3D Model-based segmentation of individual brain structures from magnetic resonance imaging data

by

D. Louis Collins

Department of Biomedical Engineering
McGill University, Montreal

A Thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

Abstract

This thesis addresses a specific problem of model-based segmentation; namely, the automatic identification and delineation of gross anatomical structures of the human brain based on their appearance in magnetic resonance images (MRI). The approach developed in this thesis depends on a general, iterative, hierarchical registration procedure and a 3-D digital model of human brain anatomy that contains both volumetric intensity-based data and geometric atlas data that co-exist in a brain-based stereotaxic coordinate system. The model contains features derived from an MRI atlas of gross neuroanatomy, that is the result of an intensity average of 305 brains created with an automatic stereotaxic registration procedure developed here.

The objective of this thesis is achieved by inverting the traditional segmentation strategy. Instead of matching geometric contours from an idealized atlas directly to the MRI data, segmentation is achieved by identifying the spatial transformation that, under certain constraints, best maps corresponding features between the model and a particular volumetric data set. After automatic recovery of the linear registration transform, the 3-D non-linear transformation is recovered by estimating the local deformation fields, recursively selected by stepping through the entire target volume in a 3D grid pattern, using cross-correlation of invariant intensity features derived from image data. This registration process is performed hierarchically, with each step in decreasing scale refining the fit of the previous step and providing input to the next. When completed, atlas contours defined in the model are mapped through the recovered transformation to segment structures in the original data set and identify them by name.

Experiments for registration and segmentation are presented using simple phantoms, a realistic digital brain phantom as well as human MRI data. The algorithm is used to estimate neuro-anatomical variability, to automatically segment cerebral structures and to create probabilistic representations of the same structures. Validation with manual methods shows that the procedure performs well, is objective and its implementation robust.

Résumé

Ce travail aborde le problème particulier de segmentation par modèle, et plus précisément l'identification et le contourage automatiques des principales structures anatomiques du cerveau humain, en se basant sur leur apparence dans les images de résonance magnétique (IRM). L'approche développée dans cette thèse repose sur une procédure de recalage global, itératif et hiérarchique, ainsi que sur un modèle anatomique 3D digitalisé du cerveau humain. Ce dernier contient à la fois des données volumétriques issues des images IRM et des données géométriques représentant un atlas, celles-ci étant exprimées dans un système de coordonnées stéréotaxiques lié au cerveau. Ce modèle présente des éléments provenant d'un atlas IRM de neuro-anatomie, qui résulte d'un moyennage en intensité de 305 cerveaux créé avec la méthode de recalage stéréotaxique automatique présentée dans cette thèse.

L'objectif de ce travail est réalisé en renversant la stratégie traditionnelle de segmentation. Au lieu de recaler directement les contours géométriques d'un atlas idéal sur les données IRM, la segmentation est réalisée en définissant la transformation spatiale qui, sous certaines conditions, recale le mieux des éléments homologues du modèle et des données volumétriques d'un individu particulier. Après récupération automatique de la transformation linéaire de recalage, la transformation 3D non linéaire est obtenue en estimant les champs de déformation locaux. Ceux-ci sont définis de façon récursive en parcourant complètement le modèle stéréotaxique, le long d'une grille tridimensionnelle, à l'aide de fonctions de "cross-correlation" appliquées à des éléments caractéristiques des images, invariants en rotation et translation. Ce recalage est réalisé de façon hiérarchique ; chaque nouvelle étape raffinant le recalage de l'étape précédente et fournissant les données pour l'étape suivante. Les contours de l'atlas, définis dans le modèle stéréotaxique, sont ensuite envoyés sur les données IRM de l'individu, à l'aide de cette transformation afin d'extraire les structures de ces images et de les identifier par leur nom.

Les exemples de recalage et de segmentation sont présentés sur de simples fantômes, sur un fantôme digital réaliste de cerveau, ainsi que sur des données IRM chez l'homme. L'algorithme est utilisé pour estimer la variabilité neuro-anatomique, pour segmenter automatiquement des structures cérébrales et pour créer des représentations probabilistiques de ces mêmes structures. La validation, réalisée par comparaisons avec des méthodes manuelles, montre que l'algorithme se comporte correctement, de façon objective et est robuste.

Dedication

To Isabelle and Anne-Sophie, for love, support and a truly real deadline.

Acknowledgements

First and foremost, I would like to thank my thesis supervisors Drs. A.C. Evans and T.M. Peters for their friendship, guidance, support, patience and encouragement throughout my studies, and in particular during the preparation of this thesis. Their dedication and leadership have created a stimulating and enjoyable research environment.

During my tenure at the Montreal Neurological Institute I have worked with many different people and have made many friends. In particular I wish to thank Sean Marrett who has given me more that I will ever be able to return, Chris Henri for an example that I have tried to follow, Bruce Pike for guidance and proving that it can be done, David MacDonald for programming expertise, assistance and great libraries, Atsushi Takahashi for MR expertise, Vasco Kollokian for providing more CPUs and disk space than anyone should ever have, Peter Neelin for on-line answers to almost any question, Kate Hanratty for magical results when dealing with administrative difficulties, Sylvain Milot for manual registrations, Colin Holmes for manual segmentations and Cecile Sorlie for translating the abstract. Thanks everybody!

I would like to thank my parents for support, both financial and moral, without which this thesis would not be.

Finally, I would like to express my deepest gratitude to my wife, Isabelle, for love, support and caring. Without your patience, assistance, optimism and encouragement, I never would have been able to finish.

The author would like acknowledge the assistance provided by the Positron Imaging Laboratory of the Montreal Neurological Institute for providing data to build the stereotaxic model and also the MR division of the Department of Radiology for providing the MR volumetric data used to test the algorithm. The author gratefully acknowledges financial support by the Medical Research Council (MRC SP-30), MRC/ISG Technologies, Toronto (university-industry grant UI-11509) and a student fellowship from les Fonds de Recherche en Santé du Québec (FRSQ).

Preface

This work presented in this thesis was performed at the Montreal Neurological Institute (MNI), directed by Drs Alan Evans and Terry Peters. Herein are addressed two problems common in 3-D multi-modal neuro-imaging for brain mapping: structure segmentation and volume registration.

> **Segmentation:** the delineation of the extent of a region within an image and the assignation of an identifying label to them, grouping them into some meaningful biological structure.
>
> **Registration:** the process of applying a geometric transformation to an image volume such that it is optimally aligned with another volume according to some similarity criterion.

Registration is required to be able to directly compare two or more volumes on a voxel-by-voxel basis; segmentation is required to compare corresponding regions between data sets. This thesis examines the problem of automatic segmentation of anatomical structures from volumetric magnetic resonance imaging (MRI) data for quantitative analysis and estimation of morphometric variability in normal human brain anatomy. Segmentation involves both identification and delineation; the structure must be named and its spatial limits established.

Present classification techniques applied to this problem can identify gross tissue types such as cerebro-spinal fluid, white or grey-matter. However, these methods cannot identify individual neuro-anatomical structures. While computer-vision based or image-processing based algorithms delineate elementary regions from the image data based on measures of local intensity properties, these regions are not anatomically labelled nor do their borders always correspond to the outline of a cerebral structure. Techniques from artificial intelligence such as expert systems that combine these two methods have had moderate success, but become complex and unwieldly when dealing with many structures and are highly sensitive to the nature of a required pre-segmentation process. The majority of fully automatic segmentation methods employ some form of model-based information to better constrain the problem. Conventional model-based segmentation methods attempt to fit geometric constructs, such as points and lines from contours of a pre-defined atlas, to the raster data of the medical image. Of the few successful techniques, none have been subjected to a thorough validation and evaluation, so the problem remains open.

The segmentation task presented in this thesis turns the conventional approach upside down. Instead of matching atlas contours directly to the image data, a model was created that contains both volumetric raster data as well as geometric atlas data. Both components are registered within a standard brain-based coordinate system know as *stereotaxic space*[1]. The segmentation strategy is separated into two stages: a registration step followed by a delineation step. The goal of the registration process is to find the optimal spatial transformation between the prior stereotaxic model and the individual MRI data set by maximizing the match (or overlap) of features derived from the data set and those in the model. Structure delineation is then achieved by applying the inverse of the recovered transformation to contours defined in the model, thereby outlining structures in the native data.

This organization yields two important benefits over previous methods. Since raster features from the data are being matched to similar features in the model, the resulting approach circumvents matching problems that confound traditional model-based segmentation techniques that attempt to fit one form of data representation, geometric contours of the atlas, directly to another, raster data of the image. The two representations are inherently different in contrast, noise and gradient properties. More importantly, the segmentation strategy described here results in an atlas-independent segmentation. The geometric contours defined in the atlas are not used to determine the match between data and model. Therefore, any atlas defined in the stereotaxic space can be used for segmentation, thereby allowing for the co-existence of multiple atlases, each of which is simultaneously mappable to the native MR image volume.

The matching of a data set with the segmentation model is based on a general intra-modality registration algorithm, which is applied recursively to decreasingly blurred versions of the original data to achieve its final match. Chapter 4 presents this multi-scale methodology that enables two arbitrary 3-D MRI brain image-volumes to be automatically registered together. Since a simple linear transformation model cannot account for non-linear morphometric variability between subjects, nor for anatomical differences between a subject and the model, a non-linear version of the registration technique is presented in Chapter 5. Local deformations are estimated at each node of a regularly sampled 3-D lattice defined on the volume of data. Summation of all local deformation vectors yields the global non-linear or "warping" transformation function.

While the registration technique is fundamentally independent of a model, a prior segmentation model is needed to guide the subsequent segmentation process. Chapter 6

[1] *Stereotaxic* refers to the representation of 3-D brain structures in a standardized brain-based coordinate system (Talairach *et al.*, 1967; Talairach and Tournoux, 1988) such that all brains have the same orientation, position and size in the three orthogonal directions. It is primarily used for comparison of brains from different individuals. *Stereotactic* is a related term, meaning "to touch in space" and is used to describe a class of neurosurgical procedures that use an external rigid frame to establish a common coordinate system for both imaging and surgery.

describes the model created for this thesis which contains both volumetric (raster) and geometric (polyhedral) data and serves a dual purpose. For both registration methods, the raster component serves as a target defined in the stereotaxic coordinate system. For segmentation, the geometric polyhedra defined on the model are used to delineate structures on volumetric data mapped into the standardized coordinate space. Application of the inverse transformation on these contours segment the data in its original space. Since the segmentation model takes no part in the registration, we are free to use other data representations for modelling individual structures. For example, the left head of caudate can be modelled as a set of labelled voxels in a binary-valued volume, where the value 1 represents caudate and 0 represents background. Other structures can be represented in a similar fashion. This representation was used for validation of the segmentation procedure.

Registration and segmentation experiments are described in chapters 7, 8 and 9. Chapter 7 outlines experiments in linear registration, demonstrating that the automatic registration method yields results that are as accurate as manual methods, while being more robust and reproducible. Statistical analysis of manually identified landmark points from a group of 17 data volumes, mapped into stereotaxic space, shows that there exists an inter-subject neuro-anatomical variabililty of 3.9mm, when measured as the 3-D full-width-half-maximum of an equivalent isotropic Gaussian probability distribution. This spread is due to anatomical variability across the normal population left unaccounted for by linear registration. The experiments presented in chapter 8 show that the non-linear registration technique can reduce this misregistration to less than 1mm on simulated data, even in noisy conditions. When applied to real data, the anatomical variability is reduced from 3.9mm to 2.5mm. The remaining variability is mostly due to observer-dependent errors in landmark identification. Chapter 8 concludes with the creation and validation of a 3-D anatomical variability map. Chapter 9 describes experiments in segmentation, and shows that simulations on a complex digital brain phantom results in automatically delineated structures whose volumes are recovered with less than 2% error, and that the structures overlap by more than 98%. When applied to real data and compared to manually identified structures, the measured volumes agree to within 4% and volume overlap is better than 90%. Probabilistic representations of caudate, thalamus, putamen, insular cortex, ventricles and corpus callosum are shown. The thesis ends with a discussion and conclusion in chapter 10.

Original contributions

The following are believed to be significant contributions made by this thesis:

1. atlas-independent segmentation by re-posing of the model-based segmentation paradigm into one involving a registration step followed by a delineation step.
2. removal of the extra level of abstraction in common segmentation methods; instead of directly fitting geometric constructs to raster data, the MRI data is fit to similar features of the model.
3. notion of operating explicitly in a standardized stereotaxic space based on an anatomical coordinate system, where the model used to guide the segmentation contains both polyhedral (i.e. geometric atlas contours) and volumetric (tomographic, image-based features) information in the same space.
4. automatic linear stereotaxic mapping procedure, obviating the need for the existing manual, time-consuming, subjective methods for spatial standardization required for brain morphometrics such as measuring and comparing structures between brains.
5. automatic non-linear stereotaxic mapping procedure, that alters the shape of a given brain volume so that it conforms to a standard thus automating part of the morphometrics task. The need for separate linear and non-linear registration procedures is required for stereotaxy and segmentation, and will be described in chapter 3.
6. experimental validation of the linear, non-linear and segmentation methods with simulated phantom data and real MRI volumes.
7. creation of a 3-D high-resolution digital brain phantom.
8. creation of a 3-D anatomical morphometric variability map.
9. creation of an initial probabilistic neuro-anatomical atlas.
10. insights into the extent of morphometric variability in the normal adult population.

Publications

Portions of this thesis have already been published as papers or conference proceedings.

1. D. L. Collins, P. Neelin, T. M. Peters, and A. C. Evans. Automatic 3D inter-subject registration of MR volumetric data in standardized talairach space. *Journal of Computer Assisted Tomography*, 18(2):192–205, March/April 1994.

2. D. L. Collins, T. M. Peters, and A. C. Evans. An automated 3D non-linear image deformation procedure for determination of gross morphometric variability in human brain. In *Proceedings of Conference on Visualization in Biomedical Computing.* SPIE, 1994.

3. D. L. Collins, T. M. Peters, and A. C. Evans. Multiresolution Image Registration and Brain Structure Segmentation. in *EMBS Satellite Symposium on Medical Imaging*, Proc IEEE, Rennes France, 105-110, Nov 1992.

4. D. L. Collins, T. M. Peters, W. Dai, and A. C. Evans. Model based segmentation of individual brain structures from mri data. In *Proceedings of Conference on Visualization in Biomedical Computing*, pages 10–23. SPIE press, Oct. 1992.

5. D. L. Collins, T. M. Peters, and A. C. Evans. Automatic multiresolution registration of mri volumetric data. In *SMRM11: Society Magnetic Resonance in Medicine 11th Annual Meeting*, page 4214, Berlin, Germany, Aug. 1992. Society Magnetic Resonance in Medicine.

Other work directly related to this thesis has been published as collaborations with other authors (Dai *et al.*, 1991a and 1991b; Evans *et al.*, 1989b, 1990b, 1990a, 1991a, 1991b, 1992b and 1992a)

Related work has been published in:

- computerized region of interest atlas (Evans *et al.*, 1989c; Marrett *et al.*, 1989a),

- manual registration of volumetric data from different modalities for functional mapping (Marrett *et al.*, 1989b; Evans *et al.*, 1989a, 1992c, 1993b, 1993b and 1993),

- development of stereoscopic neurosurgery planning software (Peters *et al.*, 1989, 1990 and 1992),

- merging of digital subtraction angiography with magnetic resonance and computed tomography imaging (Henri *et al.*, 1989, 1990, 1991b and 1991c),

- 3-D reconstruction of cerebral vasculature (Henri *et al.*, 1991a and 1993),

- image classification for detection of multiple sclerosis lesions: manually (Collins, 1990) and automatically (Kamber *et al.*, 1992 and 1994), and

- registration and analysis of magnetic resonance spectroscopy data (Collins *et al.*, 1992b; Arnold *et al.*, 1992; Preul *et al.*, 1992).

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Modern imaging modalities such as computed tomography (CT) and magnetic resonance imaging (MRI) have facilitated the investigation of complex spatial relationships between multiple structures within the brain. The three-dimensional (3-D) nature of tomographic imaging removes the blurring due to the unavoidable overlap of different anatomical structures that occurs in conventional two-dimensional (2-D) radiographic projection imaging techniques. Other imaging modalities such as positron emission tomography (PET) and single photon emission computed tomography (SPECT) permit the *in vivo* measurement of a wide variety of functional parameters on a regional basis in the human brain such as local hemodynamics, metabolism, pharmokinetics, tissue pH and the distribution of chemotherapeutic agents (Phelps *et al.*, 1986). Unfortunately, qualitative interpretation of these functional images is hampered by poor spatial resolution, low counting statistics, low contrast between different brain structures of interest and often the distribution of the radio-label does not reflect the underlying anatomy. Therefore, there is a general need to include external anatomical information for its interpretation and analysis. Two related solutions have been proposed in the past: 1) Merge complementary anatomical information provided by CT or MRI. 2) Overlay a brain atlas on the functional data, using it as an anatomical guide.

Registration (also described as correlation, matching, or image fusion) is required for both solutions. In a first step, the geometric transformation required to spatially align the two data sets is estimated. Afterwards, the digital nature of the data permits one of them (image or atlas) to be resampled so that it corresponds to the other. The two images can then be compared on a voxel-by-voxel basis to give an anatomical interpretation to the functional data. Due to the potential informative gains and the relative simplicity of the procedures involved, registration techniques have seen rapid growth in most areas of medical imaging. For example, correlative imaging has been used for disease diagnosis (Levin *et al.*, 1988), longitudinal monitoring of disease progression or remission (Arnold *et al.*, 1992), pre-operative evaluation and surgical planning (Kelly *et al.*, 1983; Kall *et al.*, 1987; Peters *et al.*, 1990), radio-surgery and radiotherapy treatment planning (Schad *et al.*, 1987; Sontag *et al.*, 1986), mapping of functional neuroanatomy of sensorimotor and cognitive processes (Fox *et al.*, 1985; Evans *et al.*, 1992c; Bohm *et al.*, 1992) as well as the analysis of neuro-anatomical variability among normal brains (Evans *et al.*, 1992a).

An essential step in the analysis completed for any of these examples is the identification of anatomical structures and their sub-units. For qualitative analysis, a human investigator implicitly segments the data into its structural components based on personal knowledge of anatomy derived from textbooks and atlases. The clinician builds a mental representation of the spatial relationships present within the images. For quantitative analysis of specific regions, explicit segmentation is required to separate and identify structures, and is often achieved using a brain atlas as a guide for anatomical localization and functional interpretation (*e.g.*, Talairach *et al.*, 1967, 1988). A human observer begins to manually identify anatomical structures, by extracting features (*e.g.*, landmark points, edges and regions of similar intensity) from the tomographic images and matches them with corresponding features from similar images in an anatomical atlas. The border between two structures in the volumetric data set is defined by mapping the corresponding contour from the atlas back onto the image. The position of the border is refined by comparing it, and its neighbouring structures, to the contours in the atlas. Unfortunately, manual outlining of structures is tedious, difficult and time consuming. Errors are due to subjectivity in atlas slice selection, structure interpretation, poor software interface design, poor eye-hand coordination, low tissue contrast, image degradation caused by artefacts and noise, and edge blurring due to partial volume effects (tissue mixing within a single voxel) and, perhaps most importantly, to the fact that tomographic slices are often not scanned at the same orientation as the atlas. These difficulties must be overcome to achieve robust, objective quantitative analysis. This results in the following problem statement:

> **Problem:** Given volumetric magnetic resonance image data, develop a procedure to automatically identify and delineate structures in the human brain that will facilitate neuro-anatomical quantitative analysis and permit characterization of morphometric variability across subjects.

This problem statement leads to the following specific objectives:

Objectives:

- Develop completely automatic model-based segmentation method for neuro-anatomical structures.
- Ensure method is atlas-independent.
- Advance the state-of-the-art in quantitative model-guided segmentation.
- Automate the stereotaxic transformation.
- Refine a conceptual framework and create tools for comparison of structural anatomy within and between subjects.
- Establish proof of principle for automated analysis of anatomical variability.

I have chosen to limit the scope of the thesis by the following constraints:

Scope:

- Deal only with one modality: MRI is the modality of choice for neuro-anatomical studies.
- Deal only with normal young adult population.
- No pathology addressed.
- Limited to gross pathology as seen in MRI (with resolution of approximately $1mm^3$).
- Cerebral cortex is not addressed in detail since is a completely different problem, complex enough to warrant a research project in its own right within this laboratory (MacDonald *et al.*, 1994).

A number of existing techniques have already been applied to this problem. Tissue classification methods only identify gross components of the brain and do not identify individual brain structures. While computer vision-based algorithms delineate elementary regions within the image data, these regions are not anatomically labelled nor do their borders always correspond to the outline of cerebral structures. Techniques from artificial intelligence such as expert systems that combine these two methods have had moderate success, but they cannot always correct errors introduced at the low-level segmentation stage and are therefore heavily dependent on this pre-segmentation step. The solution resides in the use of external information to guide the segmentation process (referred to as "model-based segmentation").

Digital atlases, usually consisting of a collection of geometric contours, have been used by many groups to improve significantly quantitative analysis of functional data by substantially reducing bias and size of systematic errors in region of interest definition (*e.g.*, Bohm *et al.*, 1985, 1991, 1992, Greitz *et al.*, 1991a, Seitz *et al.*, 1990, Evans *et al.*, 1988, 1991a). Typically, the atlas is first globally registered to the data, then each structure of the atlas is manipulated individually to customize the fit to the individual data volume.

The use of such atlases also facilitates comparisons between different individuals or groups through spatial standardization; by applying the inverse atlas transform to the data volume, it is thus registered to the standard reference space and can be directly compared to other volumes mapped into this space. This powerful concept forms the basis for the work presented here and is described by the term *stereotaxy*. This term is used to describe the representation of 3-D brain structures in a standardized brain-based coordinate system such that all brains have the same orientation, position and size in the three orthogonal directions (Talairach *et al.*, 1967, 1988) . It should be noted however, that a non-negligible amount of normal neuro-anatomical variability remains even after this normalization. If left unaccounted for, it can be the source of significant positional difference for homologous landmarks between subjects. Chapters 7 and 8 deal with the characterization, and where appropriate, the removal of this residual variability.

In manual techniques using an atlas and in automatic model-driven approaches, it is assumed that the target (atlas or digital model) and the data to be segmented are topologically equivalent but that internal structures are mutually deformed. This means that even though the absolute positions of structures may vary between subjects, their relative positions remain (loosely) stable and that there exists a one-to-one mapping between structures in the data and those in the atlas.[1] This assumption, and the following example, form the basis for the segmentation strategy developed in this thesis. Suppose a segmented dataset, **A**, exists, and that each of its voxels are tagged with some neuro-anatomical label. In order to segment another brain volume, **B**, in an equivalent manner, the non-linear spatial warping transformation that represents the one-to-one mapping from **A** to **B** must be recovered. Each of the labels in **A** is then mapped onto **B**, thus segmenting **B** . Hence, the segmentation problem is expressed as a registration problem.

The strategy for automatic segmentation developed in this thesis uses a digital standardized anatomical brain model to eliminate subjectivity in manual interpretation of structure definition and takes advantage of image processing techniques to extract features from the data that are robust in the presence of noise. This segmentation strategy is possible only with a two component brain model developed here that contains both volumetric (raster) information to drive the registration process and geometric (polyhedral) information used to guide the segmentation process.

The automatic segmentation procedure begins by extracting features from the given data set for matching with those stored in the raster component the brain model. Afterwards, a generic, fully automatic, hierarchical, iterative, non-linear registration procedure (described in chapter 5) is applied to identify the spatial transformation that, under certain constraints, best registers corresponding features between the model and the given volumetric data set. The transformation is recovered by detecting local deformations required to map each region of the model onto its homologue in the data in hierarchical fashion, starting with large regions and reducing their size at each iteration. The largest deformations are recovered first, based on heavily blurred copies of both data and model. The global warp is then refined by the addition of smaller deformations estimated from less blurred data. Since raster features from the data are being matched to similar features in the model, many difficulties that confound other existing techniques (Bajcsy and Kovacic, 1989) that attempt to fit geometric contours directly to raster data are avoided. In the last step of the segmentation algorithm, the geometric atlas contours defined in the model are mapped through the recovered transformation to segment structures in the original data set and identify them by name.

This strategy has one powerful advantage over other model-based segmentation methods. The design presented here makes for atlas-independent segmentation. The geometric

[1]This assumption is justified when dealing with normal neuro-anatomy, as is the case herein. Abnormal anatomy is discussed in chapter 10.

contours defined in the atlas are not used in any way to determine the match between data and model. The non-linear transformation is established from the raster data in the brain model. Therefore, any atlas defined in the stereotaxic space can be used for segmentation.

In this manner, the underlying structure of the MRI is used as a frame, onto which the atlas contours are attached. Thus, any structure can be defined in the atlas and need not follow edges present in the tomographic image volume. For example, functional areas whose edges do not correspond to an explicit anatomical border, but whose limits are correlated with the underlying anatomical substrate, may not only be defined in the atlas, but they can be segmented with this methodology.

The outlines of structures visible in the MRI volume are explicitly defined where a distinct contrast difference exists between them in the tomographic image. The position of the remaining contour is estimated from two sources: 1) the most probable location described in the model and 2) points positioned on both contour portions previously found and on nearby structures. For structures not directly visible in the MR image, their position, orientation, scale and shape are inferred from the data in the model and constrained by the position of neighbouring structures.

Summary

The motivation for this thesis stems from the need to have a fully automatic procedure to identify and delineate meaningful anatomical regions of the human brain from volumetric MRI data. The automatic method presented herein goes beyond existing techniques in significant ways while retaining some of their best features. The method is objective and its hierarchical implementation robust. The use of a standardized brain model yields consistent results and its separation into constituent raster and geometric parts allows for atlas independent segmentation.

Even though the methods developed in this thesis are applied to the normal human brain and rely on a model of same, the techniques presented here can be applied to other structures. The necessary models can be constructed using the methodology described in chapter 3. Also, the non-linear registration model described in chapter 5 can be used to register or track deformable or movable structures such as the lung or liver.

While the need for averaging across subjects to obtain significant responses is removed with new imaging techniques such as functional MRI (fMRI), the machinery for cross-subject image comparison in a standard reference frame is still necessary for drawing conclusions about the variability of functional systems across individuals. The tools developed in this project are necessary to determine whether morphological features of human brain anatomy serve as markers for cyto/myeloarchitecture. The deformation maps created by this procedure will also provide primary data for the statistical integration of morphological variability into the digital cerebral atlas.

The following chapter summarizes related work done by others in this field. This is followed by a chapter that discusses the issues related to stereotaxy and the use of a standardized space for brain morphometrics. Chapters three, four and five present methodology for linear registration, non-linear registration and segmentation, respectively. Experiments with simulated data and real MRI volumes are covered in chapters six, seven and eight for the same three topics. Please note that unless otherwise specified, all figures is this thesis represent 2-D slices of 3-D volumes and all processing for registration and segmentation is performed entirely in 3-D. This point must be stressed, since the majority of previous work has been completed only in 2-D, and other authors simply allude to the 3-D solution. Although computationally more difficult, the 3-D algorithms make the problem tractable, since a 2-D procedure can never account for data out of the plane of interest. The thesis ends with conclusions and suggestions for further research for linear registration, non-linear registration and segmentation that are offered in chapter ten.

Chapter 2

Literature Review

2.1 Introduction

For the purposes of this thesis, we shall define the terms classification and segmentation:

- *Classification* involves labeling pixels in terms of different classes by grouping pixels that have similar characteristics, based on the measurement or estimation of different features known to exist in the image using low-level operators with small areas of support (Duda and Hart, 1973). It does not demand spatially contiguous voxels within any single class.

- *Segmentation* is the parcellation of the input image into meaningful contiguous groups of voxels. The word "meaningful" indicates a task-dependent definition. For this project, the segmentation task involves the identification and delineation of commonly-recognized structures in the human brain that are labelled as such with neuroanatomy textbooks.

The measurements from two or three different MR imaging scans of the same brain are usually enough for existing simple classification algorithms using only image processing techniques or statistical pattern recognition methods such as thresholding, multi-feature classification, and automatic clustering (*e.g.* (Duda and Hart, 1973; Puliti and Tascini, 1989; Amamoto *et al.*, 1990; Herrmann *et al.*, 1988; Delapaz *et al.*, 1985; Kübler and Gerig, 1990; Gerig *et al.*, 1990; Gerig *et al.*, 1991a; Bonar *et al.*, 1993) to characterize each voxel into a unique tissue class (see Fig. 2.1 for an example). Local image structure such as edges are used to improve the separation between distinct regions (Feehs and Arce, 1987; Panda and Rosenfeld, 1978; Chuang and Udupa, 1989; Bomans *et al.*, 1990; Zucker and Hummel, 1981). (A good review of MR image classification algorithms can be found in (Bezdek *et al.*, 1993).) In contrast to a classification algorithm that may label pixels as grey-matter, white-matter, cerebro-spinal fluid (CSF), muscle or fat, a segmentation procedure should automatically identify and establish the spatial limits in 3-D of cerebral structures such as the ventricles, thalamus, caudate nucleus or hippocampus.

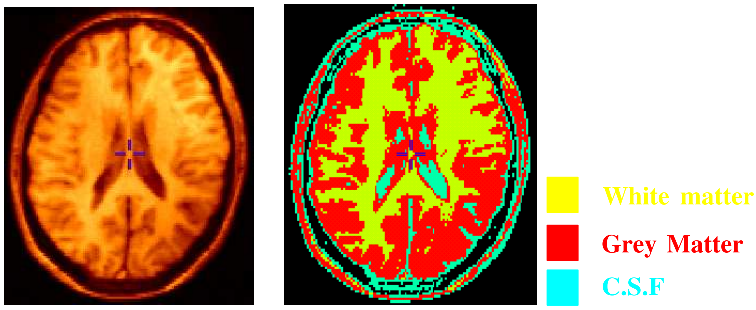

Figure 2.1: Classification example

The image on the left shows a typical transverse MR image through the level of the ventricles. The image on the right is the result of application of a minimum distance classifier, identifying tissue types of grey-matter, white-matter, and CSF.

While classification techniques permit the calculation of total volumes for the different classes, each statistic cannot be broken down into separate components (*e.g.*, amount of grey matter in left temporal lobe or CSF volume in right lateral ventricle) because only image-based information is used for classification (only features computed locally are being used) and they do not take into account more global information, such as the voxel's spatial position. The calculation of such measures requires knowledge of the brain anatomy combined with segmentation to calculate these values. (See Fig 2.2.)

The work presented here is fundamentally different than the data-driven approaches, *e.g.*, Pizer *et al.* (1989,1990), in basic computer vision research, where regions are defined as "reasonable" on the basis of feature homogeneity within the raw data, rather than some external, in this case neuro-anatomical, definition. This project is concerned with the segmentation of cerebral structures from MR image volumes of the normal brain. The image data represents only one measurement concerning the underlying neuroanatomy, and by itself is not enough to discriminate between similar adjacent, but anatomically distinct regions that are differentiated on the basis of histology, cyto-architecture and connectivity, cyto-chemistry or function. Data from external sources is required to guide and constrain the segmentation process, in order to achieve the required goal.

2.2 Manual segmentation methods

The most basic and primitive approach to image segmentation relies on the human expert to identify and outline structures. Computer programs have been developed to assist in the task and relieve some of the tedium involved in painting and contouring regions (e.g, Kennedy *et al.*, 1989, Collins, 1990), however it is the user who locates and defines regions of interest, based on personal anatomical knowledge.

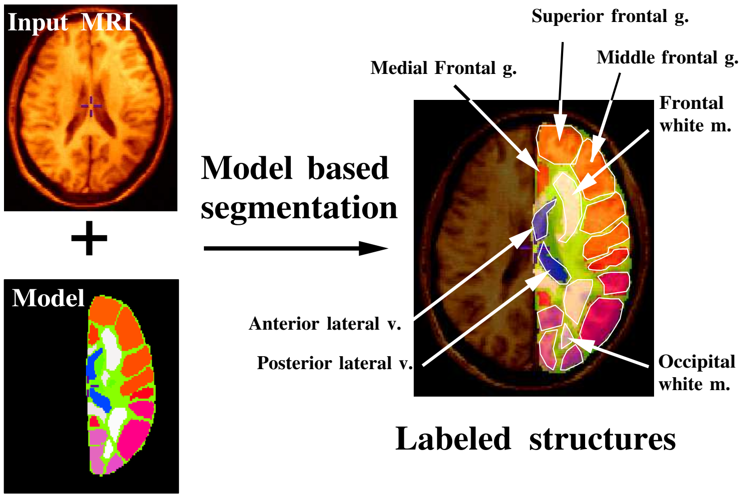

Figure 2.2: Model-based segmentation

This diagram shows the paradigm for model-based segmentation. A predefined anatomical is fit to an input MRI in order to segment individual structures from the brain volume.

Structural information can be stored explicitly in digital brain atlases, often represented as a collection of geometric outlines (Bohm *et al.*, 1983; Evans *et al.*, 1988). Some atlas-based systems have been developed to assist in anatomy teaching or to serve as a reference for radiologists or surgeons (Höhne *et al.*, 1988; Tiede *et al.*, 1993; Kazarnovskaya *et al.*, 1991). Others have been developed to serve as a guide for the interpretation of functional images by linearly mapping the atlas to the data or vice-versa (Schifter *et al.*, 1993; Berlangieri *et al.*, 1993; Lehmann *et al.*, 1991), however in these examples the atlas is not used to directly segment the data. Fitting techniques must be used to delineate structures and to serve a method for measuring morphometric change (due to growth, evolution, disease progression, and other shape modifications) by deforming the atlas to volumetric data.

The group from the Karolinska hospital (Bohm *et al.*, 1983, 1985, Seitz *et al.*, 1990, Greitz *et al.*, 1991a, 1991b) has developed a digital brain atlas that can be applied to a data set to simplify structural identification. Manually chosen global 3-D scaling, rotation and positioning are used first to remove the positional variance naturally occuring among different subjects and fit the atlas to the brain image volume. A number of additional parameters (*e.g.*, off-axis scaling, skew, scoliosis and second order scaling) are then modified by the user to achieve a non-linear fit. The goodness of fit is judged visually by the user, who must make compromises between areas that fit well and those that don't, due to the limited number of parameters used to define the non-linear transformation.

Evans *et al.* (1988, 1991a and Marrett *et al.*, 1989a) have developed a similar system with a region of interest (ROI) atlas containing 120 sub-structures of the normal human brain[1]. In a first step, homologous pairs of landmarks are manually chosen on both the dataset and atlas. These landmarks are used to define the best affine transformation that minimizes (in the least-squares sense) the residual between point pairs. Two methods are then possible to address the remaining differences due to morphological variability between subjects and the atlas: one manual and the other semi-automatic.

In the first case, manually determined scaling, rotation and translation are applied to each of the ROIs. When needed, editing of contour points achieves the final fit. The authors have shown this method reduces inter-observer variability in structure delineation (Evans *et al.*, 1988,1991a). In the second case, the morphological differences are dealt with using a well-behaved non-linear mapping that determines the smoothest continuous spatial deformation that transforms one set of 3-D coordinates onto another set of homologous points. We have implemented the so-called "thin-plate" spline (TPS) procedure (Duchon, 1976; Bookstein, 1989) and have extended it to 3-D (Evans *et al.*, 1991b). The procedure decomposes the overall deformation into a series of principal warps of decreasing geometric scale, exactly fitting the homologous points and interpolating between them. This is mathematically analogous to the bending energy required to deform a thin metal sheet so that a set of points on the sheet have a defined height above corresponding points on a flat surface. Unfortunately, these methods have not been practical for routine use as a deformation/warping model because of the subjectivity involved in selecting the precise location and number of points that will define the non-linear deformation.

2.3 Automated segmentation methods

Automatic medical image matching and structure identification is a difficult task, due to the anatomical variability between patients, the distinct physical properties measured by the imaging modalities, differences in subject positioning and variability of acquisition parameters such as slice thickness and pixel size. A wide variety of methods has been proposed to solve the registration problem for different applications. Most segmentation methods follow a two-step paradigm: features are extracted from the data set and then a mapping is found to assign labels (such as structure names) to them.

This mapping is typically found using one of two techniques: 1) The first builds a symbolic mapping between the extracted features (usually small homogeneous regions) and an iconic model of the structure(s) to be segmented. Expert rule-based systems are often used to achieve this mapping. These are described in the next section. 2) The other type of algorithm calculates a spatial transformation function that best maps the features of one data set into the other. These routines differ in the dimensionality (2-D, 3-D or 4-D), the types of features being matched (points, lines, surfaces, volumes), the number of parameters used to define the transformation (selected from translations, rotations, scalings and shears, amongst others), the method used to find the transformation (direct

[1]This atlas is described in more detail in Chapter 6, since it forms the basis of the segmentation model used in this thesis.

solution, or iterative search) and the amount of interaction needed on the part of the user. Such registration algorithms are described in the following sections.

2.3.1 Rule-based systems

Anatomical knowledge can be stored explicitly along with segmentation heuristics in semantic form such as an 'if-then' rule. Raya *et al.* (1989,1990) apply a simple rule-based system to the low-level classification of MRI brain scans to extract large structures such as brain, extra-cerebral CSF and ventricles. Chen and Sontag (1989) have developed a more complex knowledge-based expert system using a blackboard data structure and input from multiple modalities to segment finer structures, *e.g.*, thalamus, putamen, and globus pallidus.

Dellepiane *et al.* (1986, 1987, 1991, Vernazza *et al.*, 1987, Serpico *et al.*, 1987) have a relatively complex system that starts with basic image processing techniques (geometric and amplitude correction, edge preserving smoothing followed by the creation of an edge image and then region-growing segmentation) to create elementary regions (ERs - similar to Pizer's reasonable regions) from MR image of the head. A knowledge database containing atlas information and rule-based expert-system is then used to identify neuroanatomical structures by grouping the ERs. Independently, Dhawan *et al.* (1988, 1990) have built a very similar system for segmentation of abdominal CT images and for MR images of the brain (Arata *et al.*, 1991; Dhawan and Arata, 1992). In the latter, the model database (in the form of structure masks) is dynamically updated with each additional validated segmentation.

While these systems have achieved some measure of success, the results are heavily dependent on the preprocessing and segmentation processes. Since such methods separate the structure identification into two tasks – pre-segmentation followed by rule application – they are not always able to correct errors introduced at the low-level segmentation stage. Davis and Taylor (1991) address these difficulties with an expert-system that uses model-driven, objective-specific forward chaining to find required objects, with graph-based backtracking to resolve errors. The rules actually specify what low-level image processing methods to apply to extract needed features.

Other algorithms do not explicitly employ if-then rules to drive the segmentation. Instead, anatomical constraints are implicitly incorporated into the procedure. One of the earliest examples is that of Kaneda *et al.* (1978), where contour extraction and 3-D reconstruction is used to identify dilated ventricles in CT images. The regions are extracted with a modified region-growing algorithm after thresholding. Identification is accomplished using model-guided identification on two dimensional slices and then reconstructed by tessellation in 3-D. The anatomical knowledge stored in the model is used to correct errors in the segmentation.

Brummer has used anatomical knowledge to design and implement morphological operators that discriminate between desired and unwanted structures to extract brain contours from images (Brummer *et al.*, 1991), and similar anatomical constraints were used to detect the longitudinal fissure with the Hough transform (Brummer, 1991).

Menhardt has used iconic fuzzy sets with a symbolic model to describe properties and relationships between objects in the brain to achieve the top-down analysis for automated interpretation of MR images (Menhardt, 1988b,1988a) such as the extraction of brain contours (Menhardt and Imme, 1988). The analysis is accomplished by recursive knowledge-based separation of groups of pixels from the MR tomograms using techniques from pattern recognition, image processing and artificial intelligence. Zachmann has continued this work, building an iconic model (a voxelated volume, where the value in each voxel represents the probability of existence of a structure) for classification of the fluid spaces of the brain (Zachmann, 1991a, 1991b). Kamber *et al.* (1992) have built a similar probability model of white matter, used as a mask when identifying multiple sclerosis lesions.

Similarly, Kapouleus has used proton density and T2-weighted images with a geometric model to segment MR images with MS lesions (Kapouleas and Kulikowski, 1988, 1990). The atlas is fit in two steps: the brain (as a whole) and the longitudinal fissure are found first, and used to determine brain size, position and orientation (linear transformations), then only the segmented brain surface is used to fit a deformable geometric model consisting of brain and ventricular surfaces. In these last techniques, it is necessary to register the data with the model used to guide the segmentation. Success of these methods is dependent on the quality of the registration between data and model.

In most semi-automatic methods, features are extracted manually and can be either points (Evans *et al.*, 1991a; Schiers *et al.*, 1989), surfaces (Pelizzari and Chen, 1987, 1989, 1993), volumes (Alpert *et al.*, 1990; Farber and Stokely, 1988) or a mixture of points, curves and surfaces (Zubal *et al.*, 1991). These techniques are described below.

2.3.2 Point-based registration

In point-based registration techniques, corresponding landmarks are identified in both data sets to be matched. These points may be obtained from external fiducial markers attached to a head-holder or preferably to the head itself, or from internal anatomical landmarks. Registration is achieved by identifying the transformation that maps one set of points onto the other in some optimal way, such as minimizing the least-squares residual mis-match between paired points. The transformation is typically solved using one of several numerical methods; *e.g.*, alignment of centers of gravity and principal axes (Langron and Collins, 1985; Sibson, 1978; Sibson, 1979), singular valued decomposition (SVD), QR decomposition, or LU decomposition (Press *et al.*, 1988). The manual point-based methods using internal landmarks have the advantage that no special pre-scan procedures are required and that retrospective matching of data is possible. Since explicit correspondence is available for the landmarks identified, another advantage is that both inter-subject and cross-modality registration is limited only by the user's ability to identify homologous features between subjects/modalities. This is dependent upon the local implementation of the user interface for point tagging as much as on the fundamental considerations of algorithmic robustness, image contrast and noise. The main disadvantage of these techniques resides in possible user-bias and subjectivity involved in the selection of both the number and placement of features required.

In order to avoid these difficulties, Dalton and du Boulay (1988) have built a knowledge-based system to register images by extracting invariant key geometric features from a segmented input image. These features serve as landmarks and are matched to corresponding feature landmarks of images stored in a canonical view database. While this technique is wholly dependent on the segmentation method, it may lend itself to cross-modality registration if the features extracted from the different imaging modalities correspond to the same physical location with the brain. Allain *et al.* (1992) also automatically detects landmarks. His method uses model-based constrained search techniques to locate the anterior and posterior commissures (the AC and PC points) and brain extents that are used to automatically map the brain into the stereotaxic space.

2.3.3 Surface-based registration

In many semi-automatic techniques, the user is needed to identify similar features in both data sets without the requirement of explicit point-to-point correspondence. Typically, an optimization procedure is used to fit the two data sets together, and in doing so, implicitly defines the correspondence between the two volumes. For example, surfaces manually extracted from two data sets can be used for registration. Known as the "head-and-hat" procedure, Pelizzari has used an optimization technique to find the transformation that minimizes a function estimating the distance between two surfaces where one serves as a fixed head and the other as a movable hat (Pelizzari and Chen, 1987, 1989). Niew *et al.* (1991) improves on this method by automating the surface extraction step. After reviewing existing surface-based registration methods in (Collingnon *et al.*, 1993b), the author further improved the head-and-hat technique by speeding up the optimization procedure as well as refining the distance function (Collingnon *et al.*, 1993a).

The surface-based methods have the following advantage over point-based techniques: an anatomically-trained expert user is not required to identify the features to be matched. However, care must be taken to ensure that the extracted surfaces correspond to the same structure. While the inner-skull surface can easily be extracted from MR or CT data, the same is not true for functional PET or SPECT data. A small change in threshold can change the position of the surface, and thus change the transformation.

The time required to calculate the similarity function can be reduced considerably if the data describing the surface can be compressed to contain only salient features. Thirion, Gueziec and Ayache (Thirion *et al.*, 1992, Gueziec and Ayache, 1992a, 1992b) extract high curvature lines from surfaces extracted from CT volumes of the skull that correspond to ridges around the nose and orbits as well as below the mandible. Matching is based on hashing using curvature and torsion measures. Since the curve-to-curve match is subject to stronger constraints, the implicit correspondence achieved is tighter than that obtained by the surface to surface measures.

The matching data can be compressed even further when one considers that a surface naturally defines an enclosed volume. Both Farber and Stokely (1988) and Alpert *et al.* (1990) have independently developed a method to calculate the principal axes of the volume, and have used this information to align the two data sets. Farber also calculated an equivalent transformation based on a tensor formulation, but found that the principal axis transform to be more robust (Farber and Stokely, 1988). Both techniques suffer when the entire brain is not contained within the data set because the principal axis tends to align itself in the image plane and not along the main axis of the images object.

2.3.4 Automatic linear registration

Fully automatic registration methods are similar in principle to the semi-automatic techniques; however features are automatically extracted from the image data, obviating the need for user intervention. These methods differ in the type of features extracted and the objective function used to measure the match.

The simplest methods operate directly on the pixel intensity values. Venot *et al.* (1983) maximize the number of zero-crossings in a 2-D difference scintigram image. Mintun and Lee (1990) extended this concept to 3-D to register two PET volumes of the same subject obtained on different days. Minoshima *et al.* (1992) used the same objective function for orienting a single volume. Here, the left half of the volume is subtracted from the mirrored right half to center and vertically align the transverse and coronal planes for left-right comparisons of a single PET data-set. The absolute difference of pixel intensity values is dependent on the intensity and contrast in the original image. More statistically robust objective functions have been incorporated into different matching algorithms.

Cross-correlation of pixel intensity values has been used in 2-D by Junck *et al.* (1990) and Rizzo *et al.* (1991). Junck uses an exhaustive search procedure to find the best transformation for PET-PET or SPECT-SPECT matching while Rizzo achieves a direct result by calculating the correlation in frequency space for 2-D inter-modality registration. Peli *et al.* (1997) has developed a method to speed the correlation process by automatically detecting points of interest in retinal photographs. The correlation is only calculated for small neighbourhoods centered on these points of interest to find their homologues in the target image. Once found, the point-based methods described above are used to determine the transformation.

Woods *et al.* (1993a) also uses the pixel intensity for both intra- and inter-modality 3-D registration. The chosen objective function measures the variance of the pixel-to-pixel intensity ratio and this function is minimized by a 12-parameter Newton-Raphson optimization. The principle employed here is that for intra-modality matching, the ratio should be unity with minimum variance. For cross-modality matching, this ratio will not be unity, but should be fixed for any tissue class, again with minimum variance.

Convolution with partial differential Gaussian kernels may be used to extract geometric features automatically from the raw image intensities without the need for manual definition of a region or contour. van den Elsen (1993) describes the L_{vv} operator that detects high intensity ridges or low intensity valleys from an image. This operator can be thought of as the second derivative of image intensity perpendicular to the direction of the maximum intensity gradient. When the width of the Gaussian kernel is properly selected, the skull can be extracted as a ridge in CT and a valley in MR for 2-D registration of MR and CT images of the same subject. Extension of the operator to 3-D is not trivial, since there is an infinite number of directions perpendicular to the gradient. Instead of convolution, the operator is set as a minimization problem, where the direction of second derivative maximum (for valleys) or minimum (for ridges) is found under the constraint that this direction be perpendicular to the gradient.

Maintaining the simplicity of convolution with a Gaussian kernel for feature detection, Collins *et al.* have developed a completely automatic multi-resolution 3-D intra-modality inter-subject registration procedure based on correlation of invariant features derived by convolution of the data with zeroth and first order Gaussian kernels (Collins *et al.*, 1992a, 1994b). At present, this procedure is used routinely at the Montreal Neurological Institute (MNI) to map MRI volumes into stereotaxic space. This technique is described in chapter 4 and is extensively compared to the methods of both Woods 1993a and Minoshima 1993.

2.3.5 Automated non-linear registration

While a linear transformation yields a reasonable approximation to fit one data set onto another or onto an atlas, non-linear transformations must be used to account for fine structure differences persisting after linear registration. There are many different parameterizations possible to define the type of non-linear transformation.

One group of methods employs a domain-specific *a priori* parametization of the non-linear transformation. Among the first non-linear computerized matching methods described in the literature, the "rubber mask", deals with 2-D human chromosome identification (Widrow, 1973). After linear registration, the author uses 16 manually adjustable parameters – length, width, angle and curvature per arm – to fit a typical model chromosome to a sample data set. Relieving the user of the parameter adjustment task, Fischler and Elschlager (1973) have described a 2-D face-recognition system that uses a simple model with 7 pre-defined features (hair, eyes, cheeks, nose and mouth) with connections between them characterized by spring models. Automatic template-matching is used to locate each feature with a local linear transformation, however each local match is constrained by the global deformation modeled with the springs. These two papers show the power of a simple deformation model applied to well defined problems. Their main disadvantage is that they require an *a priori* parametization of deformation and user-defined correspondence of features.

In order to address this problem, Burr (1981a,1981b) has developed a technique that uses 1) a nearest-neighbour rule on points, 2) a nearest-line with similar orientation rule on curves and 3) a nearest-region with high similarity rule on images to determine correspondence automatically for hand-drawn characters and 2-D face images. The global warp is calculated iteratively, bringing corresponding features closer together by estimating one deformation vector per feature and the warp is interpolated between features. A stiffness parameter controls the amount of interaction between neighbouring areas to reduce the effect of mismatches. The stiffness is reduced at each iteration so that global fitting is achieved first, and local fitting last. Leclerc and Benchimol (1991) have developed a similar system for registration of digital subtraction angiography images to correct for patient motion during scanning that uses local image cross-correlation but uses a different method to avoid mis-matches. An iterative coarse-to-fine fitting strategy is used by first matching points positioned on a coarse grid and then interpolating the warp onto a new grid with reduced spacing at each successive iteration.

Another method to correct for patient motion in 2-D abdominal images has been developed by Gerig *et al.* (1991b). After application of a ridge/valley detection algorithm, a user identifies the two renal contours. The kidneys can then be segmented and tracked in successive images using a Hough transform on contour segments. The transform yields displacements for each segment, which are then used to geometrically warp the images back onto the first for comparison and analysis.

Neural networks have been used by Kosugi *et al.* (1993) to find corresponding parts between a subjects brain and a standard 2-D MR image. Using 2 degrees of freedom per node, the network generates a set of deformation vectors, organized on a rectangular lattice, based on local cross-correlation of image intensity to establish a local match. The neural net achieves consensus between nodes to ensure a smooth mapping.

By themselves, linear and non-linear registration are not enough to achieve segmentation. These routines are designed to match one image onto another for motion detection, tracking and correction. However, no explicit identification of parts of the image, i.e. labelling, is implied by the process. A model or atlas must be incorporated with the registration method to build a viable segmentation scheme. These methods are discussed next.

2.3.6 Nonlinear warping for segmentation.

Many model-based segmentation techniques share the concept of deforming a model to fit the image data. Edges/surfaces are extracted from the image/volume data and are used to guide the shape of the model subject to certain constraints. For example, Terzopoulos *et al.* apply elasticity theory to construct differential equations that model the behavior of non-rigid curves, surfaces and solids as a function of time (Terzopoulos *et al.*, 1987). These techniques have been used to simulate non-rigid objects for animation (Terzopoulos and Witkin, 1988; Terzopoulos and Fleischer, 1988; Metaxas and Terzopoulos, 1992) and to model 3-D shape and non-rigid motion from single and stereo 2-D images as well as video sequences (Terzopoulos *et al.*, 1988; Terzopoulos and Metaxes, 1990). In 2-D, a "snake" is an energy-minimizing spline that is deformed from an initial model by external forces that pull it toward image features such as lines and edges while physically-based constraints imposed by the elastic model counteract them (Kass *et al.*, 1988). Improving on these techniques, Nastar and Ayache (1993) have proposed a similar deformation based on masses attached by springs, however they attach extra springs between non-neighbouring nodes to model some volumetric properties of the object and can thus maintain the initial model shape. The spring equations are integrated over time until an equilibrium is found.

These physically-based methods provide a flexible method of deformation, but two problems exist: 1) The modeling of natural phenomena usually results in a deformation that simply finds the nearest local energy minimum, not necessarily the global minimum. 2) The deformation process is modelled as a physical process. While this may be useful in practice, care must be taken not to interpret the resulting displacement map as a true physical deformation of the underlying object.

The routines described above have been used mostly to segment a single object from an image or volume. The problem addressed in this thesis is that of segmentation of multiple complex 3-D structures from volumetric data sets. While the deformation of many interconnected objects is possible by extending the techniques described above, the authors have not yet done so.

Inspired by the work of Fischler and Burr described above (Fischler and Elschlager, 1973; Burr, 1981a; Burr, 1981b), Broit has developed a 3-D non-linear registration algorithm to match volumetric data by deforming one data set onto another to address the multiple object problem (Broit, 1981). Like Burr, warping is based on local matching of corresponding grid points constrained by a model of a continuous elastic solid. This technique has been used in 2-D to match CT images to a predefined atlas image (Bajcsy and Broit, 1982,1983). The atlas contours are then mapped through the recovered deformation onto the original data for segmentation. This method was extended to 3-D atlas matching in (Dann *et al.*, 1988; Bajcsy and Kovacic, 1989) and tested with simulated data in (Gee *et al.*, 1993).

Zhengping and Mowforth (1991) have described similar work applied to 2-D MR data. The authors defined a volumetric anatomical model by manually painting the scalp, bone, CSF grey-matter and white-matter on one MRI data volume. A customized model is instantiated to match a given MR image, by colouring the voxels from each tissue class with a grey-level value determined by the scanning parameters of the volume to be registered. After an exhaustive search linear registration procedure, a single arbitrary slice is extracted from the model for input to the non-linear matching procedure. Warping of the model is based on a coarse-to-fine estimation of local deformation in the same manner as Leclerc and Benchimol (1991). Unfortunately, the method is applied on 2-D slices only, and therefore it can never fully recover the complete non-linear deformation required to match two volume together.

The non-linear segmentation method developed in this thesis and elsewhere (Collins *et al.*, 1992c, 1992d, 1994a) is an extension of the work of Burr, Bajscy and Zhengping described above. Although the work presented here is similar, our approach differs significantly in a number of respects such as the (1) type of data used, (2) the features matched, (3) the definition of the atlas, (4) the optimization procedure and (5) the amount of user interaction needed to complete the segmentation. The consequences of these differences is that unlike the previous methods, the technique developed here operates on volumetric MR data and estimates local deformations entirely in 3-D and requires no user interaction whatsoever. These algorithms are compared in detail in the discussion section of chapter 5.

2.4 Summary

The specific aim of this thesis is neuro-anatomical segmentation. Ideally, a perfect segmentation device would obviate the need for any form of post-processing since pointing the device at some location would yield the desired information, *e.g.*, this tissue belongs to the head of the caudate. The nearest one can come to the ideal device is the human visual system combined with the expert knowledge of a trained anatomist. Since medical images are an indirect representation of what we wish to observe, measuring some physical property as opposed to a unique structural identifier, the data must be subject to post-processing techniques to extract meaningful neuroanatomic labels.

In previous work described above, algorithms enhanced with high-level domain knowledge, such as anatomical information (Brummer, 1991; Menhardt, 1988b; Menhardt, 1988a), atlases in 2-D (Dalton and du Boulay, 1988; Kaneda *et al.*, 1978; Bajcsy and Broit, 1982; Bajcsy *et al.*, 1983), in 3-D (Kapouleas and Kulikowski, 1988; Kapouleas, 1990; Dann *et al.*, 1988; Bajcsy and Kovacic, 1989) or expert systems (Dellepiane *et al.*, 1987; Kapouleas and Kulikowski, 1988; Vandermeulen *et al.*, 1989; Chen and Sontag, 1989; Dhawan and Juvvadi, 1988) have not always been able to correct errors accomplished at the low-level segmentation and the latter can become quite complex and unwieldy, sometimes requiring manual intervention. In these traditional AI approaches to segmentation, the pre-segmentation step serves as a data reduction stage, compressing the information contained within the original input data to salient features which succinctly describe the data. If the feature extraction process were perfect, there would be no need for reasoning systems to manipulate the data. Unfortunately, features (edges, points, regions) derived from the image volume are indirect measures of the structure labels that we wish to apply to the brain data. In many situations, the image shows no signal change at boundaries between structures and these bounds must be inferred from neighbouring boundaries. Hence, more global information such as spatial relationships must be represented in the symbolic rule-base, since extracted features represent only local information. When these operators fail, possibly because assumptions about the nature of the image are ill-founded or that spatial relationships conflict, the whole system breaks down. Another representation of the imaging model must be used to replace the edges and regions in the data as well as the anatomical structures and their spatial relationships

In all approaches that use an explicit model or atlas, the latter is fit to the data using linear or non-linear registration. The registration methods discussed in this chapter are summarized in Table 2.1 for linear registration and Table 2.3 for non-linear methods. The model-driven techniques assume that the geometric model and the data to be segmented are topologically equivalent but that internal structures are mutually deformed with a one-to-one mapping. Segmentation is achieved by detecting local deformations to match model to data. Even though there may be a large variability in the overall volume, position and shape of individual brains, the relative position of internal structures is maintained. Thus, segmentation is viewed as a registration and interpolation problem, with structures being delineated by the deformed atlas.

Both Kapouleas (1988, 1990) and Bajcsy and Kovacic (1989) stretch and match only the cortical and ventricular surfaces, and interpolate the positions of other structures to yield "good outlines for most of the basal ganglia". However, very little validation has been published other than the work of Gee *et al.* (1993) that used classified data and simulated deformations. The technique of fitting a model to the data is powerful, however linear interpolation between cortical and ventricular surfaces is not sufficient to locate the position of internal structures. Similar to basing a decision function on a prior probability only, interpolation gives the 'best guess' of structure identification at a given location based on distant features, but it does not take into account local structure. Obviously, the use of local features will yield a better local fit.

Some of the segmentation/fitting errors of these methods may be due to incorrect assumptions that lead to errors in the design of model-based segmentation algorithms. In many cases, the models used to match and describe the data do not have the same representation as the data itself (i.e. raster image data vs. geometric line segments). Thus, features used to measure the match between model and data are not directly comparable and may be subject to error. For example, some model-based segmentation algorithms match contours from an idealized atlas to those detected by an edge operator without regard to image contrast, noise, slice thickness or sampling (Bajcsy and Broit, 1982; Bajcsy *et al.*, 1983; Dalton and du Boulay, 1988; Kaneda *et al.*, 1978; Kapouleas and Kulikowski, 1988; Kapouleas, 1990), while others attempt to match edges based on image intensity of photographs of microtomed brains with that of computed tomographic (CT) images (Dann *et al.*, 1988; Bajcsy and Kovacic, 1989). While these techniques may have some success in regions where both model and data have similar contrast, in many other regions they fail where their imaging model breaks down. To date, no method exists to sub-segment tomographic image data reliably into anatomically meaningful entities.

Linear registration methods

Name	dim	features matched	num param	sol'n type	inter-action	method	subject	modal -ity	data type	notes
Allain 1992	3	landmarks	13	proc	auto		inter	inter	M, P, A	automatic search for AP, PC and brain extents.
Alpert 1990	3	volume	6	direct	semi-auto	P.A.T.	intra	inter	C, M, P	principal axis transform (P.A.T.) on manually extracted volumes.
Collingnon 1993	3	surf-surf	6	iter	auto	optimi-zation	inter	inter	C, M, P	Modified ``Head and Hat" - improved speed, objective function.
Collins 1994	3	int+grad	12	iter	auto	optimi-zation	inter +atlas	intra	M, A	maximize cross-correlation.
Dalton 1988	3	landmarks	12	iter	auto	proc.	inter	intra	M, A	expert system used to identify landmarks.
Evans 1991	3	landmarks	9	direct	man	svd	inter +atlas	inter	C, M, P, A	matching of manually identified homologous landmarks.
Farber 1988	3	intensity	9	direct	semi-auto	P.A.T. & tensor	intra	inter	C, M, P	PAT or tensor based on manually extracted volumes.
Fox 1985	3	landmarks	9	direct	man	AC-PC based	inter +atlas	inter	M, P	manual identification of GI line to estimate AC-PC line.
Friston 1989	3	landmarks	9	direct	man	AC-PC based	inter +atlas	inter	M, P	manual identification of AC-PC line and brain extents.
Junck 1990	2	intensity	3	direct	auto	search	intra	intra	P, S	exhaustive search using cross correlation of image intensity.
Lemoine 1991b	3	landmarks	13	direct	man	AC-PC based	inter +atlas	inter	M, P, A	manual interation to define Talairach transformation.

Table 2.1: Summary of linear registration methods. Legend follows table 2.3.

Linear registration methods (cont.)

Name	dim	features matched	num param	sol'n type	inter-action	method	subject	modal-ity	data type	notes
Minoshima 1992	3	l-r int diff	5	iter	auto	search	intra	intra	P	like Venot:: but left minus right for centering:: rotational correction.
Mintun 1990	3	intensity	6	iter	auto	search	intra	intra	P	like Venot:: extended to 3D, using multi-dimensional search.
Neiw 1991	3	pts-surf	6	iter	auto	optimi-zation	intra	inter	C, M, P	Removed manual delineation step from ``Head and Hat".
Peli 1987	2	intensity	3	iter	auto	SSD	intra	intra	retina	auto detect points using local cross-correlation.
Pelizarri 1987	3	pts-surf	6	iter	semi-auto	optimi-zation	intra	inter	C, M, P	``Head and Hat". Optimization with Powell and simulated annealing.
Rizzo 1991	2	intensity	3	direct	auto	xcorr	intra	inter	C, M, P	Frequency space cross-correlation.
Schiers 1989	3	landmarks	12	direct	man	fitting	inter	inter	all	linear or polynomial fit to manually identified landmarks.
Venot 1983	2	volume	3	iter	auto	optimi-zation	intra	intra	scinto-grams	maximize zero-crossings in difference image.
Woods 1993	3	intensity	12	iter	auto	optimi-zation	inter	inter	M,P	minimizes variance of intensity ratio with Newton-Raphson opt.
van den Elsen 1993	3	L_{vv}	6	iter	auto	optimi-zation	inter	inter	C,M,P	uses ``chamfer distance" minimization.
Zubal 1991	3	all	12	?	semi-auto	?	inter	inter	all	

Table 2.2: Linear registration methods (cont.) Legend follows table 2.3.

Non-linear registration methods

Name	dim	features matched	objective function	transform type	sol'n type	inter-action	subject	object	data type	notes
Bajscy 1989	3	Heukle	cross-corr	elastic	iter	semi-auto	inter	brain	CT	multi-scale match to coloured atlas
Burr 1981a	2	intensity	abs diff	elastic	iter	auto	inter	faces	photo	deformation modeled with descreasing stiffness at each iteration
Collins 94	3	|grad|	cross-corr	plastic	iter	auto	inter	brain	MR	multi-scale match to average MR atlas
Fischler 1973	2	masks	abs diff	masses+ springs	iter	auto	inter	images	photo	predefine structure masks attached with springs
Friston 1991	2	intensity	line integral	plastic	direct	auto	intra	brain	PET	match line integral along rows, then along columns
Gerig 1991b	2	edges	hough xform	plastic	direct	semi-auto	intra	abdo-mon	MR	user id of 1st roi, hough transform tracks motion
Kosugi 1993	2	local intensity	cross-corr	plastic	iter	auto	inter	brain	MR	neural-net selects local deforamtion yielding best neighbourhood consensus
Leclerc 1991	2	local intensity	cross-corr	elastic	iter	auto	intra	vessels	DSA	multi-scale match to correct for patient motion
Lemoine 1991a	3	landmarks	lsq fit	tri-variate polynomial	direct	man	inter	brain	CT, MR	use polynomial for C^1 continuity of previous piecewise linear transform
Widrow 1973	2	intensity	binary overlap	fixed parameters	iter	man	inter	iamges	photo	``Rubber-mask:'' manual adjustment of parametized deformation
Zhenping 1991	2	intensity DoG	cross-corr	plastic	iter	auto	inter	brain	MR	multi-scale match to MR model

Table 2.3: Summary of non-linear registration methods. Legend on following page.

The following abbreviations are used for the Tables 2.1, 2.2 and 2.3:

- Features matched: *surf-surf* is surface to surface, *int+grad* is intensity and gradient magnitude, *DoG* is difference of Gaussians.
- Solution type: *proc* is procedural, *iter* is iterative.
- Interaction: *auto* is automatic, *semi-auto* is semi-automcatic, *man* is manual.
- Method: *P.A.T.* is principal axis transform, *proc* is procedural, *svd* is singular valued decomposition.
- Subject: *inter* is inter-subject, *intra* is intra-subject, *atlas* is atlas to subject.
- Modality: *inter* is inter-modality, *intra* is intra-modality.
- Data-type: *M* is MRI, *P* is PET, *C* is CT, *A* is atlas.
- Objective function: *cross-corr* is cross-correlation, *abs-diff* is absolute difference, *lsq* is least squares.

Chapter 3

Stereotaxy: background & theory

The purpose of this chapter is fourfold: 1) to describe the rationale for, and the design of, the Talairach stereotaxic coordinate system, 2) to present the implementation of the MNI stereotaxic coordinate system that addresses problems associated with the original Talairach implementation, 3) to show the MRI atlas created using this locally defined stereotaxic mapping transformation, 4) to describe the differences between segmentation for anatomical parcellation and stereotaxy for the analysis of anatomical morphometric variability.

3.1 Introduction

The methodology presented in this thesis originates from a basic concept first proposed by Talairach *et al.* (1967,1988) for stereotactic neurosurgery. His goal was to provide external reference material, such as an anatomical atlas, for guidance in stereotactic procedures – a blind surgical procedure where the target is approached from a small twist-drill hole in the skull. The guiding principle is to *establish a standard anatomically-based coordinate system within the brain* so that locations from the atlas may be mapped back into the patient brain space to predict the position of sub-structures in the basal ganglia for surgical treatment. Anatomical landmarks visible in the different medical imaging modalities are used to define the coordinate system.

In the neurosurgical context, stereotactic means *to touch in space* and refers to peri-operative procedures in a single patient. The derived term *stereotaxic* means *to organize in space* and is used in this thesis to describe the use of a brain-based coordinate system for analysis of populations of brains, most often from normal subjects. A standardized method for identification of structure location and position is achieved so that regions of interest can be compared between brains using these coordinates. In current brain mapping research, 3-D brain volumes are transformed into this coordinate system and resampled onto a common sampling grid, such that all brains have the same orientation and size, making comparisons of different populations on a voxel-by-voxel basis possible.

The transformation to this coordinate system also provides a means for enhancement of functional signals by averaging images in this space (Fox *et al.*, 1988). This paradigm allows information (anatomical, metabolic, electrophysiological, chemical, architectonic) from different brains to be spatially organized and catalogued by mapping all brains into the same coordinate system (Fox *et al.*, 1994). Finally, in the original Talairach spirit, the coordinate corresponding to a particular structure, as defined by an atlas in this coordinate system, can be used to predict its location in a brain volume of a given subject mapped into the same space. The latter forms the basis of the segmentation procedure presented in chapter 6, since structure outlines defined in the digital model described below are used to predict the location and orientation of structures within any given brain.

3.2 Segmentation and Stereotaxy

The methodologies developed here depend on a stereotaxic segmentation model that serves two purposes: 1) it defines the neuro-anatomical structures that will be identified (i.e., labelled) and delineated by the segmentation algorithm and 2) it also defines a standard coordinate system in which to estimate spatial neuro-anatomical variability across individuals. While determination of structure labels and estimation of positional variability are closely related, they aim to achieve two very distinct goals:

- Segmentation describes the identification and assignment of a label to each brain volume voxel according to some predefined neuro-anatomical segmentation model. Conceptually, labelling requires that a perfect match be found between an unlabelled source volume and the model defining the segmentation, so that labels from the target volume can then be transferred to voxels within the source. The mapping function must have a large number of degrees of freedom to establish a one-to-one correspondence between source and target volumes, theoretically as many as the total number of voxels.

- In contrast, the analysis of anatomical variability may be achieved by mapping a group of labelled volumetric data sets into a common frame of reference with a limited number of well-defined degrees of freedom (i.e., stereotaxy). The distance between homologous points from different subjects is then used to estimate their positional variance with respect to the chosen coordinate system and mapping. Since this amounts to establishing a convention, these measurements are only meaningful if the number of degrees of freedom and the dimensions of the 3-D target coordinate space are well-defined.

Implicit in this analysis is the notion that perfect segmentation has been performed, by the previous segmentation step, and that a one-to-one correspondence exists between the

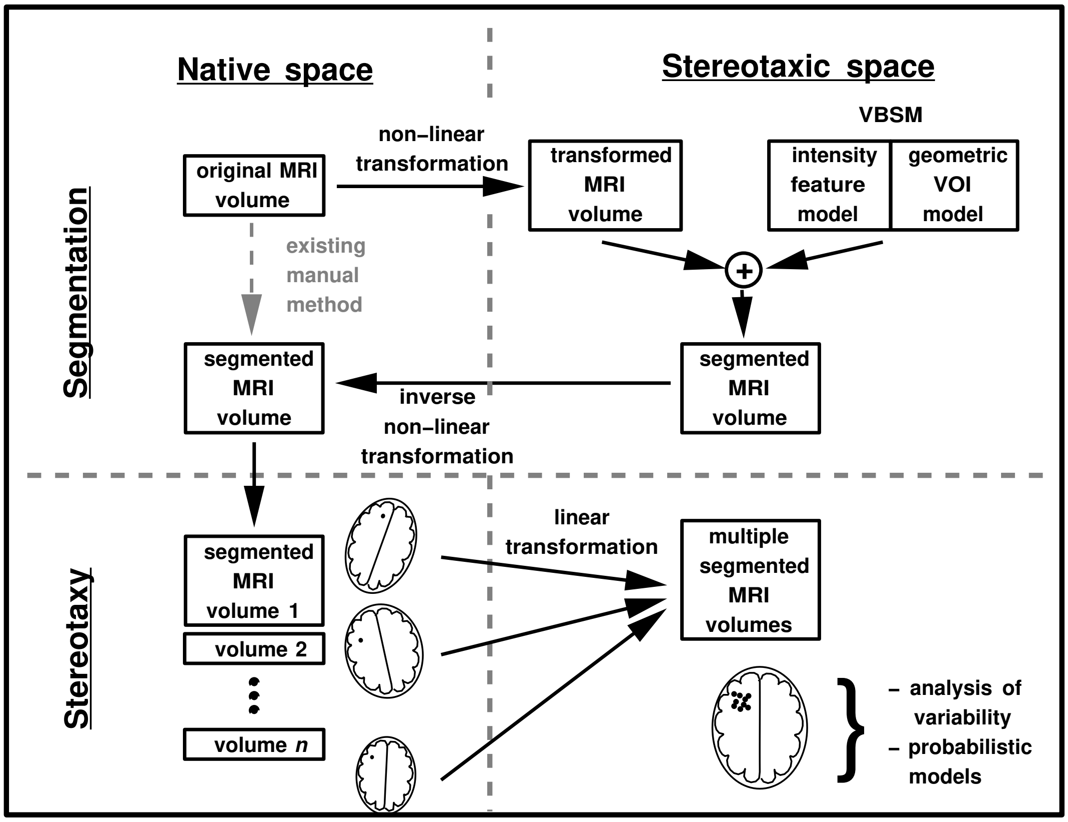

Figure 3.1: Differences between segmentation and stereotaxy.

Segmentation is achieved by estimating the non-linear transformation required to optimally map the original MRI volume to the intensity features of the VBSM. Once resampled into the stereotaxic space, the MRI volume is implicitly segmented by the VOI atlas that is co-resident with the intensity features in the VBSM. Structures are identified by the VOI labels, and their outlines defined by the VOI's geometric contours. In fact, any brain volume, registered in the stereotaxic space is segmented by the VOI atlas. The inverse non-linear transformation is applied to both the labels and contours to segment the original MRI volume in native space. In principle, a *perfect* non-linear deformation which maps the source onto the VBSM also yields a *perfect* segmentation of the source brain.

Stereotaxy requires segmented volumes to be mapped into the stereotaxic space using a well-defined transformation with a fixed number of degrees of freedom in order to standardize all positional information to a predefined brain size and orientation. The difference in position of a given anatomical landmark or segmented structure in the standard coordinate system allows for the estimation of anatomical morphometric variability as well as the creation of probabilistic models of anatomy.

two brains. Figure 3.1 summarizes the differences between segmentation and stereotaxy.

The determination of the *best* reference frame for stereotaxy is non-trivial as it depends on region of interest (e.g, cortex, deep brain structures or a specific lobe) and the difference function to be measured (*e.g.*, based on image intensity, gradient magnitude, or landmark correspondence). One linear coordinate system for the analysis of anatomical variability could be defined as that which would minimize the positional variance (or residual) of all points mapped into the space. This so-called minimum variance frame could also be defined by maximizing the overlap of equivalent neuro-anatomical labels across all brains, resulting in a slightly different coordinate system. Therefore, all reasonable frames are essentially equivalent, although some are more convenient than others. A brain-based coordinate system very similar to that proposed by Talairach and Tournoux (1988) has been chosen here, since it is a well-circulated published atlas and provides a common anatomically standardized coordinate system for precise and unambiguous reporting of the location of points of interest within the human brain.

3.2.1 Segmentation

Within this stereotaxic coordinate system, a segmentation model has been created that contains both intensity-based and geometric data. The volumetric data are represented at different spatial scales (as in pyramid techniques) where a number of features are associated with each voxel at the particular scale (*e.g.* grey level intensity and gradient magnitude). The geometric data consists of a collection of 3-D polyhedral objects that represent important anatomical structures within the human brain. Information regarding each structure (name, tissue type, etc...) is also stored with each structure in the atlas.

The segmentation procedure consists of matching the particular subject data set to the stereotaxic segmentation model. This process is approached as a feature detection (at a local level) and registration (to the global model) problem which involves finding the transformation that best maps volumetric features of the given data set to those in the model brain volume. The method uses the non-linear warping method described in chapter 5 to register a given brain volume with the model, based on the estimation of local deformations derived from local neighbourhood correlation of invariant features calculated from image data (Collins *et al.*, 1992c, 1992a, 1994a). Once the two data sets are in optimal registration, the contours defined in the model are directly applicable to the transformed brain volume, since the intensity-based data and the polyhedral data are co-extensive in the model. Conversely, the inverse transformation may be applied to the stereotaxic segmentation model to outline structures back in the native data volume.

3.2.2 Stereotaxy

The method for estimation of anatomical variability is based on the analysis of the differences in mapping between homologous point pairs, mapped into stereotaxic space with both the linear and non-linear transformation models. For example, a transformation allowing only three degrees of freedom to represent translations in the $x-$, $y-$ and $z-$directions is not likely to provide much information other than how consistently subjects are positioned within the scanner. Transformations allowing six degrees of freedom, with the three additional parameters used to map rotations about the coordinate axes, allow alignment of both position and orientation of brain volumes. Hence, measurement of different structures based on homologous points allows analysis of size variability across a population. If an additional scaling parameter is incorporated, then the analysis of size variability can be normalized to a standard gross brain size. Notwithstanding noise considerations, additional parameters in the transformation should yield better fits of data to the model and allow more specific variability studies to be accomplished.

The following sections describe the stereotaxic coordinate system, the creation of the segmentation model, the automatic stereotaxic mapping procedure and finally, the segmentation procedure.

3.3 The stereotaxic coordinate system

3.3.1 Talairach definition

The Talairach stereotaxic coordinate system (Talairach *et al.*, 1967; Talairach and Tournoux, 1988) is based on the identification of the line contained in the inter-hemispheric plane, passing through the superior aspect of the anterior commissure (AC) and the inferior edge of the posterior commissure (PC): the so-called AC-PC line (see Fig. 3.2). The second axis is defined by a vertical line in the midline plane, perpendicular to the AC-PC line, passing through the posterior margin of the AC (VAC-line). The intersection of these two lines on the AC serves as the origin. The third axis runs laterally through the origin, perpendicular to the AC-PC line and the VAC-line. The extent of the standard space is defined by the smallest bounding box that completely contains the cortex. Talairach's normalized proportional grid is established on the given brain within the bounding box. It is divided into 12 sub-volumes: 2 divisions laterally (left, right), 2 vertically (above and below the AC-PC line) and 3 in the anterior-posterior direction (from posterior limit to PC, from PC to AC and from AC to anterior limit). This piecewise linear transformation was an attempt to model non-linear differences between brains which would not be adequately handled by simple linear rescaling along the orthogonal axes.

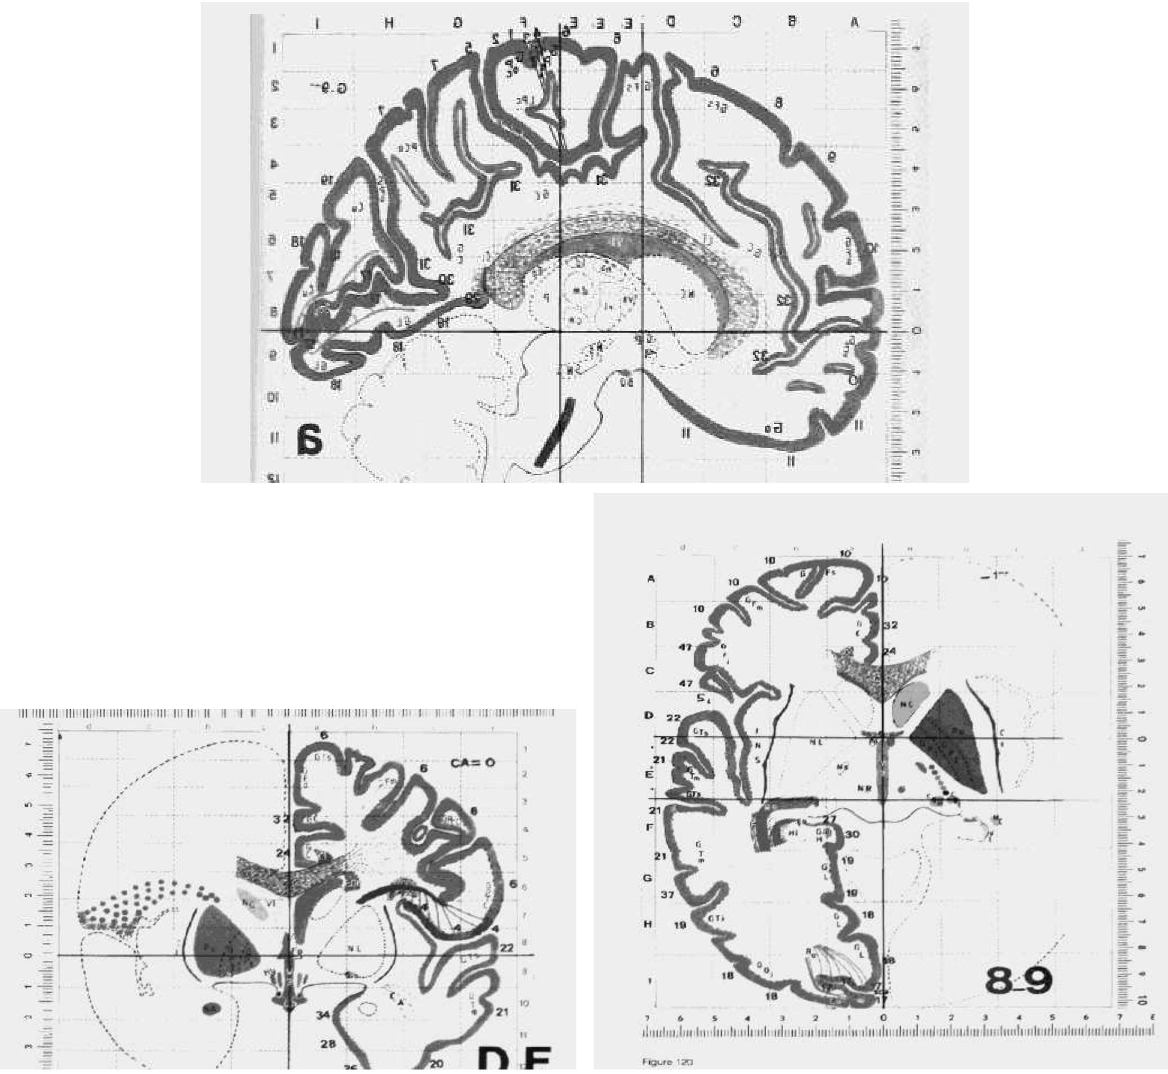

Figure 3.2: The Talairach stereotactic coordinate system.

These three images show a transverse (z=-1mm), sagittal (x=-9mm) and coronal (y=0mm) slice through the Talairach atlas.

3.3.2 Methods for stereotaxic mapping

A number of different methods have been published to map a given data set into the Talairach brain-based coordinate system (Fox *et al.*, 1985; Friston *et al.*, 1989; Evans *et al.*, 1992c; Lemoine *et al.*, 1991b). Manual methods used to transform brain volumes into the Talairach stereotaxic space are based on the estimation of the AC-PC line and brain extents. Since the AC and PC points are relatively close together (approximately 25mm apart) and difficult to identify reliably even on MRI, significant errors can be introduced which are magnified at the level of the cortex. Originally, Fox *et al.* (1985) inferred the location of the AC-PC line from that of the glabella-inion (GI) line, visible on a lateral skull radiograph in combination with a CT-like transmission scan acquired on the PET scanner. Friston *et al.* (1989) used a sagittal PET image to establish this line directly. Owing to the difficulty of identification of the line from AC and PC alone when the images are noisy, they used linear regression through the commissures and other points in the midline plane. Our group (Evans *et al.*, 1992c) uses a similar technique with MRI data, fitting 5 well-separated midline landmarks with a least-squares approximation to the AC-PC line on the MR images, which were previously co-registered with a PET volume of the same subject (Evans *et al.*, 1989b; Evans *et al.*, 1991a).

The extents of the brain in the lateral and vertical directions are assumed to lie on the perpendicular bisectors of the AC-PC line. While this approach is not strictly equivalent to the Talairach bounding box definition, experience has shown that the extreme cortical edges often lie near the bisector. The St.-Louis group (Fox *et al.*, 1985) scale each slice to fit the corresponding atlas slice. Friston *et al.* (1991) use two 1D plastic transformations on each slice to map the edges of the cortex to that of the atlas. In these last two methods, there exists a possibility that the stereotaxic mapping in not continuous in the $z-$direction, since the in-plane scale changes from slice to slice without any vertical constraints. This is not the case with the implementation presented here, the same $x-$ and $y-$scale is used for the whole volume.

With a few exceptions (*e.g.* Lemoine *et al.* (1991b)), most groups have avoided the awkward piecewise linear approach of Talairach, preferring simply either linear or a true non-linear mapping. A single affine transformation is used at the MNI to map data volumes into stereotaxic space. The present 9 parameter linear transformation can be separated into translation, rigid body rotation and anisotropic scaling along preselected axes (the AC-PC line, the VAC-line and the left-right lateral line). This model was selected since the Talairach space has traditionally involved rescaling along the stereotaxic coordinate axes. Since only 9 parameters are used, the transformation does not include a shearing (or off-axis scaling) term. Recently, Woods *et al.* (1993a) have reported a 12-parameter algorithm which allows non-orthogonal scaling axes. The inter-related issues of (1) the number of allowed degrees of freedom needed to adequately capture the observable variance, (2) computational efficiency and (3) robustness are addressed further in the discussion of chapter 4.

To summarize the manual stereotaxic mapping procedure used at the MNI:

1. Manually identify five mid-line landmarks: inferior margin of both genu and splenium of corpus callosum, inferior margin of thalamus, superior aspect of cerebellum and the occipital pole.

2. Use least-squares estimation to fit AC-PC line to landmarks.

3. Identify anterior and posterior brain limits on the AC-PC line.

4. Identify brain height on vertical perpendicular bisector of AC-PC line.

5. Identify brain width from left and right extents of horizontal perpendicular bisector of AC-PC line.

6. Resample brain volume into stereotaxic space using orientation and scaling determined above.

3.3.3 Problems with the Talairach atlas

The Talairach coordinate system and atlas form the most commonly used frame of reference for reporting cognitive activation stimuli (Fox *et al.*, 1985). However, it is not without limitations for brain mapping. The atlas is the result of a detailed study of a single human post-mortem specimen. The subject, a 60 year old right-handed European female, is considerably older than the young normal subjects most often studied in *in-vivo* cognitive research. Also, the subject is not necessarily representative of the average normal population, since an individual brain, even though normal, may represent an extreme of normal variability. Moreover, while the atlas shows transverse, sagittal and coronal planes, positional information is not always consistent between orthogonal planes. Finally, sampling is only 4mm between slices.

The Talairach model uses an awkward piece-wise transformation in an attempt to reduce anatomical variability. While this system works reasonably well for structures near the origin (AC) of the coordinate space such as the basal ganglia, its accuracy decreases as distance from the origin increases. As indicated by Talairach 1988, the errors can be particularly large for gyri of the cerebral cortex, where it can reach magnitudes of 1-2cm in the AP direction for the central sulcus, for example. This is due to both the limitations of the affine transformation as well as the significant amount of variation in human cortical morphology. Unfortunately, the atlas contains no rigorous analysis of anatomical variability – especially important in cortex where most activation foci occur.

Despite these practical difficulties, the Talairach coordinate system remains the gold standard for anatomical interpretation of many cognitive functional activation studies due to the power of the basic Talairach concept. However, the precise anatomical localization of focal activation derived from PET using the Talairach atlas alone may lead to over-interpretation of the results. For example, Drevits *et al.* (1992) have pointed out that an activation foci from an anxiety study was associated with the temporal pole when using the Talairach atlas for anatomical interpretation. However, when matched MRIs were used, it was found that the activation corresponded to flexion of the temporalis muscle due to jaw clenching during the study.

3.3.4 Automation of stereotaxic mapping

While validation studies in 37 MRI volumes of young normals show that manual estimation of the AC-PC line by least-squares fitting is a robust technique (Evans *et al.*, 1992c), the landmarks must be selected manually. Besides being time-consuming, the method is highly dependent on the quality of the software used to locate landmarks in the 3-D volume. Furthermore, another problem exists with manual stereotaxic mapping methods. Even though most manual procedures follow the Talairach definition for manual identification of the AC-PC line and extreme limits of the brain volume, they differ on implementation details, as described above in section 3.3.2. Therefore, there are methodological differences as well as subjective differences among observers for the coordinate reported for a particular point of interest within the Talairach stereotaxic space. Standardization of methodology is a key factor in identifying a particular cerebral location.

An objective, automatic technique is needed to address these problems. The registration method presented in detail in chapter 4 achieves automatic registration of two brain volumes by 3-D cross-correlation between features extracted from the two data sets. The automatic stereotaxic mapping procedure applies this general registration technique, using an MRI atlas described in the next section to serve as the target volume. The transformation recovered is applied to the given data set to resample it within the stereotaxic space. This procedure is described in greater detail in section 4.6 and is experimentally validated in chapter 7.

3.4 Stereotaxic MRI atlas

In this section, a model is described that is the first step in the evolution from the original single-brain Talairach stereotaxic space (Talairach and Tournoux, 1988) to a voxel field operational space defined by the minimum variance frame that produces the maximal overlap of neuro-anatomical labels across all structures. Since there is significant morphometric variability between individuals (Steinmetz and Seitz, 1991; Evans *et al.*, 1991b), the initial model was established from 241 MRI data sets from young healthy normals (age 23.4$\pm$4.0; 170 males; 71 females) (Evans *et al.*, 1992a), rather than using the brain from a single subject. These brain volumes are from a data-base of MRIs of normal volunteers, acquired as part of an on-going brain mapping programme at the MNI. The data sets were acquired on a Philips Gyroscan 1.5 Tesla superconducting magnet system using a multi-slice spin-echo acquisition, yielding 64 non-overlapped T1-weighted (TR=400ms, TE=30msec) 2mm thick image planes. Each slice is stored on a 256$\times$256 matrix with $1 \times 1 \times 2$mm voxels. The model was created in two stages, the second removing some of the subjectivity involved when defining the first stage model.

Stage I:

The first stage model is based on a voxel-to-voxel intensity average of all data sets. This was accomplished by registering the data sets to the standardized model space by manually identifying the AC-PC line, the interhemispheric fissure and the brain extents using the method described above in section 3.3.2 (Evans *et al.*, 1993b; Evans *et al.*, 1992c). After registration, each individual MRI volume was then intensity normalized so that each volume had the same mean intensity and resampled on voxel grid parallel with the AC-PC, VAC and lateral lines of the Talairach coordinate system, with spacing of 1.34mm x 1.72mm x 1.50mm in X, Y and Z directions. The entire ensemble of resampled brain volumes was averaged to create the mean stereotaxic MRI brain (see Fig. 3.3).

All major anatomical features were clearly defined in the average. The mean positions of primary fissures and sulci, *e.g.*, central, occipito-parietal, inter-hemispheric, calcarine, cingulate, sylvian, were well-defined while the more variable secondary and tertiary sulci were less evident. The cingulate, occipito-parietal and calcarine sulci appear highly conserved (see Fig. 3.3-c). The mean MRI also exhibits details not usually apparent on individual brain volumes, since in regions near the origin of the stereotaxic space, the noise reduction due to image averaging outweighs the resolution loss caused by anatomical variability. For example, the dorsomedial nucleus of the thalamus is visible and the slight contrast between the optic radiations and adjacent white matter is evident (Fig. 3.3-a). Comparison of left and right hemispheres, preformed by reflecting one hemisphere about the midline of stereotaxic space and subtracting the images, reveals some interesting features consistent with previous work (Geschwind and Levitsky, 1968; Rubens *et al.*, 1976; Witelson, 1977; Galaburda *et al.*, 1978; Chui and Damasio, 1980). There is clear evidence of left occipital and right frontal petalias, i.e., anatomical asymmetries (le May and Kido, 1978; Weinberger *et al.*, 1982).

Stage II:

After initial validation experiments with the automatic stereotaxic transformation algorithm, it was applied to reduce the impact of subjectivity in the choice of landmarks used to build the initial model of Stage I. In the second stage, a new MRI atlas brain volume was created by automatically mapping each brain into the standard coordinate system from its original position in the native data base. After the required stereotaxic transformation was applied to 305 data sets (239 males, 66 females, mean age 23.4 $\pm$ 4.1 years), each volume was intensity normalized and then resampled, using tri-cubic interpolation, to a 256x256x160 raster with voxel dimensions of 0.67mm $\times$0.86mm $\times$ 0.75mm, resulting in a sharpened average MRI volume (experiments described in chapter 7 show the automatic stereotaxic mapping procedure to be more robust than the manual method used to create the initial MRI atlas). Figure 3.4 shows a close-up view of the stereotaxic MRI atlas.

Figure 3.5 shows the Talairach atlas overlaid on the mean MRI atlas. Overall, both

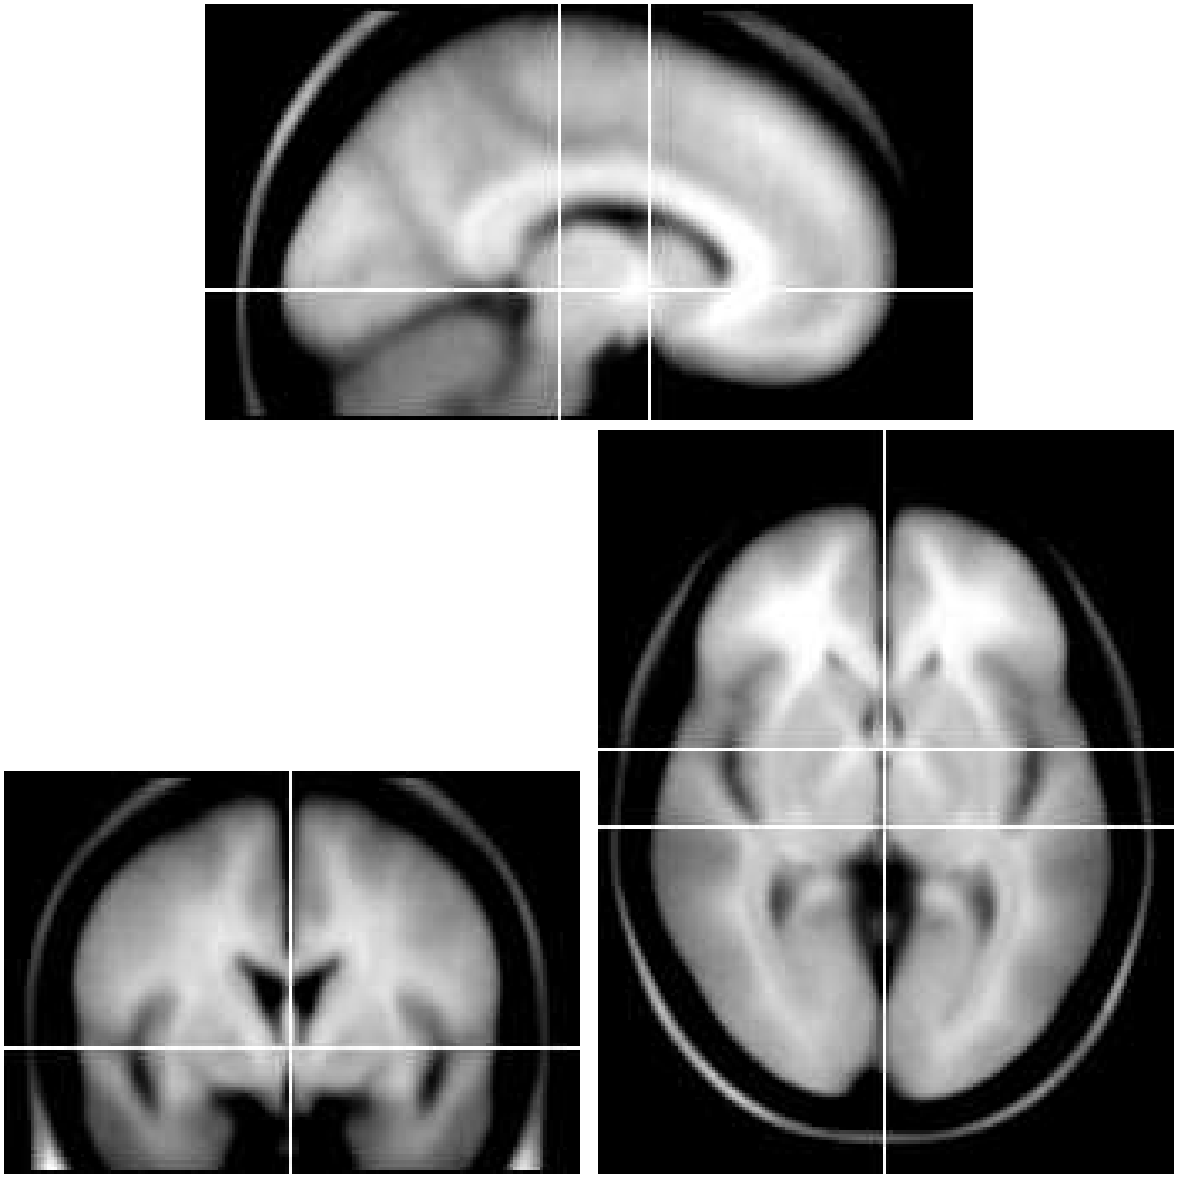

Figure 3.3: The MNI stereotaxic coordinate system.

These three images show a transverse (z=-1mm), sagittal (x=-9mm) and coronal (y=0mm) slice through the mean MRI atlas, corresponding to the Talairach atlas slices shown in Fig. 3.2.

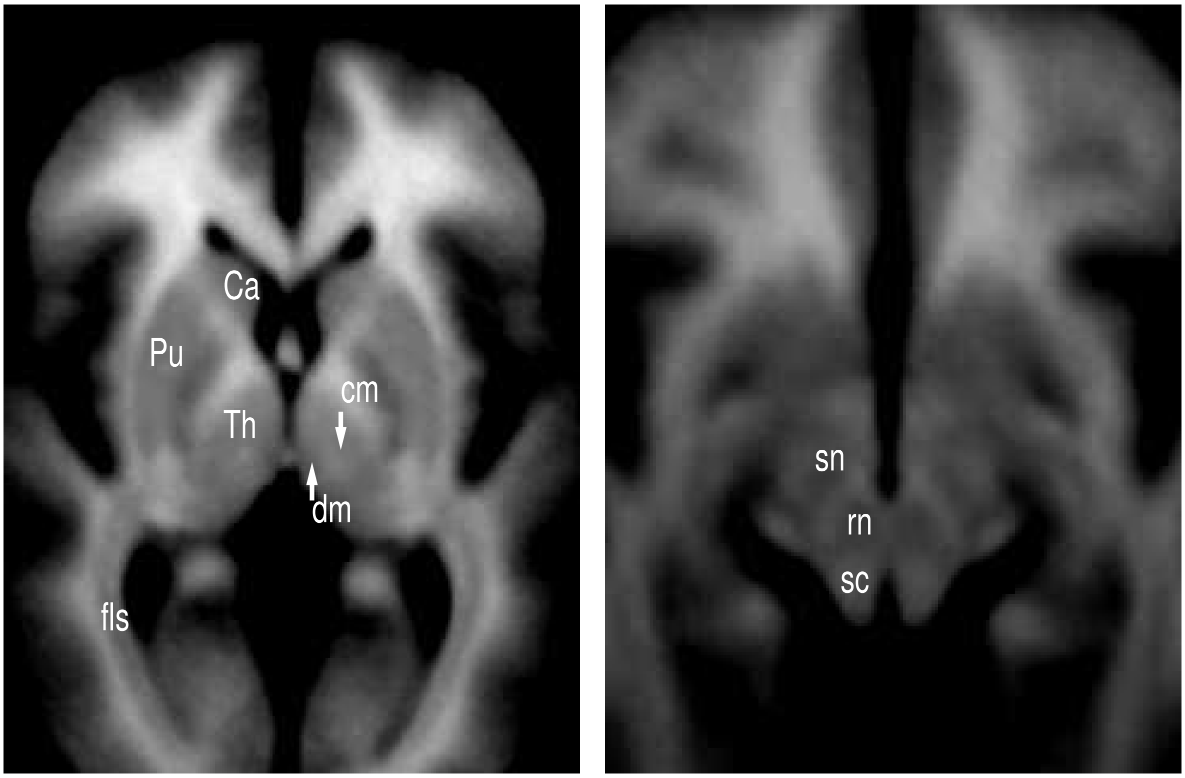

Figure 3.4: Close up of stereotaxic MRI atlas following Stage II averaging procedure.
Two transverse slices through the stereotaxic MRI atlas at z=0mm (left) and z=-10mm (right). These images show detail not normally apparent on a single MRI, but are brought out by noise reduction with image averaging. On the left can be seen both the centro-medial (cm) nucleus and dorso-medial (dm) nucleus of the thalamus. The connections between the putamen (Pu) and the thalamus (Th) can be made out as well. The increase in contrast-to-noise allows differentiation of white matter tracts in the occipital lobe such as the *fasiculus longitudinalis superior* (fls). The image on the right is positioned 1cm lower, and is enlarged by 25%, compared to the left. Here the red nucleus (rn), substantia nigra (sn) and the superior colliculus (sc) can be seen.

data sets fit each other quite well, however there exist some regions where the MRI atlas does not perfectly match the Talairach atlas. For example, the area of the inferior frontal lobe is a few millimeters lower on the Talairach than the mean MRI atlas, the cerebellar tentiorium is 1.4cm higher on the Talairach definition; on the coronal view, the temporal lobes extend lower by 7-8mm on the mean MRI atlas than on the Talairach atlas; the sylvian fissure is positioned slightly higher and more posteriorly in the Talairach atlas than the mean MRI atlas. These differences are due to normal anatomical variability in the Talairach brain. Even though the subject was considered normal, it does not represent the average anatomy, since it contains structures that are different from the mean brain created from over 300 subjects. These points must be kept in mind when using only the Talairach atlas to give a neuro-anatomical location for functional data obtained from other brains.

One could argue that the mean MRI-atlas does not fit the Talairach atlas because we did not use the complete 12 component piece-wise transformation specified by Talairach. While this is true, most groups (with the exception of Lemoine *et al.* (1991b)) use a single affine transformation to map functional data into stereotaxic space (*e.g.*, Fox *et al.*, 1985 and Friston *et al.*, 1989). Consequently, the limitations described above apply for most brain-mapping applications that use the Talairach atlas.

Ideally, this atlas should be defined in the minimum-variance frame described in section 3.2. Unfortunately, this is not possible to determine without estimates of neuro-anatomical variability over a large number of data sets. However, the model described here will be used to acquire this necessary information in order to define the "best" coordinate system for stereotaxy, in a data-driven fashion. Once defined, the MRI-atlas can be re-constructed in this frame, and all algorithms developed here re-applied with the new MRI-atlas. While this is beyond the scope of this thesis, the methods and tools developed here will serve to estimate the neuro-anatomical variability in a large number of subjects, and thus define the minimum-variance frame.

3.5 Summary

By establishing a standard anatomically-based coordinate system within the brain, Talairach formed a powerful concept that allowed external neuro-anatomical atlases to be used to predict the locations of specific structures within the brain. More importantly, the transformation to the stereotaxic space permits the analysis of anatomical variability in a well-define coordinate system.

In contrast to the single-brain atlas, the average MRI brain serves as a large n-brain, low resolution atlas, where the blurring in the atlas serves as a visual indicator of morphometric variability and a warning against over-interpretation of functional data in anatomical terms. This data base is a first step in the construction of a 3-D probability map of gross neuroanatomy across the normal population. The average brain volume may be used not only as an atlas, but this new MRI atlas replaces the previous one and serves as

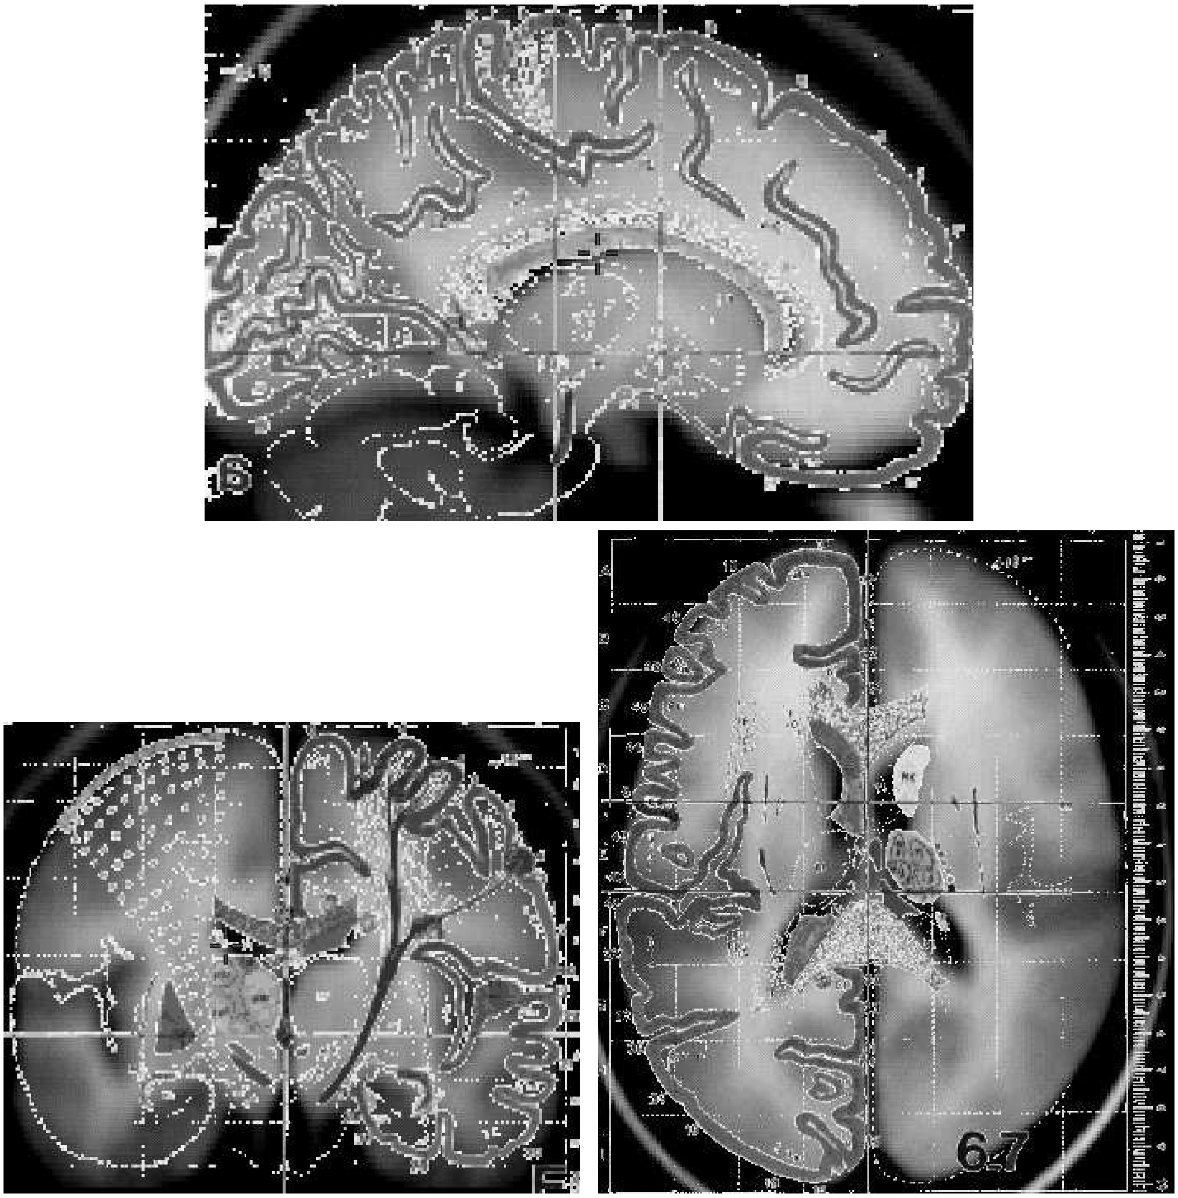

Figure 3.5: Mean MRI atlas with superposed Talairach atlas.

These three images show a sagittal (x=-11mm), coronal (y=15mm) and transverse slice (z=18mm) through the mean MRI atlas with the Talairach atlas overlaid on it. Even though the fit between the two atlases is good overall, there exist local differences due to normal non-linear neuro-anatomical morphological variability in the single brain used by Talairach.

a target volume for the automatic stereotaxic mapping procedure.

Chapter 4

Linear registration:theory and implementation

In both diagnostic and research applications, the interpretation of magnetic resonance (MR) images of the human brain is facilitated when different data sets can be compared by visual inspection between each other or to equivalent anatomical planes of an atlas. Quantitative analysis with pre-defined atlas templates often requires the initial alignment of atlas and image planes. Unfortunately, the axial planes acquired during separate scanning sessions are often different in their relative position and orientation, and these slices are not usually coplanar with those in the atlas.

This chapter describes the development of a completely automatic method to register a given volumetric data set to another based on multi-scale, three dimensional (3-D) cross-correlation. Once the data set is resampled by the transformation recovered by the algorithm, slices from the target volume can be directly super-imposed on the corresponding slices of the resampled volume allowing direct voxel-to-voxel comparisons. When the target volume is an atlas, the atlas slices can be overlaid on the resampled data, to be used as an aid in neuro-anatomical identification or as a basis for morphometric analysis.

4.1 Introduction

Many neurological studies of a single subject require the precise superposition of tomographic images from the same modality (*intra-subject intra-modality* registration). This may be necessary for clinical diagnosis, therapy planning, detection of anatomical or functional changes from previous scans and outcome evaluation. Intra-subject intra-modality registration can also be used to detect subject motion between frames during the course of a long PET acquisition. *Intra-subject inter-modality* registration is useful for integration of complementary information from different scanners. For example, the analysis of functional data from PET is enhanced when related to the underlying anatomy as defined by MR or CT where regions of interest defined on the anatomical scan are overlaid on the functional data. Both MR and CT can be used together for planning skull-base surgery where CT shows bony structures and MR soft tissues as in the paper by Hill *et al.* (1991). Applications of *inter-subject intra-modality* registration include the assessment of morphometric variability over a large number of subjects (Steinmetz and Seitz, 1991; Evans *et al.*, 1992a), or the analysis of subtle cognitive activation foci measured from cerebral blood flow (CBF) volumes derived from positron emission tomography (PET) within a standardized space (Fox *et al.*, 1985; Fox *et al.*, 1988; Evans *et al.*, 1992c)[1]. Direct measurement and comparison on a pixel-by-pixel basis is possible only when the data sets are properly registered.

Registration consists of determining the transformation that best maps corresponding features from one data set into another. The term *best* depends on the definition of a similarity criterion between two data sets and on the number of degrees of freedom allowed for the mapping function. Once found, the transformation can be used to resample the first data set in the coordinate space of the second, thus aligning them geometrically so that corresponding morphological features of both data sets are assigned to the same spatial location. In the context of this thesis, registration is the first step in the segmentation procedure, mapping features of a segmentation model onto the data set.

The proliferation of registration algorithms in the literature is due, in part, to the selection of different transformations (*e.g.* rigid body (Sibson, 1978; Marrett *et al.*, 1989a), piece-wise linear (Lemoine *et al.*, 1991b) or non-linear (Bookstein, 1989; Evans *et al.*, 1991b)), the matching of different features (*e.g.* principal axis (Alpert *et al.*, 1990), homologous points (Evans *et al.*, 1989a; Evans *et al.*, 1991a; Hill *et al.*, 1991), points to surfaces (Pelizzari *et al.*, 1989), homologous surfaces (Jiang *et al.*, 1992), corresponding high-curvature lines on surfaces (Gueziec and Ayache, 1992a)), the modalities registered (*e.g.* MRI-MRI (Collins *et al.*, 1994b), PET-PET (Minoshima, 1993; Friston *et al.*, 1991; Woods *et al.*, 1992), MRI-CT (Hill *et al.*, 1991), MRI-PET (Evans *et al.*, 1989a; Evans *et al.*, 1991a; Pelizzari *et al.*, 1989; Woods *et al.*, 1993a; Alpert *et al.*, 1990)), and the measure of the registration error that was used to determine the best match (see van den Elsen *et al.*, 1993 or Evans, 1993 for comprehensive reviews.) In general, intra-subject image matching is non-trivial because of different volume imaging parameters (*e.g.* slice thickness, pixel size, inter-slice gap, angulation, and resolution) and different patient positioning during the separate scanning sessions. The problem is even more difficult, both to define and to solve, for inter-subject registration because of dissimilarity in brain sizes and shapes.

The registration method presented in this chapter is applicable to both intra- and inter-subject intra-modality registration[2]. The algorithm uses an iterative optimization procedure to minimize the difference between image intensity features of the first data set and the second. The following sections describe 1) the transformation model used to

[1]Note that even though high resolution imaging techniques such as fMRI make it possible to measure activation within a single subject, it will still be necessary to compare results among individuals in order to fully understand normal cognitive function within a population.

[2]While the algorithm is also applicable to intra-subject inter-modality registration (such as MRI-PET), this aspect does not fall within the scope of anatomical segmentation as proscribed for this thesis, and thus will not be examined further.

map one data set on to another (§4.2) 2) the features extracted from the data (§4.3), 3) the similarity function used to measure the match (§4.4), and 4) the optimization method used to find the transformation model parameters (§4.5). After the general methodology is presented, the automatic stereotaxic mapping procedure is described in section 4.6. Practical aspects are considered and the procedure is compared to two other automatic registration procedures in section 4.7.

4.2 Transformation model

Either linear or non-linear transformations can be used to map the coordinate system of one data set into another. The linear case is addressed here, while non-linear transformations are dealt with in the following chapter.

Linear transformations for volumetric registration are either rigid or affine, where the former is a subset of the latter. A transformation is called rigid if the distance between any two points in the first volume is preserved in the transformed coordinate system, as is the case when moving a rigid body in space. This transformation can be decomposed into a translation component to center one data set on the other and a rotation component to align their orientations. Mirror-reflections can also be included in this definition of a rigid-transformation. However, for the practical task of matching 2 brain data sets, reflections are not used.

The distance constraint is dropped for affine transformations. However, straight lines in one data set continue to be mapped to straight lines in the other. Also, parallel lines remain parallel in the mapped volume. The affine transformation can be decomposed into translation and rotation as before, with the addition of scaling and shearing. Using matrix formulation and homogeneous coordinates, a point $(x, y, z, 1)$ is mapped to the point $(x', y', z', 1)$ using the homogeneous coordinate transformation:

$$(x', y', z', 1) = [\mathbf{A}](x, y, z, 1)'. \tag{4.1}$$

Where $\mathbf{A}$ is a 4×4 affine transformation matrix containing 12 independent elements and results from the concatenation of 4 matrices, representing translation, rotation, scaling and shear:

$$\mathbf{A} = [\mathbf{Sh}][\mathbf{Sc}][\mathbf{R}][\mathbf{T}]. \tag{4.2}$$

The matrices $\mathbf{T}$, $\mathbf{R}$, $\mathbf{Sc}$, $\mathbf{Sh}$ are defined as follows:

$$\mathbf{T} = \begin{bmatrix} 1 & 0 & 0 & tx \\ 0 & 1 & 0 & ty \\ 0 & 0 & 1 & tz \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{4.3}$$

where tx is the translation in x between the centroid of the two volumes in the 'real-world' coordinate system, while ty and tz are the corresponding translations in y and z.

$$\mathbf{R} = [\boldsymbol{\Theta}][\boldsymbol{\Phi}][\boldsymbol{\Psi}] \tag{4.4}$$

where

$$\mathbf{\Theta} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & cos\theta & sin\theta & 0 \\ 0 & -sin\theta & cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{4.5}$$

$$\mathbf{\Phi} = \begin{bmatrix} cos\phi & 0 & -sin\phi & 0 \\ 0 & 1 & 0 & 0 \\ sin\phi & 0 & cos\phi & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{4.6}$$

$$\mathbf{\Psi} = \begin{bmatrix} cos\psi & sin\psi & 0 & 0 \\ -sin\psi & cos\psi & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{4.7}$$

and the angles θ, ϕ and ψ are clockwise rotations around the $x-$, $y-$ and $z-$axes, respectively. The scaling matrix **S** is defined by:

$$\mathbf{Sc} = \begin{bmatrix} sx & 0 & 0 & 0 \\ 0 & sy & 0 & 0 \\ 0 & 0 & sz & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{4.8}$$

where sx, sy and sz are scaling factors along each of the axes.

$$\mathbf{Sh} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ a & 1 & 0 & 0 \\ b & c & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{4.9}$$

where a, b and c are the three free parameters left to define the shear matrix **Sh**. The shear parameters convert a rectangular parallelepiped into an oblique one.

In practice, the number of parameters (between 6 and 12) actually used to characterize the transformation is dependent on the type of registration required. For most cases of intra-subject registration, one can assume that the structure to be matched, the brain in this case, has not changed size or shape. Therefore a rigid-body transformation consisting of only translation and rotation is necessary to achieve registration. This transformation requires the estimation of only six parameters (three translations tx, ty and tz, and three rotations θ, ϕ and ψ).

More parameters are necessary to account for the increased degrees of freedom in inter-subject registration. In the simplest case, a total of seven transformation parameters would result from the addition of a single uniform scaling parameter to account for different brain sizes ($sx = sy = sz$). For most of the work described here, anisotropic orthogonal scaling along the coordinate system axes has been chosen (in the spirit of the Talairach coordinate system), yielding a total of nine parameters.

Other authors use all 12 parameters, including shear or, equivalently, non-orthogonal or off-axis scaling (Woods *et al.*, 1993a). In principle, the extra parameters should permit a better fit. Experiments in chapter 7 show that for a selected set of landmarks, in practice there is no significant decrease in the registration residual given the increased fitting power of the transformation with three additional parameters.

4.3 Feature detection and scale space

Most registration algorithms begin by extracting salient features from the image data that will be used in the evaluation of the matching process. This extraction is the goal of feature detection algorithms and is usually accomplished by convolution of an operator with the image data. Local peak values in the convolved array indicate the presence of structure similar to the desired feature and can be used to compare one image to another. There are a large number of ad hoc features-detectors described in the literature that could be chosen for evaluation in the similarity function. For example, edge enhancement was first accomplished with simple mask operators by Roberts (1965). The Sobel operator is slightly more complex, as it was designed to approximate the discrete gradient function in the vertical and horizontal directions (Duda and Hart, 1973; Levine, 1985). Bajcsy *et al.* has used Hueckel edge detection basis functions (Hueckel, 1973) in their papers published in 1983 and 1987, and then used Hermite polynomials in (Bajcsy and Kovacic, 1989). Until recently, there were no fundamental grounds to select one of these operators over another. However, there are some constraints on the volumetric features that limit the choice. The features must be insensitive to small perturbations in shape and robust to noise. Most importantly, the features detected must be intrinsic to the object, and dependent neither on the position nor on the orientation of the object within the image. For example, the feature detected at a given point on the object should have the same value if the object is moved or rotated in the image. In Fig. 4.1 the partial derivative in the horizontal direction for an ellipse is shown. When the ellipse is vertical, the detected feature is equal to zero at the superior or inferior point. However, on the rotated ellipse, the superior pole is almost entirely bright and the inferior pole dark. Thus, this feature is not intrinsic to the object, but dependent on the orientation of an object's edge with respect to the image coordinate system.

An additional constraint is imposed on the features used for registration. In order to minimize problems associated with local minima in the solution hyper-surface, the optimization algorithm used to find the registration parameters is performed in a hierarchical fashion so that gross structures are used to begin the registration, and smaller details are included in processing as needed to refine the fit. In this way, the possibility of the optimization procedure settling into a local minimum is lessened. Therefore, features must be extracted from the data in a multi-scale manner, where *scale* corresponds to its intuitive meaning of size. In order to observe an image at different scales, a set of spatially tuned operators is required. Recently, a construct termed *scale space* was developed by Koenderink and van Doorn (1987) and is defined "as a family of images generated in a

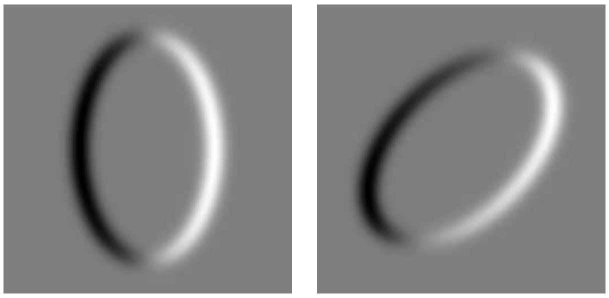

Figure 4.1: Partial derivatives of rotated object
Both images show the partial derivative, in the horizontal direction, of the same object rotated by 45 degrees. One can easily see that the upper and lower vertex of the ellipse does not have the same intensity signature, and thus does not have the same feature values, in the same region of the rotated object. Hence, this feature is not orientation invariant and cannot be used for registration where similar regions are matched together because they have similar feature values.

continuous manner from a given image by means of a convolution with a suitable spatial kernel of size determined by a one-dimensional parameter σ" (Llacer *et al.*, 1993). The absence of *a priori* geometrical knowledge imposes a number of constraints on the kernel used in the construction of the scale-space. The operator must be:

- **linear:** to allow for superposition of multiple input signals.
- **spatially shift invariant:** there must be no preferred position in the field of view for feature detection, thus allowing convolution of the operator with the image for detection of the feature.
- **rotationally invariant:** there must be no preferred orientations or directions for feature detection as there are for simple edge detectors such as Sobel or Hueckel operators that prefer horizontally or vertically oriented edges.
- **scale invariant:** there should be no preferred scale.

If the operator does not satisfy these constraints, then the feature detected at a particular point on the object will not necessarily have the same value when the object is rotated or moved in the image. Since object registration is achieved by maximizing feature similarity, it is imperative that corresponding regions of the two objects have similar values. Convolution with an isotropic Gaussian kernel satisfies these requirements, where the σ of the Gaussian operator is indicative of the spatial scale (tar Haar Romeny *et al.*, 1991; Koenderink and van Doorn, 1987). In this thesis, full-width-half-maximum (FWHM$= 2.36\sigma$) is used as the measure of spatial scale. Let

$$G(\vec{x}, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2}(\frac{\vec{x}}{\sigma})^2}, \tag{4.10}$$

be the Gaussian operator and let $L_0(\vec{x})$ be the 3-D image intensity function. Then, the convolution, represented by $*$, of L_0 with G is:

$$L(\vec{x}, \sigma) = \{L_0 * G(.;\sigma)\}(\vec{x}, \sigma). \tag{4.11}$$

Since the Gaussian operator is separable, we can write:

$$L(\vec{x}, \sigma) = \{((L_0 * G_x(.;\sigma)) * G_y(.;\sigma)) * G_z(.;\sigma)\}(\vec{x}, \sigma). \tag{4.12}$$

where G_x, G_y and G_z are the 1D Gaussian kernels along the x, y and z axes. This simplifies the processing of the data and is described below in the section on practical considerations.

Convolution with this operator is well-behaved. There is no introduction of sporadic detail with increasing blurring, i.e., no new details appear at a given scale if they did not exist at previous smaller scales. Unfortunately, convolution with the Gaussian kernel by itself is not enough for a complete local description of the image, since convolution with different values of σ will simply result in blurred copies of the original data. The local structure needed for both registration and segmentation can be extracted using differential operators. Scaled partial differentiation in a cartesian coordinate system is well defined to be the convolution of the original data with the corresponding partial derivative of the zeroth-order Gaussian at that scale:

$$L_{I_{i \ldots i_n}}(\vec{x}, \sigma) = \{L_0 * G_{I_{i \ldots i_n}}(.;\sigma)\}(\vec{x}, \sigma). \tag{4.13}$$

Fortunately, the differential Gaussian operators are well-behaved even at small scales; the spatial integration inherent in the convolution more than compensates for the noise amplification that usually accompanies differentiation of noisy images.

4.3.1 Invariant features used

Both registration and segmentation are dependent on *invariant* features (as described above). Two features were selected for evaluation in the similarity function. The first is the blurred intensity: a zeroth-order invariant feature. The value of $\partial L/\partial x$ is inherently dependent on the selection of the coordinate axes, and is therefore not an invariant feature. Fig. 4.1 shows the horizonal derivative for the same object, rotated by 45 degrees. The apex of the ellipse does not have the same intensity signature, because $\partial L/\partial x$ is orientation dependent. However, the magnitude of the 3-D intensity gradient is independent of the position and orientation of the coordinate axes and it is used as the first order invariant feature. This feature is defined as:

$$|\nabla| = \sqrt{(\partial L/\partial x)^2 + (\partial L/\partial y)^2 + (\partial L/\partial z)^2} \tag{4.14}$$

These two features, blurred intensity and gradient magnitude, can be calculated at every voxel in the 3D image volume and were found to be sufficient for linear registration. Figure

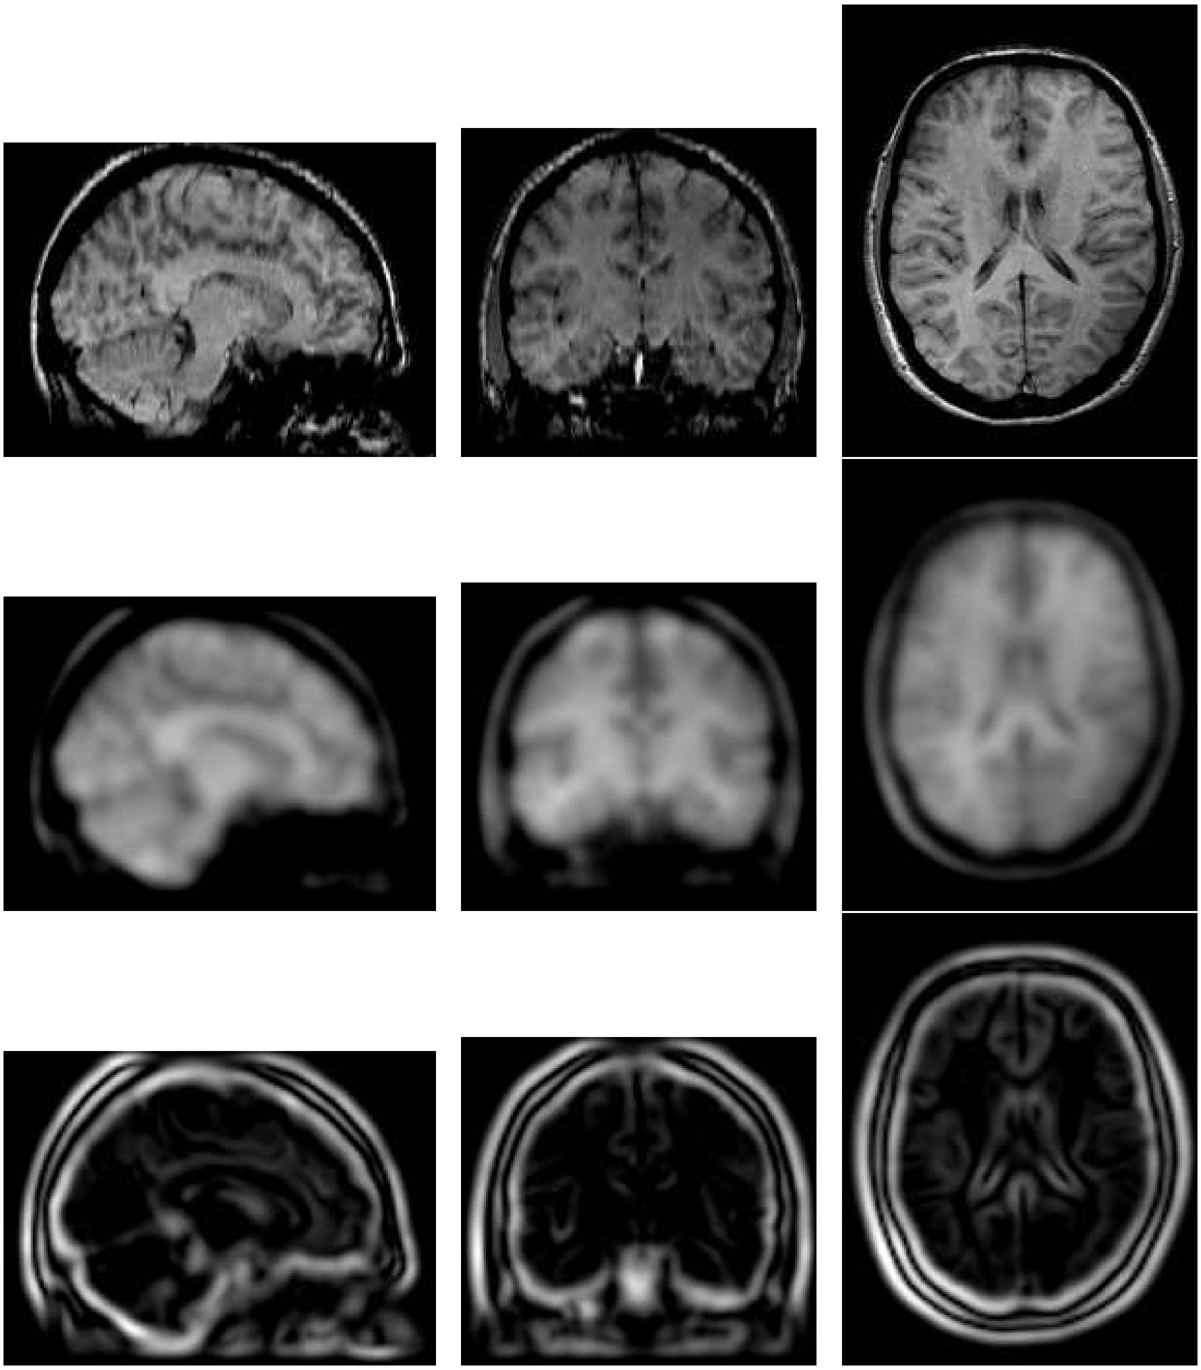

Figure 4.2: 8mm features used for matching
Tomographic planes through the volumetric features of a single data set at a scale of FWHM=8mm . On the top row from left to right are sagittal, coronal and transverse slices through the original data. The middle row shows the same slices blurred with a 3-D Gaussian kernel. The last row shows the corresponding gradient magnitude data.

4.2 shows slices through the original data, and the features extracted at the FWHM=8mm scale.

Since the Gaussian blurring is defined by the scale parameter σ, the operator is not scale invariant. Therefore, the features must be calculated anew for each scale of interest in the registration problem. This will be described in more detail below.

4.4 Similarity function

In the method presented here, registration is achieved by identifying the affine transformation $\mathbf{T}$ that maximizes the similarity between features derived from the voxels in the two data sets, $\mathcal{V}_s$ and $\mathcal{V}_m$, the volumetric data for the subject and the target model, respectively. $\mathbf{T}$ is the concatenation of 3 coordinate transformations: from subject-voxel-to-world in the subject, from world-to-world for registration, and from world-to-target-voxel to map a point from the subject into the target model volume:

$$\mathbf{T} = \mathbf{S}^{-1}\mathbf{A}\mathbf{M}, \tag{4.15}$$

where "world" corresponds to the real world coordinates of the source or target objects. Matrices $\mathbf{S}$ and $\mathbf{M}$ define the known world-to-voxel coordinate system mapping for the subject and model, respectively, and $\mathbf{A}$ is the required (or unknown) world-to-world affine registration matrix.

Like the other registration procedures, (*e.g.*, Rizzo *et al.* (1991)), correlation is used as a measure of goodness-of-fit between the transformed volume and the target volume. At a given scale step, the correlation value is evaluated on a set of voxel positions, $\mathcal{L}$, organized on a 3-D cubic lattice, with spacing between each lattice point equal to half of the FWHM distance of the current scale.

The spacing between lattice points used to sample the two data sets must be sufficiently small to avoid aliasing artefacts. These errors are especially evident when examining the shape of the correlation function in the neighbourhood of the correct transformation. When spacing is too large, the objective function is not smooth through the region of parameter space near the best transformation. Moreover, the function does not monotonically decrease towards the correct answer, but instead has large discontinuous jumps in value as one approaches the correct value for the given parameter.

Let $\vec{x}$ be an element of $\mathcal{L}$. The points used to calculate the cross-correlation are those that are within the volume of the subject's data set *and* that map through $\mathbf{T}$ into unmasked voxels in the model volume; i.e.,

$$\mathcal{L} = \{\vec{x} \mid \forall \vec{x}, \vec{x} \in \mathcal{V}_s \ \wedge \ \mathbf{T}\cdot\vec{x} \in \mathcal{V}_m\}. \tag{4.16}$$

The normalized cross-correlation value, R, between the two volumes, $\mathcal{V}_s$ and $\mathcal{V}_m$ for a given transformation $\mathbf{T}$ is defined:

$$R(\mathcal{V}_s, \mathcal{V}_m; \mathbf{T}) = \frac{\sum_{\vec{x}\in\mathcal{L}} f(\mathcal{V}_s,\vec{x}) f(\mathcal{V}_m, \mathbf{T}\cdot\vec{x})}{(\sum_{\vec{x}\in\mathcal{L}} f^2(\mathcal{V}_s,\vec{x}))^{\frac{1}{2}} (\sum_{\vec{x}\in\mathcal{L}} f^2(\mathcal{V}_m,\mathbf{T}\cdot\vec{x}))^{\frac{1}{2}}}, \tag{4.17}$$

where $f(\mathcal{V}, \vec{x})$ is the interpolated feature value from the volume $\mathcal{V}$ at voxel position $\vec{x}$, and the summation is done over all elements $\vec{x} \in \mathcal{L}$. R takes on a maximum value of 1.0 when the two volumes are in perfect registration. In the registration algorithm, this function is evaluated at each step of the optimization procedure described in the following section.

4.5 Optimization

The 3-D image registration task is set up as an optimization problem to identify the required transformation while minimizing the computational complexity usually associated with an exhaustive grid search approach. The optimization is performed at different spatial resolutions to minimize the problems associated with local minima in the solution hyper-surface, starting with very blurred data and increasing detail at each step by using less blurred images. This is similar to pyramid techniques where the optimal solution for the match at one level is passed down to the next (more complex) level (*e.g.* Pizer *et al.*, 1988, Kalvin and Peleg, 1989, Sokolowska and Newell, 1986). The following sections describe the optimization at one scale step (§4.5.1), the starting conditions (§4.5.2), the selection of scale steps used in the optimization (§4.5.3) and registration strategy (§4.5.4).

4.5.1 One scale step

At each step in decreasing scale, the optimization procedure refines the transformation parameters for registration of the source volume onto the target. A simplex multi-dimensional optimization procedure is used to maximize the normalized cross-correlation value (eq. 4.17) of the blurred intensity volume (the zeroth-order invariant feature) at the current scale. The resampling process used to calculate the correlation value employs tri-linear interpolation on super-sampled data to minimize sampling artifacts (the volumes have an imposed Gaussian blur, with a FWHM larger than the 2 mm -sized voxels used to represent the volumetric features).

The registration procedure can provide up to 12 updated resampling parameters (three scales, three rotations, three translations and three shears) that are used to optimize **A**, and thus optimize **T** in eq. 4.15. There are two benefits obtained by posing the optimization in terms of the actual transformation parameters instead of in terms of the elements of the transform matrix. The first is that the type of transformation can be specified easily, and only the number of parameters required for its definition need be optimized, thus achieving a savings of CPU time and reducing the possibility of error due to the more complex hypersurface of the objective function of all 12 parameters of an affine transform. The second benefit is that the transformation parameters are nearly orthogonal in the multi-dimensional optimization space when the center of gravity of the transformed volume is used as the center of rotation and scaling, *e.g.*, a small change in the rotation parameters does not require a complete re-optimization along the scaling or translation axes.

4.5.2 Starting conditions

The optimization procedure requires an initial estimate of the transformation parameters that are to be optimized. The better the estimate, the faster the optimization procedure will converge. In order to maintain the fully automated characteristic of the registration procedure, a principal axis transform, similar to that described by Alpert *et al.* (1990), is used to determine the starting parameters for the optimization. This automatic procedure eliminates the need for any user intervention.

A weighted principal-axis method is used to calculate the covariance matrix and the center of gravity of the blurred intensity feature volumes that are used in the first scale-step optimization. The rotation angles are extracted from the principal axis vectors in order to establish the approximate transformation parameters. In cases of partial volume coverage (e.g, where the scanned volume covers only a part of the brain), only the centers of gravity are used to establish the initial transformation and the rotation parameters are set to zero since they cannot be reliably extracted from the covariance matrix.

4.5.3 Scale selection

A discrete image contains features with a limited range of scales. The smallest or *inner* scale is bounded by the pixel size and imaging point spread function, and defines the finest level of detail (or highest frequencies) that can be represented in the image. Similarly, the largest or *outer* scale is bounded by the field of view and limits the size of gross structures (or lowest frequencies). When extracting characteristic features from data, the proper scale for measurement must be determined, since each feature is defined by an inherent scale. Traditionally, a factor of 2 is used to step through scale space, starting at the original pixel size and sampling at 2, 4, 8, 16 and 32 times the original pixel size.

Which of these scales should be used to drive the registration process? With a typical inplane pixel size of 1mm, the resolution of the MRI image is approximately 2.3mm as measured by the FWHM of a Gaussian approximation to the imaging point spread function. Slice thickness and spacing determine the resolution out-of plane, yielding out-of-plane resolutions of 4.7 to 7mm for the data used in this thesis. These values fix the lower limit for the inner scale.

Data blurred to have isotropic resolution of 32, 16 and 8 mm were used sequentially to calculate the multi-scale registration. We found that the most blurred data could be eliminated from the procedure without penalty. Therefore, data blurred to the FWHM=16 mm scale is use to start the registration with the principal axis transform. The first optimization process used the same data. When the optimal solution is obtained for the FWHM=16 mm scale, it is used as a starting point for the same process calculated with the blurred intensity feature volumes at FWHM=8 mm scale ($\sigma = 3.4$ mm).

4.5.4 Registration strategy

Other authors (*e.g.*, Pratt, 1978, p557) have remarked that correlation of blurred intensity signals provides relatively poor discrimination between objects of different shapes, but of similar size or energy content. This has also been found to be true here, since ending the optimization at the 8mm FWHM scale shows a tendency for the angulation about the $z-$axis to be biased by approximately 4 degrees, misaligning the longitudinal fissure in the transverse plane and the $z-$scale was over-estimated by as much as 20% in some cases. This over-estimation occured when the source volume was smaller in the z-direction than the target. The optimization procedure compensated for the missing data by stretching the source volume in the z-direction.

A plot of the objective function verses parameter value for correlation of blurred intensity volumes shows a very flat peak in the region of the correct value for the transformation parameter (see Fig. 4.3). The flatness makes it difficult for the optimization procedure to quickly reach the correct answer. In comparison, the peak based on correlation of the gradient magnitude at the same scale is much sharper, making it easier for the optimization to arrive at the right answer.

With this in mind, one more refinement was added to the optimization process. The result of the FWHM=8mm blurred intensity fit is used as the starting point for optimization using *gradient magnitude* volume (first-order invariant feature) at the same scale. Since this feature emphasizes boundaries between tissue types and is less sensitive to low frequency variations in absolute intensities, it makes the algorithm more robust when dealing with data sets where image sensitivity is not constant across the entire object due to radio-frequency inhomogeneity during the image acquisition.

In summary of the registration process, the optimal solution obtained at each step is used as input to the next, where it is refined. Both intensity and gradient magnitude features are used:

1. calculation of image features (for FWHM=16 and FWHM=8mm).

2. principal axis transformation on FWHM=16 mm intensity volumes.

3. optimization on FWHM=16 mm intensity volumes.

4. optimization on FWHM=8 mm intensity volumes.

5. optimization on FWHM=8 mm gradient magnitude volumes.

The linear transformation does not become significantly more accurate by continuing to FWHM=4 or 2 mm scale. This is due to non-linear morphometric variability which limits the precision of the transformation recovery to that found at 8mm.

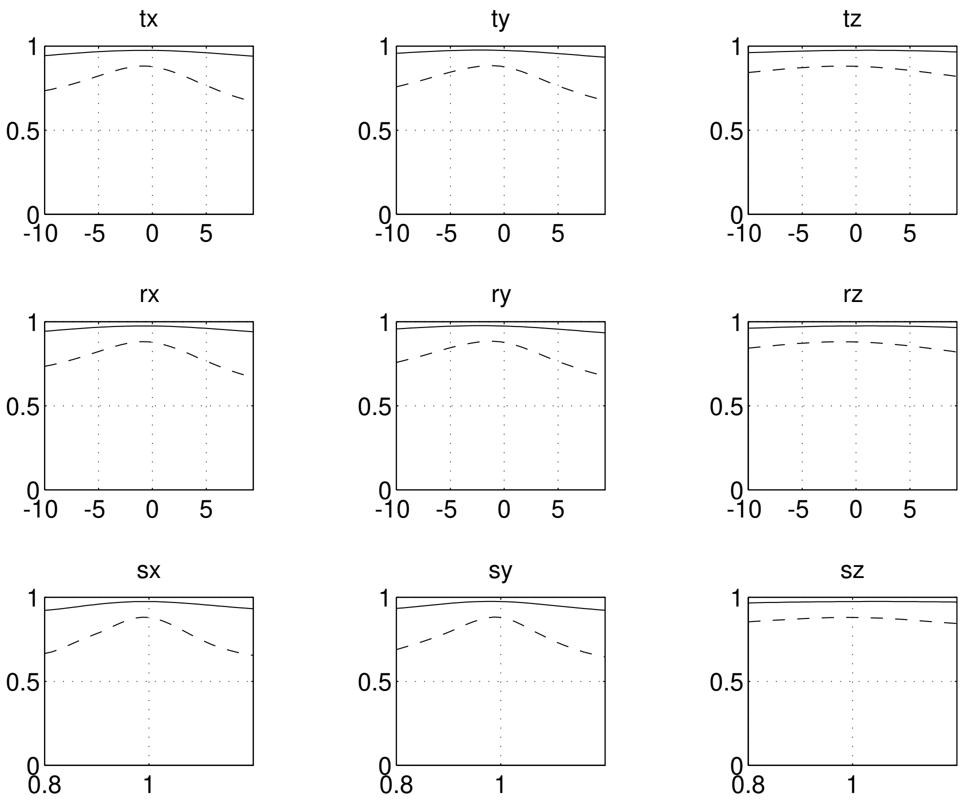

Figure 4.3: Comparison of image intensity vs gradient-based fitting.

These graphs show the correlation value (vertical axis) plotted against parameter error for blurred intensity (solid line) and gradient magnitude (dotted line). Seventeen data sets were deliberately mis-registered, varying one parameter at a time. The objective function based on blurred intensity shows almost no deviation

4.6 Automatic stereotaxic transformation

As noted in the introduction of this chapter, the use of a stereotaxic coordinate space addresses a problem common to studies that involve inter-subject comparisons between large numbers of data sets, where it is necessary to compare many subjects, regardless of the size, position or orientation of the original volumetric data set. This need is most evident in the assessment of morphometric variability over a large number of subjects (Evans *et al.*, 1992a; Steinmetz and Seitz, 1991), or the analysis of subtle cognitive activation foci measured from cerebral blood flow (CBF) volumes derived from positron emission tomography (PET) (Fox *et al.*, 1985; Evans *et al.*, 1992c). The registration technique developed in this chapter may be used to automate the stereotaxic transformation procedure, however two additional steps are required to ensure robust estimation of the transformation:

1) Ideally, the matching procedure calculates the correlation of similar features between the two data sets to determine the registration transformation. Since the resolution of the target MRI brain model is lower than an individual MRI volume because of the inherent blurring in the average due to morphometric variability, the individual data sets are pre-blurred using a 3-D smoothing kernel. A 6mm FWHM isotropic Gaussian blurring kernel was used, because initial manual estimates of anatomical variability indicated values of 6-7mm (Evans *et al.*, 1991b).

2) The existing correlation technique identified the best transformation for brain plus scalp, which will sometimes bias the overall scale in the resulting transformation when a subject has a thicker or thinner skull than the average or other errors introduced by bright signals from fat in the scalp in some patients. Since it is the optimal overlap of corresponding brain voxels that is of interest here, a brain mask was manually defined in stereotaxic space and added one more step to the multi-resolution correlation procedure. The best fit at the 8 mm scale is used as the starting point for another optimization procedure calculated at the same scale, but using only voxels that fall within the brain mask.

In summary of the registration process, the optimal solution obtained at each step is used as input to the next, where it is refined:

1. pre-blur data with FWHM=6.0mm ($\sigma = 2.55$mm) kernel.
2. extract features from data with kernels of FWHM=16 and 8mm .
3. principal axis transformation on FWHM=16mm intensity volumes.
4. optimization on FWHM=16mm intensity volumes.
5. optimization on FWHM=8mm intensity volumes.
6. optimization on FWHM=8mm intensity volumes with brain mask.

7. optimization on FWHM=8mm gradient magnitude volumes with brain mask.

The experiments in chapter 7 show that this completely objective automatic technique has accuracies similar to that of landmark-based registration, or surface to surface matching techniques, while having the additional advantage that explicit manual intervention is not required to identify such points or contours.

4.7 Discussion

4.7.1 Practical considerations

Blurring of the intensity volume is achieved by 3-D convolution of the data volume with a 3-D Gaussian kernel. There are a number of methods that have been applied to calculate the blurred intensity and gradient magnitude features for a given scale such as recursive linear filtering (Monga and Deriche, 1989), 3-D spline interpolation, or local Gaussian filtering using a small finite kernel (such as a $7 \times 7 \times 7$ convolution kernel). Since non-zero order differentials are very sensitive to noise, we have decided not to blur the data by convolution with a small finite kernel (say 5^3 or 7^3 voxels) or by spline interpolation using a small area of support. Instead, multiplication in the Fourier domain is used as the transform equivalent of convolution. A 3-D FFT algorithm (Press *et al.*, 1988) is applied to the data volume which is then multiplied, voxel-by-voxel, with the 3-D FT of a Gaussian kernel centered at the origin of FT space. An inverse FFT is applied to the product, resulting in a blurred intensity volume. The partial derivative volumes are also calculated in the Fourier domain, multiplying corresponding data rows of the FT of the blurred data by $-2\pi i$ before taking the inverse transform. Since blurring and differentiation are completed in the Fourier domain, the long tails of both the Gaussian and derivative kernels can be maintained without additional computational cost as would be the case if the convolution were done in the spatial domain. Also, ringing artefacts are avoided since neither kernel is truncated to be smaller than the number of samples in the data volumes.

The original MRI data does not have isotropic resolution prior to convolution, usually being poorer in the z-direction. Therefore, in theory, a 2-D in-plane Gaussian blurring should be applied to restore isotropy before convolution with the 3-D kernel. However, this is a moot issue, since the blurrings add in quadrature and the out-of-plane resolution is still smaller than the kernel sizes used here. In future, as the resolution of acquired data improves, this will no longer be a problem.

4.7.2 Multi-resolution methodology

There are two benefits achieved using the multi-step, multi-resolution strategy. 1) The multi-step process lessens the likelihood of settling into a local minimum in the optimization process. An approximate solution is found at the first step and it is refined in each sequential step. 2) The multi-step multi-resolution approach offers a substantial computational savings. While there was an eight-fold increase in time required to calculate the correlation coefficient when passing from one scale level to the next, the linear transformation can be recovered using the multi-resolution method in less than one quarter the time taken when starting at the highest resolution level (where the solution found was not always correct because of local minima). At its current un-optimized state, the routine requires approximately 30 minutes on an SGI Indigo^2 Extreme, a 60.5 SpecFP Unix machine, to calculate the features and register a specific volume into stereotaxic space.

4.7.3 Relation to other registration techniques

Other research studies of 3-D intra-modal registration have employed different objective functions and matching strategies. In this section the work of two algorithms of direct relevance to the present work are discussed in detail. The work of Woods *et al.* (1992, 1993b, 1993a) and that of Minoshima *et al.* (1992) is described and then compared to the methods developed here on the basis of 1) the features used, 2) in the objective function, 3) the parameters used in the fitting model and 4) the optimization procedure used to calculate the result.

Woods *et al.* (1992) developed an algorithm that achieves six-parameter (three translations, three rotations) registration of PET data by minimizing the variance of the ratio of one image to the other on a voxel-by-voxel basis. This method assumes that voxel values in the two data sets are related by a single multiplicative factor when registered. A threshold is used to eliminate zero-division errors and to restrict voxels considered to those within the brain volume. It is not clear how the level of this threshold affects the sensitivity of the matching.

An innovative modification to the algorithm allows PET-MRI registration, where pixels are grouped based on MRI intensity value, and the variance is calculated for each group and the weighted sum of the group variances is minimized (Woods *et al.*, 1993a). However, non-trivial manual intervention is required to edit and remove the non-brain regions from the MRI data, rendering the procedure semi-automatic at present. Inter-subject MRI-MRI registration has been accomplished by increasing the number of transformation parameters from 6 to 12 (Woods *et al.*, 1993a).

Like others, (Kapouleas and Kulikowski, 1988; Brummer *et al.*, 1991; Allain *et al.*, 1992), Minoshima *et al.* (1992) have designed a two stage registration process that begins by identifying the inter-hemispheric plane. The procedure finds the plane, defined by two rotation angles and a lateral translation, that maximizes the number of sign changes (or zero-crossings) in the subtraction of the mirrored left side from the right. In a second step, the volume is aligned into stereotaxic space using empirical rules to automatically locate four points along the AC-PC line (Minoshima, 1993). These rules are dependent on the PET tracer distribution, but are appropriate for CBF images.

Features: Both Woods and Minoshima calculate their objective functions based directly on the intensity values in the original image. The technique presented here calculates invariant features from the images using both blurred image intensity and gradient magnitude. Our routine has been found to be more robust using these two features than when using intensity alone when dealing with partial volume coverage or with data sets affected by radio-frequency (RF) inhomogeneity artefacts (see experiments in chapter 7.

Objective function: The technique of Woods minimizes the variance of the voxel-to-voxel ratio. That of Minoshima maximizes the number of zero-crossings and this method maximizes the correlation value.

Fitting parameters: The method presented here, and that of Minoshima, use a 9 parameter transformation (3 translations, 3 rotations and 3 scales) to map a brain into stereotaxic space. Woods *et al.* use a 12 parameter affine fit, where three more parameters allow off-axis scaling. In principle, the extra parameters should permit a better fit, unless the noise characteristics negate the advantage of the extra parameters.

The objective function used here does not require the difference in overall image intensity to be modeled as an additional separate parameter, thus simplifying the optimization. This feature is similar to the approach of Woods *et al.* , but differs from that of Minoshima and from other approaches that attempt to minimize the summed absolute differences in pixel intensities.

Optimization procedure: Minoshima *et al.* use a dual-level grid search technique. A global coarse search is used first and is followed by a finer step adjustment of the parameters in the neighbourhood of the initial transformation. Woods *et al.* first used a derivative-based technique where the parameter with the largest magnitude first derivative is adjusted at each iteration (Woods *et al.*, 1992; Woods *et al.*, 1993a). In their newer method (Woods *et al.*, 1993b), new values are calculated for all parameters on each iteration. Woods also uses a hierarchical technique, starting with data samples spaced at every eighty-first voxel within the volume. After the initial convergence, the resulting parameters are used to restart the algorithm sampling every twenty-seventh voxel, then every ninth voxel, every third voxel and finally every voxel of the data set. Our optimization is achieved in a hierarchical multi-scale, multi-resolution strategy. Using recursive refinement, the optimization procedure is applied to data at different resolutions, beginning with very blurred data. The result is used as a starting point for a new optimization using less blurred data.

While all methods use multi-level search techniques, the two former methods employ sub-sampling to define coarser grids but do not impose blurring kernels on the original data, and therefore do not use a well-behaved (linear, invariant (tar Haar Romeny *et al.*, 1991)) multi-scale approach as described previously. Noise in the data may cause spurious detail within the sub-samples, adversely affecting the optimization routine. However, the inherent blur of the PET data may keep these algorithms from falling into local minima. The higher spatial frequencies of MRI data, both real and noise-related, cause many local minima in the objective function hypersurface. Hence, the global minima in the objective function for MRI-MRI correlation may be missed if a true multi-scale approach is not applied in the optimization process.

4.8 Summary

In this chapter, an objective and automatic procedure to register two or more volumetric data sets together has been described. The experiments presented in chapter 7 show that 3-D multi-resolution correlation to be an effective tool to determine the affine component of transformation between two data sets. While the technique has benefits and accuracies similar to those of landmark-based registration or surface-to-surface matching techniques, it does not require explicit manual intervention to identify such points or contours.

The registration method developed here can be applied to automate the linear stereotaxic mapping procedure and it can be applied retrospectively since it relies on automatically detected edges in the data that correspond to internal anatomical structures rather on than fiducial markers. As well as characterizing the stereotaxic space, the use of the standardized model provides information on the location of specific structures, obviating the need to manually identify, edit and remove scalp, skull or meninges from the MRI volumes.

Currently, in many registration methods, the AC-PC line is often explicitly identified, based on the technique originally developed by Talairach to localize structures in the basal ganglia and brain stem. However, cortical landmarks may be more important for many brain mapping projects. No assumptions are made here about the AC-PC line and its relation to the brain, nor are assumptions needed regarding left-right symmetry. The approach registers a given data set with stereotaxic space using a generic 3-D fitting routine, matching all structures in the brain simultaneously, and in doing so, does not depend fundamentally on the AC-PC points. The 3-D masking incorporated in the routine can easily be used to limit the data fitted to the regions of interest. For example, either occipital lobe structures or peri-ventricular structures could be masked and fitted after an initial transformation using the whole brain volume.

This technique can be readily extended to other modalities. Other features can be added to the volumetric information stored in the model using the same methodology as that described above for MRI. For example, the average intensity and gradient magnitude values derived from positron emitted tomography (PET) cerebral blood flow (CBF) data can be stored in the model and used to register against other PET CBF studies. This achieves two goals. The first permits cross-registration of PET volume data sets. The second provides a procedure to automatically register data sets from different modalities. Since MRI and PET volumes can be independently registered to the same target space, the transformation between the two volumes is known implicitly. The capacity to compare and correlate brain maps across modalities and individuals in an objective, reproducible fashion will greatly enhance our understanding of normal and pathological brain states.

Chapter 5

Non-linear registration:theory and implementation

Registration using only an affine transformation cannot account for non-linear morphometric variability between subjects. This has profound implications for applications employing stereotaxic systems which map individual brains into a common reference frame: quantitative neuro-radiology, stereotactic neurosurgery and cognitive mapping of normal brain functions with PET where point localization with an accuracy of less than 5mm can mean failure of the process. In this chapter, an automatic procedure is presented to address these issues by generating the necessary non-linear transformation in the form of a spatial warping field, where a 3-D deformation vector is stored for each voxel in the field. The recovery of this global non-linear warp is based on the recursive estimation of local linear deformations, again using cross-correlation of invariant features derived from image data. Not only does this method serve to accomplish non-linear registration between data sets, it serves as the basis for the segmentation procedures described in the next chapter.

5.1 Introduction

5.1.1 Background

New imaging modalities and techniques, *e.g.* PET, SPECT, functional MRI (fMRI), magnetoencephalography (MEG), and electroencephalography (EEG) have made it possible to map functional areas of the human brain with respect to anatomy. Two aspects of this work require integration of data between individuals: 1) The low signal associated with cognitive activation (*e.g.* a subtle change in cerebral blood flow as measured by PET) requires averaging between subjects to improve statistical significance of measured changes (Fox *et al.*, 1988; Evans *et al.*, 1992c). 2) High resolution imaging techniques such as fMRI make it possible to measure activation within a single subject. However, it will still be necessary to compare results between individuals in order to fully understand the normal anatomical variability in the location of focal sites of physiological responses underlying normal cognitive operations. Both averaging and comparison require precise point-to-point correspondence between brains, and each has typically been accomplished by mapping the volumetric data into a stereotaxic brain-based coordinate system (Talairach and Tournoux, 1988). Until recently, most centers have used linear transformations only (Fox *et al.*, 1985; Friston *et al.*, 1989; Minoshima *et al.*, 1992; Woods *et al.*, 1992). However, previous work (Talairach and Tournoux, 1988; Steinmetz and Seitz, 1991) has shown that even after linear mapping, there is variability of up to 1.5 cm in the position of cortical structures, which represented a significant source of error when mapping activation foci. We have previously shown (Evans *et al.*, 1991b), that on average for points throughout the brain (cortical and sub-cortical), there is a 6-7mm anatomical variability (defined as the 3-D r.m.s. distance between a given landmark and its homologue in the target volume) in position not accounted for by linear registration.

5.1.2 Issues

In order to account for these morphological differences in normal anatomy, a registration procedure that allows more degrees of freedom than a simple linear affine transformation must be used. Stated simply, the problem is this: *Given two volumetric data sets of the human brain, find the non-linear spatial distortion function to map all points of the first brain to their homologues in the second.* Implicit in this statement is the assumption that homologous points exist, that they can be identified reliably, and that there exists a one-to-one mapping between all points of the first brain and their homologs in the target. Hence, non-linear registration is much more complex than it first appears.

Depending on the technique used to find the spatial distortion function, homologous points are often required either to establish or to evaluate it. The points selected to create homologous sets are those that serve as anatomical landmarks and usually have biologically meaningful labels. These landmarks must be defined with respect to their local neighbourhood, i.e., the intersection of a line with a plane (*e.g.*, anterior commissure on mid-sagittal plane), the point of contact between two structures (e.g, inter-thalamic adhesion), a point of extreme curvature of a surface (*e.g.*, occipital pole) or the center of a structure (*e.g.*, center of head of caudate). Many of these definitions are idealized, and cannot be realized in practice without additional constraints. For example, the definition of a landmark at the inter-thalamic adhesion must be refined to include a notion of centroid, since the left and right thalami are not connected at a single point. Talairach does not actually define the origin of the stereotaxic coordinate system on the anterior commissure, but at the intersection of two lines: one passing horizontally through the superior aspect of the AC and the other vertically through the posterior aspect of the AC. Thus, when identifying the 3-D coordinate of a landmark, a small difference in interpretation of its definition may cause a positional error of 1-2mm between observers (*inter-observer error*). Repeatability trials yield another source of error that stems from the inability of a single observer to identify exactly the same landmark in a given brain at different times (*intra-observer error*).

The difficulty in reliable identification of homologous points is far more complex when dealing with multiple subjects because of normal natural morphometric variability between brains (*inter-subject variability*), where there may not exist a complete one-to-one correspondence. It is this variability that is of interest for morphometric analysis and must be quantified in stereotaxy. The combination of the two observer-dependent errors with the inter-subject variability renders exact homology between brains unattainable in both practice and in theory. Nevertheless, an engineering approach allows us to refine our methodology continuously and to overcome observer errors and to make useful inferences about neuro-anatomical variability during the pursuit of the idealized mapping.

The use of different brains also confounds the notion of a one-to-one mapping function. While there may exist a one-to-one transformation for a given set of landmarks, it is not necessarily unique for that set, nor need it be the same for a different set of landmarks defined on the same two brains. There may not even exist a one-to-one mapping for particular structures, *e.g.*, the para-cingulate sulcus is more prominent on the left than the right hemisphere, being absent twice as often on the right side (Paus *et al.*, 1994). Since the structure may not exist, how does one determine the one-to-one mapping for that region? These facts lead to the conclusion that no completely objective, general, direct measurement of non-linear registration quality is attainable. However, a task-dependent measure may be defined with respect to the given operation requiring higher-order warping for registration. Given that a general theoretical approach to the non-linear registration problem appears impracticable, and based on the success of the linear registration algorithm described in the previous chapter, a similar empirical strategy was undertaken for the development of the non-linear registration strategy to address the misregistration remaining after linear alignment between data sets. The experimental results described in chapters 8 and 9 demonstrate the practical utility of this approach.

5.1.3 Proposed method

Non-linear warping using non-parametric models has been proposed as a means to characterize the transformation that most smoothly maps one data set into the second. The goal was to develop a technique that reduced the registration residual due to anatomical variability, despite the difficulties described above. Both automatic and manual warping models are based on the assumptions that different brains are topologically equivalent, and that the application of non-linear deformations to one data set can bring it into correspondence with another or with a model.

One existing technique applied to the problem was the thin-plate spline, a manual non-linear warping method based on homologous landmarks (Duchon, 1976; Bookstein, 1989; Evans *et al.*, 1991b). Its main advantage is that achieves exact matching of the landmark points and yields a continuous spatial deformation function that is decomposed into orthogonal warps of decreasing energy. Since landmarks must be identified, explicit correspondence is available at a fixed number of points. However, because the procedure yields an exact fit at the landmarks, errors in their identification are not damped, but are propagated over a neighbourhood around the error. Even with modification of the algorithm to produce a least-squares fit, the method still depends on user intervention and thus is time consuming, requiring more than one hour of a neuro-anatomical expert's time to identify 40 landmarks, even with real-time 3-D display software. Finally, the subjectivity involved in selecting the precise location and the number of points that will define the non-linear deformation does not make these methods practical for routine use as a deformation/warping model.

This has lead to the consideration of fully automatic, non-parametric, objective, non-linear mapping techniques (Bajcsy *et al.*, 1983; Friston *et al.*, 1991; Minoshima, 1993; Collins *et al.*, 1992c; Collins *et al.*, 1992d; Collins *et al.*, 1994a). These alternative approaches avoid the potentially time-consuming landmark tagging by replacing it with automatic feature-matching. Instead, correspondence is estimated using a similarity criterion and constrained with an elastic model, making the procedure tolerant of errors in local point correspondence.

The non-linear registration technique described here is a straightforward extension of the linear registration method of the previous chapter, except that the target volume is reduced to a small neighbourhood of the whole brain, recursively selected by stepping through the entire target volume in a 3-D grid pattern, and the local transformation is limited to 3 translations instead of a full 12 parameter transformation. The global non-linear warp needed to map the data set onto the model is composed of a set of local deformations that are derived from the 3 translations, with one deformation vector being defined for each voxel position. Each deformation is based on the estimation of the best local transformation that maximizes the correlation of features from the neighbourhood surrounding the given voxel to its corresponding neighbourhood in the target data set.

The following inter-related sections define 1) the representation of the non-linear transformation, 2) features used in the matching process, 3) the strategy used to recover the global warp from the source and target MRI data sets, and 4) the notion of correspondence. The chapter concludes with a section comparing the method developed here with two related, previously published techniques.

5.2 Representation of the non-linear transformation

5.2.1 Deformation function

Using the same notation as the previous chapter, let $f(\mathcal{V}_s, \vec{x})$ represent the intensity of source volume $\mathcal{V}_s$ at position $\vec{x}$. Let $\mathcal{V}_s$ be transformed into $\mathcal{V}_d$ so that each point in $\mathcal{V}_d$ matches its homologue in the target model $\mathcal{V}_m$. The transformation from source to target space is affected by resampling the source volume by the non-linear spatial warping function, $\mathbf{N} : \Re^3 \rightarrow \Re^3$.

$$f(\mathcal{V}_d, \mathbf{N}(\vec{x})) = f(\mathcal{V}_s, \vec{x}). \tag{5.1}$$

The mapping $\mathbf{N}$ is divided into two components: one affine and one non-linear. (See Fig. 5.1). The affine component, $\mathbf{A}$, is represented by a homogeneous matrix and is recovered using the method described in the previous chapter. After removal of this component, the residual differences between $\mathcal{V}_s$ and $\mathcal{V}_d$ are due to morphometric differences between source and target data sets. The goal of the procedure is to estimate the non-linear component, $\mathbf{D}$, representing the additional deformation required to match the transformed source point (i.e., after affine transformation) to its homologue in the target volume:

$$\mathbf{N}(\vec{x}) = \mathbf{A} \cdot \vec{x} + \mathbf{D}(\mathbf{A} \cdot \vec{x}), \tag{5.2}$$

where $\mathbf{N}$ represents the world-to-world transformation from subject to model. Equation 4.15 is rewritten to include this term such that $\mathbf{T_{NL}} = \mathbf{S}^{-1}\mathbf{NM}$ now becomes the total non-linear transformation from the subject voxel space to the target voxel space. We define a deformation function in the following manner:

> A deformation function is a continuous, invertible mapping from $\Re^3$ to $\Re^3$ that accounts for the residual mismatch between two volumes after linear registration.

For illustration, a 1D profile from an example 3-D deformation function is shown in Fig. 5.2, where t steps along an arbitrary parametric line L through the volume. The slope of $\mathbf{A}$ indicates the global scaling along L. If $d\mathbf{A}/dt > 1.0$, there is an overall increase in scale, and conversely, if the slope is less than 1.0, there is an overall shrinkage. In the regions where $d\mathbf{D}(t)/dt > 0.0$ (such as in regions a and c in Fig. 5.2), there is a local expansion and where $d\mathbf{D}(t)/dt < 0.0$, as in region b and d, there is a local compression (see Fig. 5.3. When $d\mathbf{D}(t)/dt = 0.0$, there is only a linear translation component to the transformation. Since $\mathbf{A}$ represents the best affine transformation, the overall scale should be accounted for. Therefore, $\mathbf{D}(t) \approx 0$ at either end of the mapping domain (or field of view FOV) and thus establishes the boundary-value at the limits of the field. (Note that $\mathbf{D}(t)$ is not strictly equal to zero, since the global scale applies to the entire 3-D object and not only to L.)

In order to map the information from the target space onto the source data set (*e.g.*, geometric contours from a standardized model onto a patient's MRI) it is necessary to invert the deformation field. For $\mathbf{N}(t)$ to be invertible, it must be monotonically increasing, requiring that $\mathbf{N}$ must have a positive-definite derivative. Therefore, the slope of $\mathbf{N}$ is bounded by:

$$0 < d\mathbf{N}/dt < \infty. \tag{5.3}$$

Substituting for $\mathbf{N}$, this leads to:

$$-d\mathbf{A}/dt < d\mathbf{D}/dt < \infty. \tag{5.4}$$

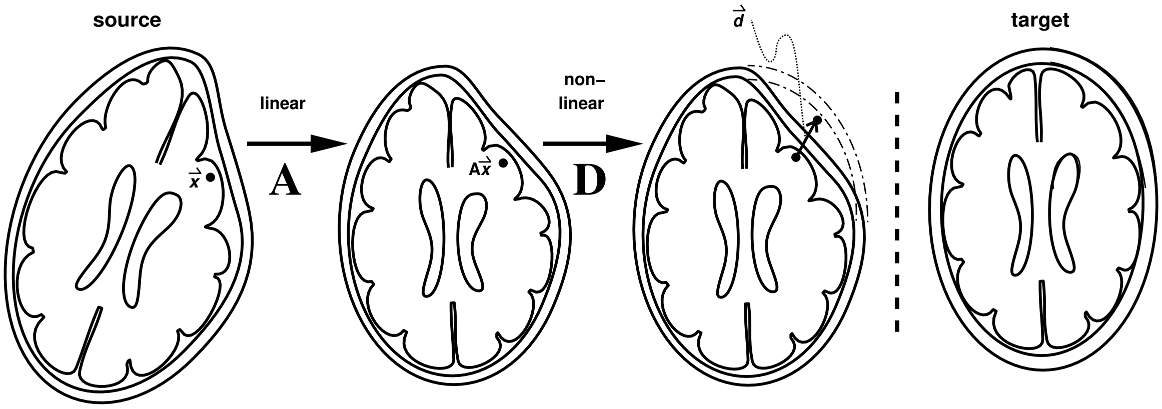

Figure 5.1: Displacement required after linear transformation.

The source image on the left must be registered to the target image on the right. The first step involves a linear transformation $\mathbf{A}$, that accounts for global position, orientation and scale. Under the linear transformation, a point $\vec{x}$ is mapped to $\mathbf{A} \cdot \vec{x}$ in the target volume space. A non-linear transformation $\mathbf{D}$ must be estimated, where an additional deformation $\vec{d}$ maps $\vec{x}$ to its homologue, for all points $\vec{x}$ of the source volume.

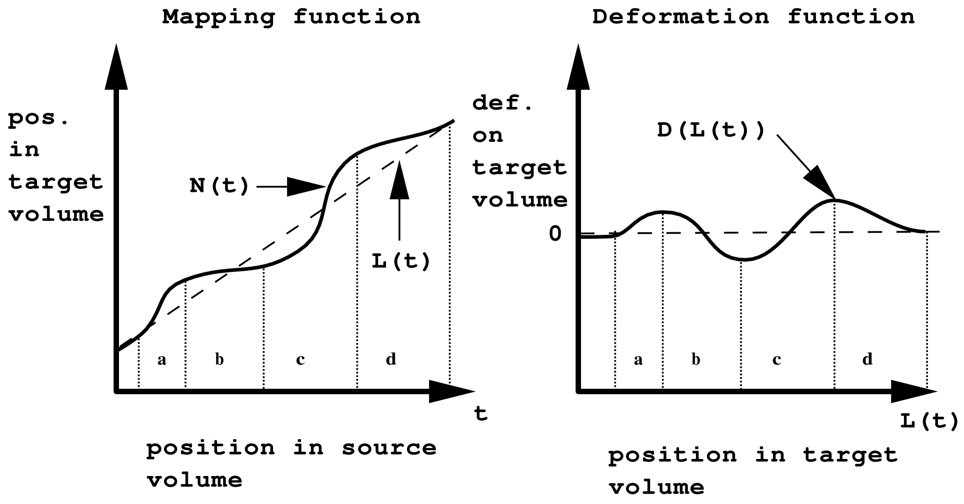

Figure 5.2: 1D profile through deformation field.

The graph on the left illustrates how points of the source volume are mapped into the target space. The graph on the right shows the deformation required at each point in the target space. $\mathbf{N}(t)$ is the deformation mapping function along the line. $\mathbf{L}$ is the linear component of $\mathbf{N}$. $\mathbf{D}(t)$ is the non-linear component of $\mathbf{N}$. There is expansion in regions **a** and **c** and contraction in regions **b** and **d**.

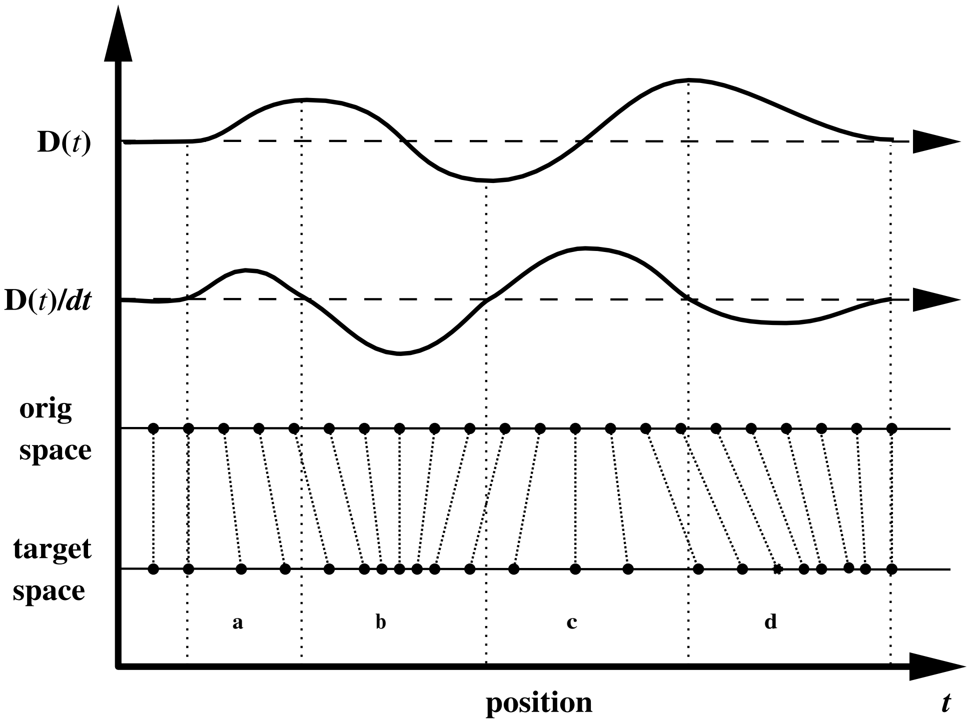

Figure 5.3: 1D deformation field.

This 1D schematic shows how points in the original data are mapped into the target space by a non-linear deformation function $\mathbf{D}$. Note that slope of $d\mathbf{D}/dt$ cannot be less than $-d\mathbf{A}/dt$, otherwise the points will overlap. Regions **a,b,c,d** as in Fig. 5.2.

Thus, the rate of change of the deformation function is bounded below by the perpendicular to the slope of $\mathbf{A}$ and above by ∞. These limits have a physical interpretation: the deformation is constrained so that it cannot compress two distinct points together or allow an overlap, nor can it induce a tearing of the field. The magnitude of the deformation is limited by the size of the FOV since the deformation must have approximately zero magnitude at both the boundaries of the field, and once the magnitude of the function exceeds FOV, it cannot return to zero without causing the data to overlap on itself. These functional constraints can be limited further to more realistic values using anatomical information. For example, the maximum deformation size may be estimated from neuroanatomical and morphological data. Also, the amount of compression and expansion can be limited by the local anatomy; the size of a normal individual's temporal lobe cannot be twice as big, for example, as that measured on average once the brain is normalized into standard space. Limits such as these, along with that of the boundary value, are used to constrain the estimation of the deformation function below.

5.2.2 Deformation field

A tri-variate polynomial, such as that used by Lemoine *et al.* (1991a), was not chosen to represent the deformation function because more degrees of freedom were necessary. Instead, the deformation function is represented explicitly by a deformation field with the following definition:

> A deformation field is a vector valued volumetric data set representing a band-limited deformation function.

In the implementation of the algorithm, the deformation function $\mathbf{D}$ is represented by a deformation field, where a 3-D displacement vector is stored for each voxel position in the field. Three scalar volumes are stored: dx, dy and dz, representing the x, y and z-components of the 3-D displacement vector. For a given arbitrary (x, y, z) position in the domain of the deformation function, the value of the corresponding 3-D displacement is given by interpolation in each component volume, yielding the three necessary values for the 3-D vector.

Since each component of the deformation field is stored volumetrically, only band-limited deformation functions can be properly represented. However, even though the maximum rate of change is functionally limited by ∞, the rate determined for real data is finite. The limits are imposed by actual anatomical morphology and by the inner scale defined by the resolution of the MR images.

The recovery of the deformation function, presented in the next section, is based on local estimation of displacement vectors, one for each voxel. Each estimated vector is a point sample from the deformation function and represents the best distortion for the given point and its surrounding neighbourhood based on the cross-correlation residual, subject to the allowed degrees of freedom. In this way, the recovered deformation field is

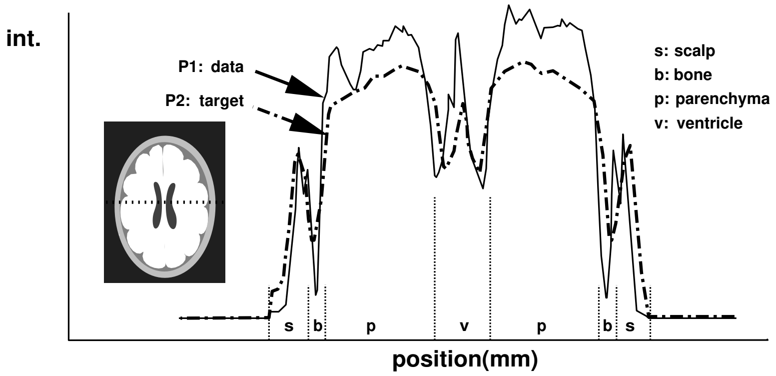

Figure 5.4: 1D profiles of 2 different T1-weighted data sets.

These two profiles correspond to the image intensity along a line running from left to right through two registered data sets at the level of the ventricles (inset). After linear registration, there remains significant positional difference at the ventricles and skull-brain interface.

forced to be a band-limited version of the true function, since the estimation process acts as a low-pass filter. Therefore, the maximum rate of change of the deformation field is limited by the smallest scale of estimation.

5.3 Estimation of deformation - 1D case

The goal of the non-linear registration procedure is to recover the deformation field $\mathbf{D}$, given only $\mathbf{A}$, the intensity values of $\mathcal{V}_s$ and $\mathcal{V}_m$ and the constraints described above. A 1D example is used to explain the principles of the methodology in this section. The full 3-D procedure is then detailed in the following section.

Fig. 5.4 shows two 1D intensity profiles $P1$ and $P2$ through two volumetric data sets. The application of the recovered non-linear transformation should align all labelled regions of $P1$ with $P2$ so that each region of $P1$ can be directly compared to the corresponding region of $P2$. If the two profiles are labelled *a priori* , then explicit correspondence is available and it is relatively trivial to calculate an interpolating transformation to register homologous regions together. Since the data to be matched are not already segmented and labelled, explicit correspondence does not exist. Therefore, correspondence between local regions must be established from the data itself, and is accomplished by maximizing similarity of image-derived features.

Examination of the 1D profiles of Fig. 5.4 reveals that corresponding regions have similar relative intensities and positions. Correspondence must be established on a local basis; a match for a small region of the first image must be found in the second image. The displacement required to achieve this match is the local deformation that must be stored

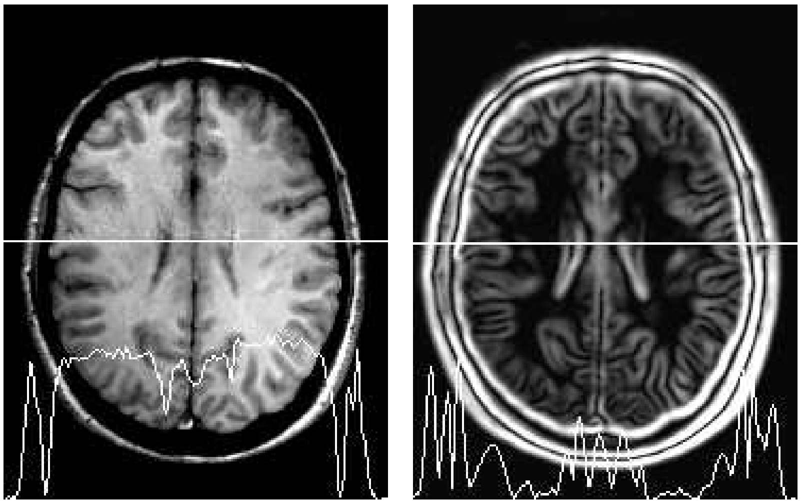

Figure 5.5: Edges reveal image structure
The image on the left shows a typical transverse MR image with an an intensity profile of the indicated row. The image on the right shows the corresponding gradient magnitude image calculated after blurring with a FWHM=4mm Gaussian kernel. Its corresponding profile is also shown.

as part of the deformation field. Unfortunately, a small neighbourhood of raw image intensity does not exhibit a great deal of local image structure for matching. Other, more specific and robust features must be extracted and used to drive the registration and these features must satisfy the same constraints as those used to accomplish the linear registration, i.e., they must be linear, rotationally invariant and shift invariant. Once again, convolution with differential Gaussian operators is used to extract local structure from the raw intensity data. Extrema of the first derivative data indicate the presence of edges in the raw images (see Fig 5.5). These edges can be used to align local regions of the data.

To begin the example, let $p_1' = T(p_1)$, where $p_1 \in P1$, the data profile in Fig. 5.4. Once the global position and scale are accounted for by the linear transformation $\mathbf{A}$, a point p_1' will be at some distance d from its homologue $p_2 \in P2$, the target profile. (Note that the vector $p_2 - p_1'$ is precisely the deformation to be recovered.) If the upper limit, d_{max}, in magnitude of all possible d's is known (*e.g.*, the value can be estimated from a knowledge of the normal range of anatomical variability), then the search on $P2$ for the homologue of p_1 can be limited to a neighbourhood of radius d_{max} centered on p_1'. Therefore, each deformation vector can be estimated using only information from a neighbourhood of fixed size surrounding each point pair $(p,\ \mathbf{A}(p))$ with a form of classical template matching[1] (Duda and Hart, 1973). The template is defined by the gradient magnitude feature in the

[1]Since only local information is used to estimate each deformation vector, in principle the procedure can be easily vectorized to run on parallel machinery.

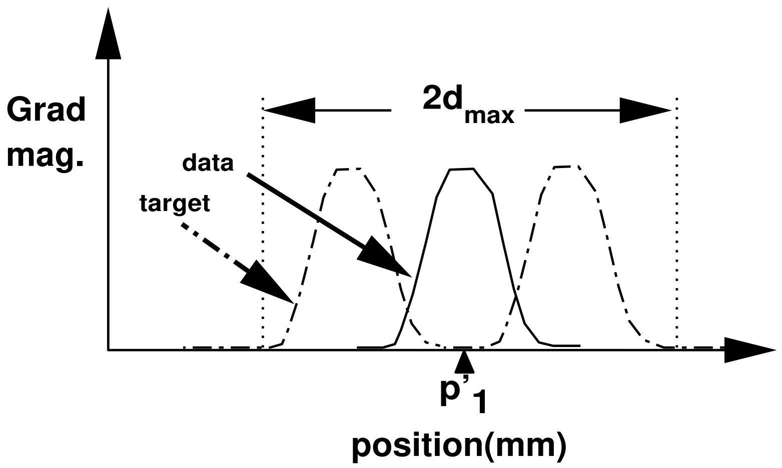

Figure 5.6: Correspondence problem.

The solid line represents the feature template from the source data set. The broken line shows feature neighbourhood of the target data set. Given only this information, it is not possible to select which of the two peaks is best.

local neighbourhood of p_1. It is overlaid on P2, and moved from $p'_1 - d_{max}$ to $p'_1 + d_{max}$. The values of the template are compared to the features of P2 and a match is declared where the template best correlates with P2. The offset from p'_1 is the local deformation vector at p'_1 for the global non-linear transformation.

A problem exists when using only local information to estimate the deformation for some point p, and is illustrated in Fig. 5.6. Under the linear transformation, the gradient peak at p_1 is mapped to p'_1, exactly between two equally likely target peaks. Since explicit correspondence is not available, the process is unable to select which of the two possible matches is correct since both states will result in an equal measure of similarity within a data window limited by $2d_{max}$. More global information is needed to determine which of the two states is preferable. Even if the data window was extended past the limits indicated by d_{max}, it would be possible to construct a similar example where two other deformations would yield the same similarity measure and further information, beyond the window defined by d_{max}, would be needed to make the decision.

The solution to this problem resides in the fact that data blurred at a higher scale contain information from a wider region of support. Thus, lower resolution data are used to bias the estimation at the current scale towards the correct answer. The strategy does not attempt to recover the complete deformation in one step, using data from only one scale. Instead, the non-linear transformation is derived in a iterative fashion where the procedure estimates the deformation field using gradient magnitude estimated at different spatial scales. Starting with very blurred features from both source and target data, the largest deformations are recovered first. Each successive iteration refines the registration estimated at the previous scale by the addition of smaller deformations estimated from less blurred data.

Besides using the data in a multi-scale fashion, this strategy also recovers the deformation itself in a hierarchical manner. The procedure begins by estimating an initial band-limited deformation field, $\mathbf{D}_1$, at a low resolution (with large voxel spacing) that accounts for part of the residual mismatch not dealt with by $\mathbf{A}$. Once completed, $\mathbf{D}_1$ is added to $\mathbf{A}$ to form $\mathbf{N}_1$. When $P1$ is resampled by $\mathbf{N}_1$ and compared to $P2$, the overall similarity measure will have increased since each deformation in $\mathbf{D}_1$ has increased the local similarity at that scale between corresponding neighbourhoods. However, since $\mathbf{D}_1$ is of low resolution, many regions will not be in perfect correspondence. In order to account for these mismatched regions, the next recovery step estimates $\mathbf{D}_2$ at twice the resolution as $\mathbf{D}_1$ based on data blurred at half the previous scale. $\mathbf{D}_2$ is added to $\mathbf{N}_1$ to form $\mathbf{N}_2$. This process is repeated for each scale of interest.

As in the linear case, the scales of interest (measured with the FWHM) are 2, 4, 8, 16mm. It was found that scales greater than 32mm blur the data too much to be of any use in the estimation of the deformation. Experiments in section 8.3.3 show that the deformation recovered at the scale of 2mm does not significantly improve the registration between different subjects. This is due to the errors caused by non-equivalent topology between brains, that outweigh the benefits of calculating the deformation at the lower scale. Therefore, this scale step is eliminated from the estimation process. The final deformation field is estimated at the 4mm scale and is sampled every 2mm. This field is tri-cubically interpolated to have deformation vector per voxel of the source data, to map that voxel into the target space.

5.4 Estimation of deformation - 3-D case

At each scale space step in the deformation estimation, the goal of the optimization procedure is to identify the non-linear transformation $\mathbf{D}$ that maximizes the similarity between features derived from the voxels in the two data sets: the volumetric data for the subject, $\mathcal{V}_s$, and the model, $\mathcal{V}_m$. These sections describe the 3-D non-linear registration procedure for a given scale step. The complete recovery strategy is summarized in section 5.4.6.

5.4.1 Lattice definition

The deformation function is recovered hierarchically at a number of different steps, each successive step refining the estimation of the previous one. Each step attempts to recover the deformation function for a given scale, where the scale is a measure of the resolution of the band-limited deformation function to be recovered. As in the previous chapter, the deformation field is evaluated on a set of voxel positions, $\mathcal{L}$, organized on a 3-D cubic lattice. If the FWHM of the current scale step is used to measure resolution, then the voxel spacing of the deformation field must be no greater than FWHM/2 to recover the function without aliasing, i.e., the usual Nyquist sampling limit.

A correlation coefficient is estimated at each of the lattice nodes using the same principle as the linear registration, however the correlation is based on a summation over a set of voxels in the local neighbourhood of the particular node. Hence, $R()$ is the normalized correlation value between the local neighbourhood of $\vec{x}$ in $\mathcal{V}_s$ and the corresponding neighbourhood of $M(\vec{x})$ in $\mathcal{V}_m$:

$$R(\mathcal{V}_s, \mathcal{V}_m; \mathbf{N}, \vec{x}) = \frac{\sum_{v \in \mathcal{N}_{\vec{x}}} f(\mathcal{V}_s, v) f(\mathcal{V}_m, \mathbf{N}(v))}{(\sum_{v \in \mathcal{N}_{\vec{x}}} f^2(\mathcal{V}_s, v))^{\frac{1}{2}} (\sum_{v \in \mathcal{N}_{\vec{x}}} f^2(\mathcal{V}_m, \mathbf{N}(v)))^{\frac{1}{2}}}, \tag{5.5}$$

where $\mathcal{N}_{\vec{x}}$ is the local neighbourhood of $\vec{x}$ with diameter=$\frac{3}{2}$FWHM, $f(\mathcal{V}, v)$ is the interpolated feature value from the volume $\mathcal{V}$ at voxel position v, and the summation is performed over all voxel elements $v \in \mathcal{N}_{\vec{x}}$. A local neighbourhood diameter of $\frac{3}{2}$FWHM was chosen since a diameter of at least FWHM/2 was needed to completely cover the source and target domains. A diameter of FWHM was initially used to guarantee complete coverage, even in areas subject to expansion or stretching and to ensure overlap between successive neighbourhoods to maintain continuity of the deformation field. Initial experiments showed that this was not enough, and $\frac{3}{2}$FWHM was the smallest diameter that resulted in robust estimation of deformation vectors. Larger diameters such as 2·FWHM or 5/2·FWHM enforce greater continuity between regions at the expense of possibly missing small local deformations.

The normalized similarity value, S, between the two volumes, $\mathcal{V}_s$ and $\mathcal{V}_m$ for a given transformation $\mathbf{N}$ is defined as the sum of all local correlations:

$$S(\mathcal{V}_s, \mathcal{V}_m; \mathbf{N}) = \frac{1}{n} \sum_{\vec{x} \in \mathcal{L}} R(\mathcal{V}_s, \mathcal{V}_m; \mathbf{N}, \vec{x}), \tag{5.6}$$

where n is the number of elements in $\mathcal{L}$. R, and like-wise S, take on a maximum value of 1.0 when the two volumes are in perfect registration.

5.4.2 Local neighbourhood features

The estimation of a local deformation vector is based on the correlation of features from both the source and target volumes. For the non-linear registration method, the geometrically invariant feature of image-intensity gradient-magnitude is used so that feature values derived from the volumetric data are not dependent on the original position, scale and orientation of one data set with respect to the other. Since gradient extraction is dependent on scale, this parameter must be chosen in relation to the resolution of the estimated deformation function.

It is not clear how the scale parameter should be chosen. A very large value is of no use, since it will blur out all structural detail from the volumetric data. Conversely, a very small value will extract a high amount of structure and increase the probability of local mis-matches, as was shown in Fig. 5.6. The size of the Gaussian blurring kernel applied to the volumetric data was chosen to be equal to the resolution of the deformation field estimated at the current scale step. Scale factors yielding gradient magnitude data with twice the resolution and with half the resolution were also tried. Only small differences in the recovery of the deformation field were observed, when compared to when the gradient magnitude data was at the same scale.

The local neighbourhood of $\vec{x}$ is defined by a sub-lattice of nodes. For data blurred with σ=FWHM/2.35, it is necessary to have samples spaced at FWHM/2 to recover the underlying band-limited function. However, only a limited region of the gradient magnitude data is being sampled by the sub-lattice. Twice as many samples (in each direction) are required to fit the gradient data, since the gradient changes at a twice the rate of the blurred data. Therefore, the cubic sub-lattice of diameter $\frac{3}{2}$FWHM requires 7^3 ($=(2n-1)^3$), where $n = 4$ samples were needed to represent the blurred data) samples to represent the local neighbourhood.

5.4.3 Node thresholding

An intensity threshold is used to eliminate lattice nodes with near zero-value gradient intensity, since no deformation can be estimated in these regions. The threshold was arbitrarily set to 10% of the average gradient magnitude. In this fashion, nodes that fall in the middle of a homogeneous region, away from any edge structures, are not used to define a local neighbourhood and consequently, are not used to estimate a local deformation vector. The deformation for these nodes is interpolated from neighbouring lattice points in a smoothing step described below. It is important to note that the threshold is applied only to select or reject a particular node for estimation of the local deformation, and it is not used to eliminate any voxels from the calculation of local correlation for any selected node. Therefore, changing the value to 20%, for example, slightly thins most of the edges seen in the gradient magnitude image, leaving them behind so that they can be used to estimate the deformation. However, low contrast edges may be lost, and thus the deformation in these regions must be interpolated from neighbouring nodes.

5.4.4 Estimation of local deformation

It is clear that equation 5.6 is maximized when each of the terms in the summation are at a maximum. Since the transformation $\mathbf{N}$ is stored such that there is one deformation vector for each node $\vec{x}$, equation 5.6 is maximized by optimization at each node of L. The optimization is similar to that performed for the linear registration procedure, however only three parameters are optimized instead of nine, since only three are required to define the local deformation vector $\vec{d_i}$ that maximizes for correlation of the local neighbourhood of $\vec{x_i}$ with its homologue in the target volume. Hence, the goal is to find $\vec{d_i}$ that maximizes $R(\mathcal{V}_s, \mathcal{V}_m; (\mathbf{N} + \vec{d_i}), \vec{x_i})$. (See Fig. 5.7.)

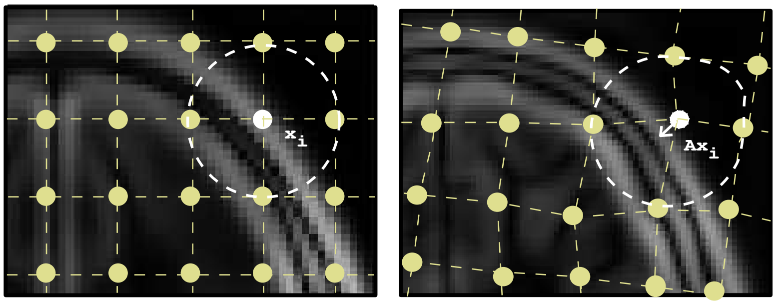

Figure 5.7: Local node deformation strategy.

Theses images show a part of the 3-D lattice overlaid on the gradient magnitude data. Under the current transformation, the node x_i, is mapped to $A \cdot (x_i)$ on the model. A small displacement towards the edge of the scalp is necessary to maximize the neighbourhood correlation for this node. This found through optimization of a three translational parameters dx, dy and dz.

The local neighbourhood of $\vec{x_i}$ is specified by the ensemble of interpolated feature values from each sub-lattice node defined above. Before optimization, the corresponding neighbourhood in the target volume is defined by transforming each sub-lattice node by the current non-linear transformation $\mathbf{N}$ and interpolating the feature value on the target volume.

The vector $\vec{d_i}$ is found using a three-dimensional Simplex optimization procedure. The starting values of the simplex are initialized to $(0, 0, 0)$, $(r, 0, 0)$, $(0, r, 0)$ and $(0, 0, r)$, where r is the radius of the initial search space of the simplex. The value of r is set to the minimum of d_{max} and FWHM. The value of FWHM is chosen because the functional limits on the rate of change of the deformation field constrain the maximum deformation to be smaller than $\pm$FWHM for the current scale step.

At each step of the optimization procedure, the coordinates of the target sub-lattice are modified, the feature values re-interpolated, and the correlation between the two sub-lattices calculated. The optimization procedure stops when the normalized difference between the maximum and minimum correlation values, evaluated at the simplex vertices, is smaller than a preset tolerance. The largest tolerance value possible should be used to avoid ineffective correlation function evaluations. Values ranging from 10^{-6} to 10^{-1} were tried, and it was found that a value near 10^{-2} worked well for the test volumes used. The tolerance of 0.01 caused approximately 130 function evaluations on average per node $\vec{x_i}$. Multiplying or dividing this tolerance by a factor of ten changed the average number of function evaluation by ± 20 only, without affecting the resulting non-linear registration result. Instead of using 0.1, the value of 0.01 was kept since tolerances slightly greater than 0.1 caused local mis-matches in some test volumes.

5.4.5 Iterative refinement

Since the deformation at a single node is part of a continuous global warp, it affects not only that node, but all neighbouring nodes as well. Therefore, the neighbourhood correlation values for the nodes in the vicinity of $\vec{x_i}$ will also change as a result, making the terms in the summation of equation 5.6 inter-dependent. Consequently, an iterative approach must be used, where a fractional value of $\vec{d_i}$ is stored for each node at each iteration.

If this fractional value is small (less than .1) then many iterations are necessary to bring the two data sets into alignment. If the value is large (0.5 - 1.0) then there is a possibility that the deformation field recovered will not be continuous (*e.g.*, if one node matches to a point that is not its homologue). Experimentally, a value between 0.2 and 0.5 has been found to represent an acceptable compromise, allowing a smooth deformation field to be recovered without necessitating a large number of iterations.

At the end of each iteration, the stored fractional value of $\vec{d_i}$ is added to the existing warp. Since some mis-matches may have occurred during the iteration, a smoothing step is applied to the deformation field to remove local outliers.

5.4.6 Non-linear registration summary

The following pseudo-code summarizes the non-linear registration strategy.

non-linear registration:

foreach scale **do**

- define a 3-D lattice on the source volume
- initialize **N** for this scale
- **foreach** iteration **do**
  - initialize **D** to zero
  - **foreach** node $\vec{x_i}$ **do**
    - define sub-lattice around $\vec{x_i}$
    - map sub-lattice through **N**
    - call simplex optimization to get $\vec{d_i}$
    - store fraction of $\vec{d_i}$ in **D**
  - update **N** with **D**

5.5 Speed considerations

Running the non-linear registration algorithm on an MRI volumetric data set of the brain originally required approximately 300 hours of CPU time on a SGI Indigo[2] Extreme (a 60.5 SpecFP Unix machine). Optimization of the program was necessary to bring this amount of time down to a reasonable level. Three approaches were employed to increase algorithm speed:

1. Reduce the amount of time required to estimate a single deformation.
2. Reduce the number of nodes where deformation must be estimated.
3. Reduce the number of iterations required for convergence.

These methods are described below:

Node estimation: Although a number of factors entered into the cost of an estimation of the local correlation, they were all proportional to the number of nodes in the local sub-lattice. Changing the local neighbourhood from a cube to a sphere of the same diameter reduced the number of sub-nodes almost by half, without affecting the resulting correlation results.

Number of nodes: Two methods were used to reduce the number of nodes participating in the the deformation estimation:

1) The overall number of nodes could be reduced by raising the threshold on the gradient magnitude. In this case, the deformation is not estimated everywhere, and must be interpolated from neighbouring regions. Thus a balance must be found between estimation and interpolation. By selecting the threshold manually, one can choose which edges in the data will participate in the estimation and which regions will be estimated. Originally, this threshold value was set to 1% of the average gradient magnitude to reject only nodes with no gradient magnitude. The value of 10% described above eliminates some of the data due to noise and maintains almost all of the edges detected in the volumetric data.

2) A voxel mask can be used to limit the number of nodes. Since the main interest is in fitting the brain, a mask defining only brain voxels was defined on the target volume – eliminating all nodes outside the brain. Note that regions outside the brain are used to constrain the fitting for low resolution deformation fields. The mask was used only at the high resolution step, after the overall brain shape has been accounted for.

Number of iterations: Two procedures have been implemented to reduce the number of iterations needed to complete the optimization. The first is applied to each iteration of the optimization algorithm, and the second modifies the multi-scale aspect of the non-linear registration strategy.

1) After an iteration over all nodes is completed, a "mini-iteration" is executed. In this pass, the deformation is estimated only at nodes where the local vector from the previous estimation step was greater than some threshold. Approximately 10% of the nodes satisfy a threshold set equal to the magnitude of the average deformation plus one standard deviation. Inserting a "mini-iteration" between each true iteration allows us to skip the estimation of nodes where the additional deformation will most likely be null. Even though a single iteration now requires 10% more time to calculate, the total number of iterations may be reduced by half without appreciable loss in accuracy.

2) Since the optimization procedure terminates sooner when the initial conditions are close to the final result, and since calculating the fit at a scale of Smm requires less than one quarter the time taken for the scale $S/2$mm, an extra scale-step was added at the beginning of the multi-scale procedure. This step recovers the deformation at a scale of 24mm, however it uses the data that is blurred for the 16mm scale step. The result of this estimation is used as input to the procedure at the 16mm step. Hence, this step can be considered an initialization procedure for the previously mentioned non-linear registration strategy.

These modifications along with some internal coding changes reduced the time required by a factor of ten. The non-linear registration program now requires approximately 30 hours to run to completion.

5.6 Discussion

This non-linear warping method is based on the assumption that there exists an equivalent topology between brains so that a 1-to-1 mapping can be estimated to account for differences between the two data sets. While not strictly true, the validity of this assumption depends on the spatial scale of comparison. When blurred at 16mm FWHM, all structures of the brain are topologically equivalent because only major structures (temporal lobe, ventricles, longitudinal fissure) are apparent. Smaller features are simply not visible. At 8mm FWHM, regions within the interior of the brain (*e.g.*, basal ganglia, ventricles, brain stem) and the major gyri and sulci remain topologically equivalent, while secondary and tertiary gyri may no longer be equivalent.

This procedure can be applied to blurred data of the whole brain volume to correct overall shape (as if all brains were forced to fit inside the same skull). At scales smaller than 8mm, this method is still able to register structures as long as there is a 1-to-1 mapping between them such as the ventricles, basal ganglia and major sulci. On the cortex, however, the topology is not consistent for secondary and tertiary gyri. This is a completely different problem, complex enough to warrant a separate research project in our laboratory. A solution to this problem has been proposed by MacDonald *et al.* (Mac-

Donald *et al.*, 1994) where the cortical surface is extracted and mapped onto a simple parametric space with the same topology, *e.g.*, a sphere.

5.6.1 Relation to other techniques

Other non-linear matching methods have employed different objective functions, matching strategies and types of data. In this section we describe the work of two authors that address most closely the automatic image-matching problem considered in this thesis and compare their techniques with ours on the basis of 1) the manual interaction required, 2) the features used, 3) the type of non-linear fitting model, 4) the objective function, and 5) the type of data used.

Broit (1981) has developed a non-linear registration procedure for CT data of the brain that models the physical deformation process of an elastic solid being warped from the position of the first data set to that of the target. After defining a 3-D grid on both volumes, warping is based on matching the local neighbourhoods of corresponding grid points. Edge-like features, derived from the image, are used to derive deformation force vectors that move each grid point, while elastic constraints attempt to resist them. The differential equations describing the system are integrated over time until an equilibrium is found. In order to avoid local minima, and speed the fitting process, the optimization is completed in a hierarchical fashion by changing the elastic properties of the model. At first, the model is rigid and the registration process essentially calculates the linear registration transformation. At each successive iteration, the model is made to be less and less stiff, thus accounting for large deformations first and refining the fit with smaller deformations in successive iterations.

Zhengping and Mowforth (1991) have developed a similar technique to match cerebral MR volumetric data, however the non-linear registration is accomplished only on transverse slices, after 3-D linear matching. Even though this method is 2-D, it is included for comparison because it is applied to the brain segmentation problem and could be directly extended to 3-D (although it has not been done as of yet). Following a coarse-to-fine regime, the deformation is recovered in a multi-scale fashion, based on the difference of Gaussians (DoG) filtered data at different resolutions.

Manual interaction required: In the procedure developed by Broit, the skull and scalp must be manually edited from each of the transverse CT images. While this step can be simplified by use of thresholding and region growing, the user must decide on the threshold levels and select non-brain regions for removal. After linear registration, Zhengping claims that it is sometimes necessary to select the appropriate 2-D slice for non-linear registration. If extended to 3-D, this operation could possibly be eliminated. In comparison, the method presented here is completely automatic; no user intervention is required to prepare or manipulate the data.

Features: All three methods use edges to drive the registration, however they differ in their method of edge detection. Broit used Hueckel basis functions (Hueckel, 1973) to extract edges from the data. Even though these detectors approximate derivative functions, they are not rotationally invariant. Therefore, the edge strength differs depending on its orientation. A derivative approximation is also used by Zhengping. The DoG operator is a type of band-pass filter, and is calculated by subtracting data blurred with a wide Gaussian filter from the same data blurred with a narrow filter. The selection of the two fixed filter widths can favor edges of one strength and size over another.

We have used a mathematically sound procedure to directly calculate the partial derivatives of blurred volumetric data in Fourier space, thus taking into account the full Gaussian kernel and thus avoiding truncation artefacts. The gradient magnitude calculated from the partial derivatives is invariant to both rotation and position. Therefore, the features detected by this procedure are intrinsic to the object, and don't depend on the position or orientation of the object in the image as is the case for the other two methods.

Since the recovery procedure is essentially edge-driven for all three methods, the resulting deformation field are all dependent on the contrast of the original images. If a structure of interest has sufficient contrast with its surrounding to yield an edge (as seen in the gradient magnitude data), then the deformation for that area will be directly calculated, based on the position of that edge. The deformations for regions of interest that do not have edges apparent in the gradient magnitude volumes are interpolated from nearby edges. Their actual location is completely dependent on the type of interpolant chosen, which in turn is dependent on the non-linear transformation model.

Non-linear model: All three authors use a non-linear transformation model that is distributed in nature. These models enhance their representational power over parametric deformation models such as those proposed by Fischler and Elschlager (1973), Widrow (1973), Lemoine *et al.* (1991a) or Gerig *et al.* (1991b). Every point in the deformation field contributes 3 degrees of freedom (2 in the case of Zhengping) to the global warp allowing a large continuum of deformations to be specified.

The non-linear transformation modeled by Broit is based on differential equations that describe the behavior of elastic solids. This model is applied to all points of the volume, balancing the need for deformation, as indicated by correlation of the image data, against the restraining forces of the elastic model. While the use of a physical elastic properties provides a flexible model of deformation, it is not necessarily a faithful representation of anatomical variability. The data is deformed like a piece of homogeneous elastic material, without consideration of real morphometric differences. Neither we, nor Zhengping use a physical model to represent the deformation process. We simply employ smoothing constraints to ensure a continuous deformation field that does not stretch or compress the data beyond the functional limits described above.

Objective function: The objective function is closely linked to the transformation model. Each of these methods uses cross-correlation over small neighbourhoods to estimate the local deformation needed to bring one region in line with another, increasing their similarity. Broit balances this deformation force with elastic constraints, so that the goal of the optimization is to minimize the cost function, *cost = cost(deformation) - cost(similarity)*. The paper from Zhengping does not define the objective function used. Here, the deformation forces are are constrained only by the limits imposed by the functional representation of the deformation function. Smoothness is guaranteed by the iterative nature of the estimation process.

Data: The type of data used for registration differs between Broit and Zhengping. Broit used CT data for registration. Like Zhengping, we use MRI data for registration. The use of MR or CT has direct ramifications on the type and position of features that can be extracted from the data. Since CT has much lower grey-white contrast than MR, the edges extracted correspond almost entirely to the ventricular and cortical surfaces. Many of the other subtle differences seen on MR are simply not visible in CT. For example in CT, cortical edges do not present as detailed an enfolding of the sulci as apparent in MR, nor is the contrast as great between basal ganglia structures such as the putamen and globus pallidus. Thus, the CT images are not as complex and this renders matching of CT data a much simpler problem. Broit therefore uses only the cortical and ventricular surfaces for non-linear mapping, interpolating the deformation between them.

Broit further simplified the problem by manual extraction of the brain in a preprocessing step. Not only is the cortical surface essentially defined by hand, and there is no possibility of mis-match between scalp and cortex since the former is removed. The computer is left to extract the ventricular surface – not a difficult task, since it is easily achieved by thresholding the CT data. Non-linear registration by edge fitting is a powerful tool, however interpolation between cortex and ventricle leaves to be desired. Like Zhengping, our technique requires no editing of the data before matching and all edges in the data are used to derive the deformation field.

5.7 Summary

In this chapter, a technique has been described to estimate a non-linear deformation field necessary to register two MRI data sets that differ in second-order anatomy due to morphometric variability. The estimation is accomplished in a hierarchical multi-scale fashion, where the registration is performed at different spatial scales, refining the match at each stage. This deformation field can be applied to one data set by resampling, making all points of the first data set match with their homologues in the second. The method presented here goes beyond existing techniques in significant ways while retaining some of their best features, being objective, fully automatic, and robust against noise. This procedure forms the core of the segmentation method described in the next chapter.

Chapter 6

Stereotaxic segmentation:theory and implementation

The purpose of this chapter is twofold: 1) to describe the rationale for, and the design of, the stereotaxic segmentation model, and 2) to present the segmentation method that merges this model with the registration methods presented in the previous chapters.

6.1 Introduction

In this thesis, the goal of neuro-anatomical segmentation is to identify and delineate regions from MRI volumetric data that correspond to biologically meaningful structures. Hence, a segmented volume consists of a set of labelled voxels. A pre-defined atlas, most often represented as a collection of geometric objects, is used to determine which anatomical structures are to be segmented. The atlas guides the segmentation process by specifying both the spatial extent of each structure and relationships between them. The automatic segmentation procedure presented here is based on the following premise:

> For both manual and automatic model-based segmentation methods, it is assumed that at some level of representation, the brain to be segmented and the brain atlas used for guidance are topologically equivalent, but that internal structures are mutually deformed.

This implies that there exists a one-to-one mapping between a subject's brain image volume and the atlas model. Segmentation is achieved by deforming the atlas, through the spatial warping function that represents the one-to-one mapping, onto the given brain data set.

The segmentation procedure developed in this thesis requires a two-component segmentation model to reverse the conventional model-based segmentation strategy. Instead of fitting the geometric contours of a brain atlas directly to the raster data of an image volume, features derived from the grey-level voxel intensities of the source data set are registered to similar intensity features stored in the first component of the brain segmentation model. Hence, the necessary spatial warping function between source and target model is recovered. The second part is formed by a brain atlas that is co-extensive with the first model component. It is similar to most conventional atlases, consisting of labelled regions and their associated contours. It is mapped through the inverse of the recovered transformation onto the native source data volume to complete the model-based segmentation process.

The major advantage of this procedure over previous techniques is it that the atlas is not used explicitly to determine the non-linear registration transform. Therefore, the segmentation method is completely atlas independent since any atlas defined in the stereotaxic brain space can be used for segmentation. The following sections describe the segmentation model and strategy in more detail.

6.2 Volumetric Brain Segmentation Model

An ideal model for the segmentation of brain images must meet a number of requirements. Foremost, the model must have a completely 3-D representation so that positional information is consistent throughout the cerebral volume. Anatomical structures in the brain are three-dimensional, therefore the model must represent their 3-D boundary definitions in a manner that is flexible, easy to manipulate and modify. The model must contain structures whose features can be identified in the data volumes on which the segmentation process will operate. Landmark information other than brain region boundaries (*e.g.*, optic chiasma, ventricular horns, ventricular notch, sinus cavities, cortical sulci, etc.) should also be included to assist in the segmentation process. In order to address problems associated with the original Talairach atlas, the model must incorporate information from a large population of young normal subjects and thus represent normal anatomical variability. Finally, since knowledge of anatomy is continuously evolving, the design of the model should not be dependent on any one existing atlas.

These requirements are met with the stereotaxic model designed here, known as the Volumetric Brain Segmentation Model (VBSM). It serves two purposes: 1) it defines the neuro-anatomical structures that will be identified and delineated by the segmentation algorithm, and 2) it defines a standard coordinate system in which to estimate spatial anatomical variability. The VBSM consists of both volumetric (intensity, gradient magnitude, etc...) and geometric (points, surfaces) data that co-exist in the stereotaxic coordinate system. The volumetric data are used to drive the matching process while the geometric atlas is used to delineate and segment structures. The former are arranged in hierarchical fashion so that gross structures are used to begin the segmentation, and smaller details are included in processing, as needed, to refine the fit. A pyramid data structure is used to represent the volumetric information at different levels of scale. The following sections describe the two model components.

6.2.1 Volumetric MRI model

Instead of using a single brain to define the VBSM, the MRI atlas described in section 3.4 is used as the source for all derived features, thus addressing the main drawback of the Talairach atlas. Since the registration algorithms developed in the last two chapters are used for the segmentation procedure described below, the features normally calculated afresh for correlation between two individual data sets are simply stored for the MRI atlas. Stable features must be extracted from the average MRI data set so that it can be used as a target in the automatic matching procedure. Convolution with a 3-D Gaussian kernel was used to blur the average data set and create three smoothed intensity volumes corresponding to values of FWHM=4, 8 and 16 mm. The corresponding partial derivative (i.e, gradient) volumes were created for the x, y and $z-$directions, yielding three more volumes each of the three scales. Finally, the magnitude of the gradient was calculated at each scale yielding a total of 15 feature volumes. Only the voxel intensity and gradient magnitude volumes at each of the three scales are stored with the MRI atlas to be used as invariant features for both the linear and non-linear matching processes. (Recall that the individual partial derivative volumes are not used, since they do not represent invariant features, see Fig. 4.1.) The MRI atlas and its corresponding feature volumes make up the VBSM-IIF (Invariant Intensity-based Features) and are used to drive the registration processes. Fig. 6.1 shows tomographic slices through the IIF for features at the FWHM = $8mm$ scale.

6.2.2 Geometric anatomical model

While the MRI atlas and the associated feature volumes of the VBSM-IIF serve to define the target space for registration, a companion model known as the VBSM-VOI is required to delineate and label structures for neuro-anatomical segmentation in stereotaxic space. Without loss of any generality, any atlas can be used to define the structures to be segmented in stereotaxic space, since the recovery of the spatial transformation to map a given data set into this space is dependent on the VBSM-IIF only. A 3-D 18-slice volume of interest (VOI) atlas has been previously created to model the human brain (Evans *et al.*, 1988; Marrett *et al.*, 1989a). The atlas was refined by incorporation of more slices in the $z-$ direction in the paper by Evans *et al.* (1991a). This atlas contains 120 3-D polyhedral objects representing the outer surface of important neuro-anatomical structures most often identified for quantitative analysis. The atlas has been used as a guide to assist manual segmentation, where it has been shown to reduce intra-observer variability in structure definition (Evans *et al.*, 1988; Evans *et al.*, 1991a).

In order to generate the 3-D atlas contours, a single normal brain was scanned using 64 contiguous 2mm-thick T1-weighted (TR=400ms,TE=30ms) MRI planes parallel to the AC-PC (anterior-commissure posterior-commissure) plane. The MR images exhibited excellent contrast between grey and white matter as well as clear definition of the ventricles. Using standard neuro-anatomical atlases for reference (Talairach *et al.*, 1967; Talairach

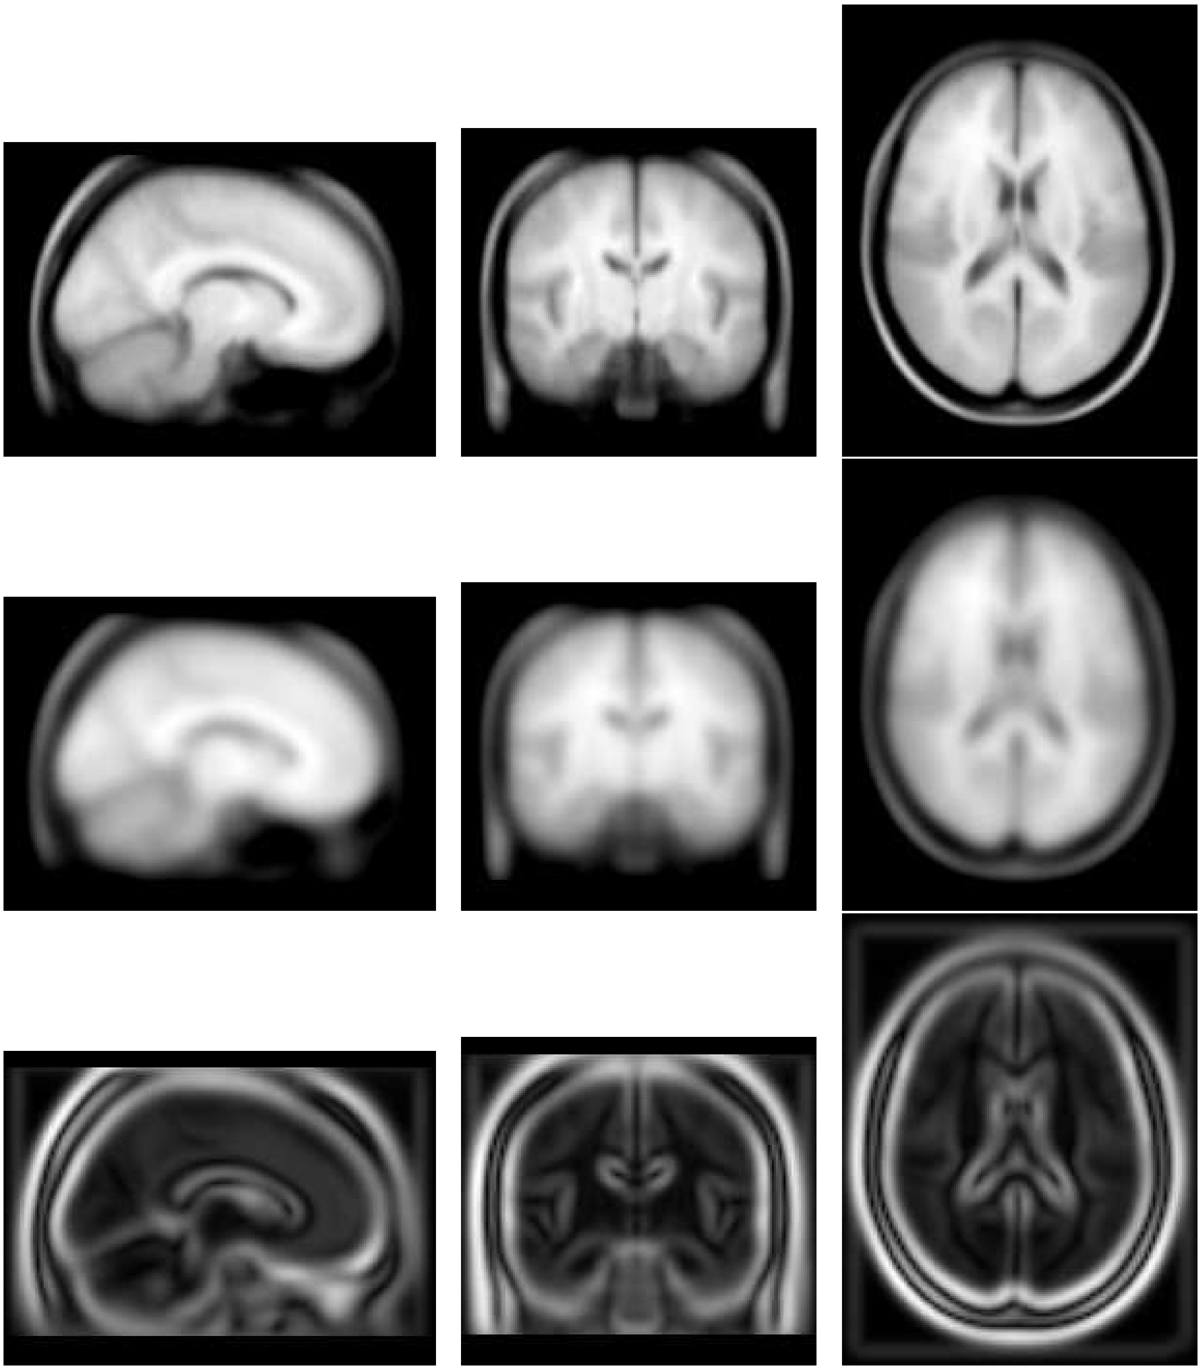

Figure 6.1: 8mm VBSM-IIF.
Tomographic planes through the VBSM-IIF at the FWHM=8mm scale. On the top row from left to right are sagittal, coronal and transverse slices through the original 305 brain MRI atlas. The middle row shows the same slices blurred with a 3-D Gaussian kernel of FWHM=8mm ($\sigma = 3.4$mm). The last row shows the corresponding gradient magnitude data at the same scale.

and Tournoux, 1988; Hanaway *et al.*, 1980; Kretchmann and Weinrich, 1986; Matsui and Hirano, 1978), a neuro-anatomically trained expert outline each of 60 brain regions in each hemisphere on every MRI plane where they were present (Evans *et al.*, 1988). These contours were then tessellated to form a set of polyhedra which can be re-sliced to generate arbitrary 2-D ROI templates and used for 2-D slice-based segmentation tasks (Marrett *et al.*, 1989a; Evans *et al.*, 1989d).

In order to facilitate manipulation with existing software, the atlas was voxelized onto a $1 \times 1 \times 1$ mm raster, i.e., all voxels within a region defined by the tessellated 3-D contours were painted with a unique voxel value. This value is used as an index into a list of structure identifiers. The volumetric representation of the atlas was then mapped into the stereotaxic space manually in a two step process: the atlas was first linearly fit before being non-linearly adjusted in stereotaxic space.

1) In order to accomplish the first step, a number of homologous landmarks were identified to calculate the linear transform needed to map the atlas into the stereotaxic coordinate system. The atlas was then resampled along the axes of the VBSM volumetric data set so that it could be directly compared to the latter.

2) Since the atlas was defined on a single subject, it did not perfectly outline structures on the average model and therefore had to be warped manually to fit the model in a second step. A non-linear warping procedure was used to address this misregistration and find the transformation that most smoothly maps the first set into the second, and is based on the *thin-plate spline* (Duchon, 1976; Bookstein, 1989; Evans *et al.*, 1991b) for interpolating surfaces over scattered data (see page 10). In order to avoid pixel drop-out, i.e. missing samples in the target volume which occur when mapping source voxels into target space in a forward transform, the (numerically determined) inverse transformation was applied to the positions of the points in the target volume, in order to locate the point in the source volume that is to be mapped to it. The target volume was then resampled using nearest-neighbour interpolation to maintain continuity among the voxel labels from the non-linearly warped source volume.

For this task, homologous points were chosen in an iterative fashion. In a first pass, fifty-four anatomical landmarks were chosen on both the atlas and the average MRI volume. These points were distributed on the cortical surface, cerebellar surface, ventricular surface, centers of basal ganglia structures and points within the brain stem. The atlas volume was resampled and compared visually to the average MRI. In following iterations, existing landmarks were adjusted and additional landmarks were placed in regions that were not matched satisfactorily in order to bring them into registration.

The VBSM-VOI in stereotaxic space now implicitly segments and delineates structures of interest for any image volume mapped into that space. Figure 6.2 shows sagittal planes though one hemisphere of the VBSM-VOI. Transverse slices though the MRI atlas with the overlaid VBSM-VOI are shown in Fig. 6.2. It is important to emphasize that the VBSM-VOI takes no part in the registration process, and is only used to define the structures to be segmented. Hence, any atlas defined on the average MRI volume in stereotaxic space

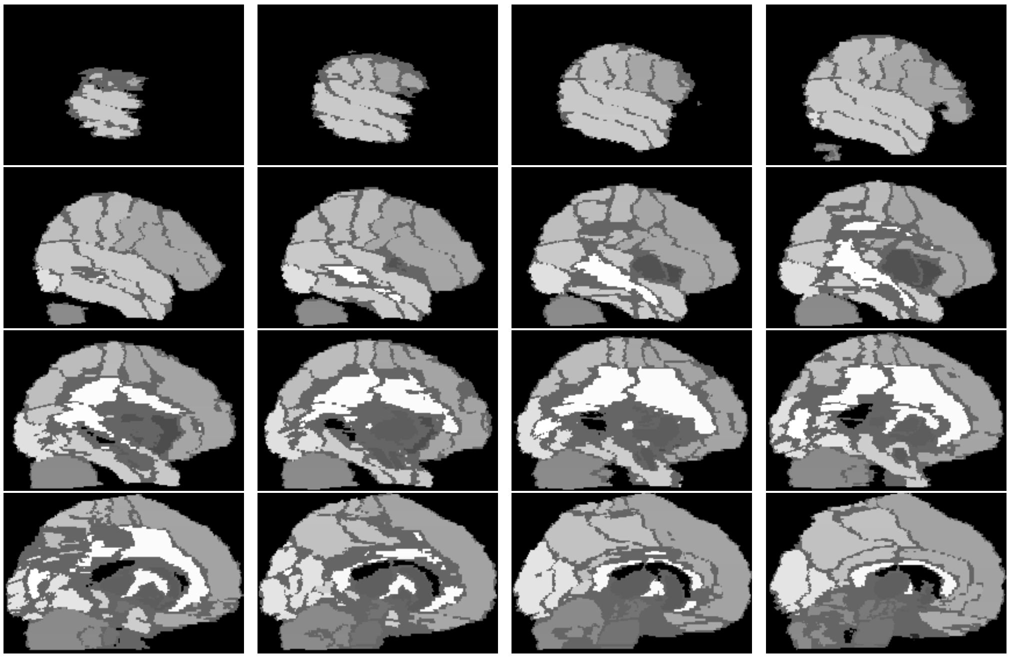

Figure 6.2: VBSM-VOI.
Sagittal tomographic planes through one hemisphere of VOI atlas of the VBSM from most lateral to 6 mm from the midline (x=66 to 6mm, in steps of 4mm). Note that while 2-D slices are shown, each VOI is defined in 3-D.

can be used to accomplish the segmentation and multiple atlases can be used together and compared.

6.3 Stereotaxic segmentation

After application of the automatic linear stereotaxic mapping procedure, a given brain volume is resampled into the standard space. The geometric VBSM-VOI yields the best approximation for structure position and orientation in the stereotaxic coordinate system, within the limits of an affine transformation. By applying the inverse transformation to the atlas contours, neuro-anatomical structures can be outlined in the native data set. Even though many regions are correctly identified (i.e., structure names are assigned to correct regions), their borders are not properly delineated by the contours mapped through a simple linear transformation (see Fig. 6.4). For the most part, these errors are due to the limitations of the affine transformation and to the significant amount of inter-subject morphometric variation. Higher order transformation functions for the stereotaxic mapping must be used to account for these complexities.

Two methods have been implemented in our lab to incorporate non-linear warping

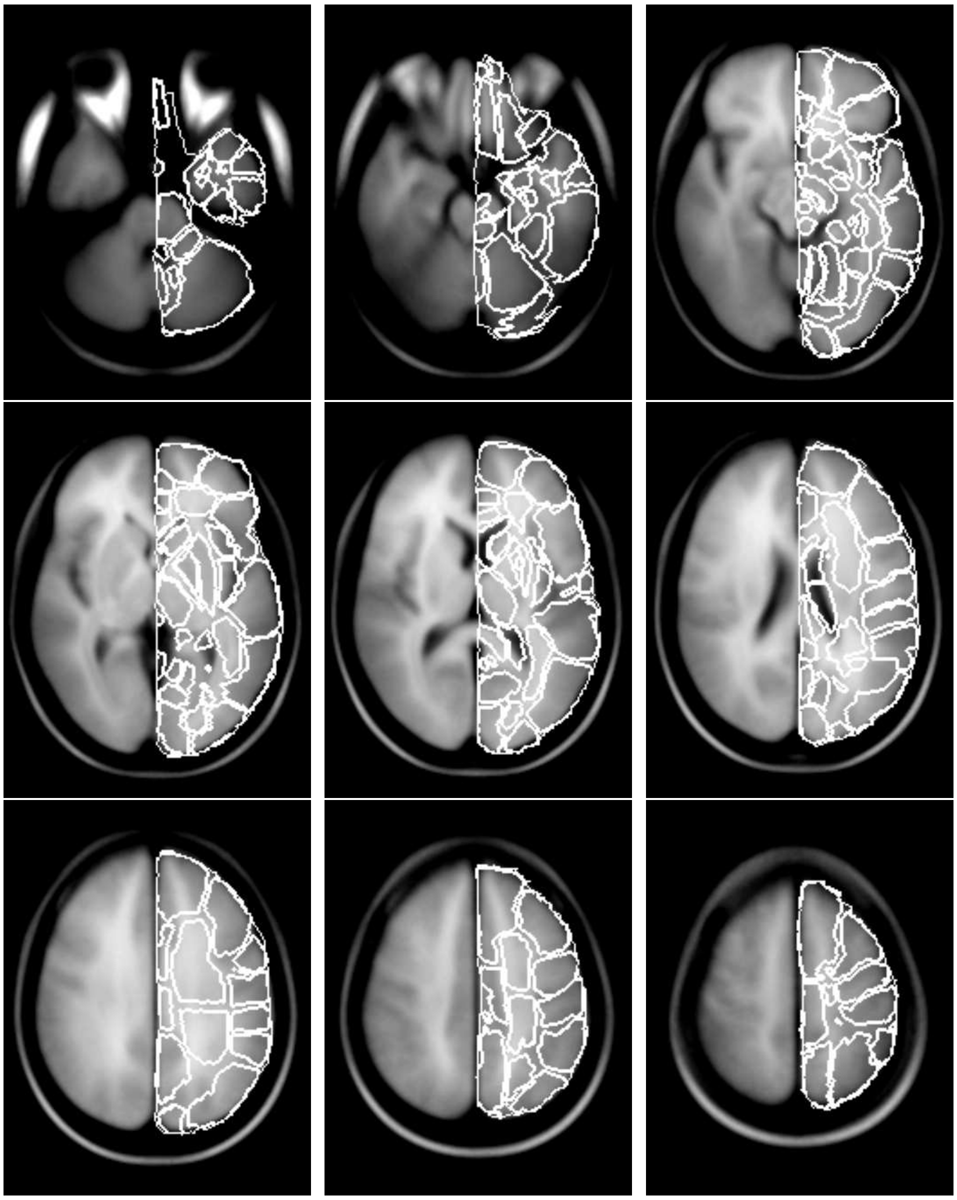

Figure 6.3: VBSM-VOI overlaid on the MRI-atlas.
Transverse slices through the VBSM showing the VBSM-VOI overlaid on the VBSM-Atlas (z=-33 to 55mm, in steps of 11mm).

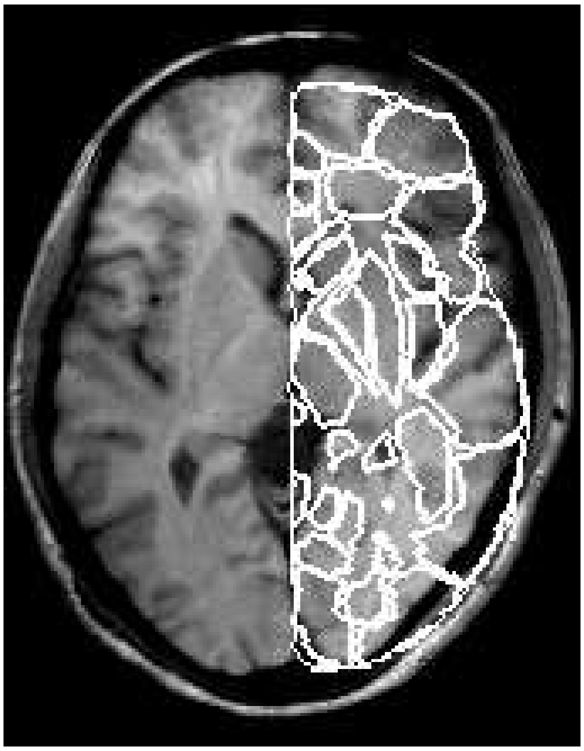

Figure 6.4: Segmentation using linear transformation only

This image shows the typical segmentation of a single transverse slice of an MRI volume, after the contours of the VBSM-VOI have been mapped through the linear transformation. While the centroid of each region of interest is near the structure that it is to segment (thus identifying it correctly), the contours do not delineate the regions well. This is due to the limitations of the affine transformation when trying to account for anatomical variability. The most notable differences are in the region of the occipito-parietal cortex and in the insular cortex regions.

into the atlas mapping procedure. While both methods use the geometric VOI atlas of the VBSM, the first requires manual intervention and the second is fully automatic. The first approach did not take explicit advantage of stereotaxic concepts and was used to match atlas to brain regardless of the original orientation of either. The second method, now adopted as our laboratory standard and forming the major theme of this thesis, was developed in large measure to overcome deficiencies that became evident in the first approach.

6.3.1 Landmark-driven segmentation

In order to fit the atlas to a given brain, and thus use the contours defined in the atlas to segment structures in the data volume, the thin-plate spline procedure was implemented to deform the entire 3-D atlas under landmark-driven constraints (Evans *et al.*, 1991b). For initial analysis, 26 landmark points were identified in each of 16 MRI volumes. The points were distributed throughout the basal ganglia, brain-stem, ventricles and cortex and each point was identified by its Talairach atlas coordinates for consistent reference between subjects. The non-linear transformation defined the 26 points on each target MRI/ MRI atlas pair was used to deform atlas and fit it on the subjects MRI volume. Fig. 6.5 shows a typical slice through one of the matched VOI/MRI volumes at the level of the peri-ventricular grey-matter regions. The good quality of the fit in the central brain regions is due to the number of clearly identifiable landmarks in the peri-ventricular region (caudate, thalamus, AC, PC, genu and splenium of the corpus callosum, anterior and posterior poles of the lateral ventricles). However, the degradation of the fit near the cortex is the result of the lack of landmarks in that area, due to the difficultly in reliably identifying cortical points.

Mapping the VOI atlas onto a single brain volume by both linear and non-linear solutions allows a direct comparison of the two solutions for that volume and is indicative of the extent to which non-linear techniques are necessary for automatic identification of brain regions. However, since each brain is of different size and orientation, pooling of results across subjects is not straightforward. By first fitting with a non-linear solution, and mapping back into stereotaxic coordinate space with the inverse linear solution, all measurements could be established in the master frame and the two methods compared. The centre-of-gravity (COG) for each of the structures of the VOI atlas was used to assess the magnitude of the non-linear deformation required throughout the brain volume. The results shown in table 6.1, indicate that a substantial component of normal anatomical variation cannot be accommodated by the linear model alone. The overall 3-D r.m.s. COG shift of 6-7mm is considerable and problematic not only for segmentation, but for all applications based on the classical stereotaxic concept.

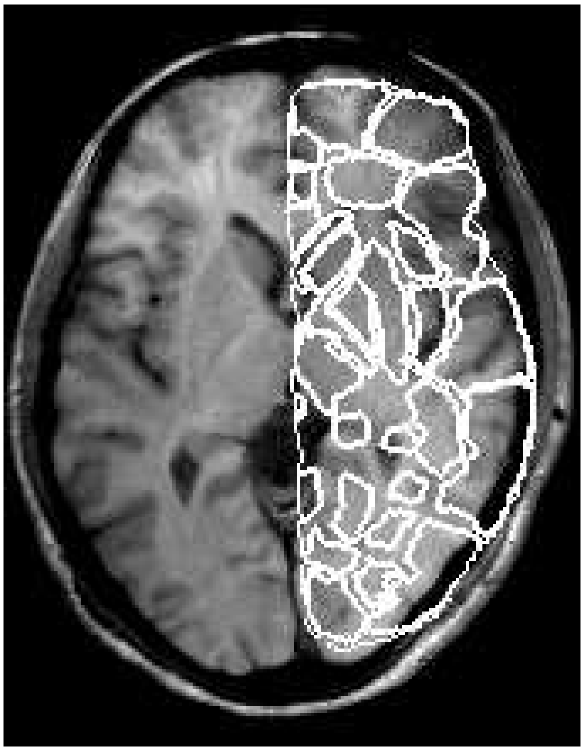

Figure 6.5: Segmentation using 3-D thin-plate spline transformation

This image demonstrates segmentation achieved by mapping the contours of the VBSM-VOI have been mapped through a manually defined TPS transformation. The segmentation of basal-ganglia structures has improved over the linear segmentation due to the number of landmarks in that region. Also, landmarks at the anterior temporal poles and in the insular region have improved the segmentation near the insular cortex. The lack of landmarks in the area of the occipito-parietal cortex has left the TPS to assume its minimum energy configuration in that area, which does not correspond to a proper segmentation in that region. This indicates the dependence of this type of manual warping technique on the number and position of the chosen landmarks.

Distance (R)	SD	Region
6.34	3.15	Superior frontal gyrus
7.69	2.27	Middle frontal gyrus
5.86	2.40	Inferior frontal gyrus
5.83	2.14	Pre-central gyrus
5.05	2.35	Post-central gyrus
4.50	1.91	Superior temporal gyrus
4.47	1.65	Middle temporal gyrus
5.19	2.07	Inferior temporal gyrus
5.38	2.19	Amygdala
5.22	2.26	Hippocampus
9.18	3.52	Head of Caudate Nucleus
7.14	2.99	Putamen
7.80	3.15	Globus Pallidus
7.63	2.94	Thalamus
6.52	3.16	Total (N.B. 60 structures)

Table 6.1: Estimate of residual anatomical variability.

The root mean squared distance (mm) of regional 3-D centre-of-gravity from target position following linear warping of VOI atlas to target space indicates residual anatomical variability not handled by the linear model (see text). 60 regions in total (14 shown) ; 16 subjects. (Evans *et al.*, 1991b)

6.3.2 Automatic segmentation

The prohibitively large number of landmarks necessary, the subjectivity inherent in their selection, and the dependence of the resultant deformation on their distribution forced us to look into fully automatic techniques for atlas-based segmentation. Automatic feature matching is proposed in this thesis as an alternative approach which avoids the potentially time-consuming landmark tagging. In this context, local properties of the image, such as edges and zones of relatively homogeneous intensity, are extracted automatically as features by neighbourhood operators. They are matched to similar features of the stereotaxic VBSM, using the local correlation and non-linear deformation procedure described in the previous chapter. Once mapped into stereotaxic space, the labels in the geometric VBSM atlas identify neuro-anatomical structures by name and the atlas contours delineate each structure region. Application of the inverse transformation segments structures in the native space. Fig. 6.6 shows a typical segmentation of a transverse MR image using the automatic method.

Matching a source MRI volume to raster data of the VBSM is key to the reversal of the traditional model-based segmentation strategy. Instead of fitting contours from the atlas to the MRI data to achieve segmentation, similar image-based volumetric features are matched between the two data sets and borders of the neuro-anatomical regions are inferred from the geometric information in the VBSM. This strategy reversal achieves three important benefits over previous methods:

1. The extra level of abstraction needed to fit geometric contours to dissimilar raster data is removed, since both model and data have the same representation as well as similar contrast, sampling and noise characteristics.

2. Subtle edges derived from the MR volume are registered with corresponding edges in the segmentation model, eliminating the need for interpolation between cortical and ventricular structures and allowing the positions and surfaces of anatomical regions to be inferred from the model.

3. Most importantly, the segmentation procedure presented here is atlas-independent. The atlas data of the model is not used in any way to estimate the registration transformation between data and model. Therefore, <u>any</u> atlas defined on the stereotaxic brain volume can be used to achieve the segmentation. Therefore, it is possible that more than one label may be associated with a single voxel to allow for multiple co-existing atlases, for a hierarchical nomenclature, or to account for partial volume effects due to multiple structures intersecting at a single voxel.

This segmentation paradigm is very general. Not only is it atlas independent, but the procedure does not require the MRI atlas to be used as the source for the invariant features in the VBSM-IIF. For example, an MRI brain volume from a single subject can be used to define the target space, and the required feature volumes derived and stored for it. The non-linear registration procedure recovers the spatial deformation function

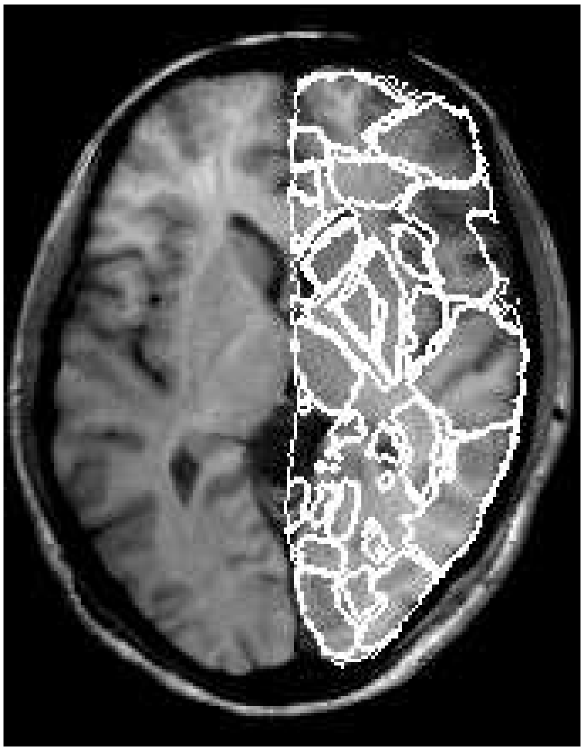

Figure 6.6: Segmentation using automatic non-linear deformations

This image shows the marked improvement in segmentation achieved by mapping the contours of the VBSM-VOI have been mapped through a the inverse of the automatically recovered non-linear deformation function. Segmentation of all structures is improved over both automatic linear segmentation and manual non-linear segmentation. Since local neighbourhood correlation is used in all regions of the volume, the automatic procedure does not suffer from a lack of information concerning the deformation as is possible with the manual procedure. The use of a non-linear deformation function can account for normal anatomical variability, unlike the automatic linear procedure. One must note that the quality of the segmentation depend on the quality of the VBSM-VOI. Since the segmentation procedure is independent of atlas used for structure delineation, improvements in the atlas can be immediately applied to segment the data, simply by remapping the new atlas through the existing non-linear transformation.

to map all voxels from any source data set on this target data set. If structures on the target brain are outlined and identified, they can be used to segment structures from other brains registered with it by mapping voxel labels through the inverse transformation from the target volume to the second data set. This method functions well with the following caveat: it suffers from the same drawback as the original Talairach atlas, namely all segmentation results must be interpreted in function of the single selected target brain. However, using this type of target permits the definition of some cortical structures that cannot be resolved on the 305-brain MRI atlas due to high morphometric variability.

6.4 Summary

Considering the number of different forms of brain atlases in the literature (Talairach and Tournoux, 1988; Hanaway *et al.*, 1980; Kretchmann and Weinrich, 1986; Matsui and Hirano, 1978) that deal with gross neuroanatomy and/or functional neuroanatomy, it is clear that there is no universally accepted model for either type of atlas. Many brain structures can be easily identified by their approximate center or main axis, but the exact delineation of their borders is problematic. The segmentation method proposed here is independent of the atlas definition. A given MRI brain volume is warped to match similar volumetric features in a standardized space. Since these features are derived from an average MRI volume, they are completely independent of the segmentation atlas. Therefore, any atlas defined in this space can be used to accomplish the segmentation and replacement of one atlas model with another requires a minimum of subjective effort. Hence, as higher resolution digital atlases become available, the algorithm remains applicable and even more functional. Furthermore, the ability to correlate any volume or model within the stereotaxic coordinate system allows the direct comparison of different previously published brain atlases.

Chapter 7

Linear registration: experiments and results

7.1 Introduction

The experiments in this chapter are divided into three parts; 1) evaluation of the linear registration algorithm using simulated data and transformations, 2) validation of the method with real MRI data and 3) evaluation of the automatic stereotaxic mapping procedure and description of the refinement of the mean MRI brain for its use in the VBSM. Note that all simulations and registrations are completed in 3D using volumetric data.

7.2 Simulated data

7.2.1 Ellipsoid phantom

A program was written to generate simple 3-D voxelated phantoms of ellipsoidal and rectangular objects with arbitrary orientation, size, slice thickness, contrast and nesting, and was used to build a simple volumetric brain model containing three concentric ellipsoids for the skin surface, cortical surface and grey-white interface, and two smaller ellipsoids to represent the ventricles. The parallelepiped represents an arbitrary object. Fig. 7.1 shows a 3-D rendering of the phantom used to test the algorithm in the first experiment. The long axis of the ellipsoid measures 100mm.

Source data sets for matching experiments were created by resampling the model volume with a randomly generated affine transformations onto a 128x128x128 matrix, with 1mm^3 voxel size. The transformations contained random translations ($\pm$15 mm) and random rotations ($\pm$15 degrees) on each axis. Fig.7.2 shows a transverse slice through the model and a typical resampled data set. The difference image in Fig. 7.2-d show a

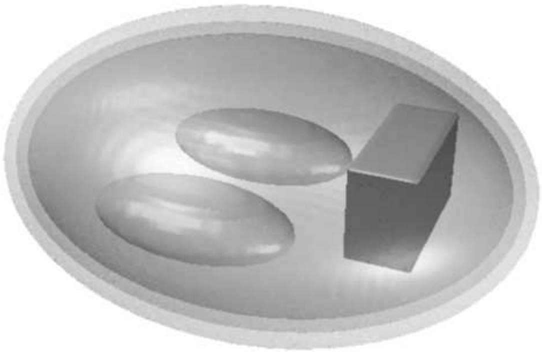

Figure 7.1: 3-D rendering of ellipsoid phantom.

The simple ellipsoid phantom contains two large concentric ellipsoids to model the cortex and grey/white matter interface. Two smaller oblong ellipsoids represent the lateral ventricles. This phantom is used to demonstrate the linear registration algorithm.

qualitative measure of the resulting registration.

Twenty random 3D linear transformations were applied to the model and two quantitative measures were used to evaluate the resulting match. The first measures differences in the recovered transformation parameters and the second measures the r.m.s. distance between homologous points. In each of the twenty cases, after application of the multiresolution registration scheme, the r.m.s. translational error from the transformation mapping was less than 0.5mm (less than half a pixel) and r.m.s. rotational error was less than 0.2 degrees around each axis.

In order to estimate mis-registration at different locations, a set specific points, organized on a 3-D lattice with 8mm spacing between nodes, were mapped through the known transform and through the recovered transform. The difference in position defines the *point registration error* and indicates the quality of the registration. On average over the 20 random transformations, the r.m.s. point registration error was 0.14mm. The maximal point registration error at the edges of the field of view was ± 0.95mm, and at the edge of the geometric model, ± 0.77mm. From these measures, we can conclude that the multi-resolution procedure can recover the transformation required to register a simple object in a noiseless data set. A more complex phantom is used in the next section.

7.2.2 Digital Brain Phantom.

In order to realistically evaluate the performance of the algorithm with respect to image intensity noise, contrast and different transformations, it was necessary to create a phan-

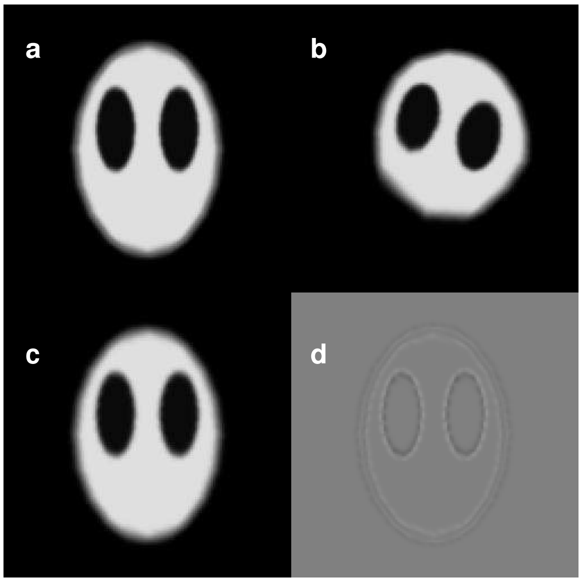

Figure 7.2: 2-D transverse slices through 3-D ellipsoid phantom.

a) Transverse slice through the model phantom. b) Same slice through same data set, after application of a random transformation. c) Shows the data set of (b) re-transformed by the recovered transformation. d) Difference image between target (a) and restored data (c). The pixel values in the difference image have been multiplied by a factor of ten to enhance the voxel intensity difference.

tom that was more complex than simple concentric ellipsoids. A detailed brain phantom was created by identifying different structures and tissue types from an MR data set of a single volunteer. The resulting labeled volume could then be endowed with the appropriate contrast for each tissue type, and noise added, to resemble a true MR acquisition. This data set provides a gold-standard with which to evaluate the linear registration algorithm.

The MR data used to create the phantom consisted of 3 volumetric data sets representing T1-, T2-, and proton density-weighted (PD-weighted) data. The T1-weighted data were acquired with a 3-D volume spoiled Gradient Echo sequence (TE=7ms, TR=23ms, flip angle 30°) yielding 180 sagittal slices with $1mm^3$ voxels. The T2-weighted and PD-weighted volumes were acquired in a second 2-D multi-slice dual echo sequence (TR=3300ms, TE=35ms for PD and TE=125ms for T2). Slice thickness was increased to 2mm to reduce noise in these two data sets. Since the three volumes were acquired with two acquisitions, there was possibility of subject movement between scans, which was in fact verified by visual comparison of the T1 and T2 data sets. Therefore, the T2-weighted volume was first registered to its T1-weighted counterpart by identification of 20 homologous landmarks on both volumes using the landmark-matching method described in (Evans *et al.*, 1989d). Once completed, the T1-weighted data set was registered to the standardized space using the same technique. All three data sets were then mapped into stereotaxic space with the same transformation and resampled using a $1mm^3$ voxel size.

After 120 voxel samples of grey-matter, white matter, CSF, muscle, fat, skin, bone and background were manually identified, a minimum-distance classifier (Duda and Hart, 1973) was used to label each voxel with one of the aforementioned classes. Since some voxels outside the skull were labelled as white, grey or CSF, a manually defined mask was used to change the values of voxel labels within the brain. A similar procedure was used to differentiate background from bone voxels. Typical slices from the resulting classified volume are shown in Fig. 7.3. While there are some small classification errors, *e.g.*, a slight overestimation of grey matter at the cortex, the phantom is more than adequate to represent of the complexity of structures within the brain for evaluation of the non-linear registration algorithm. It should be noted that it does not matter that some voxels from the original MRI data were mis-labelled by the minimum-distance classifier. The labelled volume is now a phantom defined as <u>truth</u> and is no different in principle than the arbitrarily defined ellipsoid phantom. It has the single advantage that it is a complex simulated data set which closely mimics the convoluted nature of a real brain.

In order to simulate MR data from a particular scanning protocol, the classified regions were *colourized* according the mean intensity value for the given tissue type. For example, a T1-weighted phantom has all grey matter voxels mapped to an intensity of 438, white matter voxels to 350, CSF to 137 and bone to 38. These representative values were estimated from regions of interest taken from one MRI volume from the second group of data sets used to test the algorithm. Noise is added in the last step of the creation of the phantom. For each voxel in the phantom, the coloured intensity is augmented by a randomly generated noise value from a Rayleigh distribution with a user-specified mean. Fig. 7.4 shows a typical transverse slice through the original data volumes used to create

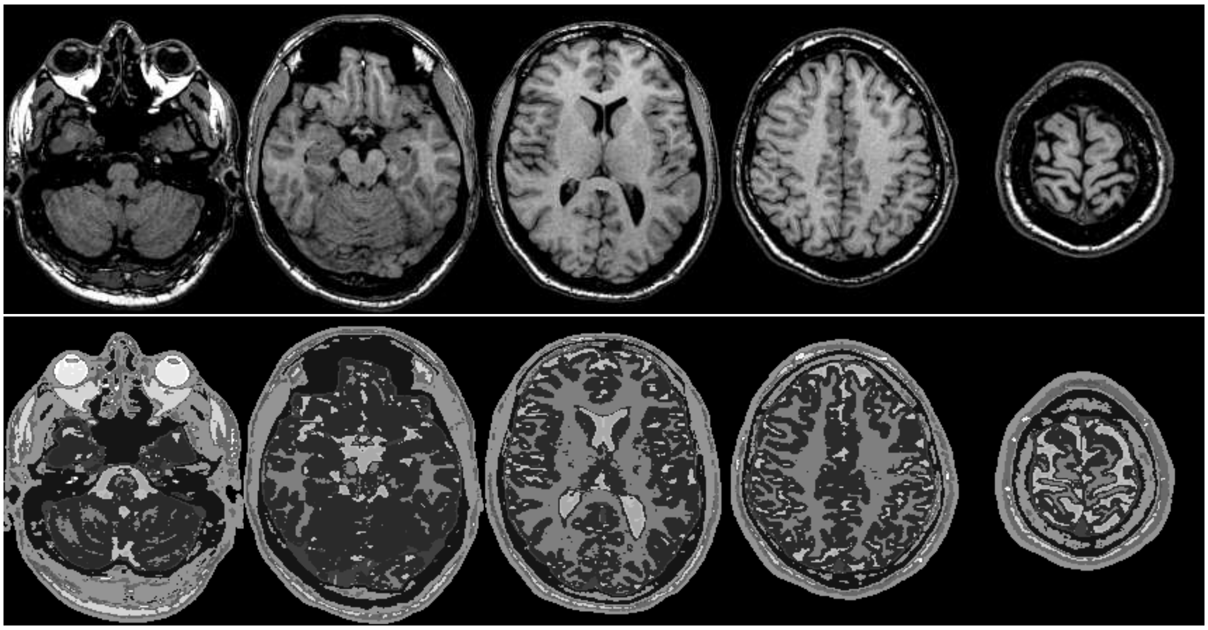

Figure 7.3: Classification result.

The first row shows 5 transverse images from the original T1-weighted volume. The second row shows the corresponding slices classified into grey-matter, white matter, ventricular CSF, sulcal CSF, muscle, skin, bone, fat and background.

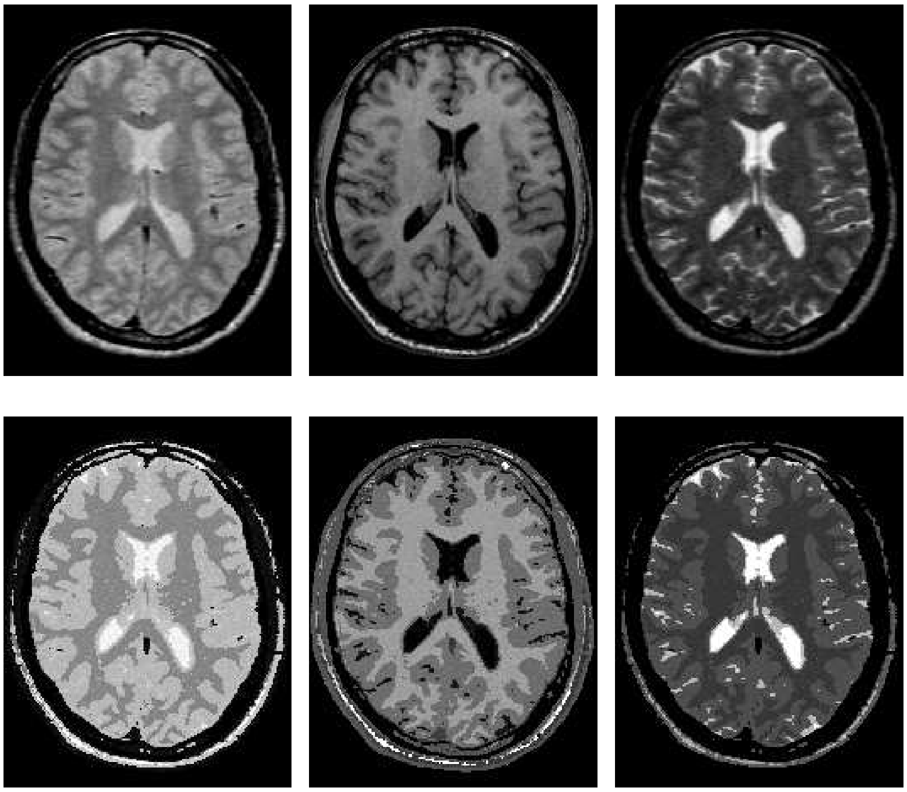

Figure 7.4: Comparison of real MRI data with phantom.
Transverse image through both real MRI data (top row) and the simulated phantom without added noise (bottom row) at the level of the lateral ventricles. From left to right: proton density, T1 and T2-weighted images. The difference is that in the latter case, all voxels within a given class are known and given the same intensity, prior to the addition of noise.

the phantom along with the corresponding simulated data.

7.2.3 Registration error vs noise

In order to test algorithm performance with respect to noise, the following methodology was used. Random transformations were applied to the brain phantom followed by the addition of noise to create source/target pairs, where the transformed volume was identified as the source and the original brain phantom as the target.

The transformation was created by generating random transformation parameters within a fixed range; the translations along the x, y and $z-$axes varied ± 10 mm, rotations around the three axes varied ± 10 degrees and scale along the axes varied from 0.9 to 1.1. These parameters were used to create the transformation matrix used to resample the brain phantom. After resampling, Rayleigh distributed noise was added to the

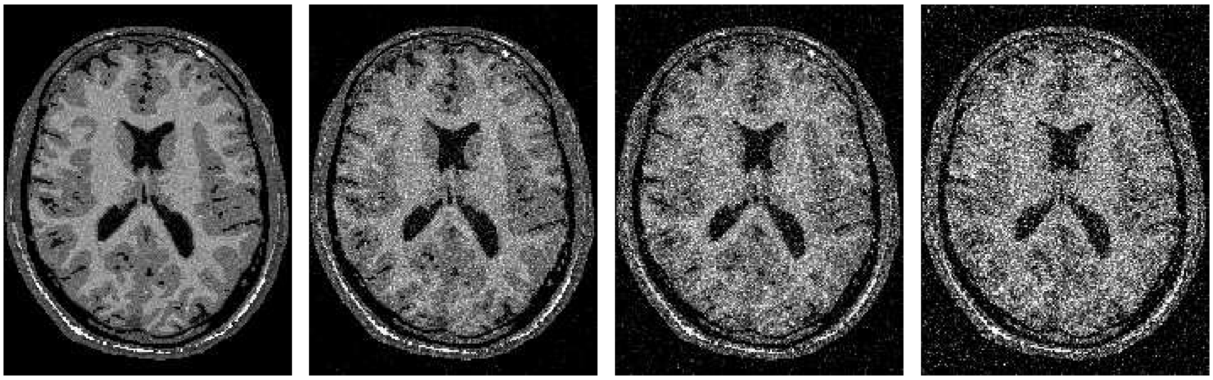

Figure 7.5: Noisy brain phantom
These images show a transverse slice through the T1-weighted digital brain phantom at noise levels of 20%, 40%, 60% and 80%, from left to right. Compare these with the equivalent noise-free version in the bottom-left corner of Fig. 7.4.

transformed volume at levels from 0% to 80% in intervals of 10% of the non-zero signal of the input volume. Note that typical MR images have a signal-to-noise ratio at least 10:1, corresponding to the 10% noise level. Fig. 7.5 shows the same transverse slice through four noise levels used in this experiment, where the level is defined to be the standard deviation of the Rayleigh distribution.

For each noise level, ten source/target pairs were generated and the automatic registration procedure applied to recover the random transformation. As for the first phantom experiment, two measures of error were calculated: parameter error and the r.m.s. between homologous points. The graph in Figure 7.6 shows that the parameter error remains quite small up to 40% noise. At that level, average scaling error is less than 0.03±0.28%, translation error is 0.08±0.16mm and rotational error is -0.06±0.67 degrees. At 50% noise, the registration of only one of the ten source/target pairs failed, while the others were recovered with error similar to the 40% level.

The second measure of registration error is the registration residual and indicates how well homologous points are mapped to each other. For this experiment, 10 points were chosen (5 cortical and 5 internal: frontal pole, occipital pole, L/R central sulcus, vertical apex, splenium of corpus callosum, L/R head of caudate, L/R posterior pole of lateral ventricle). These points were mapped through the forward random transformation and then back through the recovered transformation onto the original brain phantom. As shown in Fig. 7.7, for each of the source/target pairs, the 3-D r.m.s. residual is less than 0.8mm until 50% noise has been added to the transformed brain volume. At 40% noise, the registration r.m.s. residual is 0.81±0.50mm and is consistent with the parameter errors measured above. At 50% noise, the error increases to 3.76±3.50mm.

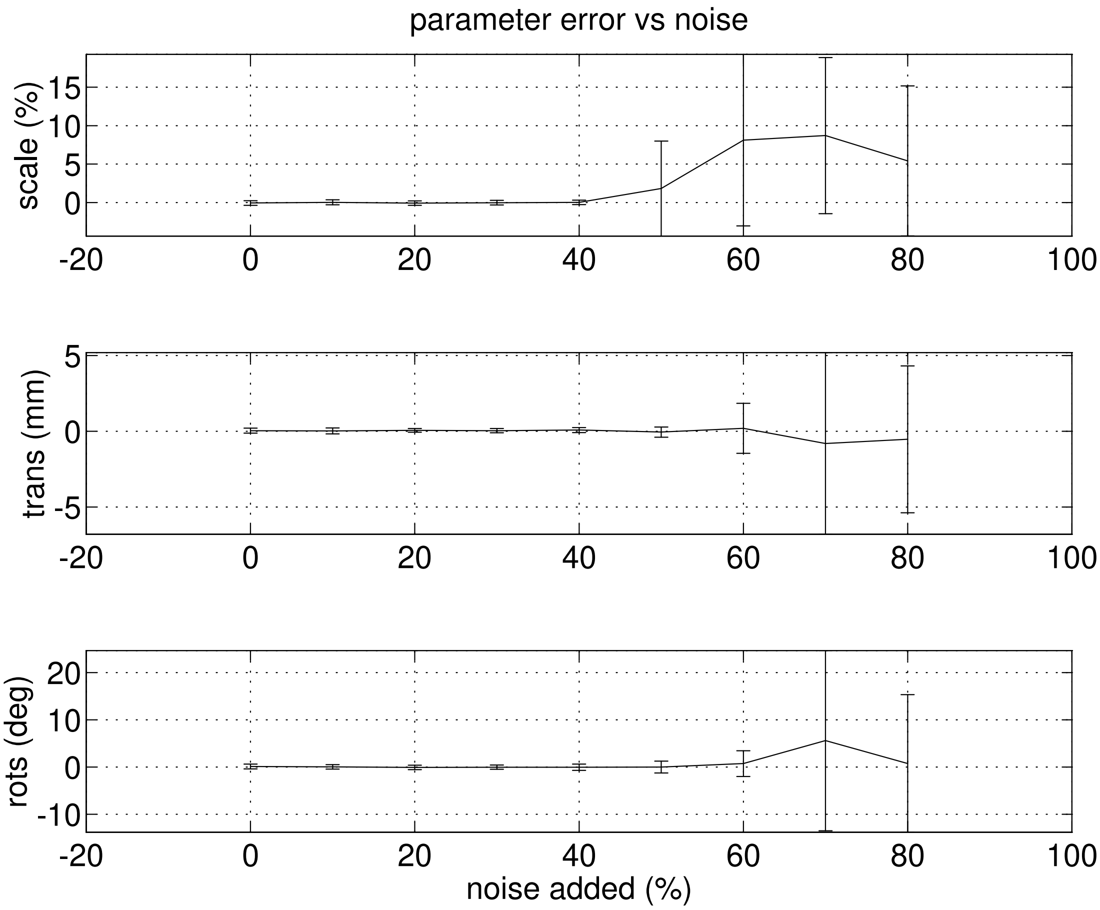

Figure 7.6: Parameter error vs noise.

These graphs indicate that the registration algorithm can deal with noise levels up to 40%. This corresponds to a signal-to-noise ratio of approximately 3:1. Since typical MRI data has a ratio of at least 10:1, the algorithm is applicable to any realistic MR volume used for neuro-scientific studies.

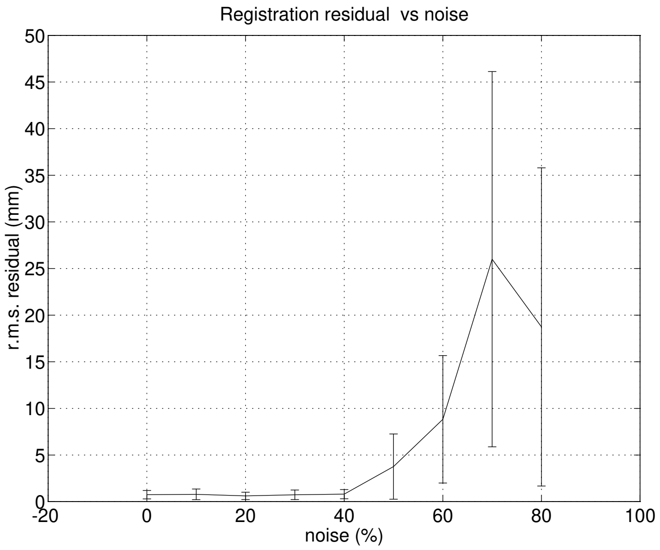

Figure 7.7: Residual error vs noise.

This graph indicates that the registration algorithm is robust against noise. The r.m.s. registration residual is 0.81mm or less until 40% noise has been added to the simulated data.

7.3 Real Data

7.3.1 Data acquisition

After simulation studies, the registration algorithm was tested with real MRI data. For this task, a set of 17 brain volumes were acquired from young volunteers (27 ± 3 years; 13 males, 3 females). These volumes form the core data base of real data for most of the analysis accomplished in this thesis.

All MRI studies for this chapter were performed on a Philips Gyroscan S15-HP 1.5 Tesla superconducting magnet system. A 3-D gradient echo acquisition, with T1-enhancement was selected for this study (TR=75ms, TE=14.1ms, FOV=250mm, slice thick=3mm, number of slices = 56, flip angle=60°) yielding a total scanning time of 37 minutes. The data set was reconstructed on a 256x256x56 matrix with 12 bits per voxel and a voxel size of 0.98x0.98x3.0mm. The imaging parameters were selected to maximize grey/white matter contrast and the signal to noise ratio to facilitate both manual and automatic structure identification while minimizing the scanning time and covering the entire brain.

A non-invasive form of head immobilization, based on a pre-formed foam pillow and a velcro strap across the forehead, was used allowing for simple setup. The absence of specific pressure points made this a relatively comfortable procedure for the short period in the scanner. The subjects (all graduate students and technicians associated with the lab) were cooperative during scanning. There were negligible artefacts in the data.

The images were reconstructed on a local VAX system and transferred via EtherNET to the NeuroImaging Laboratory. For display and analysis in the creation of the VBSM, the 12 bit/voxel intensity range was reduced to 8 bits without significant degradation of image contrast.

7.3.2 Data preprocessing.

In dealing with real data, there are a number of problems with machine-related imperfections in the image. These are dealt with by preprocessing steps as follows.

Image intensity corrections.

A procedure has been implemented to correct for the RF field inhomogeneity that is evident in the patient data as a slightly brighter diagonal band from the upper left to the lower right corner of each transverse image. This artefact is typical of the linearly polarized RF coil used in the MR scanner when the imaging experiments were performed. This artefact can be modeled as a multiplicative intensity variation. There are a number of methods in the literature that address this problem: unsharp masking, phantom correction, homomorphic filtering (Bracewell, 1978; Pratt, 1978) and unsharp blurring (Gohagan *et al.*, 1987).

The intensity inhomogeneity in each slice is addressed using the unsharp blurring technique, a method similar to homomorphic filtering except that processing involves dividing images by their neighbourhood average rather than subtracting logarithms. The algorithm has been implemented within the MSP program running on the SPARC 1+ as a preprocessing step (Collins, 1990) before the volume is brought into stereotaxic space for comparison with the model.

This preprocessing step is no longer strictly necessary, since the final steps of the registration procedure are based on gradient magnitude data and are therefore much less sensitive to this type of intensity artefact.

Geometric corrections.

In some systems, geometric distortions in the images require correction before processing. However, for the system used here, measurements of the dimensions of an MRI calibration phantom indicate negligible distortions through the central portion of the MRI imaging field, thus eliminating the possibility of significant geometric artifacts (Drangova, 1987; Peters *et al.*, 1988).

	RMS error				*standard deviation*		
	x-axis	y-axis	z-axis	avg.	x-axis	y-axis	z-axis
exp. 1							
rot. (deg)	0.0766	0.0958	0.0773	0.0832	0.0487	0.0612	0.0466
trans. (mm)	0.0959	0.1027	0.1081	0.1022	0.0576	0.0675	0.0639
exp. 2							
rot. (deg)	0.0787	0.1010	0.0874	0.0890	0.0500	0.0640	0.0513
trans. (mm)	0.1055	0.1034	0.1218	0.1102	0.0628	0.0678	0.0760
scale	0.0012	0.0011	0.0018	0.0014	0.0009	0.0007	0.0011

Table 7.1: Simulation of intra-subject registration.

Twenty random transformations were applied to each of 17 subjects. The automatic linear registration procedure was used to recover the transformation. The Table shows the errors in affine parameters for rotation and translation only (Exp. 1), and with scaling (Exp. 2).

7.3.3 Simulated intra-subject registration

In an initial experiment to establish intra-subject registration, 20 random linear transformations (only translations and rotations) were applied to each of the seventeen MRI volumes of the head from the normal volunteers. Errors were measured in the same fashion as for the ellipsoid and the digital brain phantoms described earlier. The errors in the parameter recovery are summarized in Table 7.1. The average r.m.s. rotational error was 0.083 degrees while the average r.m.s. translational error was 0.102 mm. The r.m.s. point registration error was less than 0.2 mm. The results were better than that obtained with the phantom, because in that case partial volume effects were not accounted for.

This experiment was then repeated with anisotropic scaling included with the previous rotations and translations in the 20 random linear transforms. The results of the nine parameter registration are still very good as shown in the second half of Table 7.1. The average r.m.s. error for rotation increased slightly to 0.089 degrees. Translational error increased from 0.102 to 0.110 mm. The average r.m.s. scaling error was 0.14%. The point registration error, as defined above, was still less than 0.2 mm. These measurements determine the lower bound on the error using this form of 3-D registration.

7.3.4 Inter-subject registration

For inter-subject registration, each of the 17 brain volumes in turn was identified as the target and the other 16 volumes registered to it. Figure 7.8-a shows the same transverse slice through nine of the seventeen volumes used in the experiment. Even though the laser cross-hairs were used for alignment on the scanner, this method of pre-scan registration achieve little success. In Fig. 7.8-b, the upper-left volume was selected as a target, and

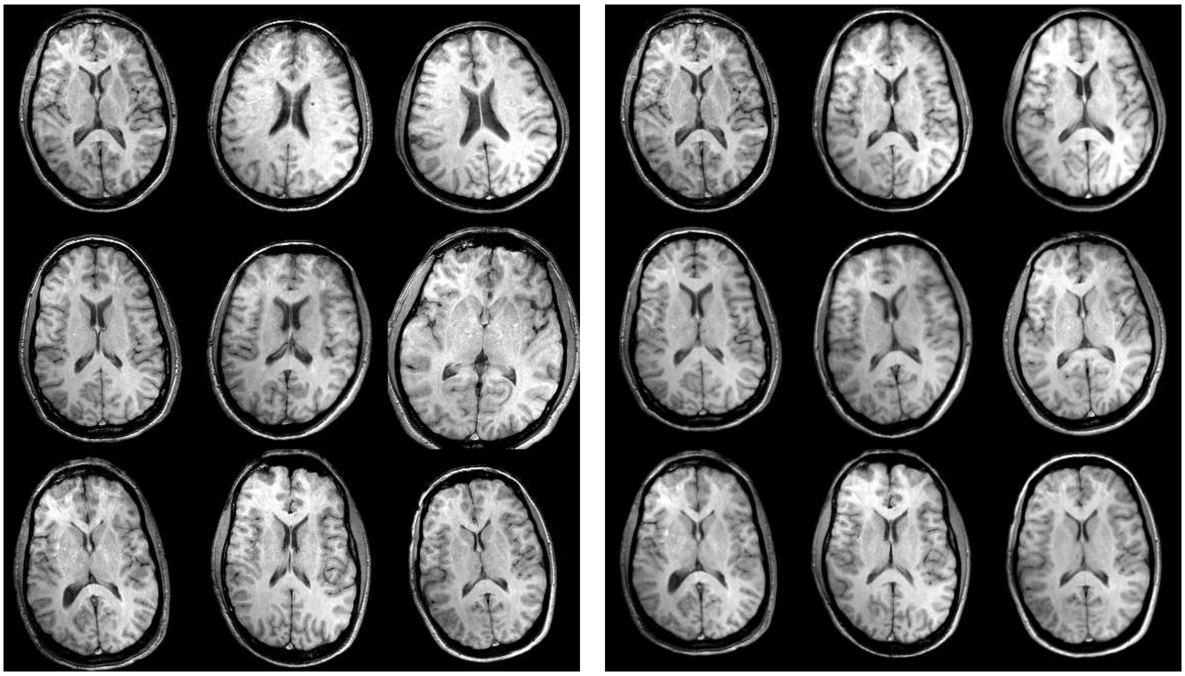

Figure 7.8: Before and after inter-subject registration
The group of images on the left shows the same transverse slice through 9 of the seventeen volumes used for inter-subject experiments in this chapter, before application of the automatic registration procedure. In the group on the right, the upper-leftmost volume was used as a target and the other sixteen (only eight shown) linearly registered to it. One can see that the automatic registration procedure accounts for the position, orientation and scaling needed to align the data sets.

the other 16 volumes registered to it, eight of which are shown in the figure. Visually, the automatic procedure appears to have recovered the transformation necessary to register the volumes together. Here, the registration error can not be estimated as before, since the source and target volumes represent different brains. Therefore, there is no generalized correct registration solution since the quality of the match will depend on the task requiring the registration. Other objective measures that can estimate the registration residual based on information derived from the two volumes must therefore be determined. One such measure is similar to the point registration error and is described below.

Landmark points and registration residual

One of the goals of registration is the alignment of similar structures. Therefore, the 3-D distance separating similar structures after registration is an objective, quantitative estimate of the quality of the registration. Anatomical landmarks are often used to establish this measure because they can be identified across a number of different subjects and a large number of them can give a reasonable estimate of the goodness-of-fit in registration. In this thesis, the *registration residual* is defined as the distance, in the target space (be it the coordinate space of an individual data set, or the stereotaxic space), between a specific landmark point and its homologue mapped into the same space by the recovered transformation. This is not only registration error, but an indication of the second order difference in anatomy between subjects.

For this and the following experiments, software was developed to allow real-time tri-plane display of tomographic data so that the user may roam through each volume in 3-D (Evans *et al.*, 1989a; Evans *et al.*, 1991a) and identify 48 anatomical landmark points on each of seventeen data sets. These landmarks were distributed over the cerebrum, cerebellum and brain stem and are separated into five groups (some points occur in more than one group):

1. longitudinal structures(n=14): (*e.g.* the ventral aspect of the genu and splenium of the corpus callosum, the anterior and posterior commissures, the inter-thalamic adhesion, the occipito-cerebellar junction at the mid-line, the inferior ventral aspect of the pons-midbrain cleft and the most superior and most inferior aspects of the vermis.)
2. basal ganglia (n=11): (*e.g.* the ventral aspect of the genu and splenium of the corpus callosum, the anterior and posterior commissures, the inter-thalamic adhesion, the center of the head of caudate and the center of the thalamus.)
3. ventricular structures (n=5): (*e.g.* both anterior horns, both posterior horns and the posterior notch of the 4th ventricle.)
4. cortical structures (n=18): (*e.g.* the most anterior point of the temporal poles, the most posterior aspect on the occipital poles, the point of intersection of major fissures such as the central sulcus with the longitudinal fissure and the parieto/occipital sulcus with the longitudinal fissure and the temporo-cerebello-occipital notch.)
5. cerebellar structures (n=18): (*e.g.* the occipito-cerebellar junction at the mid-line, the most superior, the most inferior and the most posterior aspects of the vermis, the most superior, the most inferior, the most posterior and the most lateral aspects of both cerebellar lobes and the centers of the cerebellar peduncles.)

These points represent a compromise between (i) the requirement for a large number of neuro-anatomically well identified landmarks throughout the brain volume, and (ii) the practical limitations on identifying reliable structures, particularly on the cortical surface.

Automatic inter-subject registration

Returning to the inter-subject registration experiment, over the 272 trials (17 targets $\times$ 16 data sets), the linear transformation recovered by the multi-resolution registration procedure yielded a 3-D r.m.s. registration residual of 6.65 mm when estimated over the 48 points above. This value is in good agreement with the mean anatomical variability of 6.52 mm in Table 6.1 of section 6.3.1, determined with a different technique, different subjects and different landmark points.

Manual inter-subject registration

In order to compare this result to manual landmark-based registration, the 48 landmarks were submitted to a homologous landmark registration procedure based on the *Procrustes* algorithm (Langron and Collins, 1985; Sibson, 1978; Sibson, 1979). This procedure requires two sets of points and finds the *best* solution in the least squares sense by minimizing the r.m.s. distance between all paired points. The optimal transformation is found in two stages. First, the translational component, **T**, of the transformation is obtained by calculating the difference between the COGs of the two clouds of points. Both sets of points are then translated so that their centroids coincide with the origin. Secondly, the rotational component, **R**, of the transformation is determined by minimizing the residual,

$$||P_1 - RP_2|| \tag{7.1}$$

where P_1 and P_2 are the $3 \times N$ matrices representing N points in the source and target volumes, respectively. The solution to this problem is obtained by calculating the singular value decomposition (SVD) of $P_2'P_1$, where $'$ indicates matrix transposition. If,

$$U(M'P_1)V = \sigma = diag(\sigma_1 ... \sigma_N), \tag{7.2}$$

is the SVD of $P_2'P_1$, then **R** above is given by $U'V$. A global scale is extracted from the ratio of the trace of $P_2\mathbf{R}P_1'$ divided by the trace of P_2P_2'.

Comparison manual vs automatic inter-subject registration

With isotropic scale the Procrustes r.m.s. residual was 5.94mm. This value obtained for direct landmark-matching is slightly smaller than the 6.65mm obtained by the automatic indirect image-matching approach. This is expected, since the points used to define the Procrustes transformation are the same as those used to measure the error. Nevertheless, the values indicate that the automatic feature-matching approach performs as required.

7.4 Stereotaxic mapping

As described in section 4.6, the automatic registration procedure can be used to map individual data sets into stereotaxic space prior to analysis. Four experiments were performed to evaluate the performance of the general algorithm when applied to this particular task. These experiments

- demonstrate the improved registration using the gradient magnitude compared with just using voxel intensity;
- compare the automatic method to a manual registration based on mapping of homologous landmarks between the two volumes;
- compare the automatic method to a second manual registration procedure, where the user begins by identifying points along the AC-PC line and then determines rotation parameters by locating points in the inter-hemispheric plane and scaling by estimating the brain extents;
- explore the algorithm's robustness with respect to missing data.

A typical automatic stereotaxic registration is shown in Fig. 7.9.

For the results quoted below, the following coordinate system convention is used: $x-$axis in the LR (left-right) direction (positive towards the right), $y-$axis in the PA (posterior-anterior) direction (positive anteriorly), $z-$axis in the CC (caudo-cranial) direction (positive superiorly) and the origin was located at the intersection of the center of the anterior commissure with the inter-hemispheric plane.

7.4.1 Gradient magnitude vs. voxel intensity

For each data set, three transformations were computed at a final scale of FWHM=8 mm . The first was based on the intensity fit alone (noted *head* in Table 7.2), the second on the brain-masked intensity fit (noted *brain*) and the third on the masked gradient fit (noted *edge*). For each transformation, the centroid and the average standard deviation for the different groups of landmark points were calculated and are shown in Table 7.2.

When comparing the edge-based registration to that accomplished by masked-intensity correlation for each group of points, *i*) mapping by the transformation recovered by the edge fit reduced their spread as indicated by the standard deviation in stereotaxic space; *ii*) the bias of points away from the mid-line ($x = 0$) was found to be reduced as well, *iii*) the bias in rotation about the z-axis indicated in Table 7.2 was removed, aligning the longitudinal fissure with that in the model in all cases tested, and *iv*) that the bias in scale in the same direction was removed in the three cases where it was over-estimated.

7.4.2 Automatic registration vs. land-mark based technique

The hypothesis tested here was that the automatic multi-resolution methodology could estimate the transformation necessary to bring a volumetric T1-weighted MR data set into stereotaxic space with errors similar to, or smaller than, those estimated by manual registration methods that use homologous landmark matching.

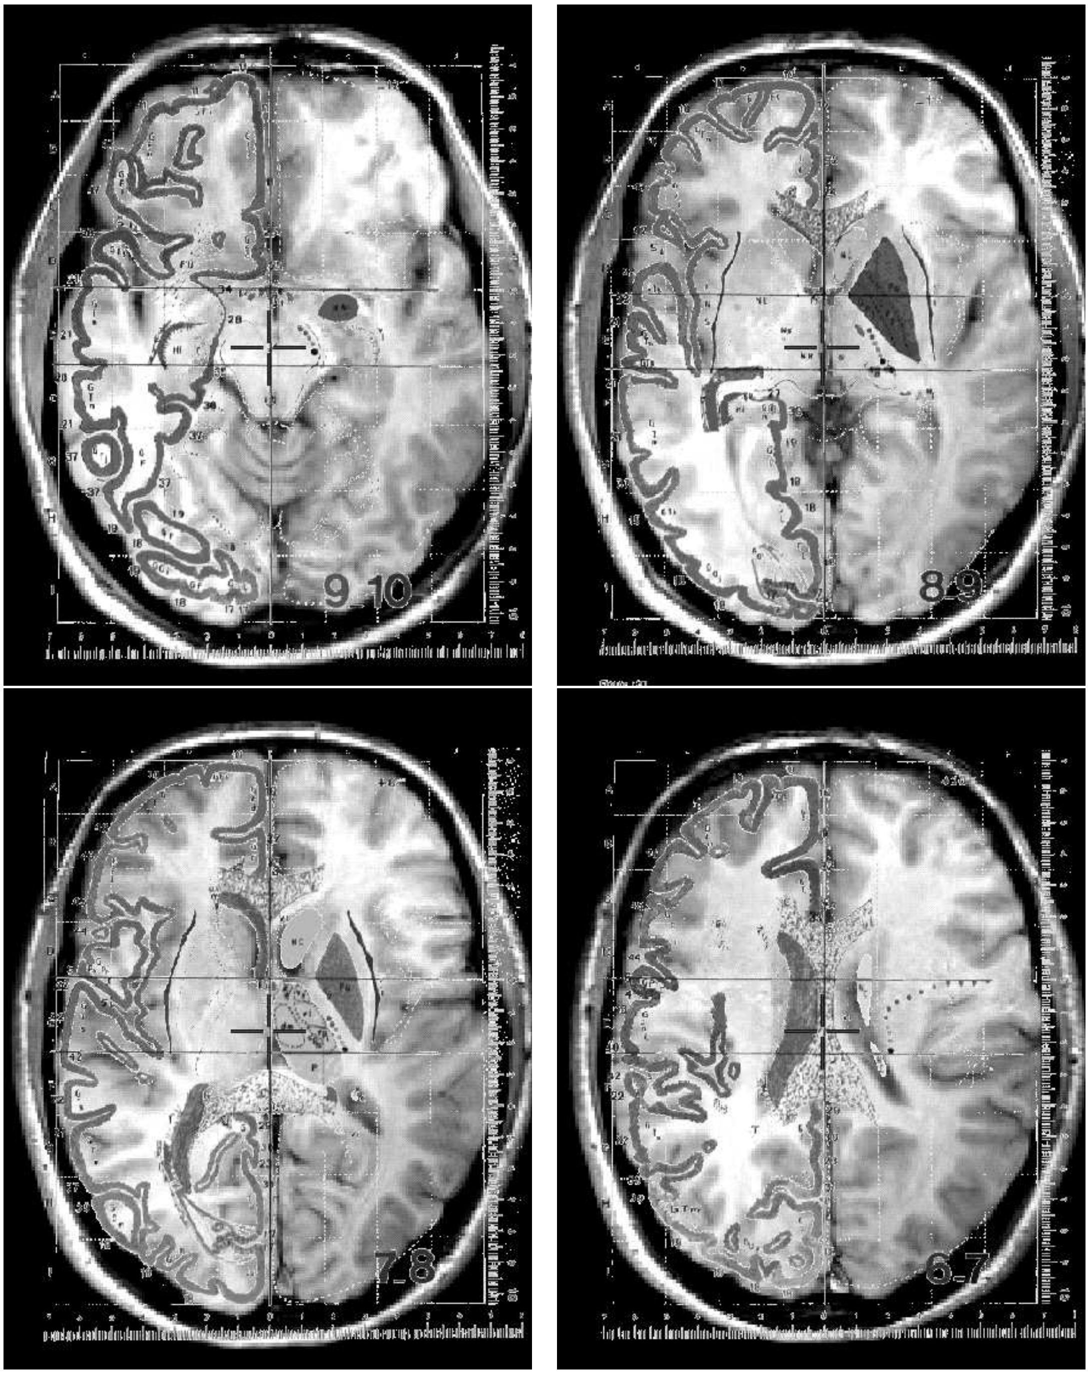

Figure 7.9: Typical automatic stereotaxic registration.
Four transverse planes (z=-10,0,10,20mm) through a typical data set after automatic registration with the stereotaxic Talairach space with the corresponding atlas slices overlaid. While these images show the quality of the automatic alignment, they also indicate the need for some form of non-linear mapping to complete the alignment of specific regions near the edge of the cortex and internal structures such as the ventricles.

total of 48 points	*mean (mm)*			*standard deviation*		
	head	brain	edge	head	brain	edge
x:	1.2	1.6	-0.1	2.37	2.08	2.00
y:	-38.6	-37.7	-36.5	4.57	4.11	4.07
z:	-3.2	-3.5	-5.5	5.05	4.11	3.62
3-D:				7.21	6.17	5.80

Table 7.2: Comparison of fitting strategies.

Automatic intra-subject registration (16 subjects) in stereotaxic space. Comparison of full volume intensity correlation (head) vs. brain-masked intensity correlation (brain) vs. brain-masked gradient correlation (edge) based on registration residuals of symmetrically (left-right) disposed landmark points. Hence $\bar{x}$ should $= 0$ ($\bar{y}$ and $\bar{z}$ have no significance). Note the reduced 3-D r.m.s. distance among homologous landmarks when comparing head, brain and edge fitting strategies.

For each of sixteen data sets, two linear transformations were calculated: one by the automatic method (T_{auto}) and the other that minimized, in the least-squares sense, the distance between manually identified homologous landmark point pairs (T_{man}) (Golub and van Loan, 1983) described in section 7.3.4. Both transformations are defined by nine parameters: three scale factors, three translations and three rotations. In order to directly compare the transformation parameters, the center of rotation and scaling was constrained in T_{man} to be equal to that found in T_{auto}. This constraint was somewhat biased against T_{man}, since the true least squares solution imposes its own center. However, the center of the of landmark points was close to the one imposed, since the points used are well distributed around the cerebral volume. There was only a 1.1 mm r.m.s. difference in 3-D position between the mapping of these points between the constrained and unconstrained transformations. Hence, there is no important difference between the the rotation and scaling parameters between the constrained and true least-squares solutions.

For each trial, three difference measures were calculated. The first two compared the automatic with manual methods without defining which was more correct, while the third measured how well either method accomplished the required task of registering target points. These measures were based on:

1. differences between the transformation parameters (rotation, translation and scaling) recovered by T_{auto} and T_{man};

2. average difference in position of landmark points (listed in section 7.3.4) mapped into stereotaxic space by T_{auto} compared to T_{man};

3. a comparison of registration residual for mapped landmark points with respect to target points for the two methods.

Parameter differences

Using a two-tailed Student's t-test for paired samples, there was no significant difference in the parameters recovered by the automatic algorithm when compared to the manual method ($p > 0.1$ for all parameters except $y-$scale, where $p > 0.05$; deg. freedom = 15; For the x, y and $z-$axis, respectively: Student's t=(0.927, 0.030, 0.351) for rotation, t=(0.189, 0.770, 0.335) for translation and t=(0.891, 1.773, 1.082) for scaling.

When the registration parameters for T_{man} and T_{auto} are compared for each data set, the r.m.s. rotational difference was 2.1, 0.75 and 1.48 degrees around the x, y and $z-$axes. The inter-hemispheric fissure strongly constrained rotation about the y and $z-$axes. Most other strong features, such as the brain and ventricular surfaces, are aligned tangentially rather than radially to the x-axis of rotation, and hence have a less constraining influence on this parameter. The r.m.s. difference in translation was 0.49 , 0.95, and 1.0 mm in the x, y and $z-$directions. Once again, the gradient edges along the longitudinal fissure fixed the LR position. The r.m.s. change in scaling was less than 2.5 percent along the x and $y-$axes and 4.8 percent along the $z-$axis. The scaling along the $z-$axis is worse since the cortical surface at the top of the head was missing on some data sets, due to the lack of MR slices in that region.

Mapping differences

The difference in mapping through T_{man} and T_{auto} into stereotaxic space was estimated by measuring the average distance between landmark points (listed in section 7.3.4) mapped forward by the two transformations (see Table 7.3). Note that this does not measure how well either method map the points to their target position, which is the subject of the next section.

For 48 points, across all sixteen data sets, the average r.m.s. residual differences between mappings were 1.87, 1.21 and 2.20 mm in x, y and $z-$directions, respectively (3.09 mm in 3-D). Not surprisingly, the smallest residuals were for structures near the transformation center where, for instance, the basal ganglia landmarks showed average r.m.s. values of 1.1, 0.7 and 1.4 mm in x, y and z (1.91 mm in 3-D). The largest were for cortical structures: 2.2, 1.4 and 2.5 mm , for x, y and z (3.61 mm in 3-D).

Misregistration with target points

For inter-subject registration, it is well known that even after transformation (manual or automatic), there remains a considerable difference in position for landmark points (Evans *et al.*, 1991b; Talairach and Tournoux, 1988; Steinmetz and Seitz, 1991). For each of the 48 landmarks, a target point was defined as the centroid of each ensemble of 17 landmarks (one from each subject) mapped into stereotaxic space by the automatic transformation. The average r.m.s. residual for landmark points mapped through T_{man} was 5.21mm in 3-D (3.72, 1.81 and 3.13mm , in x, y and $z-$directions, respectively). (The 3-D residual for

structure/group points	*mean 3-D dist.* (mm)
basal ganglia (n=11):	1.91 ± 0.93
ventricular (n=5):	2.49 ± 1.16
longitudinal (n=14):	2.61 ± 1.14
cerebellar (n=18):	2.85 ± 1.55
cortical (n=18):	3.61 ± 1.36
total (n=48):	**3.09 ± 1.15**

Table 7.3: Mapping differences: automatic vs. manual.

Average r.m.s. difference in position for points mapped into stereotaxic space with the manual (least-squares) method and the automatic correlation technique (16 subjects). Groups of points are described in section 7.3.4

the unconstrained least-squares, i.e. no fixed origin, was 5.05 mm .) The corresponding value for T_{auto} was slightly larger at 6.13mm in 3-D (4.39, 2.08 and 3.73mm , in x, y and z). Not surprisingly, the smallest r.m.s. residuals were for points in the basal ganglia (2.7 mm vs 2.9 mm , for T_{man} and T_{auto} respectively) and the largest were for cortical structures (6.8 mm vs 7.8 mm). (See Table 7.4 for details.)

7.4.3 Automatic registration vs. manual AC-PC based techniques

The hypothesis tested in the second experiment was similar to the first: that the automatic method was comparable to manual techniques used over the last five years at the Montreal Neurological Institute to map over 400 MR data sets into stereotaxic space (Evans *et al.*, 1992c), as described in section 6.2.

Two volumetric goodness-of-fit measures have been calculated for each of 60 brains transformed into stereotaxic space. One measure was the correlation statistic (as described in Appendix A) and the other was an r.m.s. difference between normalized voxel intensities, where voxel values for each volume are replaced by a z-value:

$$z = \frac{x - \bar{x}}{s} \tag{7.3}$$

where $\bar{x}$ and s are the sample mean and the standard deviation for voxels within the brain-masked region. The z-score difference (or just z-score) is defined as the voxel-to-voxel difference in z-value between the two data sets. This measure, a difference of z-values, is used instead of a z of the intensity difference (z-of-difference) because it allows the comparison of volumes with different mean intensity values and contrasts. While it would be possible to normalize the z-of-difference to the mean intensity of one volume,

structure/group points	*automatic 3d dist* (mm)	*manual 3d dist* (mm)
basal ganglia (n=11):	2.87 ± 0.96	2.70 ± 0.77
ventricular (n=5):	4.32 ± 1.68	4.14 ± 1.75
longitudinal (n=14):	4.99 ± 1.45	4.09 ± 1.02
cerebellar (n=18):	5.99 ± 1.88	4.91 ± 1.59
cortical (n=18):	7.74 ± 1.74	6.78 ± 1.46
total (n=48):	$\mathbf{6.13 \pm 1.09}$	$\mathbf{5.21 \pm 0.90}$

Table 7.4: Misregistration: automatic vs. manual.

Average r.m.s. misregistration of homologous landmark points mapped into stereotaxic space. Comparison of the automatic correlation method vs. manual landmark based registration (16 subjects). Even though the difference in group means is significant ($p < 0.01$, $t = 5.2$), note that the manual result is optimized for the given set of landmark points.

the other volume, or both, this would bring the z-of-difference back to the difference-of-z used here.

In order to establish a comparison between methods, these two measures were calculated for each data set in stereotaxic space with respect to the brain-masked model, one each for the automatic and manually transformed volumes. When estimated in over 60 data sets, there was a higher average correlation value for the automatically transformed volumes, compared to their manually transformed counterparts (0.984 ± 0.0023 vs. 0.965 ± 0.019). While this measure was biased towards the automatic registration method which uses a cross-correlation measure to achieve the fit, and therefore should give a higher mean value, the standard deviation of the measured residual across the 60 automatic transformations was more than eight times smaller than that for the manual transformation, indicating a more stable procedure. The r.m.s. z-score difference was not as biased towards one method or the other. When estimated over the same 60 volumetric MRIs, this measure was slightly lower for the automatic technique compared to the manual method: 0.734 ± 0.064 vs 0.784 ± 0.093. Note again the smaller (i.e better) standard deviation of the z-score for automatically registered data, although now only a factor of 1.5, indicating that this approach is more stable than the manual technique while yielding comparable registrations.

The z-score measures intensity variation not only due to misregistration, but also to noise in the original data and to sampling errors incurred in the transformation to stereotaxic space. The effects of resampling and noise were reduced by imposing a blur (FWHM=8mm) on the MRIs, after mapping into stereotaxic space, and re-computing the comparisons for the corresponding blurred intensity and gradient magnitude volumes. For the cross-correlation measure, the automatic vs manual results remained essentially the same: 0.996 ± 0.0015 vs. 0.985 ± 0.014 for blurred intensity data; and 0.958 ± 0.0085

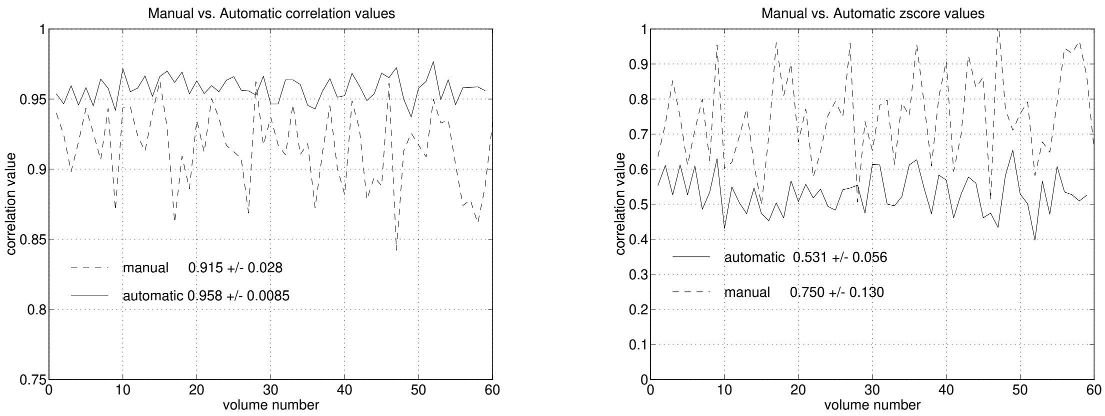

Figure 7.10: Comparison of manual and automatic stereotaxic registration
The graph on the right shows the correlation value for 60 stereotaxically transformed volumes and the MRI atlas. The correlation value was calculated on gradient magnitude (FWHM=8mm) data. The dashed line corresponds to the manual transformation, and the solid line corresponds to the automatic transformation. Not only does the automatic transformation yield higher correlation values, but the standard deviation is smaller as well, indicating a more stable technique. (Note the non-zero ordinate axis, used to accentuate the difference between the manual and automatic correlation value.) The graph on the right shows the z-score value calculated for the same data.

vs. 0.915 ± 0.028 for gradient data (see Fig. 7.10-a). The z-score difference was found to be significantly improved, since it was no longer measuring differences due to image noise. With the imposed blur, the z-scores are much lower for the automatic technique than manual method (automatic vs manual): 0.409 ± 0.076 vs. 0.535 ± 0.13 for intensity data; 0.531 ± 0.056 vs. 0.750 ± 0.13 for gradient data (see Fig. 7.10-b). These numbers are consistent with the automatic approach being more stable than the manual technique.

In order to provide a context for these measurements, z-scores were calculated for twenty stereotaxic data sets, that had been deliberately misregistered by varying one of the nine parameters at a time. In each case, the translation parameters were varied from -10 to 10 mm , the rotation parameters from -10 to 10 degrees and the scale from -20% to +20%. From the graphs, shown in Fig 7.11, one can estimate the amount of parameter error for a given z-score.

For example, in the case of the manual registrations, the average z-score for gradient magnitude data was 0.75. This corresponds to a translational error of 6 mm in the x and y-directions or 8 mm in the z-direction with respect to the stereotaxic space, rotational errors of 2 to 3 degrees or scaling errors of 6% too small or 8% too large in the x and

y-directions or 12% too small or 20% too large in the z-direction. For the automatic approach with an average z-score of 0.53, the corresponding errors would be 1-2mm for translation, 1-2 degrees for rotation and 2% for scale in x and y and 3-4

Note that this does not lead to the conclusion that the error in registration is on the order of 6mm in x or y, nor 8mm in z. The magnitude of these values is due to inter-subject differences.

7.4.4 Algorithm robustness

Finally, the robustness of the algorithm to missing data, i.e., when the data set to be registered covers only part of the model in stereotaxic space, was examined.

For each experiment, a number of slices were removed from the top, from the bottom or from both top and bottom of five original MRI volumes. The resulting data were then used as input to the automatic registration routine. As can be seen in the graph in Fig. 7.12, this r.m.s. misregistration stays relatively stable with up to 30% removed from the top, 60% removed from the bottom, or 50% removed from top and bottom.

It should be noted that the MRI volumes used in this experiment were scanned at different angles than that in the model and extended lower than the model to include the cerebellum and the nasal sinus cavities. Therefore, slices from the bottom of the data set that have no corresponding structure in the model may be dropped without penalty, and this accounts for the stability of the algorithm for the first twenty percent of bottom slices removed. However, the registration remains stable with up to 60% of slices removed from the bottom, corresponding to keeping slices superior to the line joining the occipital lobe and the lower third of the frontal lobe (line c in Fig. 7.13).

7.5 Measures of anatomical variability

7.5.1 Inter-subject variability estimated from landmark points

In stereotaxy, the analysis of anatomical variability requires both a well-defined coordinate system to establish a standard frame of reference <u>and</u> a transformation with a limited number of parameters to map volumetric datasets to this coordinate system. Therefore, to normalize statistical results across subjects and allow direct point-by-point comparisons between subjects, data sets were mapped into stereotaxic space using both the manual and automatic methods described above. The difference in position of a given landmark across a number of brains, after normalization to a predefined brain size and orientation, was used as a measure of neuro-anatomical variability. Unfortunately, point correspondence is difficult to establish not only because of real anatomical variability, but also because subjectivity in landmark identification, i.e., landmark interpretation and positioning error. In order to reduce the effect of subjective decisions on the point residual, five neuro-

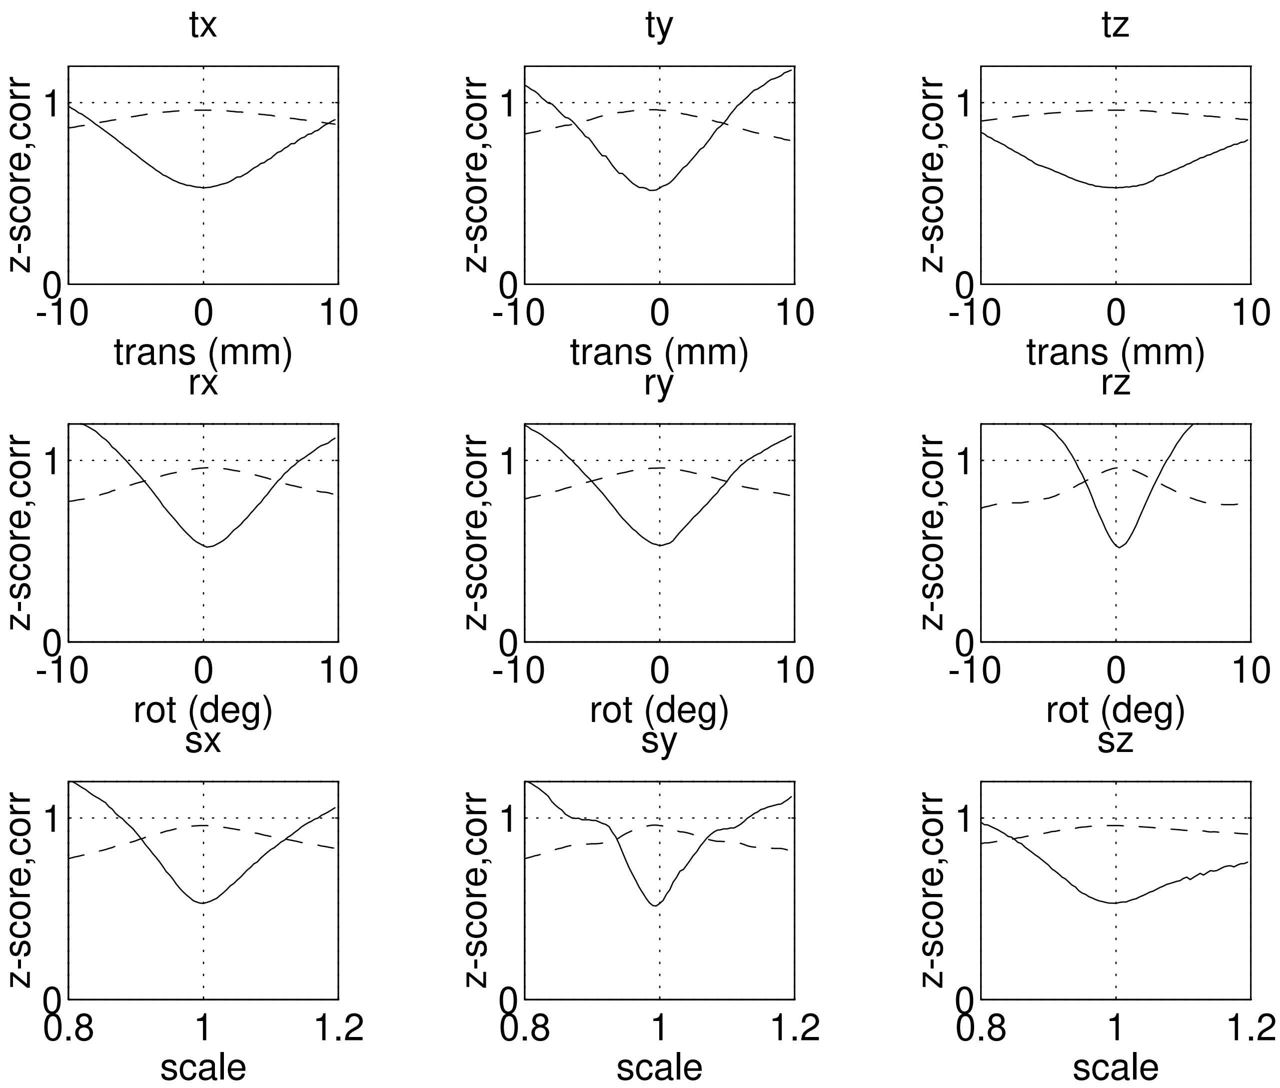

Figure 7.11: Correlation value and z-score vs parameter error.

Plot of correlation value and z-score vs parameter error for 8mm gradient magnitude data. Twenty data sets were deliberately mis-registered, varying one parameter at a time. The solid line corresponds to the z-score, the dashed line to the correlation measure. Note that rot-z is steeper than rot-x and rot-y, since the edges corresponding to the inter-hemispheric fissure strongly constrained rotation around the z-axis. The scale and position along the x and y-axis is constrained more than that along z since the limits of the cortex, in the LR (x) and AP (y) directions, were always apparent on the transverse images.

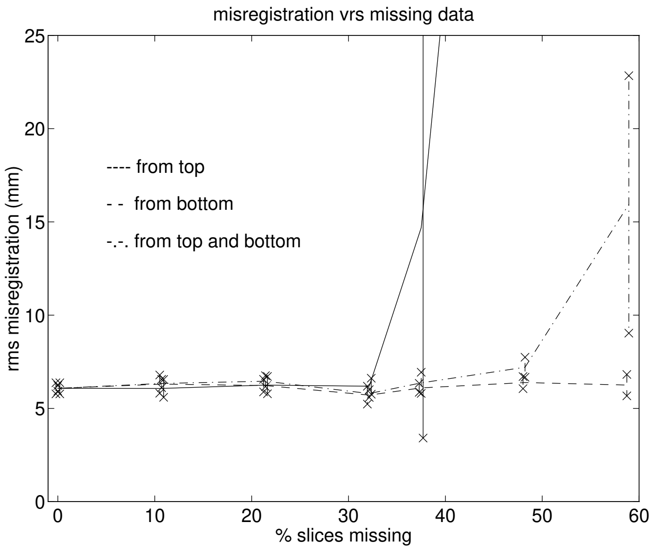

Figure 7.12: Misregistration error vs slices missing.

Solid line corresponds to slices removed from the top, dashed line to slices removed from the top and bottom, dash-dotted line to slices removed from the bottom.

anatomically trained physicians were asked to define a list of landmarks and identify them on the 17 MRI brains used in the previous experiments. The observers agreed on a search strategy for identification of each landmark. Then they followed that strategy for each landmark in each brain to generate an ensemble of $5 \times 17 \times 34$ landmarks which could be analyzed to separate the contribution of landmark mis-matching from sources due to inter-subject variability (ISV), inter-observer variability (IOV) and a residual variability (RV) term associated with intra-observer variability. This work has recently been submitted by Sorlié *et al.* (1994).

Landmark identification

The 34 anatomical landmarks selected were distributed over the cerebrum, the cerebellum and the brain stem, and can be separated into two broad categories:

- Points on cortical structures (n=15)
  (enthorinal sulcus, collateral sulcus, intersection ascendant-horizontal sulci, central sulcus at interhemispheric fissure, olfactory sulcus, parieto-occipital sulcus, marginal ramus, occipital lobe).

- Points on internal structures (n=19)
  (red nucleus, ventral beginning caudate, end splenium, curvature chiasma, tip an-

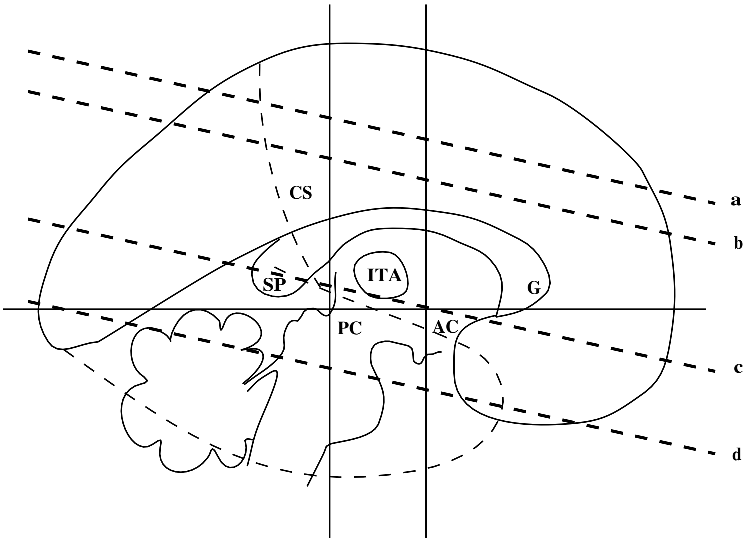

Figure 7.13: Robustness to missing data.

Mid-sagittal slice showing approximate levels of data removed for testing. Registration remains stable with slices removed above line b, removed below line c, or removed both above a and below d. PC, posterior commissure; AC, anterior commissure; CS , central sulcus; SP, splenium; G, genou; ITA, inter-thalamic adhesion.

terior genu, anterior and posterior commissures, inter-thalamic adhesion, rostral pons at interpedicular, fastigium, superior cerebellar peduncle, inferior and superior colliculus).

As before, this set of landmark points represents a compromise between the need for a large number of points, their ease of identification and the aim of not specifying redundant anatomical information by selecting points very close to each other simply because they were easy to identify.

Registration and analysis methodology

The automatic image registration algorithm was used to recover the stereotaxic transformation T_{auto} required to map each volume (and its set of landmark points) into stereotaxic space. Anatomical variability was estimated by measuring the distribution of landmark points after affine transformation to stereotaxic space. For completeness, the procedure was repeated for different degrees of freedom. The probability distribution of a landmark location as a function of position along the x, y and z-axes was modelled as a Gaussian, centered on the mean landmark position pooled over all data sets and all observers. The FWHM of the Gaussian was the measure chosen to express the variability.

The width of the distribution measured for a given landmark in stereotaxic space, averaged over all observers and all data sets, is a combination of variabilities including intra-observer error, inter-observer error and inter-subject variability. A statistical analysis of variance was used to estimate and remove the observer-dependent variabilities and then to infer the desired measurement of the true inter-subject variability. In the following analysis, the notation ξ_{ij} refers to the position along the ξ-axis ($\xi = x$,y or z) for a single landmark tagged on the i-th subject ($i = 1, n$; $n = 17$) by the j-th observer ($j = 1, m$; $m = 5$). A dot is used to represent the mean position when averaged over observers $\xi_{i.} = \frac{\sum_j^m (\xi_{ij})}{m}$, over subjects $\xi_{.j} = \frac{\sum_i^n (\xi_{ij})}{n}$ or both $\xi_{..} = \frac{\sum_i^n \sum_j^m (\xi_{ij})}{n \cdot m}$. These points are shown schematically in Fig. 7.14.

Thus, the coordinate $\xi_{..}$ represents the grand average estimate of the true location of any given landmark since it is averaged over all subjects and all observers. The coordinate $\xi_{i.}$ is the best guess of that landmark position on a given subject, since it is the centroid of the points identified by the five anatomists. The difference between these two points is defined by $\alpha_i = \xi_{i.} - \xi_{..}$ and is indicative of the anatomical deviation from the average of this landmark in subject i. Another difference of interest is defined by $\beta_j = \xi_{.j} - \xi_{..}$; it is an estimate of the bias of observer j, compared to the best estimate $\xi_{..}$.

These definitions allow the derivation of expressions that characterize the various components of the variance across estimates of each landmark coordinate. The *total*

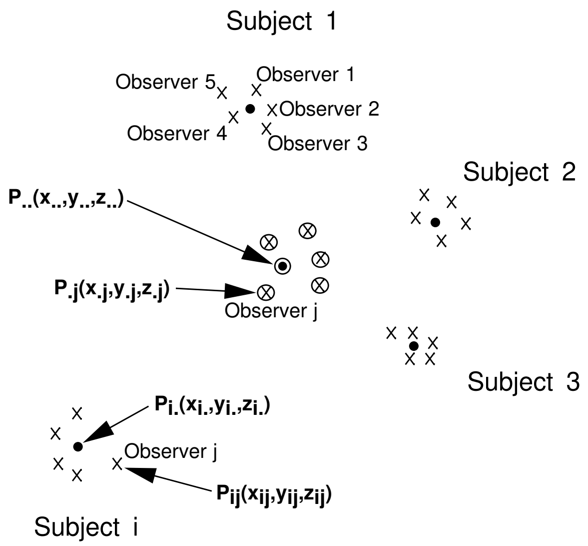

Figure 7.14: Point definitions for statistical analysis.

Four types of points are used in the statistical analysis:

P_{ij}: landmark point tagged by observer j on subject i.
$P_{i.}$: centroid over observers for subject i $(j = 1, 5)$.
$P_{.j}$: centroid over subjects for observer j $(i = 1, 17)$.
$P_{..}$: centroid over all observers and all subjects. $(i = 1, 17;\ j = 1, 5)$

variability is simply the standard deviation measured over all ξ_{ij}:

$$TV_\xi = \sqrt{\frac{\sum_i^n \sum_j^m (\xi_{ij} - \xi_{..})^2}{(n-1)(m-1)}}. \tag{7.4}$$

We can derive equations that separate the inter-subject variability from the observer-dependent variabilities. We can model the difference between any landmark and the "true" position, $\xi_{ij} - \xi_{..}$, as a sum of terms α_i, β_j and ϵ_{ij} such that,

$$TV_\xi = \sqrt{\frac{\sum_i^n \sum_j^m (\alpha_i + \beta_j + \epsilon_{ij})^2}{(n-1)(m-1)}}. \tag{7.5}$$

where α_i and β_i were defined above, and $\epsilon_{ij} = \xi_{ij} - \xi_{i.} - \xi_{.j} + \xi_{..}$.

The *residual variability* (RV) is defined as a measure of the consistency of an observer when tagging the same point in the same subject. Thus this term RV equals 0 if the observer is consistent and will differ from 0 according to his inconsistency. The residual variability RV_ξ is written for each of the coordinate system axes ξ as:

$$RV_\xi = \sigma_\xi^\epsilon = \sqrt{\frac{\sum_i^n \sum_j^m (\xi_{ij} - \xi_{i.} - \xi_{.j} + \xi_{..})^2}{(n-1)(m-1)}}. \tag{7.6}$$

The *inter-subject variability* (ISV) is the measure of neuro-anatomical variability that is to be separated from the total variability. It is defined for a given landmark, as a measure of the difference between its "mean" location, $\epsilon_{..}$, and its 17 "individual" positions in the corresponding 17 subjects volumes, $\epsilon_{i.}$ minus a term that is a function of the residual variability. The inter-subject variability measured along each coordinates system axis ξ is then expressed as:

$$ISV_\xi = \sigma_\xi^\alpha = \sqrt{\frac{\sum_i^n \sum_j^m (\xi_{i.} - \xi_{..})^2}{n-1} - \frac{RV_\xi^2}{m}}. \tag{7.7}$$

The *inter-observer variability* (IOV) is a measure of the error made by the observers when tagging the same landmark. It is, for a given landmark, the distance between the centroid for the points tagged by one observer, $\epsilon_{.j}$ in all the subjects and its "mean" location, $\epsilon_{..}$, minus a term including the residual variability. It thus gives an estimate of the inter-observer variability. For each of the coordinates system axes ξ, the inter-observer variability is given by:

$$IOV_\xi = \sigma_\xi^\beta = \sqrt{\frac{\sum_i^n \sum_j^m (\xi_{.j} - \xi_{..})^2}{m-1} - \frac{RV_\xi^2}{n}}. \tag{7.8}$$

These variability measures in x-, y- and z-directions for all the landmark points can be considered as an estimate of the standard deviation on the points, which are assumed to come from the same distribution. The 3-D average FWHM value of each variability for every landmark point was then calculated using the following equation, but replacing σ_{xi} with the appropriate σ^{α}, σ^{β} or σ^{ϵ}:

$$\text{FWHM} = 2.35\sqrt{(\sigma_x^2 + \sigma_y^2 + \sigma_z^2)/3}. \tag{7.9}$$

Results of variability estimation

These equations provide the means to estimate neuro-anatomical variability, based on the landmarks selected above. Six automatic transformations were evaluated that differed in the number of degree of freedom allowed in the optimization where 3, 6, 7, 9, 10 or 12 parameters were used. The 3-parameter transformation includes only translations along the 3 cardinal axes. Three rotations around the cardinal axes are added to form the 6-parameter transform. The use of a single scaling factor is added in the 7-parameter transformation, or the use of three scaling along each axis for the 9-parameter fit. A single shear defined in the longitudinal plane is added for the 10 parameter fit, while the addition of 3 different shears (one along each axis) leads to the 12-parameter transformation. Figure 7.15 graphs the IOV, ISV and RV for the six transformations. As expected, ISV decreases with added degrees of freedom, however there does not seem to be a significant improvement between the 9-, 10- and 12-parameter transformations. The estimate of IOV remains stable as it should, since it does not depend on the type of stereotaxic transformation employed. The same holds true for RV, however it is larger than either IOV or ISV. The flat behavior of IOV and RV curves, coupled with the expected reduction in ISV with increasing parameters, supports the assumption of additive, orthogonal errors and the use of the model described by eq. 7.5. The average ISV over all points is 3.9mm for the 9 parameter transformation, when measured as a 3D FWHM. As expected, the ISV is greater for cortical points (mm) than for internal points (mm).

7.5.2 Anatomical blurring kernel

Recently, Evans *et al.* (1993a) presented an alternative method to characterize anatomical variability. They modelled the variability as a spatially invariant blurring kernel that when applied to a single individual makes it appear similar to the mean MRI atlas. The optimal kernel width was found by comparing the 3-D Fourier transform of the true average MRI volume with that of a single brain. The width estimated by this technique corresponds to a 3-D FWHM of 4.08mm. The average ISV of 3.9mm corresponds well with this estimate.

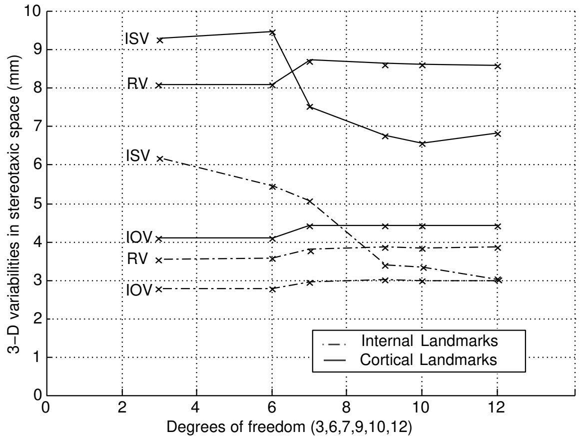

Figure 7.15: Variability vs degree of freedom.

Average inter-subject (ISV), inter-observer (IOV) and residual (RV) variabilities, defined as 3D FWHM, plotted vs varying number of degrees of freedom (3, 6, 7, 9, 10, 12) for internal and cortical landmark points. Not surprisingly, cortical points are subject to larger variabilities than internal points.

7.6 Discussion

7.6.1 Comparison to manual registration methods

When compared to manual landmark based methods, there was no bias in the registration parameters recovered by the automatic technique. As stated previously, the 3-D r.m.s. difference in position for landmark points mapped into stereotaxic space automatically when compared to manual mapping was 3.1 mm .

Both techniques yielded similar r.m.s. residuals between specific target landmark points and their homologues, in standard space: 5.19 mm for manual and 6.13 mm for the automatic technique. While the misregistration error appears significantly greater for the automatic procedure ($p < 0.01$, $t = 5.2$ for comparison of group means), its measurement was biased in favor of the manual technique, since the points used to define the transformation were the same as those used to measure their error.

Based on the smaller standard deviations for both correlation (0.0023 vs. 0.019 for raw data) and r.m.s. normalized difference measures (0.056 vs. 0.13 for gradient data) when using the automatic procedure, we conclude that it is more stable than the manual registration method. The lower z-scores (0.41 vs. 0.54 for gradient data) for the automatic technique also indicate that it was superior to the manual one.

7.6.2 Algorithm robustness

The algorithm has been shown to be robust to missing data and to noise. Registration experiments that included noise show that the algorithm can recover simulated transformations with up to 40% noise before the registration residual begins to increase. This is explained by the blurring due to convolution with a Gaussian kernel before gradient extraction.

Because the final fit is *edge*-driven by the gradient magnitude, it is important that the input data contain robust, reliable edges to be registered with the model. At the FWHM=8 mm scale, one can see (on Fig. 4.2) strong edges at the cortex, insula, longitudinal fissure and ventricular surfaces; these are the structures that dominate the registration process at this scale.

In order to constrain the scale and position of the $x-$axis, the data should cover the longitudinal fissure and the lateral aspects of the cortical rim. The anterior and posterior aspects of the cortex must be visible to fix the scale and position of the $y-$axis.

The scale in the $z-$direction was the most difficult to obtain. While the ventricular boundaries play a large role in fixing the $z-$position, they do not extend for a large distance in that direction, and small errors in registration magnify scaling errors at the edges of the brain unless either the top or the bottom limit of the cortex are also present.

As indicated in Fig. 7.13, a large amount of data could be missing from the bottom (below line c), since the ventricles along with superior aspect of the cortex will fix the CC-position and scale. The same was true when slices are missing from the top (above line b), as the ventricles and the floor of the cerebral volume serve to fix the parameters. ‘

7.6.3 Multi-resolution methodology

There are two benefits achieved using the multi-step multi-resolution strategy. 1) The multi-step process lessens the likelihood of settling into a local minimum in the optimization process. An approximate solution is found at the first step and it is refined in each sequential step. 2) The multi-step multi-resolution approach offers a substantial computational savings. There was an eight-fold increase in time required to calculate the correlation coefficient when passing from one scale level to the next. The linear transformation can be recovered using the multi-resolution method in less than one quarter the time taken when starting at the highest resolution level (where the solution found was not always correct because of local minima). At its current un-optimized state, the routine requires approximately 30 minutes on an SGI Indigo[2] Extreme, a 60.5 SpecFP Unix machine, to calculate the features and register a specific volume into stereotaxic space.

7.7 Summary

A fully automatic registration method has been presented that yields results comparable to manually-based techniques, while eliminating the drawbacks involved in manual intervention such as reproducibility and inter-observer variability. Multi-resolution cross correlation has been shown to be an effective tool to determine the affine component of the transformation between source and model for both simulated and real data. While this technique has benefits and accuracies similar to that of landmark-based registration, or surface to surface matching techniques, it has the advantage that explicit manual intervention is not required to identify such points or contours. The algorithm can be applied retrospectively since it relies on automatically detected edges in the data that correspond to internal anatomical structures rather on than fiducial markers.

This registration method has been used to automate the stereotaxic mapping procedure. As well as characterizing the stereotaxic space, the use of the standardized model obviated the need to manually identify, edit and remove scalp, skull or meninges from the MRI volumes.

Chapter 8

Non-linear registration: experiments and results

8.1 Introduction

This chapter presents the validation experiments for the non-linear registration procedure and the evaluation of its performance with simulated data with respect to noise and with different real data sets. The experiments presented here are divided into three parts. The first uses simulated data and transformations to demonstrates the functionality of the non-linear registration with a simple bar phantom, a simple ellipsoid phantom and the 3-D digital brain phantom introduced in the previous chapter. Real MRI data are used in the second group of experiments to show that the algorithm is capable of reducing the inter-subject neuro-anatomical variability in stereotaxic space by 35%, when the measure is based on landmark points. The non-linear registration technique is compared to a manual landmark-based method and shown to be more robust. In the last experiment, an automatic method for estimating anatomical variability is shown and correlated with manual results of the previous chapter.

8.2 Simulated data

8.2.1 Bar phantom

The recovery of a deformation field is illustrated by applying a non-linear warp to a 3-D rectilinear parallelepiped to create a "bent" parallelepiped. The warp was applied with a 3-D thin-plate spline defined by 9 points; 1 at each corner of the bar, and 1 in the middle. The registration procedure was then used to recover the applied transformation (see Fig. 8.1). The hierarchical optimization routine recovered the 3-D deformation field (shown as a quiver plot in Fig. 8.1, where the projection onto the 2-D slice of the 3-D

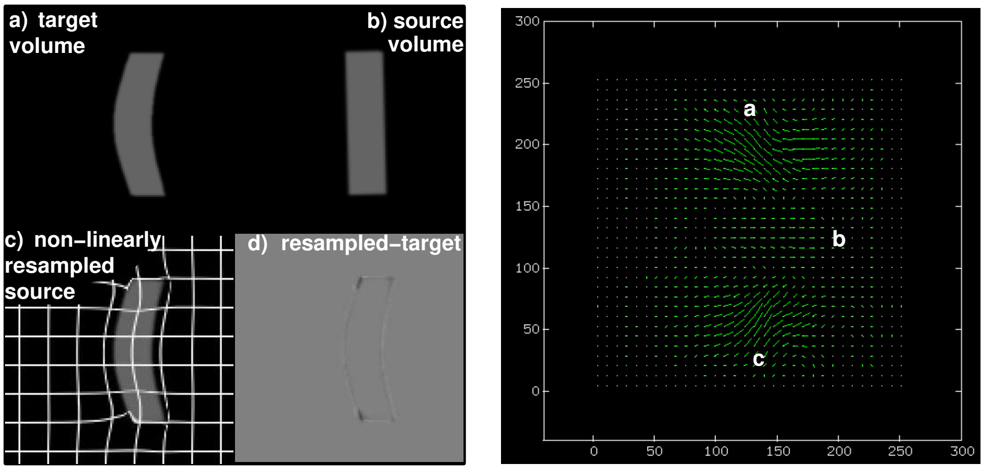

Figure 8.1: Non-linear simulation: bar phantom
On the left, the non-linear registration routine was applied to the straight bar in b) using the bent target in a). c) shows the application of the recovered non-linear warp (indicated by the grid) to the source volume. In d) the subtraction of the resampled volume from the target indicates very good registration. On the right is a vector representation of the deformation field. Each local 3-D deformation vector from the 2-D slice shown is projected onto that plane and represented by a small vector, originating from the center of each voxel in the deformation field. The regions **a** and **c** show the diagonal vectors pushing on the ends of the bar. The horizontal vectors of region **b** are pushing the middle of the bar towards the left.

local deformation vector is represented by small arrows, originating from the center of each voxel in the deformation field) required to map the source onto the target volume. The difference image shows that a good registration has been achieved.

8.2.2 Ellipsoid phantom

In order to examine the non-linear registration with a more complex example, the thin-plate spline warping algorithm developed for point to point non-linear registration (see page 10) was used to deform the ellipsoid phantom originally described in section 7.2.1. The automatic non-linear matching algorithm was then applied to fit the original volume onto its warped counter-part. The top row of Fig. 8.2 shows a slice through the center the source volume, the pear-shaped deformed volume as well as the volume resampled through the recovered transformation. At each hierarchical step in the non-linear registration procedure, a deformation field is stored at that scale. The second row of Fig. 8.2 shows an intensity difference image, formed by the subtraction of the resampled source from the target volume, at three of the four scales used in the recovery process. The last difference image indicates that the registration is not completely perfect. This is due to

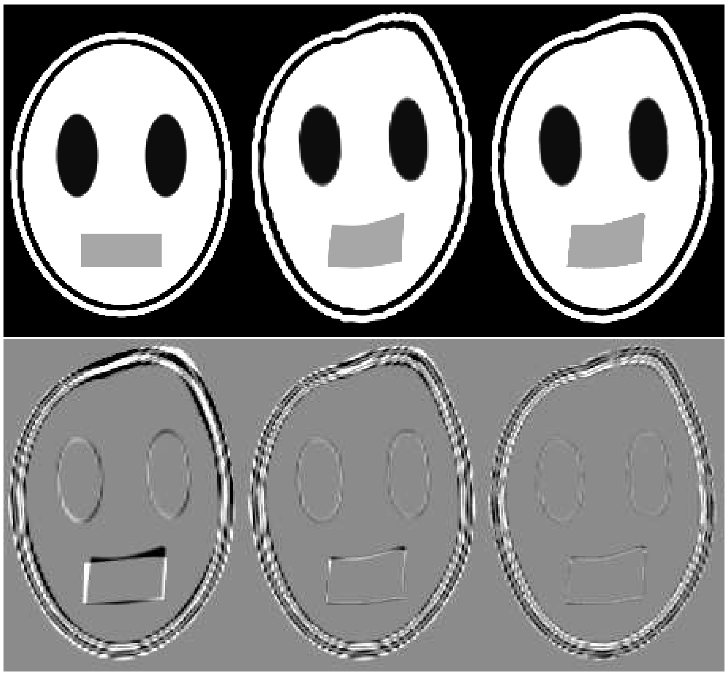

Figure 8.2: Non-linear registration of ellipsoid phantom.

These images show the result of the non-linear registration procedure applied to the ellipsoid phantom. In the top row, the first image is the source phantom. The second image shows the warped target and the third shows the result of application of the recovered deformation to the source volume. The bottom row shows difference images between the resampled source and the target at the first, second and forth scale steps (16mm, 8mm and 2mm). One can see that the error is reduced at each scale step. The difference images have their range set to -0.5 to 0.5, where 1.0 was the maximum intensity of the original data sets.

at least two reasons. The first is that linear interpolation on the original volume is used to resample it through the recovered transformation and form the warped result. Small interpolation errors are accentuated by the high contrast at the edge of each structure. Also, at each iteration of the non-linear registration procedure, only a fraction of the required local deformation is added to the global warp. Thus, an edge will tend to align onto its corresponding edge, without ever reaching it.

The difference images in Fig. 8.2 show that the non-linear registration procedure can recover the deformation needed to match the intentionally warped volume. On average, voxels were subject to a 4.14mm warp in the TPS-applied deformation. The r.m.s. distance of landmark points sent through the forward and inverse transformations from their original locations was 1.78mm. The relatively large error remaining is due to the fact that the registration method is based on gradient matching. Areas where there are no gradients to indicate the true local deformation are simply interpolated from more distant areas. Even though the objects may be well registered (as indicated by the difference images), their surfaces may slide on each other, therefore possibly increasing the point-to-point residual. When more complex structure is available on the object surface, the registration on the surface of the object is further constrained and the residual error is reduced. This effect is apparent in the next experiment, using a digital brain phantom.

8.2.3 Brain phantom

Realistic evaluation of the non-linear registration procedure begins with the use of the brain phantom described in section 7.2.2. The 3-D phantom was used to create three realistic volumes, representing T1, T2 and PD-weighted scans containing 10% Rayleigh-distributed noise. The first column of Fig. 8.3 shows transverse slices through the three original volumes. A set of landmarks was defined for input to the TPS algorithm to warp each volume into topologically equivalent, but spatially non-equivalent, target data sets. The non-linear registration procedure was then applied to each original/warped volume pair to recover the deformation needed to transform the original so that it would coincide perfectly with the target. There is no difference between recovering the transformation from source to target, or from target to source, since the non-linear deformation field is constrained to be invertible. The second column of Fig. 8.3 shows a transverse slice through the data set that was mapped through the recovered transformation. The corresponding difference images are presented in the third column.

Since the deformation is recovered in a hierarchical fashion, the original volume was mapped through the transformation recovered at each scale step. Figure 8.4 shows the intensity difference images for the linear and non-linear registrations. As the scale of the deformation is reduced, the registration improves. By the 4mm scale, there appears to be less than half a voxel mis-registration in most regions of the volume.

In these simulations, three measures are possible for quantitative evaluation of the non-linear registration. These include the registration error on homologous points and image-based z-score difference used in chapter 7, as well as the difference between the applied and recovered non-linear deformation. The 34 points described in section 7.5.1 were identified on the unwarped brain phantom. The coordinate of each landmark was mapped through the applied TPS and recovered deformation field. As before, the registration error was the 3-D r.m.s. difference in position between these two sets of points, and was equal to 3.8±2.1mm following linear registration but before non-linear registration. For the T1-weighted data, the registration error was reduced by 87% to 0.50±0.42mm. The error was slightly higher for the T2 data: 0.54±0.48mm, and even higher for the PD-weighted volumes: 0.94±1.30. These differences are discussed below.

As in chapter 7 the z-score difference of the gradient magnitude volumes was also used as a goodness-of-fit measure for the non-linear registration. The recovered transformation was applied to the FWHM=4mm gradient magnitude data so that it could be compared

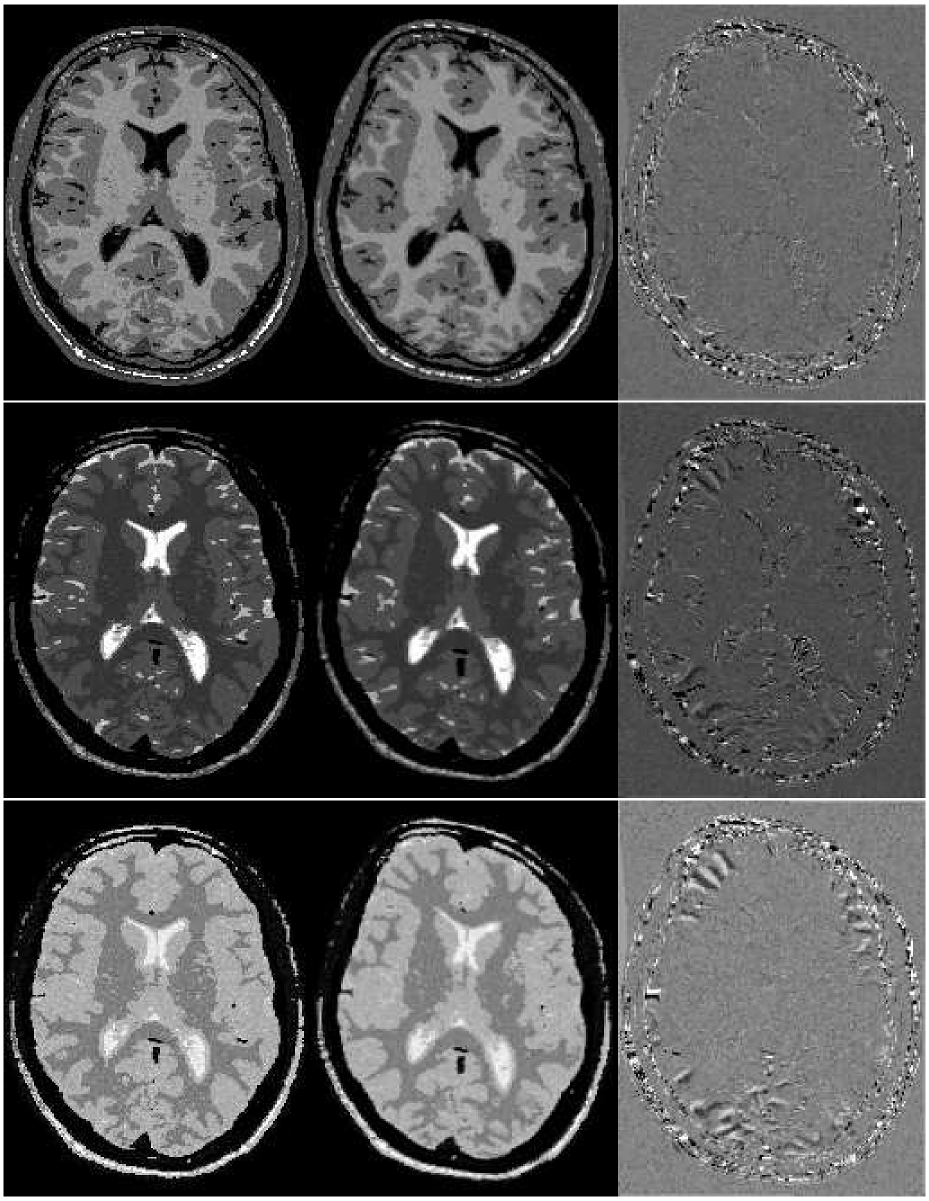

Figure 8.3: Non-linear registration for T1, T2 and PD-weighted volumes.

These images show the results of non-linear registration for T1 (first row), T2 (second row) and PD-weighted (third row) volumes, each with 10% added noise. The first column shows a single transverse slice through the original 3-D data. Each volume was warped with the same TPS interpolant to form spatially deformed targets. The non-linear registration procedure applied, and the recovered deformation was used to resample the original volume so that it resembles the target (second column). The voxel-by-voxel intensity difference is shown in column three to judge the quality of the registration. Since a brain-mask is used at the last step of the deformation recovery process, the fit inside the brain is better than the scalp region.

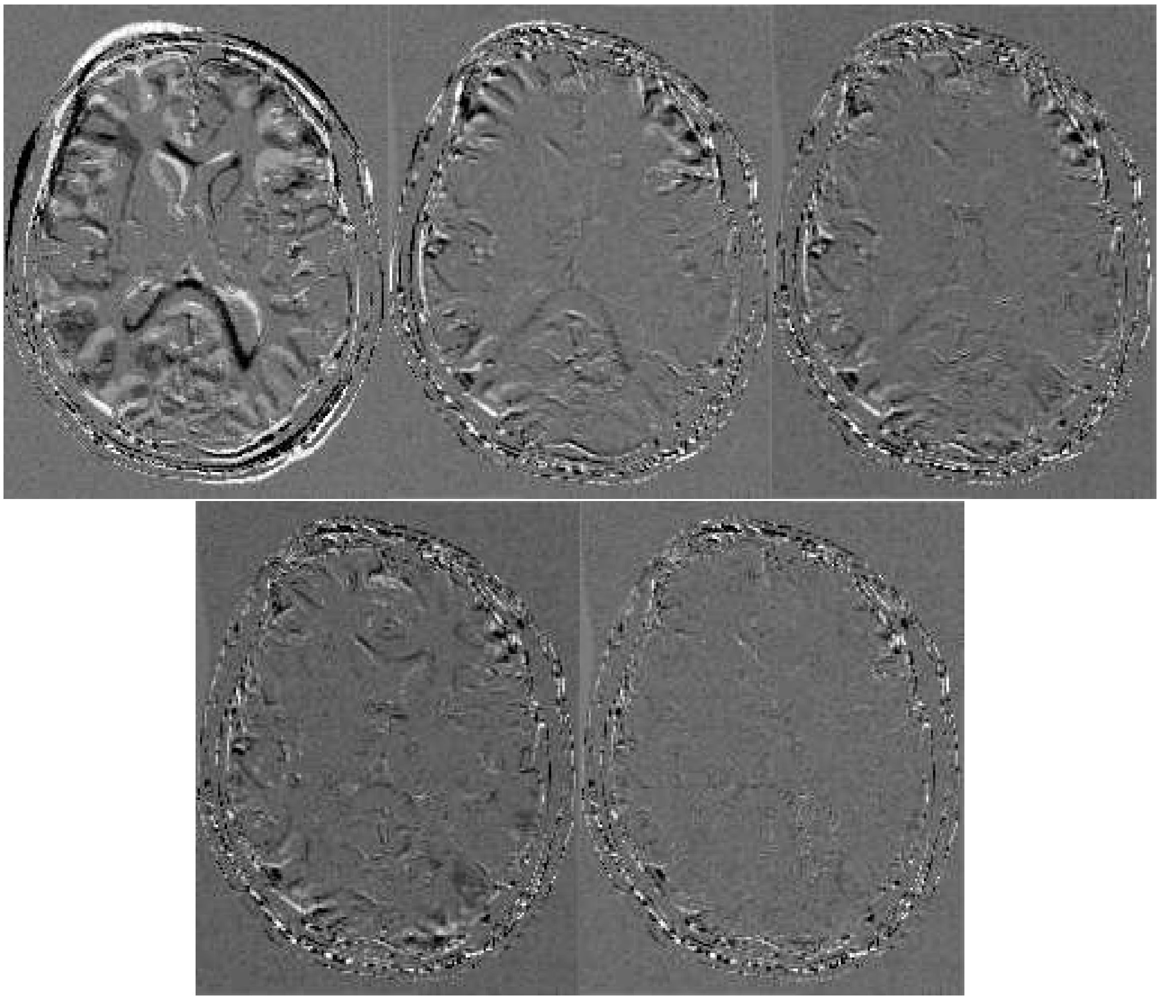

Figure 8.4: Deformation recovery at different scales.

Intensity difference images for the non-linear registration of the T1 original/warped volume pair. From top left to bottom right: Linear, 24mm, 16mm, 8mm and 4mm non-linear transformations used to resample the original volume into the target space. Each successive step refines the fit of the previous, evident in the reduction of the voxel intensity difference.

	w/mask		no mask	
data	*lin*	*non-lin*	*lin*	*non-lin*
T1:	0.558	0.159	0.481	0.168
T2:	0.486	0.075	0.439	0.113
PD:	0.424	0.070	0.430	0.134
average:	**0.489**	**0.101**	**0.450**	**0.135**

Table 8.1: Z-scores for simulated deformations on T1, T2 and PD data. This table shows the z-score values calculated for the T1, T2 and PD-weighted data, mapped through the recovered linear (*lin*) and non-linear (*non-lin*) transformations. Since the fit over the brain region is of interest, the z-scores were calculated with and without a brain mask. On average, the non-linear transformation improves the registration over the brain region by almost 80%.

to the equivalent target data. A mask was used to limit the calculation to voxels within the brain. On average, the non-linear registration algorithm improved the z-score by 79% for voxels with the brain over the linear transformation.

The registration error describes misregistration only at the landmark points selected and does not estimate error throughout the volume. On the other hand, the z-score was evaluated over all voxels, but it does not specifically measure the true deformation, i.e., regions may be misregistered, but if they have the same z-score they will not be flagged as such. However, since the transformation applied was simulated, the deformation vector for each voxel of the source volume is known and can be directly compared to the deformation vector recovered by the automatic procedure. This measure is termed the *recovery error* and is equivalent to the registration error estimated at every point in the volume. This value is a true estimate of the ability of the non-linear registration algorithm to recover an applied deformation. On average, for all points in the volume, there was an r.m.s. recovery error of 4.17$\pm$2.38mm after linear registration. For the T1-weighted data, this value was reduced to 0.63$\pm$0.60mm. The recovery based on T2 or PD-weighted data was higher: 0.80$\pm$0.70mm and 1.15$\pm$1.11mm respectively. These values correspond well to the errors on landmark position estimated above.

Interestingly, the T1-weighted volume appears to have the best recovery. This is probably due to the contrast difference between sulcal CSF and grey matter that is not as evident in the PD or T2-weighted volumes. This extra contrast yields more edges where the local deformation can be directly estimated instead of interpolated from neighbouring regions. While these extra edges improve the difference images, the recovery error, and the registration residual on landmark points, the z-score difference is greater for the T1-weighted data than for the T2 or PD results quoted in the previous paragraph. Since there are more edges in the T1 data compared to the T2 or PD volumes, there is more chance for misregistration, therefore increasing the z-score difference.

Non-linear registration error vs noise

Since the recovery of the applied deformation was quite good for 10% noise, one may ask what is the behavior of the non-linear registration algorithm in the presence of greater levels of noise. Using the T1-weighted images, noise was added to the unwarped data at levels from 10% to 80% in steps of 10%. The linear component of the transformation was fixed to that recovered in the noise-free linear registration. Fig. 7.5 in the previous chapter showed example slices of the noisy source data. Voxel intensity difference images in Fig. 8.5 show the quality of the registration for different levels of noise. In order to separate intensity differences due to noise from those due to registration error, the noise-free source volume was mapped through the recovered transformation to be compared to the target for the error measures described here.

The graphs in Fig. 8.6 show that the algorithm is very stable throughout the range of noise levels tested by the algorithm. The reasons for this stability is due to the inherent blurring of the Gaussian kernel used to define the gradient magnitude feature. As in the previous section, the z-score measure was calculated for the gradient magnitude volumes of both the resampled source and the target data at the FWHM=4mm scale, with and without a brain mask. For the linear fit between the warped target and the linearly transformed, but un-warped source volumes, the z-score was 0.558 and 0.481 with and without a brain mask respectively. This value is reduced by approximately 70% to 0.159 and 0.168 for the 10% noise level, after application of the recovered non-linear transformation. Even at the 80% noise level, the z-score was reduced significantly to 0.196 and 0.217, with and without the mask. The z-score measures how well edges in both volumes overlap after application of the recovered non-linear transformation. However, this measure does not estimate the point-to-point error, since one surface (as defined by high gradient magnitude values) could conceivably slide along its equivalent in the other volume without penalty. The registration error addresses this problem, estimating the misregistration at previously selected landmark points. As above, the 34 landmarks were mapped through both the applied TPS transformation and the recovered deformation field. The registration error measured as the r.m.s. difference in 3-D position between these two mappings ranged from 0.64±0.42mm for 10% noise to 0.82±0.48mm for 80% noise. For comparison, when mapped through the linear component only of the registration, the r.m.s. residual was 4.37mm. Therefore, at the levels of noise expected in typical MRI volumes (i.e., less than 10%), these simulations show that the non-linear deformation algorithm recovers more than 85% of the residual not accounted for by the linear model.

The recovery error defined above was calculated for each noise level. As shown in Fig. 8.6-c, the recovery error is similar to the residual error; ranging from 0.62±0.59mm for 10% noise to 0.85±0.66mm for 80% noise for the entire volume. Note the agreement with the graph of residual error (Fig. 8.6-b), measured on 34 landmark points.

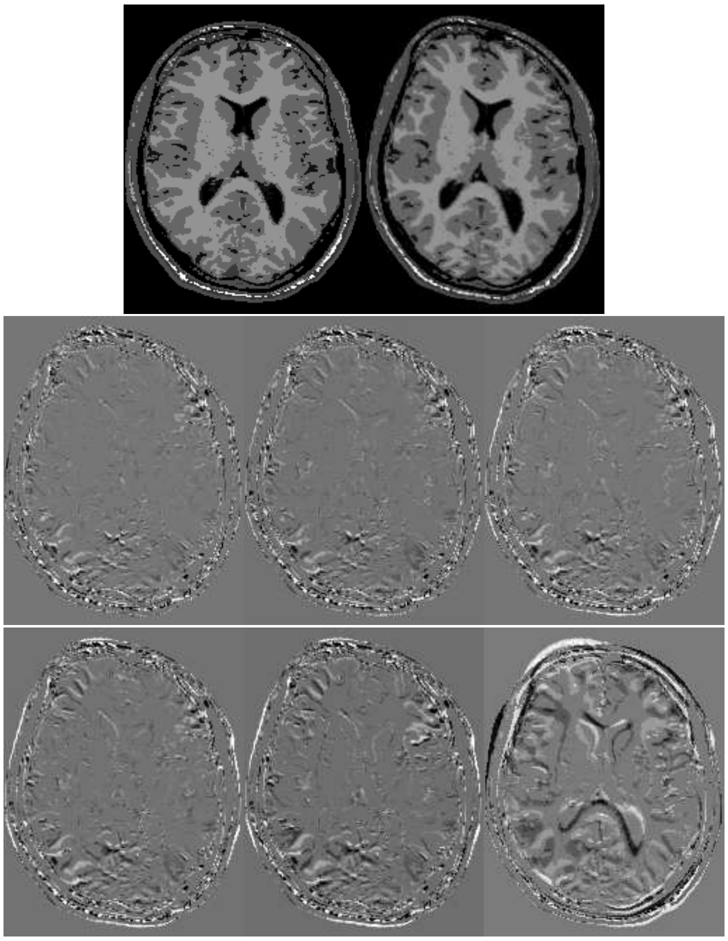

Figure 8.5: Non-linear registration in presence of noise.

The first row shows one transverse slice through the center of the noise-free source volume, and the corresponding slice in the warped target volume. Noise was added at levels from 0% to 80% in steps of 10% to the source volume and the non-linear registration procedure applied. In order to estimate the quality of the registration, the noise-free source was resampled through the recovered transformation and intensity difference volumes were formed with the target, so that the voxel intensity in the difference image would result from registration error only (and not from differences in noise). The second and third rows show the corresponding slice through the intensity difference image of the noise-free source, resampled by the transformations estimated on noisy data at levels of 0%, 20% and 40% (2nd row, left to right), 60% and 80% (3rd row). For comparison, the bottom right image shows the intensity difference after linear registration only.

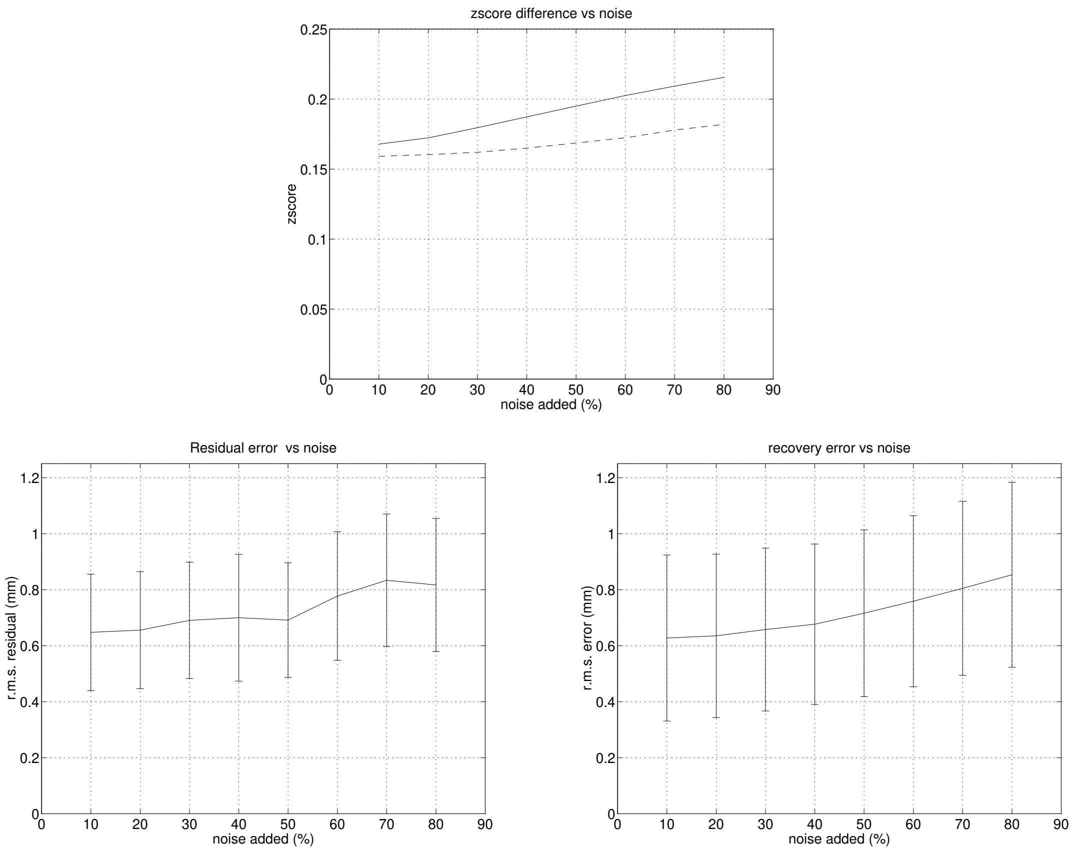

Figure 8.6: Quantitative non-linear error measures vs noise.

These graphs show three different quantitative measures of non-linear registration error. (**a**,**b**,**c** top, bottom left and right). In **a**, the solid line shows the z-score normalized gradient intensity difference for the whole volume, and for the brain only (dashed line). The graph in **b** shows the r.m.s. residual error for 34 landmarks defined on the phantom. The recovery error is shown in **c**, measuring the r.m.s. difference in deformation between the applied transformation and the recovered transformation for all brain voxels. The tick marks in **b** and **c** indicate one standard deviation. Notice the agreement between the residual error and the recovery error, indicating that the landmarks used to establish the residual error are representative of the whole brain.

8.3 Real MRI data

Real data were used to complete the validation of the non-linear registration procedure. For this task, the seventeen volumetric MRI studies described in the previous chapter were again used. No preprocessing to correct for image intensity variations or geometric artefacts was applied to the data before application of the non-linear registration procedure. As stated previously, the use of gradient fields greatly reduces the impact of RF inhomogeneity and spatial distortions are minimal. Three registration experiments were completed. The first used a single subject as a target brain, and the other 16 brains were registered to it. In the second and third experiments, the non-linear registration algorithm was incorporated into the stereotaxic mapping procedure in order to reduce the residual anatomical variability in stereotaxic space not accounted for by the linear transformation model.

8.3.1 Inter-subject non-linear registration

In order to demonstrate the operation of the general non-linear registration procedure between subjects, a single data set was selected as a target, and the other 16 brains registered to it using the methodology outlined in chapter 5. Figure 8.7 shows one transverse slice through four of the seventeen brains after linear registration with the left-most volume serving as the target. Since each volume has a different mean signal value, the volumes were first normalized for mean intensity and variance, thus minimizing error in the difference image from this source. Even so, the difference images show structure indicating remaining intensity variation not accounted for by the simple normalization procedure or by possible errors in registration. Most importantly, actual morphometric differences between subjects may be the cause of the high frequency structure. The result of application of the non-linear registration is shown in Fig. 8.8 for the same three brains. The grey-level intensity of the structure seen in the difference images of Fig. 8.8 is less than that of Fig. 8.7 indicating that a better registration has been achieved, especially near the cortical surface and in the neighbourhood of the ventricles.

Two quantitative measures were used to objectively evaluate the improvement of the non-linear over the linear registration. The first was the z-score and the second was based on the registration residual of landmark points. The recovery error cannot be calculated directly, since the non-linear transformation required to match source to target for real data is not known as for the previous simulation experiments. The z-score measure was calculated for the gradient magnitude volumes of both the individual target volume and each resampled source volume at the FWHM=4mm scale using a mask to calculate the statistic over the brain region only. When evaluated over the sixteen resampled source/target pairs the mean z-score was 0.656 ± 0.035 for the linear transformation and 0.414 ± 0.020 for the non-linear transformation. The 37% difference in z-score indicates a substantial improvement in the registration.

Landmarks (defined in section 7.5.1) identified in each individual brain were mapped

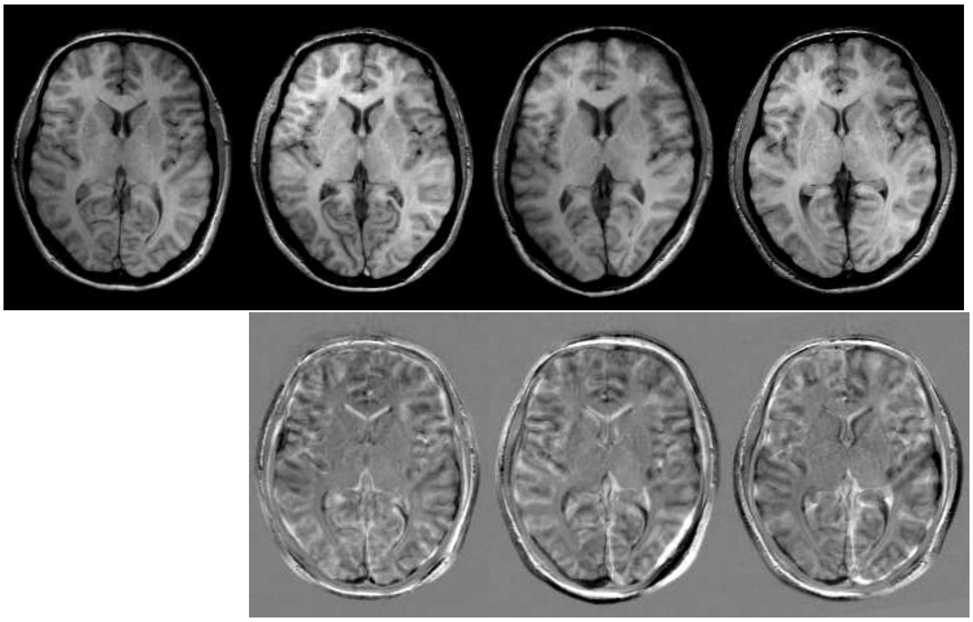

Figure 8.7: Linear registration result.
In the top row, the leftmost volume was used as a target and the other three linearly registered to it. In the previous chapter, it was shown than there was no significant difference ($p > 0.1$) between the transformation parameters recovered by the automated algorithm and those recovered by hand using homologous landmark for linear matching. The images on the bottom are difference images, where each volume has been subtracted from the target volume on a voxel by voxel basis after normalization for slice intensity and variance. (The range of the difference image is from -2.0 to 2.0 standard deviations in intensity.) The structure seen in the difference images is primarily due to non-linear morphological variability between brains. It is this variability that must be accounted for in the non-linear registration method.

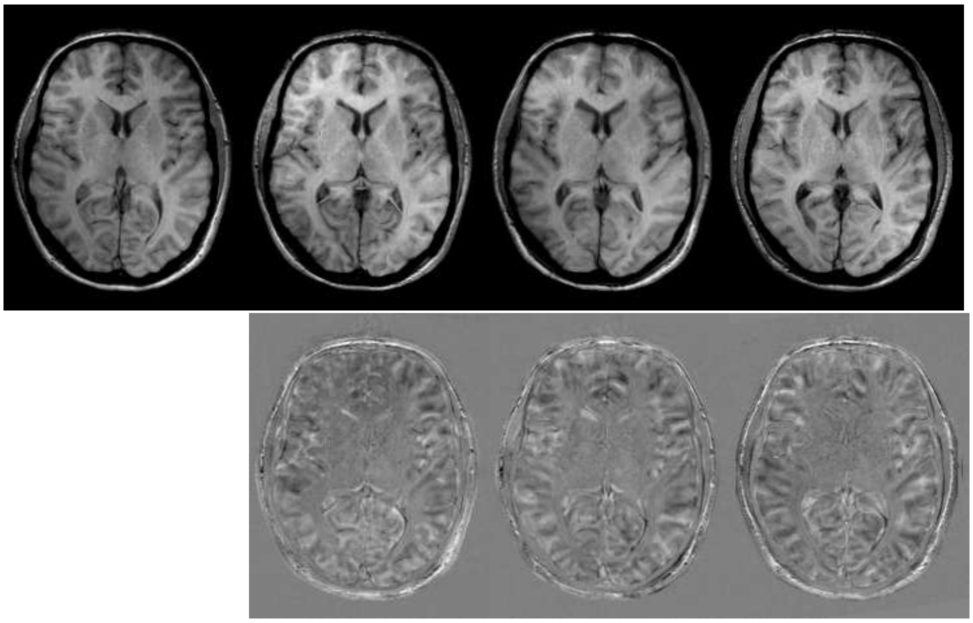

Figure 8.8: Non-linear registration result.
The top row shows the result of non-linear registration to the left-most volume. Much less structure is evident in the difference images of the bottom row, especially near the ventricles, than in the previous figure. This indicates better registration. The intensity range of the difference image is the same as in Fig. 8.7

through the recovered deformations and compared to the corresponding landmarks identified in the target brain. When mapped through the linear transformation only, the r.m.s. registration residual evaluated over all points and all subjects was 4.9±2.5mm. This value was only reduced to 4.2±2.5mm for the non-linearly registered volumes. Hence, despite the apparent improvement in registration in the difference images of Fig. 8.8 over those in Fig. 8.7 and the significant reduction in z-score, the improvement in landmark fitting is only 14.3%. This is explained as follows.

The simulations described above show an improvement of at least 70% in registration residual. After linear transformation, the registration residual is dependent on inter-subject anatomical differences as well as observer-dependent errors. With the difference images and z-score values, it is certain that the anatomical registration is improved by the non-linear algorithm, however the magnitude of the observer-dependent errors may be too great when compared to the inter-subject variability to show a substantial decrease in the registration residual. Experiments completed in section 7.5.1 show that the inter-observer variability is 2.8mm FWHM compared to the non-linear anatomical variability of 3.9mm FWHM when estimated over the same landmarks (see section 7.5.1). This may cause an absolute lower limit in the registration residual, given that only five expert observers identified landmark points.

8.3.2 Stereotaxic mapping with Real MRI data

MRI Atlas target volume

In section 4.6, the linear automatic stereotaxic mapping procedure was described. This procedure was used to map data volumes into a well-defined coordinate system with a transformation having a fixed number of degrees of freedom, thus retaining non-linear anatomical variabilities in order to quantify them. These variabilities must be removed for labelling in segmentation or for increasing the signal-to-noise ratio for functional experiments. Therefore, in order to address this residual anatomical variability after application of the affine transform, the non-linear registration algorithm was added to refine the stereotaxic mapping procedure. For each brain, the non-linear registration algorithm was applied hierarchically at four scales, down to a scale of FWHM=4mm with a voxel size of $2 \times 2 \times 2$mm for the resulting deformation field. A typical non-linear registration for a single brain with the MRI-atlas is shown in Fig. 8.9 and may be compared to the linear registration of the same brain. Morphometric differences in anatomy were accounted for by the non-linear registration procedure, and structures were brought into better alignment with the MRI-atlas, *e.g.*, right parieto-occipital lobe, both anterior temporal poles, both sylvian fissures and insular regions. Note how asymmetries in the occipital lobes have been removed. The example in this figure also indicates the need for a truly 3-D approach to the non-linear registration problem. The alignment of the region of the insular cortex, among others, was possible only by deforming data from above or below into the slice shown in Fig. 8.9, and this type of warp is possible only with a 3-D algorithm.

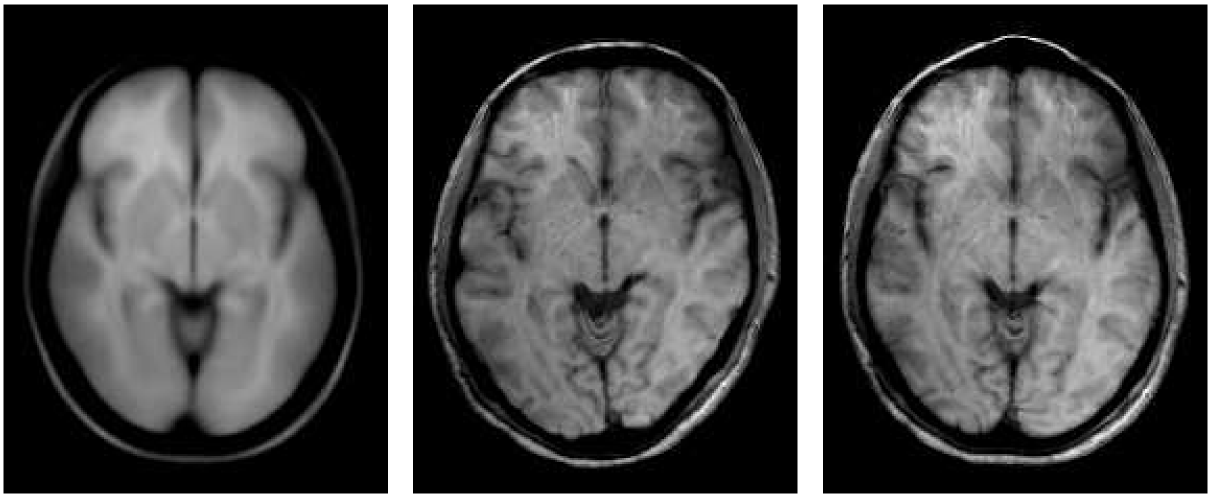

Figure 8.9: Typical stereotaxic transformation: linear vs non-linear. These three transverse images show the result of the linear (**b**) and non-linear (**c**) stereotaxic transformation to the MRI-atlas (**a**). The non-linear registration pulls the right occipital lobe out, so that it matches the MRI atlas. Since the entire procedure is performed in 3-D, the regions of the insular cortex are also brought into alignment from slices above and below. This type of registration problem can never be handled by the techniques that use 2-D warping only.

In order to evaluate the improvement in registration, the seventeen brain volumes used above were mapped into stereotaxic space with both the linear and non-linear registration procedures. The brain-masked z-score gradient intensity (FWHM=4mm) difference between the MRI atlas and each stereotaxically transformed individual was used to measure the quality of the registration for both mapping procedures. The average z-score over the 17 brains for the linear transformation was 0.647±0.032. This measure was reduced by 26% to 0.479±0.014 for the non-linear registration procedure. Figure 8.10-a shows how the z-score is reduced at each scale step in the non-linear procedure. These values for matching against the 305-brain MRI-atlas may be compared with those for matching between two individual brains of 0.656±0.035 (linear) and 0.414±0.020 (non-linear) from page 136. The values are larger here because the target average MRI model does not have the high frequency structure contained in a single-subject target.

The goal of the non-linear registration procedure is to reduce the inter-subject variability remaining after linear transformation. Using the methods described in section 7.5.1, the inter-subject variability (ISV), inter-observer variability (IOV) and residual variability (RV) were estimated. Figure 8.10-b graphs the values of ISV, IOV and RV for each scale step of the non-linear registration. There is a 35% reduction in ISV, from 3.9mm for linear to 2.5mm for the non-linear registration.

Figure 8.11 merges ISV estimates from the linear experiments of section 7.5.1 with those presented here. The graph shows how the inter-subject variability is reduced with additional parameters in the spatial transformation function used to map data from their native space to the stereotaxic coordinate system. However, even though the z-score measure shows improved registration for decreasing scale (Fig. 8.10-a), inter-observer errors in landmark identification limit the improvement measurable by ISV.

After mapping each MR volume non-linearly to stereotaxic space, an average intensity MR brain was created, following the same methodology as that used to build the MRI-atlas for linear and the non-linear stereotaxic mapping. Figure 8.12 shows matched slices through both the linear and non-linearly mapped data set averages. Details in the non-linearly mapped volume are much clearer. For example, the optic nerve is detectable directly behind the orbits in the non-linear average; the gyri of the insular region are aligned; the contrast between caudate, putamen and globus pallidus are enhanced; the region of Heschl's gyrus is aligned such that the individual gyri are identifiable; and many of the cortical gyri are aligned.

This procedure could be applied to the 305-brain data base in order to improve the MRI atlas. In principal, this procedure could be repeated again and again, each time increasing the high-frequency content of the average brain data set and tending toward an ideal high-resolution brain. It is important to note that this process maintains the order-independence in creation of the average brain; each new average is dependent only on the previous MRI-atlas, created independently from any single brain data-set, other than the original brain used for the Talairach atlas.

Not surprisingly, the non-linear average appears sharper than the linear average. This

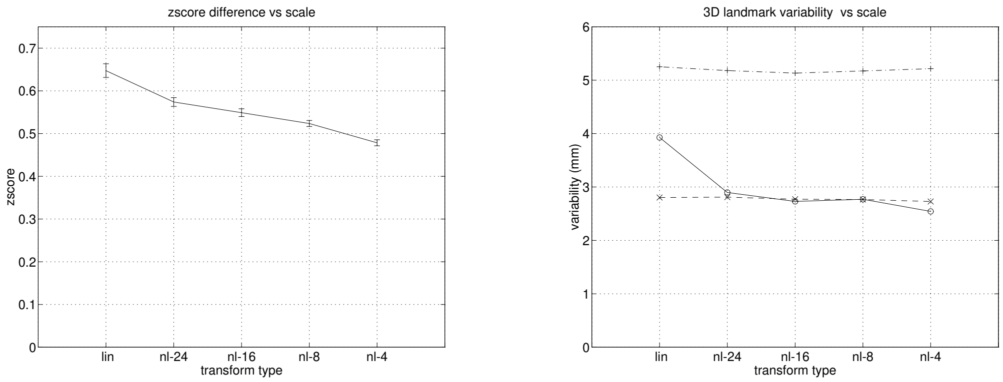

Figure 8.10: Error measures with respect to scale of deformation field. The graph on the left (**a**) shows the average (17 subjects) z-score value for the 9-parameter linear transformation (lin) and the non-linear transformation at scales of 24mm (nl-24), 16mm (nl-16), 8mm (nl-8) and 4mm (nl-4). The tick marks represent one standard deviation on the mean z-score value. The graph on the right (**b**) shows the average 3-D point variability measured as a 3D FWHM for all 34 landmark points sent through each of the transformations: solid line, inter-subject variability (ISV); dashed line, inter-observer variability (IOV); dashed-dot line, residual variability (RV). The flat behavior of IOV and RV continues to support the assumption of additive, orthogonal errors and the use of the model described in eq. 7.5.

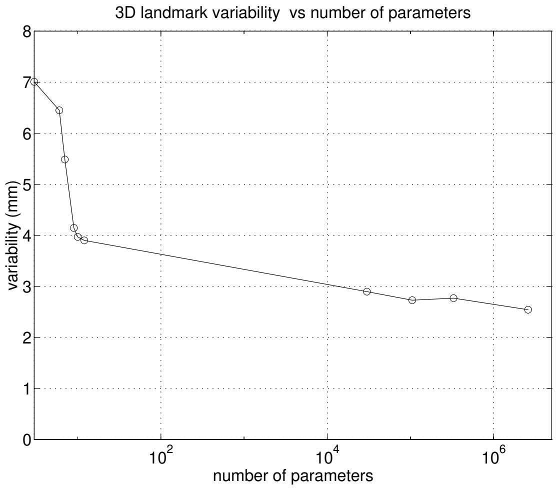

Figure 8.11: 3-D point variability vs number of parameters in transformation.

This graph shows the decreasing 3-D point inter-subject (17 subjects) variability measured as a 3D FWHM as a function of the number of parameters in the stereotaxic transformation function. The first data points on the left of the graph represent the 3,6,7,9,10 and 12-parameter linear transformation. The four data points on the right are the result of application of the non-linear deformation field recovered at 24, 16, 8 and 4mm scales using approximately 3x10^4, 1.0x10^5, 3.3x10^5, and 2.6x10^6 parameters, respectively. The graph seems to reach an asymptote near 2.5mm. This is due to observer errors in identifying landmarks, since z-score values show that the registration between the two volumes has not ceased to improve.

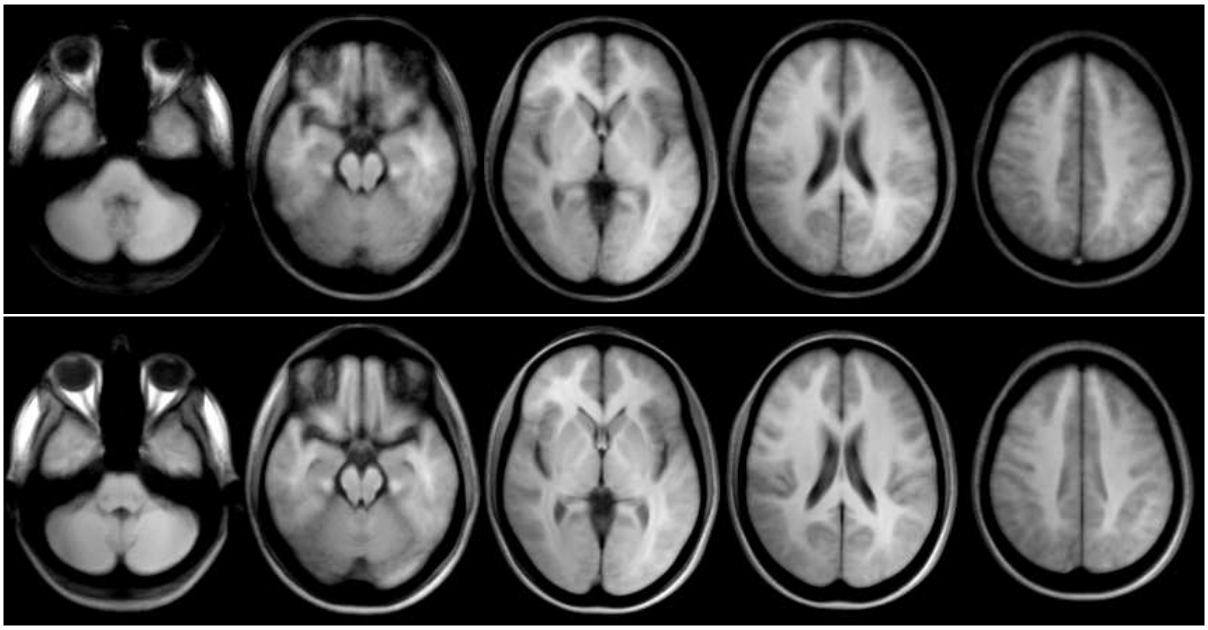

Figure 8.12: Comparison of linear and non-linear average Volumes. In the top row, 5 transverse slices through the linear average of the 17 data sets is shown. The bottom row shows the same slices through the non-linear average. Much more detail can be seen in the non-linear average.

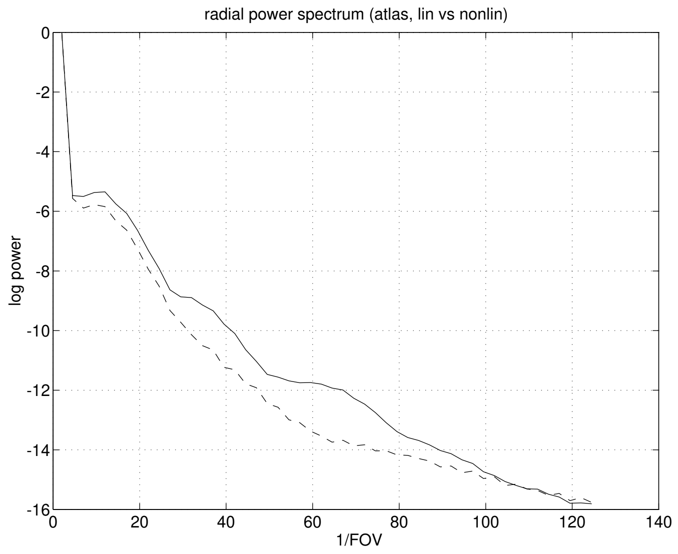

Figure 8.13: Radial spectral power density function of linear/non-linear averages.

This graph shows the $\log_{10}$ normalized average radial power spectrum for the linear average (dashed line) and for the non-linear average volume (solid line). See text for definition of *radial* power spectrum. There is consistently more power at each frequency for the non-linear average. The flat "step" in the region of 5-15 cycles/FOV is due to the overall brain shape.

is confirmed by examining the frequency power spectrum from the two volumes[1]. Figure 8.13 shows the normalized radial spectral power density function of both the linear and non-linear average volumes. This distribution function is calculated as follows: The power spectrum is calculated by taking the Fourier transform of the average data volume and multiplying it by its complex conjugate. Instead of taking 1D cuts along the kx, or ky axes, an average radial spectrum was created. Circular annuli, corresponding to frequency windows, were defined on the power spectrum data. The sum of all voxels within an annulus was divided by the number of voxels in the annulus, yielding the average power for that frequency bin. The set of all frequency bins formed a power spectrum histogram that was the average radial spectral power density function. The spectrum was then normalized so that the integral under the curve was equal to 1.0. Figure 8.13 shows consistently more power at each frequency for the non-linear average.

[1]The power spectrum was estimated on a 2-D slice through the data set instead of in 3-D due to memory limitations and the size of the volumes.

Single brain in stereotaxic space

As shown on page 141, the 3-D point variability based on ISV has been reduced by 35%, from 3.9mm to 2.5mm, by addition of the non-linear component to the stereotaxic mapping procedure. Simulations presented in section 8.2.3 have shown that a corresponding reduction of more than 70% was possible when a one-to-one mapping was guaranteed in simulated data. Since the reduction of ISV on real data was not as great, the non-linear registration procedure is not achieving the required goal of recovering the complete one-to-one mapping required to map an individual brain into stereotaxic space. The simulation experiments were also based on data where IOV=0mm and RV=0mm, so the simulation results were expected to be better even if the real data had a one-to-one mapping.

The smaller reduction in ISV with real data is also due to the inherent blurring of the target MRI atlas caused by residual anatomical variability of all the data sets which were mapped into stereotaxic space with a linear transform. The model does not exhibit many sharp edges, nor does it contain representations of the secondary cortical structures. The MRI atlas is a large-n low resolution data set, and does not have the high frequency components needed to tighten the mapping of homologous points. Therefore it is not be possible to map all corresponding regions from an individual brain into stereotaxic space in a well-defined fashion. The hypothesis tested in this section is that a single individual brain in stereotaxic space will provide the necessary high-resolution structure to reduce landmark point variability and the z-score difference. Even though an individual was selected arbitrarily as a target, the stereotaxic methodology allows all results to be recalculated in an equivalent fashion for the single stereotaxic brain, yielding information on the potential gains of a sharper stereotaxic model.

The brain volume selected to represent a high-resolution target in stereotaxic space was that which yielded the highest correlation with the mean MRI volume after linear transformation. The target model was created by first mapping the single data set through the linear stereotaxic transformation. The features needed for the non-linear registration procedure were calculated and stored. A customized brain mask, necessary to speed up the final stages of the non-linear estimation process, was created by manually editing the standard brain mask to ensure that it fully covered the subject's brain in stereotaxic space. This step is needed since only a linear transformation was used in the first step, and normal variation on brain shape may allow some areas of cortex to extend beyond the standard mask, rendering them inaccessible to the non-linear procedure. With these steps completed, the 305-brain VBSM-IIF was replaced by an equivalent data structure for the chosen individual.

The experiment described in section 8.3.2 was repeated, using the single brain volume as the target. The z-score gradient intensity (FWHM=4mm) difference, 3-D point variability and frequency power spectrum analysises were completed for the remaining sixteen data sets mapped through the recovered non-linear deformation. Note that the z-score difference is calculated between each of the sixteen volumes and the individual data set selected as the target, instead of the MRI-atlas as was done for the previous ex-

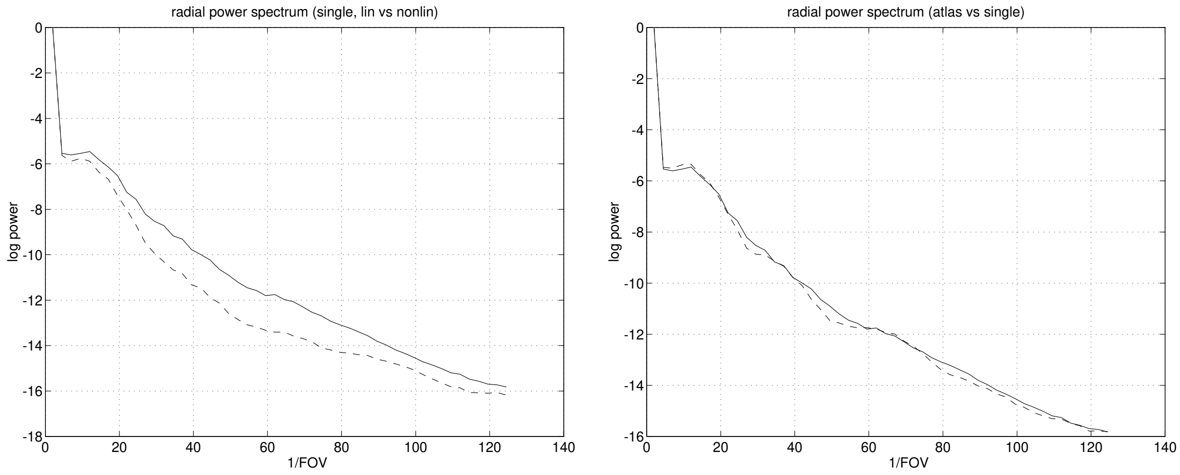

Figure 8.14: Radial spectral power density function comparisons. The graph on the left shows the $\log_{10}$ normalized average radial power spectrum for the linear average (dashed line) and for the non-linear average volume using the single brain as a target. As before, the "step" from 5 to 15 cycles/FOV is due to overall brain shape. Once again, there is consistently more power at each frequency for the non-linear average. On the right, the normalized average radial power spectra are compared for the non-linear averages using the MRI atlas as a target (dashed line) and a single brain as a target (solid line). No significant difference in frequency content is demonstrated.

periment. Since the value of the z-score measure is dependent on the target selected, the absolute values of the z-scores are not directly comparable. However, the improvement of non-linearly transformed data over linearly transformed data can be compared, since the target volume bias is removed.

When mapped through the non-linear transform, the average (subjects=16) brain-masked z-score improved by 37%, from 0.656±0.035 for linear to 0.414±0.020 for non-linear. This improvement was greater than that measured when using the MRI atlas as a target (26%, from 0.647±0.032 to 0.479±0.014, for linear and non-linear, respectively). On the other hand, the inter-subject variability estimated over the 34 landmarks passed from 4.1mm for linear to 2.5mm for non-linear, a 39% improvement. This result was not significantly better than the previous experiment on page 141, where ISV was reduced by 36%, starting at 3.9mm for linear and decreasing to 2.5mm for the non-linear transformation. This result is confirmed by the radial spectral power density function. Figure 8.14 shows that there is no significant difference in the frequency distribution of the averages formed by using a single brain as a target, or by using the MRI atlas as a target.

These results lead to the conclusion that the single individual brain chosen did not provide the high-resolution structure necessary to reduce the inter-subject variability in stereotactic space. One other explanation may be that the non-linear registration process was ended too soon, and that the deformation estimated at a 2mm FWHM scale would improve the fit. This hypothesis was tested in the next section.

8.3.3 High-resolution deformation field estimation

The non-linear registration process operates in a multi-scale fashion, ending at a scale of 4mm FWHM, with samples every 2mm in the x, y and $z-$directions. To test the hypothesis that the registration does not get better by going to higher resolution, the following experiment was performed. Since the volumetric data sets used in the previous sections had a slice thickness of 3mm, higher-resolution data were needed to test the hypothesis.

The recent arrival of a new Philips Gyroscan ACS 1.5 Tesla superconducting magnet system at the MNI permitted the acquisition of isotropic high-resolution volumes with $1 \times 1 \times 1$mm voxels that covered the whole head with a scan time of less than fifteen minutes. Eleven subjects were scanned using a T1-weighted 3-D spoiled gradient-echo acquisition with sagittal volume excitation (TR=18, TE=10, flip angle=30^o, 140-180 sagittal slices). This sequence achieves a high grey/white matter contrast with good signal-to-noise characteristics (15:1).

One of these volumes was selected as the target and the other ten data sets non-linearly registered to it, using the same methodology as in section 8.3.1 where the fitting ended at the 4mm scale. Since the non-linear registration process would require an additional 40-50 hours of computation time per volume to calculate the fit at the 2mm scale, the deformation at this scale step was estimated only on a sub-volume of $58 \times 96 \times 90\text{mm}^3$ centered on the basal ganglia. Figure 8.15 shows three tomographic slices through one of the data sets, cropped at the limits of the sub-volume.

The z-score measure was calculated for the gradient magnitude volumes of both the individual target volume and each resampled source volume at the FWHM=2mm scale. The z-score was calculated for the linearly transformed and each of the non-linearly transformed gradient magnitude data sets, over the region defined by the sub-volume. The graph shown in Fig. 8.16 shows that on average over the 10 data sets, the zscore decreases at each step of the registration process. However, the improvement from the 4mm scale to the 2mm scale was not statistically significant, using a two-tailed Student's t-test for paired samples (t=0.206, p¿0.8). Hence, there was no evidence to justify the significant increase in computation time to calculate the non-linear deformation at the 2mm scale.

8.4 Comparison to manual warping

The thin-plate spline (TPS) interpolant described in section 2.2 has been proposed as an effective means to transform a single data set onto another in a non-linear fashion

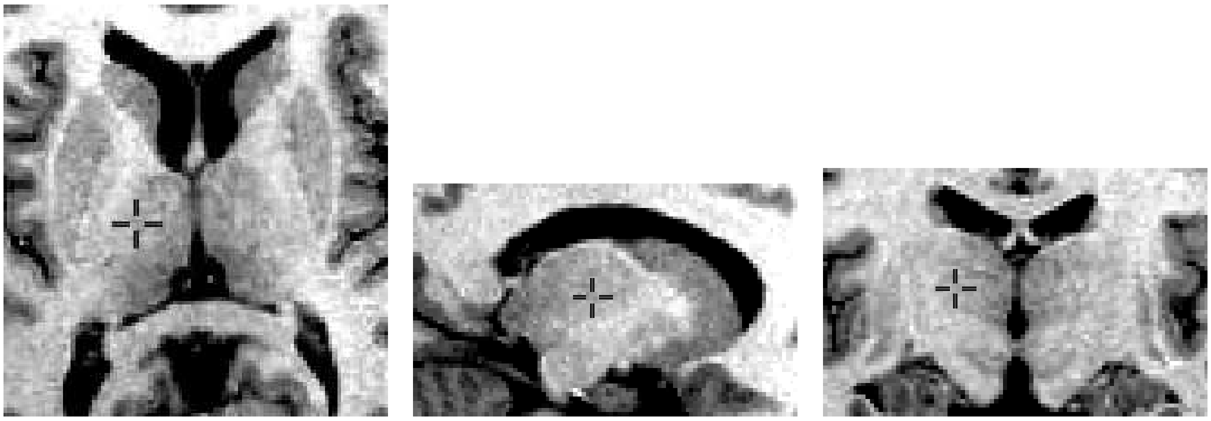

Figure 8.15: Sub-volume used for estimation of 2mm deformation
From left to right, these images show a transverse (z=5.5mm), sagittal (x=-14) and coronal (y=-14) slice through the sub-volume on one subject used to estimate the high-resolution deformation field.

Figure 8.16: Zscore difference vs scale for high-resolution deformation
This graph shows the average (10 subjects) z-score value for the 9-parameter linear transformation (lin) and the non-linear transformation at scales of 24mm (nl-24), 16mm (nl-16), 8mm (nl-8), 4mm (nl-4) and 2mm (nl-2). The tick marks represent one standard deviation on the mean z-score value. Note that the registration improves at each scale of the non-linear fitting process. However,

(Duchon, 1976; Bookstein, 1989; Evans *et al.*, 1991b). It therefore provides a useful point of comparison for the automated non-linear 3-D procedure. This procedure was used to compare the automatic non-linear transformation to a manual method. Since the TPS algorithm requires a set of homologous points to be identified on both the source and target volumes, the 34 landmark points identified in section 7.5.1 were used to defined the 3-D interpolant. The individual brain selected to define the target space for experiment in section 8.3.2 was used here. The manually-defined transform was applied to the remaining 16 data volumes, transforming them into the stereotaxic space. After resampling, a voxel intensity average data set was created for comparison with the average created with the automatic non-linear registration procedure in the previous experiment. Figure 8.17 shows 5 transverse slices through both data sets. These average images clearly indicate better registration for the automatically transformed volumes. The 3-D landmark point variability cannot be used for quantitative measurement of the manual registration quality, since the TPS interpolant exactly fits the data at the landmark points, hence reducing the variability to zero at these points. However, both z-score and the 3-D power spectrum analysis can be used to compare the two techniques. Figure 8.18 shows a comparison of the radial spectral power density function for average volumes created with the automatic non-linear registration procedure and with the manual non-linear transform. There is consistently more power at each frequency for the automatic non-linear average.

The average brain-masked z-score of the resampled 4mm intensity gradient volumes for the TPS transformation was 0.730 ± 0.034. This value was much higher than the z-score estimated for the automatic non-linear transformation (0.479 ± 014). In fact, the z-score for the TPS transformation was greater than the z-score for the linearly transformed volumes (0.647 ± 0.032). The manually-defined non-linear results are highly dependent on the points selected to define the TPS. The variability remaining in the cortical region of the manual non-linear average is due to the lack of landmarks in that region. The 34 selected landmark points are certainly not the “best” set of landmarks that could have been chosen. Ideally, one should select landmarks in those regions with highest anatomical variability and many landmarks should be chosen for better local fitting. Even though five neuro-anatomically trained physicians identified the landmark points, they found it very difficult to define and then locate points on the cortex and other regions of high variability. Due to the nature of the TPS interpolant, regions away from the center of the object (like the cortex) that are not “tied down” with landmarks are subject to possible lever-type forces from other landmark points. However, in practice, the identification of many of these landmarks is impractical. The automatic non-linear registration procedure sacrifices exact correspondence that is available with the manual technique, however the automatic method yields a good estimate of point correspondence throughout the whole brain volume. These results suggest that the approach described here is superior to the manual non-linear warping procedure.

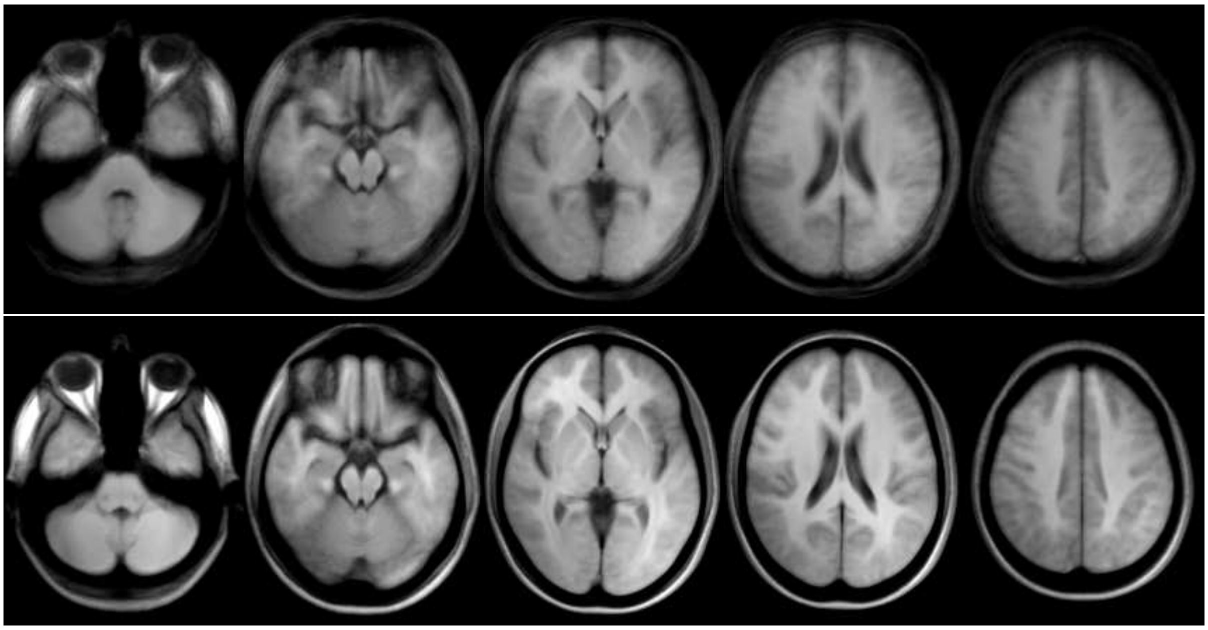

Figure 8.17: Manual vs automatic non-linear registration.
The manual non-linear registration procedure based on the thin-plate spline was compared to the automatic non-linear registration procedure. Using the same volumetric data set as a target in stereotaxic space, the sixteen other data sets were registered to it using both methods. The top row shows 5 transverse slices through the average of 16 data sets mapped into stereotaxic space using the manually-based transformation. The bottom row shows the same slices through the automatically non-linear transformed data set average.

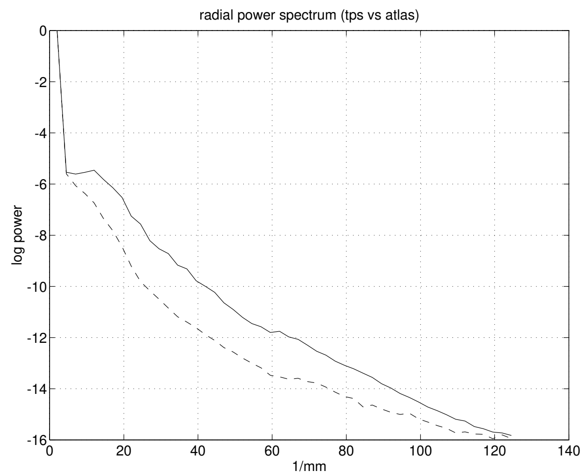

Figure 8.18: Radial spectral power density function of manual/automatic averages.

This graph shows the $\log_{10}$ normalized average radial power spectrum for the average volume created with the manual thin-plate spline method (dashed line) and for the average volume created with the automatic non-linear registration procedure (solid line). Each frequency shows more power for the automatically produced average volume.

8.5 Average anatomical variability map

The result of application of the non-linear registration procedure yields a volumetric deformation field. For a given subject, each voxel of the deformation field represents the magnitude and direction of the warping vector needed to locally match the individual to the target after affine transformation into stereotaxic space. In the volumetric representation of the inverse function, the field is stored such that each vector represents the deformation that must be applied to a given stereotaxic position to match the homologous position in the subject's brain volume after application of the linear stereotaxic transformation. Hence, a vector is stored for each voxel in stereotaxic space that is a measure of the non-linear spatial difference between the MRI-atlas and the subject, after normalization for brain shape, size and orientation.

When the deformation vectors are combined from a large number of subjects at each voxel, the resulting volume indicates the amount normal anatomical variability with respect to the MRI-atlas for each point. Specifically, the standard deviation at each voxel position is computed separately for each component volume (x, y and z) of the deformation volume. Hence three volumes, σ_x, σ_y and σ_z are produced from the 17 deformation fields. The automatic estimate of inter-subject variability is calculated as an average 3-D

FWHM value using a definition similar to eq. 7.9 at each voxel:

$$\text{FWHM} = 2.35\sqrt{(\sigma_x^2 + \sigma_y^2 + \sigma_z^2)/3}. \tag{8.1}$$

Fig. 8.19 shows such an map of morphometric variability of the human brain, averaged over 17 normal subjects.

The anatomical variability map is not symmetric on left and right sides. The left frontal lobe and right parieto-occipital lobe appear to be more variable than their counter parts. As indicated in Fig. 8.19, the area near the posterior superior temporal gyri appear to be more variable on the left than on the right. There can be a relatively large change of variability over a short distance. For example, in the sagittal plane along the longitudinal fissure, the anterior aspect of the inter-thalamic adhesion has a variability of 1.8mm, however the posterior aspect of the structure has a mean deformation value of 4.7mm with the distance between the two points being less than 10mm.

For this map to be valid it must be compared to an existing manual estimation technique. The manual estimates of ISV described in section 7.5.1 were used for two comparisons, since they were completed on the same 17 data volumes. 1) In those experiments, the average inter-subject variability, manually estimated over 34 points, was 3.9mm measured as a 3D FWHM. The corresponding measure estimated automatically for all brain voxels in the morphometric variability volume was 4.2mm. 2) Using the "grand average" coordinate of each of the 34 landmarks from section 7.5.1, the automatically determined variability was extracted from the volume and correlated with the manually determined ISV estimates, both measured as 3D FWHM. Figure 8.20 shows a plot of manual verses automatic ISV values. The regression coefficient was 0.867, demonstrating good correlation between both methods at the 1% significance level. The regression line had a slope of 0.80, indicating that the automatic procedure may be slightly underestimating the variability. This is due to two reasons: 1) as described in section 8.2.2, at each step of the non-linear registration procedure, only a fraction of the required local deformation is added to the global warp, so that an edge will never completely reach its target with the limited number of iterations allowed. 2) More importantly, the deformation is estimated only on voxels with a gradient magnitude that exceeds a preset threshold. Therefore, the deformation is not estimated in areas that appear homogeneous in MRI, but is interpolated from nearby edges. Since the weighting is inversely proportional to distance, voxels in these regions can only under-estimate the true variability.

8.6 Summary

In warping techniques based on landmark matching (Bookstein, 1989; Evans *et al.*, 1991b), correspondence is defined explicitly by manually defining homologous points. The problem with these methods resides in the choice of which landmarks to use, how many to use and the subjective error in identifying equivalent landmarks in different brains. In the automated method presented here, explicit point-to-point correspondence is not required.

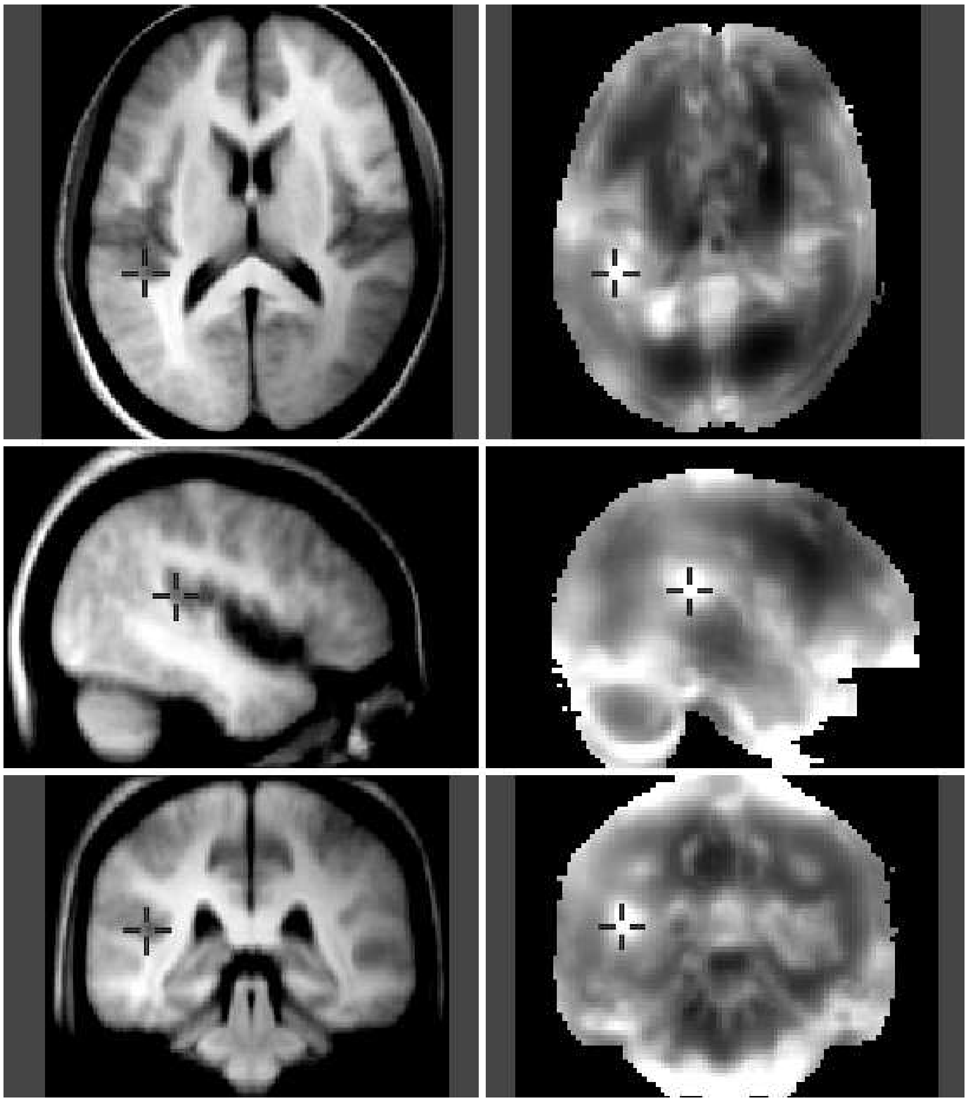

Figure 8.19: Average variability map.

The images on the left show three slices through the average intensity volume of the 17 subjects mapped into stereotaxic space non-linearly. The three on the right show the corresponding slices through the average variability map, calculated from the deformation vectors from each subject for each point in the standard space. From top to bottom are transverse (z=14mm), sagittal (x=-44mm) and coronal (y=-37mm) images. The average variability for all brain voxels is 4.21mm (3D FWHM). The regions of largest neuro-anatomical variability were posterior poles of the lateral ventricles, the region near the fourth ventricle, the cingulate sulcus (slightly more on the left than the right), the inferior frontal lobe and the area just above the splenium of the corpus callosum.

The cross marker is near the region of the planum temporale, the most posterior aspect of the superior temporal gyri, a region known to be variable, measured here to be 6.3mm 3-D FWHM. It is interesting to see that in this region, the left side of the brain appears more variable than the right.

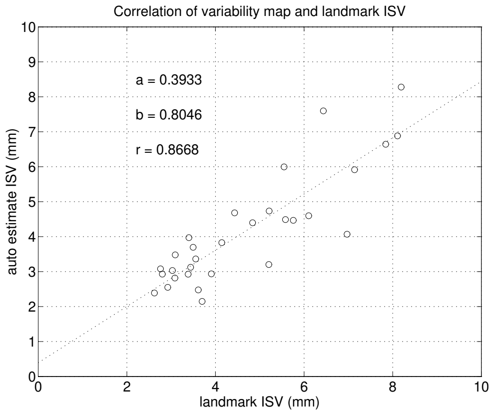

Figure 8.20: Correlation of automatic and manual ISV estimates

This graph shows the correlation between the automatic estimate of neuro-anatomical variability and the manual estimate based on ISV (section 7.5.1 for each of the 34 landmarks. The regression coefficient is 0.867, and is significant at the 1% level.

It is derived by maximizing local correlation between source and target volumes with respect to a cost function. This achieves surface-to-surface correspondence since a surface from one data set, defined by high values of the gradient magnitude, is constrained to fit on the closest surface within the target volume. While the radial component (perpendicular to the surface) of the deformation is is well constrained, displacements in the tangential direction do not affect the similarity function as much. However, experimental results show that even under noisy conditions this is not a practical limitation. The registration residual is reduced significantly, indicating that anatomical variability is reduced in the stereotaxic reference frame.

Experiments with simulated data have shown that the non-linear registration procedure can reduce the registration error by at least 70%, resulting in mis-registration of less than 1mm for homologous points. T1-weighted data appears to be better suited than T2 or PD-weighted data for non-linear matching due to the increased contrast between grey-matter and CSF that yields more edges at which the algorithm can estimate the necessary deformation vector. Inter-subject registration experiments with real data indicate that the registration residual was decreased by only 14.3%, however other measures showed a substantial improvement over the linear registration. Inter-observer variability, on the order of 2-3mm, forces a lower limit on the estimation of the residual based on landmark data

The non-linear registration algorithm was incorporated into the stereotaxic mapping procedure, in order to reduce inter-subject variability in the standardized space. Using the statistical landmark-based tools developed in chapter 7, the estimated inter-subject variability was reduced from 3.9mm to 2.5mm, again limited by the inter-observer variability. Mean MR intensity volumes were created by averaging stereotaxically transformed volumes. The automatic non-linear average appeared much sharper than the linear average or the manual non-linear average. These results were confirmed using the z-score difference measure and power spectral density function methods.

Finally, an anatomical variability map was created and compared to the estimates of ISV from chapter 7. A regression coefficient of 0.867 indicates significant correlation between the manual and automatic estimation methods.

Chapter 9

Segmentation: experiments and results

9.1 Introduction

In this chapter, the registration strategies are employed to effect the primary goal of this thesis: neuro-anatomical segmentation of structures of the human brain from MRI volumetric data. Recall that the segmentation procedure inverts the common segmentation strategy by using non-linear registration to compute the spatial deformation required to map intensity based features of a source volume onto similar features stored in the VBSM. Delineation of structures is accomplished by mapping the VBSM-VOI through the inverse transformation, thus outlining structures in their native acquisition space.

The experiments in this chapter are divided into three parts. 1) Simulated data is used to demonstrate the validity of the new segmentation strategy. 2) Real data is used in the second part to evaluate the segmentation method with respect to the primary goal set for the thesis. 3) The segmentation procedure is used to generate automatically some examples of how neuro-anatomical structures can be represented as probability functions in stereotaxic space.

9.2 Simulated data

9.2.1 Ellipsoid phantom

The simple ellipsoid phantom of the previous two chapters was the first data set used to demonstrate the feasibility of the segmentation procedure. As in section 8.2.2, the TPS warping algorithm was used to deform the original ellipsoid phantom and create a warped target volume. The non-linear registration algorithm was then applied to the original phantom to recover the applied deformation. The geometric contours defined on

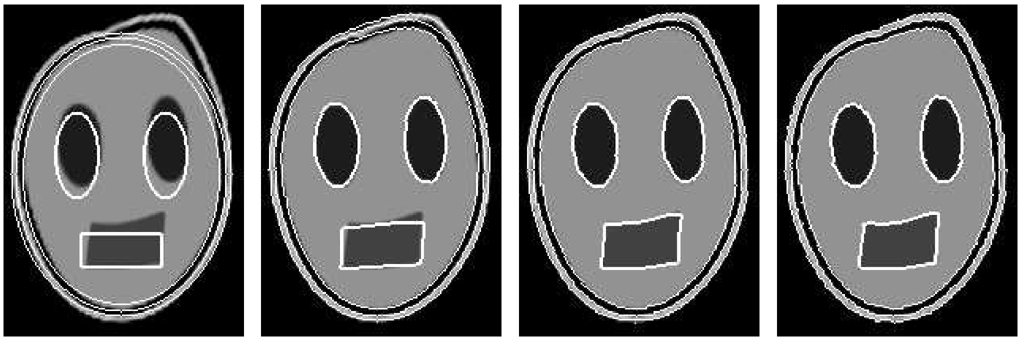

Figure 9.1: Segmentation of ellipsoid phantom.
These images show the result of the segmentation procedure applied to the ellipsoid phantom. From left to right, the ellipsoid contours are mapped through the linear, and the first, third and last scale steps (24mm, 8mm and 2mm) of the non-linear transformation. The first non-linear step begins to take account of the general brain shape, and finds the positions of the ventricles. By the third step in scale, the upper right region of the volume is almost properly segmented as well as the corners of the "mouth", however there appears to be a misregistration of approximately 1mm between the object edge and the segmented contour. This error is corrected, and a perfect segmentation is completed by the last scale step.

the phantom were mapped through the recovered non-linear phantom and overlaid on the target data set. Figure 9.1 shows the resulting segmentation for the linear transformation and the first, third and last scale steps (24mm, 8mm and 2mm) of the non-linear transformation. One can see that there is a significant improvement in segmentation evident at the first non-linear result when compared to the linearly transformed contours. The successive steps in scale refine the segmentation even further, finishing with an almost perfect segmentation of the phantom.

Two quantitative error measures are used to evaluate the quality of the segmentation. The first simply measures the percent difference in absolute volume between the true and the segmented structure. This value is defined as:

$$\delta(T,S) = \frac{V_T - V_S}{V_T} \times 100.0\%, \tag{9.1}$$

where δ is the percent difference in absolute volume, V_T is the volume of the true structure and V_S is the volume of the automatically segmented structure. While this value indicates differences in size between the segmented and true structures, it does not measure the quality of the segmentation, since two structures can have the same volume and not be the same shape, nor be in the same position. The second error measure is the percent volume overlap, Δ, defined as:

$$\Delta(T,S) = \frac{|T \cap S|}{|T|} \times 100.0\%, \tag{9.2}$$

where T is the set of voxels in the true structure, S is the set of voxels in the automatically segmented structure, $\cap$ indicates volumetric intersection and $|\cdot|$ returns the volume of the set of voxels. The structure defined in the VBSM is defined by a binary label volume, i.e., a voxel=1.0 if the voxel is part of the structure, otherwise the voxel=0.0. Tri-linear interpolation is used when mapping the binary model structure through the inverse of the recovered spatial deformation function onto the original source volume. Therefore after mapping, voxel labels defining the segmented structure can have values between 0.0 and 1.0. Consequentially, equation 9.2 is implemented by a function similar to the correlation measure of equation 4.17:

$$\Delta(T,S) = \frac{\sum_{\mathbf{x}\in\mathcal{L}} f(\mathcal{T},\mathbf{x}) f(\mathcal{S},\mathbf{x})}{\left(\sum_{\mathbf{x}\in\mathcal{L}} f^2(\mathcal{T},\mathbf{x})\right)} \times 100.0\%, \tag{9.3}$$

where $\mathcal{T}$ and $\mathcal{S}$ are the label volumes representing the true and segmented structures, respectively, $f(\mathcal{V},\mathbf{x})$ is the interpolated value (between 0.0 and 1.0) from the label volume $\mathcal{V}$ at voxel position $\mathbf{x}$, the summations are done over all elements $\mathbf{x} \in \mathcal{L}$ and $\mathcal{L}$ is as previously defined in chapter 4 in eq. 4.16 with spacing between nodes equal to the voxel size of the original source volume.

When the label volumes are binary valued, Δ exactly measures the ratio of the voxels of the segmented structure that overlap the true structure. If the edges of the segmented structure are blurred, due to partial volume effects after resampling, Δ has a tendency to increase slightly. This value is 100.0 when the segmented and true structures completely overlap. In fact, if the segmented structure is larger than the true structure, Δ can still equal to 100.0. Therefore, one measure forward $\Delta_f = \Delta(T,S)$ and one measure backward $\Delta_b = \Delta(S,T)$ were calculated, and the minimum of the two was used to estimate the segmentation results.

Five objects were used in the definition of the ellipsoid phantom and they can also be employed to test the segmentation. These are the skin surface, the cortical surface and three internal structures: two ventricles and one rectangular object (the eyes and mouth of the phantom). For the segmentation described above, the values of δ and Δ were calculated for each object. The results in table 9.1 show that on average the estimated segmented structure volume has less than 1% error and structure overlap is better than 99%, indicating that the objects are well delineated by the contours mapped through the deformation field recovered by the non-linear registration method.

9.2.2 Brain phantom

A more realistic evaluation of the automatic segmentation employed the digital brain phantom created in section 7.2.2. The set of landmark points identified in chapter 8 were used to define a TPS transformation that was applied to the T1-weighted phantom data set to produce a topologically equivalent, but spatially warped target data set. Ten percent noise was added to the source volume.

structure	*linear*			non-linear		
	δ	Δ_f	Δ_b	δ	Δ_f	Δ_b
skin:	7.1%	92.8%	99.6%	-0.2%	100.0%	99.9%
brain:	6.1%	93.4%	99.1%	0.6%	99.6%	100.3%
ventricle (R):	1.7%	84.8%	85.6%	0.5%	99.3%	100.6%
ventricle (L):	-2.1%	79.6%	77.3%	0.5%	99.1%	100.7%
rectangle:	23.0%	71.1%	90.3%	-1.9%	99.4%	98.3 %
average:	**8.0%**	**84.3%**	**90.4%**	**0.7%**	**99.5%**	**99.9%**

Table 9.1: Segmented structure volume difference and overlap.

This table shows the percent volume difference and the percent volume overlap for the individual structures used to define the ellipsoid phantom. These values are normalized to the true structure volume. The average percent difference in volume is calculated in the last row for δ. Note that since Δ is normalized by a factor dependent on T only, it can take on a value slightly greater than 100.0 due to numerical errors and finite sampling.

Five structures were defined on the original phantom data using a voxel painting feature within a software program, developed at the MNI, that allowed simultaneous tri-plane (transverse, coronal and sagittal) viewing of volumetric data. These structures included the head of caudate (left and right) as well as the lateral and third ventricles. The unwarped phantom and the manually segmented structures formed the model atlas data. Together, they served as the target for the segmentation procedure.

The manually segmented structures were mapped through the applied transformation to create the set of known structures in the warped data volume that must be segmented by the registration/delineation procedure. These structures were used to evaluate the segmentation algorithm using the percent difference and the percent overlap measures defined above. Figure 9.2 shows the result of the segmentation procedure. The contours defined by the painted structures were mapped back through the inverse of the recovered transformation and overlaid on the warped data. The structures in the segmented data are outlined well by the warped contours.

Quantitative evaluation of the brain phantom segmentation is summarized in table 9.2. Since a classified data set was used to create the brain phantom, the set of grey-matter and white-matter voxels were known. These too were mapped through the applied and recovered transformation for evaluation and were included in table 9.2. On average, δ was reduced by more than half, from 4.2% to 1.7%, and Δ was improved 38%, from 70.4% to 97.4% when using non-linear as opposed to linear matching. Not only is there an improvement in the estimated structure volume, but the structures are properly delineated as well. The results for these structures indicate that the segmentation method functions on realistic data.

In order to assess the quality of the non-linear registration based Δ in terms of mis-

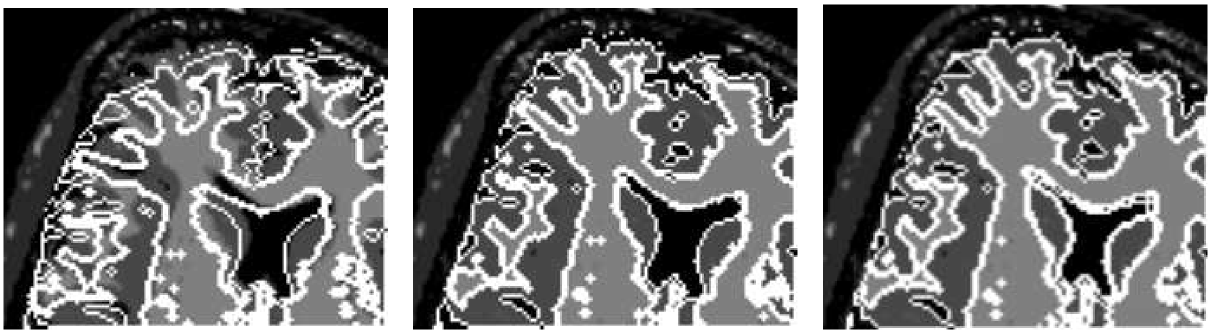

Figure 9.2: Segmentation of digital brain phantom.

Three transverse images of the digital brain phantom through the ventricles at the level of the basal ganglia are shown here with the contours of the resulting segmentation superposed. The contours correspond to the head of caudate, the lateral ventricles and the grey and white matter boundaries. On the left, the contours are mapped through the linear transformation only. In the center are the true contours, generated by mapping the original atlas contours through the applied transformation. On the right, are the contours resulting from the segmentation algorithm. Almost no difference is visible between the two sets of contours.

	linear			non-linear		
structure	δ	Δ_f	Δ_b	δ	Δ_f	Δ_b
caudate:	1.3%	71.5%	65.7%	4.9%	96.4%	101.8%
ventricle:	7.0%	79.1%	77.6%	-0.0%	99.1%	99.4%
grey:	3.4%	77.5%	72.0%	1.0%	97.6%	97.0%
white:	5.0%	71.1%	66.1%	-0.7%	96.3%	95.6%
average:	**4.2%**	**74.8%**	**70.4%**	**1.7%**	**97.4%**	**98.5%**

Table 9.2: Segmented structures from brain phantom.

This table shows the percent volume difference, δ, and the percent volume overlap, Δ, for the individual structures defined on the digital brain phantom and segmented from a manually warped data set. These values are normalized to the true structure volume. The average percent difference in volume is calculated in the last row for δ.

registration distance, a single caudate was taken and deliberately mis-registered by applying a translation along one of the three axis. The value of Δ was calculated for each original/misregistered pair for offset values varying from -5mm to 5mm, in steps of 0.25 mm along the x, y and z coordinate axes. The results are shown in Fig. 9.3. The volumetric correlation (as defined in equation 4.17) on intensity values was calculated as well and plotted for comparison; Δ corresponds well with the correlation coefficient at each offset position. A value of Δ greater than 90% corresponds to less than 1mm misregistration along the x-axis. The same misregistration on either the y or z axes alone corresponds a Δ of approximately 95%. Therefore, the value of 96.4% found for caudate in this experiment indicates that there is less than 1mm misregistration in the segmentation of the simulated data. These values correspond well to the residual error and deformation recovery error measurements described in section 8.2.3 and shown in Fig. 8.6.

9.3 Real Data

9.3.1 Data acquisition

Segmentation experiments were carried out with new volumetric MRI data acquired for the experiment described in section 8.3.3. Since the RF sensitivity of the new machine was appropriately uniform, no preprocessing of the data was necessary before application of the automatic registration algorithms.

9.3.2 Manual structure identification

Testing of the segmentation algorithm requires a set of gold-standard structures to be defined and segmented from the 11 data volumes. The voxel painting program described above was used by a trained neuro-anatomist to identify individual anatomical structures in the brain volumes. The voxels labeled by this procedure serve as the gold-standard definition for each structure on each brain, and were used to evaluate the automatic segmentation result. The left and right head of caudate were chosen for testing, because of their inherent importance in many aspects of current brain research and because they present a relatively well-defined object for manual labelling.

9.3.3 Automatic structure segmentation

Without loss of generality, the segmentation algorithm allows one brain to serve as a target in stereotaxic space and structures identified on it to define the anatomical atlas that is used for segmentation of all other data sets. Therefore one brain volume, and its associated voxel labels, were selected as the model and the other ten data sets segmented automatically with it. The methodology outlined in section 8.3.2 was used to build the single-brain target VBSM-IIF and VBSM-VOI.

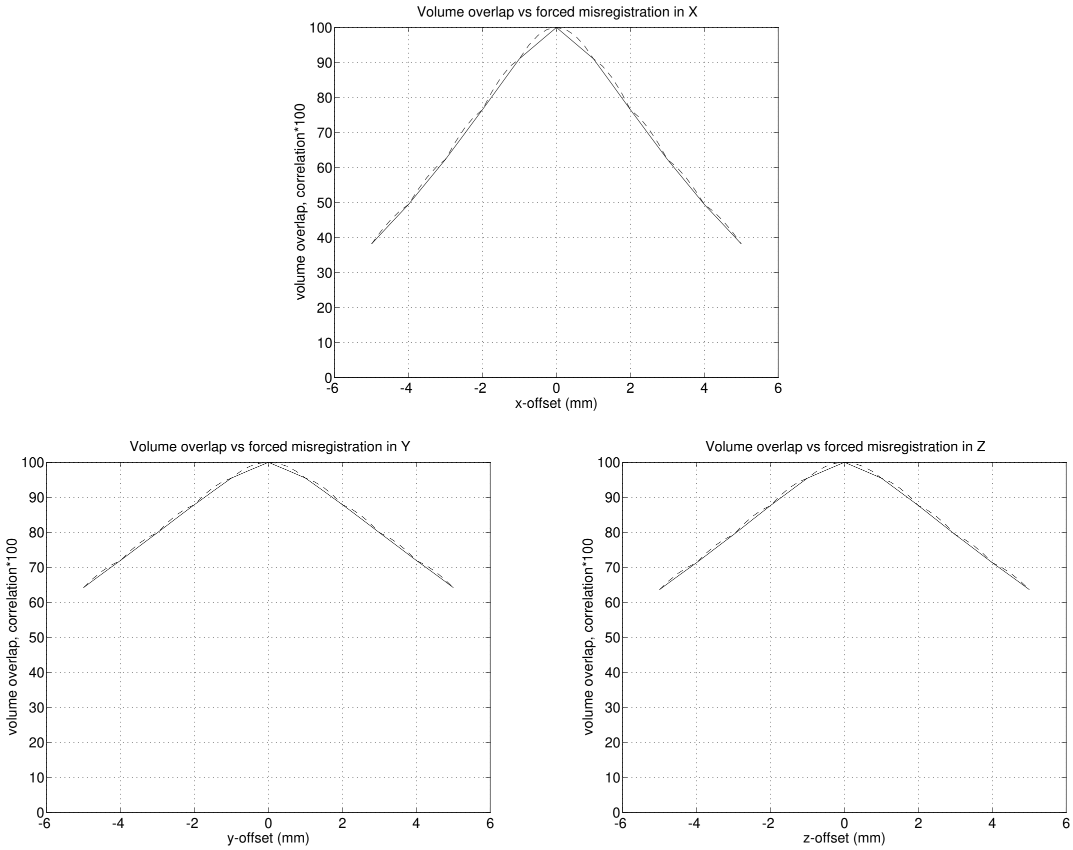

Figure 9.3: Volume overlap and correlation value vs registration error. Plot of Δ (solid line) and correlation value (dashed line) for a single segmented structure (left head of caudate). The data set was deliberately misregistered, varying one translation component (x, y or z) at a time. Note that the graph for x-offset is steeper due to the shape of the caudate. Since the structure is narrow in the x-direction, taller in the y direction and even longer in the z direction, a 1mm x-offset reduces Δ more that an y or z-offset of the same size.

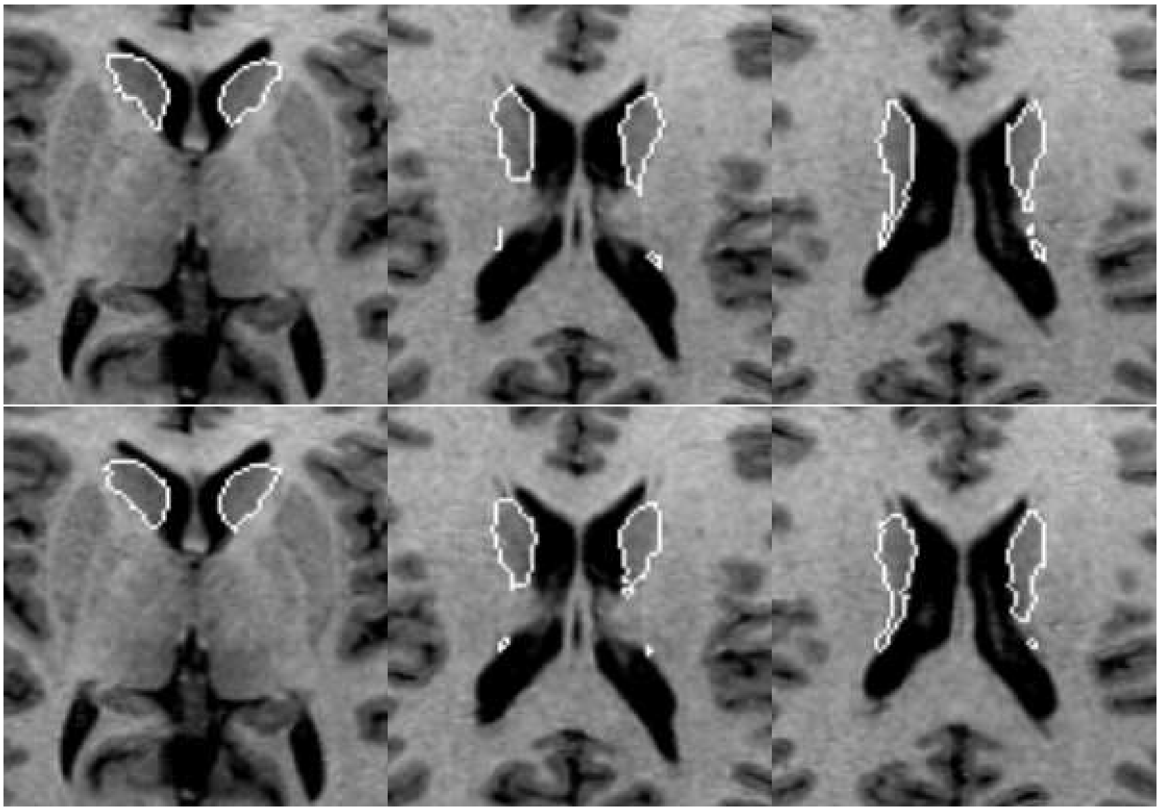

Figure 9.4: Typical segmentation of real MRI data.

Transverse slices through a single brain at the level of the ventricles and basal ganglia. The manually identified contours for the caudate are shown in the top row. The corresponding automatic segmentation is shown in the bottom row. Notice that there is good agreement between the manually and automatically-defined contours.

After application of the non-linear registration algorithm to each of the individual/model pairs, the inverse of the recovered transformation was used to resample the structure labels of the target volume onto the individual data set. Figure 9.4 shows three transverse slices through the ventricles and basal ganglia of one of these segmented data sets. Contours corresponding to the head of caudate are overlaid on the images for both the manual and automatic segmentation. Comparison of the two sets of images shows that the segmentation algorithm succeeds in identifying the boundaries of the structure, based on its definition in the model brain.

The results in Fig. 9.5 show how the values of δ and Δ change with increasing scale step. Table 9.3 shows that the overlap increases from 75.6% to 88.6% on average over the ten caudates. The percent volume difference of the right caudate is greater than in the left, after non-linear segmentation. Since the right caudate is thinner and occupies slightly less volume than the left (5285mm^3 vs 5482mm^3, on average of the manually segmented

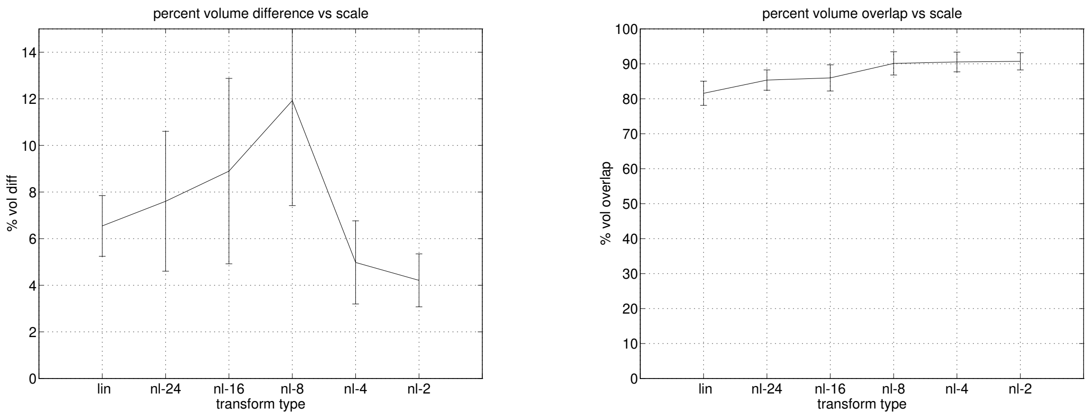

Figure 9.5: Volume difference and overlap for segmentation of caudate. The graph on the left shows the value of $|\delta|$ averaged over ten brains for the automatically segmented head of caudate for the linear transformation (lin) and the non-linear transformation at scales of 24mm (nl-24), 16mm (nl-16), 8mm (nl-8), 4mm (nl-4) and 2mm (nl-2). The height of the tick marks represent one standard deviation on the mean magnitude of δ. The graph on the right shows the value of Δ for the same data.

data), misregistration between the automatic and manually segmented structures cause more important errors for the right side. While the value of Δ increased to more than 97% for the simulations using the brain phantom, Δ increases to only 90% with real data. There are two reasons for this lack of improvement. The main cause may be due to intra-observer variability in the definition of the caudate on different subjects. For example, the anatomist found it quite difficult to identify the anterior-ventral aspect of the caudate near the region of the *nucleus accumbens*, sometimes including it in the structure definition and sometimes not. This region represents 1-2% of the volume of the head of the caudate. Also, since the structures were painted manually, some ventricular voxels were included on the medial wall of the body of the caudate and may account for an additional 3-4% of the difference. The second reason may be due to bias incurred by using the manually segmented caudate from the single individual. Errors in manual segmentation of this structure are carried through the non-linear transformation, yielding errors in the segmentation. This will be true in general for any atlas defined on the basis of a single brain.

9.4 Probabilistic structure definition

One of the goals of this thesis was to develop a technique that could be used to automate the process of neuro-anatomical variability estimation. One aspect of this task involves

structure	*linear*			non-linear		
	δ	Δ_f	Δ_b	δ	Δ_f	Δ_b
caudate (L):	7.3%	77.5%	77.7%	3.7%	89.5%	90.5%
caudate (R):	6.6%	85.6%	73.5%	6.9%	92.0%	86.6%
average:	**7.0%**	**81.6%**	**75.6%**	**5.3%**	**90.8%**	**88.6%**

Table 9.3: Segmented structures from real MRI data.

Percent volume difference, δ, and the percent volume overlap Δ, for the individual structures segmented manually and automatically. These values are normalized to the manually identified structure volume. The average percent difference in volume is calculated in the last row for δ.

the creation of a stereotaxic probabilistic atlas, so called, because each of the voxels in the atlas represent the probability of finding a given structure at that position after mapping a given dataset into stereotaxic space with an affine transformation. In theory, a perfect non-linear transform would render all brains to be identical, and there would be no probabilistic distribution. Segmented data from a number of brain volumes must be averaged in order to build probabilistic definitions of specific structures. This required individual structure segmentation from each brain volume participating in the average and until now, was only possible with manual segmentation techniques. For example, an average thalamic atlas was created by meticulously outlining the left and right thalami on MRI volumes of 200 subjects (Absher, 1993). Each labelled thalamus was mapped into stereotaxic space individually, using the 9 parameter linear transformation identified for the particular subject. The ensemble of left thalami were then averaged together on a voxel by voxel basis. The same was done for all right thalami. Once completed, each voxel value represents the probability of finding the thalamus at that coordinate in standardized space. The intensity profile perpendicular to the surface of the average structure is a measure of the positional variability of that part of the thalamus. This method is manually intensive and prohibitively time-consuming to replace each of the structures in the VBSM-VOI atlas with its probabilistic representation. Furthermore, intra- and inter-observer variations in labelling strategy would confound the overall goal. Completely automatic and accurate regional volume segmentation such as that described in this thesis is required to address these problems.

9.4.1 Evaluation of method

In this section, the hypothesis that the segmentation method developed in this thesis can be used to build a probabilistic structure that is not significantly different from that created manually is tested. The manually segmented head of caudate from the experiment in section 9.3.3 was used to compare the manual and automatic methods.

As described above, creation of the probabilistic structure requires mapping of the segmented structures into stereotaxic space using an affine transformation with a fixed number of parameters. The following procedure was repeated for left and right caudates, segmented manually and automatically, thus yielding four average volumes. For each caudate, the automatically recovered 9-parameter affine transformation identified for the subject was used to map it into stereotaxic space. After transformation and resampling, all caudates for each group (*e.g.*, manual-left) were averaged together on a voxel-by-voxel basis. The resulting volumes were multiplied by a constant so that each voxel represents a percentage value, indicating the probability of caudate for each voxel.

Figure 9.6 shows sagittal slices through the left probabilistic caudate for both manual and automatic averages for qualitative comparison. When visualized using the tri-plane display, the automatic average appears to be slightly lower and more posterior than the manual average. This fact is confirmed by the difference image in Fig. 9.6. When evaluated between the two volumes, δ=-1.3%, indicating that the automatic average is slightly larger than the manual one. The value of Δ_f=100.0, however Δ_b=94.9.

These statistics are broken down for different probability levels and shown in Fig. 9.7. The value of δ decreases with increasing probability values. This is due to a number of factors. The manually created structure is slightly more spread out, so that at low probability levels, it occupies a greater number of voxels. However, as the probability threshold is increased, it is the automatically defined structure that is larger. The larger area occupied by the manually-defined version may be due observer variability in structure definition. The graph of Δ shows that the volumetric overlap is almost constant between the two representations. The sharper decrease between probabilities of 80% to 100% is due to the small size (approximately 3cm^3) of the structures at those levels. A forced 1-D misregistration of 1mm along the x-axis of the manual probabilistic caudate thresholded at the 100% level yields a Δ of 75.8. Similarly, a 1D translation along the z-axis yields a Δ of 85.9 and the same translation along y gives 83.6. For the 50% probability, where the structure is 10.7cm^3 in volume, a 1mm translation along the x, y and z-axis give Δ values of 85.7, 91.35 and 90.9, respectively. Hence from these values, one may infer that the automatically and manually created probabilistic caudate are very similar, differing by no more than 1mm in any single direction.

9.4.2 Application of method

Given the validation results determined on the caudate, we conclude that the automatic segmentation procedure can therefore be used to produce probabilistic structures of gross anatomy of the human brain. Therefore, a number of other structures were identified on the target brain by the neuro-anatomist. These included the corpus callosum, insular cortex, lateral ventricles, thalamus and putamen. These labelled volumes were used to segment these structures on the 10 data sets that were averaged to create the probabilistic caudate.

Using the automatic procedure described above in section 9.4.1, probabilistic represen-

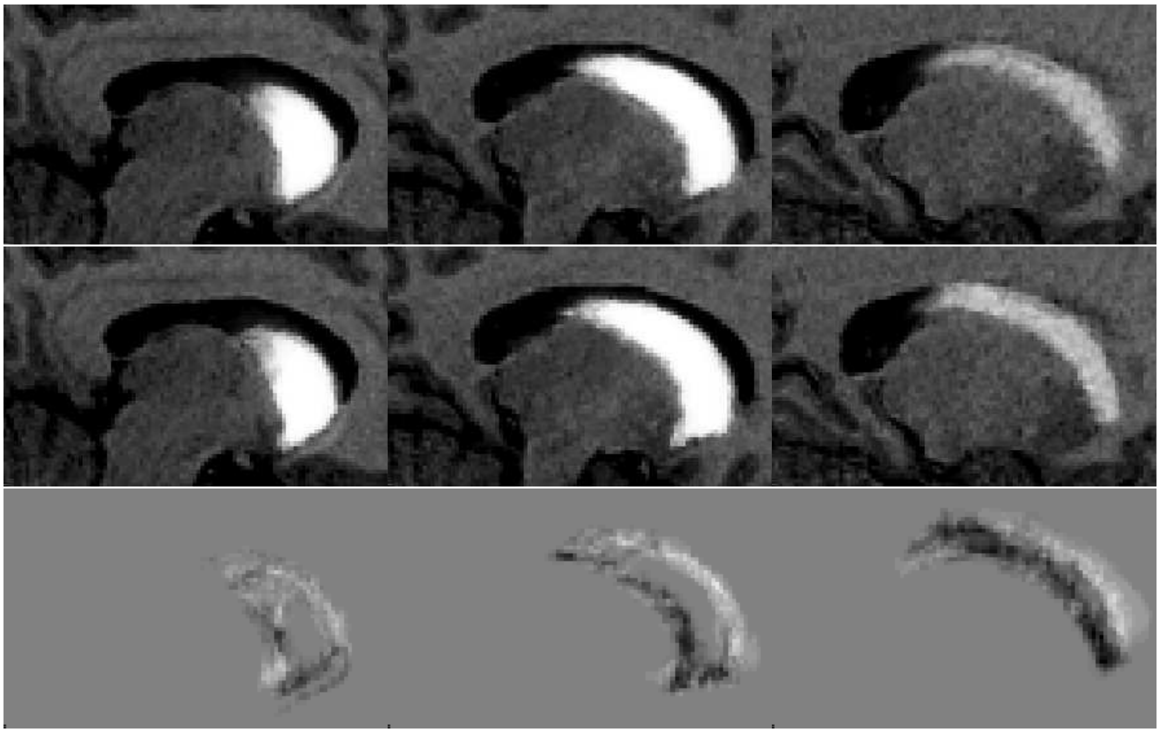

Figure 9.6: Probabilistic caudate.

Three sagittal slices (x=10,15,20mm) through the individual subject used as a target, showing the probabilistic caudate overlaid on the darker MRI. The manually created average is on the top and the automatically created one is in the middle the bottom. The difference image formed by manual – automatic in the bottom row, shows that the latter is slightly lower (1mm) and more posterior than the former.

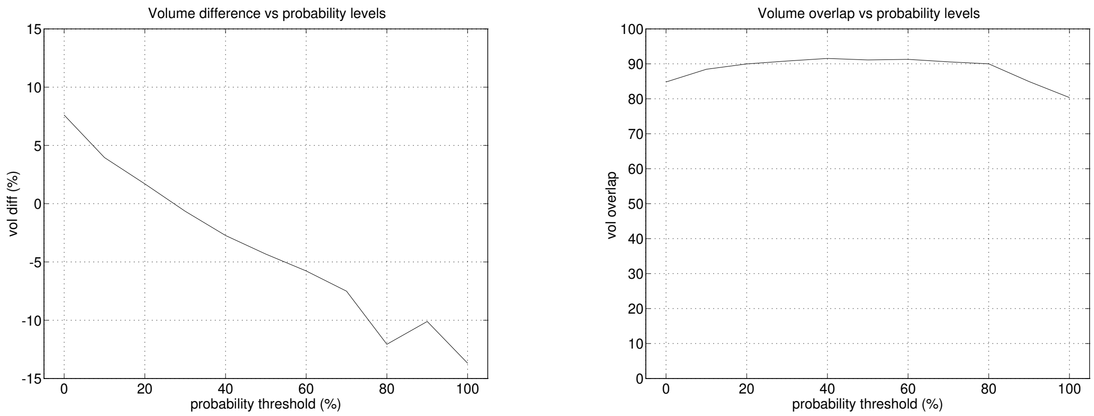

Figure 9.7: Volumetric comparisons at different probability levels. These graphs compare δ and Δ for the manual and automatic probabilistic caudate at various probability threshold levels. The graph on the left shows that δ decreases with increasing probability, indicating that the manually defined structure is more spread out at lower threshold values. The graph on the right shows that Δ between the two representations is relatively constant at 90% overlap.

tations of the corpus callosum (Fig. 9.8), insular cortex (Fig. 9.9), ventricles (Fig. 9.10), and the basal ganglia structures of the thalamus, caudate and putamen (Fig. 9.11) were created. These structures form the building blocks that will be used to refine the VBSM-VOI so that it represents neuro-anatomical variability.

9.5 Summary

A fully automatic segmentation strategy has been presented to address the goal set forth at the beginning of this thesis. The reversal of the standard segmentation strategy, from matching geometric contours directly on image data, to one of registration followed by delineation has been show to be valid. The structures identified by this method are comparable to those segmented manually and the experimental results on both simulated and real MRI data have shown that the method is accurate and robust.

The segmentation method developed here has direct application in the quantitative study of neuro-anatomical variability in the normal population, the determination of inter-group differences, the detection of subtle abnormalities, and the tracking of normal development. It also has direct utility in studies of functional neuro-anatomy with macroscopic imaging techniques such as PET and fMRI. For PET, averaging across subjects in a standardized space is used to detect subtle cognitive activation foci measured from cerebral blood flow (CBF) volumes. The removal of the non-linear individual neuro-anatomical

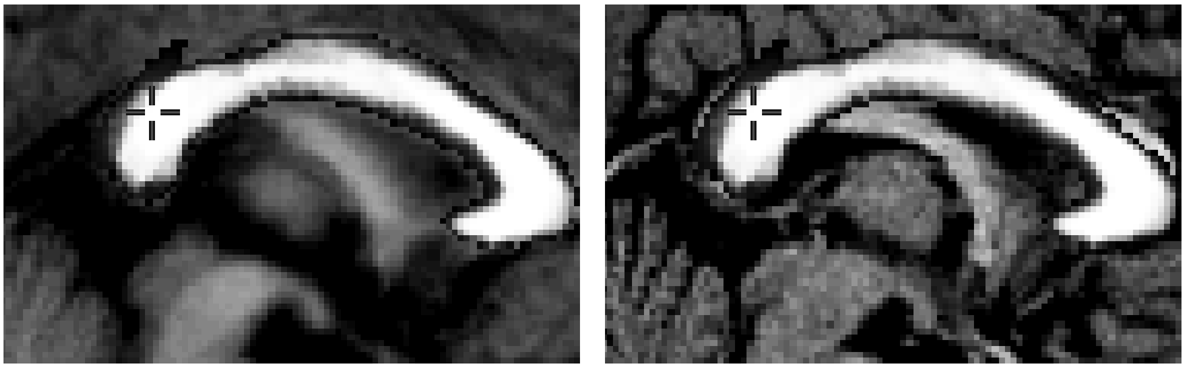

Figure 9.8: Probabilistic corpus callosum

These sagittal images shows the probabilistic corpus callosum (CC) overlaid on the mean intensity MRI data set of ten subjects, linear mapped into stereotaxic space (left) and on the target MRI data set used to define the segmentation. The probabilistic CC has been thresholded at 20%, so that only voxels that have more than a 20% chance of belonging to CC are overlaid on the MRIs. Note that there is a small amount at the superior aspect of the genu of the CC on the target that extends above the average CC. This is due to the neuro-anatomical variability of the selected individual target brain with respect to the average anatomy.

differences will potentially remove the blurring due to morphometric variability and potentially increase the signal-to-noise ratio, and enhance the detectability of even smaller signals. For fMRI, the automatic segmentation procedure can be used to objectively define an anatomical region associated with some cognitive process. The analysis of the fMRI data within these regions may provide clues as to the existence and quantification of neuro-functional variability between subjects.

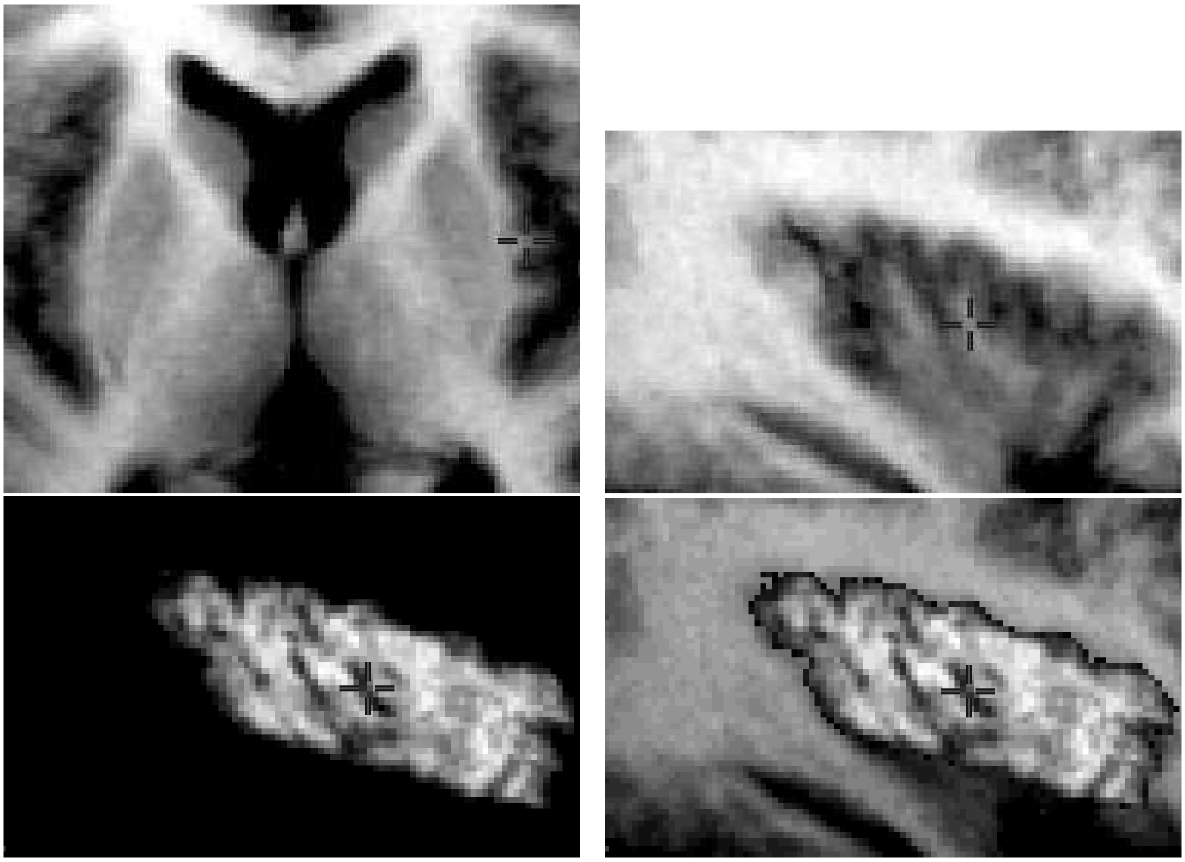

Figure 9.9: Probabilistic insular cortex

These images shows the probabilistic insular cortex and the corresponding slices of the mean intensity MRI data set of the 10 subjects, linearly mapped into stereotaxic space. The probabilistic structure has been thresholded at 25%. The cross upper-left transverse image (z=5mm) of the mean MRI data set shows the sagittal position through the insular cortex for reference for the three other images. The upper-right and lower-left show the sagittal (x=39mm) slice through the mean MRI volume and the probabilistic insular cortex data set, respectively. The probabilistic structure is overlaid on the mean MRI in the lower right. Surprisingly, there is not a lot of variability, and the insular cortex appears conserved in the mean MRI and even more so in its probabilistic representation. Note the agreement between the two volumes: the cross is overlaid in white matter in the mean MRI and this area appears dark in the probabilistic volume, indicating low probability of insular cortex (i.e., grey matter).

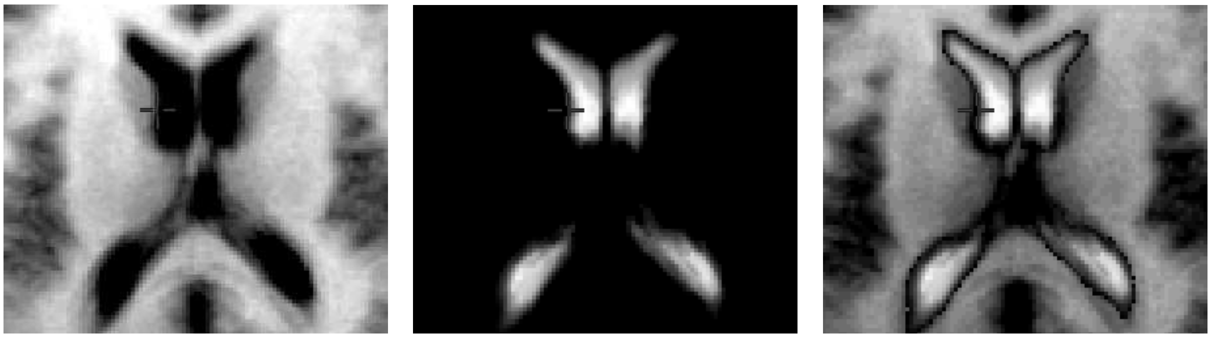

Figure 9.10: Probabilistic lateral ventricles

These three images show a transverse slice (z=16mm) through the mean intensity MRI data set (10 subjects), linearly mapped into stereotaxic space (left), the probabilistic representation of the ventricles (middle) and the two overlaid together (right). The probabilistic structure is thresholded at the 20% level. As expected, the brightest regions of the probabilistic ventricles correspond to the central region of the two lateral ventricles.

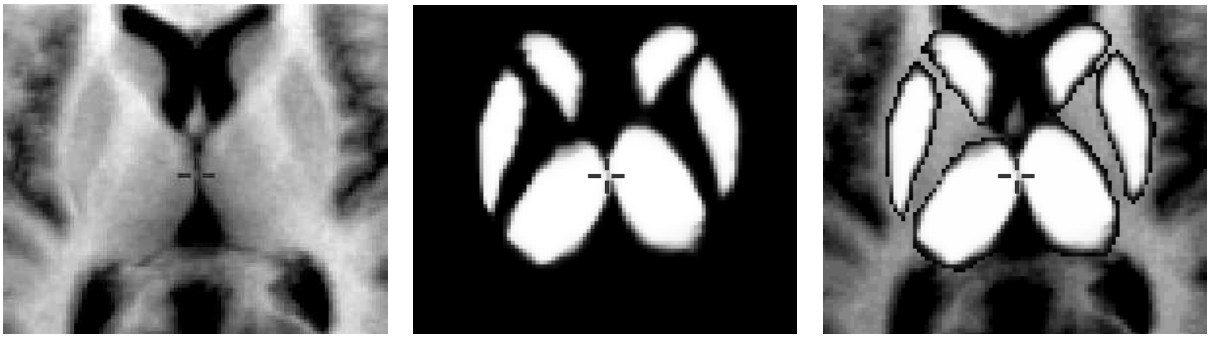

Figure 9.11: Probabilistic basal ganglia

These three images show a transverse slice (z=7mm) through the mean intensity MRI data set (10 subjects), linearly mapped into stereotaxic space (left), the probabilistic representation of the head of the caudate, thalamus and putamen (middle) and the two overlaid together (right). The probabilistic structure is thresholded at the 20% level.

Chapter 10

Discussion

10.1 Rationale

This thesis has addressed a fundamental problem that confounds many neuro scientific endeavors using modern medical imaging modalities: how does one account for neuro-anatomical morphological differences in order to make quantitative comparisons between individuals, or between groups, for a given analysis or experiment? One solution is found in current brain mapping research where 3-D brain image volumes are transformed into the brain-based stereotaxic coordinate system and, using a linear mapping, are resampled onto a common sampling grid, such that all brains have the same position, orientation and size. Hence, the transformation to stereotaxic space removes global variability and permits limited voxel-by-voxel comparisons between data sets, or between the data set and an atlas defined in that coordinate system; limited in the sense that all results must be interpreted with respect to the residual non-linear neuro-anatomical variability which was not removed.

Estimation of this residual variability is not straightforward. It requires evaluation of the variance on position between homologous points from different subjects, after these subjects have been mapped into a common frame of reference with a limited number of well-defined degrees of freedom. Implicit in this analysis is the notion that all points in the brain have been identified exactly, and that there exists a one-to-one correspondence between all brains. This is clearly untrue and one is forced to adopt indirect and inferential techniques to obtain usable information about quantitative neuro-anatomical variability.

Previous attempts to assess neuro-anatomical variability have been limited to a few structures, based on manual analysis of a small number of excised brains while others have used CT and MRI. Many of these studies have been purely qualitative or semi-quantitative, classifying the type of variability in the cortex (Ono *et al.*, 1990) or assessing differences in volume to determine general anatomical asymmetries in the brain (Chui and Damasio, 1980; Galaburda *et al.*, 1978; le May and Kido, 1978), or of specific structures such as the sylvian fissure (Rubens *et al.*, 1976), temporal lobe (Witelson, 1977), temporal speech regions (Geschwind and Levitsky, 1968) or left-right differences in occipital and frontal lobes (Weinberger *et al.*, 1982).

The task of manual structure delineation is difficult on medical images, such as MRI, due to low tissue contrast, partial volume effects and mis-orientation of the scanned slices with those in standard atlases. These difficulties are compounded when the shape of a given structure is different from an atlas, or from subject to subject due to normal anatomical morphometric variability. This can result in inter- and intra-observer observer difference in shape definition due to subjective structure interpretation that cause subjective decisions to be made on each brain when drawing boundaries, for example. Observer-dependent differences in segmentation strategy confound the goal of probabilistic structure creation, since methodological differences in analysis for each brain may be indistinguishable from the true neuro variability. Furthermore, manual labelling is time consuming (*e.g.*, left and right head of caudate required 2 hours on average per subject for one observer), making structure identification on a very large number of brains impractical, signalling the necessity for a fully automated segmentation technique.

The first automated procedures applied to the segmentation problem included classification techniques. While these methods could identify gross tissue types such as cerebrospinal fluid, white or grey-matter, they were unable to achieve the delineation required to separate individual neuro-anatomical structures. Other algorithms are also data driven, extracting "reasonable-regions" that may be reasonable for a computer-vision model, but are not meaningful in terms of anatomy. Expert system-based methods have had moderate success, but become complex and unwieldly when dealing with many structures and are highly sensitive to the nature of a required pre-segmentation process. Model-based information must be used to better constrain the problem. A landmark-driven 3-D non-linear deformation procedure for matching a model to different brains was described in section 2.2. However, as described in section 8.4, this method requires the manual identification of homologous points in different brains and is time-consuming and highly dependent on the landmarks selected and the interpolant used. Many automatic model-based algorithms fit a model to the image using some optimization criteria with physical constraints to force contours in the model to match edges estimated from the data. Of the few successful techniques (Broit, 1981; Bajcsy *et al.*, 1983; Terzopoulos and Witkin, 1988; Zhengping and Mowforth, 1991; Nastar and Ayache, 1993; Kosugi *et al.*, 1993), fewer have been applied in 3-D (Dann *et al.*, 1988; Bajcsy and Kovacic, 1989; Terzopoulos and Metaxes, 1990),none have been subjected to a thorough validation and evaluation.

10.2 Summary of the methodology

The problem statement expressed in chapter 1 is repeated here:

> **Problem:** Given volumetric magnetic resonance image data, develop a procedure to automatically identify and delineate structures in the human brain

that will facilitate neuro-anatomical quantitative analysis and permit characterization of morphometric variability across subjects.

Thus, addressing the problems of existing manual and automatic segmentation methods has formed the basis of this thesis. The segmentation procedure developed here inverts the intuitive segmentation approach. Instead of fitting geometric contours from an atlas directly to the raster image data, a volumetric brain segmentation model (VBSM) was created that contains both geometric atlas data (VBSM-VOI) and volumetric raster data (VBSM-IIF). Both components are registered together by nature of the model definition in the brain-based coordinate system known as stereotaxic space. The unique methodology presented here separates segmentation into two stages: a registration step followed by a delineation step. The optimal spatial transformation between an individual data set and the model is found in the registration task by maximizing the overlap between volumetric features derived from the data set and those stored in the VBSM-IIF. The inverse of the recovered transformation is used to complete the segmentation process. Structure delineation is achieved by applying the inverse function to the VBSM-VOI, thereby outlining structures in the native data.

The two stage approach is broken up into three steps: linear registration, non-linear registration and segmentation. The linear registration procedure used optimization over transformation parameters to maximize the cross-correlation of invariant image intensity-base features between two volumetric data sets. When combined with VBSM used as a target, the technique was used to automate the stereotaxic mapping procedure. Experiments completed in chapter 7 showed that the automated procedure was comparable to existing manual stereotaxic mapping methods, and yielded more stable results. The automated method is completely objective and has benefits and accuracies similar to that of landmark-based registration, or surface to surface matching techniques, while having the additional advantage that explicit manual intervention is not required to identify such points or contours. This eliminates the well-known drawbacks of manual techniques such as reproducibility and inter-observer variability. When the procedure was used to map data volumes into stereotaxic space, inter-subject variability measured by identifying homologous landmark points was estimated to be 4.11mm.

The goal of the non-linear registration procedure developed in chapter 5 was to reduce this variability in stereotaxic space by estimating the non-linear spatial warping function required to map all points from one subject's brain to their homologues in a second brain. The procedure represents the non-linear transformation in the form of a spatial warping field, where a 3-D deformation vector was stored for each voxel in the field. The recovery of the global warp is a straightforward extension of the linear registration method, except that the target volume is a small neighbourhood of the whole brain, recursively selected by stepping through the entire target volume in the 3-D grid pattern defined by the deformation field. Experiments on a simulated brain phantom showed that the procedure could recover more than 85% of the residual not accounted for by the linear model (see section 8.2.3), and was robust against noise (section 8.2.3, Fig. 8.6). Application of the non-linear procedure to the set of real MRI volumes used in the inter-subject variability study, reduced this variability by 35%, from 4.11mm to 2.50mm. Furthermore, the deformation fields created for each subject were averaged together to form a 3-D anatomical variability map within stereotaxic space (section 8.5). The values from this map were highly correlated ($r = 0.867$) with the statistical estimates of inter-subject variability based on manually-identified landmark points (Fig. 8.20).

The combination of linear and non-linear registration procedures forms the engine that drives the unique segmentation method developed in this thesis. Segmentation becomes simply a by-product of the non-linear transformation required to map one data set onto the model. Because the VBSM-VOI that defines the structures to be segmented takes no part in the registration process, the segmentation procedure is completely atlas-independent. Therefore, any atlas defined within the stereotaxic coordinate system of the VBSM can be used to achieve segmentation. It is possible to permit multiple co-existing atlases, for a hierarchical nomenclature or for comparison between them. Experiments with the digital brain phantom presented in chapter 9 show that the automatically segmented structures overlapped the true ones by more than 97% (see section 9.2.2, table 9.2). Comparison of the automatic segmentation algorithm with manual structure identification on real data was completed for the left and right head of caudate in 10 volumetric data sets and resulted in an average overlap of almost 90%, indicating an average misregistration of less than 1mm (see section 9.3.3, Fig. 9.4 and table 9.3). The automatically segmented structures were used to create a probabilistic representation of the caudate. This probabilistic structure overlapped the manually-generated one by 95% (section 9.4.1, Fig. 9.7), indicating that the automatic segmentation procedure can be used to build probabilistic structures given a structure definition on the target model volume and a number of deformation field from individual subjects.

In summary, these results support the contention that segmentation can be defined as registration followed by delineation and that the reversal of the traditional model-based segmentation techniques can lead to a practical and robust mechanism for outlining specific brain regions in 3-D. In this thesis, a method for fully automatic, atlas-independent, model-based segmentation was developed, evaluated and shown to performed well on real data. The examples presented in this thesis have been exploratory and insights acquired by application of the methodology to different structures will determine the true value of the method.

10.3 Caveat emptor

The segmentation method is subject to the following caveats. Even though the goal of the non-linear registration algorithm was to account for as much anatomical variability as possible, it is impossible to eliminate all variability, because of non-equivalent topology in some brain regions. Therefore it is not clear that there can exist a completely objective *correct* answer that can be used to evaluate the algorithm. In particular, while the major cortical convolutions are present in all brains, the secondary and tertiary gyri are much more variable. Hence, at some scale, the notion of absolute correspondence breaks down. This problem may ultimately require a completely different strategy in which the cortical mantle is "unfolded" and inter-subject matching is achieved my maximizing the overlap of these unfolded sheets with a 2-D algorithm. This work has been the subject of the PhD project of another student at the MNI (MacDonald *et al.*, 1994). Eventually, these two methods must be combined into one.

The VBSM was based on the 305-brain MRI-atlas that was derived from the averaging of many T1-weighted data sets. Even though the registration process functions well for T1-weighted data, the procedure may not work for T2 or PD-weighted images when using this same model because the edges derived from other data types will not correspond to the edges found in T1-weighted model data. Nevertheless, experiments in chapter 8 have shown that the non-linear registration procedure is relatively independent of the data on which it works since simulations with T2 and PD-weighted brain phantom yielded results comparable to those obtained on T1-weighted data. Other average models can be created using the strategy described in section 3.4.

Some structures will never be properly segmented by this method. For example, there are connections between the head of the caudate and the nearby putamen that pass through the internal capsule. These connections appear in random locations between the structures. Since their positions are random, they don't adhere to the equivalent topology assumption required for proper segmentation. However, the segmentation procedure can be used to identify the region where these connections are most likely by identifying such a region in the VBSM-VOI.

No attempt has been made to address the complications brought about by pathology within the anatomy. Nevertheless, the region of a tumour could be masked and ignored by the algorithm in order to see how it adapts to the remaining anatomy.

10.4 Applications and suggestions for future work

10.4.1 Applications

Only initial validation has been completed of the segmentation procedure. Other structures must be evaluated to validate the automatic method. As part of the ongoing neuro-psychology projects at the MNI, many neuro-anatomical structures are manually segmented and analyzed. This set of user-defined structures will provide the data-base necessary for testing of the segmentation procedure using the methods described in section 9.3.3.

One of the applications of the linear stereotaxic mapping procedure was to increase the signal-to-noise ratio for functional activation paradigms. The transformation required to map the MRI volume to stereotaxic space was applied to the co-registered PET data (Evans *et al.*, 1992c) and data sets from multiple subjects were then averaged to enhance the small signal associated with cognitive activation. By using the non-linear stereotaxic transformation, the anatomical region thought to give rise to the functional signal can be aligned. Application of the same transformation to the PET data may further increase the activation signal. Once anatomical variability is eliminated, it may be possible to determine the existence of functional variability in the location of focal sites of physiological responses underlying normal cognitive operations.

An international consortium has recently been created for a research project entitled "A probabilistic reference atlas of the human brain" which has been funded by the U.S. Human Brain Map Project initiative, a multi-agency program lead by the National Institute of Mental Health. The specific goals of the project are to develop techniques to assess neuro-anatomical variability and to build a probabilistic atlas of gross neuro-anatomy for the human brain. A major portion of the rationale for the automatic image segmentation component for that project was derived from the work in progress at the MNI and presented here. The probabilistic caudate and the anatomical variability map created in the previous chapter are specific examples of the general strategy to be employed in the ICBM consortium, but at this stage only represent proof of principle. These procedures must be repeated on a large data base of subjects.

10.4.2 Enhancement of the registration methodology

Local deformation constraints

The non-linear registration procedure of chapter 5 uses a single constraint on the amount of deformation allowed at each node. The use of stereotaxic space permits this constraint to be tailored for each point in the VBSM by storing a vector quantity for each voxel that reflects the local anatomical variability. The program should now be modified to take advantage of this data, instead of the single constant value used presently.

Other matching features

Other features should be evaluated for use in the registration procedure. While the invariant features of intensity and gradient magnitude were sufficient to achieve both linear and non-linear registration, other invariant features such as strength of curvature may be used to improve the non-linear registration. The present algorithm used zeroth and first order gradient magnitudes and the incorporation of second and higher-order derivative invariants is straightforward. It is possible that other types of blurring functions may yield interesting results, particularly the class of edge preserving filters based on anisotropic diffusion (Perona and Malik, 1990).

Simulations with the realistic digital brain phantom showed very good results. It would be interesting to examine the behavior of algorithm running only on classified data, instead of on real data. It is possible that with classified data, a hierarchy of edges could be built, so that the non-linear registration begins by fitting the cortical edges only. Once completed, the ventricles could be added and the fitting procedure restarted. Afterwards, the grey-white border could be added in and the fitting procedure restarted once again. It would appear that this type of directed strategy may speed the overall fitting process.

Non-equivalent topology:

While the scope of this thesis was limited to cerebral structures that are found below the cortex, in order to maintain the validity of the equivalent topology assumption, initial experiments have shown that major structures on the cortex (*e.g.*, interhemispheric fissure, central sulcus, and sylvian fissure) can be identified with the segmentation algorithm. A source of anticipated difficulty will arise from other more variable structures, particularly in secondary and tertiary cortical structures, where topological non-equivalence is more the rule than the exception. For example, in a study of the cingulate sulcus for 238 hemispheres from the 305-MRI database at the MNI, Paus *et al.* (1993) have identified two common patterns of either a single cingulate sulcus or two, with a transition zone between limbic cortex and neocortex in the latter pattern. These alternatives can be accommodated by regarding each pattern not as two mutually exclusive models, but as probabilistically-weighted sub-types of the two-sulcus pattern. The segmentation atlas will therefore have to modified to reflect the more complex pattern. The occurrence of a simpler pattern will be reflected in a lower probability value for labelling the additional sulcus. Similar strategies will be employed where multiple folding patterns have been observed (Ono *et al.*, 1990).

Discontinuous spatial deformation function

The recovery of the deformation is based on the assumption that there exists a one-to-one mapping between homologous points in different brains. Since neighbouring points within a particular brain structure of the source data set should be mapped to neighbouring points in the target, the non-linear deformation is constrained to be continuous throughout the domain of the brain. However, there exist neighbouring points in unconnected structures (such as on opposite sides of a sulcus, or on either side of the longitudinal fissure) that do not need to be mapped to neighbouring points in the target. Therefore, it may be desirable to allow discontinuities in the transformation at internal brain surfaces, e.g surfaces that separate the cerebellum from the occipital lobe or that separate the temporal lobe from the inferior frontal lobe. While this may complicate the recovery of the global transformation, defining piece-wise continuous regions where the current algorithm can be applied is certainly possible. However, the main problem will then reside in the inversion of the dis-continuous spatial deformation transformation required for segmentation.

Merging with other non-linear matching procedures

In the Positron Imaging Laboratory at the MNI there are two other on-going projects that address non-linear matching and segmentation of brain structures. The first addresses the problem of cortical mantle extraction and segmentation based on the simultaneous deformation of multiple surfaces from a model combined with curvature matching (MacDonald *et al.*, 1994). The second project deals with the matching of previously extracted sulci from a given data volume with a set of model sulci, using a force-based deformation technique without explicit pre-determined correspondence (Luo and Evans, 1994). The procedure is iterative, matching major sulci first and then adding in smaller sulci as the fit between data and model improves. Since the scope of the thesis was limited to brain structures beneath the cortex, these two algorithms complement the one developed here, and ultimately the three should be merged together.

In particular, the extracted surfaces could be used to delimit piece-wise continuous regions of the brain where the non-linear registration algorithm could be applied. The sulcal matching algorithm can also be applied to vessel structures, and with the fit of the cortical sulci, provide additional constraints for registration.

10.4.3 Enhancement of the segmentation model

There are at least four ways to enhance the VBSM: 1) apply the non-linear registration procedure to the 305 brain data base to enhance MRI atlas; 2) include the 3-D spatial variability field in the model, for use as a local constraint in the non-linear registration procedure; 3) add other modalities such as PET, CT or PD-weighted MRI to the VBSM and 4) redefinition of the present VOI labelling with probabilistic representation of all structures. These enhancements are described below.

Non-linear registration to enhance MRI-atlas

The initial stereotaxic model was created by manually mapping 305 MRI data sets in to the stereotaxic coordinate space using a linear transformation. As described in section 3.4, the automatic stereotaxic mapping procedure was applied to each of these volumes in order to rebuild the average MRI model, removing the initial subjectivity involved in the identification the points used to define the transformation. It was found that some features in the average model were significantly sharper, allowing the differentiation of structures not previously visible, *e.g.* separation of white matter tracts of the optic radiations and in grey-matter nuclei of the brain-stem. Using the same boot-strap paradigm, the non-linear stereotaxic mapping procedure should be applied to all brains in the data base, so as to bring into focus details of the anatomy now blurred by morphometric variability. The method could conceivably be repeated a number of times, with the average getting sharper at each iteration until the residual difference between iterations reaches some stopping criteria. Note that the creation of the atlas in this fashion remains order-independent and that the result would be a higher-resolution MRI-atlas to be used as the basis of a new refined VBSM.

Local stiffness parameters

A second enhancement to the model should be the incorporation of average local spatial variability measures in the form of elasticity or stiffness parameters to address new problems posed by the use of the non-linear warping procedure. Once an average anatomical variability map is created with enough subjects, the tolerance for local deformation, which is reflected in the local spatial variability, could be stored in order to tailor the model on a voxel-by-voxel basis. These parameters should be used as constraints, limiting the amount of deformation permitted at any voxel location during the fitting process.

Other modalities

The linear registration technique can be readily extended to other modalities. Other features can be added to the VBSM-IIF and stored in the VBSM using the same methodology as that described for MRI. For example, the average intensity and gradient values derived from positron emitted tomography (PET) cerebral blood flow (CBF) data can be stored in the model and used to register against other PET CBF studies. This achieves two goals. The first permits cross-registration of PET volume data sets. The second provides a procedure to automatically register data sets from different modalities. Since MRI and PET volumes can be independently registered to the same target space, the transformation between the two volumes is known implicitly. A 243-brain average stereotaxic PET volume has been created within the Positron Imaging Laboratory within the MNI. This data set will be incorporated into the VBSM.

Probabilistic VBSM-VOI

The last improvement involves the VBSM-VOI. The present atlas is based on contours derived from a single individual and is not, in its present form, representative of normal anatomical variability. The composite mean MRI-intensity atlas is useful as a qualitative index of local anatomical variability, but it is insufficient as a quantitative tool. The methodology presented in chapter 9 can be used to replace all structures of the VBSM-VOI with their probabilistic representation.

10.5 Summary

In conclusion, in response to the primary goal restated at the beginning of this chapter, a fully automatic model-based segmentation procedure for the identification and delineation of neuro-anatomical structures has been developed and validated on both simulated and real data. The reversal of the tradition segmentation procedure has resulted in an atlas-independent segmentation scheme that can allow multiple atlases to be used and take advantage of improvement in structure definition without recalculating the spatial transformation required to achieve the segmentation.

In this thesis, a conceptual framework has been constructed, based on the concept of stereotaxic space, that has permitted the creation of an automatic method for comparison of structural anatomy within and between subjects. The tools developed here have been brought to the proof of principle stage for the automated analysis of anatomical variability. Application of these procedures results in objective, reproducible delineation of gross neuro-anatomical structures and their use will extend our knowledge of quantitative anatomical variability.

References

ABSHER, J. (1993), 'A probabilistic atlas of the thalamus', *Work in progress*.

ALLAIN, P., TRAVÈRE, J. and BARON, J. (1992), 'Multimodality: 3D MRI analysis as a primary tool for PET acquisition referenced to an anatomical data base', In *3D Advanced Image Processing in Medicine*, Proceedings of the IEEE, Nov., pp. 201–203.

ALPERT, N., BRADSHAW, J., KENNEDY, D. and CORREIA, J. (1990), 'The principal axis transformation - a method for image registration', *Journal of Nuclear Medicine*, Vol. 31, No. 10, Oct., p. 1717.

AMAMOTO, D.Y., KASTURI, R. and MAMOURAIN, A. (1990), 'Tissue-type discrimination in magnetic resonance images', In *10th International Conference on Pattern Recognition*, IAPR, Atlantic City, NJ, June, pp. 603–607.

ARATA, L., DHAWAN, A.P., BRODERICK, J. and GASKIL, M. (1991), 'Model-based analysis of MR images of the brain', *IEEE Engineering in Medicine and Biology*, Vol. 13, No. 3, pp. 1331–1332.

ARNOLD, D., RIESS, G., MATTHEWS, P., COLLINS, D., FRANCIS, G. and ANTEL, J. (1992), 'Quantitation of disease load and progression in multiple sclerosis by means of proton magnetic resonance spectroscopy', In *SMRM11: Society Magnetic Resonance in Medicine 11th Annual Meeting*, Society Magnetic Resonance in Medicine, Berlin, Germany, Aug., p. 1911.

BAJCSY, R. and BROIT, C. (1982), 'Matching of deformed images', In *Proceedings of the 6th International Conference on Pattern Recognition*, IEEE, Munich, Germany, Oct 19-22, pp. 351–353.

BAJCSY, R. and KOVACIC, S. (1989), 'Multiresolution elastic matching', *Computer Vision, Graphics, and Image Processing*, Vol. 46, pp. 1–21.

BAJCSY, R. and SOLINA, F. (1987), 'Three dimensional object representation revisited', In *Proceedings of the 1st International Conference on Computer Vision*, IEEE, London, June, pp. 231–240.

BAJCSY, R., LIEBERSON, R. and REIVICH, M. (1983), 'A computerised system for the elastic matching of deformed radiographic images to idealized atlas images', *Journal of Computer Assisted Tomography*, Vol. 7, No. 4, pp. 618–625.

BERLANGIERI, S.U., SCHIFTER, T., HOFFMAN, J.M., HAWK, T.C., HAMBLEN, S.M. and COLEMAN, R.E. (1993), 'An interactive, computer-based atlas of neurologic positron emission tomographic studies for use in teaching', *AJR Am J Roentgenol*, Vol. 160, No. 6, Jun, pp. 1295–8.

BEZDEK, J.C., HALL, L.O. and CLARKE, L.P. (1993), 'Review of MR image segmentation techniques using pattern recognition', *Med Phys*, Vol. 20, No. 4, Jul-Aug, pp. 1033–48. Review.

BOHM, C., GREITZ, T., KINGSLEY, D., BERGGREN, B.M. and OLSSON, L. (1983), 'Adjustable computerized stereotaxic brain atlas for transmission and emmision tomography', *American Journal of Neuroradiology*, Vol. 4, May/June, pp. 731–733.

BOHM, C., GREITZ, T., KINGSLEY, D., BERGGREN, B.M. and OLSSON, L. (1985), 'A computerized individually variable stereotaxic brain atlas', In GREITZ, T. (Ed.), *The metabolism of the human brain studied with positron emission tomography*, Raven Press, New York, pp. 85–91.

BOHM, C., GREITZ, T., SEITZ, R. and ERIKSSON, L. (1991), 'Specification and selection of regions of interest (rois) in a computerized brain atlas', *J Cereb Blood Flow Metab*, Vol. 11, No. 2, Mar, pp. A64–8. Review.

BOHM, C., GREITZ, T. and THURFJELL, L. (1992), 'The role of anatomic information in quantifying functional neuroimaging data', *J Neural Transm Suppl*, Vol. 37, No. 37, pp. 67–78. Review.

BOMANS, M., HÖHNE, K.H., TIEDE, U. and RIEMER, R. (1990), '3D segmentation of MR images of the head for 3D display', *IEEE Transactions on Medical Imaging*, Vol. 9, No. 2, june, pp. 177–183.

BONAR, D.C., SCHAPER, K.A., ANDERSON, J.R., ROTTENBERG, D.A. and STROTHER, S.C. (1993), 'Graphical analysis of MR feature space for measurement of CSF, gray-matter and white-matter volumes', *Journal of Computer Assisted Tomography*, Vol. 17, No. 3, pp. 461–470.

BOOKSTEIN, F.L. (1989), 'Principal warps: Thin-plate splines and the decomposition of deformations', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. PAMI-11, No. 6, pp. 567–585.

BRACEWELL, R.W. (1978), *The Fourier Transform and its Applications*, McGraw-Hill, New York.

BROIT, C. (1981), *Optimal registration of deformed images*, PhD thesis, University of Pennsylvania, Philadelphia.

BRUMMER, M.E. (1991), 'Hough transform detection of the longitudinal fissure in tomographic head images', *IEEE Transactions on Medical Imaging*, Vol. 10, No. 1, Mar.

BRUMMER, M.E., MERSEREAU, R.M., EISNER, R.L. and LEWINE, R.R.J. (1991), 'Automatic detection of brain contours in MRI data sets', In COLCHESTER, A.C.F. and HAWKES, D.J. (Eds.), *Information Processing in Medical Imaging*, IPMI, Wye, UK, July, p. 188.

BURR, D.J. (1981a), 'A dynamic model for image registration', *Comput. Graphics and Image Process.*, Vol. 15, February, pp. 102–112.

BURR, D.J. (1981b), 'Elastic matching of line drawings', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. PAMI-3, No. 6, November, pp. 708–713.

CHEN, L.S. and SONTAG, M.R. (1989), 'Representation, display and manipulation of 3D digital scenes and their medical applications', *Computer Vision, Graphics, and Image Processing*, Vol. 48, pp. 190–216.

CHUANG, K.S. and UDUPA, J.K. (1989), 'Boundary detection in grey-level scenes', In *NCGA Conference Proceedings V. I*, NCGA, Philadelphia, April 17-20, pp. 112–117.

CHUI, C. and DAMASIO, A. (1980), 'Human cerebral asymmetries evaluated by computed tomography', *J Neurochem Neurosurg Psychiat*, Vol. 43, pp. 873–878.

COLLINGNON, A., GéRAUD, T., VANDERMEULEN, D., SUETENS, P. and MARCHAL, G. (1993a), 'New high-performance 3D registration algorithms for 3D medical images', In *Proceedings of SPIE Medical Imaging*, SPIE, p. 410.

COLLINGNON, A., VANDERMEULEN, D., SUETENS, P. and MARCHAL, G. (1993b), 'Surface-based registration of medical images', In *Proceedings of SPIE Medical Imaging*, SPIE, p. 385.

COLLINS, D.L. (1990), *MSP - multiple sclerosis analysis package - user guide*, McGill University, Montreal, Canada.

COLLINS, D.L., PETERS, T.M. and EVANS, A.C. (1992a), 'Automatic multiresolution registration of MRI volumetric data', In *SMRM11: Society Magnetic Resonance in Medicine 11th Annual Meeting*, Society Magnetic Resonance in Medicine, Berlin, Germany, Aug., p. 4214.

COLLINS, D., TAKAHASHI, A. and ARNOLD, D. (1992b), 'Determination of metabolic lesion size and severity on proton MR spectroscopic images using semi-automatic classification and mean distances in feature space', In *SMRM11: Society Magnetic Resonance in Medicine 11th Annual Meeting*, Society Magnetic Resonance in Medicine, Berlin, Germany, Aug., p. 4271.

COLLINS, D.L., PETERS, T.M., DAI, W. and EVANS, A.C. (1992c), 'Model based segmentation of individual brain structures from MRI data', In *Proceedings of Conference on Visualization in Biomedical Computing*, SPIE press, Oct., pp. 10–23.

COLLINS, D.L., PETERS, T.M. and EVANS, A.C. (1992d), 'Multiresolution image registration and brain structure segmentation', In *3D Advanced Image Processing in Medicine*, Proceedings of the IEEE, Nov., pp. 105–109.

COLLINS, D.L., PETERS, T.M. and EVANS, A.C. (1994a), 'An automated 3D non-linear image deformation procedure for determination of gross morphometric variability in human brain', In *Proceedings of Conference on Visualization in Biomedical Computing*, SPIE.

COLLINS, D.L., NEELIN, P., PETERS, T.M. and EVANS, A.C. (1994b), 'Automatic 3D inter-subject registration of MR volumetric data in standardized talairach space', *Journal of Computer Assisted Tomography*, Vol. 18, No. 2, March/April, pp. 192–205.

DAI, W., MARRETT, T.S., COLLINS, D.L. and EVANS, A.C. (1991a), 'Three-dimensional non-linear warping of a computerized volume-of-interest brain atlas for morphometric analysis', In *XVth International Symposium on Cerebral Blood Flow and Metabolism*, Miami, June.

DAI, W., MARRETT, D.L.C.T.S. and EVANS, A.C. (1991b), 'Three-dimensional warping of a computerized brain atlas', In *Annual Symposium of the IEEE Engineering in Medicine and Biology Society*, Orlando, June.

DALTON, B.L. and DU BOULAY, G. (1988), 'Medical image matching', In *Proceedings of SPIE Medical Imaging*, SPIE, pp. 456–464.

DANN, R., HOFORD, J., KOVACIC, S., REIVICH, M. and BAJCSY, R. (1988), 'Three-dimensional computerized brain atlas for elastic matching: Creation and initial evaluation', In *Medical Imaging II*, SPIE, Newport Beach, Calif., February, pp. 600–608.

DAVIS, D.N. and TAYLOR, C.J. (1991), 'A blackboard architecture for automating cephalometric analysis', *Med Inf (Lond)*, Vol. 16, No. 2, Apr-Jun, pp. 137–49.

DELAPAZ, R., BERNSTEIN, R., JITENDRA, D. and CHANG, P. (1985), 'Tissue characterization with MRI using advanced digital processing techniques ('fuzzy clustering')', In *Proceedings of the fifth Annual Meeting of SMRM*, Society of Magnetic Resonace in Medicine, pp. 229–230.

DELLEPIANE, S., SERPICO, S. and VERNAZZA, G. (1986), 'Approximate reasoning and knowledge in NMR image understanding', In *8th International Conference on Pattern Recognition*, IEEE, Paris, France, October, pp. 943–946.

DELLEPIANE, S., SERPICO, S., VERNAZZA, G., BRUZZONE, S. and REGAZZONI, C. (1987), 'An expert system for the recognition of anatomical organs from 3D data base', In *IEEE Ninth annual conference of the Enginnering in Medicine and Biology Society*, pp. 933–934.

DELLEPIANE, S., G. VENTURI and VERNAZZA, G. (1991), 'Segmentation and recognition of MR pathological images', In *Proceedings of Computer Assisted Radiology*, pp. 600–605.

DHAWAN, A.P. and ARATA, L. (1992), 'Knowledge-based multi-modality three-dimensional image analysis of the brain', *American Journal of Physiologic Imaging*, Vol. 7, No. 3-4, Jul-Dec, pp. 210–9.

DHAWAN, A.P. and JUVVADI, S. (1988), 'Knowledge-based analysis and understanding of 3D medical imaging', In *Proceedings of SPIE Medical Imaging*, SPIE, Newport Beach, Calif., February, pp. 422–428.

DHAWAN, A.P. and JUVVADI, S. (1990), 'Knowledge-based analysis and understanding of medical images', *Computer Methods and Programs in Biomedicine*, Vol. 33, pp. 221–239.

DRANGOVA, M. (1987), *Stereotactic neurosurgical planning using magnetic resonance imaging: Image distortion evaluation*, Master's thesis, Medical Physics Unit, McGill University, Montreal Neurological Institute, September.

DREVITS, W., VIDEEN, T., MACLEOD, A., HALLER, J. and RAICHLE, M. (1992), 'PET images of blood flow changes during anxiety: Correction', *Science*, Vol. 256, p. 1696.

DUCHON, J. (1976), 'Interpolation des fonctions de deux variables suivant le principe de la flexion des plaques minces', *R.A.I.R.O., Analyze numérique*, Vol. 10, pp. 5–12.

DUDA, R. and HART, P. (1973), *Pattern Recognition And Scene Analysis*, Wiley, New York.

EVANS, A.C. (1993), 'Multi-modality correlation and analysis', In WAGNER, H. (Ed.), *Principles of Nuclear Medicine*, ??? in press.

EVANS, A.C., BEIL, C., MARRETT, S., THOMPSON, C.J. and HAKIM, A. (1988), 'Anatomical functional correlation using an adjustable MRI based atlas with PET', *Journal of Cerebral Blood Flow and Metabolism*, Vol. 8, No. 4, pp. 813–830.

EVANS, A., MARRETT, S., COLLINS, L. and PETERS, T. (1989a), 'Anatomical-functional correlative analysis of the human brain using three dimensional imaging systems', In *Proceedings of SPIE Medical Imaging*, SPIE, Newport Beach, California, Jan. 29 - Feb. 3, pp. 264–274.

EVANS, A.C., MARRETT, S., COLLINS, L. and PETERS, T.M. (1989b), 'Anatomical-functional correlative imaging in three dimensions', Abstract, May. Presented at: International Symposium on Cerebral Blood Flow and Metabolism, Bologna, Italy.

EVANS, A.C., MARRETT, S., COLLINS, L. and PETERS, T.M. (1989c), 'A customizable volume-of-interest (VOI) atlas for three dimensional PET data analysis', Society of Nuclear Medicine Symposium, St. Louis, June.

EVANS, A.C., BEIL, C. and MARRETT, S. (1989d), 'MRI-PET correlation using an adjustable ROI atlas', *Journal of Nuclear Medicine*, Vol. 28, No. 6, p. 1071.

EVANS, A.C., PETERS, T.M., COLLINS, D.L., HENRI, C.J., MARRETT, T.S., PIKE, G.B. and DAI, W. (1990a), '3-D correlative imaging and segmentation of cerebral anatomy, function and vasculature', In *12th Annual International Conference of IEEE Engineering in Medicine and Biology Society*, IEEE, Philadelphia, June.

EVANS, A.C., PETERS, T.M., COLLINS, D.L., HENRI, C.J., MARRETT, T., PIKE, G.B. and DAI, W. (1990b), '3-D correlative imaging and segmentation of cerebral anatomy, function and angiography', *IEEE Trans Eng Med Biol*, Vol. 12, No. 3, pp. 1297–98.

EVANS, A.C., MARRETT, S., TORRESCORZO, J., KU, S. and COLLINS, L. (1991a), 'MRI-PET correlation in three dimensions using a volume-of-interest (VOI) atlas', *Journal of Cerebral Blood Flow and Metabolism*, Vol. 11, No. 2, Mar, pp. A69–78.

EVANS, A.C., DAI, W., COLLINS, D.L., NEELIN, P. and MARRETT, T. (1991b), 'Warping of a computerized 3-D atlas to match brain image volumes for quantitative neuroanatomical and functional analysis', In *Proceedings of the International Society of Optical Engineering: Medical Imaging V*, SPIE, San Jose, California, 27 February – 1 March.

EVANS, A.C., COLLINS, D.L. and MILNER, B. (1992a), 'An MRI-based stereotactic atlas from 250 young normal subjects', *Soc.Neurosci.Abstr.*, Vol. 18, p. 408.

EVANS, A.C., PETERS, T.M., COLLINS, D.L., HENRI, C.J., MARRETT, T., PIKE, G.B. and DAI, W. (1992b), '3-D correlative imaging and segmentation of cerebral anatomy, function and vasculature', *Automedica*, Vol. 14, No. 1, pp. 65–80.

EVANS, A.C., MARRETT, S., NEELIN, P., COLLINS, D.L., WORSLEY, K., DAI, W., MILOT, S., MEYER, E. and BUB, D. (1992c), 'Anatomical mapping of functional activation in stereotactic coordinate space', *NeuroImage*, Vol. 1, No. 1, pp. 43–53.

EVANS, A.C., COLLINS, D.L., MILLS, S.R., BROWN, E.D., KELLY, R.L. and PETERS, T.M. (1993a), '3d statistical neuroanatomical models from 305 mri volumes', In *Proc. IEEE-Nuclear Science Symposium and Medical Imaging Conference*, pp. 1813–1817.

EVANS, A.C., PETERS, T.M., COLLINS, D.L., NEELIN, P. and GABE, C. (1993b), 'Image registration based on discrete anatomical structures', In MACIUNAS, R. (Ed.), *Interactive image-guided neurosurgery*, American Association of Neurological Surgeons, Chap. 5, pp. 63–80.

FARBER, T.L. and STOKELY, E.M. (1988), 'Orientation of 3D structures in medical images', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 10, No. 5, September, pp. 626–633.

FEEHS, R.J. and ARCE, G.R. (1987), 'Multidimensional morphological edge detection', In *SPIE: Visual Communications and Image Processing II*, SPIE, pp. 285–292.

FISCHLER, M. and ELSCHLAGER, R. (1973), 'The representation and matching of pictorial structures', *IEEE Transactions on Computers*, Vol. c-22, No. 1, January, pp. 67–92.

FOX, P.T., PERLMUTTER, J.S. and RAICHLE, M.E. (1985), 'A stereotactic method of anatomical localization for positron emission tomography', *Journal of Computer Assisted Tomography*, Vol. 9, No. 1, pp. 141–153.

FOX, P.T., MINTUN, M.A., REIMAN, E.M. and RAICHLE, M.E. (1988), 'Enhanced detection of focal brain responses using intersubject averaging and change-distribution analysis of subtracted PET images', *Journal of Cerebral Blood Flow and Metabolism*, Vol. 8, pp. 642–653.

FOX, P.T., MIKITEN, S., DAVIS, G. and LANCASTER, J.L. (1994), 'BrainMap: A database of functional brain mapping', In THATCHER, R.W., HALLETT, M., ZEFFIRO, T., JOHN, E.R. and HEURTA, M. (Eds.), *Functional Neuroimaging, technical foundations*, Academic Press, San Diego, Ca., pp. 95–105.

FRISTON, K., PASSINGHAM, R. and NUTT, J. (1989), 'Localization in PET images: Direct fitting of the intercommissural (AC-PC) line', *Journal of Cerebral Blood Flow and Metabolism*, Vol. 9, No. 5, pp. 690–695.

FRISTON, K., FRITH, C., LIDDLE, P. and FRACKOWIAK, R. (1991), 'Plastic transformation of PET images', *Journal of Computer Assisted Tomography*, Vol. 15, No. 1, pp. 634–639.

GALABURDA, A., LEMAY, M., KEMPER, T. and GESCHWIND, N. (1978), 'Left-right asymmetries in the brain', *Science*, Vol. 199, pp. 852–6.

GEE, J.C., REIVICH, M. and BAJCSY, R. (1993), 'Elastically deforming 3D atlas to match anatomical brain images', *J Comput Assist Tomogr*, Vol. 17, No. 2, Mar-Apr, pp. 225–36.

GERIG, G., KIKINIS, R. and JOLESZ, F. (1990), 'Image processing of routine spin-echo MR images to enhance vascular structures: Comparison with MR-angiography', In HÖHNE, K., FUCHS, H. and PIZER, S. (Eds.), *3D Imaging in Medicine: Algorithms, Systems, Applications*, NATO ASI Series, Travemünde, Germany, June 25-29, pp. 121–132.

GERIG, G., MARTIN, J., KIKINIS, R. and KUBLER, O. (1991a), 'Segmentation of dual-echo MR head data', In *Proceedings of Computer Assisted Radiology*, pp. 606–611.

GERIG, G., KIKINIS, R., KUONI, W., VON SCHULTHESS, G.K. and KUBLER, O. (1991b), 'Semiautomated ROI analysis in dynamic MR studies. part 1: Image analysis tools for automatic correction of organ displacements', *Journal of Computer Assisted Tomography*, Vol. 15, No. 5, pp. 725–732.

GESCHWIND, N. and LEVITSKY, W. (1968), 'Human brain: left-right asymmetries in temporal speech regions', *Science*, Vol. 161, pp. 181–187.

GOHAGAN, J.K., SPITZNAGEL, E.L., MURPHY, W.A. and ET AL., M.A.V. (1987), 'Multispectral analysis of MR images of the breast', *Radiology*, Vol. 163, pp. 703–707.

GOLUB, G. and VAN LOAN, C. (1983), *Matrix Computations*, Johns Hopkins University Press, Chap. 12.

GREITZ, T., BOHM, C., HOLTE, S. and ERIKSSON, L. (1991a), 'A computerized brain atlas: construction, anatomical content, and some applications', *J Comput Assist Tomogr*, Vol. 15, No. 1, Jan-Feb, pp. 26–38.

GREITZ, T., HOLTE, S., BOHM, C., ERIK'¡SSON, L., SEITZ, R., ERICSON, K., NYBÄCK and STONE-ELANDER, S. (1991b), 'A data base library as a diagnostic aid in neuroimaging', *Neuroradiology*, Vol. 33, No. suppl.

GUEZIEC, A. and AYACHE, N. (1992a), 'Smoothing and matching of 3D space curves', In *Proceedings of Conference on Visualization in Biomedical Computing*, SPIE press, p. 259.

GUEZIEC, A. and AYACHE, N. (1992b), 'Smoothing and matching 3D space curves', In *3D Advanced Image Processing in Medicine*, 14th annual proceedings of the IEEE EMBS, Nov., pp. 161–165.

HANAWAY, J., SCOTT, W. and STROTHER, C. (1980), *Atlas of the human brain and orbit*, WArren H. Green, Inc., St.-Louis.

HENRI, C., COLLINS, L., PETERS, T., EVANS, A. and MARRETT, S. (1989), 'Three-dimensional interactive display of medical images for stereotactic neurosurgery planning', In *Proceedings of SPIE Medical Imaging*, SPIE, Newport Beach, California, Jan. 31 - Feb. 3, pp. 67–74.

HENRI, C.J., PIKE, G.B., COLLINS, D.L. and PETERS, T.M. (1990), 'Three-dimensional display of cortical anatomy and vasculature: MR angiography versus multi-modality integration', In *Proceedings of SPIE Medical Imaging*, SPIE, Feb. 4-5, pp. 172–182.

HENRI, C.J., COLLINS, D.L. and PETERS, T.M. (1991a), 'Reconstruction of 3-D branching structures', In COLCHESTER, A.C.F. and HAWKES, D.J. (Eds.), *Lecture Notes in Computer Science*, Springer, Heidelberg.

HENRI, C., COLLINS, L. and PETERS, T. (1991b), 'Multi-modality image integration for stereotactic surgical planning', *Medical Physics*, Vol. 18, No. 2, Mar/Apr, pp. 167–177.

HENRI, C.J., PIKE, G.B., COLLINS, D.L. and PETERS, T.M. (1991c), 'Three-dimensional display of cortical anatomy and vasculature: magnetic resonance angiography versus multimodality integration', *Journal of Digital Imaging*, Vol. 4, No. 1, pp. 21–27.

HENRI, C.J., COLLINS, D.L. and PETERS, T.M. (1993), 'Optimization of projection geometry for few-view reconstruction of sparse objects', *Medical Physics*, Vol. 20, No. 5, pp. 1537–1547.

HERRMANN, A., LEVIN, D.N. and HU, X. (1988), 'Segmentation and 3D display of MR images of brain lesions', In *Proceedings of the Seventh Annual Meeting of SMRM*, Society of Magnetic Resonace in Medicine, August 20-26, p. 199.

HILL, D., HAWKES, D., CROSSMAN, J., GLEESON, M. and COX, T. (1991), 'Registration of MR and CT images for skull base surgery using point-like anatomical features', *British Journal of Radiology*, Vol. 64, No. 767, pp. 1030–1035.

HÖHNE, K.H., BOMANS, M., TIEDE, U. and RIEMER, M. (1988), 'Display of multiple 3D-objects using the generalized voxel-model', In *Proceedings of SPIE Medical Imaging*, SPIE, pp. 850–854.

HUECKEL, M. (1973), 'A local visual operator which recognizes edges and lines', *Journal of the Association of Computing Machines*, Vol. 18, No. 1, pp. 113–125.

JIANG, H., ROBB, R. and HOLTON, K. (1992), 'A new approach to 3D registration of multimodality medical images by surface matching', In *Proceedings of Conference on Visualization in Biomedical Computing*, SPIE press, p. 196.

JUNCK, L., MOEN, J.G., HUTCHINS, G.D., BROWN, M.B. and KUHL, D.E. (1990), 'Correlation methods for the centering, rotation and alignment of functional brain images', *Journal of Nuclear Medicine*, Vol. 31, pp. 1220–1276.

KALL, B.A., KELLY, P.J. and GOERSS, A.J. (1987), 'Comprehensive computer assisted data collection, treatment planning and interactive surgery', In *Proceedings of SPIE Medical Imaging*, SPIE, pp. 509–514.

KALVIN, A. and PELEG, S. (1989), 'A 3D multi-resolution segmentation algorithm for surface reconstruction from CT data', In *Medical Imaging III*, SPIE, Newport Beach, Calif., February, pp. 173–182.

KAMBER, M., COLLINS, D.L., SHINGHAL, R., FRANCIS, G.S. and EVANS, A.C. (1992), 'Model-based 3D segmentation of multiple sclerosis lesions in dual-echo MRI data', In *Proceedings of Conference on Visualization in Biomedical Computing*, SPIE, Chapel Hill, North Carolina, 13-16 October, pp. 590–600.

KAMBER¡, M., SHINGHAL, R., COLLINS, D., FRANCIS, G. and EVANS, A. (1994), 'Model-based 3D segmentation of multiple sclerosis lesions from magnetic resonance brain images', *IEEE Transactions on Medical Imaging*, p. (accepted for publication).

KANEDA, Y., FUJII, S., KOBASHIRI, Y., YOSHIRDA, M. and MATSUO, M. (1978), 'Pattern recognition and three dimensional construction from CT images', In *Proceedings of the International Conference on Cybernetics and Society*, IEEE, Tokyo and Kyoto, Japan, Nov 3-7, pp. 281–284.

KAPOULEAS, I. (1990), 'Segmentation and feature extraction from magnetic resonance brain image analysis', In *10th International Conference on Pattern Recognition*, IAPR, Atlantic City, NJ, June, pp. 583–590.

KAPOULEAS, I. and KULIKOWSKI, C.A. (1988), 'A model-based system for the interpretation of MR human brain scans', In *Medical Imaging II*, SPIE, Newport Beach, Calif., February, pp. 429–437.

KASS, M., WITKIN, A. and TERZOPOULOS, D. (1988), 'Snakes: active contour models', In *First International Conference on Computer Vision*, IEEE, London, England, pp. 259–268.

KAZARNOVSKAYA, M.I., BORODKIN, S.M., SHABALOV, V.A., KRIVOSHEINA, V.Y. and GOLANOV, A.V. (1991), '3-D computer model of subcortical structures of human brain', *Comput Biol Med*, Vol. 21, No. 6, pp. 451–7.

KELLY, P., KALL, B. and GOERSS, S. (1983), 'Stereotactic CT scanning for the biopsy of intracranial lesions and functional neurosurgery', *Applied Neurophysiology*, Vol. 46, pp. 193–199.

KENNEDY, D.N., FILIPEK, P.A. and VERNE S. CAVINESS, J. (1989), 'Anatomic segmentation and volume calculations in nuclear magnetic resonance imaging', *IEEE Transactions on Medical Imaging*, Vol. 8, No. 1, March, pp. 1–7.

KOENDERINK, J. and VAN DOORN, A. (1987), 'Representation of local geometry in the visual system', *Biol. Cybern.*, Vol. 55, pp. 367–375.

KOSUGI, Y., SASE, M., KUWATANI, H., KINOSHITA, N., MOMOSE, T., NISHIKAWA, J. and WATANABE, T. (1993), 'Neural network mapping for non-linear stereotactic normalization of brain MR images', *Journal of Computer Assisted Tomography*, Vol. 17, No. 3, pp. 455–460.

KRETCHMANN, H.J. and WEINRICH, W. (1986), *Neuroanatomy and cranial computed tomography*, Georg Thieme Verlag, Stuttgart.

KÜBLER, O. and GERIG, G. (1990), 'Segmentation and analysis of multidimensional data-sets in medicine', In HÖHNE, K., FUCHS, H. and PIZER, S. (Eds.), *3D Imaging in Medicine: Algorithms, Systems, Applications*, NATO ASI Series, Travemünde, Germany, June 25-29, pp. 63–82.

LANGRON, S. and COLLINS, A. (1985), 'Perturbation theory for generalized procrustes analysis', *J. R. Statist. Soc.*, Vol. 47, No. 2, pp. 277–284.

LE MAY, M. and KIDO, D. (1978), 'Asymmetries of the cerebral hemispheres on computed tomograms', *Journal of Computer Assisted Tomography*, Vol. 2, pp. 471–476.

LECLERC, V. and BENCHIMOL, C. (1991), 'Automatic elastic registration of DSA images', In *Proceedings of Computer Assisted Radiology*, pp. 719–723.

LEHMANN, E.D., HAWKES, D.J., HILL, D.L., BIRD, C.F., ROBINSON, G.P., COLCHESTER, A.C. and MAISEY, M.N. (1991), 'Computer-aided interpretation of SPECT images of the brain using an MRI-derived 3D neuro-anatomical atlas', *Med Inf (Lond)*, Vol. 16, No. 2, Apr-Jun, pp. 151–66.

LEMOINE, D., BARILLOT, C., GIBAUD, B. and PASQUALINI, E. (1991a), 'A 3D C^1 stereotactic deformation model to merge multimodality images and atlas data', In *Proceedings of Computer Assisted Radiology*, pp. 665–668.

LEMOINE, D., BARILLOT, C., GIBAUD, B. and PASQUALINI, E. (1991b), 'An anatomical-based 3D registration of multimodality and atlas data in surgery', In COLCHESTER, A. and HAWKES, D. (Eds.), *Information Processing in Medical Imaging*, IPMI, Wye, UK, July, p. 188.

LEVIN, D.N., PELIZZARI, C.A., CHEN, G.T., CHEN, C.T. and COOPER, M.D. (1988), 'Retrospective geometric correlation of MR, CT and PET images', *Radiology*, Vol. 169, pp. 817–823.

LEVINE, M. (1985), *Vision in Man and Machine*, McGraw Hill, New York.

LLACER, J., TER HAAR ROMENY, B.M., FLORAK, L.M.J. and VIERGEVER, M. (1993), 'Representation of medical images by visual response functions', *IEEE Engineering in Medicine and Biology*, Vol. 12, No. 1, Mar., pp. 40–47.

LUO, S. and EVANS, A. (1994), 'Matching sulci in 3D space using force-based deformations', *In preparation.*

MACDONALD, D., AVIS, D. and EVANS, A.C. (1994), 'Multiple surface identification and matching in magnetic resonance images', In *Proceedings of Conference on Visualization in Biomedical Computing*, SPIE.

MARRETT, S., EVANS, A., COLLINS, L. and PETERS, T. (1989a), 'Volume of interest (VOI) atlas for the analysis of neurophysiological image data', In *Proceedings of SPIE Medical Imaging*, SPIE, Newport Beach, California, Jan. 29 - Feb. 3, pp. 467–477.

MARRETT, T., COLLINS, L., PIKE, G., KU, S., PETERS, T. and EVANS, A. (1989b), 'Volumetric MRI/PET image correlation display and analysis', In *S.M.R.M. Eighth Annual Meeting*, Society of Magnetic Resonance in Medicine, August.

MATSUI, T. and HIRANO, A. (1978), *An atlas of the human brain for computerized tomography*, Igako-Shoin, Tokyo.

MENHARDT, W. (1988a), *Image analysis using fuzzy sets*, PhD thesis, Hamburg University, Hamburg.

MENHARDT, W. (1988b), 'Iconic fuzzy sets for MR images segmentation', In *Medical Imaging*, NATO/ASI, Pavoa de Varzim, September.

MENHARDT, W. and IMME, M. (1988), 'Contour detection using fuzzy sets', In *Signal Processing IV: Theories and Applications*, EURASIP, pp. 1633–1636.

METAXAS, D. and TERZOPOULOS, D. (1992), 'Dynamic deformation of solid primitives with constraints', In CATMULL, E.E. (Ed.), *Computer Graphics (SIGGRAPH '92 Proceedings)*, July, pp. 309–312.

MINOSHIMA, S. (1993), 'Automated detection of the intercommissural (AC-PC) line for stereotactic localization of functional brain images.', *Journal of Nuclear Medicine*, Vol. 34, No. 2, pp. 322–329.

MINOSHIMA, S., BERGER, K., LEE, K. and MINTUN, M. (1992), 'An automated method for rotation correction and centering of 3D functional brain images', *Journal of Nuclear Medicine*, Vol. 33, No. 8, pp. 1579–1585.

MINTUN, M.A. and LEE, K.S. (1990), 'Mathematical realignment of paired PET images to enable pixel-by-pixel subtraction', *Journal of Nuclear Medicine*, Vol. 31, p. 816.

MONGA, O. and DERICHE, R. (1989), '3D edge detection using recursive filtering: applications to scanner images', In *International Conference on Computer Vision and Pattern Recognition*, IEEE, San Diego, Calif., January, pp. 28–35.

NASTAR, C. and AYACHE, N. (1993), 'Non-rigid motion analysis in medical images: a physically based approach', In *Proceedings of Information Processing in Medical Imaging*, Flagstaff, Arizona, June, pp. 17–32.

NIEW, H.M., CHEN, C.T., LIN, W.C. and PELIZZARI, C.A. (1991), 'Automated three-dimensional registration of medical images', In *Proceedings of SPIE Medical Imaging*, SPIE, pp. 259–264.

ONO, M., KUBIK, S. and ABERNATHEY, C. (1990), *Atlas of Cerebral Sulci*, Georg Thieme Verlag, Stuttgart.

PANDA, D.P. and ROSENFELD, A. (1978), 'Image segmentation by pixel value in (gray level, edge value) space', *IEEE Transactions on Computers*, Vol. C-27, No. 9, pp. 875–879.

PAUS, T., PETRIDES, M., EVANS, A.C. and MEYER, E. (1993), 'Role of the human anterior cingulate cortex in the control of oculomotor, manual, and speech responses: a positron emission tomography study', *Journal of Neurophysiology*, Vol. 70, No. 2, Aug, pp. 453–69.

PAUS, T., TOMAIUOLO, F., OTAKI, N., MACDONALD, D., EVANS, A., PETRIDES, M., ATLAS, J. and MORRIS, R. (1994), 'Human cingulate and paracingulate sulci: pattern, variability, asymmettry, and probabilistic map', *Human Cortex*, p. submitted.

PELI, E., AUGLIERE, R.A. and TIMBERLAKE, G.T. (1997), 'Feature-based registration of retinal images', *IEEE Transactions on Medical Imaging*, Vol. 6, No. 3, Sep, pp. 272–278.

PELIZZARI, C.A. and CHEN, G.T.Y. (1987), 'Registration of multiple diagnostic imaging scans using surface fitting', In BRUINVAS ET AL, I. (Ed.), *The Use of Computers in Radiation Therapy*, Elsevier Science Publishers B.V., North-Holland.

PELIZZARI, C.A., CHEN, G.T.Y., SPELBRING, D.R., WEICHSELBAUM, R.R. and CHEN, C.T. (1989), 'Accurate three-dimensional registration of CT, PET and MRI images of the brain', *Journal of Computer Assisted Tomography*, Vol. 13, No. 1, pp. 20–26.

PELIZZARI, C.A., LEVIN, D.N., CHEN, G.T.Y. and CHEN, C.T. (1993), 'Image registration based on anatomical surface matching', In MACIUNAS, R. (Ed.), *Interactive image-guided neurosurgery*, American Association of Neurological Surgeons, Chap. 4, pp. 47–62.

PERONA, P. and MALIK, J. (1990), 'Scale-space and edge detection using anisotropic diffusion', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 12, No. 7, July, pp. 629–639.

PETERS, T.M., DRANGOVA, M., CLARK, J.A., PIKE, G.B. and OLIVIER, A. (1988), 'Image distortion in mr imaging for stereotactic surgery', In *SMRM7: Society Magnetic Resonance in Medicine 7th Annual Meeting*, Society Magnetic Resonance in Medicine, San Francisco, Ca, August 22-26.

PETERS, T.M., CLARK, J.A., PIKE, G.B., HENRI, C., COLLINS, L., LEKSELL, D. and JEPPSSON, O. (1989), 'Stereotactic neurosurgery planning on a personal-computer-based work station', *Journal of Digital Imaging*, Vol. 2, No. 2, pp. 75–81.

PETERS, T.M., HENRI, C.J., PIKE, G.B., CLARK, J.A., COLLINS, D.L. and OLIVIER, A. (1990), 'Integration of stereoscopic DSA with 3D image reconstruction for stereotactic planning', *Stereotactic and Functional Neurosurgery*, Vol. 55, pp. 471–476.

PETERS, T.M., HENRI, C.J., COLLINS, D.L., LEMIEUX, L., PIKE, G.B. and OLIVIER, A. (1992), 'Stereotactic neurosurgery planning using integrated three-dimensional stereoscopic images', In KELLY, P.J. (Ed.), *Computers in Stereotactic Neurosurgery*, Blackwell, Cambridge Mass., Chap. 19, pp. 259–270. book chapter.

PHELPS, M., MAZZIOTTA, J. and SCHELBERT, H. (1986), *Positron Emission Tomography and Autoradiography*, Raven Press, New York.

PIZER, S.M., GAUCH, J.M. and LIFSHITZ, L.M. (1988), 'Interactive 2D and 3D object definition in medical images based on multiresolution image descriptions', In *Proceedings of SPIE Medical Imaging*, SPIE, pp. 438–444.

PIZER, S.M., FUCHS, H., LEVOY, M., ROSENMAN, J., DAVIS, R.E. and RENNER, J.B. (1989), '3D display with minimal predefinition', In *Proceedings of Computer Assisted Radiology*, pp. 723–736.

PIZER, S., CULIP, T. and FREDERICKSON, R. (1990), 'Toward interactive object definition in 3D scalar images', In HÖHNE, K., FUCHS, H. and PIZER, S. (Eds.), *3D Imaging in Medicine: Algorithms, Systems, Applications*, NATO ASI Series, Travemünde, Germany, June 25-29, pp. 83–106.

PRATT, W.K. (1978), *Digital Image Processing*, John Wiley and Sons, New York.

PRESS, W.H., FALNNERY, B.P., TEUKOLSKY, S.A. and VETTERLING, W.T. (1988), *Numerical recipes in C*, Cambridge University Press, Cambridge.

PREUL, M., COLLINS, D., FEINDEL, W. and ARNOLD, D. (1992), 'Metabolic characterization and discrimination of human intracranial tumours in vivo using proton spectroscopic imaging and classification in feature space', In *SMRM11: Society Magnetic Resonance in Medicine 11th Annual Meeting*, Society Magnetic Resonance in Medicine, Berlin, Germany, Aug., p. 1953.

PULITI, P. and TASCINI, G. (1989), 'A grey-level image segmentation method', In *Medical Imaging III*, SPIE, Newport Beach, Calif., February, pp. 633–639.

RAYA, S.P. (1989), 'Rule-based segmentation of multi-dimensional medical images', In *Conference Proceedings Volume I*, NCGA, Philadelphia, April 17-20, pp. 193–198.

RAYA, S.P. (1990), 'Low-level segmentation of 3D magnetic resonance brain images - a rule based system', *IEEE Transactions on Medical Imaging*, Vol. 9, No. 3, September, pp. 327–337.

RIZZO, G., PASQUALI, P. and GILARDI, M. (1991), 'Multimodal biomedical image integration: Use of a cross-correlation technique', *ieeeemb*, Vol. 13, No. 1, pp. 219–220.

ROBERTS, L. (1965), 'Machine perception of three-dimensional solids', In TIPPETT, J. (Ed.), *Optical and Electro-Optical Information Processing*, MIT Press, Cambridge, Mass., pp. 159–197.

RUBENS, A., U.W., M. and HUTTON, J. (1976), 'Asymmetry of the lateral (sylvian) fissures in man', *Neurology*, Vol. 26, pp. 620–624.

SCHAD, L.R., BOESECKE, R., SCHLEGEL, W., HARTMANN, G.H., STURM, V., STRAUSS, L.G. and LORENZ, W.J. (1987), 'Three dimensional image correlation of CT, MR, and PET studies in radiotherapy treatment planning of brain tumors', *Journal of Computer Assisted Tomography*, Vol. 11, No. 6, pp. 948–954.

SCHIERS, C., TIEDE, U. and HöHNE, K. (1989), 'Interactive 3D registraton of image volumes from different sources', In *Proceedings of Computer Assisted Radiology*, pp. 666–670.

SCHIFTER, T., TURKINGTON, T.G., BERLANGIERI, S.U., HOFFMAN, J.M., MACFALL, J.R., PELIZZARI, C.A., TIEN, R.D. and COLEMAN, R.E. (1993), 'Normal brain F-18 FDG-PET and MRI anatomy', *Clin Nucl Med*, Vol. 18, No. 7, Jul, pp. 578–82.

SEITZ, R.J., BOHM, C., GREITZ, T., ROLAND, P.E., ERIKSSON, L., BLOMQVIST, G., ROSENQVIST, G. and NORDELL, B. (1990), 'Accuracy and precision of the computerized brain atlas programme for localization and quantification in positron emission tomography', *J Cereb Blood Flow Metab*, Vol. 10, No. 4, Jul, pp. 443–57.

SERPICO, S., STURARO, F., VERNAZZA, G. and DELLEPIANE, S. (1987), 'Fuzzy-reasoning approach to understanding of 2D NMR images', In *IEEE Ninth annual conference of the Enginnering in Medicine and Biology Society*, pp. 1980–1981.

SIBSON, R. (1978), 'Studies in the robustness of multidimensional scaling: Procrustes statistics', *J. R. Statist. Soc.*, Vol. 40, No. 2, pp. 234–238.

SIBSON, R. (1979), 'Studies in the robustness of multidimensional scaling: Perturbation analysis of classical scaling', *J. R. Statist. Soc.*, Vol. 41, No. 2, pp. 217–229.

SOKOLOWSKA, E. and NEWELL, J. (1986), 'Multi-layered image representation: Struture and application in recogition of parts of brain anatomy', *Pattern Recognition Letters*, Vol. 4, September, pp. 223–230.

SONTAG, M., CHENG, E., REYNOLDS, R., TALTON, D. and WALLACE, R.E. WAXLER, G. (1986), '3D graphics in radiation-therapy treatment planning', *Medical Physics*, Vol. 13, No. 4, pp. 574–574.

SORLIÉ, C., COLLINS, D.L., WORSLEY, K.J. and EVANS, A.C. (1994), 'An anatomical variability study based on landmarks', *Human brain mapping*, p. submitted.

STEINMETZ, H. and SEITZ, R. (1991), 'Functional anatomy of language processing: Neuroimaging and the problem of individual variability', *Neuropsychologia*, Vol. 29, No. 12, pp. 1149–1161.

TALAIRACH, J. and TOURNOUX, P. (1988), *Co-planar stereotactic atlas of the human brain: 3-Dimensional proportional system: an approach to cerebral imaging*, Georg Thieme Verlag, Stuttgart, New York.

TALAIRACH, J., SZIKLA, G. and TOURNOUX, P. (1967), *Atlas d'anatomie stereotaxique du telencephale*, Masson, Paris.

TAR HAAR ROMENY, B., FLORACK, L.M., KOENDERINK, J.J. and VIERGEVER, M.A. (1991), 'Scale space: its natural operators and differential invariants', In COLCHESTER, A.C.F. and HAWKES, D.J. (Eds.), *Information Processing in Medical Imaging*, IPMI, Wye, UK, July, p. 239.

TERZOPOULOS, D. and FLEISCHER, K. (1988), 'Modeling inelastic deformation: Viscoelasticity, plasticity, fracture', In DILL, J. (Ed.), *Computer Graphics (SIGGRAPH '88 Proceedings)*, August, pp. 269–278.

TERZOPOULOS, D. and METAXES, D. (1990), 'Dynamic 3D models with local and global deformations: Deformable superquadrics', In *Proc. 3rd I.C.C.V.*, Osaka, Japan, p. 606.

TERZOPOULOS, D. and WITKIN, A. (1988), 'Physically based models with rigid and deformable components', *IEEE Computer Graphics and Applications*, Vol. 8, No. 6, November, pp. 41–51.

TERZOPOULOS, D., PLATT, J., BARR, A. and FLEISCHER, K. (1987), 'Elastically deformable models', *ACM Computer Graphics*, Vol. 21, No. 4, July, p. 205.

TERZOPOULOS, D., WITKIN, A. and KASS, M. (1988), 'Constraints on deformable models: recovering 3D shape and non-rigid motion', *Artificial Intelligence*, Vol. 36, No. 1, pp. 91–123.

THIRION, J.P., GOURDON, A., MONGA, O., GUEZIEC, A. and AYACHE, N. (1992), 'Automatic registration of 3D CAT-scans images using surface curvature', In *3D Advanced Image Processing in Medicine*, 14th annual Proceedings of the IEEE EMBS, Nov., pp. 85–90.

TIEDE, U., BOMANS, M., HOHNE, K.H., POMMERT, A., RIEMER, M., SCHIEMANN, T., SCHUBERT, R. and LIERSE, W. (1993), 'A computerized three-dimensional atlas of the human skull and brain', *AJNR Am J Neuroradiol*, Vol. 14, No. 3, May-Jun, pp. 551–9.

VAN DEN ELSEN, P. (1993), *Multi-modality matching of brain images*, PhD thesis, Utrecht University, Netherlands.

VAN DEN ELSEN, P., POL, E. and VIERGEVER, M. (1993), 'Medical image matching - a review with classification', *IEEE Engineering in Medicine and Biology*, Vol. 12, No. 1, Mar., pp. 26–39.

VANDERMEULEN, D., SEUTENS, P., GYBELS, J., MARCHAL, G. and OOSTERLINCK, A. (1989), 'Delineation of neuroanatomical objects using deformable models', In *Proceedings of Computer Assisted Radiology*, pp. 644–650.

VENOT, A., LEBRUCHEC, J.F., GOLMARD, J.L. and ROUCAYROL, J.C. (1983), 'An automated methods for the normalization of scintigraphic images', *Journal of Nuclear Medicine*, Vol. 24, pp. 529–531.

VERNAZZA, G., SERPICO, S.B. and DELLEPIANE, S.G. (1987), 'A knowledge-based system for biomedical image processing and recognition', *IEEE transactions on circuits and systems*, Vol. CAS-34, No. 11, November, pp. 1399–1416.

WEINBERGER, D., LUCHINS, P., MORIHISA, J. and WYATT, R. (1982), 'Asymmetrical volumes of the right and left frontal and occipital regions of the human brain', *Neurology*, Vol. 11, pp. 97–102.

WIDROW, B. (1973), 'The "rubber mask" technique – 1. pattern measurement and analysis', *Pattern Recognition Letters*, Vol. 5, pp. 175–197.

WITELSON, S. (1977), 'Anatomical asymmetry in the temporal lobes: its documentation, phylogenesis and relationship to functional asymmetry', *Ann N.Y. Acad Sci*, Vol. 299, pp. 328–354.

WOODS, R., CHERRY, S. and MAZZIOTTA, J. (1992), 'Rapid automated algorithm for aligning and reslicing PET images', *Journal of Computer Assisted Tomography*, Vol. 16, No. 4, pp. 620–633.

WOODS, R., MAZZIOTTA, J. and CHERRY, S. (1993a), 'MRI-PET registration with automated algorithm', *Journal of Computer Assisted Tomography*, Vol. 17, pp. 536–546.

WOODS, R., MAZZIOTTA, J. and CHERRY, S. (1993b), 'Automated registration of mri and pet studies', *Journal of Cerebral Blood Flow and Metabolism*, Vol. 13, No. Suppl. 1, p. S842.

ZACHMANN, H. (1991a), 'Interpretation of cranial MR-images using a digital atlas of the human head', *IEEE Transactions on Medical Imaging*, pp. 99–110.

ZACHMANN, H.H. (1991b), 'Interpretation of cranial MR-images using a digital atlas of the human head', In *Proceedings of Computer Assisted Radiology*, pp. 283–288.

ZHENGPING, J. and MOWFORTH, P.H. (1991), 'Mapping between MR brain images and a voxel model', *Med Inf (Lond)*, Vol. 16, No. 2, Apr-Jun, pp. 183–93.

ZUBAL, G., TAGARE, H., ZHANG, L. and DUNCAN, J. (1991), '3D registration of intermodality medical images', In *Proceedings of Computer Assisted Radiology*, pp. 293–294.

ZUCKER, S.W. and HUMMEL, R.A. (1981), 'A three–dimensional edge operator', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. PAMI-3, No. 3, pp. 324–331.