified as temporal horn. The awkward nature of such purely voxel-based segmentations is reflected by the fact that some slices in the central areas of the temporal horns do not contain any voxel in the segmentation, again illustrating the need for analysing PVE. Volume estimates for these segmentations are obtained by multiplying the number of marked voxels times the voxel volume, and listed in Table 12.

5.6.5 *PVE analysis: conservative bounds and Monte Carlo simulation*

Following the scheme described in Chapter 4, mixture models are built in order to construct statistical distributions for the PVE as a function of voxel intensity. Conservative bounds based on adding voxel-based confidence limits (cf. section 4.5.1) are computed. From these, upper and

lower volume estimates are obtained. The relative improvement of these bounds can be measured as the difference in width (upper minus lower) with respect to the voxel-based bounds obtained in section 5.6.4. Results are listed in Table 13 in the following section.

Next, Monte Carlo simulations are performed in order to build statistical distributions for the volumes of the temporal horns. As mentioned in section 4.8, the uncertainty in the estimate of the tissue parameters is modelled explicitly by computing such parameters several times and computing their means and standard deviations (listed in Appendix C, Table 19). The mixture model is then built and the Monte Carlo simulation initiated. 10000 iterations were applied in our experiments. Sample means and standard deviations were computed and are listed in the next section (Table 14).

5.7 Results

Table 11 lists upper and lower bounds on volumes computed using voxel-based methods, as described in section 5.6.4. For subsequent comparison, voxel-based volume estimates of segmentations performed by an expert using the region growing tool directly on the intensity images are shown in Table 12. Table 13 lists results using the method described in section 4.5.1 to establish conservative bounds. Volume bounds using local confidence levels of 80% and 90% are listed. It should be borne in mind that such local confidence levels, when propagated to compute the bounds, produce an overestimation, thus providing volume bounds of much greater confidence. Finally, Table 14 shows the results obtained via Monte Carlo simulation (cf. 4.5.2). The end-product of such process is a statistical distribution on the volume, in this case following a Gaussian model. The mean and standard deviation of such Gaussian are indicated, together with confidence bounds obtained by adding or subtracting to the mean 3 times the value of the standard deviation ($\pm 3\sigma$).

Data set	Left temporal horn		Right temporal horn	
	Lower	Upper	Lower	Upper
006_SC	141.06	1036.23	263.67	1476.56
007_SC	113.38	822.66	87.01	878.03
008_SC	216.21	1078.42	97.56	631.49
010_SC	247.85	1236.62	152.93	924.17
024_SC	123.93	996.68	152.93	1379.00
025_SC	98.88	762.01	123.93	1095.56
033_SC	80.42	539.21	316.41	1396.14
039_SC	146.34	1202.34	47.46	487.79
104_NC	101.51	850.34	77.78	610.40
105_NC	21.09	225.44	130.52	893.85
106_NC	104.15	693.46	159.52	949.22
107_NC	50.10	452.20	233.35	1058.64
112_NC	199.07	1150.93	139.75	817.38
114_NC	84.37	557.67	118.65	711.91
117_NC	292.68	1409.33	201.71	1152.25
118_NC	75.15	512.84	72.51	568.21

Table 11 Voxel-based upper and lower bounds on volume (cf. section 5.6.4). Measurements are expressed in mm³.

Data set	Left	Right
006_SC	300.59	462.74
007_SC	427.15	346.73
008_SC	242.58	123.93
010_SC	490.43	373.10
024_SC	543.16	500.98
025_SC	287.40	301.90
033_SC	188.53	561.62
039_SC	665.77	433.74
104_NC	304.54	196.44
105_NC	189.84	428.47
106_NC	264.99	299.27
107_NC	141.06	416.60
112_NC	332.23	226.76
114_NC	160.84	193.80
117_NC	738.28	477.25
118_NC	179.30	167.43

Table 12 Voxel-based volume estimates (in mm³) obtained from segmentations performed by a neurologist expert in the use of the region growing tool.

Data set	Left temporal horn				Right temporal horn			
	80% Low	80% Up	90% Low	90% Up	80% Low	80% Up	90% Low	90% Up
006_SC	315.31	410.03	304.88	425.25	479.75	603.50	465.94	623.53
007_SC	260.17	386.96	247.38	481.38	273.25	417.95	258.46	522.30
008_SC	355.89	472.81	344.00	539.73	199.69	270.32	192.35	298.21
010_SC	413.65	628.95	390.57	744.17	310.88	483.29	291.72	572.80
024_SC	301.77	510.42	248.46	557.54	369.51	629.62	321.34	692.17
025_SC	225.15	363.10	189.50	385.26	275.42	464.70	233.30	495.82
033_SC	157.22	222.07	100.13	229.72	465.34	611.75	356.62	628.94
039_SC	302.51	566.11	199.57	600.12	149.89	243.29	83.01	254.80
104_NC	178.32	290.07	123.84	300.84	141.87	217.47	90.35	225.11
105_NC	91.24	119.73	87.98	142.25	272.49	370.43	263.48	455.30
106_NC	196.65	273.26	177.18	283.55	267.70	360.54	249.23	372.97
107_NC	127.30	184.32	120.81	193.45	372.34	488.63	359.69	507.18
112_NC	349.94	448.07	336.88	461.43	240.90	306.00	231.98	315.01
114_NC	188.70	252.92	181.91	290.30	224.31	295.02	217.25	336.41
117_NC	481.15	647.50	461.35	675.23	353.65	481.24	339.24	502.48
118_NC	149.70	212.59	97.79	220.79	143.70	212.35	100.65	220.69

Table 13 Conservative bounds on volume, in mm³. 80% and 90% confidence values are shown. These confidence values correspond to local bounds; global confidence on the measurements is much higher (cf. section 4.5.1).

Data set	Left temporal horn				Right temporal horn			
	Mean	Std. dev.	3std Low	3std Up	Mean	Std. dev.	3std Low	3std Up
006_SC	364.99	2.98	356.05	373.93	542.85	3.42	532.59	553.11
007_SC	314.04	4.07	301.83	326.25	334.47	4.26	321.69	347.25
008_SC	400.11	2.85	391.56	408.66	227.72	2.15	221.27	234.17
010_SC	502.81	5.10	487.51	518.11	380.45	4.63	366.56	394.34
024_SC	421.01	3.91	409.28	432.74	519.04	4.29	506.17	531.91
025_SC	305.35	3.36	295.27	315.43	388.88	4.03	376.79	400.97
033_SC	196.09	2.28	189.25	202.93	555.55	3.54	544.93	566.17
039_SC	486.43	5.32	470.47	502.39	213.49	2.97	204.58	222.40
104_NC	254.75	2.89	246.08	263.42	192.10	2.40	184.90	199.30
105_NC	102.34	1.31	98.41	106.27	312.28	2.55	304.63	319.93
106_NC	241.44	3.40	231.24	251.64	324.57	3.69	313.50	335.64
107_NC	155.76	2.04	149.64	161.88	433.83	2.90	425.13	442.53
112_NC	427.57	3.12	418.21	436.93	289.06	2.50	281.56	296.56
114_NC	214.52	1.96	208.64	220.40	253.26	2.10	246.96	259.56
117_NC	571.42	3.94	559.60	583.24	422.19	3.37	412.08	432.30
118_NC	193.57	2.36	186.49	200.65	192.06	2.51	184.53	199.59

Table 14 Monte Carlo results. The end-product of the Monte Carlo simulation is a Gaussian distribution on volume. The mean and standard deviation of such distribution are indicated for each subject, both for the left and right temporal horns. $\pm 3\sigma$ confidence bounds are also listed. All measurements are indicated in mm³.

5.8 *Analysis of results*

The first striking result is the size of voxel-based bounds (Table 11). As noted earlier, this is due to the fact that the temporal horns are a very thin structure, so their intersection with each slice is very small relative to the voxel size. This means that most of the voxels they traverse are PVE voxels. Figure 52 shows upper and lower bounds computed using the different methods employed in this study. These results correspond to the left temporal horn of patient 006_SC. Note that voxel bounds are much wider than those obtained with techniques that analyse PVE.

Width of the confidence intervals, given by the difference between upper and lower bounds, undergoes a remarkable reduction when making use of an explicit model for PVE. Figure 52 clearly shows this fact. Widths expressed as a percentage of the mean Monte Carlo volume are listed in Table 15, and averages over all patients are plotted in Figure 53. Conservative PVE bounds considerably reduce the confidence interval, from 228.45% average width using voxel bounds to 55.75% for a 90% local confidence level, or 35.27% for 80% confidence. Note that in all our experiments 80% confidence bounds consistently contain the mean volume (or ground truth, for validation studies reported earlier). Monte Carlo estimates produce further reduction by another order of magnitude, to 6.08%.

The problem of studying anatomical characteristics of the temporal horns from MRI data has been considered intractable to date. This becomes clear if we take a look at the “expert” results reported from voxel-based segmentations using the region growing tool, methodology that exemplifies techniques used to date to segment anatomical structures from MR data. Visual inspection of the results (cf. rendering in Figure 46) suffices to prove that such segmentations are incomplete. Differences with respect to the mean Monte Carlo volumes are wide and follow no predictable pattern (mean difference is 21.79%, with a standard deviation of 23.01). No conclusive studies about asymmetry of the temporal horns have been reported to date, due to the lack of appropriate tools. The results reported in our experiments show confidence bounds in the range of 6% of the estimated volume. Thus, we have succeeded in transforming an intractable problem into one that is feasible.





Figure 52 Volume bounds, in mm³, for the left temporal horn of patient 006_SC. The top graph shows results for voxel-bounds, conservative bounds to 90% and 80% local confidence levels, and $\pm 3\sigma$ bounds computed from the results of the Monte Carlo simulation. For clarity, the bottom graph shows a close-up of the results obtained using the last three methods. The mean value of the Monte Carlo volume distribution is marked as a dark dot.

Data set	Left temporal horn				Right temporal horn			
	Voxel	Con90%	Con80%	MCarlo	Voxel	Con90%	Con80%	MCarlo
006_SC	245.26	32.98	25.95	4.90	223.43	29.03	22.80	3.78
007_SC	225.86	74.51	40.37	7.78	236.50	78.88	43.26	7.64
008_SC	215.49	48.92	29.22	4.27	234.47	46.49	31.02	5.66
010_SC	196.65	70.32	42.82	6.09	202.72	73.88	45.32	7.30
024_SC	207.30	73.41	49.56	5.57	236.22	71.45	50.11	4.96
025_SC	217.17	64.11	45.18	6.60	249.85	67.51	48.67	6.22
033_SC	233.97	66.09	33.07	6.98	194.35	49.02	26.35	3.82
039_SC	217.09	82.34	54.19	6.56	206.25	80.47	43.75	8.35
104_NC	293.95	69.48	43.87	6.81	277.26	70.15	39.35	7.50
105_NC	199.68	53.03	27.84	7.68	244.44	61.43	31.36	4.90
106_NC	244.08	44.06	31.73	8.45	243.31	38.12	28.60	6.82
107_NC	258.15	46.64	36.61	7.86	190.23	34.00	26.81	4.01
112_NC	222.62	29.13	22.95	4.38	234.43	28.72	22.52	5.19
114_NC	220.63	50.53	29.94	5.48	234.25	47.05	27.92	4.98
117_NC	195.42	37.43	29.11	4.14	225.15	38.67	30.22	4.79
118_NC	226.11	63.54	32.49	7.32	258.10	62.50	35.74	7.84

Table 15 Widths of confidence intervals established by the different methods used in this study. All values are expressed as percentages of the mean value of the Monte Carlo simulation.




Figure 53 Average width of confidence intervals on temporal horn volumes. From left to right, voxel bounds, 80% and 90% conservative bounds, and $\pm 3\sigma$ bounds on Monte Carlo results. Values are expressed as percentages of the mean value of the Monte Carlo simulation.

5.9 Clinical conclusions

We now analyse the results from a clinical standpoint. Symmetry is the focus of this study, so it is necessary to define a suitable measure. We define the following normalised symmetry coefficient:

$$S = \frac{L - R}{L + R}$$

where L denotes the volume of the left temporal horn, and R is the volume of the right temporal horn.

Data set	Sym. coef.
006_SC	-0.20
007_SC	-0.03
008_SC	0.27
010_SC	0.14
024_SC	-0.10
025_SC	-0.12
033_SC	-0.48
039_SC	0.39
104_NC	0.14
105_NC	-0.51
106_NC	-0.15
107_NC	-0.47
112_NC	0.19
114_NC	-0.08
117_NC	0.15
118_NC	0.00

Table 16 Symmetry coefficients for all subjects.

A t-test (Chatfield, 1983) on this data may be applied to test for significant differences between schizophrenics and normal controls. The t-test determines the probability of two samples being drawn from distributions with the same mean, and is especially suited for small number of sample values. Using a two-tailed t-test, the probability of both sets of symmetry coefficients (SC and NC) being drawn from a distribution with the same mean is 0.6. This means that no significant differences in asymmetry of temporal horns between controls and schizophrenics were found.

The reader may be wondering, however, whatever happened to the confidence intervals. A huge effort has been made to yield bounds on the volumes of left and right temporal horns of the subjects under study. However, symmetry coefficients have been developed purely from the mean

Monte Carlo values. We now proceed to propagate the standard deviations obtained from such Monte Carlo process in order to establish an estimate of the standard deviation of the symmetry coefficients shown above.

Volume distributions for left and right temporal horns follow Gaussian distributions, in virtue of the central limit theorem:

$$L \leftarrow G(\mu_L, \sigma_L), R \leftarrow G(\mu_R, \sigma_R)$$

Since linear combinations of Gaussian distributions are Gaussian, we deduce that:

$$L - R \leftarrow G(\mu_L - \mu_R, \sqrt{\sigma_L^2 + \sigma_R^2}), L + R \leftarrow G(\mu_L + \mu_R, \sqrt{\sigma_L^2 + \sigma_R^2})$$

Let us introduce the following change of notation to simplify the following steps. Denote $L - R$ as $A \leftarrow G(\mu_A, \sigma)$, and $L + R$ as $B \leftarrow G(\mu_B, \sigma)$. Note that both distributions have the same standard deviation ($\sigma = \sqrt{\sigma_L^2 + \sigma_R^2}$). Now, the symmetry coefficient can be defined as:

$$S = f(A, B) = \frac{A}{B}$$

Such a expression is unbiased (Chatfield, 1983), i.e.

$$\mu_S = \frac{\mu_A}{\mu_B} = \frac{\mu_L - \mu_R}{\mu_L + \mu_R}$$

This means that the values in Table 16 are in fact the correct means of the distribution for the symmetry coefficients for the subjects in the study.

In order to approximate the standard deviation of the symmetry coefficient, σ_S , we make use of the following formula from standard error-propagation theory (Chatfield, 1983):

$$\sigma_{f(A,B)} = \sqrt{\left(\frac{\partial f}{\partial A}\right)^2 \sigma_A^2 + \left(\frac{\partial f}{\partial B}\right)^2 \sigma_B^2}$$

For the particular case of the symmetry coefficient:

$$\sigma_S = \sqrt{\left(\frac{1}{B}\right)^2 \sigma^2 + \left(-\frac{A}{B^2}\right)^2 \sigma^2} = \sigma \cdot \sqrt{\frac{1}{(L+R)^2} + \frac{(L-R)^2}{(L+R)^4}}$$

Table 17 lists the mean and standard deviation of the symmetry coefficients for all subjects in the study.

Data set	Mean sym	Std sym
006_SC	-0.20	0.005
007_SC	-0.03	0.009
008_SC	0.27	0.006
010_SC	0.14	0.008
024_SC	-0.10	0.006
025_SC	-0.12	0.008
033_SC	-0.48	0.006
039_SC	0.39	0.009
104_NC	0.14	0.008
105_NC	-0.51	0.008
106_NC	-0.15	0.009
107_NC	-0.47	0.007
112_NC	0.19	0.006
114_NC	-0.08	0.006
117_NC	0.15	0.005
118_NC	0.00	0.009

Table 17 Mean and standard deviation of the symmetry coefficients computed on all data sets.

The standard deviations listed in Table 17 are very small, less than 0.01, for values of the symmetry coefficient between -1 and 1 . This means that we can be very confident on the mean values computed for the symmetry coefficients. This result was the expected one, since when performing random effects analysis such as the one above cross-subject variance includes, and therefore reflects, the intra-subject variance. Thus, the result of the t-test reported at the beginning of this section still holds, i.e. **there are no significant group differences in left-right symmetry of the temporal horns between the schizophrenics and normal controls analysed in this study.**

5.10 Discussion

The study presented in this chapter shows how the morphometric tools developed in this thesis significantly improve the resolution of measurements obtained from MRI data. Narrow confidence bounds allowed for the temporal horns to be accurately measured and analysed, problem which remained intractable to date. Additionally, the statistical framework allowed for a propagation of uncertainties, which resulted in estimates of the accuracy of the symmetry coefficients and subsequent analysis of the significance of the final results.

Clinically, this study is an important contribution towards assessing anatomical differences between schizophrenics and healthy subjects. However, a note of warning should be stated, for several reasons. First, the number of subjects included in the study was rather low, and there were differences in sex and handedness between them (cf. Table 10). Some studies have reported that such factors may be crucial in the incidence of schizophrenia on the anatomy of the patient, particularly relating symmetry differences. Another issue to be taken into account is the ambiguity of the definition of the posterior boundary of the temporal horn. The need to define arbitrary cutting planes to separate the temporal horns from the bodies of the lateral ventricles may affect the results. Strict anatomical criteria should be established in order to allow for geometrical definitions of boundaries between structures to be made. Finally, it should be reminded that volume is not a powerful shape descriptor. Temporal horns with similar volumes may have very different shapes, and this is crucial for the study of symmetry.

Clearly, similar neurological studies based on MRI could benefit from the use of the tools here presented. In particular, the study of symmetry differences in schizophrenia is not limited to the temporal horns. Differences in other structures have been reported, and it would be interesting to investigate them. MS lesions, as well, could be segmented and measured with these techniques¹⁸. It should be obvious, however, that most of the tools developed in this thesis have a broader application, not limited to neurological studies from MRI. The use of such techniques to other application fields or other imaging modalities is foreseeable.

¹⁸ MS may present some problems, however. Particularly, boundaries between white matter and lesions may not be sharp, but blurred, since plaques are demyelinated white matter. Additionally, lesions may have texture, complicating the establishment of intensity models for them.

Conclusions and Future Work

6.1 Contributions of this thesis

The several tools presented in this thesis have been created with specific practical applications in mind. In particular, the need for morphometric tools to analyse neurological data contained in MRI volumes was identified, and Chapters 1 and 2 provided clinical and technical background in the matter.

A framework for morphometric studies was described in Chapter 3. Although some individual components of the method are *off-the-shelf* existing techniques, such as the EM algorithm for bias field correction and the simplex mesh for 3D segmentation, there were a number of novel contributions. A new shape model for modelling complex biological surfaces based on the combination of the simplex mesh and triangular Gregory-Bézier patches was established, showing good flexibility and modelling power even in extremely complex examples such as the cortex (section 3.11). A new feature for detecting boundaries between tissues in MRI, based in the identification of PVE voxels from the probability maps obtained via the EM algorithm was introduced in section 3.3. Additionally, the combination of statistical voxel-based segmentation methods and geometric modelling via active shape models is interesting, and the work presented in Chapter 3 is a contribution to the scant literature on the topic. However, we consider that the most important element of our morphometric framework is the estimation of confidence intervals by bounding the segmentation using an inner and outer surface encompassing the detected PVE voxels. This uncertainty area is used to bound volume measurements in the extensive validation study of section 3.10, which also provides some valuable insight into the problem of lack of ground truth for the validation of these sort of techniques, and proposes some solutions.

The volume contained between the inner and outer surface bounds is not insignificant. In fact, even with current high resolution protocols, not less than 20% of the volume of an object is contained in PVE voxels, and this number increases to 60% or more for structures with high surface area relative to the volume enclosed, such as the ventricles. The PVE prediction scheme developed in Chapter 4 allows for a considerable reduction of this uncertainty. This method is novel not only in its algorithms and mathematical formulation, but also in its rigorous statistical approach to modelling PVE. The use of such a framework allows for a substantial reduction in measurement uncertainty derived from MRI data sets.

The clinical study reported in Chapter 5 describes the application of our techniques to a challenging clinical investigation: the assessment of differences in the asymmetry of the temporal horns in schizophrenia. This study proves that the original aims of this thesis were realised, by constructing new morphometric methods with sufficient accuracy and resolution to assist in neurological studies and help to prove hypotheses relating to schizophrenia and multiple sclerosis. This study is in fact an important step in this direction, and therefore a big clinical contribution.

6.2 *Future work*

6.2.1 *Improved tissue models*

The key idea behind the improvements reported in this thesis in order to obtain sub-voxel accuracy is that of using prior knowledge about the intensity distributions of different tissue types in particular acquisition protocols. The analysis of image samples allows for the construction of rudimentary models, but a rigorous physics-based derivation of tissue models is desirable. In essence, it should be possible to determine the shape of the intensity distribution for a given tissue type knowing its chemical properties and the acquisition times of the protocol used for imaging the sample.

The introduction of new tissue models in the PVE framework poses no problem. The current implementation uses such priors to create a mixture model and then to sample α^i given the

intensity I^i . Except for the case of Gaussian distributions, where analytical expressions have been derived, the usual approach is to build the mixture model by means of a Monte Carlo simulation.

6.2.2 *Priors for PVE prediction*

The model adopted for the prior on the proportion of each tissue in a PVE voxel, $p(\alpha^i)$, consisting of a uniform distribution in the range $[0, 1]$, is correct subject to a set of assumptions. As discussed in section 4.2, only PVE voxels must be included in the samples. Otherwise, small deviations from 0 or 1 in α^i for pure voxels accumulate and affect significantly the results. It was suggested in section 4.2 that it would be possible to include pure voxels if $p(\alpha^i)$ were modelled as a distribution with “tails”, pushing values close to 0 or 1 to absolute 0 or 1. However, the shape of these tails does matter. Tails which are too wide affect PVE voxels, whereas too narrow tails do not have the desired corrective effect. A set of priors with the size of tails could be built by studying synthetic images. The shape of the distribution is a function of the number of pure voxels with respect to the number of PVE voxels. It should therefore be possible to establish a model relating the width of the tails of $p(\alpha^i)$ to the ratio (surface area)/volume.

Another crucial assumption relating to $p(\alpha^i)$ is that of independence. It was argued in Chapter 4 that the range of surfaces traversing voxels is virtually infinite, so that no reliable model can be built. There are, however, a set of simplifications that can be made in order to establish neighbour relations in α^i . If the surface passing through voxel i is assumed to be smooth at a scale greater than the voxel size, it can be argued that values of α^i should bear some similarity to the α values at neighbouring PVE voxels. A Markov random field could then be employed to model such constraints. It should be noted, however, that the assumption of small variation on α values across neighbours does not hold for very small structures, such as the temporal horns, analysed in Chapter 5.

6.2.3 *Use of anatomical knowledge*

It cannot be stressed too often that the way forward in medical computer vision is to make use of anatomical (and physiological) knowledge. Computer vision is a new research field, maturing at a

fast pace but still in its early stages of development; medicine is not. Anatomical knowledge is available, and common clinical applications focus on localised regions of the human body. This makes feasible the construction of anatomical priors.

For example, the previous section can be conceived as an attempt to model the piece-wise continuous nature of brain tissues. Several approaches can be taken to model anatomical knowledge, including deformable atlases (Christensen et al., 1993; Subsol et al., 1995; Lester and Arridge, 1999) or heuristics such as the ones described in section 4.9. Modelling and application of anatomical knowledge to current techniques is a challenging research, and no doubt a difficult task. However, the expected improvement in performance fully justifies the effort.

6.2.4 *Surface fitting based on PVE constraints*

As part of the segmentation framework described in Chapter 3, PVE voxels are detected from the probability maps resulting from the EM algorithm; then inner and outer surfaces are constructed bounding these PVE voxels. Clearly, this is an extremely conservative approach. Chapter 4 dealt with PVE voxels and established a statistical framework based on their intensities and the intensity distributions of tissues in order to estimate the proportion of each tissue in them. It seems obvious that the next step in this line of research is to extend the segmentation framework in order to obtain truly sub-voxel fitting of the surface model.

An important contribution towards this end was offered in section 4.12. The technique we described allows for a rendering of a surface model onto a voxel grid. Thus, surface-based PVE estimates, α_s , are computed as the proportion of the volume of a voxel to each side of the surface. This, in addition to the voxel-based PVE estimation framework described in the first sections of Chapter 4, provides the basic tools to create a sub-voxel fitting framework. In particular, additional force constraints could be added to the active surface framework as to minimise the difference between α_s^i and intensity-based estimate α^i for every voxel i traversed by the surface:

$$\sum_i \|\alpha^i - \alpha_s^i\|^2 \quad (6.1)$$

In such framework, α_s^i values depend on the parameters of the surface model, which are the variables to be updated in order to minimise (6.1). Effectively, this updates the position of the surface as to fit to the image intensities, i.e. fits the surface based on PVE constraints.

6.2.5 *Developing the concept of “statistical shape”*

A further extension, of a more global character, is conceivable. The shape model mentioned in the previous section may represent just one discrete instance of a statistical continuum, defined by the likelihood of the surface given the voxel image. Clearly, a more powerful shape model is required in order to capture this statistical nature. Little work has been done in the field of statistical shape (Pennec, 1996; Dryden and Mardia, 1997). The desired properties of such a model are as follows:

- **Local modelling of uncertainty.** A parameterisation like the one employed for our model establishes a one to one correspondence between the parameter space and the three-dimensional space in which the surface (or volume) is embedded. This model is not appropriate for the application in mind. However, an extension that allows for a mapping of a point in parametric space into a statistical distribution of spatial locations is possible (see Figure 54). Alternatively, implicit representations such as level sets provide a natural representation of these processes. In this case, the usual approach consists of representing a surface as the kernel of a function defined in 3D space. Such implicit functions could be modified so that the value of the function at a spatial location indicates the probability of the surface traversing that point. The most likely surface would then be defined as the loci of the local maxima of such function (Figure 55).
- **Neighbour-level interaction.** Local geometric properties such as continuity in curvature or tangent plane (G^1 continuity) act as constraints on the location of the surface. Independence across neighbours cannot be assumed. This further complicates the formulation of a statistical shape model. A possible solution is to include such constraints in the implicit probability function mentioned above. Markov random fields could be employed to such effect.




Figure 54 Statistical shape modelling via parametric surfaces. Each point in parametric space corresponds to a statistical distribution of possible 3D locations, as opposed to a single 3D point, which is the common approach.




Figure 55 Implicit representations such as level sets allow for a natural representation of uncertainty if the implicit function f quantifies the probability of the surface traversing a particular point in 3D space.

- **Propagation of local uncertainty to global magnitudes.** Mechanisms for building global shape descriptors from the surface definition are essential for morphometric applications. The statistical nature of local position should be propagated to such descriptors. A model in the lines of the probability density functions obtained via Monte Carlo simulations in section 4.5.2 seems particularly suited.

This is a very challenging and interesting issue, with an endless range of applications in different fields. The establishment of an appropriate shape model with the characteristics sketched above could revolutionise the way shape modelling is conceived of, and allow for a more realistic representation of reality and a better integration of statistics and exact sciences.

6.2.6 *Anatomy-based global parameterisation*

Symmetry appears to be the key issue in schizophrenia. Its study is clearly a comparative process. Symmetry must be defined with respect to a plane or surface. A “mirror” of the structure on one side of such surface may be obtained and compared with the structure at the other side. If we consider other comparative processes, for example the comparison of the shape of an organ across patients, or the time evolution in a single patient, a similar line of action (without mirroring) should be taken. Therefore, we consider the more general case of comparing two “homologous” structures.

In order to compare two structures, it is necessary to identify corresponding regions. The usual approach is to perform a registration and then study the deformation field in order to obtain shape descriptors of the difference in shape. This technique is used in methods based on deformable atlases (Christensen et al., 1993; Lester and Arridge, 1999). The process of identifying corresponding structures would be greatly simplified if the surfaces being registered contained in-built anatomical information.

A re-parameterisation of the surface may be applied so anatomically meaningful locations on both structures correspond to the same values in their corresponding parametric spaces. Thus, establishing correspondences is conceived as an intrinsic process, rather than an extrinsic point-to-point matching problem. Appendix D offers a more detailed explanation. This technique bears similarities with the natural way of human recognition based on recognising landmarks. In our case, we employ this technique to establish correspondences between surfaces (or volumes, adding an extra parameter to map the extra dimension). The method was implemented and tested in some simple examples. A more detailed description and algorithmic details are offered in Appendix D. Note that this method relies on the prior localisation of useful landmarks on both

surfaces. Although this is not difficult to establish in structures like the cortex, other organs have a rather featureless shape, which complicates the definition of reference landmarks.

6.2.7 *New shape descriptors for morphometric studies*

To my knowledge, not much work has been done to date in establishing morphometric descriptors for particular applications. Particularly, most studies relating schizophrenia and symmetry have been performed using standard measures such as volume, area in a slice, or length of a curve across the structure of interest. Although it is not fair to say that these traditional shape descriptors are of little use, it should be clear that they do not portray complex concepts like symmetry or “global shape”.

The previous section provided some insight into the problem of establishing correspondences, and proposed an anatomy-based parameterisation. This scheme would be a great aid for shape description, since quantitative measurements would relate directly to anatomically meaningful concepts. The subsequent definition of particular shape descriptors is a very interesting and complex problem, and this line of research should be followed in the future.

6.2.8 *Study of sampling limitations*

The work presented in this thesis highlights the importance of the resolution of the sampling grid used for imaging, and attempts to overcome some of its limitations by building a model of PVE in order to obtain sub-voxel accuracy. It would be convenient, however, to have a formal model of the accuracy we expect to obtain as a function of the voxel size. Alternatively, the inverse question could be asked, i.e. which voxel size would be necessary to obtain accurate quantitative measurements (with a confidence interval $\pm\epsilon$) of a structure with certain shape characteristics. A combination of the Nyquist theorem and the PVE model based on tissue properties could provide an answer. The construction of such sampling model may be performed in a theoretical manner or by experimental simulation of synthetic phantoms.

This is an interesting issue with important practical consequences. It would provide radiologists with a solid base to determine what sort of structures they can reliably measure with a certain hardware. Additionally, it would help on the design of optimised MRI machines.

6.3 *Concluding remarks*

The field of medical computer vision enjoys an impressive momentum due to its evident social, economic and human impact. The subject is rapidly maturing and new techniques for processing medical images are being continuously reported. Additionally, new imaging modalities such as functional MRI or elastography allow for the study of function in addition to anatomy, thus opening new and exciting possibilities of research.

In parallel, medical robotics is progressing at a formidable pace, and a number of robots are being used routinely in surgery at several institutions. Pre-operative planning tools and intra-operative image acquisition devices allow for a better understanding of the patient's anatomy and reduce considerably risks during surgery. Additionally, some techniques have been devised to display "enhanced reality" models showing the pre-operative plan overlaid on top of images acquired in real time during surgery.

Twenty years ago, medical vision and medical robotics could only be conceived as an exercise of science-fiction. Today, many of these utopian ideas are being realised in laboratories across the world. These tools will be ported to hospitals in the very near future, and they will produce a breathtaking revolution in clinical practices. Prepare for it. It is going to be exciting, and we all will be part of it.

Triangular Gregory-Bézier patches

A.1 Definition

Triangular patch models are used to build a smooth surface interpolating a triangulations. In this section we will focus on how to build a G^1 -continuous surface, i.e. a surface ensuring continuity of the tangent plane across patches (inside each patch, the surface will be C^1 -continuous). Most commonly, constraining continuity across the common boundary of two patches is quite straightforward, whereas establishing such constraints at the locations of the vertices of the triangles, where several patches meet, is quite complicated. Triangular Gregory-Bézier (tGB) patches avoid such problems, since G^1 -continuity may be enforced by solely posing edge (common boundary) constraints.

TGB patches are a degenerate case of quartic triangular Bernstein-Bézier (tBB) patches, derived from them by: i) degenerating the three common boundaries by one degree, and ii) replacing each of the three inner control points by a rational combination of two other points. Thus, tGB patches are cubic rational surfaces defined by 15 control points (Figure 56). Their expression, in barycentric coordinates, is as follows:

$$\begin{aligned} \mathbf{GB}(u, v, w) = & u^3 \mathbf{P}_0 + v^3 \mathbf{P}_1 + w^3 \mathbf{P}_2 + 12u^2 v w \mathbf{P}_{211} + 12u v^2 w \mathbf{P}_{121} + 12u v w^2 \mathbf{P}_{112} + \\ & 3u^2 v(1-w) \mathbf{P}_{01} + 3u v^2(1-w) \mathbf{P}_{02} + 3v^2(1-u) w \mathbf{P}_{11} + \\ & 3(1-u) v w^2 \mathbf{P}_{12} + 3u(1-v) w^2 \mathbf{P}_{21} + 3u^2(1-v) w \mathbf{P}_{22} \end{aligned} \quad (\text{A.1})$$

where $0 \leq u, v, w \leq 1$, $u + v + w = 1$ and:

$$\mathbf{P}_{211} = \frac{w \mathbf{P}_{211}^v + v \mathbf{P}_{211}^w}{w + v}, \quad \mathbf{P}_{121} = \frac{u \mathbf{P}_{121}^w + w \mathbf{P}_{121}^u}{u + w}, \quad \mathbf{P}_{112} = \frac{v \mathbf{P}_{112}^u + u \mathbf{P}_{112}^v}{v + u}$$




Figure 56 Triangular Gregory-Bézier (tGB) patch, defined by 15 control points.

A.2 G^1 -continuity constraints

To obtain the G^1 continuity connection between two adjacent patches, we need to express the first order derivative of the patch $\mathbf{GB}(u, v, w)$ along the common boundary, say $\Gamma(v) = \mathbf{GB}(0, v, 1-v)$, in the *radial direction* which connect any point of $\Gamma(v)$ to its opposite vertex. With this radial direction $\theta_u = (1, 0, 0) - (0, v, 1-v) = (1, -v, v-1)$, the directional derivative takes the form:

$$\left. \frac{\partial \mathbf{GB}(u, v, w)}{\partial \theta_u} \right|_{u=0} = 4 \left(\frac{3}{4} (1-v)^3 (\mathbf{P}_{21} - \mathbf{P}_2) + 3(1-v)^2 v (\mathbf{P}_{112}^u - \mathbf{P}_{12}) + 3(1-v) v^2 (\mathbf{P}_{121}^u - \mathbf{P}_{11}) + \frac{3}{4} v^3 (\mathbf{P}_{02} - \mathbf{P}_1) \right) \quad (\text{A.2.1})$$

Similarly, the radial directional derivative vectors along the two other boundaries are as follows:

$$\left. \frac{\partial \mathbf{GB}(u, v, w)}{\partial \theta_v} \right|_{v=0} = 4 \left(\frac{3}{4} (1-w)^3 (\mathbf{P}_{01} - \mathbf{P}_0) + 3(1-w)^2 w (\mathbf{P}_{211}^v - \mathbf{P}_{22}) + 3(1-w) w^2 (\mathbf{P}_{112}^v - \mathbf{P}_{21}) + \frac{3}{4} w^3 (\mathbf{P}_{12} - \mathbf{P}_2) \right) \quad (\text{A.2.2})$$

$$\left. \frac{\partial \mathbf{GB}(u, v, w)}{\partial \theta_w} \right|_{w=0} = 4 \left(\frac{3}{4} (1-u)^3 (\mathbf{P}_{11} - \mathbf{P}_1) + 3(1-u)^2 u (\mathbf{P}_{121}^w - \mathbf{P}_{02}) + 3(1-u) u^2 (\mathbf{P}_{211}^w - \mathbf{P}_{01}) + \frac{3}{4} u^3 (\mathbf{P}_{22} - \mathbf{P}_0) \right) \quad (\text{A.2.3})$$

It should be noticed that the expressions for these first order derivatives do not share any inner control point. In practice, this means that G^1 continuity can be guaranteed simply by constraining the connection along the boundary between two adjacent patches, avoiding the tedious process of considering continuity at corners. G^1 continuity across patches is ensured by enforcing a coplanarity constraint between the two radial first derivative vectors

$$\left. \frac{\partial \mathbf{GB}_R(u_R, v_R, w_R)}{\partial \theta_{u_R}} \right|_{u_R=0} \quad \text{and} \quad \left. \frac{\partial \mathbf{GB}_L(u_L, v_L, w_L)}{\partial \theta_{u_L}} \right|_{u_L=0}$$

(where θ_{u_R} and θ_{u_L} take values in $\{u, v, w\}$, depending on the orientation of the patch), and the first order derivative vector $\mathbf{\Gamma}^{(1)}(v)$ of the common boundary (see Figure 57).




Figure 57 Enforcing G^1 continuity in the boundary of neighbouring tGB patches.

This coplanarity condition can be expressed as follows:

$$\alpha(v) \left. \frac{\partial \mathbf{GB}_R(u_R, v_R, w_R)}{\partial \theta_{u_R}} \right|_{u_R=0} + \beta(v) \left. \frac{\partial \mathbf{GB}_L(u_L, v_L, w_L)}{\partial \theta_{u_L}} \right|_{u_L=0} + \gamma(v) \mathbf{\Gamma}^{(1)}(v) = 0, \quad 0 \leq v \leq 1 \quad (\text{A.3})$$

where $\alpha(v)$, $\beta(v)$, and $\gamma(v)$ are polynomial functions of v . $\beta(v)$ and $\gamma(v)$ are chosen as linear functions, $\alpha(v)$ being chosen as constant. A discussion of these choices is provided in (Schmitt et al., 1991). Therefore:

$$\alpha(v) = 1, \quad \beta(v) = \beta_0(1-v) + \beta_1 v, \quad \gamma(v) = \gamma_0(1-v) - \gamma_1 v$$

To simplify, we use the following new notation for the control points of the two patches which appear in the coplanarity condition (see Figure 58):

$$\begin{aligned} \mathbf{P}_1^R, \mathbf{P}_{11}^R, \mathbf{P}_{12}^R, \mathbf{P}_2^R &\Rightarrow \mathbf{C}_i & \mathbf{P}_2^L, \mathbf{P}_{12}^L, \mathbf{P}_{11}^L, \mathbf{P}_1^L &\Rightarrow \mathbf{C}_i \\ \mathbf{P}_{02}^R, \mathbf{P}_{121}^R, \mathbf{P}_{112}^R, \mathbf{P}_{21}^R &\Rightarrow \mathbf{P}_i^R & \mathbf{P}_{21}^L, \mathbf{P}_{112}^L, \mathbf{P}_{121}^L, \mathbf{P}_{02}^L &\Rightarrow \mathbf{P}_i^L \end{aligned} \quad i = 0, 1, 2, 3$$




Figure 58 Notation change for the control points of the two patches \mathbf{GB}_L and \mathbf{GB}_R .

Now, by replacing functions $\alpha(v)$, $\beta(v)$ and $\gamma(v)$ in the coplanarity condition and developing it on the Bernstein basis of degree 4, we obtain the following set of G^1 constraints on the control points:

$$(\mathbf{P}_0^R - \mathbf{C}_0) + \beta_0(\mathbf{P}_0^L - \mathbf{C}_0) + \gamma_0(\mathbf{C}_1 - \mathbf{C}_0) = 0 \quad (\text{A.4.1})$$

$$(\mathbf{P}_1^R - \mathbf{C}_1) + \frac{1}{4}(\mathbf{P}_0^R - \mathbf{C}_0) + \beta_0(\mathbf{P}_1^L - \mathbf{C}_1) + \frac{1}{4}\beta_1(\mathbf{P}_0^L - \mathbf{C}_0) + \frac{1}{2}\gamma_0(\mathbf{C}_2 - \mathbf{C}_1) + \frac{1}{4}(\gamma_0 - \gamma_1)(\mathbf{C}_1 - \mathbf{C}_0) = 0 \quad (\text{A.4.2})$$

$$(\mathbf{P}_2^R - \mathbf{C}_2) + (\mathbf{P}_1^R - \mathbf{C}_1) + \beta_0(\mathbf{P}_2^L - \mathbf{C}_2) + \beta_1(\mathbf{P}_1^L - \mathbf{C}_1) + \frac{1}{4}\gamma_0(\mathbf{C}_3 - \mathbf{C}_2) + \frac{1}{2}(\gamma_0 - \gamma_1)(\mathbf{C}_2 - \mathbf{C}_1) + \frac{1}{4}\gamma_1(\mathbf{C}_0 - \mathbf{C}_1) = 0 \quad (\text{A.4.3})$$

$$(\mathbf{P}_2^R - \mathbf{C}_2) + \frac{1}{4}(\mathbf{P}_3^R - \mathbf{C}_3) + \beta_1(\mathbf{P}_2^L - \mathbf{C}_2) + \frac{1}{4}\beta_0(\mathbf{P}_3^L - \mathbf{C}_3) + \frac{1}{2}\gamma_1(\mathbf{C}_1 - \mathbf{C}_2) + \frac{1}{4}(\gamma_1 - \gamma_0)(\mathbf{C}_2 - \mathbf{C}_3) = 0 \quad (\text{A.4.4})$$

$$(\mathbf{P}_3^R - \mathbf{C}_3) + \beta_1(\mathbf{P}_3^L - \mathbf{C}_3) + \gamma_1(\mathbf{C}_2 - \mathbf{C}_3) = 0 \quad (\text{A.4.5})$$

Note that constraint (A.4.1) and (A.4.5) involve only the control points situated along the boundaries of the two patches, while the two couples of inner control points $(\mathbf{P}_1^L, \mathbf{P}_1^R)$ and $(\mathbf{P}_2^L, \mathbf{P}_2^R)$ appear only in the three remaining constraints. Hence, we determine the control points in two successive steps: first, constraints (A.4.1) and (A.4.5) are used to define the control points situated along the boundaries, then the other three constraints are used to define the two couples of inner control points.

A.3 Determination of control points situated along the boundaries

In order to constrain the many degrees of freedom of equation (A.4), constraints on the normal vector the resulting surface must have at a given vertex are imposed. In the framework presented

in Chapter 3 such constraints were naturally derived from the definition of normal vector at a simplex mesh vertex. (Schmitt et al., 1991) propose an alternative method of establishing normal vectors at the vertices of the triangulation by computing the direction of minimum moment of inertia on the set of neighbours to a given vertex.

Let us consider an edge of the 3D triangular mesh with their two 3D points \mathbf{C}_0 and \mathbf{C}_3 . Let us call \mathbf{n}_0 and \mathbf{n}_1 their normal vectors estimated as described above. We want to determine the supplementary control points \mathbf{C}_1 and \mathbf{C}_2 in order to define a 3D cubic Bézier curve which could be a “good” patch boundary candidate corresponding to this edge. Let us call \mathbf{a} and \mathbf{b} the unit tangent vectors of the Bézier curve in \mathbf{C}_0 and \mathbf{C}_3 respectively. The control points \mathbf{C}_1 and \mathbf{C}_2 can then be parameterised as follows (see Figure 59a):

$$\mathbf{C}_1 = \mathbf{C}_0 + \lambda_1 \mathbf{a}, \quad \mathbf{C}_2 = \mathbf{C}_3 + \lambda_2 \mathbf{b} \quad (\text{A.5})$$

To completely specify \mathbf{C}_1 and \mathbf{C}_2 , we first determine the unit tangent vectors \mathbf{a} and \mathbf{b} and then the value of the two parameters λ_1 and λ_2 . To build \mathbf{a} , we first obtain a point \mathbf{M} by intersecting the tangent plane defined at \mathbf{C}_0 with the projection ray passing through the point \mathbf{C}_3 . Then \mathbf{a} is obtained by normalising the vector $\overrightarrow{\mathbf{C}_0\mathbf{M}}$ (see Figure 59).




Figure 59 Determination of the boundary control points.

We construct \mathbf{b} in a similar way. To obtain a very smooth boundary we then define λ_1 and λ_2 by minimising the following integral:

$$\int_0^1 \left(\frac{\partial^2 \mathbf{B}(v)}{\partial v^2} \right)^2 dv$$

where $\mathbf{B}(v)$ is the cubic Bézier curve defined by the control points $\mathbf{C}_i, i=0, \dots, 3$, with \mathbf{C}_1 and \mathbf{C}_2 parameterised as in (A.5). After some calculations, we obtain the following expressions for λ_1 and λ_2 :

$$\lambda_1 = \frac{2\mathbf{a} \cdot \mathbf{c} - (\mathbf{a} \cdot \mathbf{b})(\mathbf{b} \cdot \mathbf{c})}{4 - (\mathbf{a} \cdot \mathbf{b})^2} \quad \lambda_2 = \frac{-2\mathbf{b} \cdot \mathbf{c} + (\mathbf{a} \cdot \mathbf{b})(\mathbf{a} \cdot \mathbf{c})}{4 - (\mathbf{a} \cdot \mathbf{b})^2}$$

where $\mathbf{c} = \mathbf{C}_3 - \mathbf{C}_0$.

The points $\mathbf{P}_0^L, \mathbf{P}_0^R, \mathbf{P}_3^L$ and \mathbf{P}_3^R can be defined with the same scheme performed at the adjacent boundaries (see Figure 58). The constants $\beta_0, \beta_1, \gamma_0$ and γ_1 are then entirely determined by the constraints (A.4.1) and (A.4.5) as follows:

$$\beta_0 = \frac{[(\mathbf{P}_0^R - \mathbf{C}_0) \times (\mathbf{C}_1 - \mathbf{C}_0)] \cdot \mathbf{n}_0}{[(\mathbf{C}_1 - \mathbf{C}_0) \times (\mathbf{P}_0^L - \mathbf{C}_0)] \cdot \mathbf{n}_0} \quad \gamma_0 = \frac{[(\mathbf{P}_0^R - \mathbf{C}_0) \times (\mathbf{P}_0^L - \mathbf{C}_0)] \cdot \mathbf{n}_0}{[(\mathbf{P}_0^L - \mathbf{C}_0) \times (\mathbf{C}_1 - \mathbf{C}_0)] \cdot \mathbf{n}_0}$$

$$\beta_1 = \frac{[(\mathbf{P}_3^R - \mathbf{C}_3) \times (\mathbf{C}_2 - \mathbf{C}_3)] \cdot \mathbf{n}_3}{[(\mathbf{C}_2 - \mathbf{C}_3) \times (\mathbf{P}_3^L - \mathbf{C}_3)] \cdot \mathbf{n}_3} \quad \gamma_1 = \frac{[(\mathbf{P}_3^R - \mathbf{C}_3) \times (\mathbf{P}_3^L - \mathbf{C}_3)] \cdot \mathbf{n}_3}{[(\mathbf{P}_3^L - \mathbf{C}_3) \times (\mathbf{C}_2 - \mathbf{C}_3)] \cdot \mathbf{n}_3}$$

The determination of the control points along the boundaries according to constraints (A.4.1) and (A.4.5) is thus completed.

A.4 Determination of the inner control points

By rewriting the constraints (A.4.2), (A.4.3), and (A.4.4) in matrix product form, we have:

$$\begin{pmatrix} 3/4 & 3\beta_0/4 & 0 & 0 \\ 1/2 & \beta_1/2 & 1/2 & \beta_0/2 \\ 0 & 0 & 3/4 & 3\beta_1/4 \end{pmatrix} \begin{pmatrix} \mathbf{P}_1^R \\ \mathbf{P}_1^L \\ \mathbf{P}_2^R \\ \mathbf{P}_2^L \end{pmatrix} = \begin{pmatrix} \mathbf{X}_1 \\ \mathbf{X}_2 \\ \mathbf{X}_3 \end{pmatrix} \quad (\text{A.6})$$

where each $\mathbf{X}_i, i=1,2,3$ is a linear combination of the already fixed control points which can be easily deduced from these three constraints. There are two cases to consider:

1. When $\beta_0 \neq \beta_1$, the rank of the 3×4 matrix is 3. We can therefore freely fix one of the 4 points (3 degrees of freedom) and find a unique solution to the 3 other points by constraints (A.6).
2. When $\beta_0 = \beta_1$, system (A.6) is reduced to only 2 independent constraints: the first and third ones. In this case, each pair of inner control points $(\mathbf{P}_i^L, \mathbf{P}_i^R), i=1,2$, has one constraint to

satisfy. We can thus freely fix one point of each couple and deduce the other point by using the corresponding constraint equation.

We present here how the two couples of points are determined in case 1. Case 2 can be treated in a similar way. In order to minimise undesired fluctuations, the approach taken is to compute control points that approximate the shape of a simpler, smoother Bernstein Bézier patch of only one cubic degree. This kind of patch has only one inner control point, which we determine as a barycentric combination of the boundary control points (see Figure 60):

$$\mathbf{CP} = \frac{1}{12}(\mathbf{P}_{012} + \mathbf{P}_{021} + \mathbf{P}_{201} + \mathbf{P}_{102} + \mathbf{P}_{120} + \mathbf{P}_{210}) + \frac{1}{6}(\mathbf{P}_{300} + \mathbf{P}_{030} + \mathbf{P}_{003})$$




Figure 60 Determination of initial inner control points.

Then the equivalent quartic tBB patch is obtained by using a degree elevation technique. Its three inner control points are obtained by the following expressions:

$$\mathbf{I}_1 = \frac{1}{4}(\mathbf{P}_{201} + \mathbf{P}_{210} + 2\mathbf{CP}), \quad \mathbf{I}_2 = \frac{1}{4}(\mathbf{P}_{021} + \mathbf{P}_{120} + 2\mathbf{CP}), \quad \mathbf{I}_3 = \frac{1}{4}(\mathbf{P}_{012} + \mathbf{P}_{102} + 2\mathbf{CP})$$

We define the two couples of inner control points $(\mathbf{P}_i^L, \mathbf{P}_i^R), i=1,2$ along a common boundary by a minimisation method. The goal is to locate these control points as close as possible to their corresponding (closest) tBB control points, $(\mathbf{I}_i^L, \mathbf{I}_i^R), i=1,2$, while enforcing the constraints (A.6). If we consider case 1 ($\beta_0 \neq \beta_1$), one of the four inner control points can be fixed freely. Let

it be \mathbf{P}_1^L . By rewriting (A.4.2), (A.4.3), and (A.4.4), the remaining points are expressed as a function of \mathbf{P}_1^L as follows:

$$\mathbf{P}_1^R = -\beta_0 \mathbf{P}_1^L + \Delta_1, \quad \mathbf{P}_2^L = \mathbf{P}_1^L + \Delta_2, \quad \mathbf{P}_2^R = -\beta_1 \mathbf{P}_1^L + \Delta_3, \quad (\text{A.7})$$

where

$$\Delta_1 = (1 + \beta_0) \mathbf{C}_1 - \frac{1}{4} (\mathbf{P}_0^R - \mathbf{C}_0) - \frac{1}{4} \beta_1 (\mathbf{P}_0^L - \mathbf{C}_0) - \frac{1}{2} \gamma_0 (\mathbf{C}_2 - \mathbf{C}_1) - \frac{1}{4} (\gamma_0 - \gamma_1) (\mathbf{C}_1 - \mathbf{C}_0),$$

$$\Delta_2 = (\mathbf{C}_2 - \mathbf{C}_1) - \frac{1}{4} (\mathbf{P}_0^L - \mathbf{C}_0) + \frac{1}{4} (\mathbf{P}_3^L - \mathbf{C}_3),$$

$$\Delta_3 = (\mathbf{C}_2 + \beta_1 \mathbf{C}_1) + \frac{1}{2} \gamma_1 (\mathbf{C}_2 - \mathbf{C}_1) + \frac{1}{4} \gamma_0 (\mathbf{C}_2 - \mathbf{C}_3) + \frac{1}{4} \beta_1 (\mathbf{P}_0^L - \mathbf{C}_0) - \frac{1}{4} \beta_0 (\mathbf{P}_3^L - \mathbf{C}_3),$$

are vectors already known.

We then proceed to minimise distances to the computed tBB points. The function to be minimised is:

$$\sum_{i=1}^2 \left[\left| \mathbf{P}_1^L - \mathbf{I}_1^L \right|^2 + \left| \mathbf{P}_1^R - \mathbf{I}_1^R \right|^2 \right]$$

By replacing (A.7) into this equation and differentiating it with respect to \mathbf{P}_1^L we obtain:

$$\mathbf{P}_1^L = \frac{\mathbf{I}_1^L + \mathbf{I}_2^L - \Delta_2 + \beta_0 (\Delta_1 - \mathbf{I}_1^R) + \beta_1 (\Delta_3 - \mathbf{I}_2^R)}{2 + \beta_0^2 + \beta_1^2}$$

\mathbf{P}_1^R , \mathbf{P}_2^L , and \mathbf{P}_2^R can then easily be obtained from (A.7).

Algorithms

B.1 Oct-tree recursive subdivision algorithm

This algorithm is used in section 3.9.1 to create synthetic data sets. Such data sets are used in section 3.10.1 as ground truth for the validation study.

```

Function OCT-SUBDIVIDE (v)
  If (volume of v < threshold_vol) Then
    volume_inside = volume_inside + volume of v / 2
    volume_outside = volume_outside + volume of v / 2
  Return
  Else If (all corners of v are inside) Then
    volume_inside = volume_inside + volume of v
  Return
  Else If (all corners of v are outside) Then
    volume_outside = volume_outside + volume of v
  Return
  Else
    Divide v into 8 cubes of equal
      volume (vi, i=1..8)
    For i=1 To 8 Do
      OCT-SUBDIVIDE (vi)
    End For
  End If
End Function

Algorithm OCT-TREE
  For n=1 To number of voxels Do
    OCT-SUBDIVIDE (voxel n)
  End For
End Algorithm

```

Figure 61 Oct-tree recursive subdivision algorithm in voxel space, used to generate MRI synthetic phantoms for tests in section 3.10.1. Refer to section 3.9.1 for details.

B.2 Quad-tree recursive subdivision algorithm

This is employed in section 4.9 to create PVE masks from inner and outer surface bounds. The algorithm “marks” voxels traversed by a surface defined as a parametric map.

```
Function QUAD-SUBDIVIDE (pr)
  r = 3D rectangle given by the 3D coordinates
    of the corners of pr
  If (area of r < threshold_area) Then
    Return
  Else If (all corners of r belong to the same
    voxel v) Then
    Mark v
  Else
    Divide pr into 4 rectangles of equal
      area (pri, i=1..4)
    For i=1 To 4 Do
      QUAD-SUBDIVIDE (pri)
    End For
  End If
End Function

Algorithm QUAD-TREE
  Establish a partition of the parametric space
   $[0,1] \times [0,1]$  into N rectangles of equal area
  For n=1 To N Do
    QUAD-SUBDIVIDE (parametric rectangle n)
  End For
End Algorithm
```

Figure 62 Algorithm for the quad-tree recursive subdivision of the parametric space of the surface. This algorithm is used to sample the surface in order to create a voxel mask indicating the voxels the surface passes through.

Clinical Study: Tissue Parameters

Data set	CSF mean	CSF std	GM mean	GM std	WM mean	WM std
006_SC	2.68	0.16	3.79	0.05	4.12	0.04
007_SC	2.17	0.38	3.33	0.10	3.64	0.07
008_SC	2.59	0.20	3.66	0.07	4.01	0.05
010_SC	2.83	0.20	3.69	0.09	4.02	0.08
024_SC	2.55	0.41	3.76	0.14	4.08	0.07
025_SC	2.67	0.32	3.76	0.12	4.09	0.07
033_SC	2.27	0.65	3.59	0.07	3.92	0.04
039_SC	2.53	0.48	3.79	0.10	4.15	0.07
104_NC	2.87	0.28	3.97	0.05	4.29	0.04
105_NC	2.52	0.15	3.61	0.08	3.90	0.05
106_NC	2.63	0.22	3.75	0.05	4.05	0.05
107_NC	2.42	0.19	3.42	0.06	3.74	0.06
112_NC	2.67	0.23	3.81	0.06	4.16	0.03
114_NC	2.68	0.19	3.68	0.09	4.07	0.04
117_NC	2.69	0.18	3.66	0.06	4.02	0.06
118_NC	2.52	0.30	3.60	0.06	3.99	0.03

Table 18 Tissue parameters as estimated by sampling intensities on a slice. These log-transformed values are used in the bias correction process.

Data set	CSF μ mean	CSF μ std.	CSF σ mean	CSF σ std.	GM μ mean	GM μ std.	GM σ mean	GM σ std.	WM μ mean	WM μ std.	WM σ mean	WM σ std.
006_SC	13.81	0.60	2.61	0.38	44.86	0.59	3.13	0.93	62.04	0.66	3.18	0.68
007_SC	8.63	0.57	2.17	0.97	27.32	0.54	3.01	0.33	38.69	0.30	3.21	0.55
008_SC	12.88	0.45	2.44	0.18	38.60	0.52	2.85	0.62	55.45	0.38	2.76	0.36
010_SC	16.08	0.66	3.85	0.71	38.88	1.21	3.80	0.94	54.70	1.87	4.59	1.41
024_SC	14.46	0.70	3.45	1.11	43.96	0.71	4.46	0.90	61.47	1.44	3.58	0.76
025_SC	13.72	1.17	3.74	0.85	43.99	0.85	4.25	0.43	62.75	1.45	4.11	0.41
033_SC	11.40	0.67	3.79	0.88	35.82	0.54	2.35	0.27	50.90	0.42	1.86	0.16
039_SC	13.82	0.79	5.50	0.47	45.69	1.05	3.95	0.48	64.86	1.09	4.79	0.51
104_NC	17.65	1.56	5.84	0.74	52.62	0.74	2.37	0.44	73.12	0.54	2.91	0.31
105_NC	12.68	0.21	2.00	0.15	36.42	0.42	2.71	0.43	49.82	0.46	2.26	0.16
106_NC	13.81	0.98	3.88	1.48	42.02	0.53	3.46	0.65	56.94	0.65	2.97	0.36
107_NC	11.58	0.25	1.85	0.15	30.62	0.31	2.29	0.21	42.20	0.28	2.41	0.24
112_NC	16.20	0.80	3.25	0.42	46.35	0.71	3.59	0.78	63.87	0.59	2.70	0.57
114_NC	14.24	0.41	2.64	0.22	40.52	0.45	2.98	0.46	57.78	0.58	2.28	0.23
117_NC	15.15	0.47	2.79	0.23	38.91	0.24	2.86	0.47	54.67	0.81	3.35	0.06
118_NC	12.52	0.70	3.38	0.50	38.02	1.15	3.07	0.52	54.15	0.43	2.05	0.22

Table 19 Tissue parameters including an explicit model of the uncertainty in their estimation from samples selected from images. Means (μ) and standard deviations (σ) for the tissue distributions are modelled as Gaussians. The means and standard deviations for these Gaussians are listed in the table.

Anatomy-Based Global Parameterisation

D.1 Introduction

The concept of establishing a global, anatomically meaningful surface parameterisation was introduced in section 6.2.6. As described previously in this thesis (section 4.10), any surface can be expressed as a parametric map of the form:

$$\mathbf{S}(u, v) = (x(u, v), y(u, v), z(u, v)) \quad u, v \in [0, 1] \quad (1)$$

Interestingly, this definition may be extended to parameterise a 3D volume by adding another parameter:

$$\mathbf{V}(u, v, w) = (x(u, v, w), y(u, v, w), z(u, v, w)) \quad u, v, w \in [0, 1] \quad (2)$$

Although the definition of tGB patches makes use of 3 parameters, one of them is expressed as a function of the other two. Thus, the tGB formula can be simplified to the form of (1). Let us focus on the case of a surface, and leave volumetric definitions in the form of (2) for later discussion.

The mapping \mathbf{S} defines the way in which correspondences between the parameter space $U = \{(u, v) \mid 0 \leq u, v \leq 1\}$ and the 3D space \mathfrak{R}^3 are established. Many mappings may be established for the same surface in \mathfrak{R}^3 . In the most common case, the particular details of the mapping are irrelevant to the application at hand, and the parameterisation is just used as a mathematical formalism and a way of indexing the points in the surface. However, an intelligent choice of parameterisation, taking into account the application in mind, may be highly rewarding. Let us illustrate this fact with the simple example of establishing a parameterisation for a circle. A possibility is to use cartesian coordinates. In this case, indexing of the points in the circle is cumbersome, and the definition of the domain of the parameters is complicated. On the other hand, a parameterisation based on the use of spherical coordinates provides a much more intuitive, simpler indexing of the surface.

It is the purpose of this appendix to define a parameterisation suitable for surfaces representing biological structures. In particular, we are interested in a mapping with the following characteristics:

- **Global parameterisation.** The surface model used in our work is based on the combination of triangular patches. A useful parameterisation should be transparent to this fact, i.e. global coordinates must naturally translate into local patch coordinates, and viceversa.
- **Anatomical interpretation.** The location of anatomical structures should be defined in the parametric space, as well as in physical 3D space. This means that a partition of the parametric space into regions which are mapped into anatomical structures in 3D must be possible.

The concept of globality in the parameterisation is important from a practical standpoint, and could be considered as an implementation issue. On the other hand, the definition of an anatomically meaningful parameterisation is a very powerful novel idea. A number of important advantages derive from establishing such a parameterisation:

- Easy indexing of anatomical structures. This allows for an intuitive manipulation of the shape model.
- Comparison of homologous structures is conceived as an intrinsic process with the aide of the parameter space, as opposed to an extrinsic deformation process in 3D. This helps circumventing problems derived from the considerable anatomical variability of certain structures, such as the brain cortex.
- This sort of model seems extremely suited for posterior shape description, since anatomy is in-built in it.

The following section describes how a global parameterisation is obtained from the mesh of tGB patches and establishes the algorithmic framework to translate global coordinates into local patch coordinates and viceversa. Next, the process of establishing a global parameterisation based on a set of anatomical landmarks is described. Finally, extensions to parameterise volumetric objects are discussed.

D.2 Correspondence between global parameters and local patch coordinates

Let us first consider the problem of establishing a mapping \mathbf{L} from global parameter values to local tGB patch coordinates. For the moment, let us assume that the global coordinates of the three vertices of the patch are known (the process of setting global coordinate values to the vertices will be described in the next section.). Let $\mathbf{P}_0 = (x_0, y_0)$, $\mathbf{P}_1 = (x_1, y_1)$, and $\mathbf{P}_2 = (x_2, y_2)$ be such vertices, defined in global coordinates. The goal is to determine the barycentric coordinates (u, v, w) of a point \mathbf{P} , given its global coordinates $\mathbf{P} = (x, y)$. Once a point in the triangular parametric domain of the tGB patch is obtained, the 3D point can be computed using the tGB formula. Thus, $\mathbf{tGB} \circ \mathbf{L}$ defines the correspondence between global coordinate values and 3D points.

The mapping \mathbf{L} is given by the following formulae:

$$\begin{aligned} u &= \frac{y(x_2 - x_1) + x(y_1 - y_2) - x_2 y_1 + x_1 y_2}{y_0(x_2 - x_1) + y_1(x_0 - x_2) + y_2(x_1 - x_0)} \\ v &= \frac{y(x_0 - x_2) + x(y_2 - y_0) - x_0 y_2 + x_2 y_0}{y_0(x_2 - x_1) + y_1(x_0 - x_2) + y_2(x_1 - x_0)} \\ w &= \frac{y(x_1 - x_0) + x(y_0 - y_1) - x_1 y_0 + x_0 y_1}{y_0(x_2 - x_1) + y_1(x_0 - x_2) + y_2(x_1 - x_0)} \end{aligned}$$

If \mathbf{P} is contained into the triangle defined by $(\mathbf{P}_0, \mathbf{P}_1, \mathbf{P}_2)$, then $u + v + w = 1$. This can be used as a test to determine which tGB patch contains \mathbf{P} .

The inverse process, i.e. finding the global parameters of a set of local tGB coordinates, can be performed using the barycentric expression:

$$\mathbf{P} = u\mathbf{P}_0 + v\mathbf{P}_1 + w\mathbf{P}_2$$

The discussion above assumes that the global coordinates at the vertices of the tGB mesh are known. The following section describes how these coordinates can be set and constraints on their values at particular locations enforced.

D.3 *Re-parameterisation based on anatomical landmarks*

We should bear in mind that the ultimate purpose of this section is to obtain an anatomically meaningful parameterisation. In order to enforce such anatomical meaning, landmarks are a key element. Such landmarks may be points, curves or surfaces, and we will be interested in setting the values of the global parameters at such locations. Whether these landmarks are obtained using an automatic or a manual method depends on the application and existing techniques. The automatic identification of anatomical landmarks is a challenging problem, out of the scope of this thesis. For our purposes, we took a manual approach and developed a program to “draw” landmarks in a triangulation (and associate parameter constraints to them).

Once constraints along certain landmarks have been established, we can simulate a diffusion process (a Dirichlet problem) to compute the values of the parameters in other locations of the surface. The method bears some similarity to the work of (Brechtbuhler et al., 1995). Discretising the diffusion equation (Boas, 1983) and taking into account the constraints at the landmarks, the following system is posed:

$$\begin{aligned} \mathbf{P} &= b_{\mathbf{P}} && \text{if } \mathbf{P} \text{ is a landmark} \\ \left(\sum_{i=1}^{n_{\mathbf{P}}} \frac{1}{d_i} \right) \mathbf{P} - \sum_{i=1}^{n_{\mathbf{P}}} \left(\frac{1}{d_i} \mathbf{N}(\mathbf{P})_i \right) &= 0 && \text{otherwise} \end{aligned}$$

where:

$\mathbf{N}(\mathbf{P})_i$ are neighbours of \mathbf{P} , i.e. vertices connected to \mathbf{P} in the triangulation

d_i is the distance from \mathbf{P} to $\mathbf{N}(\mathbf{P})_i$

$b_{\mathbf{P}}$ is the value enforced at \mathbf{P} , following the constraints at the landmarks

$n_{\mathbf{P}}$ is the number of neighbours of point \mathbf{P}

D.4 *Extension to volumetric objects*

The framework described in the previous section is not only valid for surfaces. Volumetric shape models could benefit from it. In fact, it is in volumetric applications where the concept of an anatomy-based parameterisation seems more convenient. For example, it could permit establishing intuitive *and* quantitative descriptions of the location or shape of a structure with

respect to other relevant organs. This would help developing appropriate shape descriptors for symmetry, and providing intuitive characterisation of the location of MS lesions.

Volumetric global coordinate systems could be defined as sets of three coordinates (x, y, z) . For example, in the case of the brain, x could indicate distance from the inter-hemispheric fissure, y could measure distance from the central sulcus, and z could be proportional to the distance from the corpus callosum. Additional constraints could be established at certain landmarks, such as the ventricles, sulci and internal nuclei.

Bibliography

- Amit, Y. Graphical shape templates for automatic anatomy detection with applications to MRI brain scans. *IEEE Transactions on Medical Imaging*, 16(1):28-40, 1997.
- Axel, L.; Costantini, J.; and Listerud, J. Intensity correction in surface-coil MR imaging. *AJR-Am-J-Roentgenol*, 148(2):418-20 issn: 0361-803x, 1987.
- Aylward, S.; Coggins, J.; Cizadlo, T.; and Andreasen, N. The effects of magnetic resonance image inhomogeneities on automated tissue classification. In procs. of *AAAI Spring Symposium on Applications of Computer Vision to Medical Image Processing*, pp. 207-210, 1994a.
- Aylward, S.; Coggins, J.; Cizadlo, T.; and Andreasen, N. Spatially invariant classification of tissues in MR images. In procs. of *VBC '94*, 1994b.
- Bardinet, E.; Cohen, L.D.; and Ayache, N. A parametric deformable model to fit unstructured 3D data, Research report 2617, INRIA, 1995.
- Barnhill, R.E. Smooth interpolation over triangles. In: Barnhill, and Reisenfeld, eds. *Computer Aided Geometric Design*, 1974.
- Besl, P.J., and McKay, N.D. A method for registration of 3D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239-256, 1992.
- Bilder, R.M.; Wu, H.; Bogerts, B.; Degreef, G.; Ashtari, M.; Alvir, J.M.J.; Snyder, P.J.; and Lieberman, J.A. Absence of regional hemispheric volume asymmetries in first-episode schizophrenia. *Am. J. Psychiatry*, 151:1437-1447, 1994.
- Blake, A.; Curwen, R.; and Zisserman, A. A framework for spatiotemporal control in the tracking of visual contours. *International Journal of Computer Vision*, 11(2):127-145, 1993.
- Blake, A., and Zisserman, A. *Visual Reconstruction*: MIT Press, 1987.
- Boas, M.L. *Mathematical Methods in the Physical Sciences*: Ed. John Wiley & Sons, 1983.
- Boehm, W., and Prautzsch, H. *Geometric Concepts for Geometric Design*: Ed. A. K. Peters, 1994.
- Bookstein, F.L. *Morphometric Tools for Landmark Data: Geometry and Biology*: Cambridge University Press, 1992.
- Bookstein, F.L. Shape and the information in medical images: a decade of the morphometric synthesis. In procs. of *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, pp. 2-12, 1996.
- Brady, M., and Asada, H. Smoothed local symmetries and their implementation, Technical report AI-757, MIT Artificial Intelligence Lab., 1984.
- Brady, M.; Li, F.; and Xie, Z. Texture segmentation from non-parametric statistical analysis of wavelet local energy. *IEEE PAMI*, to appear.

- Brechtbuhler, C.; Gerig, G.; and Kubler, O. Parametrization of closed surfaces for 3-D shape description. *Computer Vision and Image Understanding*, 61(2):154-70, 1995.
- Brechtbuhler, C.; Gerig, G.; and Szekeley, G. Compensation of spatial inhomogeneity in MRI based on a parametric bias estimate. In procs. of *Visualization in Biomedical Computing, 4th International Conference, VBC '96*. Springer Verlag, Berlin, Germany, pp. 141-6, 1996.
- Burbeck, C.A., and Pizer, S.M. Object representation by cores: identifying and representing primitive spatial regions. *Vision Res.*, 35(13):1917-1930, 1995.
- Chatfield, C. *Statistics for Technology*: Chapman & Hall, 1983.
- Choi, H.S.; Haynor, D.R.; and Kim, Y. Partial volume tissue classification of multichannel magnetic resonance images-a mixel model. *IEEE Transactions on Medical Imaging*, 10(3):395-407, 1991.
- Chover, M. *Un Modelo de Desplazamiento, Deformación y Textura, para la Representación de Características Macrogeométricas Sobre Superficies Paramétricas*. PhD thesis. Departamento de sistemas informáticos y computación, Universidad Politécnica de Valencia, 1996.
- Christensen, G.E.; Rabbitt, R.D.; and Miller, M.I. A deformable neuroanatomy textbook based on viscous fluid mechanics. In procs. of *Conference on Information Sciences and Systems*. Prince, and Runolfsson, eds. pp. 211-216, 1993.
- Colchester, A.C.F.; Zhao, J.; Holton-Tainter, K.S.; Henri, C.J.; Maitland, N.; Roberts, P.T.E.; Harris, C.G.; and Evans, R.J. Development and preliminary evaluation of VISLAN, a surgical planning and guidance system using intra-operative video imaging. *Medical Image Analysis*, 1(1):73-90, 1996.
- Collins, D.L.; Le Goualher, G.; Venugopal, R.; Caramanos, A.; Evans, A.C.; and Barillot, C. Cortical constraints for non-linear cortical registration. In procs. of *Visualization in Biomedical Computing, 4th International Conference, VBC '96 Proceedings*. Springer Verlag, Berlin, Germany, pp. 307-316, 1996.
- Cootes, T.F.; Hill, A.; Taylor, C.J.; and Haslam, J. Use of active shape models for locating structures in medical images. *Image and Vision Computing*, 12(6):355-65, 1994.
- Crow, T.J. Temporal lobe asymmetries as the key to the etiology of schizophrenia. *Schizophrenia bulletin*, 16(3):433-443, 1990.
- Crow, T.J. The relationship between morphologic and genetic findings in schizophrenia: and evolutionary perspective. *Schizophrenia*, 38:15-25, 1995.
- Crow, T.J.; Ball, J.; Bloom, S.R.; Brown, R.; Bruton, C.J.; Colter, N.; Froth, C.D.; Johnstone, E.C.; Owens, D.G.C.; and Roberts, G.W. Schizophrenia as an anomaly of development of cerebral asymmetry. *Arch. Gen. Psychiatry*, 46:1145-1150, 1989.
- Dawant, B.M.; Zijdenbos, A.P.; and Margolin, R.A. Correction of intensity variations in MR images for computer-aided tissue classification. *IEEE Transactions on Medical Imaging*, 12(4):770-81, 1993.
- DeCarlo, D., and Metaxas, D. Adaptive shape evolution using blending. *Proceedings, Fifth International Conference on Computer Vision*:834-9, 1995.
- Delingette, H. Simplex meshes: a general representation for 3D shape reconstruction, Research report 2214, INRIA, 1994.

- DeLisi, L.E.; Tew, W.; Xie, S.; Hoff, A.L.; Sakuma, M.; Kushner, M.; Lee, G.; Shedlack, K.; Smith, A.M.; and Grimson, R. A prospective follow-up study of brain morphology and cognition in first-episode schizophrenic patients: preliminary findings. *Biol. Psychiatry*, 38:349-360, 1995.
- Dempster, A.P. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society*, B 39(1):1-38, 1977.
- Dryden, I., and Mardia, K.V. *Statistical Shape Analysis*: J. Wiley & Sons, 1997.
- Eberly, D.; Gardner, R.; Morse, B.; Pizer, S.; and Scharlach, C. Ridges for image analysis. *Journal of Mathematical Imaging and Vision*, 4(4):353-73, 1994.
- Fang, T.P., and Piegl, L.A. Delaunay triangulation in three dimensions. *IEEE Computer Graphics and Applications*, 15(5):62-69, 1995.
- Gauch, J.M., and Pizer, S.M. The intensity axis of symmetry and its application to image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(8):753-770, 1993.
- Gilles, S.; Brady, M.; Declerck, J.; Thirion, J.P.; and Ayache, N. Bias field correction of breast MR images. In procs. of *Visualization in Biomedical Computing, 4th International Conference, VBC '96*. Springer Verlag, Berlin, Germany, pp. 153-8, 1996.
- Gonzalez Ballester, M.A. Visualizador de superficies paramétricas con texturas, iluminación y sombreado, Degree thesis, Universitat Jaume I, 1996.
- Gonzalez Ballester, M.A.; Zisserman, A.; and Brady, M. Combined statistical and geometrical 3D segmentation and measurement of brain structures. In procs. of *Workshop on Biomedical Image Analysis*. Santa Barbara, Ca, USA, pp. 14-23, 1998a.
- Gonzalez Ballester, M.A.; Zisserman, A.; and Brady, M. Measurement of brain structures based on statistical and geometrical 3D segmentation. In procs. of *Medical Image Computing and Computer Assisted Intervention MICCAI'98. First International Conference*. Springer Verlag, Berlin, Germany, pp. 499-508, 1998b.
- Gonzalez Ballester, M.A.; Zisserman, A.; and Brady, M. Segmentation and measurement of brain structures in MRI including confidence bounds. *Medical Image Analysis*, to appear.
- Gonzalez, R.C., and Woods, R.E. *Digital Image Processing*. Addison-Wesley Publishing Co., 1992.
- Griffin, L.; Colchester, A.; Roll, S.; and Studholme, C. Hierarchical segmentation satisfying constraints. In procs. of *British Machine Vision Conference (BMVC'94)*. Sheffield. Hancock, E., ed. BMVA Press, pp. 135-144, 1994.
- Griffin, L.D. The intrinsic geometry of the cerebral cortex. *J-Theor-Biol*, 166(3):261-73 issn: 0022-5193, 1994.
- Guillemaud, R., and Brady, M. Estimating the bias field of MR images. *IEEE Transactions on Medical Imaging*, 16(3):238-51, 1997.
- Harrison, T.R., ed. *Principles of Internal Medicine*. McGraw-Hill, 1998.
- Hayton, P. *Analysis of Contrast Enhanced Dynamic Breast Magnetic Resonance Images*. DPhil thesis. Department of Engineering Science, University of Oxford, 1998.

- Hong, Q., and Terzopoulos, D. Triangular NURBS and their dynamic generalizations. *Computer Aided Geometric Design*, 14(4):325-47, 1997.
- Hoppe, H.; DeRose, T.; Duchamp, T.; Halstead, M.; Jin, H.; McDonald, J.; Schweitzer, J.; and Stuetzle, W. Piecewise smooth surface reconstruction. *Computer Graphics Proceedings. Annual Conference Series*, 94(512):295-302, 1994.
- Ivins, J., and Porrill, J. Statistical snakes: active region models. In *procs. of BMVC94. 5th British Machine Vision Conference*. Guildford, UK. BMVA Press., pp. 2, 1994.
- Jacob, G. *Quantifying regional left ventricular function using spatio-temporal tracking techniques*. DPhil thesis. Dept. of Engineering Science, University of Oxford, 1999.
- Johnstone, E.C.; Crow, T.J.; Frith, C.D.; Husband, J.; and Kreel, L. Cerebral ventricular size and cognitive impairment in chronic schizophrenia. *Lancet*, 2:924-926, 1976.
- Kamber, M.; Collins, D.L.; Shinghal, R.; Francis, G.S.; and Evans, A.C. Model-based 3D segmentation of multiple sclerosis lesions in dual-echo MRI data. *SPIE Visualization in Biomedical Computing*, 1808:590-600, 1992.
- Kapur, T.; Grimson, W.E.L.; Kikinis, R.; and Wells, W.M. Enhanced spatial priors for segmentation of magnetic resonance imagery. In *procs. of Medical Image Computing and Computer Assisted Intervention MICCAI'98. First International Conference*. Springer Verlag, Berlin, Germany, pp. 457-68, 1998.
- Kapur, T.; Grimson, W.E.L.; Wells, W.M.I.; and Kikinis, R. Segmentation of brain tissue from magnetic resonance images. *Medical Image Analysis*, 1(2):109-127, 1996.
- Kass, M.; Witkin, A.; and Terzopoulos, D. Snakes: active contour models. In *procs. of First International Conference on Computer Vision (ICCV'87)*, pp. 259-268, 1987.
- Krishnamurthy, V., and Levoy, M. Fitting smooth surfaces to dense polygon meshes. *Computer Graphics Proceedings. SIGGRAPH '96. ACM, New York, NY, USA*, 528:313-24, 1996.
- Laidlaw, D.H.; Fleischer, K.W.; and Barr, A.H. Partial-volume Bayesian classification of material mixtures in MR volume data using voxel histograms. *IEEE Transactions on Medical Imaging*, 17(1):74-86, 1998.
- Lawrie, S.M., and Abukmeil, S.S. Brain abnormality in schizophrenia. *British Journal of Psychiatry*, 172:110-120, 1998.
- Leon-Garcia, A. *Probability and Random Processes for Electrical Engineering*. Addison-Wesley Publishing Co., 1994.
- Lester, H., and Arridge, S.R. A survey of hierarchical non-linear medical image registration. *Pattern Recognition*, 32:129-149, 1999.
- Lipson, P.; Yuille, A.; O'Keefe, D.; Cavanaugh, J.; Taaffe, J.; and Rosenthal, D. Deformable templates for feature extraction from medical images, Technical report 89-14, Harvard Robotics Laboratory, 1989.
- Listerud, J.; Lenkinski, R.E.; Kressel, H.Y.; and Axel, L. The correction of nonuniform signal intensity profiles in magnetic resonance imaging. *J-Digit-Imaging*, 2(1):2-8 issn: 0897-1889, 1989.

- Lohmann, G.; Kruggel, F.; and Von Cramon, D.Y. Automatic detection of sulcal bottom lines in MR images of the human brain. In *procs. of Information Processing in Medical Imaging, 15th International Conference, IPMI'97*. Springer Verlag, Berlin, Germany, pp. 369-74, 1997.
- Loop, C. *Generalized B-spline surfaces of arbitrary topological type*. PhD thesis. University of Washington, 1992.
- Lorensen, W.E., and Cline, H.E. Marching cubes: a high resolution 3D surface construction algorithm. *Computer Graphics (Procs. SIGGRAPH'87)*, 21(3):163-169, 1987.
- Maes, F.; Vandermeulen, D.; Suetens, P.; and Marchal, G. Computer-aided interactive object delineation using an intelligent paintbrush technique. In *procs. of Computer Vision, Virtual Reality and Robotics in Medicine (CVRMED'95)*. Springer-Verlag, pp. 77-83, 1995.
- Maintz, J.B.A., and Viergever, M.A. A survey of medical image registration. *Medical Image Analysis*, 2(1), 1998.
- Marais, P.C. *The Segmentation of Sparse MRI Images*. DPhil Thesis. Department of Engineering Science, University of Oxford, 1999.
- Maurer, C.R., and Fitzpatrick, J.M. A review of medical image registration. In: Maciunas, R.J., ed. *Interactive Image-guided Neurosurgery*: American Assoc. of Neurological Surgeons, 1993. pp. 17-44.
- McCarley, R.; Wible, C.G.; Frumin, M.; Hirayasu, Y.; Levitt, J.J.; Fischer, I.A.; and Shenton, M.E. MRI anatomy of schizophrenia. *Biol. Psychiatry*, 45:1099-1119, 1999.
- McInerney, T., and Terzopoulos, D. Deformable models in medical image analysis: a survey. *Medical Image Analysis*, 1(2), 1996.
- Monga, O., and Benayoun, S. Using partial derivatives of 3D images to extract typical surface features. *Computer Vision and Image Understanding*, 61(2):171-89, 1995.
- Montagnat, J., and Delingette, H. Volumetric medical images segmentation using shape-constrained deformable models. In *procs. of CVRMED-MRCAS'97*. Springer-Verlag, pp. 13-22, 1997.
- Naf, M.; Kubler, O.; Kikinis, R.; Shenton, M.E.; and Szekely, G. Characterization and recognition of 3D organ shape in medical image analysis using skeletonization. In *procs. of IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*. San Francisco, 1996.
- Nastar, C. *Modèles physiques déformables et modes vibratoires pour l'analyse de mouvement non-rigide dans les images multi-dimensionnelles*. PhD thesis. École Nationale des Ponts et Chaussées, 1994.
- Niessen, W.J. *Multiscale Medical Image Analysis*. PhD Thesis. University of Utrecht, 1997.
- Pennec, X. *L'Incertitude Dans les Problèmes de Reconnaissance et de Recalage. Application en Imagerie Médicale et Biologie Moléculaire*. PhD Thesis. École Polytechnique, 1996.
- Prima, S.; Thirion, J.P.; Subsol, G.; and Roberts, N. Automatic analysis of normal brain dissymmetry of males and females in MR images. In *procs. of Medical Image Computing and Computer Assisted Intervention MICCAI'98. First International Conference*. Springer Verlag, Berlin, Germany, pp. 770-9, 1998.
- Rajapakse, J.C.; Giedd, J.N.; and Rapoport, J.L. Statistical approach to segmentation of single-channel cerebral MR images. *IEEE Transactions on Medical Imaging*, 16(2):176-86, 1997.

- Rinck, P.A. *Magnetic Resonance Imaging*. Blackwell Scientific Publications, 1993.
- Rodriguez, P.L. *MRI Indications for the Referring Physician*. Aurora publishing co., 1996.
- Roll, S.R.; Colchester, A.C.F.; Summers, P.E.; and Griffin, L.D. Intensity-based object extraction from 3D medical images including a correction for partial volume errors. In procs. of *BMVC94. Proceedings of the 5th British Machine Vision Conference*. BMVA Press, Guildford, UK, pp. 1, 1994.
- Ronfard, R. Region-based strategies for active contour models. *International Journal of Computer Vision*, 13(2):229-51, 1994.
- Santago, P., and Gage, H.D. Quantification of MR brain images by mixture density and partial volume modeling. *IEEE Transactions on Medical Imaging*, 12(3):566-74, 1993.
- Schmitt, F.; Xin, C.; and Wen Hui, D. Geometric modelling from range image data. In procs. of *European Computer Graphics Conference and Exhibition (EUROGRAPHICS '91)*. North Holland, Amsterdam, Netherlands, pp. 317-28, 1991.
- Scott, G.L. The alternative snake - and other animals. In procs. of *3rd Alley Vision Conference*, pp. 341-347, 1987.
- Sled, J.G.; Zijdenbos, A.P.; and Evans, A.C. A nonparametric method for automatic correction of intensity nonuniformity in MRI data. *IEEE Transactions on Medical Imaging*, 17(1):87-97, 1998.
- Solloway, S.; Taylor, C.J.; Hutchinson, C.E.; and Waterton, J.C. Quantification of articular cartilage from MR images using active shape models. In procs. of *Computer Vision ECCV '96. 4th European Conference on Computer*. Springer Verlag, Berlin, Germany, pp. 2, 1996.
- Soltanian Zadeh, H.; Windham, J.P.; and Yagle, A.E. Optimal transformation for correcting partial volume averaging effects in magnetic resonance imaging. *IEEE Transactions on Nuclear Science*, 40(4):1204-12, 1993.
- Souhami, R.L., and Moxham, J., eds. *Textbook of Medicine*. Churchill Livingstone, 1994.
- Stoddart, A.J., and Baker, M. Reconstruction of smooth surfaces with arbitrary topology adaptive splines. In procs. of *Computer Vision ECCV'98. 5th European Conference on Computer Vision*. Springer Verlag, Berlin, Germany, pp. 2, 1998.
- Stoddart, A.J.; Hilton, A.; and Illingworth, J. Slime: a new deformable surface. In procs. of *BMVC94. 5th British Machine Vision Conference*. Guildford, UK. BMVA Press, pp. 1, 1994.
- Subsol, G.; Thirion, J.P.; and Ayache, N. General scheme for automatically building 3D morphometric anatomical atlases: application to a skull atlas. In procs. of *MRCAS'95*. Baltimore, USA, 1995.
- Szekely, G.; Kelemen, A.; Brechbuhler, C.; and Gerig, G. Segmentation of 2-D and 3-D objects from MRI volume data using constrained elastic deformations of flexible Fourier contour and surface models. *Medical Image Analysis*, 1(1):19-34, 1996.
- Terzopoulos, D., and Metaxas, D. Dynamic 3D models with local and global deformations: deformable superquadrics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7), 1991.
- Terzopoulos, D.; Witkin, A.; and Kass, M. Symmetry-seeking models and 3D object reconstruction. *International Journal of Computer Vision*, 1(3):211-21, 1987.

- Thacker, N.; Jackson, A.; Zhu, X.P.; and Li, K.L. Accuracy of tissue volume estimation in NMR images. In *procs. of Procs. of MIUA'98, Leeds, UK, 1998.*
- Thirion, J.P., and Gourdon, A. Computing the differential characteristics of iso-intensity surfaces. *Computer Vision and Image Understanding*, 61(2):190-202, 1995.
- Thirion, J.-P., and Calmon, G. Deformation analysis to detect and quantify active lesions in 3D medical image sequences, Research report 3101, INRIA, 1997.
- Thompson, P.M., and Toga, A.W. Detection, visualization and animation of abnormal anatomic structure with a deformable probabilistic brain atlas based on random vector transformations. *Medical Image Analysis*, 1(4):271-294, 1997.
- Tiede, U.; Hoehne, K.H.; Bomans, M.; Pommert, A.; Riemer, M.; and Wiebecke, G. Investigation of medical 3D-rendering algorithms. *IEEE Computer Graphics and Applications*, 10(2):41-53, 1990.
- Tiede, U.; Schiemann, T.; and Hoehne, K.H. Visualization blackboard: visualizing the visible human. *IEEE Computer Graphics and Applications*, 16(1):7-9, 1996.
- Tincher, M.; Meyer, C.R.; Gupta, R.; and Williams, D.M. Polynomial modeling and reduction of RF body coil spatial inhomogeneity in MRI. *IEEE Transactions on Medical Imaging*, 12(2):361-5, 1993.
- van den Elsen, P.A.; Pol, E.J.D.; and Viergever, M.A. Medical image matching - a review with classification. *IEEE Engng Med. Biol.*, 12:26-39, 1993.
- Van Leemput, K.; Maes, F.; Vandermeulen, D.; and Suetens, P. Automatic segmentation of brain tissues and MR bias field correction using a digital brain atlas. In *procs. of Medical Image Computing and Computer Assisted Intervention MICCAI'98. First International Conference.* Springer Verlag, Berlin, Germany, pp. 1222-9, 1998.
- Vemuri, B.C., and Radisavljevic, A. Multiresolution stochastic hybrid shape models with fractal priors. *ACM Transactions on Graphics*, 13(2):177-207, 1994.
- Webb, S. *The Physics of Medical Imaging*: Institute of Physics Publishing, 1988.
- Wells, W.M., III; Grimson, W.E.L.; Kikinis, R.; and Jolesz, F.A. Adaptive segmentation of MRI data. *IEEE Transactions on Medical Imaging*, 15(4):429-42, 1996.
- Whitaker, R.T. Algorithms for implicit deformable models. In *procs. of Fifth International Conference on Computer Vision*, pp. 822-7, 1995.
- Wicks, D.A.; Barker, G.J.; and Tofts, P.S. Correction of intensity nonuniformity in MR images of any orientation. *Magn Reson Imaging*, 11(2):183-96, 1993.
- Williams, P.L., and Warwick, R., eds. *Gray's Anatomy*: Churchill Livingstone, 1980.
- Wilson, D.L., and Noble, J.A. An adaptive segmentation algorithm for extracting arteries and aneurysms from time-of-flight MRA data. *IEEE TMI*, to appear.
- Xiaolan, Z.; Staib, L.H.; Schultz, R.T.; and Duncan, J.S. Segmentation and measurement of the cortex from 3D MR images. *Medical Image Computing and Computer Assisted Intervention MICCAI'98. First International Conference*:519-30, 1998.

Zhang, Y.; Smith, S.; and Brady, M. Segmentation of brain MR images using Markov random field. In procs. of *Medical Image Understanding and Analysis (MIUA'99)*. Oxford, UK. BMVA Press, 1999.

Zhu, S.C., and Yuille, A.L. Region competition and its analysis: a unified theory for image segmentation, Technical report 95-07, Harvard Robotics Lab, 1995.