REFERENCE ONLY

UNIVERSITY OF LONDON THESIS

Degree PhD Year 2003 Name of Author CUNNAM, K.M.

COPYRIGHT

This is a thesis accepted for a Higher Degree of the University of London. It is an unpublished typescript and the copyright is held by the author. All persons consulting the thesis must read and abide by the Copyright Declaration below.

COPYRIGHT DECLARATION

I recognise that the copyright of the above-described thesis rests with the author and that no quotation from it or information derived from it may be published without the prior written consent of the author.

LOANS

Theses may not be lent to individuals, but the Senate House Library may lend a copy to approved libraries within the United Kingdom, for consultation solely on the premises of those libraries. Application should be made to: Inter-Library Loans, Senate House Library, Senate House, Malet Street, London WC1E 7HU.

REPRODUCTION

University of London theses may not be reproduced without explicit written permission from the Senate House Library. Enquiries should be addressed to the Theses Section of the Library. Regulations concerning reproduction vary according to the date of acceptance of the thesis and are listed below as guidelines.

A. Before 1962. Permission granted only upon the prior written consent of the author. (The Senate House Library will provide addresses where possible).

B. 1962 - 1974. In many cases the author has agreed to permit copying upon completion of a Copyright Declaration.

C. 1975 - 1988. Most theses may be copied upon completion of a Copyright Declaration.

D. 1989 onwards. Most theses may be copied.

This thesis comes within category D.

☑ This copy has been deposited in the Library of UCL

☐ This copy has been deposited in the Senate House Library, Senate House, Malet Street, London WC1E 7HU.

Orientation Matching for Diffusion Tensor Image Registration

Kathleen M. Curran
UCL Department of Computer Science

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

of the

University of London.

Department of Computer Science

University College London

UMI Number: U591903

All rights reserved

INFORMATION TO ALL USERS
The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript
and there are missing pages, these will be noted. Also, if material had to be removed,
a note will indicate the deletion.

UMI U591903
Published by ProQuest LLC 2013. Copyright in the Dissertation held by the Author.
Microform Edition © ProQuest LLC.
All rights reserved. This work is protected against
unauthorized copying under Title 17, United States Code.

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

Abstract

This thesis develops a registration algorithm specifically for diffusion-tensor (DT) images. The proposed approach matches the tensor orientations to find the registration transformation. Early results show that local optimisation does not find the global minimum in registration of DT-MR brain images. Therefore, a global optimisation registration technique is also implemented. This thesis proposes several new similarity measures for DT registration and provides a comparison of them along with several others previously proposed in the literature. The thesis also proposes several new performance evaluation measures to assess registration quality and develops a performance evaluation framework that uses directional coherence and landmark separation.

Experiments with direct optimisation demonstrate increased local minima in tensor registration objective functions over scalar registration. Using registration with global optimisation, this thesis compares the performance of scalar-derived similarity measures with those derived from the full tensor. Results suggest that similarity measures derived from the full tensor matrix do not find a more accurate registration than those based on the derived scalar indices. Affine and higher-order polynomial registration is not reliable enough to make a firm conclusion about whether diffusion tensor orientation matching improves the accuracy of registration over registration algorithms that ignore orientation. The main problem preventing a firm conclusion is that the local minima problem persists despite the use of global optimisation, causing poor registration of the regions of interest.

Acknowledgements

First and foremost, I would like to thank Dr. Daniel Alexander from the Computer Science (CS) Department at University College London (UCL) for being such an excellent supervisor. Without his constant encouragement and support, scientific rigour and expertise, huge understanding and great patience, multiple proof reading and funding for equipment and conferences and a little bit of pressure I would never have completed this thesis.

Acknowledgements go to the NMR Research Unit of the Institute of Neurology, London, in particular Dr. Claudia Wheeler-Kingshott and Dr. Olga Ciccarelli. They have been very helpful in acquiring and making DT data available and taking the time to answer my questions.

I am also grateful to the Engineering and Physical Sciences Research Council for funding this research.

It would not have been possible to run all the experiments in this thesis if it was not for the use of the UCL Central Computing Cluster (CCC). I would like to acknowledge the support of the UCL Computing Resource Allocation Group (CRAG), in particular, Dr. John Brodholt, the chairman of the CRAG and Mark Dixon, the CCC administrator.

Sincere thanks go to Professor Gareth Barker, from the Centre for Neuroimaging Sciences, Kings College London and Professor David Barber, from the Department of Medical Physics, Sheffield University, for agreeing to be my PhD examiners. I would also like to thank Professor Simon Arridge from CS at UCL and Professor Andrew Todd-Pokropek, from the Medical Physics Department at UCL for examining my PhD transfer viva. I would like to thank them for pointing me in the right direction and for

their constructive criticism of my work.

Many thanks to Professor James Gee, from the Gee Lab, University of Pennsylvania and Dr Alexander for organising my student exchange in Philadelphia and to Jeffrey Duda for agreeing to spend a rainy summer in London. Special thanks to Professor Gee for his scientific guidance and to all in his lab for making me feel so welcome, especially Tessa Sundaram who showed me all the best lunch trucks and bars in Philly and introduced me to on-line chatting, which has become a wonderful distraction when trying to fix a bug or compile a thesis in LaTeX. I am also grateful to Hui (Gary) Zhang for some interesting on-line discussions.

I really appreciate the support of my friends and colleagues, especially, Emma Woodhall, who has been very supportive throughout my Ph.D., and what would I have done without 4 o'clock tea and trips to the JB, with the medical image processing group (MIPG) at UCL? In particular, I would like to thank Philip Cook for his solution to every problem and for proof reading and Dr. Jason Riley for his support and advice. I wish to sincerely thank Dr. William Crum, Dr Matthew Hall and Dr Mark Symms for proof reading and providing such thorough feedback.

Ba mhaith liom buíochas a ghabháil le mo thuismitheoirí Emmet agus Alice Curran, le mo dheirfiúr Áine agus le m'aintín May Curran. Is mór agam a gcuid tacaíochta le linn dom an tráchtas seo a scríobh. Ba mhaith liom buíochas speisialta a ghabháil le mo fear chéile Matthew DiFranco; don fhoighne a bhí aige liom, don grá a thug sé dom, don comhairle faoin dteicneolaíocht agus faoi na profaí. Matt- go raibh míle maith agat. Tiomnaím an tráchtas seo dóibh.

Contents

List of Figures

List of Tables

Notation

$\mathbf{A}$	Affine matrix.
b	Discrete cosine transform (DCT) basis function.
c	Cooling factor.
CC	Correlation coefficient.
D	Diffusion tensor (DT).
D'	Transformed D.
D_{ij}	Element of D.
$\langle\mathsf{D}\rangle$	The mean diffusivity of D.
$\mathsf{D}_1{:}\mathsf{D}_2$	The tensor scalar product of D_1 and D_2.
D	The diffusion deviatoric.
$\mathbf{e}$	Eigenvector of D.
$\mathbf{e}_1$	Principal eigenvector of D.
$\mathbf{e}_3$	Minor eigenvector of D.
$\mathbf{E}$	Matrix of eigenvectors of D.
$\mathbf{F}$	Non-singular linear transformation.
G	Gaussian smoothing operator.
H	Entropy.
I	The identity tensor.
I_1	Source image.
$\bar{I}_1$	Mean intensity of the source image.
I_2	Target image.
$\bar{I}_2$	Mean intensity of the target image.

$I_1 \cap I_2$	The overlap of I_1 and I_2.
$\mathbf{J}$	The Jacobian.
$\mathbf{K}$	Skew matrix.
M	The objective function.
MI	Mutual information.
N	Number of voxels in the image.
NMI	Normalised mutual information.
p	Probability density function of a particle displacement.
$\mathbf{r}$	Displacement.
P	Projection.
r	A random number between 0 and 1.
r_x, r_y, r_z	Components of the rotation axis.
$\mathcal{R}$	Ratio image.
$\mathbf{R}$	Rotation matrix.
RIU	Ratio image uniformity.
S	Similarity measure.
$\mathbf{S}$	Scale matrix.
SSD	Sum of squared differences.
t	Diffusion time.
$\mathbf{t}$	Translation vector.
t_x, t_y, t_z	The x, y and z components of $\mathbf{t}$, respectively.
T	Temperature.
T_0	Starting temperature.
T_f	Freezing temperature.
$\mathcal{T}$	Transformation.
$\mathcal{T}^*$	The optimal transformation.
$\mathcal{T}_x, \mathcal{T}_y, \mathcal{T}_z$	The transformed positions of x, y and z, respectively.
Tr	Trace.

$\mathbf{U}$	Deformation component.
$\mathbf{x}$	Image position.
$\mathbf{x}'$	Transformed image position.
x, y, z	Elements of $\mathbf{x}$.
x', y', z'	Elements of $\mathbf{x}'$.
X, Y, Z	The number of voxels in an image for each dimension.
α	Longitude of axis of rotation.
β	Co-latitude of axis of rotation.
$\alpha_j, \beta_j, \gamma_j$	Coefficients of the j-th basis functions for the x, y and z components of the DCT, respectively.
ν_r	Relative anisotropy.
ν_f	Fractional anisotropy.
δ_1	Relative-anisotropy difference.
δ_2	Modulus difference.
δ_3	Tensor difference.
δ_4	Normalised tensor difference.
δ_5	Tensor scalar product.
δ_7	Principal-direction difference.
δ_p	Park tensor difference.
$\mathcal{E}$	Objective function energy.
$\mathcal{E}_0$	Current energy of the objective function.
$\mathcal{E}_1$	New energy of the objective function.
λ	Eigenvalue of D.
Λ	Diagonal matrix of eigenvalues of D.
η	Correlation ratio.
θ	Rotation angle.
σ	Standard deviation.
Ψ	Skewness.
ω_i	Transformation offset.

Chapter 1

Introduction

In medical imaging, several complementary sources of information are often required, such as scans taken at different times, scans from different modalities and normalised scans from template anatomies. Registration is the process of finding the spatial relation of one image in reference to another image. It is used to superimpose the information from different sources and enables their combined interpretation.

The goal of image registration is to determine a transformation that aligns the features in one image with the corresponding features in another. To automate image registration, a search strategy or optimisation algorithm is required to find the transformation that maximises a similarity measure between the two images. The optimisation algorithm makes successively better estimates of the transformation until eventually converging to a solution. Convergence is reached when no transformation can be found that results in a better value of the similarity measure, to within a preset tolerance.

Medical applications of image registration are numerous and reviews can be found in [70] and [87]. To evaluate the progress of a disease we can register time-series data from the same patient, in one or more modalities [72]. Registering images from the same patient in different modalities allows their simultaneous interpretation. For example, combining functional positron emission tomography (PET) with magnetic resonance imaging (MRI), where soft tissues are better visualised. Registration of pre-operative images with the physical space occupied by the patient during surgery or treatment is a fundamental step in interactive, image-guided surgical techniques [50]

and treatment planning [116]. Finally, registering images to a template anatomy can be useful for comparison purposes. For example, creating digital atlases to study anatomical variability [39, 63]. This thesis investigates the use of a recent MRI technique, diffusion-tensor magnetic resonance imaging (DT-MRI), to improve the accuracy of image registration.

1.1 Motivation

Diffusion-tensor MRI measures the diffusion of water molecules in tissue and can reveal structural information irresolvable with conventional MR imaging techniques. It provides rotationally invariant statistics, which reflect the anisotropy of the tissue microstructure, and it provides an estimate of the dominant orientation of microstructural fibres.

Diffusion-tensor MR images contain complementary information to that contained in standard MR images. Specifically, DT-MR images contain information about tissue-microstructure orientation in addition to the more standard structural information and thus provide additional cues for matching. The exploitation of DT-MR data for registration should improve the accuracy of image matching for scalar data because the additional orientational information can be used both to test registration quality and improve it. To fully exploit the information in DT-MRI we need similarity measures that are sensitive to all aspects of the DT, including size, shape and orientation. Current approaches generalise existing scalar measures or examine specific aspects of the DT, but none of the existing registration algorithms that match tensor orientation have compared DT similarity measures.

This thesis tests the hypothesis that matching the tensor orientations during registration provides a transformation that aligns white matter structures more accurately than transformations obtained without matching orientations.

1.2 Research Objectives

1. To answer the question: does diffusion tensor orientation matching improve the accuracy of registration over registration algorithms that ignore orientation?

2. To determine the most effective similarity measures for DT-MR image registration.

1.3 Hypotheses

We first define the two central thesis hypotheses, C1 and C2:

- **Hypothesis (C1):** "Similarity measures derived from the full tensor matrix find a more accurate registration than those based on the derived scalar indices that do not use tensor orientational information."

- **Hypothesis (C2):** "Diffusion tensor orientation matching improves the accuracy of registration over that from registration algorithms that ignore orientation."

To test the central hypotheses, we implement a DT-MR image registration algorithm, which uses numerical optimisation to minimise an image similarity measure with respect to a parametric image transformation. Image registration algorithms often use direct optimisation techniques such as Powell's algorithm or the Levenberg-Marquardt algorithm [107]. However, such techniques do not guarantee to find the global minimum of an objective function which can cause inaccurate registration. The implementation of the registration algorithm for testing C1 and C2 leads to four specific sub-hypotheses that also require testing:

- **Hypothesis (S1):** "Powell's optimisation finds a consistent transformation for **intra-subject** registration."

- **Hypothesis (S2):** "Powell's optimisation does not find a consistent transformation for **inter-subject** registration."

- **Hypothesis (S3):** "Gaussian smoothing of the source and target images prior to registration improves the consistency of the registration transformation found using Powell's optimisation."

- **Hypothesis (S4):** "The minimum of the objective function found using gradient annealing is better than the minimum found using Powell's method."

1.4 Original Contributions

This section gives a brief summary of the original work that this thesis contributes to medical image registration and, in particular, DT-MR registration.

1.4.1 System for Diffusion Tensor Registration

In addition to the intensity information in traditional magnetic resonance (MR) images, a tensor describing local water diffusion is calculated for each voxel in DT-MR images. In Chapter 3, we outline a general framework for voxel-similarity-based image registration, which we extend to a generalised scheme in Chapter 5. The novelty of the proposed approach is that it matches the tensor orientations to find the registration transformation.

1.4.2 Orientation Matching

Standard registration techniques assume that changes in the transformation do not affect the actual data values in the image, only the spatial location. However, this assumption does not hold for DT-MRI since image transformations change the tensor orientation, see Chapter 4. Including tensor reorientation directly in the optimisation process places much greater constraints on the computed transformation, and therefore the algorithm should find a better voxel-to-voxel match, at least within anisotropic regions, rather than simply associating region boundaries.

1.4.3 Global Optimisation for Registration

Finding the global minimum of the objective function is the key to accurate image registration. However, most existing registration methods use direct optimisation schemes,

such as Powell's algorithm or Levenberg-Marquardt [107], which find a minimum local to the starting point of the optimisation. This local minimum may not be the global minimum. In Chapters 6 and 7, we show that Powell's method does not find the global minimum in image registration. Instead, we implement registration based on a global optimisation technique, simulated annealing, see Chapters 7 and 8.

1.4.4 Similarity Measure Performance Comparison

We propose several new similarity measures for DT registration in Chapter 5 and provide a performance comparison of the new similarity measures and several others previously proposed in the literature, for DT registration, see Chapter 8.

1.4.5 Registration Validation Methodology

The acute sensitivity of the DT orientation to the correctness of the registration transformation provides a unique opportunity for comparison, assessment and validation of registration techniques. We propose several new performance evaluation measures to assess registration quality and develop a validation method that uses direction coherence and landmark separation, see Chapter 8.

1.5 Thesis Outline

In Chapter 2, we briefly review the background on DT-MRI. Image registration is described in Chapter 3 and we discuss the challenges in DT-MRI registration in Chapter 4. Chapter 5 outlines the design and implementation of a registration algorithm that uses DT-MR data, including implementations for rigid, affine and higher-order polynomial registration, using both direct and global optimisation. Chapter 6 describes the registration experiments performed with direct optimisation and highlights the added difficulties, caused by local minima, in orientation matching over scalar image matching. Chapter 7 studies the global optimisation algorithm to determine appropriate settings for running the experiments in Chapters 7 and 8. Chapter 8 details experiments that address the central hypotheses of the thesis outlined in Chapter 1. We compare the effectiveness of the similarity measures, proposed in Chapter 5, both qualitatively and

quantitatively. We then describe an additional set of experiments, which compares image registration with no orientation matching to that with orientation matching. Finally, we draw conclusions in Chapter 9 and discuss possible future research directions.

Chapter 2

Diffusion Tensor Magnetic Resonance Imaging (DT-MRI)

This chapter lists common features derived from the diffusion tensor for visualisation and analysis in order to introduce terms and quantities used in later chapters.

2.1 Background

Diffusion tensor MRI [15] measures the apparent diffusion tensor D at each voxel of an image volume on the assumption that the displacements of water molecules over a fixed time t have a zero-mean Gaussian distribution [35]. With this model, the probability of a molecule undergoing displacement $\mathbf{r}$ in time t, is

$$p(\mathbf{r}) = ((4\pi t)^3|\mathsf{D}|)^{-1/2} \exp\left(-\frac{\mathbf{r}^T\mathsf{D}^{-1}\mathbf{r}}{4t}\right). \tag{2.1}$$

The Gaussian function in Eq. 2.1 has ellipsoidal contours with relative axis lengths $\sqrt{\lambda_1}$, $\sqrt{\lambda_2}$ and $\sqrt{\lambda_3}$ where λ_1, λ_2 and λ_3 are the eigenvalues of the diffusion tensor, D. The major axes of the ellipsoid are oriented with the eigenvectors of D, which is a 3×3 symmetric matrix.

It is often convenient to decompose D into its eigenvectors and eigenvalues:

$$\mathsf{D} = \mathbf{E}\Lambda\mathbf{E}^T \tag{2.2}$$

where,

$$\mathbf{E} = (\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3), \tag{2.3}$$

$$\mathbf{\Lambda} = \text{diag}(\lambda_1, \lambda_2, \lambda_3), \tag{2.4}$$

and the eigenvectors, $\mathbf{e}_1$, $\mathbf{e}_2$ and $\mathbf{e}_3$, describe the orientation of the DT and the eigenvalues, λ_1, λ_2 and λ_3, describe its size and shape, as illustrated in Figure 2.1. Isotropic DT's are prevalent in cerebrospinal fluid (CSF) and grey matter regions of the brain, prolate DTs are prevalent in white matter regions and oblate DTs arise in white matter regions. The tissue microstructure forms barriers to the mobility of water molecules

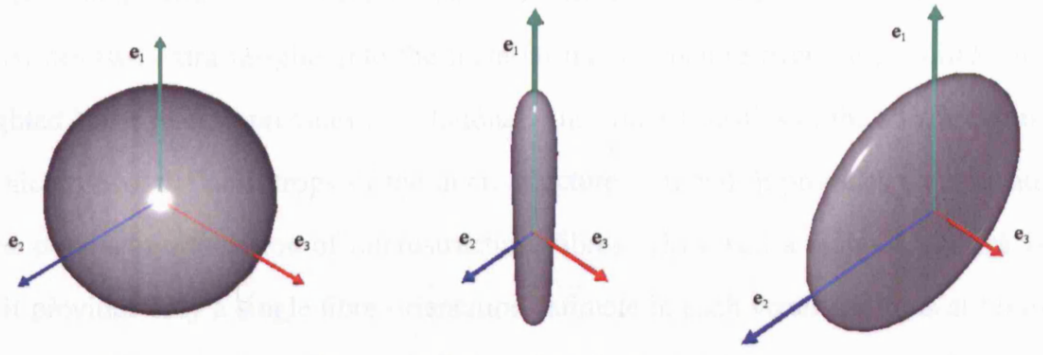

Isotropic: $\lambda_1 \approx \lambda_2 \approx \lambda_3$ Prolate: $\lambda_1 >> \lambda_2 \approx \lambda_3$ Oblate: $\lambda_1 \approx \lambda_2 >> \lambda_3$

Figure 2.1: Isotropic, oblate and prolate DTs, where λ_1, λ_2 and λ_3 are the eigenvalues of D and $\mathbf{e}_1$, $\mathbf{e}_2$ and $\mathbf{e}_3$ are the eigenvectors (adapted from [5]).

within. Different types of tissue have different microstructure and water molecules within have different characteristic scatter patterns. The apparent diffusion tensors in different tissues have characteristic shapes that reflect the microstructure of the tissue. In the brain, for example, regions such as the ventricles contain only CSF. The movement of water molecules in such regions is unhindered in all directions and diffusion is isotropic. In grey matter, mobility is reduced because the tissue contains barriers to water mobility. However, at the typical resolution of diffusion tensor imaging, diffusion also appears isotropic [15], since the barriers have no strongly preferred orientation. In fibrous tissues such as brain white matter, fibre walls in directions parallel to the axis of the fibre hinder movement of water molecules in perpendicular directions, but mo-

bility is unhindered along the axis of the fibre. Water mobility is thus anisotropic and displacements are larger on average along the fibre than across it. Thus the contours of probability, p, extend along the fibre direction. The longest axis of the prolate diffusion ellipsoid is in the fibre direction, so the principal eigenvector of D provides an estimate of the dominant microstructural fibre orientation. More complex microstructures also appear in the brain. At fibre crossings, tissue microstructure has multiple dominant fibre orientations and the contours of p extend along the direction of each fibre. When the microstructure contains two orthogonal and equally weighted fibre orientations, for example, the best-fit Gaussian model then has oblate-shaped contours (Figure 2.1).

Diffusion-tensor MRI is the most popular diffusion MRI reconstruction technique. It provides two extra insights into the material microstructure over simple diffusion-weighted MRI. First, it provides orientationally invariant statistics of the anisotropy of p, which reflect the anisotropy of the microstructure. Second, it provides an estimate of the dominant orientation of microstructural fibres. However, a major drawback is that it provides only a single fibre-orientation estimate in each voxel and fails at fibre-crossings. References [74, 129, 130] all describe alternative reconstruction techniques that do not assume p is Gaussian and can resolve the orientations of multiple fibre populations within single voxels.

2.2 DT Scalar Indices

The primary features of D that we need to account for under a spatial image transformation are the size, shape and orientation. The six independent elements of the diffusion tensor are not informative in isolation but we can derive various rotationally invariant statistics from the elements that describe particular aspects of the tensor. For example trace, anisotropy and skewness [103] are common scalar indices of the size and shape of D and are independent of the orientation of D.

The trace of the diffusion tensor is defined as

$$\mathrm{Tr}(\mathsf{D}) = \mathrm{D}_{xx} + \mathrm{D}_{yy} + \mathrm{D}_{zz} = \lambda_1 + \lambda_2 + \lambda_3 = 3\langle \mathsf{D} \rangle, \tag{2.5}$$

where D_{xx}, D_{yy} and D_{zz} are the diagonal elements of D and the mean diffusivity

$$\langle \mathsf{D} \rangle = \mathrm{Tr}(\mathsf{D})/3. \tag{2.6}$$

The trace of the diffusion tensor is proportional to the mean squared displacement of water molecules and thus indicates the mobility of water molecules within each voxel, which reflects the density of microstructural barriers.

The diffusion tensor, D, can be decomposed into isotropic and anisotropic tensors [17]:

$$\mathsf{D} = \langle \mathsf{D} \rangle \mathbf{I} + \boldsymbol{D}, \tag{2.7}$$

where $\mathbf{I}$ is the identity tensor and $\boldsymbol{D}$ is the deviatoric tensor. In Eq. 2.7, the first term on the right hand side is the isotropic part of D and the deviatoric tensor captures the anisotropic part. Basser et al. [17] take the magnitude of the anisotropic part of D and normalise it by the magnitude of the isotropic part of D to obtain a measure of the relative anisotropy

$$\nu_r = \frac{1}{\sqrt{3}} \frac{\sqrt{\boldsymbol{D} : \boldsymbol{D}}}{\langle \mathsf{D} \rangle} = \frac{\left[3 \sum_{i=1}^{3} (\lambda_i - \langle \mathsf{D} \rangle)^2 \right]^{\frac{1}{2}}}{\sum_{i=1}^{3} \lambda_i}, \tag{2.8}$$

where the tensor scalar product

$$\mathsf{D}_1 : \mathsf{D}_2 = \mathrm{Tr}(\mathsf{D}_1 \mathsf{D}_2). \tag{2.9}$$

Basser, et al. [17] also define the fractional anisotropy

$$\nu_f = \left(\frac{3}{2} \frac{\boldsymbol{D} : \boldsymbol{D}}{\mathsf{D} : \mathsf{D}} \right)^{\frac{1}{2}} = \left[\frac{\left(3 \sum_{i=1}^{3} (\lambda_i - \langle \mathsf{D} \rangle)^2 \right)}{2 \sum_{i=1}^{3} \lambda_i^2} \right]^{\frac{1}{2}}. \tag{2.10}$$

The index ν_f measures the fraction of the magnitude of D that we can ascribe to aniso-

tropic diffusion [17]. Equations 2.8 and 2.10 write ν_r and ν_f in terms of the eigenvalues of D (although it is not necessary to compute the eigenvalues to compute ν_r or ν_f), showing that both indices are second moments of the three eigenvalues normalised by a measure of the tensor magnitude. The normalisation constants ensure that each is in the range $[0, 1]$.

The shape of anisotropic DTs ranges from prolate to oblate, as illustrated in Figure 2.1. Skewness is the third moment of the three eigenvalues and is defined by Basser [13] as:

$$\Psi = \frac{\sum_{i=1}^{3} (\lambda_i - \langle \mathsf{D} \rangle)^3}{3}. \tag{2.11}$$

Prolate DTs have positive skewness and oblate DTs have negative skewness [5]. However, a better measure is skewness normalised for DT size:

$$\Psi = \left(\frac{9 \sum_{i=1}^{3} (\lambda_i - \langle \mathsf{D} \rangle)^3}{2 \sum_{i=1}^{3} \lambda_i^3} \right)^{\frac{1}{3}}. \tag{2.12}$$

2.3 Tensor Visualisation

Unlike scalar data there is a three-dimensional tensor at each voxel position in a DT-MR image. Therefore, greyscale images are not sufficient for the representation of tensor data. The tensor can be represented as an ellipsoid where the main axes lengths correspond to the eigenvalues and their direction to the respective eigenvectors, as discussed in §2.1. Figure 2.2 shows ellipsoid representation of anisotropic diffusion in the highlighted region, the corpus callosum, a white matter region in the brain. Tensors with highest anisotropy are colour-encoded red, while low anisotropy, e.g. isotropic tensors found in CSF, are colour-encoded blue.

2.4 DT-MRI Applications

Current clinical applications of DT-MRI include analysis of stroke, epilepsy, multiple sclerosis, dementia and many other white-matter diseases. References [82] and [47]

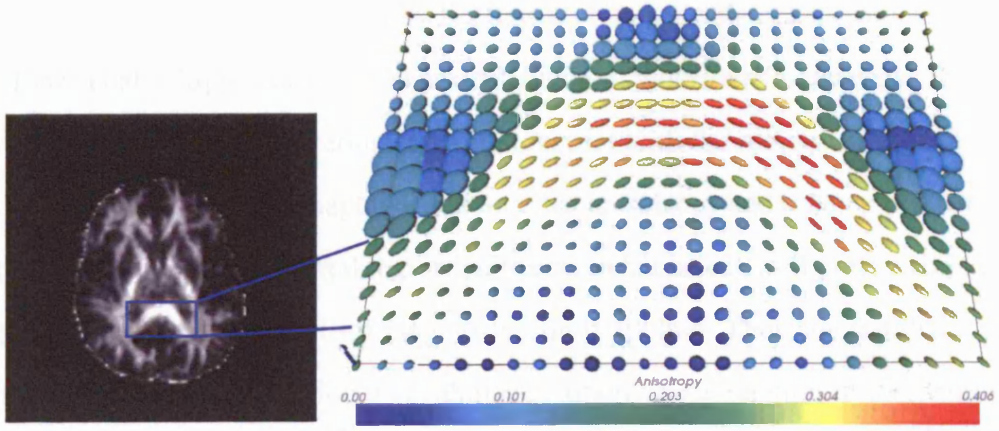

Figure 2.2: Ellipsoid representation of anisotropic diffusion in the corpus callosum, a white matter region in the brain.

give good reviews. Diffusion-tensor MR is able to characterise anisotropy and estimate fibre directions but provides no explicit connectivity information between voxels [77]. The fibre-orientation estimate from the DT principal eigenvector allows fibres to be traced through DT images. The process of tracing white matter pathways using DT-MRI data is referred to as fibre tractography and several algorithms have been proposed including [16, 33, 34, 78, 94, 101, 106]. An important clinical application of fibre tractography is in accurately determining the effect of brain tumours on white matter pathways for radiotherapy treatment and planning prior to surgery [95]. It is also used to assess differences in connectivity between different population groups [31]. Behrens et al. [19] provide the first quantitative demonstration of anatomical connectivity between human white-matter tracts and grey-matter and the first connectivity-based segmentation of grey-matter. They identified specific connections between the human thalamus and cortex using a novel probabilistic tractography algorithm with diffusion imaging data.

Xu et al. [144, 145] present a methodology for creating a probabilistic map of fibre distribution in white matter by combining DT-MR-based fibre tractography and deformable registration methods. Fibres tracked in each individual's image are spatially normalised via an elastic image-warping procedure [43]. They demonstrate on a group of normal subjects that DT-MRI, in conjunction with fibre tracking algorithms, can

produce reconstructions of fibre pathways that are reproducible across subjects.

Jones et al. [76] perform spatial normalisation of whole tensor datasets. They employ the AIR [139] co-registration software together with the reorientation techniques of Alexander et al. [5], see Chapter 4, and find that anisotropy and orientational information is well preserved in central regions of the brain but is less well preserved in more peripheral areas when a simple affine registration is applied. They suggest that a more sophisticated elastic registration algorithm may improve registration in the peripheral regions of the brain.

Chapter 3

Medical Image Registration

This chapter begins with a discussion of medical image registration techniques and their clinical applications. We then describe the essential components of registration algorithms and discuss a variety of examples in the literature.

3.1 Introduction

Image registration is well documented. Maintz and Viergever [87], Lester and Arridge [85], Hill et al. [70], Zitova and Flusser [148], Hajnal et al. [69] and Crum et al. [36] all review image registration methods. The goal of image registration is to determine the spatial alignment between multiple images with similar content, such as brain images, i.e. map corresponding structural or functional anatomy from the same or different subjects, and from the same subject acquired with the same or different modalities.

Registration is achieved by determining a transformation that moves the features of a source image to the positions of the corresponding features in a fixed target image (Figure 3.1). Registration seeks the transformation that minimises some measure of difference between the two images, or equivalently, maximises a measure of correspondence.

Registration tasks can be intra-modal, where the images being registered are from the same imaging modality, e.g. matching two single-photon emission computed tomography (SPECT) images [12]; or multimodal, where two different modalities are involved, e.g. combining an MR image and a computed tomography (CT) image [87].

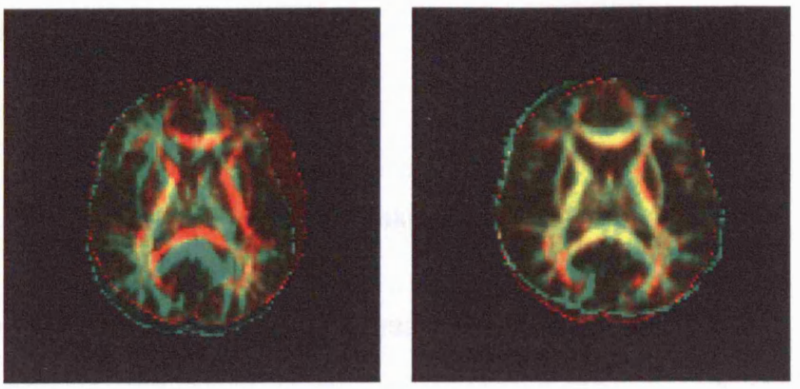

(a) (b)

Figure 3.1: Inter-subject DT-MR brain registration. (a) Pre-registration and (b) post-registration. The source (green) is superimposed on the target (red) so that the combined image is yellow where the images are well aligned.

Combining multiple images from the same individual is known as intra-subject registration. Intra-subject registration has many important applications, such as registration of time series of images for monitoring bone growth in children, degenerative brain disease such as dementia or multiple sclerosis, tumour growth and response to chemotherapy or radiotherapy [87].

Comparing images between individuals (inter-subject registration) or from multiple individuals (cross-population) also has many clinical and research applications, such as finding statistical differences between different population groups, e.g. normals and schizophrenics [1].

3.2 Components of a Registration Algorithm

Brown [26] breaks down the task of image registration into four key components:

1. The image features to be matched, e.g. image intensity or geometrical structures.
2. The search space or range of transformations to be used.
3. The search strategy to find the optimum transformation within the search space.
4. The similarity metric that measures the optimal transformation for the selected feature set.

We discuss each component in the following sections.

3.3 Image Features

Typical image features are [85]:

- Manually selected anatomical landmarks.
- Segmented structure boundaries or surface matching.
- Curvatures or spatial second derivatives of the intensity.
- Image intensity.

3.3.1 Anatomical Landmarks

Image features can be either intrinsic or extrinsic. Intrinsic registration is based on the anatomy of the patient as derived from an image. The information used to register images can be any object that is identifiable in both images as corresponding to the same anatomical location.

Due to their accuracy, extrinsic features such as fiducial markers can be used to define a gold standard for registration [135]. Fiducial markers, such as markers filled with a solution designed to be bright in CT and MR or filled with a positron emitter to allow them to be seen clearly in PET (Positron Emission Tomography), are carefully positioned on a subject prior to scanning to label key anatomical landmarks. Some limitations of this method are that the studies must be planned in advance because special protocols are necessary for registration, accurate positioning of markers can be time consuming, and in some cases the studies are invasive, e.g. stereotactic frames fixated to the head during neurosurgery [87]. For rigid registration, an average target registration error for the whole volume can be estimated from the knowledge of the landmark positions, i.e. distance between corresponding landmarks post-registration [53]. This form of analysis is not generally possible for non-rigid techniques because although the error at landmarks can be established, the error in other parts of the volume is dependent on the transformation model and must be estimated using other means.

3.3.2 Surface Matching

Boundaries, or surfaces, in medical images tend to be more distinct than landmarks, and various segmentation algorithms can successfully locate such high contrast surfaces [70]. The iterative closest point algorithm (ICP) was proposed by Besl and McKay [20] and is probably the most widely used surface matching algorithm in medical imaging applications [70]. The algorithm first identifies the closest model point for each data point and then finds the rigid body transformation that minimises the sum of square distances between corresponding features. Thompson et al. [128] extract the surfaces of the lateral ventricle and the cerebral cortex in a subject's brain and in a corresponding brain atlas automatically. A detailed survey of surface-based medical image registration can be found in [10].

3.3.3 Registration using Crest Lines

An alternative to the idea of using surface matching for registration is to use distinctive features on those surfaces instead. For this to work the surfaces must be smooth, i.e. differentiable up to third-order. Using tools of differential geometry, it is possible to define two principal curvatures k_1 and k_2 at each point on a surface, (the strength of k_1 is greater than k_2) with associated principal directions. Crest lines are the loci of the surface where the value k_1 is locally maximal in its principal direction [93]. Images can be registered by aligning the crests identified in the images. This has greatest applicability when the images are very similar, in which case there will be good correspondence between crest lines [102].

3.3.4 Voxel Similarity Based Registration Scheme

Image registration using voxel similarity or intensity based measures determines the transformation by optimising some measure calculated directly from the voxel values in the source and target images rather than from geometrical structures, such as points or surfaces derived from the images. This transformation can then be used to transform the source image into alignment with the target image within the region of overlap of the two images. This process involves interpolation and needs to take account of the

difference in sample spacing in the two images. A general outline for voxel similarity based image registration follows.

Given two images I_1 and I_2 we want to find the transformation $\mathcal{T}$ such that $\mathcal{T}(I_1)$ matches I_2. To do this, we define a similarity function:

$$M = \int_{I_1 \cap I_2} S(I_1(\mathbf{x}), I_2(\mathbf{x}))\mathrm{dx} \tag{3.1}$$

where S is a similarity measure between the measurement values at position $\mathbf{x}$ in I_1 and I_2 and $I_1 \cap I_2$ is the overlap region of the two images. For registration we want to find $\mathcal{T}$ such that

$$M(\mathcal{T}) = \int_{I_1 \cap I_2} S(\mathcal{T}(I_1(\mathbf{x})), I_2(\mathbf{x}))\mathrm{dx} \tag{3.2}$$

is maximum. To do this we:

- choose some initial $\mathcal{T}$, e.g. $\mathcal{T} = I$, where I is the identity transformation.
- use an optimisation scheme with objective function $M(\mathcal{T})$ to update the parameters of $\mathcal{T}$ iteratively until the maximum of M is found.
- once the maximum is found at say $\mathcal{T}^*$, we apply $\mathcal{T}^*$ to I_1 to obtain the registered image.

3.4 Transformation Space

In order to transform the source image I_1 to match the target image I_2, it is necessary to determine a geometric transformation, $\mathcal{T}$, or mapping from each voxel position (x, y, z), in I_2, to a corresponding position (x', y', z') in I_1. A rigid transformation, consisting of a simple translation and rotation, may be sufficient for accurate intra-subject brain registration because minimal physical deformation can occur, as the brain is constrained within the cranial cavity. However, registration of temporal sequences of the brain in a patient who has had surgical resection, e.g. frontal lobectomy (lobe removal), requires a non-rigid or more complex mapping because, in addition to relative position and orientation, the shape of the brain has changed.

Non-rigid transformations [54] are required for applications to deformable anatomy, such as the heart or lungs, and inter-subject registration. Affine transformations are widely used to correct for scaling errors or skew in data sets [70]. Often other transformations are required to estimate the displacement field of a transformation. Transformations can be represented by polynomials [120], thin-plate splines [66], harmonic functions [9, 80], elastic models [63], radial basis functions [21, 55], wavelets [6, 142], free-form deformations [114] and fluid models [37, 85].

3.4.1 Rigid Transformations

Rigid transformations preserve all distances. In three dimensions (3D), the rigid transformation group has six parameters that specify translations in the three orthogonal directions and a rotation about any axis. A rigid transformation

$$\mathcal{T}(\mathbf{x}) = \mathbf{R}\mathbf{x} + \mathbf{t} \tag{3.3}$$

is the combination of a 3×3 rotation matrix $\mathbf{R}$ and a translation vector $\mathbf{t}$, where $\mathbf{x}$ is the point to be transformed.

3.4.2 Affine Transformations

An affine transformation is given by

$$\mathcal{T}(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{t}, \tag{3.4}$$

where $\mathbf{A}$ is any real-valued 3×3 matrix. This transformation includes all rigid transformations as well as scaling and shears. In 3D, affine transformations have twelve degrees of freedom (DOF). They preserve the straightness of lines and parallelism but allow lengths and angles between lines to change.

Transformations are often written in matrix form because several transformations can be combined simply by multiplying the matrices together to form a single matrix. The affine transformation is often written as a 4×4 matrix, combining $\mathbf{A}$ and $\mathbf{t}$ and

using homogenous coordinates for $\mathbf{x}$:

$$\begin{pmatrix} x\prime \\ y\prime \\ z\prime \\ 1 \end{pmatrix} = \begin{pmatrix} a_{00} & a_{01} & a_{02} & t_x \\ a_{10} & a_{11} & a_{12} & t_y \\ a_{20} & a_{21} & a_{22} & t_z \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix}. \tag{3.5}$$

3.4.3 Polynomial Transformations

Curved transformations do not preserve the straightness of lines. In curved transformations, the simplest functional form for $\mathcal{T}$ is a polynomial [27, 66]. For example, the quadratic transformation model is defined by second-order polynomials

$$\begin{aligned} \mathcal{T}(\mathbf{x}) &= \begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} \\ &= \begin{pmatrix} a_{00} & \dots & a_{08} & a_{09} \\ a_{10} & \dots & a_{18} & a_{19} \\ a_{20} & \dots & a_{28} & a_{29} \\ 0 & \dots & 0 & 1 \end{pmatrix} \begin{pmatrix} x & y & z & x^2 & xy & xz & y^2 & yz & z^2 & 1 \end{pmatrix}^T \end{aligned} \tag{3.6}$$

whose coefficients are the 30 degrees of freedom of the transformation. This model extends naturally to higher-order polynomials such as cubic (60 DOF) and quartic (105 DOF). A 3D polynomial of order N has a total of $(N+1) * (N+2) * (N+3)/2$ independent parameters [140]. Higher-order polynomials above fifth-order tend to introduce artefacts such as oscillations [125] and are rarely used for non-rigid registration.

3.4.4 Basis Functions

An alternative to using polynomials is to use a linear combination of basis functions b_j to describe the deformation field. Common choices are trigonometric [9] or wavelet

bases [6]. The transformation from coordinates (x, y, z) to (x', y', z') is:

$$x' = x + \sum_{j=1}^{n} \alpha_j b_j(x, y, z) \tag{3.7}$$

$$y' = y + \sum_{j=1}^{n} \beta_j b_j(x, y, z) \tag{3.8}$$

$$z' = z + \sum_{j=1}^{n} \gamma_j b_j(x, y, z) \tag{3.9}$$

where α_j, β_j and γ_j are the j-th coefficients for x, y and z, respectively, and $b_j(x, y, z)$, is the j-th basis function at position (x, y, z). In statistical parametric mapping (SPM), Ashburner et al. [9] use the discrete cosine transform (DCT) basis functions. The n-th order DCT basis functions have the form

$$b(x, y, z) = \cos\left(\frac{i\pi x}{X}\right) \cos\left(\frac{j\pi y}{Y}\right) \cos\left(\frac{k\pi z}{Z}\right), \tag{3.10}$$

where max(i, j, k) = n. Thus for example, the order-zero basis function is:

$$b_0(x, y, z) = \cos\left(\frac{0\pi x}{X}\right) \cos\left(\frac{0\pi y}{Y}\right) \cos\left(\frac{0\pi z}{Z}\right) = 1 \tag{3.11}$$

where X, Y and Z are the number of voxels in the image for each of the three dimensions. The first-order basis functions are 1, $b_1(x, y, z) = cos\left(\frac{\pi x}{X}\right)$, $b_2(x, y, z) = cos\left(\frac{\pi y}{Y}\right)$, $b_3(x, y, z) = cos\left(\frac{\pi z}{Z}\right)$, $b_4(x, y, z) = b_1 b_2$, $b_5(x, y, z) = b_1 b_3$, $b_6(x, y, z) = b_2 b_3$ and $b_7(x, y, z) = b_1 b_2 b_3$. The DCT basis provides a spectral representation of the deformation field where each basis function describes a particular frequency of the deformation.

3.4.5 Spline Warps

Splines are curves that are built in segments. Each segment has control points that dictate the shape of the curve. Many registration techniques using splines are based on the assumption that a set of corresponding points or landmarks have been identified in the source and target images. At these control points, spline-based transformations

either interpolate or approximate the displacements which are necessary to map the location of the control point in the target image into its corresponding location in the source image. Between control points they provide a smoothly varying displacement field.

The control points can be determined in a number of different ways. For example, anatomical or geometrical landmarks that can be identified in both images can be used to define a spline-based mapping function that maps the spatial position of landmarks in the source image into their corresponding position in the target image [23]. Alternatively, control points can be arranged with equidistant spacing across the image, forming a regular mesh [45]. In the latter case the control points are only used as a parameterisation of the transformation and do not correspond to anatomical or geometrical landmarks.

3.4.6 Thin-Plate Splines

Thin-plate splines can be defined as a linear combination of radial basis functions and have been widely used for image registration [22, 23, 66]. There are a number of alternative choices for radial basis functions including multiquadrics and Gaussians [8, 45]. Modelling deformations using thin-plate splines has a number of advantages. For example they can be used to incorporate additional constraints such as rigid bodies [86] or directional constraints [24] into the transformation model.

3.4.7 B-Splines

The global influence that radial basis function control points have on the transformation is sometimes undesirable since it becomes difficult to model local deformations. An alternative is to use free-form deformations (FFDs) based on locally controlled functions such as B-splines, which have been used successfully for image registration [46, 114]. The basic idea of FFDs is to deform an object by manipulating an underlying mesh of control points. In contrast to radial-based splines, which allow arbitrary configurations of control points, spline-based FFDs require a regular mesh of control points with uniform spacing. FFDs are controlled locally, making them computationally efficient even

for a large number of control points.

3.4.8 Elastic Deformations

Elastic registration was first proposed by Bajcsy and Kovacic [11] for matching a brain atlas with a CT image of a human subject. Elastic registration [25, 11, 63, 61, 4] models the deforming image as an elastic body subject to external forces, which drive the registration, and internal forces (stresses), which impose smoothness constraints. A common choice for the external force is the gradient of a similarity measure such as a local correlation measure based on intensities [11], intensity differences [29] or intensity features such as edge and curvature [62]. An alternative choice is the distance between the curves [44] and surfaces [128] of corresponding anatomical structures. The major problem in elastic registration is the determination of the material properties and forces that act on the tissue [54]. An extension of the elastic-registration framework has been proposed by Davatzikos [43] to allow for spatially varying elasticity parameters. This enables certain anatomical structures to deform more freely than others.

3.4.9 Viscous Fluid Model

The elastic model is limited by the fact that highly localised deformations cannot be modelled since the deformation energy caused by stress increases proportionally with the strength of the deformation [85]. An alternative is a viscous fluid model, in which these internal forces relax as the image deforms over time [30, 84, 85, 37]. This makes fluid registration especially attractive for inter-subject registration tasks (including atlas matching) which have to accommodate large deformations and large degrees of variability. An advantage of fluid registration is that the transformation model implicitly forbids tearing. In many situations, e.g. serial MR brain registration of a subject undergoing diffuse atrophy, there is a prior expectation that the registration of one image to another should be diffeomorphic, i.e. if the transformation were applied to a real physical object then no tearing of the object would occur. The problem is that tearing can often result in a transformation that makes the images more similar despite it being physically invalid. The danger of fluid registration is that the opportunity for misreg-

istrations is increased, as fluid registrations have a vast number of degrees of freedom, generally involving the growth of one region instead of a shifting and / or distortion of another [85].

3.5 Optimisation

Registration can be formulated as an optimisation problem whose goal is to minimise an associated energy or objective function. Several numerical methods exist to optimise a function M that depends on one or more independent variables [107]. Optimisation is a means of finding the point in the parameter space, where M has a maximum or $(-M)$ has a minimum value. Usually, optimisation seeks the global minimum of the objective function, although a common problem with optimisation techniques is that they can converge to sub-optimal solutions at local minima.

The optimisation method or search strategy is critical for registration because the registration accuracy depends on the optimisation quality [75]. For registration, an optimisation method is required to find the transformation that maximises a similarity or equivalently minimises a difference measure between two or more images. In the registration optimisation problem the objective function often contains a very large number of local minima and the number increases with the number of variables. These local minima are caused by noise and artefacts in the image and correspond to sub-optimal solutions. Finding the optimal solution requires finding the global minimum.

Some of the available local optimisation algorithms are gradient-free methods, such as Powell [107] and Nelder and Mead [107]. Others require evaluation of the derivatives of the objective function, such as the conjugate gradients method [107] and are usually more efficient than gradient-free methods. Other methods, such as Newton-Raphson and Levenberg-Marquardt, require evaluation of the first and second derivatives of the objective function. All these methods offer limited reliability for registration because they often get trapped in local minima, which is a common cause of registration error and affects the accuracy of the registration [75].

Some of these minima may be very small, caused either by interpolation artifacts

or a local good match between features or intensities [104]. These small minima can be removed to some extent by blurring the images prior to the registration [70]. Gaussian smoothing should reduce the number of local minima in the objective function. Woods et al. [139] found that smoothing the data with a Gaussian filter improved registration accuracy. However, if too high a degree of smoothing is applied, they found that the registration is less accurate than when no smoothing is applied.

An alternative way to improve reliability is to use a global optimisation algorithm. Global optimisation is the task of finding the absolute best set of parameters to optimise an objective function. Since a systematic search in the function domain requires so many function evaluations as to be impractable, one must rely on a limited search, either systematic or random. The most simple algorithm is to run a direct algorithm many times and from different starting points in the hope that the global minimum will be amongst the local minima obtained over the many runs. The basin of attraction of the global minimum can be very narrow and therefore difficult to reach unless a lucky start is made close to the global minimum. A review of statistical methods for global optimisation can be found in [147]. Genetic algorithms (GA) were not originally developed for optimisation [79] and basic GA does not offer any statistical guarantee of finding a global minimum [56]. The simulated annealing (SA) algorithm [107] does guarantee to find the global minimum under ideal conditions. However, these conditions usually imply that a very long time must be spent in calculation (days or weeks or infinity). However, for objective functions with many minima, these techniques increase the chance of finding the global minimum and will usually find better minima than simple gradient-descent methods.

Jenkinson et al. [75] find that the use of Powell's optimisation [107] method together with the standard multi-resolution approach (series of sub-samplings) is not sufficient to find the global minimum reliably. A multi-resolution framework typically involves starting with low resolution images (containing only gross features) and working progressively through to higher resolutions, in an attempt to avoid local minima. However, they find that their multi-resolution method is more reliable at finding the

global minimum than several of the currently available registration packages in common usage, all of which use a single shot local optimisation algorithm [75]. Butz et al. [28] use a genetic optimisation algorithm to perform affine multi-modal, inter-subject registration. They describe the drawback of local minima when matching fine blood vessels in the retina. The global optimisation correctly matches the entire blood vessels that they were matching and not just vessel segments.

Numerical methods including Powell, the downhill simplex method, Levenberg-Marquardt, Newton-Raphson, gradient descent methods, genetic algorithms and simulated annealing are well documented in [107]. Here we briefly describe Powell's optimisation [107] and simulated annealing [107], which are the methods we concentrate on later in the thesis.

Powell's method is a powerful optimisation technique that does not require calculation of the gradient of the objective function, which makes it particularly useful for image registration where the gradients can be difficult and expensive to compute. The disadvantage of this method is that it is a local search, which is not guaranteed to find the global minimum.

In order to avoid being trapped in local minima and thus sub-optimal solutions, a global optimisation technique, simulated annealing (SA) [107] can be employed. Simulated annealing methods use an analogy with thermodynamics. Slowly cooled systems, such as materials that form crystals, find the global minimum energy state, or configuration of molecules in the crystal, whereas quickly cooled systems, i.e. quenched systems, do not. Instead, quenched materials find a local minimum, which are imperfect crystals. Simulated annealing has been reported to perform very well in the presence of a very high number of variables [73]. It is based on random evaluations of the objective function, which allow for transitions out of local minima. In practice, the procedure does not guarantee to find the global minimum, but if the function has many near-optimal solutions, it should find one. The algorithm first searches an area in the function domain where the global minimum should be present, following the gross behaviour irrespective of small local minima found on the way. It then develops finer

details, finding a good, near optimal local minimum, if not the global minimum itself.

The idea of SA as an optimisation algorithm was first introduced by Kirkpatrick et al. [81] who used the method for the optimisation of the traveling salesman problem [107] and the design of digital computer circuits. The SA algorithm employs the Metropolis algorithm [92] to reach a "steady state" at each of a decreasing sequence of temperatures. The Metropolis algorithm uses the general scheme of taking a downhill step while sometimes taking an uphill step.

The following elements are required to optimise problems with continuous N-dimensional control spaces, e.g. finding the global minimum of some objective function $\mathcal{E}(\mathbf{x})$, in the presence of many local minima, where $\mathbf{x}$ is an N-dimensional vector:

- A description of possible system configurations, where the system state is the point $\mathbf{x}$.
- A generator of random changes in the configuration, i.e. a procedure for taking a random step from $\mathbf{x}$ to $\mathbf{x} + \Delta\mathbf{x}$.
- An objective function, i.e. the value of $\mathcal{E}$ (analog of energy) whose minimisation is the goal of the procedure.
- A control parameter T (analog of temperature) and an annealing schedule, i.e. the choice of initial temperature, how many iterations are performed at each temperature, and how much the temperature is decremented at each step as cooling proceeds.

The Metropolis algorithm starts with an initial value of the objective function $\mathcal{E}_0$ (or "energy") that is the best solution to the problem found so far, and the temperature is set to the initial high temperature T_0. At each step of the iteration a new configuration is generated by applying a random perturbation to the current configuration. If the difference $\Delta\mathcal{E}$ between the new and the old objective functions is negative, i.e. the energy decreases, then the new configuration is accepted. If the difference $\Delta\mathcal{E}$ is positive, i.e. the energy increases, the new configuration can still be accepted with a certain

probability. This acceptance probability depends on a control parameter which plays the role of the temperature T. The higher the temperature, the higher the probability of accepting configurations that increase the objective function and vice versa.

The basic SA algorithm where T is the temperature is as follows:

1. set T at a high starting temperature.

2. optimise at given T using Metropolis simulation.

3. lower T.

4. go to 2 until T is less than some fixed "freezing" temperature.

Simulated annealing optimisation starts with a Metropolis simulation [92] at a high temperature. This means that a relatively large percentage of the random steps that result in an increase in the energy will be accepted. During each iteration, the temperature is lowered until the system is considered frozen. The annealing schedule determines the speed of the decrease of the temperature. Even at low temperatures the system has some chance of increasing its energy state. This allows systems to get out of local minima (given enough time).

The annealing schedule requires the definition of a cooling factor which fixes the rate at which the temperature is decreased. The annealing schedule also requires the specification of an initial temperature T_0 and final temperature T_f of the system, that is the temperature at which the system is considered frozen. In 1984 Geman and Geman [64] proved the existence of annealing schedules that guarantee convergence to the minimum energy state. We need to allow enough iterations at each temperature so that the system stabilises at that temperature; however, the number of iterations at each temperature to achieve this is exponential in the problem size. In practice we need to compromise by either running a large number of iterations at a few temperatures or a small number of iterations at many temperatures or a balance between the two. It is important to lower the temperature slowly and spend a long time at temperatures in the vicinity of the freezing point to avoid local minima [81]. If the initial temperature of

the system is too low or cooling is done insufficiently slowly the system may become quenched, forming defects or freezing out in metastable states (ie. trapped in a local minimum energy state).

3.6 Objective Functions

In image registration the objective function is the similarity between the source and target images. Similarity measures can be divided into two main categories: point-based and voxel-based. Point-based similarity measures minimise the distance between features such as points, curves or surfaces of corresponding anatomical structures and require prior feature extraction. In recent years, voxel-based similarity measures such as sums of differences, cross correlation or mutual information have become increasingly popular. These were shown in a study by West et. al. [135] to out-perform the geometric methods in terms of accuracy and have the advantage that they do not require any feature extraction process. Mis-registration can occur because the objective function gives a lower value for some incorrect transformation, i.e. the global minimum of biological correspondence is not coincident with global voxel similarity minimum, or the optimisation algorithm fails to find the transformation with the lowest objective value. Voxel intensity similarity measures include sum of squared differences, ratio image uniformity, correlation coefficient, correlation ratio and mutual information measure [70].

3.6.1 Sum of Squared Differences

For two images, I_1 and I_2 the sum of squared differences (SSD) is

$$SSD(I_1, I_2) = \sum (I_1(\mathbf{x}) - I_2(\mathbf{x}))^2 \tag{3.12}$$

where the sums are taken over the overlap of the two images, $I_1 \cap I_2$. The SSD is used for intra-modality registration using voxel similarity measures [70]. It is popular for driving non-linear registration methods because it is fast to compute and its derivative is easy to obtain [84]. However, this measure is inappropriate for cross-modality

matching because we cannot assume that tissues to be matched have similar intensities in both modalities.

3.6.2 Correlation Coefficient

The SSD measure makes the assumption that the images only differ by Gaussian noise. The correlation coefficient (CC) similarity measure assumes there is a linear relationship between the intensity values in the two images [70] and is defined:

$$CC(I_1, I_2) = \frac{\sum(I_2(\mathbf{x}) - \overline{I_2})(\mathcal{T}(I_1(\mathbf{x})) - \overline{I_1})}{\left[\sum(I_2(\mathbf{x}) - \overline{I_2})^2 \sum(\mathcal{T}(I_1(\mathbf{x})) - \overline{I_1})^2)\right]^{1/2}}, \tag{3.13}$$

where $\overline{I_1}$ and $\overline{I_2}$ are mean intensity values of the source and target images, respectively and where the sums are over the overlap of the transformed source and target images. The CC should be close to 1 for perfectly registered images.

3.6.3 Correlation Ratio

The correlation ratio [113]

$$\eta(I_1|I_2) = 1 - \frac{1}{N\sigma^2} \sum_i N_i \sigma_i^2 \tag{3.14}$$

assumes a functional relationship between intensities. It can be defined in terms of sums and sums of squares of source voxels that correspond to a number N_i of iso-intense voxels in the target image, where

$$\sigma^2 = \frac{1}{N} \sum_{I_1 \cap I_2} \mathcal{T}(I_1(\mathbf{x}))^2 - m^2, \qquad m = \frac{1}{N} \sum_{I_1 \cap I_2} \mathcal{T}(I_1(\mathbf{x})) \tag{3.15}$$

and

$$\sigma_i^2 = \frac{1}{N_i} \sum_{\substack{I_1 \cap I_2 \\ i = I_2(\mathbf{x})}} \mathcal{T}(I_1(\mathbf{x}))^2 - m_i^2, \qquad m_i = \frac{1}{N_i} \sum_{\substack{I_1 \cap I_2 \\ i = I_2(\mathbf{x})}} \mathcal{T}(I_1(\mathbf{x})). \tag{3.16}$$

3.6.4 Ratio Image Uniformity

The ratio image uniformity (RIU) intra-modality measure was proposed by Woods et al. [141] for registration of multiple PET images of the same subject and has subsequently been widely used for serial MR registration of the brain [57]. The RIU is also known as the variation of intensity ratios (VIR) and is calculated from images I_1 and I_2,

$$RIU = \frac{\sqrt{\frac{1}{N}\sum(\mathcal{R}(I_2(\mathbf{x})) - \bar{\mathcal{R}})^2}}{\bar{\mathcal{R}}} \tag{3.17}$$

where

$$\mathcal{R}(\mathbf{x}) = \frac{I_2(\mathbf{x})}{\mathcal{T}(I_1(\mathbf{x}))} \tag{3.18}$$

and

$$\bar{\mathcal{R}} = \frac{1}{N}\sum \mathcal{R}(\mathbf{x}) \tag{3.19}$$

where the sums are over N voxels in the overlap of the transformed source and target images. For each estimate of the registration transformation, a ratio image $\mathcal{R}$ is calculated by dividing each voxel in I_2 by each voxel in I_1. The uniformity of $\mathcal{R}$ is then determined by calculating the normalised standard deviation of $\mathcal{R}$. The RIU algorithm iteratively determines the transformation $\mathcal{T}$ that minimises the normalised standard deviation, i.e. maximises uniformity.

3.6.5 Mutual Information

A survey of mutual-information-based registration of medical images can be found in [105]. The mutual information MI of two images I_1 and I_2 is defined in terms of the average information content or entropies H(I_1) and H(I_2) of the images, combined with their joint entropy H(I_1,I_2), as follows [99]:

$$MI(I_1, I_2) = H(I_1) + H(I_2) - H(I_1, I_2) \tag{3.20}$$

where $H(.)$ is the Shannon entropy [118]. The Shannon entropy can be computed for an image on the distribution of the grey values of the image. A probability distribution

of grey values can be estimated by counting the number of times each gray value occurs in the image and dividing those numbers by the total number of occurrences. An image consisting of almost a single intensity will have a low entropy value; it contains very little information. A high entropy value will be yielded by an image with more or less equal quantities of many different intensities, which is an image containing a lot of information. Mutual information measures the statistical dependence of one image on another. It is the probability of accurately predicting the pixel intensities in one image given the intensities at corresponding locations in the other. The joint intensity histogram is an estimate of the joint probability distribution of image intensities [71], which disperses as misalignment increases [71].

Studholme et al. [123] propose a normalised measure of mutual information, which is less sensitive to changes in overlap:

$$NMI(I_1, I_2) = \frac{H(I_1) + H(I_2)}{H(I_1, I_2)} \tag{3.21}$$

They found a distinct improvement in the behavior of the normalized measure for rigid registration of MR-CT and MR-PET images.

3.6.6 Extreme Behaviour of Objective Functions

There are two main forms of extreme behaviour that need to be considered when optimising objective functions, both of which are associated with the fact that the images have a finite field of view. These are:

- When the overlap between images becomes very small (which occurs for example, with large translations).
- When the scales become very disparate (which occurs with large scale factors).

Both situations can lead to artificially low values of the objective function. To avoid converging in these extreme solutions, it is common to associate large objective values with these conditions, as they reflect poor registrations but are easily detected within the registration procedure.

3.7 Example Registration Algorithms

The automatic image registration (AIR) package [139] is freely available registration software that is used in many laboratories [122, 135]. The following is a brief outline of their algorithm:

- Optional smoothing or interpolation to cubic voxels.
- The source image is resampled, requiring a spatial transformation model and an interpolation model to compute voxel intensities, to match the target image.
- A threshold is applied to source and target images to exclude voxels, for example outside the head.
- Optional editing to exclude voxels outside the target image that contain no useful spatial information.
- Evaluation of the objective function and its derivatives with respect to the spatial transformation parameters.
- Minimisation to compute new spatial transformation matrix.
- Evaluation of the termination criteria.
- Save optimal transformation matrix.

In AIR, Woods et al. [139] compute the similarity measure in a very basic way by computing the similarity of the image intensity at some number of locations evenly distributed in the image volumes. If the image has millimeter dimensions μ, ν and ω and we choose to compute the similarity at a sampling resolution of n, then

$$M = n^{-3} \sum_{i=1}^{n} \sum_{j=1}^{n} \sum_{k=1}^{n} S(\mathcal{T}(I_1(\mathbf{x}_{ijk})), I_2(\mathbf{x}_{ijk})) \tag{3.22}$$

where

$$\mathbf{x}_{ijk} = \left(\frac{i \times \mu}{n}, \frac{j \times \nu}{n}, \frac{k \times \omega}{n} \right). \tag{3.23}$$

Eq. 3.22 is normalised by the number of points that actually contribute to the integral, because ones that fall below the threshold are excluded. This method may be improved by use of some of the techniques described in [107], such as Gaussian quadrature. In Woods optimisation scheme [139], they compute the derivatives of M and optimise the objective function with a quasi-Newton method [107].

Other commonly used registration programs include FMRIB's Linear Image Registration Tool (FLIRT) [75] and Statistical Parametric Mapping (SPM) [58]. The FLIRT program uses an affine global optimisation method (combining local optimisation with several search strategies, see §3.5) to minimise the correlation ratio or mutual information objective functions. The SPM program uses a least squares fitting procedure, based on spatial and intensity transformations applied to the source image to match the target [58, 9]. Ashburner et al. [9] use the discrete cosine transform basis to provide a spectral representation of the deformation field where each basis function describes a particular frequency of the deformation.

3.8 Interpolation

For discrete data, the intensity is normally defined only on a grid of discrete locations or lattice sites. In general, the geometric mapping to find which point in the source image volume corresponds to each point in the target image will not fall directly on a voxel. To evaluate the intensity at intermediate locations requires interpolation. The interpolation can be viewed as reconstructing a full continuous image from the discrete points, although to evaluate the objective function it is usually only necessary to know the intensity at the corresponding lattice sites.

Nearest neighbour interpolation simply assigns to a voxel in the output image the intensity value of the closest voxel in the input image, while trilinear interpolation calculates the distance weighted average of the eight neighbouring voxels. Typically, interpolation methods are based on a convolution of the discrete data with some continuous kernel such as trilinear, spline and windowed sinc kernels [127]. There is a trade-off in accuracy versus computational speed, with nearest neighbor and linear in-

terpolation being faster but less accurate, and spline and windowed sinc interpolation being slower and more accurate [97]. One major effect that the choice of interpolation has is to what degree the objective function becomes continuous or discontinuous. This is also affected by the boundary conditions used, such as padding with zeros or only using the overlapping volume.

3.9 Techniques for Assessing Registration Performance

This section reviews various approaches for assessing the accuracy of image registration algorithms. There are two main reasons why registration performance needs to be evaluated. Clinical diagnosis may depend on accurate registration and newly developed algorithms will require comparison with existing methods.

3.9.1 Qualitative Assessment

In many situations it is effective to assess registration quality visually, using tools such as colour overlays when the images are superimposed or crosshairs when they are side by side. Gross misregistration will be easily identified using this approach. Images of the differences between deformed source and target may show more clearly where registration was inaccurate.

3.9.2 Quantitative Assessment

It is difficult to define a registration error as a single number because it is normally spatially varying over the image. The lack of a gold standard means that we can only supply a measure of accuracy by reference to controlled phantom studies, simulations or other registration methods. These validation methods are reviewed in [87]. A complete validation study requires a great deal of effort and time. An excellent example of such a study can be found in [135], where West et al. present evaluation results for intermodality brain registration from a large number of institutions around the world. Validation methods based on external markers can be found in [52], where Fitzpatrick et al. define the target registration error (TRE), or error induced by inaccurate localisation of the fiducial markers at a given target, for predicting errors. These methods are

compared with manually identified anatomical landmark-based registration in [63, 122] and with frame-based registration in [122, 141].

If images are available for which the correct registration transformation is known, then registration accuracy can be calculated by comparing the transformation calculated by any registration algorithm with the known solution. In the case of simulated data, where one modality is simulated from the other to obtain a registration standard [38] or synthetic data [5], the gold standard can have arbitrary accuracy but the images are often not very realistic.

In [117] a biomechanical model of the human breast is used to simulate MR images of a breast subject to mechanical forces as might be experienced during biopsy or movement during dynamic contrast-enhanced imaging. For inter-modality registration invasive markers can be used to provide the gold standard. This approach has so far only been shown effective for rigid body inter-modality registration. For non-rigid registration, or intramodality registration the approach is not satisfactory because much higher accuracy is required. Various kinds of consistency test are also used in validation. For example, Holden et al. [72] measure the consistency of transformations, i.e. establishing that registration of source to target produces the same alignment as from target to source. Given three images, 1, 2, 3, there are three transformations that can be compared $\mathcal{T}_{1\to 2}$, $\mathcal{T}_{2\to 3}$ and $\mathcal{T}_{3\to 1}$. Combining all three transformations completes a circuit and should give the identity transformation for a perfect algorithm.

3.10 Clinical Applications

Registration of medical images has become a vital part of routine clinical practice. References [70, 87] review medical applications of image registration. In this section we describe some example applications, including:

- rigid registration of three-dimensional images of the same subject taken at different times using the same modality,
- rigid registration of three-dimensional images of the same subject taken at different times using different modalities,

- inter-subject registration and
- non-rigid registration.

3.10.1 Serial Image Registration

Time series of images are acquired for various reasons, such as monitoring tumour growth, post-operative monitoring of treatment or evaluating response to chemotherapy. It is common for a patient to have multiple MR images acquired at different times, in order to monitor disease progression or response to treatment. Traditional radiological reporting of these images involves viewing the printed films side-by-side on a light box. This approach can make it difficult to identify small changes in the images between scans. A more sensitive technique is to subtract the images, generating difference images, and to view the difference images to identify regions of the images that have changed. Serial MR image registration is usually applied to the brain [57, 72] and uses a rigid transformation. The SSD and CC similarity measures are widely used for serial MR registration because the images are likely to be very similar. Lemieux [83] has estimated a registration accuracy of the order of 50 μm is required to avoid a mis-registration artifact.

3.10.2 Inter-Modality Registration

The motivation for inter-modality registration is to combine complementary information about the same patient from different imaging modalities, for example MR and computed tomography (CT) [116] or MR and positron emission tomography (PET) [80, 116, 135, 141]. Registration of MR and CT is most often applied to images of the head and is used for surgery and radiotherapy planning [116]. Since MR and CT images have very different intensity distributions, image subtraction is never used for viewing the aligned images. As a result, the very high accuracy requirements of serial MR registration do not apply here. The accuracy requirement for MR-CT registration is usually governed by the accuracy of the treatment system that will subsequently be used. Since neurosurgery and radiosurgery systems are accurate to a small number of millimeters, registration accuracy of 1mm is normally sufficient [70].

Registration of PET and MR is frequently desired, to combine the functional information from a PET image with the anatomical detail from an MR image, to assist in their interpretation. Although the low resolution of PET images (6-10mm) reduces the required registration accuracy compared to MR-CT registration, there are several difficulties in this registration application. Firstly, some PET tracers are very specific to certain parts of the brain. The mutual information (MI) similarity measure has been shown to work well when registering MR images to PET fluro-deoxyglucose (FDG) images [135]. However, the accuracy of registration of the more specific tracers will be tracer-dependent and may be much lower. Secondly, for patients with large lesions, the PET images can have very large intensity abnormalities. A lesion that makes a small difference to image intensities in an MR image may make a very large difference to intensities in the PET image. However, in the validation study by West et al. [135] the registration accuracy using MI was better than 4mm, despite these intensity abnormalities.

3.10.3 Inter-Subject Registration

In order to study variability between subjects we register images from different individuals into a standard space. This type of registration is frequently called inter-subject normalisation. The most common space to align images is in the Talairach stereotactic space [126]. An affine transformation is generally used [32] but more degrees of freedom are sometimes used [9, 30, 58, 140]. This is normally an intra-modality registration problem, as it is normal to register the images of a particular subject to the images of another subject of the same modality whose images have already been put into Talairach space.

3.10.4 Non-Rigid Registration

While the head and bones are rigid, most of the rest of the body is non-rigid, and therefore non-rigid transformations must be employed. Furthermore inter-subject registration requires non-rigid transformations for all parts of the anatomy. These transformations are more complex and varied than rigid ones. The appropriate form of non-rigid

transformation for intra-subject registration is strongly affected by the mechanics of the tissue, i.e. bone and soft tissue and for inter-subject registration by the natural variation in anatomic development. Registration of pre- and post-contrast MR mamograms [114] is an example of an application area that requires non-affine registration.

Chapter 4

DT-MR Registration

In this section we highlight the differences between scalar and diffusion-tensor image registration and discuss the applications of DT registration. The two new challenges that are particular to DT-MRI registration are the need to find appropriate similarity measures for DT data and to change the orientations of the tensors as well as their location when applying image transformations. We review current DT image registration techniques, define various tensor similarity measures required to drive the registration and reorientation strategies for warping DT images and explain how orientation matching (OM) can be used to improve registration.

4.1 Transforming Tensor Images

The application of spatial transformations to DT-MR images is complicated by the fact that DTs contain orientational information, which is affected by the transformation. As a transformation is applied to a DT image, the orientation of the tensors changes. This is illustrated in Figure 4.1, where a 45 degree rotation is applied to a small region (the corpus callosum) of the DT-MR brain-image in Figure 2.2, both with and without reorientation. If the tensors in the transformed image are merely re-located from the corresponding voxels in the untransformed image, without reorientation, the directional structure in the region of the corpus callosum is lost (Figure 4.1b). As discussed by Alexander et al. [5], the DTs themselves need to be rotated similarly and this can be achieved by applying the same rigid rotation matrix to each DT in the im-

age via a similarity transform. Thus, if $\mathbf{R}$ is the rotation matrix representing the image transformation, each D, is replaced by

$$\mathsf{D}' = \mathbf{R}\mathsf{D}\mathbf{R}^T. \tag{4.1}$$

In the case of rigid transformations reorientation is straightforward, as described above. However, reorientation is not straightforward for affine or higher order transformations because we cannot simply replace D at each point in the image by applying the registration transformation to the tensor. When a shear is applied to a DT image, the required reorientation of the tensor not only depends on the applied transformation but also the original orientation of the tensors. Figure 4.2 shows a rectangular tensor field (a) and the transformed tensor field (b) when a horizontal shear is applied. Although the spatial transformation is the same for all voxels, the voxels in the bottom row are not rotated, since the tensors are parallel to the direction of the shearing. However, the tensors in the top row are rotated appropriately. When DT images are scaled, as shown in Figure 4.3, the shape of the tensors must be preserved.

In general, if the transformation is not rigid, D' from Eq. 4.1 has different size and shape (eigenvalues) to D [5]. Alexander et al. [5] assume that image transformations should affect only the orientation of the DT but the shape of the DT remains unchanged. They argue that the apparent diffusion tensor is a property of the tissue microstructure and independent of the shape or extent of the tissue region in the image. A number of reorientation strategies have been proposed for these cases, which try to find rigid transformations that reflect the local reorientation of the image due to the transformation [5].

4.2 Orientation Matching

We use an optimisation procedure to find the transformation that matches tensor size, shape and orientation most closely. This idea is illustrated in Figure 4.4. Correct association of DTs and their orientation requires a very specific transformation and places

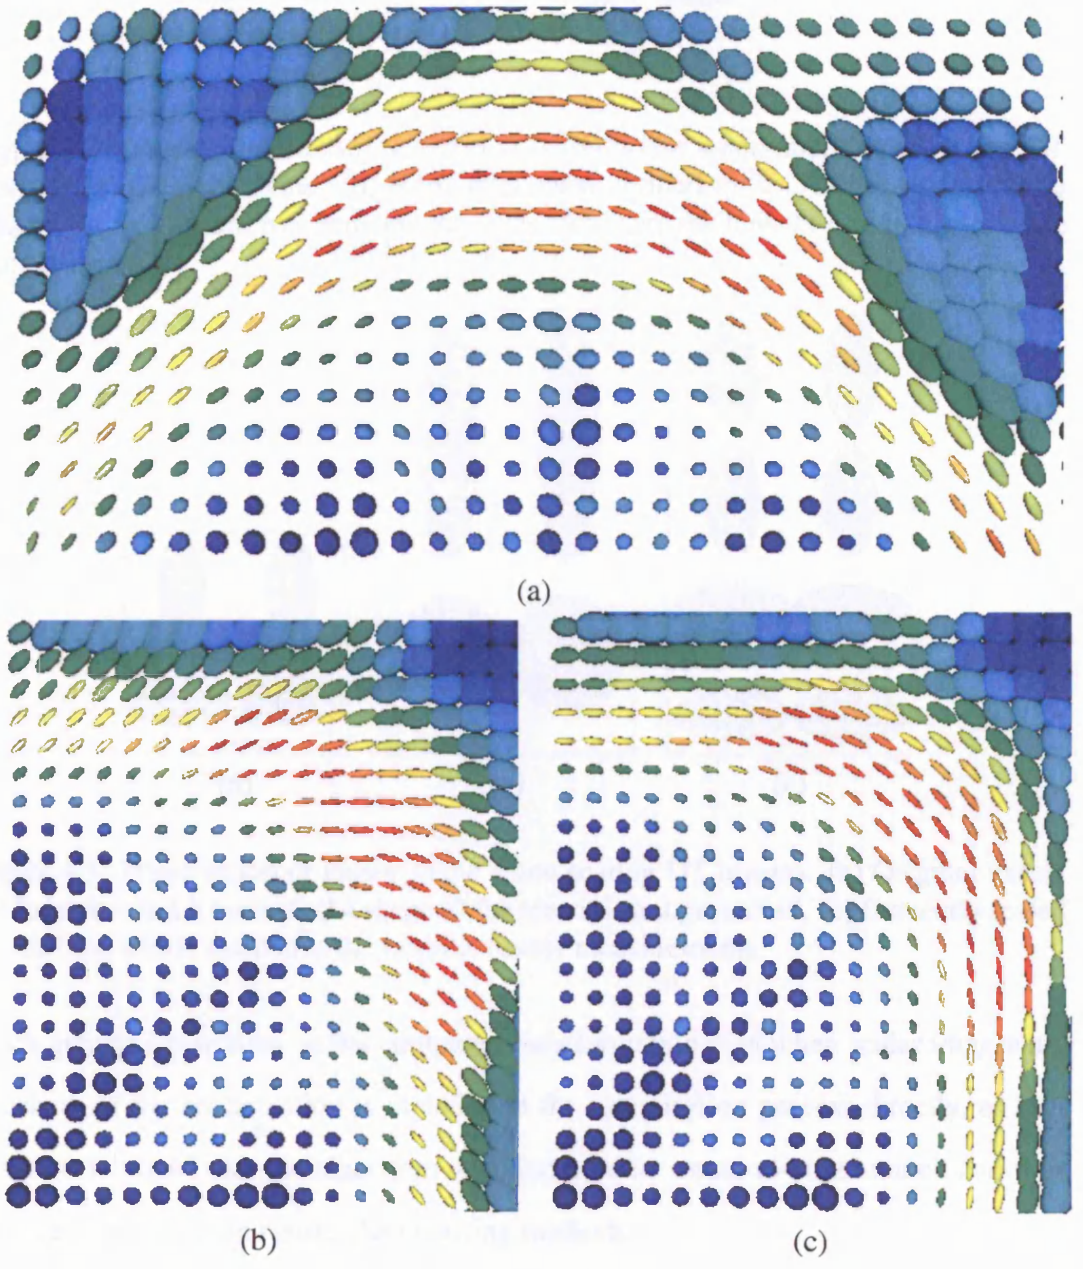

Figure 4.1: The requirement for DT reorientation with image warps. (a) Original image. (b) Tensors are re-located but not re-oriented and are no longer aligned with the fibre pathway; (c) Tensors are correctly re-located and re-oriented.

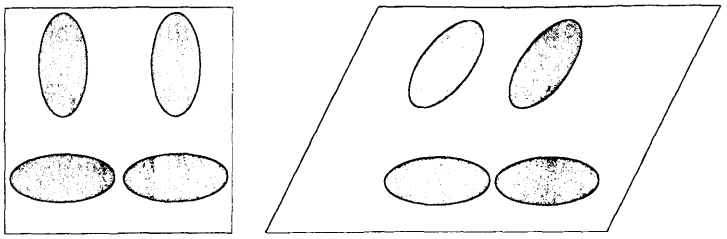

Figure 4.2: The original tensor orientation is important when applying shears to DT images. (a) Original tissue. (b) Horizontal shear applied to (a). The direction of the tensors on the bottom row remains the same, whereas the tensors on the top row are reoriented.

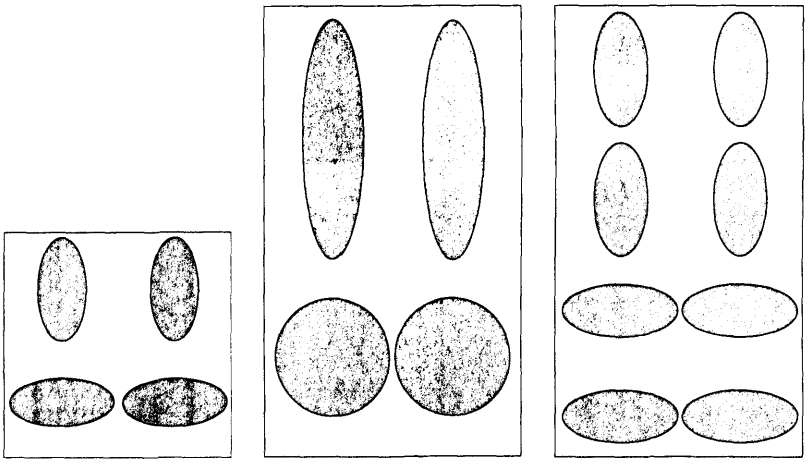

Figure 4.3: Preservation of tensor shape when scaling DT images. (a) Original tissue. (b) Incorrect result because the shape of the tensor is not preserved. (c) Correctly scaled tensor field which maintains the original tensor measurements.

much greater constraints on the computed transformation than when scalar images are matched. If the reorientation is included in the optimisation process directly, as proposed, it is hoped that the algorithm will find a better voxel-to-voxel match and thus provide more accurate results than existing methods.

The motivation for using DT data to drive registration algorithms is to avoid erroneous twists and folds in homogenous regions by searching for transformations that match orientations in these ill-defined regions. Orientation matching is illustrated in figure 4.4. Since the background and foreground regions in Figure 4.4 (a) are homogenous, applying the identity transformation or many other transformations, such as the transformation in Figure 4.4 (b), transform Figure 4.4 (a) to itself. However, if we

consider a similar image with homogenous regions of tensors, as in Figure 4.4 (c), the transformation in Figure 4.4 (b) causes local changes in the image orientation, which the identity transformation does not. Orientation matching updates the orientation of the diffusion tensors to reflect changes in image orientation, as illustrated in Figure 4.4 (d).

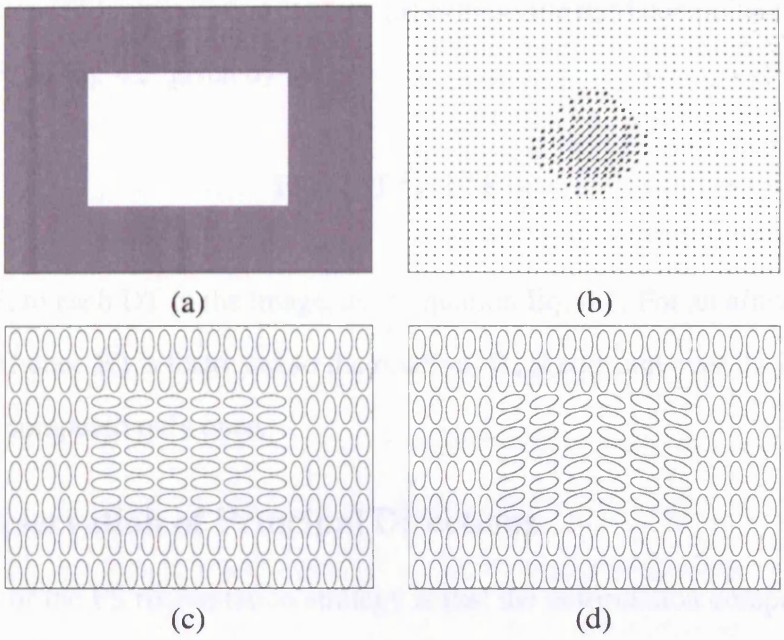

Figure 4.4: Orientation matching. Image (a) shows a scalar image that does not change under the transformation of (b). Image (c) shows a tensor image that is homogenous in the same regions as image (a). Image (d) shows the image in (c) after the transformation in (b). Adapted from [60].

4.3 Reorientation Techniques

This section outlines the reorientation strategies currently in the literature. Alexander et al. [5] present several approaches for the estimation of an appropriate reorientation of the data from the local displacement field, which describes the image transformation. The finite strain (FS) [5], preservation of principal direction (PPD) [5] and Procrustean tensor reorientation [143] strategies will be discussed in this section.

4.3.1 Finite Strain

Any non-singular linear transformation $\mathbf{F}$ can be decomposed into a rigid component $\mathbf{R}$ and a deformation component $\mathbf{U}$ [88] where

$$\mathbf{F} = \mathbf{UR}. \tag{4.2}$$

The finite strain (FS) reorientation strategy [5] extracts the rigid component of the transformation, $\mathbf{R}$, in Eq. 4.2, given by [88]

$$\mathbf{R} = (\mathbf{F}\mathbf{F}^T)^{-1/2}\mathbf{F}, \tag{4.3}$$

and applies it to each DT in the image, as in equation Eq. 4.1. For an affine transformation $\mathbf{F}$ (and hence $\mathbf{R}$) is fixed and so the rotation, $\mathbf{R}$, is constant over the entire image and need be computed only once.

4.3.2 Preservation of Principal Directions

A drawback of the FS reorientation strategy is that the deformation component, $\mathbf{U}$, of the affine transformation is discarded and does not contribute to the estimated reorientation [5]. When the deformation component of the image transformation is small, FS reorientation is a good approximation to the preservation of principal directions (PPD) method. For a global affine transformation, the FS strategy uses the same rotation for every tensor in the image. Alexander et al. [5] note that the required rotation depends on the original orientation of the tensor. The reorientation that occurs as a consequence of transformations such as shearing or stretching depends on the original orientation of the image structure, as illustrated in Figure 4.2, where the resulting change in orientation of the fibres can be seen to be dependent on their original direction [5].

The PPD strategy directly examines the effects of the transformation on the eigenvectors of the DT at each point. The PPD algorithm finds the rotation that maps $\mathbf{e}_1$, the principal eigenvector of $\mathbf{D}$, exactly to $\mathbf{Fe}_1$ and maps $\mathbf{e}_2$, the second eigenvector, as close as possible to $\mathbf{Fe}_2$. The strategy assumes that the directionality of the tissue

structure corresponds to the direction of the eigenvectors of the DT [17]. With PPD, the DT reorientation is not constant over the image for general affine transformations, unlike the FS strategy, and a separate $\mathbf{R}$ must be computed at each voxel. The PPD method preserves the principal direction of the DT through the transformation as well as the plane of the first two eigenvectors and so is effective for both prolate and oblate DTs as well as for intermediate DTs for which $\lambda_1 > \lambda_2 > \lambda_3$ [5].

The PPD strategy proposed in [5] directly examines the effects of the transformation on the eigenvectors of the DT at each point. We give an algorithm for implementing the PPD algorithm introduced in [5]. We can simplify the PPD algorithm from the original outline in [5] by considering only the eigenvectors of the diffusion tensor. Suppose we want to transform a source image I_1 by a transformation $\mathcal{T}$. At a point $\mathbf{x}$ in I_1, the diffusion tensor $\mathsf{D} = \mathbf{E}\mathbf{\Lambda}\mathbf{E}^T$, where $\mathbf{E}$ is the matrix of eigenvectors $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$ and $\mathbf{\Lambda} = \mathrm{diag}(\lambda_1, \lambda_2, \lambda_3)$ is the diagonal matrix of eigenvalues with $\lambda_1 \geq \lambda_2 \geq \lambda_3$. To compute the diffusion tensor D' at the corresponding position $\mathcal{T}(\mathbf{x})$ in the transformed image $I_2 = \mathcal{T}(I_1)$, we determine the reoriented eigenvectors $\mathbf{E}' = (\mathbf{e}'_1, \mathbf{e}'_2, \mathbf{e}'_3)$ and retain the eigenvalues so that $\mathsf{D}' = \mathbf{E}'\mathbf{\Lambda}\mathbf{E}'^T$. To find the reoriented eigenvectors, we need the Jacobian matrix $\mathbf{J} = d\mathcal{T}/d\mathbf{x}$ of the transformation at $\mathbf{x}$. We compute: $\mathbf{n}_1 = \mathbf{J}\mathbf{e}_1/|\mathbf{J}\mathbf{e}_1|$ and $\mathbf{n}_2 = \mathbf{J}\mathbf{e}_2/|\mathbf{J}\mathbf{e}_2|$, then

$$P(\mathbf{n}_2) = \mathbf{n}_2 - (\mathbf{n}_2 \cdot \mathbf{n}_1)\mathbf{n}_1.$$

Finally,

$$\mathbf{e}'_1 = \mathbf{n}_1, \tag{4.4}$$

$$\mathbf{e}'_2 = P(\mathbf{n}_2)/|P(\mathbf{n}_2)| \tag{4.5}$$

$$\mathbf{e}'_3 = \mathbf{e}'_1 \times \mathbf{e}'_2. \tag{4.6}$$

For an affine transformation, the Jacobian $\mathbf{J} = \mathbf{A}$, cf Eq. 3.4, which is the linear part of the affine transformation.

4.3.3 Procrustean Reorientation

A more recently proposed method for estimating the rotation matrix, $\mathbf{R}$, is Procrustean tensor reorientation [143]. Xu et al. [143] claim that the principal direction of the tensor is only a noisy observation of the fibre direction and that directly using this noisy observation reduces the accuracy of reorientation. The Procrustean approach [65] for tensor reorientation [143] assumes that the true underlying fibre orientation follows a statistical distribution that can be estimated from tensor measurements in the neighbourhood of the voxel under consideration. This neighbourhood approach is potentially vulnerable to the neighbourhood overlapping regions with significantly different fibre orientation. The Procrustean reorientation is computationally more complex than the PPD method and it would be interesting to compare the two reorientation techniques to determine if the increased time complexity of the Procrustean method results in a significantly more accurate reorientation.

4.4 DT Similarity Measures

In order to obtain the best possible match between two DT images, it is essential to use an appropriate similarity measure to drive the image registration. A numerical measure of similarity is obtained by comparing the data values at corresponding image locations. For scalar images the simplest approach is to use the difference in scalar intensity at corresponding image locations but many others have been proposed that generally produce better results [70]. In the case of DT image matching, a comparative measure of similarity between diffusion tensors is required to drive the registration. To exploit the data fully, we need a similarity measure that exploits all aspects of D, i.e. the size, shape and orientation of the tensor. The simplest kinds of comparisons we can make between DTs are by looking at the difference in scalar indices derived from D, such as the relative anisotropy ν_r or $\langle \mathsf{D} \rangle$:

$$\delta_1(\mathsf{D}_1, \mathsf{D}_2) = -|\nu_r(\mathsf{D}_1) - \nu_r(\mathsf{D}_2)| \tag{4.7}$$

or

$$\delta_2(\mathsf{D}_1, \mathsf{D}_2) = -|\mathrm{Tr}(\mathsf{D}_1)/3 - \mathrm{Tr}(\mathsf{D}_2)/3|. \tag{4.8}$$

We might also create combined measures of this style such as:

$$\delta_1\delta_2(\mathsf{D}_1, \mathsf{D}_2) = \xi\delta_1(\mathsf{D}_1, \mathsf{D}_2) + \psi\delta_2(\mathsf{D}_1, \mathsf{D}_2), \tag{4.9}$$

where ξ and ψ are weighting parameters that balance the effect of δ_1 and δ_2.

Similarity measures based on the derived scalar indices do not use the orientational information that is contained in a tensor. The following measures, numbered out of order for consistency in Chapter 5, are derived from the full tensor matrix and are sensitive to the relative orientation of the tensors:

- The tensor scalar product is the sum of the products of the nine corresponding elements of the tensors [17, 4], given by:

$$\delta_5(\mathsf{D}_1, \mathsf{D}_2) = \mathsf{D}_1 : \mathsf{D}_2 = \mathrm{Tr}(\mathsf{D}_1\mathsf{D}_2) = \sum_{i=1}^{3}\sum_{j=1}^{3} \lambda_{1i}\lambda_{2j}\mathbf{e}_{1i} \cdot \mathbf{e}_{2j} \tag{4.10}$$

  The tensor scalar product is acutely sensitive to the relative size of the two tensors.

- The tensor difference is a Euclidean measure between the nine corresponding elements of the two DT matrices. For two DTs, D_1 and D_2,

$$\delta_3(\mathsf{D}_1, \mathsf{D}_2) = -\sqrt{(\mathsf{D}_1 - \mathsf{D}_2) : (\mathsf{D}_1 - \mathsf{D}_2)}. \tag{4.11}$$

4.5 Current DT Registration Techniques

Many diffusion tensor similarity measures have been proposed including those that generalise existing scalar measures [115, 112] and those defined specifically for DT images [4, 2]. Alexander and Gee [4] extend an elastic matching algorithm, originally designed for scalar images, to work with DT-MR images. They describe a number of similarity

measures, derived from the full tensor, that are sensitive to orientational information, including trace and anisotropy differences and direct comparisons such as the tensor difference and tensor scalar product, defined in §4.4. Alexander and Gee [4] show that better registration performance is obtained from the use of the tensor difference (Eq. 4.11) than the tensor scalar product (Eq. 4.10) but little improvement is shown between the tensor difference and simpler measures, such as relatively anisotropy and trace difference. However, their registration algorithm does not include tensor reorientation in the optimisation and therefore does not match orientations.

Ruiz-Alzola et al. [115] have also developed a scheme for DT image registration. They have taken a different approach to that proposed here. They have attempted to generalise the idea of registration in such a way that certain similarity measures can be derived naturally - at least if a template matching scheme is used. Similar to Alexander and Gee [4], they do not explicitly embed tensor reorientation into their registration scheme, as we are proposing. They perform template matching without reorientation and only reorient the tensors (using a Finite Strain approach [5]) when the computed transformation is applied.

Guimond et al. [67] generalise their single channel demons registration algorithm [68] to multiple channel data with finite strain reorientation [5]. In the case of inter-subject registration, quantifying the accuracy of a given method is difficult. In [67] they opt for a qualitative method based on the visual inspection of tensors from the target and source images after registration. They argue that using the tensor eigenvalues during registration provides results that are comparable to the ones obtained using the full tensor, while not significantly increasing computational time. Park and colleagues [100] also use orientation matching to drive the registration and they show that full-tensor registration with PPD reorientation finds the best performance in effectively normalising the tract morphology and tensor orientation. However, Park et al. [100]

(a) (b)

Figure 4.5: (a) Two orthogonal tensors, where D_1 = diag(10, 2, 2) and D_2 = diag(2, 10, 2). (b) The two orthogonal tensors in (a) rotated by forty-five degrees.

use the sum of square differences of the six independent elements of the DT:

$$\delta_p(\mathsf{D}_1, \mathsf{D}_2) = \sqrt{\sum_{i=1}^{6} (\mathsf{D}_{1i} - \mathsf{D}_{2i})^2}. \tag{4.12}$$

as their similarity measure. This similarity measure is similar to the tensor difference defined in Eq. 4.11, which is the sum of square differences of the *nine* elements of the DT, but unlike δ_3, δ_p is not rotationally invariant, which is a serious drawback to their work. To see this problem, we calculate the tensor difference of the two tensors depicted in Figure 4.5(a). We get $\delta_3 = \delta_p = 8\sqrt{2}$, since the off-diagonal elements are all zero. However, if we calculate the tensor difference of the two tensors in Figure 4.5(b), where a forty-five degree rotation has been applied to the tensors, we get $\delta_3 = 8\sqrt{2}$ but $\delta_p = 8$. Park et al. [100] introduce evaluation measures based on alignment of extracted fibre tracts from the registered data. They acknowledge that errors in fibre tracking affect the evaluation of the registration but suggest that such errors are evenly distributed for all comparison procedures.

Curran and Alexander [40] present a voxel similarity based registration scheme and use an optimisation procedure to find the transformation that matches tensor size, shape and orientation most closely. We adjust the tensor orientations iteratively during registration according to the transformation using the FS or PPD method. We show that orientation matching improves registration results using direct optimisation

for intra-subject simulations. In [41], Curran and Alexander examine the directional coherence of the transformed-source and target images in two hand-defined regions. Experimentation with the direct optimisation scheme used in [40] within inter-subject brain-registration tasks reveals difficulties with local minima (see Chapter 6). An additional contribution of this work is the highlighting of the local minimum problem in orientation matching and the need to use global optimisation techniques, despite the dramatic increase in computation times, to overcome these difficulties. Curran and Alexander [42] show that the relative-anisotropy difference, a scalar-derived measure, finds the best affine registration when the Euclidean distance is compared between corresponding landmarks in the source and target images. However, the tensor difference, a full tensor similarity measure, performs almost as well. We cover this work in greater detail in later chapters.

Duda et al. [49] estimate a mapping between two DT-MRI slices by minimising an energy function. This function consists of three terms. One term compares tensor eigenvalues to quantify shape differences between tensors. Another ensures a smooth mapping. The third term is a neighbourhood similarity measure that examines tensor orientation patterns.

Rhode et al. [112] propose another similarity measure, which is based on mutual information. Rhode et al. [111] incorporate the multivariate version of mutual information into an existing non-rigid adaptive radial-bases registration algorithm [109], performing finite strain reorientation only after affine registration and after the non-linear registration is complete, assuming that after affine registration most structures in the source and target images will have similar orientations. They investigate if the inclusion of all DT parameters during registration provides an additional improvement over multi-channel registration based on rotationally invariant measures. They compare two multi-channel configurations. The scalar similarity measure includes the trace and relative anisotropy of the tensor. The tensor configuration contains the six unique diffusion tensor elements, as well as the amplitude parameter of the model. While the results of the simulation and cardiac-gated DT image registration experiments showed

improved accuracy when using the tensor channel combination, the results of the non-rigid inter-subject registrations did not show a significant improvement in accuracy compared with the results obtained using the scalar channel combination. This finding is similar to Guimond et al. [67], though the methods used were entirely different. Rhode et al. [111] note that because of the lack of a gold standard for inter-subject registration (the results were only visually inspected) their evaluation methods for this experiment are less reliable.

Verma and Davatzikos [132] propose a method for matching two DT images by detecting correspondences using Gabor features [90, 119]. They characterise each voxel in a DT image by a feature vector defined using Gabor filters with different scales and frequencies that are all oriented along the dominant direction of the tensors in a neighborhood about the voxel.

Zhang et al. [146] present a framework for multi-resolution piecewise affine registration of DT images that allows registration with finite strain reorientation [5] to be computed using analytical derivatives. They match general apparent diffusion coefficient (ADC) profiles rather than diffusion tensors. Zhang et al. [146] expect that a metric formulated using the full apparent diffusion coefficient profile performs as well or better than full tensor metrics. In a similar way, we would expect full tensor similarity measures to perform better than scalar measures.

A related set of registration algorithms associated with DT-MRI aligns the set of diffusion-weighted image volumes prior to reconstruction of the tensor. This thesis addresses a different problem, which is registration of separate data sets after tensor reconstruction. However, we include a brief review of these other registration techniques here for completeness.

Eddy current distortions in diffusion-weighted images [89] are corrected for by registering the diffusion-weighted images to the unweighted ("$b = 0$") image using affine image transformations [77, 110]. In diffusion-weighted imaging, the $b = 0$ image has a significantly different appearance from the diffusion-weighted images, and it is appropriate to view registration of diffusion weighted images as a multimodality

problem [77]. However, it is possible that non-linear warping of the images can occur and it is not clear how well conventional affine registration methods perform in cortical regions in newborn patients.

Symms et al. [124] use AIR [139] to perform a slice by slice registration of each diffusion weighted image to the first $b = 0$ acquisition using two-dimensional affine transformations. Nielson and colleagues [96] investigate affine and second-order polynomial registration of the diffusion-weighted images to the unweighted ($b = 0$) image in newborns, using mutual information. Nielson et al. [96] show that nonlinear distortions exist in echo-planar DTI data sets and that polynomial registration systematically reduces the presence of erroneous regions of high fractional anisotropy (FA) and that such artifacts can be identified by visual inspection of FA maps resulting from affine and polynomial registration. Their results demonstrate the need for non-affine registration for quantitative DTI. Nielson et al. [96] use the downhill simplex method [107] and emphasise the need to search for an image transformation using a "greedy" search algorithm, as the possibility that a local and not the global minimum is obtained increases for higher-order transformations [96]. Anderson et al. [7] propose a model based approach where they combine affine registration of the diffusion weighted images with reconstruction of the diffusion tensor into a single optimisation.

4.6 DT Registration Applications

In recent years DT-MRI has emerged as a clinically important magnetic resonance technique capable of detecting subtle changes in tissue microstructure that may go unnoticed in conventional T1- and T2-weighted MRI scans [14]. The development of techniques for spatial normalisation of DT-MR images also allows for cross-population studies of this data to be performed. The potential clinical applications include mapping brain connectivity and diagnosis of stroke and multiple sclerosis [76].

Statistical atlases [63, 76, 134, 145] of brain structure provide a means of characterising normal and pathological neuroanatomy of populations. They represent the characteristics of an entire population and therefore can be used to study anatomical

variability. Mapping brain connectivity assists in understanding how brain connections change with development, aging and disease. Changes in white matter structure have been reported in diseases and disorders, such as Alzheimer's Disease [133] and schizophrenia [138].

Traditional MRI techniques provide no information about brain connectivity; however, DT-MRI has shown changes in brain connectivity in studies of aging and development [95, 98]. Population differences can then be quantified by statistically comparing the respective spatially normalised fibre maps in order to identify regions in which brain connectivity is affected by disease. However, group averaging studies can only represent population trends, and not individual deviations from these trends.

An important aspect of generating statistical atlases is spatial normalisation, which is a procedure to remove gross inter-individual anatomical variability. These methods can be volume-based, which match intensities across subjects [30, 57, 58, 63], and feature-based, which extract features such as structure boundaries, sulci and gyri from MR images and subsequently determine three dimensional transformations that match these features [43, 51].

Chapter 5

Methods and Implementation

This chapter describes the design and implementation of a registration algorithm that uses DT-MR data. We describe implementations for rigid, affine and higher-order polynomial registration.

5.1 DT-MR Voxel Similarity Registration Scheme

The basic registration scheme outlined in §3.3.4 (Eq. 3.1) is appropriate for intensity images, but when the measurement at each point is more complex than just a scalar, e.g. a vector or tensor, we need to generalise this scheme. We propose a generalised registration scheme, which is similar to Eq. 3.1, but now write the similarity measure minimised to achieve registration as

$$M(\mathcal{T}) = \int_{I_1 \cap I_2} S(\mathcal{T}(I_1(\mathcal{T}(\mathbf{x}))), I_2(\mathbf{x}))\mathrm{dx}. \tag{5.1}$$

This generalised formulation affords the possibility that the transformation, $\mathcal{T}$, affects the actual values in the image as well as their location. Figure 5.2 shows schematically the flow of control of the registration algorithm. The algorithm consists of four main stages:

1. an optional data preparation stage that preprocesses both I_1 and I_2, for example we might smooth both images with a Gaussian kernel,

2. a function that spatially transforms I_1, using interpolation by transformation, $\mathcal{T}$,

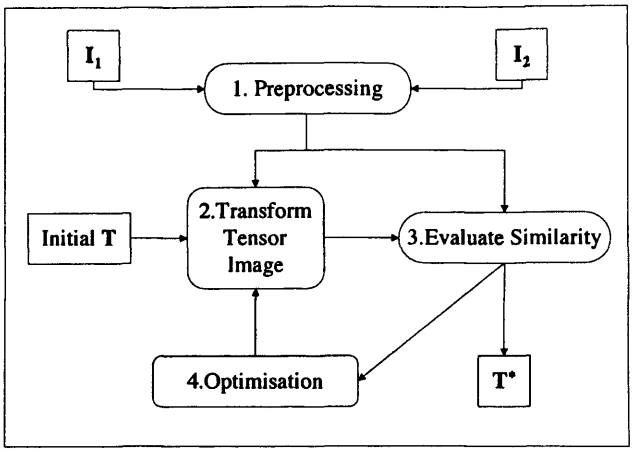

Figure 5.1: Schematic representation of the registration algorithm. The initial transformation $\mathcal{T}$ is iteratively updated to maximise the similarity of I_1 (source image) and I_2 (target image) until the best mapping ($\mathcal{T}^*$) is found.

and reorients each tensor appropriately,

3. a function that measures the similarity of I_1 and I_2,

4. an optimisation algorithm that searches for the optimal transformation $\mathcal{T}^*$, which maximises the similarity.

5.2 Preprocessing

As an optional preprocessing step, we apply Gaussian smoothing to the source and target images prior to registration. Gaussian smoothing outputs a weighted average of each voxel's neighbourhood, with the average weighted more towards the central voxels. The Gaussian smoothing operator

$$G(x,y,z) = \frac{e^{-(x^2+y^2+z^2)}}{2\sigma^2}, \tag{5.2}$$

where x, y and z are the image coordinates relative to the voxel at the centre of the neighbourhood and σ is the level of smoothing that is applied. The smoothing kernel has infinite support but in practice we truncate it at some distance from the central voxel; this kernel width may also be adjusted.

5.3 Similarity Measures

For tensor images, we can choose S in Eq. 5.1 to be any of the similarity indices described in §4.4. We also propose the following new similarity measures.

The tensor difference can be normalised, in order to emphasise differences in shape and orientation, in the following way:

$$\delta_4(\mathbf{D}_1, \mathbf{D}_2) = -\sqrt{(\mathbf{D}_1 - \mathbf{D}_2) : (\mathbf{D}_1 - \mathbf{D}_2)} / \sqrt{\mathrm{Tr}(\mathbf{D}_1)\mathrm{Tr}(\mathbf{D}_2)}. \tag{5.3}$$

and in a similar way to δ_3, the tensor scalar product can be normalised

$$\delta_6(\mathbf{D}_1, \mathbf{D}_2) = \frac{\mathbf{D}_1 : \mathbf{D}_2}{\mathrm{Tr}(\mathbf{D}_1)\mathrm{Tr}(\mathbf{D}_2)}. \tag{5.4}$$

We also propose a principal-direction difference measure given by:

$$\delta_7(\mathbf{D}_1, \mathbf{D}_2) = -\cos^{-1} |\mathbf{e}_{11} \cdot \mathbf{e}_{12}|. \tag{5.5}$$

where $\mathbf{e}_{11}$ and $\mathbf{e}_{12}$ are the principal eigenvectors of the two diffusion tensors.

5.4 Data Acquisition

Diffusion-weighted data were acquired in six healthy subjects. The diffusion protocol consisted of a single-shot diffusion-weighted echo-planar imaging sequence. The diffusion acquisition parameters were: TE 95ms; acquisition matrix 96×96 reconstructed as 128×128; f.o.v. = $220 \times 220\text{mm}^2$; pixel resolution = $2.5 \times 2.5\text{mm}^2$ reconstructed to $1.7 \times 1.7\text{mm}^2$; 3 b = 0smm^{-2} images; 60 diffusion weighted images, with gradients applied along 60 optimised diffusion directions, diffusion times of δ = 32ms and Δ = 40ms, and max gradient amplitude of $22 \times 10^{-3}\text{Tm}^{-1}$, giving a maximum b-factor b = 1000smm^{-2}; 42 contiguous axial slices; 2.5mm slice thickness; diffusion data acquisition time = 20-30 minutes depending on subject's heart rate (cardiac gating [137] was used to reduce motion artefacts due to pulsation of blood and cerebrospinal fluid).

5.5 Step Size

The step size controls the size of the array of voxels that contribute to the similarity. The experiments in this thesis look at four possible step sizes: 8, 4, 2 and 1 - which for the data sets used here (described in §5.4):

- Step size one, considers all of the voxels ($128 \times 128 \times 42$).
- Step size two, considers a quarter of the voxels ($64 \times 64 \times 42$).
- Step size four, considers a sixteenth of the voxels ($32 \times 32 \times 42$).
- Step size eight, considers a sixty-fourth of the voxels ($16 \times 16 \times 42$).

5.6 Practical Considerations

In practice we include the following heuristics to improve and stabilise the algorithm:

- We use the brain extraction tool (BET) [121] to segment foreground and background and then set all background voxels to identity tensors.
- We only calculate the similarity index for the region where the volumes overlap. In practice, we normalise the similarity, M, in Eq. 5.1 by the number, N, of non-background voxels in the overlap domain, $I_1 \cap I_2$.
- If the percentage of overlap of the source and target images is $<10\%$, a low similarity value well outside the range of plausible values is returned to avoid extreme transformations, as described in §3.6.6.

5.7 Optimisation Techniques

We implement two optimisation algorithms. The first is Powell's method [107], described in §3.5, which is a local search method. We also implement a global optimisation technique that we call gradient annealing.

5.7.1 Powell's Optimisation

The basic idea behind Powell's direction set method [107] is to break the N-dimensional minimisation down into N separate one-dimensional (1D) minimisation problems. Then, for each 1D problem a line search is implemented to find the local minimum within a given range. Furthermore, on subsequent iterations an estimate is made of the best directions to use for the 1D searches. This enables it to navigate efficiently along narrow valleys when they are not aligned with the axes. Here we follow the implementation in [107].

5.7.2 Gradient Annealing Optimisation

To avoid being trapped in local minima and thus sub-optimal solutions, we also implement simulated annealing [107] to optimise the starting point for Powell's method [107]. We call the combined method *gradient annealing* (Figure 5.2). The simulated annealing algorithm is computationally very expensive but the optimisation process is accelerated by combining a fast local optimisation method with a global optimisation technique. We note that finding optimal solutions is particularly important here, since the aim of the work is to determine which approach fundamentally provides the best image match. Spurious optima at local minima will not demonstrate the full power of a particular matching technique.

We described the concept of simulated annealing (SA) in §3.5. In this section we discuss the implementation of three of the four elements required by the simulated annealing algorithm, i.e. an objective function, a system state and a generator of random changes in the configuration. We defer discussion of the control parameter (analog of temperature) and the annealing schedule until Chapter 7.

For gradient annealing the energy state in Algorithm 1 is the starting point for Powell's direct optimisation. The objective function can use any of the similarity indices defined in §4.4 or §5.3, summed over the overlap volume. The system state is a vector $\mathbf{p}$ containing the transformation parameters we wish to optimise. We modify the parameters by adding $\log(T+1) * r$ to the current $\mathbf{p}$, where r is a random number

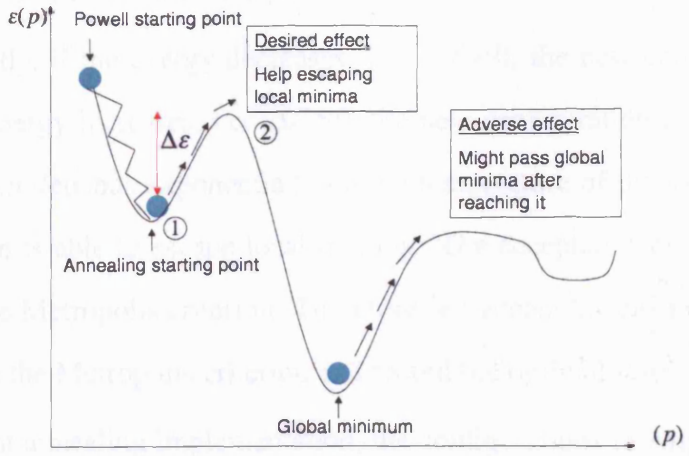

Figure 5.2: Gradient annealing in one dimension, where $\Delta\mathcal{E}$ is the change in energy, i.e. jumping from 1 to 2, required to escape local minima and p is the parameter to be optimised.

between 0 and 1 and T is the current temperature.

Algorithm 1: The simulated annealing algorithm.
Input: A set of possible configurations, an objective function, a random number r between 0.0 and 1.0 and a temperature control parameter T.
Output: Final objective function or energy state $\mathcal{E}_0$ and the optimised set of configurations.

(1) **repeat**
(2) **generate** a new configuration by perturbing the current configuration
(3) **if** $\mathcal{E}_1 < \mathcal{E}_0$
(4) **accept** $\mathcal{E}_1$ and new configuration
(5) **else if** $e^{-(\mathcal{E}_1-\mathcal{E}_0)/T} > r$
(6) **accept** $\mathcal{E}_1$ and new configuration
(7) **else**
(8) **reject** $\mathcal{E}_1$ and new configuration
(9) **end**
(10) **decrease** T
(11) **until** system is *frozen*

The SA algorithm, illustrated in Algorithm 1, consists of a pair of nested loops. The outer-most loop sets the temperature and the inner-most loop runs a Metropolis simulation at that temperature. The sampled configuration is accepted with probability depending on the difference $\Delta\mathcal{E}$ between the current objective function $\mathcal{E}_0$ and the new objective function $\mathcal{E}_1$. If $\mathcal{E}_1<\mathcal{E}_0$, we accept the new state. If $\exp(-\Delta\mathcal{E}/T)>r$, we

accept the new state, otherwise we reject it.

Consequently, if the energy decreases, i.e. $\Delta\mathcal{E}<0$, the new configuration is accepted. If the energy increases, i.e. $\Delta\mathcal{E}>0$, the new configuration is accepted with a probability which depends exponentially on the temperature of the system. As a result the algorithm is able to escape local minima. The acceptance probability is often referred to as the Metropolis criterion. Therefore, we accept or reject each set of parameters based on the Metropolis criterion and record the optimal state.

For gradient annealing implementation, the configurations in Algorithm 1 are the starting parameters for Powell's optimisation and the energy states, $\mathcal{E}_1$ and $\mathcal{E}_2$ are the values of the objective function at the end of the Powell optimisation.

5.8 Reorientation

For each of the reorientation strategies described in this section we evaluate $M(\mathcal{T})$ by iterating through all or a fraction of the voxels in the image and optimise $\mathcal{T}$ using Powell's optimisation algorithm [107] or the gradient annealing optimisation method.

The simplest possible strategy is to ignore the problem of applying reorientation to the DTs in the image, or equivalently to set $\mathbf{R}$ in Eq. 4.1 equal to the identity at every voxel. The value at each point in the transformed image is copied directly from the corresponding position in the untransformed image, as we would do for a scalar image. The no reorientation (NR) strategy is included here as a control against which to measure the effectiveness of the finite strain (FS) and preservation of principal direction (PPD) reorientation strategies, outlined in §4.3.

When the image transformation $\mathcal{T}$ is rigid, $\mathcal{T} = \{\mathbf{R}, \mathbf{t}\}$, where $\mathbf{R}$ is a rotation and $\mathbf{t}$ is a translation, as discussed in §4.1, we can use $\mathbf{R}$ to reorient the tensors directly as in Eq. 4.1. Thus we optimise

$$M(\mathcal{T}) = \int_{I_1 \bigcap I_2} S\left(\mathbf{R}^T(\mathbf{x})\mathsf{D}_1(\mathcal{T}(\mathbf{x}))\mathbf{R}(\mathbf{x}), \mathsf{D}_2(\mathbf{x})\right) \mathrm{d}\mathbf{x} \tag{5.6}$$

where $\mathsf{D}_i(\mathbf{x})$ is the DT at position $\mathbf{x}$ in image I_i.

Affine transformations are complicated by the fact that the reorientation matrix is harder to extract from $\mathcal{T}$. We use either the FS or PPD reorientation strategies, as described in §4.3. If we use the FS reorientation strategy evaluating Eq. 5.6 is analytic and differentiable unlike PPD, which is algorithmic.

For higher-order polynomial transformations the appropriate reorientation must be estimated at each point from the local displacement gradient $\mathbf{J}_{\mathcal{T}}$. We obtain a local affine transformation of the higher-order transformation by differentiating $\mathcal{T}$ with respect to x:

$$\mathbf{J}_{\mathcal{T}}(\mathbf{x}) = \frac{d\mathcal{T}}{d\mathbf{x}} \tag{5.7}$$

where

$$\mathbf{J}_{\mathcal{T}} = \begin{pmatrix} \frac{\partial \mathcal{T}_x}{\partial x} & \frac{\partial \mathcal{T}_x}{\partial y} & \frac{\partial \mathcal{T}_x}{\partial z} \\ \frac{\partial \mathcal{T}_y}{\partial x} & \frac{\partial \mathcal{T}_y}{\partial y} & \frac{\partial \mathcal{T}_y}{\partial z} \\ \frac{\partial \mathcal{T}_z}{\partial x} & \frac{\partial \mathcal{T}_z}{\partial y} & \frac{\partial \mathcal{T}_z}{\partial z} \end{pmatrix} \tag{5.8}$$

and $\mathcal{T}(x, y, z) = (\mathcal{T}_x, \mathcal{T}_y, \mathcal{T}_z)$. We then use the FS or PPD method to extract the reorientation matrix from the local affine transformation.

5.9 Transformation Parameterisation

In this section we outline how we parameterise the transformations.

5.9.1 Rigid Transformation Parameterisation

The rigid transformation consists of a rotation $\mathbf{R}$ and a translation $\mathbf{t}$, thus $\mathcal{T}(x) = \mathbf{R}\mathbf{x} + \mathbf{t}$. We parameterise rotations in three dimensions (3D) with three angles: α and β specify the axis of rotation $(r_x, r_y, r_z) = (\cos\alpha \sin\beta, \sin\alpha \sin\beta, \cos\beta)$ and θ is the angle of rotation. The rotation matrix, $\mathbf{R}$, can be obtained from the general formula for rotation through angle θ about an axis given by the unit vector (r_x, r_y, r_z), [59]:

$$\mathbf{R} = \begin{bmatrix} r_x^2(1-c_\theta)+c_\theta & r_x r_y(1-c_\theta)-r_z s_\theta & r_x r_z(1-c_\theta)+r_y s_\theta \\ r_x r_y(1-c_\theta)+r_z s_\theta & r_y^2(1-c_\theta)+c_\theta & r_y r_z(1-c_\theta)-r_x s_\theta \\ r_x r_z(1-c_\theta)-r_y s_\theta & r_y r_z(1-c_\theta)+r_x s_\theta & r_z^2(1-c_\theta)+c_\theta \end{bmatrix} \tag{5.9}$$

where cos θ and sin θ are abbreviated as c_θ and s_θ, respectively. To remove any dependence of the translation parameters on the rotation parameters, we ensure that the rotation is always performed about the centre of the image. The translation parameters are then an offset which we add subsequently. Thus the translation

$$\mathbf{t} = \mathbf{t_r} + \omega \tag{5.10}$$

where

$$\mathbf{t_r} = \begin{pmatrix} X/2 \\ Y/2 \\ Z/2 \end{pmatrix} - \mathbf{R} \begin{pmatrix} X/2 \\ Y/2 \\ Z/2 \end{pmatrix} \tag{5.11}$$

is the translation required to ensure that the centre $(X/2,\ Y/2,\ Z/2)$ of the image is stationary under the rotation, X, Y and Z are the dimensions of the image and $\omega = (\omega_x, \omega_y, \omega_z)$ is the offset. We optimise the image similarity with respect to the six parameters α, β, θ, ω_x, ω_y and ω_z.

5.9.2 Affine Parameterisation

In the three-dimensional affine transformation, $\mathcal{T}(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{t}$, where $\mathbf{A}$ is any real-valued 3×3 matrix. We use a standard decomposition of the linear matrix, $\mathbf{A}$, which decouples the parameters, splitting the transformation into three rotation, three translation, three scale and three skew parameters. Thus $\mathbf{A} = \mathbf{RSK}$, where we parameterise the rotation, $\mathbf{R}$, as in the previous section,

$$\mathbf{S} = \begin{pmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & s_z \end{pmatrix} \tag{5.12}$$

and

$$\mathbf{K} = \begin{pmatrix} 1 & k_1 & k_2 \\ 0 & 1 & k_3 \\ 0 & 0 & 1 \end{pmatrix}. \tag{5.13}$$

As in the rigid transformation case, we decouple the translation parameters by writing

$$\mathbf{t} = \mathbf{t_A} + \omega \tag{5.14}$$

where

$$\mathbf{t_A} = \begin{pmatrix} X/2 \\ Y/2 \\ Z/2 \end{pmatrix} - \mathbf{A} \begin{pmatrix} X/2 \\ Y/2 \\ Z/2 \end{pmatrix} \tag{5.15}$$

is chosen to ensure that the centre of the image is stationary under the transformation

$$\mathbf{x} \rightarrow \mathbf{A}\mathbf{x} + \mathbf{t_A}. \tag{5.16}$$

We optimise the image similarity with respect to the twelve parameters: t_x, t_y, t_z, α, β, θ, s_x, s_y, s_z, k_1, k_2 and k_3.

5.9.3 Higher Order Parameterisation

The higher order polynomial transformations extend naturally from the affine transformation, which is a first-order polynomial transformation. We implement up to order-four, which has parameter matrix

$$\mathbf{A_4} = \begin{bmatrix} & b_{00} & \dots & b_{05} & c_{00} & \dots & c_{09} & d_{00} & \dots & d_{0(29)} \\ \mathbf{A} & b_{10} & \dots & b_{15} & c_{10} & \dots & c_{19} & d_{10} & \dots & d_{1(29)} \\ & b_{20} & \dots & b_{25} & c_{20} & \dots & c_{29} & d_{20} & \dots & d_{2(29)} \end{bmatrix} \tag{5.17}$$

where $\mathbf{A}$ is the 3×3 linear transformation matrix, which we parameterise as in the affine case, the b_{ij} are the second-order polynomial coefficients, c_{ij} are the third-order polynomial coefficients and d_{ij} are the fourth-order polynomial coefficients. We decouple the polynomial parameters from the translation parameters in a similar way to the rigid and affine cases. We optimise the higher-order polynomial coefficients b_{ij}, c_{ij} and d_{ij} directly.

Chapter 6

Registration with Direct Optimisation

This chapter describes experiments performed to test the consistency of direct optimisation when registering a DT-MR brain image with itself (intra-subject) and with another DT-MR brain image (inter-subject). We use Powell's optimisation algorithm [107] to maximise image similarity. Registration is consistent if the optimisation finds the same final transformation, regardless of the initial transformation from which the optimisation starts. We define the hypotheses we wish to test, present the experiments and results, and give preliminary conclusions.

6.1 Hypotheses

We wish to test the following hypotheses.

- **Hypothesis (S1):** "Powell's optimisation finds a consistent transformation for **intra-subject** registration."

- **Hypothesis (S2):** "Powell's optimisation does not find a consistent transformation for **inter-subject** registration."

- **Hypothesis (S3):** "Gaussian smoothing of the source and target images prior to registration improves the consistency of the registration transformation found using Powell's optimisation."

6.2 Direct Optimisation Experiments

The hypotheses defined in §6.1 are tested separately for two transformation models: translations (3 Degrees of Freedom) and rigid transformations (6 DOF) §5.9.1. We perform registration with reorientation for all the experiments in this chapter. We run tests using the relative anisotropy difference, δ_1 and the tensor difference, δ_3, similarity measures in the objective function. In each test, we register a source image with a target image, using one hundred different initial transformations A_j, j = 1, ..., 100, for the optimisation, which gives one hundred registration transformations $\mathcal{T}_j$, j = 1, ..., 100. To choose the starting positions, A_j, for translation-only registration we draw each of the translation parameters, t_x, t_y and t_z, randomly from a uniform distribution on the interval [-5, 5]mm. For rigid registration we draw t_x, t_y and t_z from a uniform distribution on the interval [-5, 5]mm, the rotation angle, θ, from a uniform distribution on $[0, \pi/4]$, and the axis of rotation, $\mathbf{r}$, from a uniform distribution over the unit sphere. We obtain a uniform distribution over the sphere by drawing $\cos\alpha$ uniformly from the interval [-1, 1] and β uniformly from $[0, 2\pi]$, then set $\mathbf{r}$ = ($\cos\alpha \sin\beta$, $\sin\alpha \sin\beta$, $\cos\beta$) [91].

6.2.1 Intra-Subject Experiments

For consistent registration, Powell's optimisation [107] should find the same final transformation, regardless of the starting position of the source and target images. In the experiments in this section, the source and target images are identical to simulate the ideal intra-subject intra-model registration task. Thus, the final transformation should be the identity. In the translation case, t_x, t_y and t_z should all be zero; in the rigid case, the rotation angle, θ, should also be zero.

Experiment: HS1A

- **Objective:** To test **hypothesis S1** (intra-subject registration using direct optimisation is consistent) for **translations** using both the **relative anisotropy difference**, δ_1, and **tensor difference**, δ_3, in the objective function and with step sizes 8, 4, 2 and 1 (§5.5).

- **Results:** We found that for both similarity measures all one hundred T_j were equal to the identity transformation to within tiny numerical errors. Figure 6.1, shows an example initial translation starting position and the corresponding final position where, as expected, the final translation parameters are all zero. We repeated experiment HS1A for step sizes 8, 4, 2 and 1 and using a different image and also found that the final translation parameters were all equal to zero in 100% of trials.

- **Conclusion:** The results support **hypothesis (S1)**.

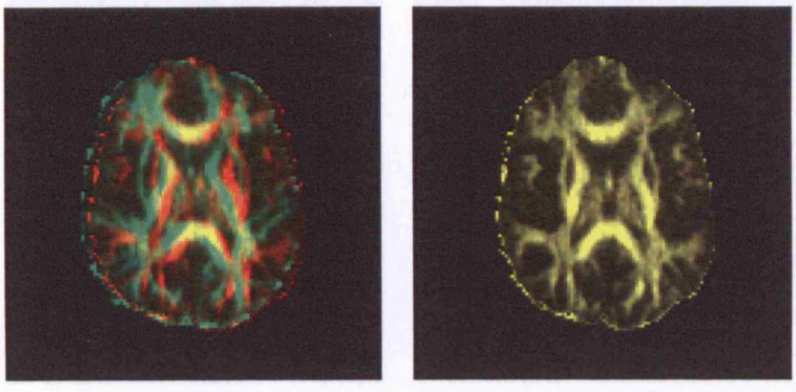

(a) (b)

Figure 6.1: These anisotropy maps illustrate intra-subject (a) pre-registration and (b) post-registration. The target image (red) is superimposed on the source image (green), in (a) and the registered image in (b), so that the image appears yellow where the images are well aligned.

Experiment: HS1B

- **Objective:** To test **hypothesis (S1)** for **rigid** transformations using both δ_1 and δ_3 in the objective function and with step sizes 8, 4, 2 and 1.

- **Results:** The optimisation finds the global minimum of the identity transformation in 98% of cases for both δ_1 and δ_3 with step 8. The four failures occurred at extreme starting positions. The initial translations were all at least 2mm in two translation directions in each failure and the initial rotation angle, θ, was close to $\pi/4$ (the range of θ in the failed cases was 30 to 45 degrees with mean 35

degrees). Figure 6.2 shows an example of one failed registration; in this typical example the final transformation is similar to the starting point.

- **Conclusion:** The results do not support **hypothesis (S1)** for rigid transformations using δ_1 or δ_3 with the range of starting parameters tested. In a further experiment, we reduced the range of the starting positions by drawing the translation parameters from [-2, 2]mm and the rotation angle from $[0, \pi/6]$. For this narrower distribution about the global minimum, we get 100% consistency for both similarity measures.

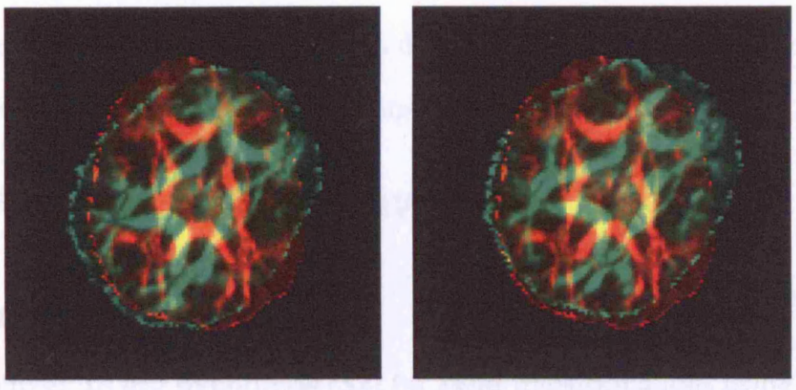

(a) (b)

Figure 6.2: These anisotropy maps illustrate an example failed rigid intra-subject (a) pre-registration and (b) post-registration.

6.2.2 Inter-Subject Consistency

The final transformation for a perfectly aligned inter-subject registration is unknown because the images are different. However, for consistent registration, the final transformation should be the same from different starting points in the optimisation.

Experiment: HS2A

- **Objective:** To test **hypothesis S2** (inter-subject registration using direct registration is **not** consistent) for **translations** using both δ_1 and δ_3 in the objective function with step sizes 8, 4, 2 and 1.

- **Results:** We found that for both similarity measures all one hundred T_j were **not** equal to the identity transformation. Figure 6.3 shows the consistency for the

translation parameter, t_x, for step sizes 8 (Figure 6.3 a and b), 4 (Figure 6.3 c and d), 2 (Figure 6.3 e and f) and 1 (Figure 6.3 g and h). We can see from figure 6.3 that although the relative-anisotropy difference, δ_1 (right column), gives more consistent registration than the tensor difference, δ_3 (left column), similarity measure, both measures show inconsistency. Although figure 6.3 (h) gives 100% registration consistency for the t_x parameter, the highest frequency for the t_y value with step size 1 is 82% (not shown). Furthermore, as expected, there is increased registration consistency as we reduce the step size from 8 in figure 6.3 (a) and (b) to 1 in figure 6.3 (g) and (h). However, even with a step size of 1, registration with Powell's algorithm does not produce the same transformation, i.e. consistent values for all three translation parameters, t_x, t_y and t_z.

- **Conclusions:** The results support **hypothesis (S2)**.

Experiment HS2B

- **Objective:** To test **hypothesis (S2)** for **rigid** transformations using both δ_1 and δ_3 in the objective function with step sizes 8, 4, 2 and 1.

- **Method:** We register the source and target image, as usual using one hundred different starting positions but draw each starting translation from [-2, 2]mm, angle, θ, from a uniform distribution on [0, $\pi/6$] and the axis, $\mathbf{r}$, from a uniform distribution over the unit sphere because it was only within this range that we achieved consistent registration for intra-subject rigid registration.

- **Results:** The rigid transformation parameters are not consistent using either similarity measure. Figure 6.4 shows the frequency of the fitted t_x values using δ_3 (left column) and δ_1 (right column), for step sizes 8 (Figure 6.4 a and b) , 4 (Figure 6.4 c and d), 2 (Figure 6.4 e and f) and 1 (Figure 6.4 g and h). Although there is increased consistency, for both similarity measures, when we reduce the step size from 8 in Figure 6.4 (a) and (b) to 1 in Figure 6.4 (g) and (h), the fitted t_x values are still inconsistent. Figure 6.4 also shows that relative-anisotropy differ-

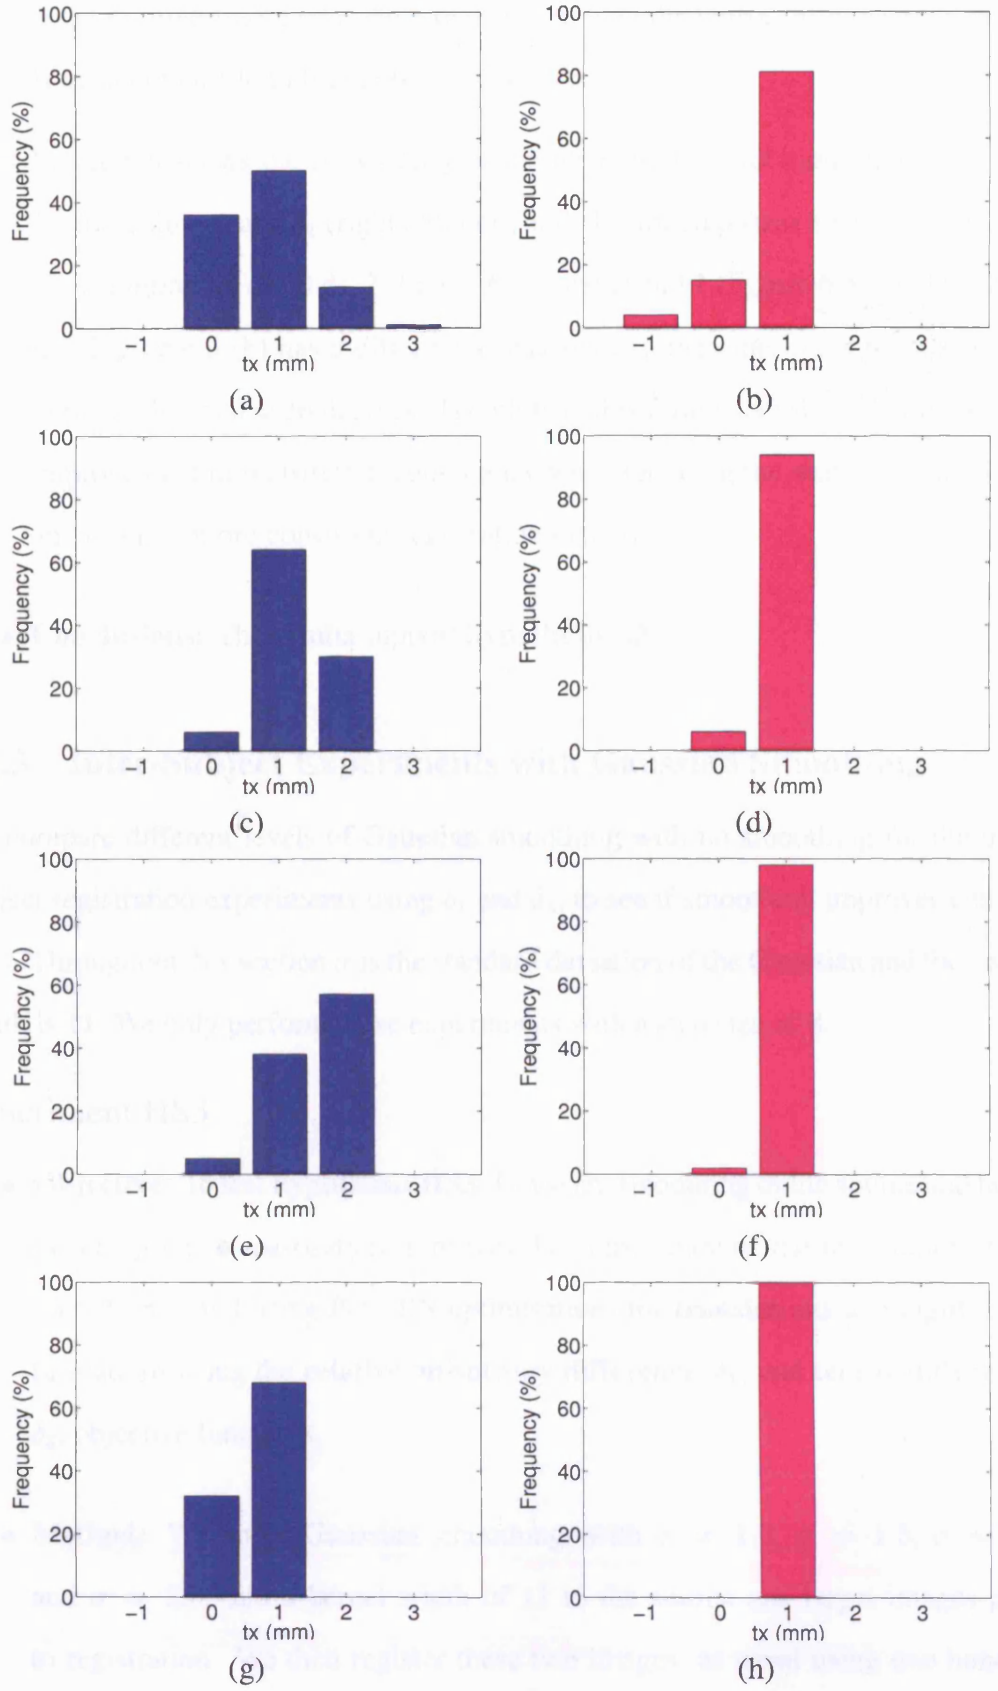

Figure 6.3: The frequency of fitted t_x translation values for **inter-subject** registration (**3DOF**) using tensor difference, δ_3 (left column), and relative anisotropy difference, δ_1 (right column). Images (a) and (b) are with a step size of **8**, (c) and (d) are with a step size of **4**, (e) and (f)) are with a step size of **2** and (g) and (h) are with a step size of **1**.

ence, δ_1, gives more consistent registration than the tensor difference, δ_3, which is most noticeable in Figure 6.4 (g) and (h).

Figure 6.5 shows the consistency of the fitted angle of rotation, θ, values using δ_3 (left column) and δ_1 (right column), and also for step sizes 8 (Figure 6.5 a and b), 4 (Figure 6.5 c and d), 2 (Figure 6.5 e and f) and 1 (Figure 6.5 g and h). Note that Figure 6.5 (b) has a different x-axis scale to the other plots because of the outlier value of $\theta = 26$ degrees. The plots in this figure also show that there is an improvement in registration consistency when reducing the step size from 8 to 1 and δ_1 gives more consistent registration than δ_3.

- **Conclusions:** The results support **hypothesis S2**.

6.2.3 Inter-Subject Experiments with Gaussian Smoothing

We compare different levels of Gaussian smoothing with no smoothing for the inter-subject registration experiments using δ_1 and δ_3, to see if smoothing improves consistency. Throughout this section σ is the standard deviation of the Gaussian and the kernel width is 11. We only perform these experiments with a step size of 8.

Experiment HS3

- **Objective:** To test **hypothesis HS3** (Gaussian smoothing of the source and target images prior to registration improves the consistency of the registration transformation found using Powell's optimisation) for **translations** and **rigid** transformations using the **relative anisotropy difference**, δ_1, and **tensor difference**, δ_3, objective functions.

- **Method:** We apply Gaussian smoothing, with $\sigma = 1.0$, $\sigma = 1.5$, $\sigma = 2.0$ and $\sigma = 2.5$, and a kernel width of 11 to the source and target images prior to registration. We then register these two images, as usual using one hundred different starting positions and drawing each starting translation uniformly from [-5, 5]mm for registration with 3DOF. For rigid registration (6DOF), we draw

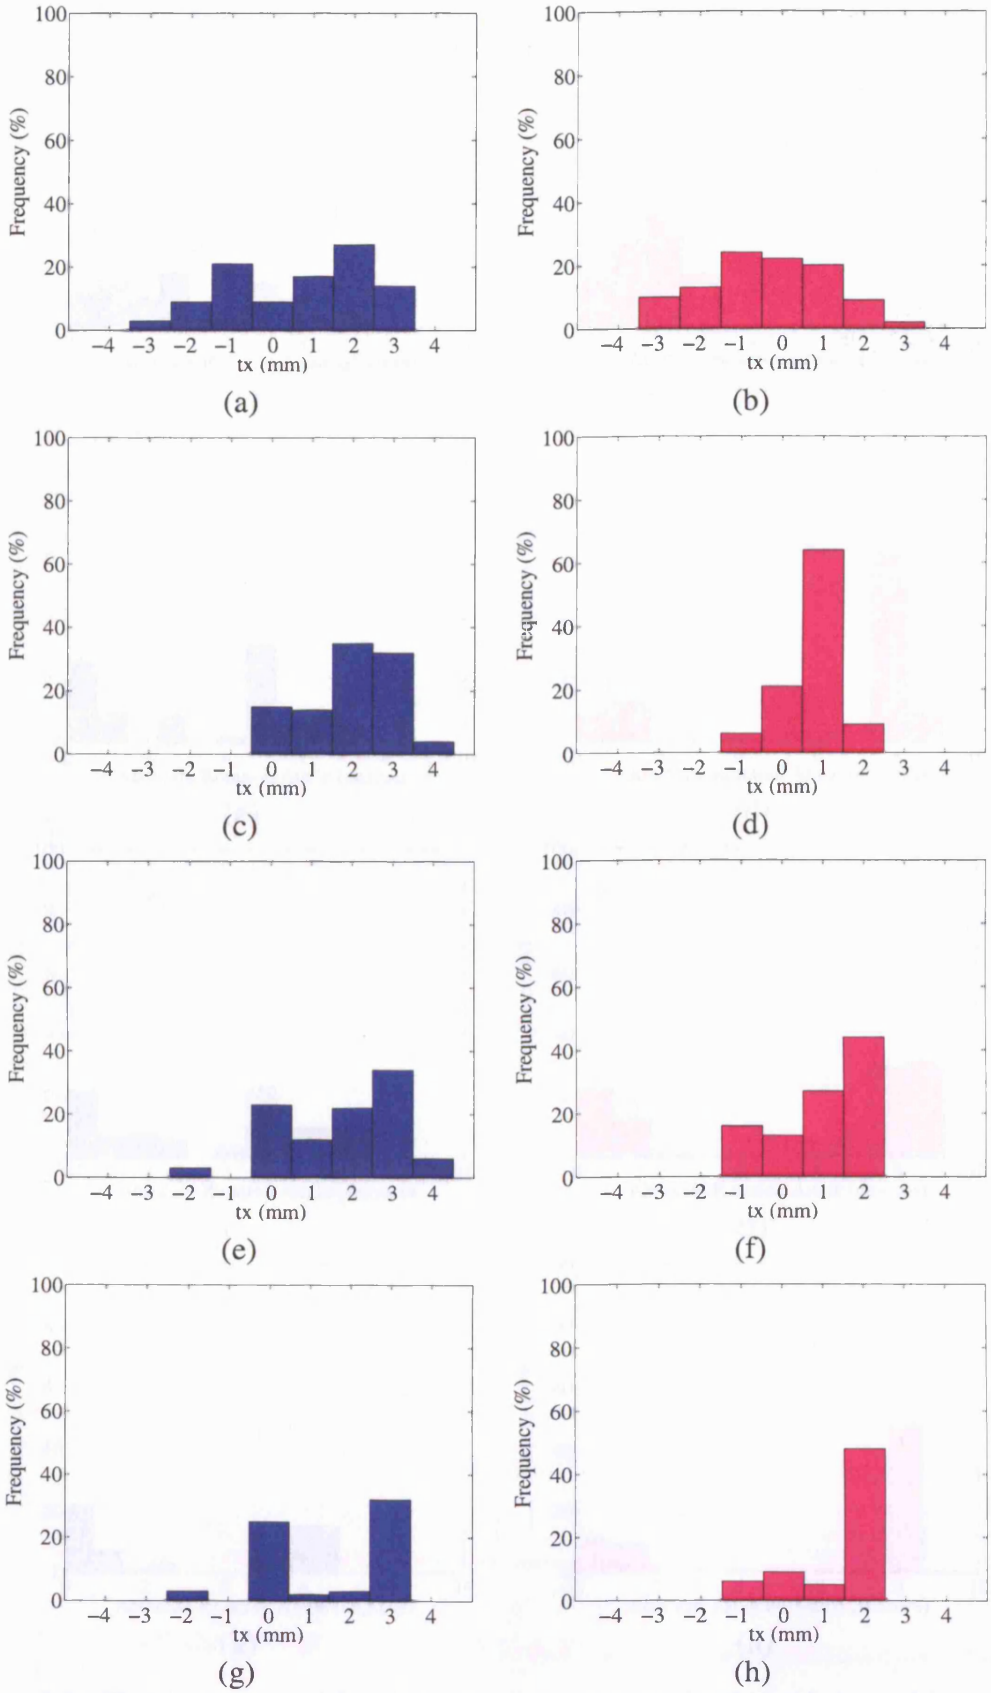

Figure 6.4: The frequency of fitted t_x translation values for **inter-subject** rigid registration (**6DOF**) using tensor difference, δ_3 (left column), and relative anisotropy difference, δ_1 (right column). Images (a) and (b) are with a step size of **8**, (c) and (d) are with a step size of **4**, (e) and (f)) are with a step size of **2** and (g) and (h) are with a step size of **1**.

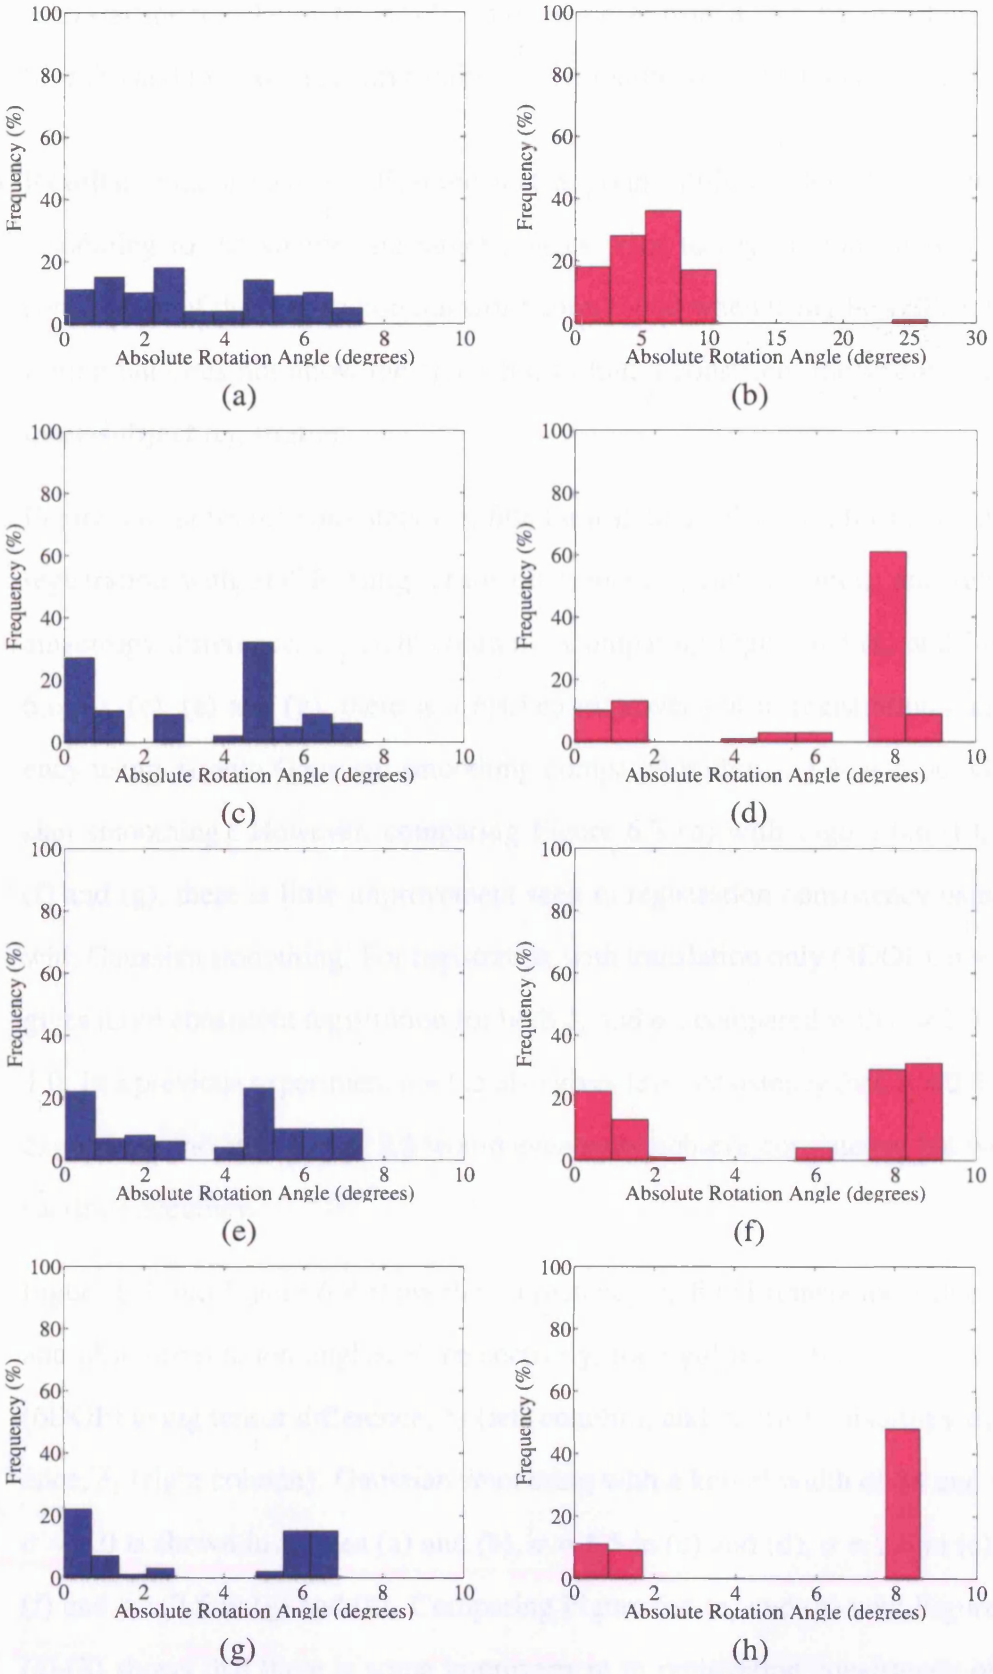

Figure 6.5: The frequency of fitted absolute rotation angle, θ, for **inter-subject** rigid registration (**6DOF**) using tensor difference, δ_3 (left column), and relative anisotropy difference, δ_1 (right column). Images (a) and (b) are with a step size of **8**, (c) and (d) are with a step size of **4**, (e) and (f)) are with a step size of **2** and (g) and (h) are with a step size of **1**.

each starting translation from [-2, 2]mm, angle, θ, from a uniform distribution on [0, $\pi/6$] and the axis, **r**, from a uniform distribution over the unit sphere.

- **Results:** Figures 6.6-6.8 illustrate that applying different levels of Gaussian smoothing to the source and target images prior to registration improves the consistency of the registration transformation found when using Powell's optimisation but does not allow the algorithm to find a consistent transformation for **inter-subject** registration.

  Figure 6.6 shows the consistency of fitted translation values, t_x, for inter-subject registration with 3DOF using tensor difference, δ_3 (left column), and relative anisotropy difference, δ_1 (right column). Comparing Figure 6.3 (a) and Figure 6.6 (a), (c), (e) and (g), there is a marked improvement in registration consistency using δ_3 with Gaussian smoothing compared with using δ_3 without Gaussian smoothing. However, comparing Figure 6.3 (b) with Figure 6.6 (b), (d), (f) and (g), there is little improvement seen in registration consistency using δ_1 with Gaussian smoothing. For registration with translation only (3DOF), $\sigma = 2.5$, gives more consistent registration for both δ_1 and δ_3, compared with $\sigma = 2.0$, 1.5, 1.0. In a previous experiment $\sigma = 0.5$ also gives less consistency than $\sigma = 2.5$. We expect that increasing $\sigma > 2.5$ would eventually achieve consistency but would sacrifice accuracy.

  Figure 6.7 and Figure 6.8 show the consistency of fitted translation values, t_x, and absolute rotation angles, θ, respectively, for rigid inter-subject registration (6DOF) using tensor difference, δ_3 (left column), and relative anisotropy difference, δ_1 (right column). Gaussian smoothing with a kernel width of 11 and with $\sigma = 1.0$ is shown in images (a) and (b), $\sigma = 1.5$ in (c) and (d), $\sigma = 2.0$ in (e) and (f) and $\sigma = 2.5$ in (g) and (h). Comparing Figure 6.4 (a) and (b) with Figure 6.7 (a)-(h) shows that there is some improvement in registration consistency of the fitted translation values, t_x, for both similarity measures, with Gaussian smoothing. However, the improvement in registration consistency is not as dramatic as

registration with only 3DOF.

Finally, comparing Figure 6.5 (a) and (b) with Figure 6.8 (a)-(h), there is little improvement in registration consistency with Gaussian smoothing for either similarity measure. Note the difference in the x-axis scale in Figure 6.8 (f). The various settings for σ give similar results and rigid registration reveals no clearly optimal σ. We show some example qualitative results in Figure 6.9. We show example inter-subject registrations for both similarity measures and compare the difference in registration results between step size 8 and 1. We also compare the inter-subject registration results using δ_1 and δ_3 with Gaussian smoothing with a kernel width of 11 and standard deviation, $\sigma = 2.5$. Figure 6.9 shows that there is little visual difference between the best registrations obtained with and without smoothing and with different step sizes.

- **Conclusions:** The results support **hypothesis (S3)**. However, even with smoothing we do not achieve 100% consistency.

6.3 Summary

- We accept hypothesis (S1) that Powell's optimisation finds a consistent transformation for **intra-subject** registration for translations and rigid transformations if the translation parameters are within 2mm and the rotation angle within $\pi/6$ of the global minimum settings.

- We accept hypothesis (S2) that Powell's optimisation does **not** find a consistent transformation for **inter-subject** registration.

- We accept hypothesis (S3) that Gaussian smoothing of the source and target images prior to registration improves the consistency of the registration transformation found using Powell's optimisation. However, even with smoothing we do not achieve 100% consistency.

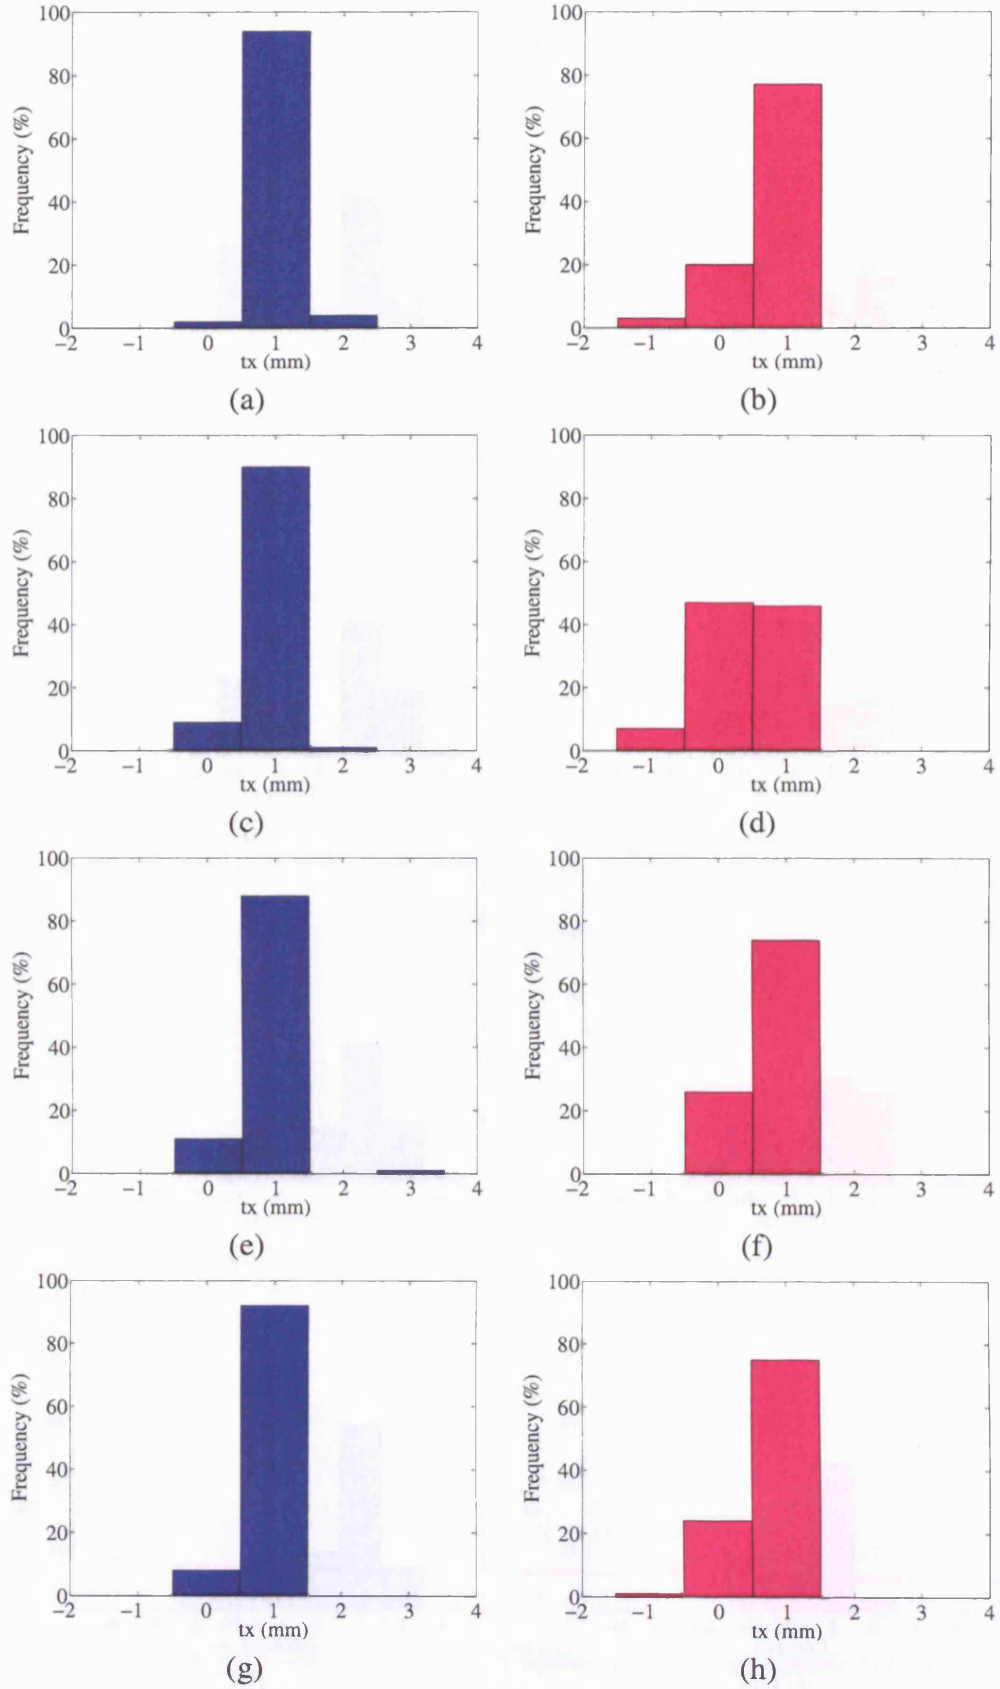

Figure 6.6: The frequency of fitted translation values, t_x, for **inter-subject** registration (**3DOF**) using tensor difference, δ_3 (left column), and relative anisotropy difference, δ_1 (right column). **Gaussian smoothing** with a kernel width of 11 was applied to the source and target images prior to registration. $\sigma = 1.0$ in Images (a) and (b), $\sigma = 1.5$ in (c) and (d), $\sigma = 2.0$ in (e) and (f) and $\sigma = 2.5$ in (g) and (h).

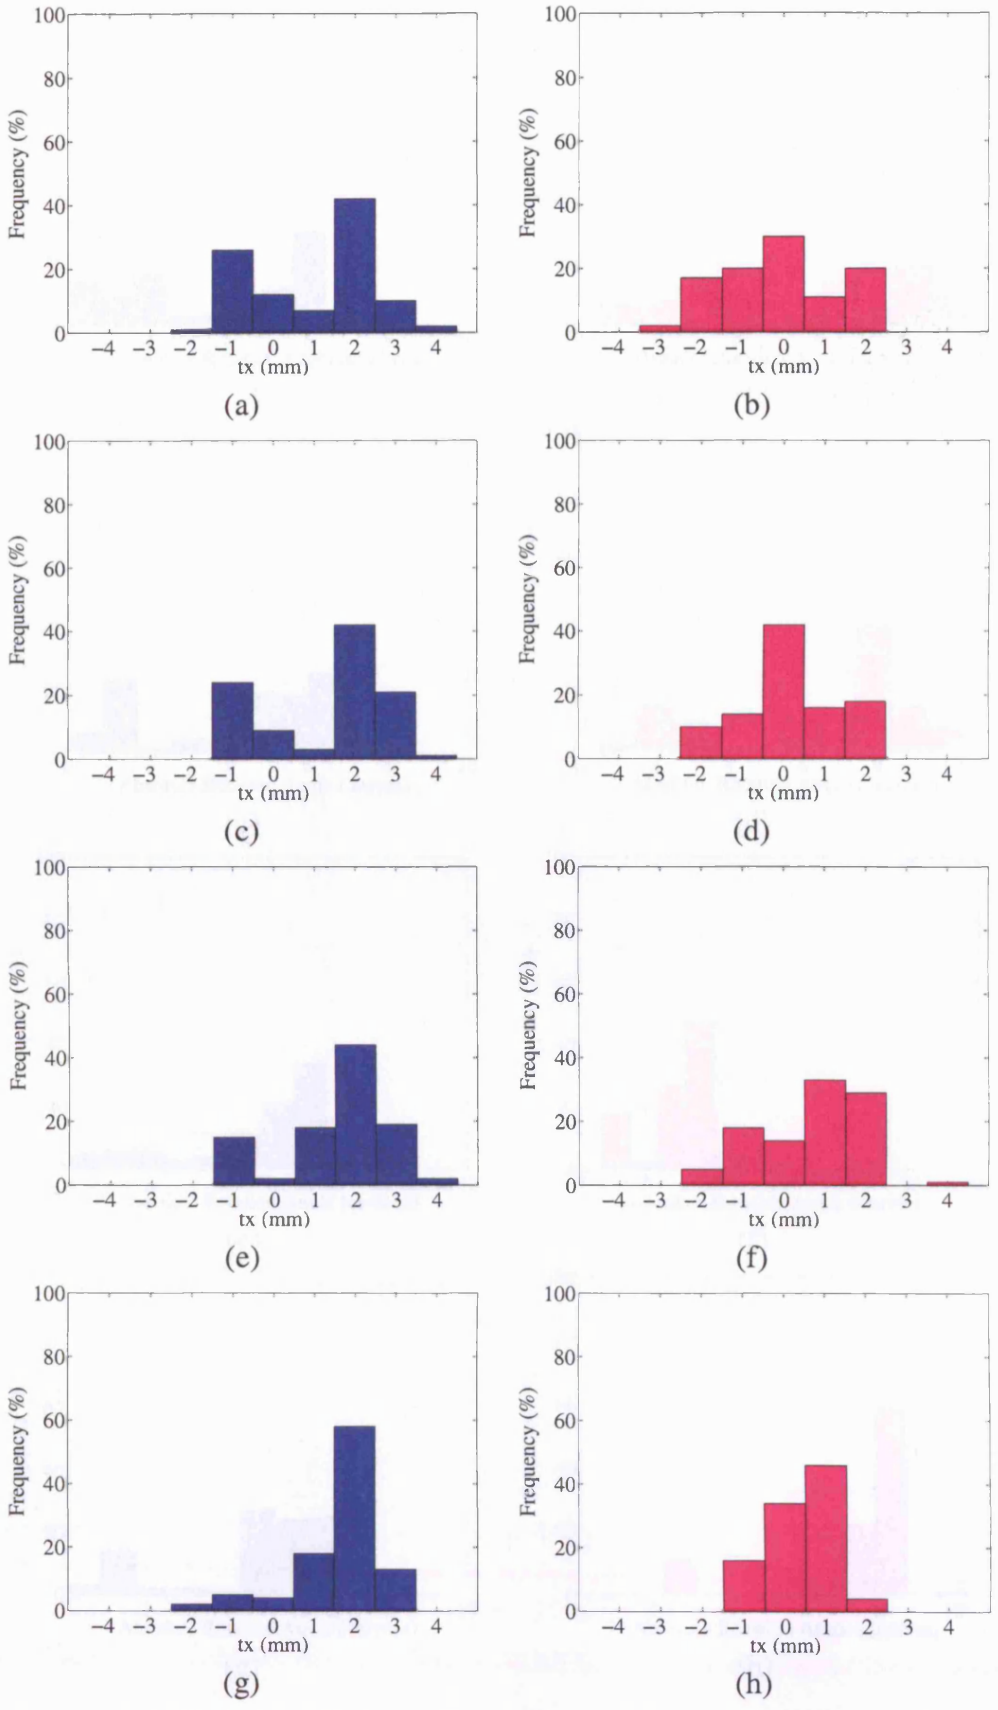

Figure 6.7: The frequency of fitted translation values, t_x, for **inter-subject** rigid registration (**6DOF**) using tensor difference, δ_3 (left column), and relative anisotropy difference, δ_1 (right column). **Gaussian smoothing** with a kernel width of 11 was applied to the source and target images prior to registration. $\sigma = 1.0$ in Images (a) and (b), $\sigma = 1.5$ in (c) and (d), $\sigma = 2.0$ in (e) and (f) and $\sigma = 2.5$ in (g) and (h).

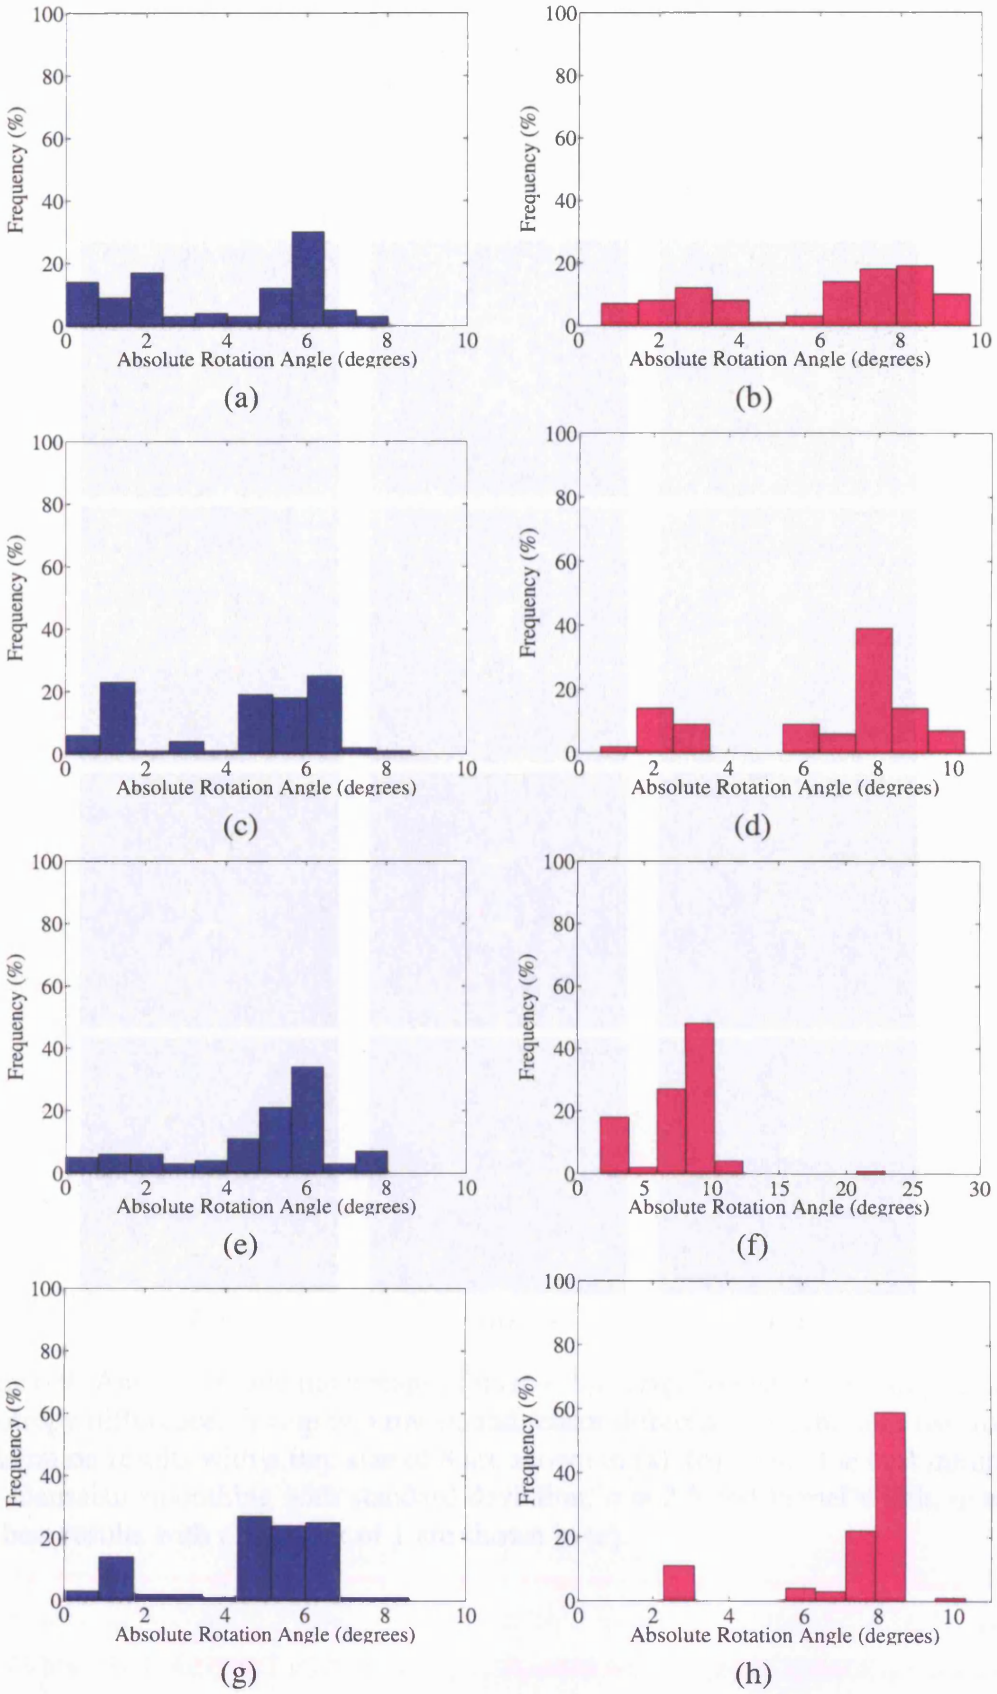

Figure 6.8: The frequency of fitted absolute rotation angle, θ, for **inter-subject** rigid registration (**6DOF**) using tensor difference, δ_3 (left column), and relative anisotropy difference, δ_1 (right column). **Gaussian smoothing** with a kernel width of 11 was applied to the source and target images prior to registration. $\sigma = 1.0$ in Images (a) and (b), $\sigma = 1.5$ in (c) and (d), $\sigma = 2.0$ in (e) and (f) and $\sigma = 2.5$ in (g) and (h).

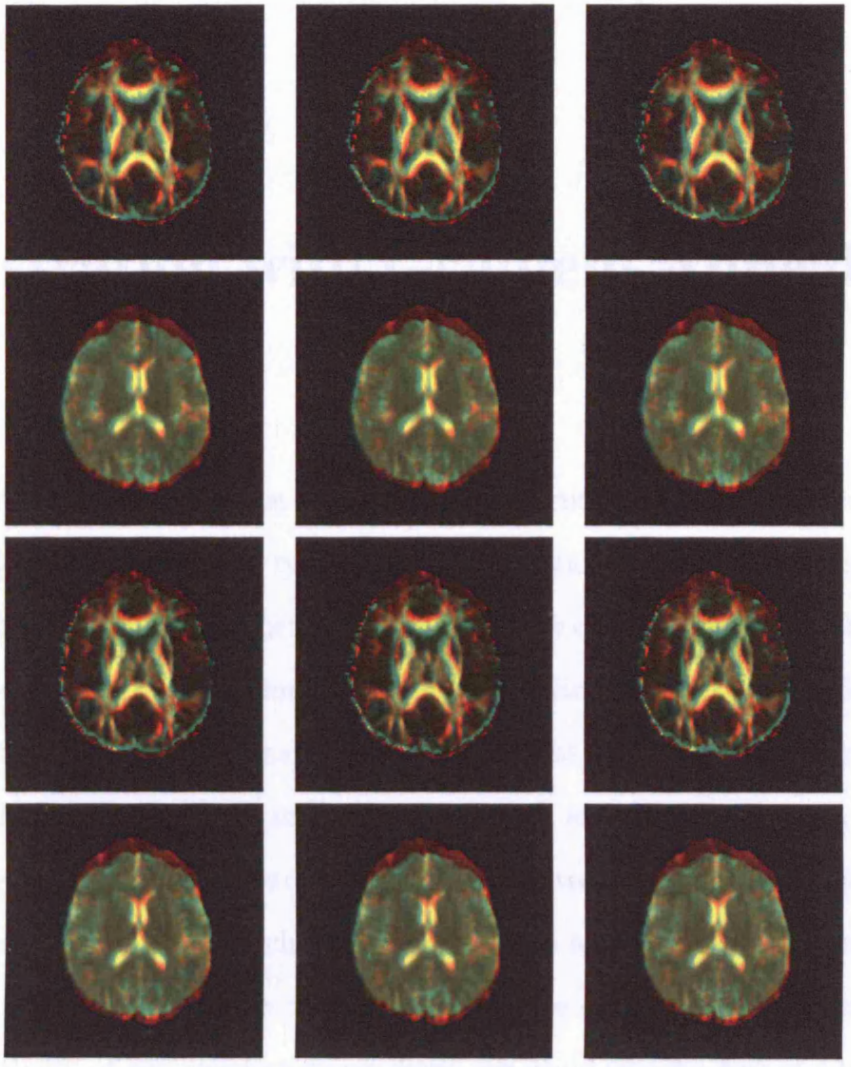

(a) (b) (c)

Figure 6.9: Anisotropy and trace maps of inter-subject rigid registration, using relative-anisotropy difference, δ_1 (top two rows), and tensor difference, δ_3, (bottom two rows). Registration results with a step size of 8 are shown in (a), (b) shows the best minimum with Gaussian smoothing with standard deviation, $\sigma = 2.5$ and kernel width, $w = 11$. The best results with a step size of 1 are shown in (c).

Chapter 7

Registration with Gradient Annealing

The previous chapter shows that Powell's direct optimisation technique does not guarantee to find the minimum of typical objective functions used in image registration. This renders such techniques ineffective for testing the central hypothesis of this thesis, i.e., that orientation matching improves registration, since any observed differences in the performance of different matching schemes might arise solely from optimisation errors. The previous chapter showed inconsistency for simple translation and rigid registration. For testing the central thesis hypotheses we need to use affine and higher-order transformations, for which we can expect even less consistency. Thus we need to use a global optimisation technique to improve the consistency of registration and thus the solidity of any conclusions we draw. We could take the best of a large number of trial runs, as in the previous chapter, but techniques such as simulated annealing offer a more principled way to search for the global minimum. The next chapter uses simulated-annealing-based registration using affine and polynomial transformations to test the central hypotheses of the thesis. This chapter studies the simulated-annealing registration algorithm, focusing on affine transformations with no reorientation, to determine appropriate settings in the algorithm for running the experiments in the next chapter. Finally, we describe a set of experiments designed to compare the gradient annealing registration algorithm with registration using Powell's algorithm directly.

7.1 Experiment Design

In theory, simulated annealing can always find the global minimum of an objective function [64]. However, computation time usually prohibits a search that guarantees to find the global minimum. In practice, we typically sacrifice the global minimum guarantee for reasonable computation time and choose an annealing schedule that gets as close as possible to the global minimum within acceptable time limits.

The longer we run the simulated annealing process the better optima we are likely to find. Here, we impose a time limit on registration optimisations in order to compute results in a reasonable amount of time and optimise the algorithm to the best performance within that fixed time frame. The gradient annealing registration algorithm has various internal parameters for which we can choose settings to optimise performance within our time limit:

- Starting temperature T_0: During annealing the temperature of the system decreases gradually from the starting temperature to a freezing point. In simulated annealing, the temperature controls the size of perturbations of the parameters to be optimised. The starting temperature therefore controls the largest perturbations that the system undergoes. If it is too low, the system may never escape minima local to the starting position. If it is too high we may waste computation time investigating remote corners of the parameter space where solutions are highly unlikely to be correct.

- Freezing temperature T_f: The freezing temperature is the temperature at which simulated annealing is stopped. In practice, it is not necessary to let the temperature reach zero because as it approaches zero the chances of accepting an uphill step are almost zero and the perturbations in the parameter values become so small that their effects on the value of the objective function become insignificant.

- Cooling factor c: The cooling factor controls the rate at which the temperature decreases. In the annealing schedule, we use $T_n = T_{n-1} \times c$, where T_n is the n-th

temperature and $0 < c < 1$. The higher the value of c, the longer it takes to reach the freezing temperature, T_f.

- Step size: The step size controls the size of the array of voxels that contribute to the similarity. The experiments in this chapter look at four possible step sizes: 8, 4, 2 and 1, see §5.5.

For each step size, we identify an appropriate freezing temperature T_f and, using this T_f, we test different combinations of starting temperatures T_0 and cooling factors c to determine the settings that find the best registration within the time-limit, which we set to approximately 24 hours of computation time on an Intel Xeon CPU running at 2.80GHz.

The final criterion we use to select the optimal settings for registration with gradient annealing is the Euclidean mean distance $\bar{d}$ between fifty-two hand-defined corresponding anatomic landmarks in the warped-source and target images.

7.2 Annealing Temperature Range

The temperature in simulated annealing determines the size of perturbations to the parameter settings in the search for the optimum. The procedure starts at high temperature where the perturbations are large and gradually reduces the temperature to focus in on the optimal solution. At present, there is no generally accepted method for finding a suitable starting temperature for an objective function. The optimum varies from one type of problem to another and also from one instance of a problem to another instance of the same problem. The literature contains some heuristic approaches, for example, Rayward-Smith et al. [108] suggest starting with a very high temperature and cooling rapidly until about 60% of solutions are accepted. The main simulated annealing process then starts with slower cooling. A similar idea, suggested in [48], is to heat the system rapidly until a certain proportion of solutions are accepted and then slow cooling can start. This is analogous to how physical annealing works in that the material is heated until it is a liquid and then cooling begins.

In order to find an appropriate temperature range for DT-MR registration, we run an initial gradient-annealing procedure over a very wide range of temperatures $[10^5, 10^{-16}]$ to match three source images to a single target image. We run the same experiment for registrations using step sizes of eight, four, two and one to determine temperature ranges separately for each step size. Experiments revealed that the parameter settings vary too widely to produce any sensible results with $T > 10^3$, see for example Figure 7.1 with $T > 600$, therefore $T_0 = 10^5$ should be sufficient to escape local minima. We choose freezing points for each step size below which the value of the objective function varies on a scale of less than 5% for at least two of the three subjects. For step size eight, Figure 7.1 thus shows that $T_f = 10^{-15}$. For step sizes four, two and one, we choose $T_f = 10^{-10}$; the figures corresponding to 7.1 for these step sizes are included in Appendix A , see Figures A.1-A.3.

We do not expect the appropriate choice of temperature range to depend strongly on the choice of similarity measure, since the freezing temperature controls the smallest scale of variation of the transformation parameters rather than the scale of the objective function. To confirm this we performed gradient annealing registration over a range of temperatures $[10^5, 10^{-16}]$ for each of the similarity measures listed in §4.4 and, as expected, for step size 8, at $T_f = 10^{-15}$, the value of the objective function varies on a scale of less than 5% for at least two of the three subjects for each similarity measure.

7.3 Initial Temperature and Cooling Factor

With the running time fixed (at 24 hours), we can determine the number of evaluations of the objective function that the simulated annealing procedure can make. The objective function of the simulated annealing process is the image similarity after optimisation with Powell's algorithm from the starting point controlled by the state of the annealing process. Thus the evaluation of the simulated annealing objective function involves one complete run of Powell's algorithm to maximise image similarity. For each step size, we compute an average runtime of Powell's algorithm using the time from the experiments in the previous section (§7.2). From the average Powell's

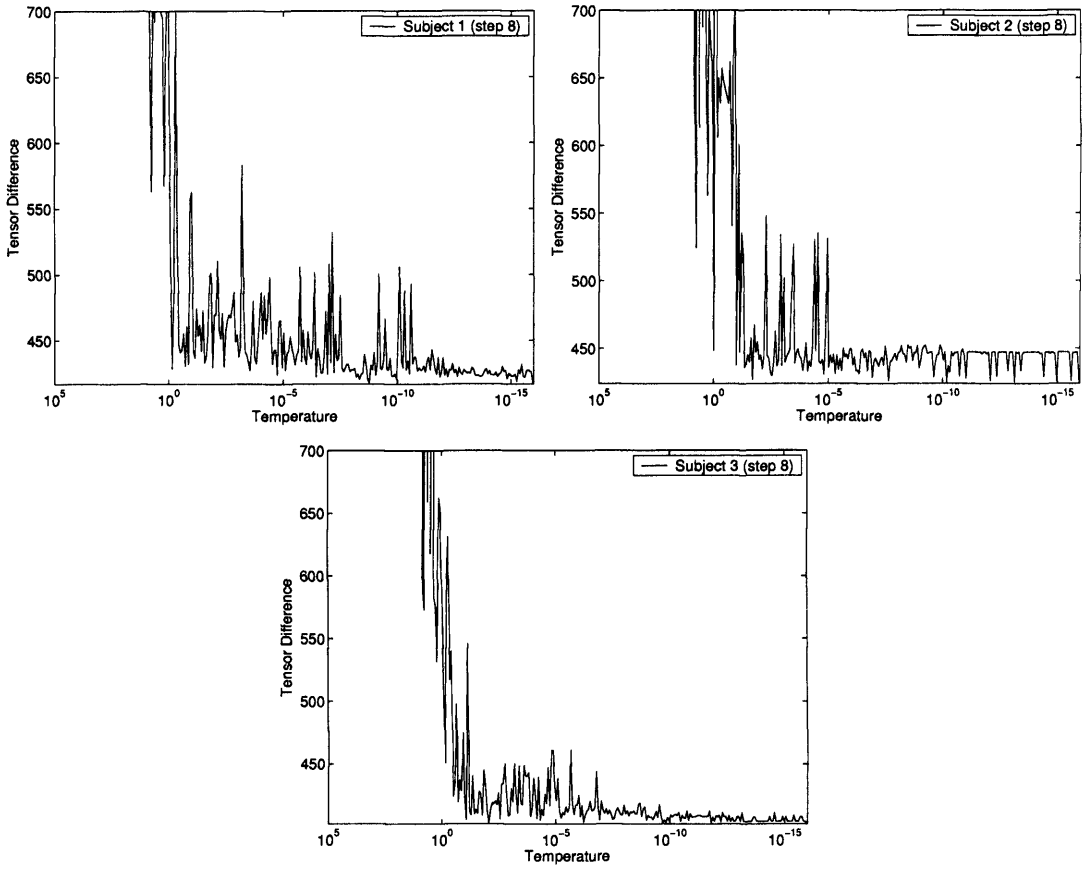

Figure 7.1: Plot of the tensor difference similarity index against temperature, where the temperature has been initialised to 10^5 and the freezing point is 10^{-16}. All three figures are from registrations with a step size of eight and each plot shows the results for three different data sets.

algorithm runtime and the total run length (24 hours) we determine the number n of allowable runs for each step size.

Since n and T_f are fixed, the cooling factor, c, and starting temperature, T_0, are dependent on each other. With starting temperature, T_0, freezing temperature, T_f, and cooling factor, c, the total number of runs is

$$n = \frac{\log T_f - \log T_0}{\log c} + 1. \tag{7.1}$$

With n and T_f fixed, we can determine T_0 from any choice of c and vice versa using Eq. 7.1. At each step size, we determine the cooling factor c that gives $T_0 = 10^5$ and then choose nine other cooling factors evenly spaced between c and one. These ten cooling factors, the corresponding ten starting temperatures and the fixed freezing temperature are the sets of parameters over which we maximise registration performance at each step size. Table 7.1 shows the final δ_3, for one subject, after the gradient annealing process terminates with various combinations of T_0 and c for each step size. The table also shows $\bar{d}$ for each computed registration. The results show that affine registration with a step size of eight, an initial starting temperature of 5.92×10^{-12} and a cooling factor of 0.9728 gives the lowest mean distance ($\bar{d}$ = 3.6mm) between the landmarks in the warped-source and target images. We note that, although the dependence of $\bar{d}$ on T_0 and c does not seem particularly smooth and the selected combination of settings does not provide a clear minimum $\bar{d}$, a similar set of experiments with the runtime set to seven days instead of 24 hours gives a similar optimal combination with a step size of eight and a high cooling factor (c = 0.9971) for the best registration.

7.4 Direct vs. Global Optimisation

In this section, we compare the performance of Powell's algorithm, a direct optimisation method, with that of the gradient annealing algorithm, proposed in §5.7.2, with the settings identified in the previous sections of this chapter for DT-MR brain registration. We first define the hypothesis we wish to test.

Table 7.1: The mean Euclidean distance $\bar{d}$ (mm) measures for ten annealing parameter combinations for step sizes eight (S8), four (S4), two (S2) and one (S1) for one subject. The mean distance for the unregistered images is 6.3mm. Note that the first T_0 for each step size is not exactly 10^5 because c has been rounded to 4 decimal places.

S8	T_0	c	δ_3	$\bar{d}$
	1.03e05	0.8639	417	3.9
	7.54e02	0.8775	417	3.8
	5.93e00	0.8911	418	3.8
	5.00e-02	0.9047	420	3.8
	4.57e-04	0.9183	420	3.9
	4.31e-06	0.9320	436	3.9
	4.49e-08	0.9456	435	3.9
	5.00e-10	0.9592	425	3.9
	5.92e-12	**0.9728**	420	**3.6**
	7.47e-14	0.9864	435	4.0

S4	T_0	c	δ_3	$\bar{d}$
	1.00e05	0.5829	422	3.9
	1.21e03	0.6246	422	3.9
	1.93e01	0.6663	422	3.8
	3.96e-01	0.7080	422	3.9
	1.00e-02	0.7497	423	3.9
	3.16e-04	0.7915	422	3.9
	1.18e-05	0.8332	422	3.8
	5.18e-07	0.8749	422	3.9
	2.63e-08	0.9166	423	3.9
	1.53e-09	0.9583	422	3.9

S2	T_0	c	δ_3	$\bar{d}$
	1.00e05	0.1778	424	3.9
	5.02e01	0.2600	423	3.9
	2.06e-01	0.3422	423	3.9
	3.00e-03	0.4245	424	3.8
	8.03e-05	0.5067	425	3.8
	3.97e-06	0.5889	424	3.7
	2.91e-07	0.6711	424	3.9
	2.89e-08	0.7533	424	3.8
	3.63e-09	0.8356	424	3.8
	5.56e-10	0.9178	424	3.9

S1	T_0	c	δ_3	$\bar{d}$
	9.31e04	0.0032	424	3.8
	8.42e-05	0.1029	424	3.8
	1.45e-06	0.2026	424	3.8
	1.31e-07	0.3022	424	3.8
	2.37e-08	0.4019	424	3.8
	6.28e-09	0.5016	424	3.8
	2.12e-09	0.6013	424	3.8
	8.43e-10	0.7010	424	3.8
	3.80e-10	0.8006	424	3.8
	1.88e-10	0.9003	424	3.8

Table 7.2: The mean similarity results, over all five subjects, for relative-anisotropy difference, δ_1, modulus difference, δ_2, tensor difference, δ_3, normalised tensor-difference, δ_4, normalised tensor scalar product, δ_6, and principal-direction difference, δ_7.

	Powell			Annealing		
	None	**FS**	**PPD**	**None**	**FS**	**PPD**
δ_1	0.068	0.068	0.067	0.063	0.063	0.063
δ_2	206	207	207	199	200	201
δ_3	464	448	470	447	424	425
δ_4	0.167	0.167	0.165	0.164	0.159	0.160
δ_6	-0.584	-0.584	-0.584	-0.584	-0.584	-0.584
δ_7	0.638	0.641	0.622	0.602	0.594	0.597

Hypothesis (S4): "The minimum of the objective function found using gradient annealing is lower than the minimum found using Powell's method."

7.4.1 Method

To test hypothesis (S4), we compare the summed similarity, for each of the candidate similarity measures, listed in §8.2.1, for the Powell and gradient annealing registrations.

7.4.2 Results

Table 7.2 compares the average final summed relative-anisotropy difference, δ_1, modulus difference, δ_2, tensor difference, δ_3, normalised tensor-difference, δ_4, normalised tensor scalar product, δ_6, and principal-direction difference, δ_7, similarity measures, over five subjects, after optimisation using Powell's method with that from the full gradient annealing optimisation. We do not present results for the tensor scalar product, δ_5, because it often failed in early testing. We therefore do not consider it further in this thesis. The summed similarity measures for each of the subjects are given in Appendix B. The results in Table 7.2 show that Powell's algorithm clearly does not find the global minimum. The global optimisation reduces the final objective function value consistently and by as much as 21%. Figure 7.2 illustrates an example, where the gradient annealing registration, using δ_3 with PPD reorientation, finds a similarity measure (δ_3 = 447) for subject two that is considerably lower than that found by Powell registration with PPD reorientation (δ_3 = 566). The global optimisation reduces the final objective

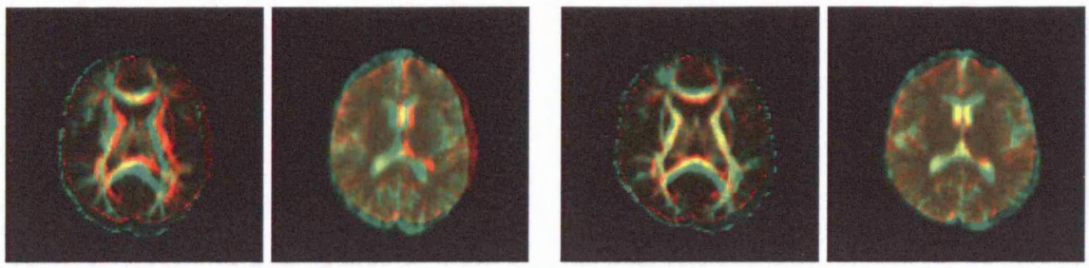

(a) (b)

Figure 7.2: Anisotropy (left) and trace (right) maps for one slice of a source (subject two) and target volumetric DT-MR image. Image (a) shows the results of the Powell registration, while (b) shows the gradient annealing registration results for the same subject.

function value consistently and by as much as 10% for δ_1, 9% for δ_2, 21% for δ_3, 7% for δ_4 and 11% for δ_7 and increases the final objective function value consistently and by as much as 0.2% for δ_6. We therefore accept hypothesis S4.

7.4.3 Summary

- We choose the following settings for the registration with gradient annealing experiments: step size eight, freezing temperature, $T_f = 10^{-15}$, starting temperature, $T_0 = 5.92 \times 10^{-12}$ and a cooling factor, c, of 0.9728 because these settings give the lowest mean distance ($\bar{d}$ = 3.6mm) between the landmarks in the warped-source and target images.
- The global optimisation reduces the final objective function value consistently and by as much as 21%.

Chapter 8

Experiments

This chapter details experiments that address the central hypotheses of the thesis outlined in Chapter 1. We compare the effectiveness of the similarity measures listed in §4.4 and §5.3 both qualitatively and quantitatively. We then describe an additional set of experiments which compares image registration with and without orientation matching. We define the hypotheses, describe the experiments and present the results.

8.1 Hypotheses

We test the following two hypotheses.

- **Hypothesis (C1):** "Similarity measures derived from the full tensor matrix find a more accurate registration than those based on the derived scalar indices that do not use tensor orientational information."

- **Hypothesis (C2):** "Diffusion tensor orientation matching improves the accuracy of registration over that from registration algorithms that ignore orientation."

8.2 Comparison of Similarity Measures

In order to compare the performance of similarity measures derived from the full tensor matrix with those based on derived scalar indices that do not use tensor orientational information (therefore test hypothesis C1), a measure of quality of the registration is required. For inter-subject brain registration, it is difficult to find a measure of quality because anatomical features may be specific to particular brains and not present in

others. We register five DT-MR images to a sixth template image using gradient annealing registration with the annealing settings determined in Chapter 7. We perform affine registration with PPD reorientation (§4.1) using each similarity measure listed in §8.2.1. We compare the mean Euclidean distance $\bar{d}$ of 52 corresponding landmarks in the warped-source and target images. Landmarks can be difficult to place accurately by eye, particularly in three-dimensional images. To place landmarks as accurately as possible, we view the trace, anisotropy and principal-direction colour maps in all three image planes simultaneously. We also perform scalar registration, minimising the Euclidean distance between landmarks, to calculate a gold standard or best possible fit of the landmarks.

8.2.1 Similarity Measures

We compare the following similarity measures that are defined in §4.4 and §5.3:

- relative-anisotropy difference, δ_1,
- modulus difference, δ_2,
- tensor difference, δ_3,
- normalised tensor-difference, δ_4,
- normalised tensor scalar product, δ_6,
- principal-direction difference, δ_7.

8.2.2 Results

The mean landmark-distances, $\bar{d}$, for each of the five subjects after registration using each similarity measure are summarised in Table 8.1. The numerical results given in Table 8.1 are plotted in Figures 8.1-8.2. The best mean result is shown in bold typeface. The gold standard distances, i.e. the distances calculated when only the landmarks were registered are $\bar{d} = 2.84$, $\bar{d} = 3.09$, $\bar{d} = 3.42$, $\bar{d} = 4.26$, $\bar{d} = 3.86$ for subjects 1-5 respectively. Table 8.1 shows that the relative-anisotropy difference,

δ_1, find mean landmark distances, $\bar{d}$, for subjects 1-4 that are close to the gold standard distances ($\bar{d} = 3.45$, $\bar{d} = 3.26$, $\bar{d} = 4.09$ and $\bar{d} = 4.78$, respectively). However, the landmark distances for the other four similarity measures are, in general, higher than the δ_1 distances. Figure 8.1 shows a plot of the mean landmark-distances for each of the five subjects after registration using each similarity measure. The plot shows that the relative-anisotropy difference, δ_1, the tensor difference, δ_3, the normalised tensor scalar product, δ_6, and the principal-direction difference, δ_7, find a mean landmark distance, $\bar{d}$, for all five subjects of less than or equal to 6mm. The normalised tensor difference, δ_4, also finds $\bar{d} < 6$mm for all five subjects, except for one subject, where $\bar{d}$ = 9.41mm. However, we can see from Table 8.1 and Figure 8.1 that the modulus difference, δ_2, has $\bar{d} > 7$mm for two subjects. Figure 8.2 shows the same results

Table 8.1: Mean landmark distances, $\bar{d}$, in mm for each of the five subjects after registration using each similarity measure. The best mean result is shown in bold typeface.

subject	δ_1	δ_2	δ_3	δ_4	δ_6	δ_7
1	3.45	7.86	4.00	3.93	3.44	3.41
2	3.26	4.15	3.83	3.98	3.88	3.63
3	4.09	3.91	4.33	4.10	4.65	4.74
4	4.78	8.75	5.22	9.41	6.00	5.24
5	5.38	4.71	4.99	5.06	5.20	5.35
mean	**4.19**	5.87	4.47	5.30	4.63	4.48

reordered to compare the performance of the similarity measures for each subject. The plot shows that subjects four and five are least accurately registered, with $\bar{d} > 4$mm for each similarity measure. For subject four, the modulus difference and normalised tensor difference show even higher landmark-distances ($\bar{d} > 8$mm). It is interesting that a relatively accurate registration ($\bar{d} < 4$mm) was found for subject one, for all similarity measures, except using δ_2 ($\bar{d} > 7$mm).

Figure 8.3 shows the pre- and post-registration results for subject two, where the mean landmark distance, $\bar{d}$, for the source and target images prior to registration is 6.30mm and the results for each similarity measure find a mean landmark distance of less than 5mm. Each of the images in Figure 8.3 show the anisotropy and trace maps

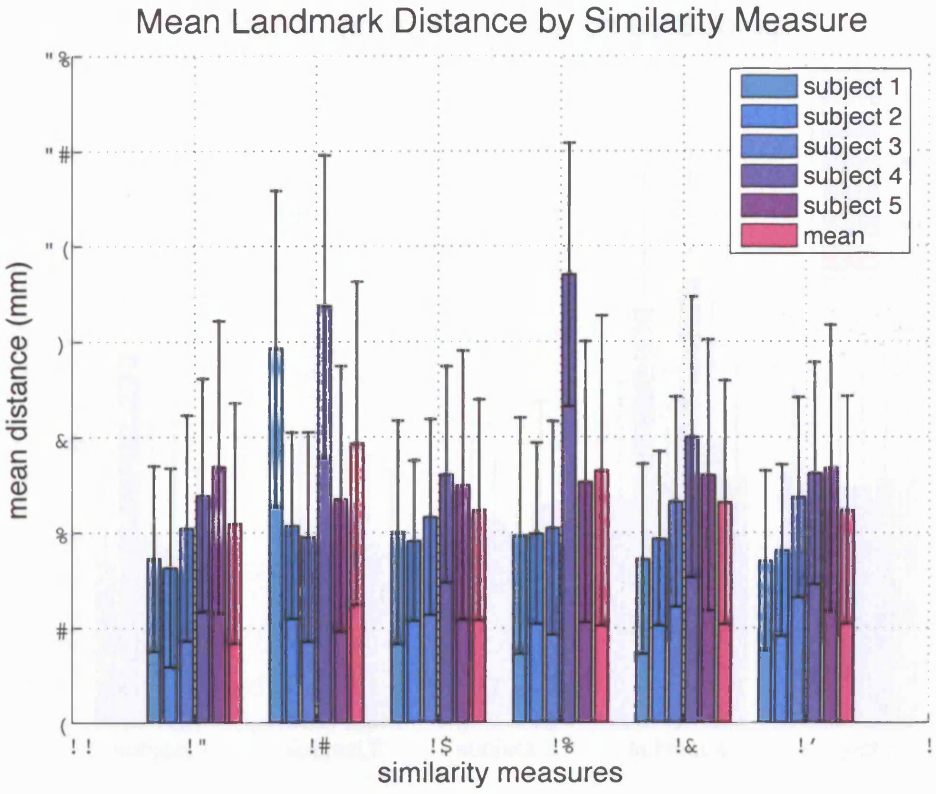

Figure 8.1: Plot of the mean landmark separations of the five subjects for the relative-anisotropy difference, δ_1, modulus difference, δ_2, tensor difference, δ_3, normalised tensor-difference, δ_4, noramalised tensor scalar product, δ_6, and principal-direction difference, δ_7, similarity measures.

for one slice of the source and target volumetric DT-MRI brain image. We can see that, qualitatively, there is an improvement between the pre-registration images and the registration results for all the similarity measures. Comparing the results of the two scalar-derived similarity measures, relative-anisotropy difference, δ_1, and modulus difference, δ_2, Figure 8.3 shows that registration using δ_1 produces a much better alignment than registration using δ_2. Interestingly, the other similarity measures that are sensitive to size, shape and orientation (δ_3, δ_4, δ_6 and δ_7) do not find as good a match as δ_1.

Figure 8.4 shows the pre- and post-registration results for subject four, where $\bar{d}$ = 10.24mm, for the source and target images prior to registration. We can see from Figure 8.4 that registration using both δ_2 and δ_4 find poor image matches. However, it is difficult to ascertain, qualitatively, if δ_1, δ_3, δ_6 or δ_7 finds the best image match.

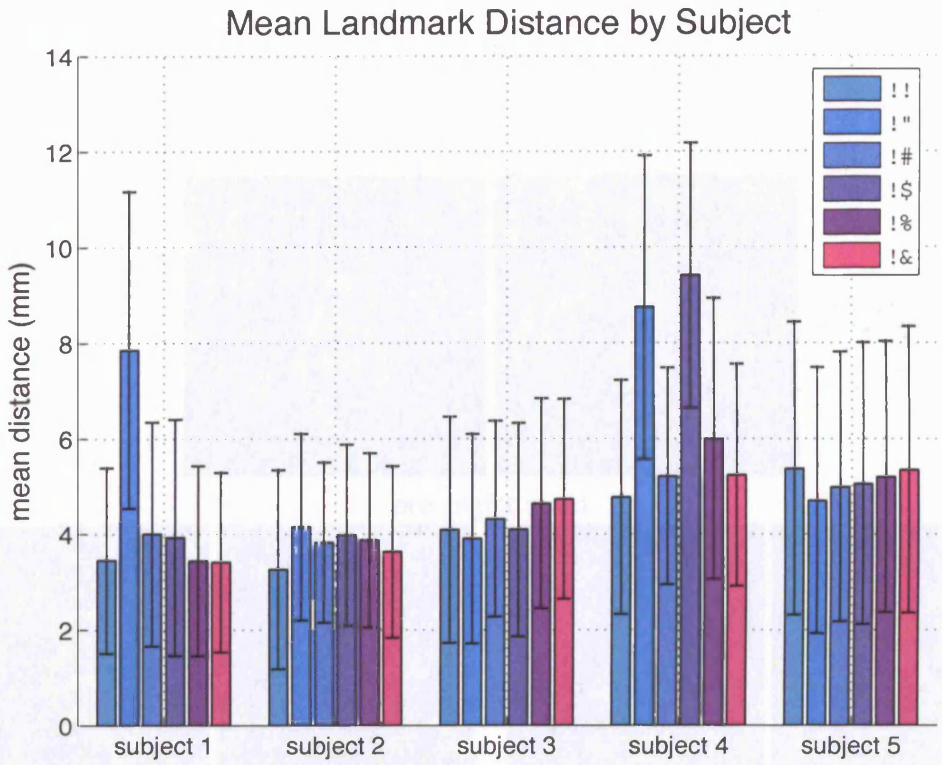

Figure 8.2: Comparison of the performance of the relative-anisotropy difference, δ_1, modulus difference, δ_2, tensor difference, δ_3, normalised tensor-difference, δ_4, normalised tensor scalar product, δ_6, and principal-direction difference, δ_7, similarity measures for each of the five subjects.

Comparing the similarity measures quantitatively (Figure 8.2), we find that for subject four relative-anisotropy difference, δ_1, finds a smaller $\bar{d}$ (4.78mm) than the tensor difference, δ_3, ($\bar{d}$ = 5.22mm), the normalised tensor scalar product, δ_6, ($\bar{d}$ = 6.00mm), and the principal-direction difference, δ_7, ($\bar{d}$ = 5.24mm).

We therefore reject hypothesis C1 (similarity measures derived from the full tensor matrix find a more accurate registration than those based on the derived scalar indices that do not use tensor orientational information), because, although δ_2 does not perform well on average, $\bar{d}$ = 5.87mm (Table 8.1), δ_1 finds the smallest mean Euclidean distance ($\bar{d}$ = 4.19mm) between the warped-source and target images for all five subjects. The tensor difference, δ_3, the normalised tensor scalar product, δ_6, and principal-direction difference, δ_7, perform similarly, with $\bar{d}$ = 4.47mm, $\bar{d}$ =4.63mm and $\bar{d}$ = 4.48mm, respectively, but do not perform as well as δ_1 despite the additional orientational in-

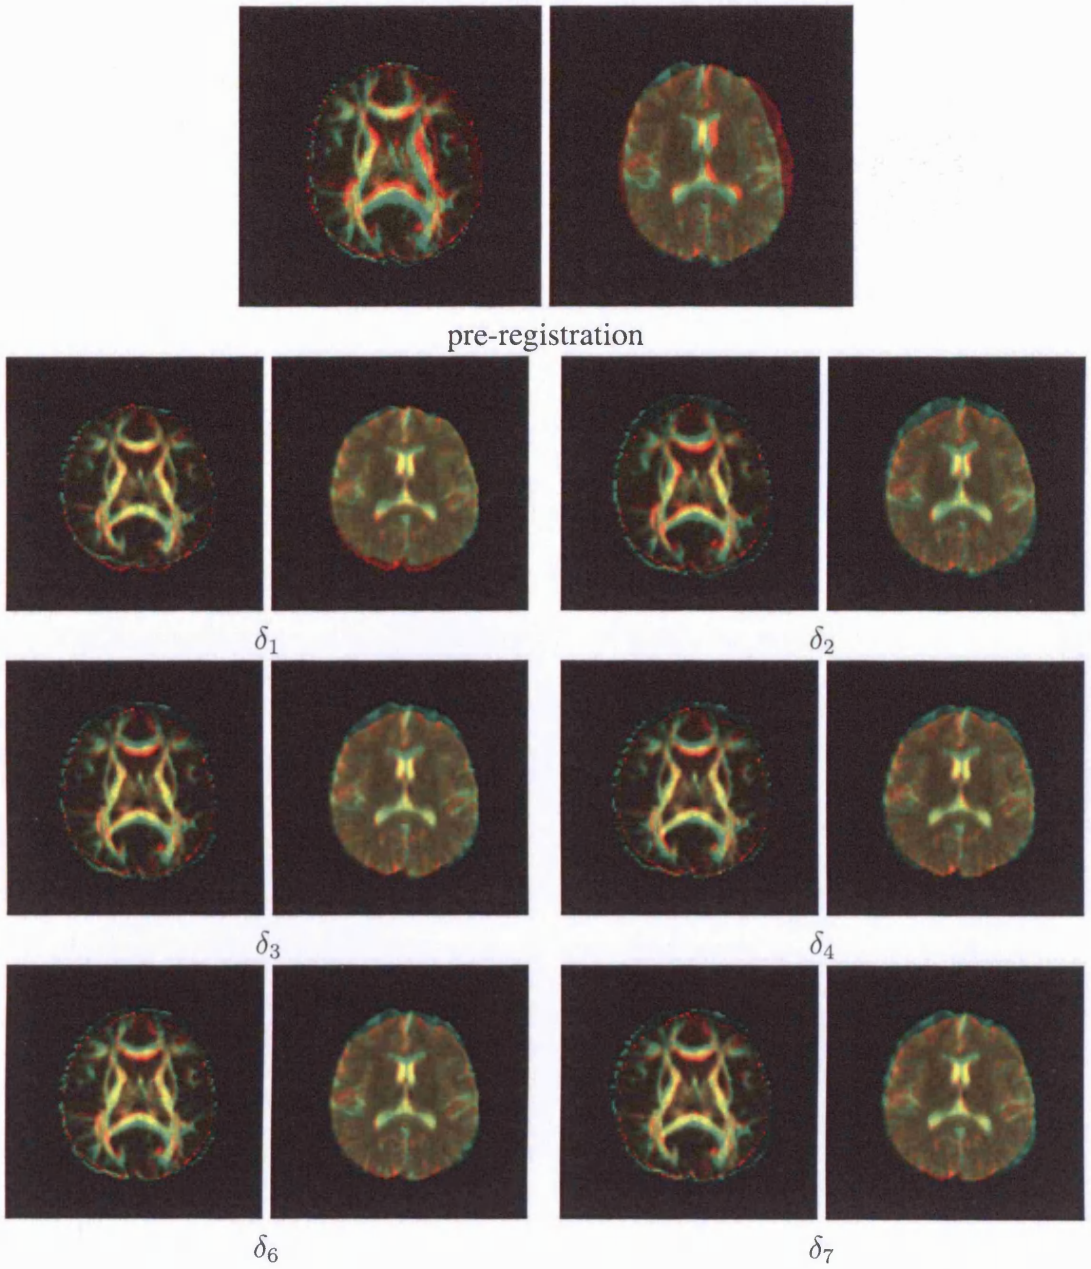

Figure 8.3: Anisotropy and trace maps for one slice of a source (subject two) and target volumetric DT-MRI image, for pre-registration, relative-anisotropy difference, δ_1, modulus difference, δ_2, tensor difference, δ_3, normalised tensor difference, δ_4, normalised tensor scalar product, δ_6, and principal-direction difference, δ_7, registration results.

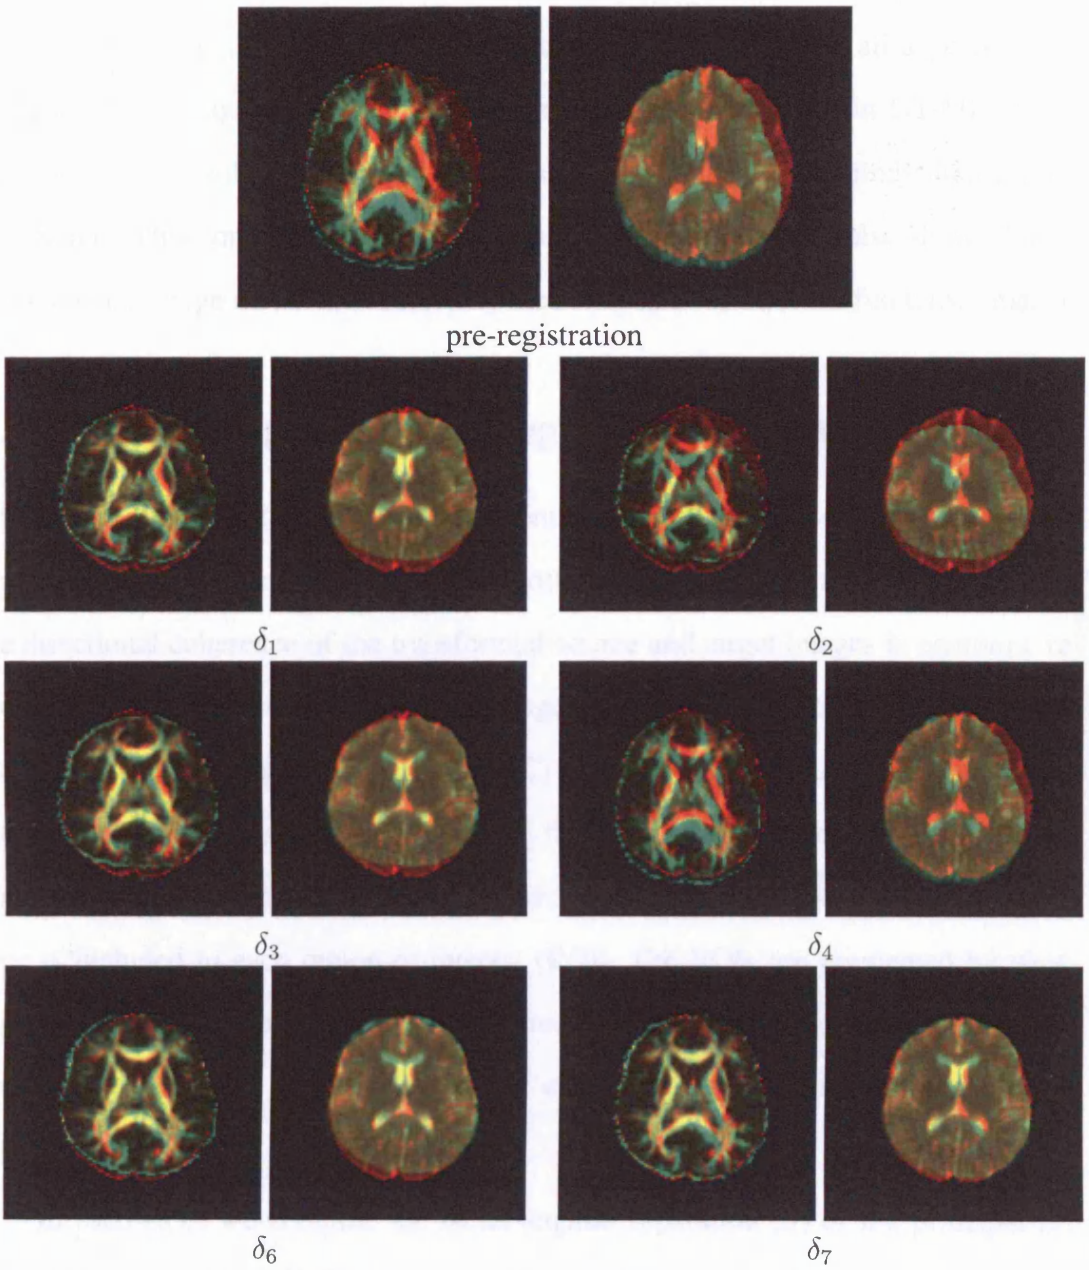

Figure 8.4: Anisotropy and trace maps for one slice of a source (subject four) and target volumetric DT-MRI image, for pre-registration, relative-anisotropy difference, δ_1, modulus difference, δ_2, tensor difference, δ_3, normalised tensor-difference, δ_4, normalised tensor scalar product, δ_6, and principal-direction difference, δ_7, registration results.

formation they exploit to guide the DT-MR registration. Furthermore, normalising the tensor-difference, δ_4, to emphasise differences in shape and orientation does not prove advantageous ($\bar{d} = 5.30$mm).

We expected to need similarity measures that are sensitive to all aspects of the diffusion tensor, including orientation, to exploit the information in DT-MRI fully; however, these results suggest that δ_1, the scalar-derived measure, finds the best image match. This conclusion supports our earlier work [42], which also shows that δ_1 finds a better image match than similarity measures derived from the full tensor matrix.

8.3 Affine Registration: Region-Based Statistics

To test hypothesis C2 (diffusion-tensor orientation matching improves the accuracy of registration over that from registration algorithms that ignore orientation), we examine the directional coherence of the transformed-source and target images in anatomic regions outlined by hand in the template image. The selected white-matter regions are the posterior limb of the internal capsule (IC) and the corpus callosum (CC), where the diffusion tensor is consistently prolate; and the pons, where the tensor is consistently oblate. The hand segmentation is conservative to ensure that only the chosen tissue type is included in each region of interest (ROI). The ROIs are segmented by viewing all three planes simultaneously. All three regions are best identified on principal direction colour maps (Figure 8.5). Small errors in tensor alignment indicate better registration.

In each ROI, we compute the mean angular separation [5] of the principal DT eigenvectors over the ROI. The mean angular separation is defined as

$$E_i = \left(\sum \sqrt{\nu_1 \nu_2}\right)^{-1} \left(\sum \left(\sqrt{\nu_1 \nu_2} \times \cos^{-1} |\mathbf{e}_{i1} \cdot \mathbf{e}_{i2}|\right)\right) \tag{8.1}$$

where ν_1 and ν_2 are the anisotropies of the two diffusion tensors, $\mathbf{e}_{i1}$ and $\mathbf{e}_{i2}$ are the i-th principal eigenvectors and the sums in Eq. 8.1 are over all the voxels in the ROI. The principal direction $\mathbf{e}_1$ is most strongly defined for prolate DTs, while the minor

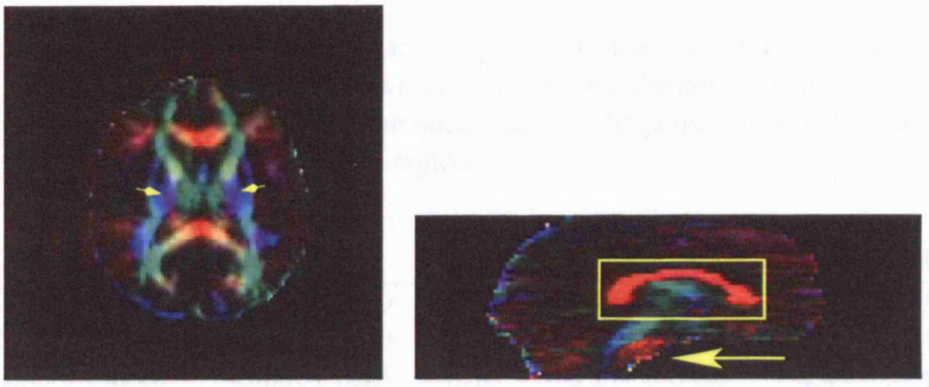

(a) (b)

Figure 8.5: Principal direction colour maps to illustrate the ROIs. Image (a) shows the posterior limb of the internal capsule as two broad symmetric dark blue fibres (arrowheads). A mid-sagittal section of the corpus callosum can be seen in image (b) as the red "C"-shaped region (rectangle). The other red-coloured region at the base of the brain in (b) is the pons (arrow).

eigenvector $\mathbf{e}_3$ is meaningful in regions of oblate DTs. We use a measure of prolateness [136] for ν in Eq. 8.1 defined by:

$$(\nu_p) = \left(\mathrm{Tr}(\mathsf{D}^2)\right)^{-\frac{1}{2}} (\lambda_1 - \lambda_2) \tag{8.2}$$

when computing the directional coherence of $\mathbf{e}_1$, E_1, and a measure of oblateness [136] for ν in Eq. 8.1 defined by:

$$(\nu_o) = \left(\mathrm{Tr}(\mathsf{D}^2)\right)^{-\frac{1}{2}} 2\,(\lambda_2 - \lambda_3) \tag{8.3}$$

when computing the directional coherence of $\mathbf{e}_3$, E_3, where $\lambda_1 > \lambda_2 > \lambda_3$ are the eigenvalues of the diffusion tensor D.

For the two prolate regions, the corpus callosum (CC) and the internal capsule (IC), we compute E_1 since we expect the principal direction, $\mathbf{e}_1$, to be aligned, while in the oblate region, the pons, we compute E_3 since we expect the minor eigenvector, $\mathbf{e}_3$, to be aligned. Table 8.2 shows the mean angular separations in each ROI for the five subjects registered using each of the reorientation strategies. Low numbers indicate good coherence and high numbers indicate poor alignment of the principal direction

Table 8.2: The average of mean angular separations for the five subjects for affine registration. The table shows E_1 for the two prolate regions, the corpus callosum (CC) and the internal capsule (IC) and E_3 for the oblate region, the pons. The results highlighted in bold are the lowest scores for each region.

	CC (E_1)			IC (E_1)			Pons (E_3)		
	NR	**FS**	**PPD**	**NR**	**FS**	**PPD**	**NR**	**FS**	**PPD**
δ_1	0.21	0.20	0.21	0.20	0.17	**0.14**	0.57	0.43	0.49
δ_2	0.37	0.36	0.41	0.20	0.23	0.37	0.67	0.63	0.70
δ_3	0.33	0.20	0.23	0.25	0.18	0.21	0.64	0.49	0.61
δ_4	0.26	0.26	0.30	0.34	0.32	0.33	0.61	0.64	0.60
δ_6	0.17	0.18	0.19	0.15	0.17	0.16	0.50	0.51	**0.37**
δ_7	**0.15**	**0.15**	0.26	0.20	0.18	0.20	0.63	0.66	0.72

between tensors in the source and target images. In the following sections we discuss the mean results for the five subjects, for each similarity measure, with each reorientation strategy and in each ROI.

8.3.1 Experiment HC2: Relative-Anisotropy Difference

Since the relative-anisotropy difference, δ_1, does not depend on the relative orientations of diffusion tensors, any difference in the results for δ_1 using no reorientation (NR), finite strain (FS) or preservation of principal directions (PPD) comes mainly from the improvement in image warping rather than better orientation matching. All three reorientation strategies find less than 0.3% difference in the mean angular separations for all three ROIs, except for two subjects (Appendix B). The registrations for subjects four and five are poorly aligned, especially the IC region on the left-side of the brain for subject four (Figure 8.4), where $E_1 = 0.40$ for affine registration with NR, $E_1 = 0.42$ with FS reorientation and $E_1 = 0.25$ with PPD reorientation. We may expect some differences in the transformations computed using the different reorientations, since different reorientation will produce tensors with different relative anisotropies at non-integer voxel locations computed using tri-linear interpolation.

8.3.2 Experiment HC2: Modulus Difference

Similarily, modulus difference, δ_2, does not depend on the relative orientations of diffusion tensors. The mean angular separation of the warped-source and target images for

the CC and pons regions are smaller using FS reorientation (E_1 = 0.36 and E_3 = 0.63) compared with no reorientation (E_1 = 0.37 and E_3 = 0.67) and PPD (E_1 = 0.41 and E_3 = 0.70). For the IC region registration with no reorientation (E_1 = 0.20) performs better than both FS (E_1 = 0.23) and PPD (E_1 = 0.37). The quality of the image registrations is not very good for this similarity measure and this poor alignment gives rise to the wide variation in angular separation for the three reorientation strategies.

8.3.3 Experiment HC2: Tensor Difference

The FS and PPD methods perform better than no reorientation for all three regions, which is the expected pattern for tensor difference, δ_3. However, registration with FS finds smaller mean angular separations (E_1 = 0.20, E_1 = 0.18 and E_3 = 0.49) than PPD (E_1 = 0.23, E_1 = 0.21 and E_3 = 0.61). Although gradient annealing attempts to find the global minimum, it is still possible to get stuck in local minima close to the global minimum. This effect may explain the unexpectedly low performance of registration using PPD.

8.3.4 Experiment HC2: Normalised Tensor-Difference

The NR and FS strategies produce similar performance (E_1 = 0.26), which is slightly better than PPD (E_1 = 0.30) for the CC region for normalised tensor difference, δ_4. The FS and PPD methods perform better (E_1 = 0.32 and 0.33, respectively) than NR (E_1 = 0.34) for the IC region and PPD (E_3 = 0.60) performs better than both FS (E_3 = 0.64) and NR (E_3 = 0.61) for the pons region. Closer visual inspection reveals that the quality of registration obtained from the affine transformation is not as good for the CC region compared with the IC region for all subjects. The poor alignment in this region gives rise to the higher than expected angular separations.

8.3.5 Experiment HC2: Normalised Tensor Scalar Product

Table 8.2 shows that registration using normalised tensor scalar product, δ_6, with NR produces similar performance to FS and PPD reorientation for the CC (E_1 = 0.17, 0.18 and 0.19, respectively) and IC (E_1 = 0.15, 0.17 and 0.16, respectively) regions. However, for the pons region, registration with PPD reduces E_3 consistently and by

as much as 50%, see Appendix B for individual subject results. The mean angular separation over all five subjects for registration with NR, FS and PPD is $E_3 = 0.50$, $E_3 = 0.51$ and $E_3 = 0.37$, respectively. Registration with PPD also performs consistently better than FS, except for one subject where $E_3 = 0.50$ and $E_3 = 0.51$, for FS and PPD, respectively. Table 8.2 shows that registration using δ_6 performs better for oblate regions than the other similarity measures tested.

8.3.6 Experiment HC2: Principal-Direction Difference

Registration using principal-direction difference, δ_7, with FS and NR produce similar performance ($E_1 = 0.15$), which is better than the PPD method ($E_1 = 0.26$) for the CC region. However, registration with FS ($E_1 = 0.18$) finds smaller average angular separations than registration with NR ($E_1 = 0.20$) and PPD ($E_1 = 0.20$) for the IC region. For the pons region, no reorientation ($E_3 = 0.63$) performs better than both registration with FS ($E_3 = 0.66$) and PPD ($E_3 = 0.72$). These higher than expected angular separations may be caused by local minima problems.

8.3.7 Qualitative Results

Here we provide a qualitative comparison of the results obtained from different similarity measures. Although some differences can be observed from registrations using different similarity measures, it is impossible to conclude which measures are the most effective by visual inspection. However, in order to give some idea of the quality of the registration that is obtained, Figure 8.6 illustrates an example where the normalised tensor-difference, δ_4, the tensor difference, δ_3, with finite strain, FS, reorientation and the relative-anisotropy difference, δ_1, find similar qualitative results.

Figure 8.6 illustrates that similarity measure, δ_3, does not find a good match for subject one. However, normalising δ_3 or applying FS reorientation improves the registration. The normalised tensor difference measure, δ_4, with FS reorientation looks similar to that without reorientation (not shown). It is interesting to note that δ_1 finds a better registration than any of the other similarity measures. Table 8.3 compares the mean angular separations for the example illustrated in Figure 8.6 for each similar-

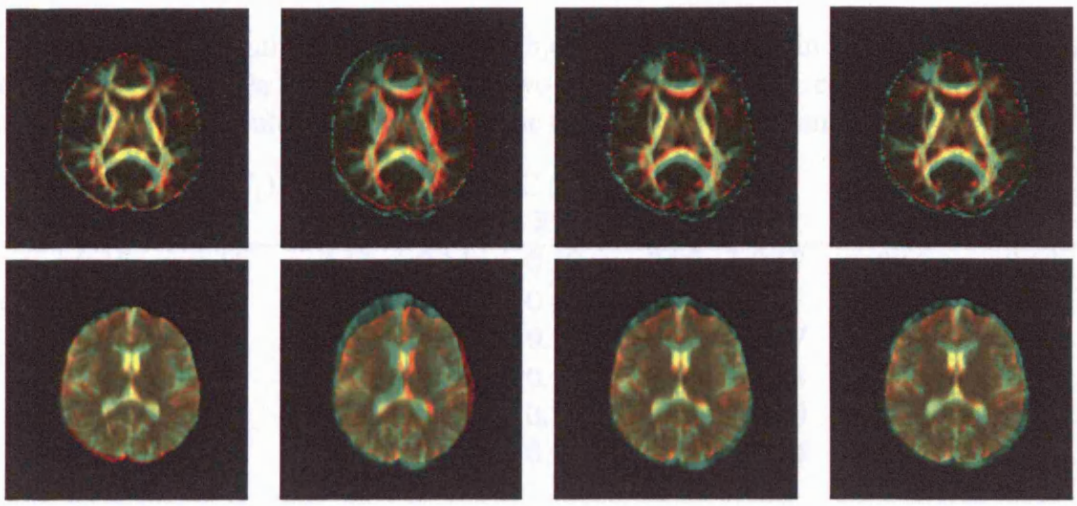

Figure 8.6: Gradient annealing registration results for subject one. Anisotropy maps (top row) and trace maps (bottom row) are shown for the relative-anisotropy-difference, δ_1, with no reorientation, NR, the tensor-difference, δ_3 with NR and with finite-strain, FS, reorientation and the normalised-tensor-difference, δ_4, with NR.

ity measure. The numerical values in bold typeface show the dramatic improvement in registration performance using the tensor difference, δ_3, with either FS or PPD reorientation compared with registration with no reorientation (NR). The mean angular separation for the CC region for registration with NR is $E_1 = 0.58$, while registration with FS and PPD improves registration performance by as much as 75%, $E_1 = 0.14$ and $E_1 = 0.29$, respectively. Although we use gradient annealing, local minima are still possible and this may explain the lower performance of PPD compared with FS in this region. Registration with FS and PPD produce similar results for the IC region, $E_1 = 0.18$, which improves registration performance by 64%, compared with registration with NR ($E_1 = 0.51$). Registration performance in the pons region is also improved by as much as 37% with FS and PPD, $E_3 = 0.61$ and $E_3 = 0.68$, respectively.

The results in Table 8.3 show that although normalising the δ_3 measure, i.e. using δ_4, produces much better registration performance for the CC, IC and pons regions, ($E_1 = 0.19$, $E_1 = 0.15$ and $E_3 = 0.74$, respectively), compared with registration using δ_3 ($E_1 = 0.58$ and $E_1 = 0.51$ and $E_3 = 0.97$, respectively), registration using δ_3 with FS reorientation performs similarly to δ_4 with no reorientation.

Table 8.3: Mean angular separation for subject one, illustrated in Figure 8.6, for affine registration. The table shows E_1 for the two prolate regions, the corpus callosum (CC) and the internal capsule (IC) and E_3 for the oblate region, the pons.

	CC (E_1)			IC (E_1)			Pons (E_3)		
	NR	**FS**	**PPD**	**NR**	**FS**	**PPD**	**NR**	**FS**	**PPD**
δ_1	0.15	0.16	0.13	0.13	0.10	0.09	0.57	0.60	0.62
δ_2	0.28	0.30	0.19	0.14	0.13	0.16	0.97	0.85	0.92
δ_3	**0.58**	**0.14**	**0.29**	**0.51**	**0.18**	**0.18**	**0.97**	**0.61**	**0.68**
δ_4	0.19	0.14	0.17	0.15	0.15	0.16	0.74	0.65	0.75
δ_6	0.19	0.17	0.15	0.08	0.20	0.09	0.59	0.50	0.51
δ_7	0.10	0.08	0.09	0.18	0.18	0.18	0.94	0.95	0.99

Table 8.2 suggests that the principal-direction difference, δ_7, often gives lowest E_1, particularly in the CC region. Table 8.3 shows that registration using δ_7 performs much better in the CC region, $E_1 = 0.08$, than any of the other similarity measures, which is reasonable to expect because minimising δ_7 minimises E_1 more directly than minimising any of the other similarity measures. Registration using δ_1 with FS reorientation finds the best directional coherence for the internal capsule region ($E_1 = 0.10$) and δ_1 with no reorientation produces the best registration performance for the pons region ($E_3 = 0.57$). The results of the similarity measure comparison (Table 8.1 and Figure 8.2) also show that δ_7 and δ_1 produce better image matches, $\bar{d} = 3.41$ and $\bar{d} = 3.45$, respectively, than the other similarity measures for subject one.

8.3.8 Conclusion

The affine registration is not reliable enough to make a firm conclusion about hypothesis (C2). Although gradient annealing attempts to find the global minimum, the problem with local minima persists, which causes poor alignment of the regions of interest, giving rise to wide variation in directional coherence results for the three reorientation strategies. These results motivate the higher-order experiments in the next section.

8.4 Higher-Order: Region-Based Statistics

To test hypothesis C2 (diffusion-tensor orientation matching improves the accuracy of registration over that from registration algorithms that ignore orientation), for higher-

order registration we also examine the directional coherence of the transformed-source and target images in anatomic regions outlined by hand in the template image, as described in §8.3. This is only an initial investigation of possible extra insights we might get from higher-order experiments. We perform affine, order-two, order-three and order-four registrations but only perform the higher-order registrations for the tensor-difference, δ_3, similarity measure.

For affine registration, we impose a time limit on registration optimisations in order to compute results in a reasonable amount of time, as described in Chapter 7. For each step size, we identify an appropriate freezing temperature, T_f, and, using this T_f, we test different combinations of starting temperatures, T_0, and cooling factors, c, to determine the settings that find the best registration within our time-limit of approximately 24 hours of computation time. We find that affine registration with a step size of eight, an initial starting temperature, T_0, of 5.92×10^{-12}, a freezing temperature, T_f, of 10^{-15} and a cooling factor, c, of 0.9728 gives the lowest mean Euclidean distance, $\bar{d}$ = 3.6mm, between the landmarks in the warped-source and target images. In order to compute higher-order registration results in a reasonable amount of time, we also select a step size of eight, $T_0 = 5.92 \times 10^{-12}$, $T_f = 10^{-15}$ but we reduce the cooling factor to $c = 0.8729$. We do not impose a time limit of 24 hours for the higher-order experiments because they would not terminate in manageable time, i.e. order-two registration takes approximately four times as long as affine gradient annealing registration.

8.4.1 Results

Table 8.4 shows the mean angular separations for affine, order-two, order-three and order-four registration using the tensor difference, δ_3, similarity measure. We first consider the two prolate regions, the corpus callosum (CC) and the internal capsule (IC). The most striking results in Table 8.4 are the unexpectedly high angular separations for affine and order-two registration with PPD reorientation. For affine registration $E_1 = 0.77$ and $E_1 = 0.71$ for the CC and IC regions, respectively. Order-two registration also produces an unusually high angular separation of $E_1 = 0.55$ for the CC region.

The serious local minima problem may explain this effect, where insufficient time is spent searching for the global minimum. In §8.3 affine registration with PPD reorientation for the same subject produces significantly smaller mean angular separations, E_1 = 0.29 and E_1 = 0.18, respectively for the CC and IC regions (Table 8.3).

Table 8.4: Mean angular separations for subject one, for higher-order registration using tensor difference, δ_3. The table shows E_1 for the two prolate regions, the corpus callosum (CC) and the internal capsule (IC) and E_3 for the oblate region, the pons. The table shows results for order-one, order-two, order-three and order-four. The results highlighted in bold are the lowest scores for each region.

	CC (E_1)			IC (E_1)			Pons (E_3)		
Ord	**NR**	**FS**	**PPD**	**NR**	**FS**	**PPD**	**NR**	**FS**	**PPD**
1	0.29	0.21	0.77	0.14	0.19	0.71	**0.38**	0.64	0.76
2	0.32	0.23	0.55	0.17	0.15	0.15	1.04	1.09	0.95
3	0.23	0.13	0.18	0.18	0.17	0.18	0.73	0.69	1.11
4	**0.10**	0.11	**0.10**	0.18	**0.13**	0.18	0.86	0.94	0.90

Order-two, order-three and order-four registration with FS and PPD produce similar or better results than NR for the two prolate regions, except for order-two registration with PPD, in the CC region, where the unexpectedly poor performance was presumably caused by local minima. Order-four registration with PPD and no reorientation (NR) produce the best alignment of the principal direction between the source and target images for the CC region, E_1 = 0.10, which is significantly better than affine registration with NR, E_1 = 0.29, but only slightly better than order-four registration with FS reorientation, E_1 = 0.11. Order-four with FS reorientation produced the smallest mean angular separation E_1 = 0.13 for the IC region but this is similar to that produced with affine registration with NR (E_1 = 0.14).

Interestingly, the best alignment for the pons region is produced with affine registration and NR (E_3 = 0.38). Closer examination of the images (Figure 8.7) reveals that although there is better alignment in the pons region for affine registration with NR (E_3 = 0.38) compared with order-four registration with FS reorientation (E_3 = 0.94), the quality of both registrations for this region is poor. Figure 8.7 (a), (b), (c) and (d) show mid-sagittal and mid-coronal principal direction colour maps of the target and source

images. Figure 8.7 (e), (f), (g) and (h) show anisotropy maps of the warped-source (green) and target (red) images superimposed so that the images appear yellow when they are well aligned. Figure 8.7 (e) and (f) show the registration for subject one for affine with no reorientation and (g) and (h) show the order-four registration for subject one with FS reorientation. Comparing (e) to (g) and (f) to (h), we can see that the order-four registration with FS reorientation finds a better image match, but it is difficult to determine visually how well the pons region is aligned. Figure 8.7 shows improvement in the alignment of the CC region when we perform order-four registration with FS reorientation compared with affine registration with NR. We compare these results qualitatively because they are the most dramatic. Affine registration with no reorientation produces the smallest angular separation for the pons region $E_3 = 0.38$, both registrations perform similarly for the IC region, but order-four registration with FS registration finds a significantly better registration, $E_1 = 0.11$, compared with affine registration with NR, $E_1 = 0.29$.

8.4.2 Conclusion

There are serious problems with local minima for affine and order-two registration with PPD reorientation. Order-two, order-three and order-four registration with FS and PPD reorientation perform similar or better than registration with no reorientation, for both prolate regions, except for the local minima problem for order-two registration with PPD reorientation. Registration with reorientation improves the directional coherence of the CC region by as much as 43% ($E_1 = 0.23$ and $E_1 = 0.13$) for order-three registration with NR and order-three registration with FS reorientation, respectively. Applying FS reorientation also improves results for the IC region for order-two, order-three and order-four by as much as 27%. Although affine registration with NR performs better than any of the other higher-order experiments in the pons region, this region is poorly aligned for all registration results. These results do not support hypothesis C2 (registration using FS and PPD should be better than using no reorientation, since these methods allow orientations to be matched properly) for higher-order registration.

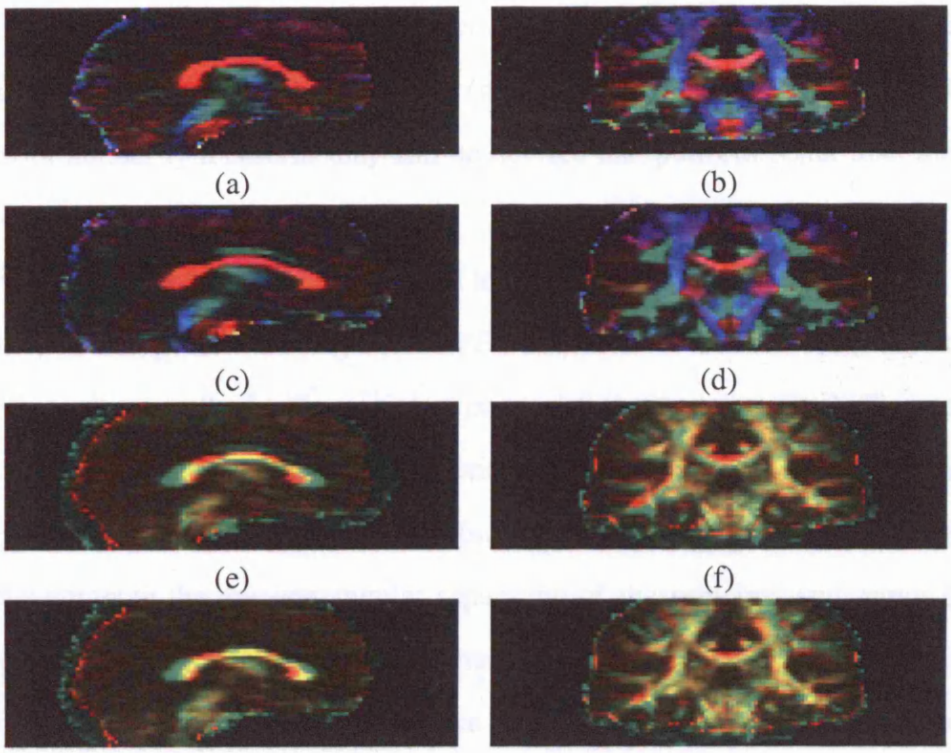

Figure 8.7: Affine and order-four polynomial registration results. Principal direction and anisotropy maps are shown for the sagittal (left column) and coronal (right column) image planes. Images (a) and (b) are slices from the target image where the pons is visible (small red region at the base of the brain) and (c) and (d) are the same slices from the source image. Images (e) and (f) show the warped-source (green) for affine registration with no reorientation (NR) and the target (red) superimposed so that when the images are well aligned they appear yellow. Images (g) and (h) show the warped-source for order-four registration with FS reorientation.

8.5 Summary

To test hypothesis HC1 (similarity measures derived from the full tensor matrix find a more accurate registration than those based on the derived scalar indices that do not use tensor orientational information), we compared the gradient annealing registration performance of several similarity measures both qualitatively and quantitatively. To compare the measures quantitatively we calculated the mean Euclidean distance $\bar{d}$ between fifty-two corresponding landmarks in the warped-source and target images. We reject hypothesis C1 because the relative-anisotropy difference, δ_1, a scalar-derived measure, finds a smaller mean distance $\bar{d}$ (4.19mm) than all the measures derived from the full tensor matrix, i.e. the tensor-difference, δ_3, ($\bar{d}$ = 4.47mm), the normalised tensor-

difference, δ_4, ($\bar{d}$ = 5.30mm), the normalised tensor scalar product, δ_6, ($\bar{d}$ = 4.63mm) and the principal-direction difference, δ_7, ($\bar{d}$ = 4.48mm). We note however that other full-tensor similarity measures may still be devised that perform better than those we have tested.

We use gradient-annealing affine and higher-order polynomial registration to test hypothesis C2 (registration using FS and PPD should be better than using no reorientation, since these methods allow orientations to be matched properly.) To assess the registration quality, we examine the directional coherence of the transformed-source and target images in anatomic regions outlined by hand in the template image. In each ROI, we compute the average angular separation of the principal and minor DT eigenvectors over the ROI. We reject hypothesis C2 because the affine and higher-order registrations are not reliable enough to make a firm conclusion about C2. The results are inconclusive because of problems with local minima and misalignments of the regions of interest. There is no guarantee that gradient annealing finds the global minimum, although it should find a minima close to the global one. It is possible that local minima prevent registration with reorientation from performing better than no reorientation and that there are many more local minima when matching with FS and PPD compared with no reorientation. The improvement in E_1, through the use of reorientation, is more significant than E_3, although it is interesting that δ_6 with PPD reorientation performs well in oblate regions. The anisotropy of oblate regions is much less than prolate regions; therefore the eigenvector directions are less clearly defined and so exhibit less consistency between corresponding regions in the images.

Chapter 9

Conclusions

This chapter gives a summary of the conclusions we draw in the thesis and discusses future diffusion-tensor registration techniques and validation methodologies.

9.1 Discussion

The motivation for using diffusion-tensor data to drive registration is that the orientational information provides new cues for matching that may improve performance. Matching orientations places much greater constraints on the registration transformation than when scalar images are matched.

The experiments with Powell's direct optimisation, described in Chapter 6, demonstrate the increased local minima in the registration objective function when using full-tensor similarity rather than scalar similarity. This observation motivated the use of a global optimisation algorithm to decide more concretely whether diffusion tensor orientation matching improves the accuracy of registration over that from registration algorithms that ignore orientation.

To avoid the local minimum problem with direct optimisation, we implemented simulated annealing [107], a global optimisation technique, to optimise the starting point for Powell's method [107]. We call the combined method *gradient annealing*. In theory, simulated annealing always finds the global minimum of an objective function [64], but in practice the time taken to find the global minimum can be infinite. We impose a time limit on registration optimisations in order to compute results in a

reasonable amount of time and optimise the algorithm to the best performance within that fixed time frame. We select the optimal settings for registration with gradient annealing by minimising the mean Euclidean distance $\bar{d}$ between fifty-two hand-defined corresponding anatomic landmarks in one example pair of warped-source and target images.

We first discuss the direct optimisation experiments that were prompted by problems with local minima and then address the two research objectives outlined in Chapter 1.

Can direct optimisation methods compute the registration transformation reliably?

Powell's optimisation finds a consistent transformation for **intra-subject** registration for translations and rigid transformations if the translation parameters in each dimension are drawn from [-2, 2]mm and the rotation angle from $[0, \pi/6]$. Although Gaussian smoothing of the source and target image prior to registration improves the consistency of the registration transformation, Powell's direct optimisation does not find a consistent transformation for inter-subject registration (Chapter 6). We find that gradient annealing reduced the objective function consistently and by as much as 21% (Chapter 7).

Are similarity measures derived from the full tensor more effective than measures based on derived scalar indices?

Using gradient annealing registration, we compare the performance of several similarity measures both qualitatively and quantitatively. To compare the measures quantitatively, we calculated the mean Euclidean distance $\bar{d}$ between fifty-two corresponding landmarks in five warped-source images and the target image. We reject hypothesis C1 (similarity measures derived from the full tensor matrix find a more accurate registration than those based on the derived scalar indices that do not use tensor orientational information) because the relative-anisotropy difference, δ_1, a scalar-derived measure, finds a smaller mean distance $\bar{d}$ = 5.8mm than all the measures derived from the full tensor matrix, for example, the tensor-difference, δ_3, ($\bar{d}$ = 6.14mm) and the principal-

direction difference, δ_7, ($\bar{d}$ = 6.2mm).

Does diffusion-tensor orientation matching improve the accuracy of registration over that from registration algorithms that ignore orientation?

To compare the potential of orientation-based methods fairly with standard methods we use gradient annealing to get closer to the global minimum of the registration objective function. Registration using finite-strain (FS) and preservation of principal direction (PPD) should be better than using no reorientation (NR), since these methods allow orientations to be matched properly. Furthermore, PPD should produce better results than FS, since it reorients tensors more accurately.

To assess the registration quality, we examine the directional coherence of the transformed-source and target images in anatomic regions outlined by hand in the template image. In each ROI, we compute the mean angular separation [5] of the principal DT eigenvectors over the ROI.

Affine registration is not reliable enough to make a firm conclusion about hypothesis C2 (diffusion tensor orientation matching improves the accuracy of registration over registration algorithms that ignore orientation). The local minima problem persists despite the use of global optimisation. These minima are caused either by interpolation artifacts or by a local good match between features or intensities, which causes poor registration of the regions of interest. Although the results show improvements with orientation matching of up to 50% in the pons region and 39% in the prolate regions, these improvements are not consistent. There is no guarantee that gradient annealing finds the global minimum, although it should find a minimum close to the global one. This may explain why registration with FS or PPD reorientation does not consistently find the best directional coherence in the ROIs. It is possible that local minima traps prevent FS and PPD reorientation from performing better than no reorientation and that there are many more local minima when matching with PPD compared with FS, and with FS compared with NR.

The higher-than-expected angular separations for the FS and PPD methods are caused by misalignments of the regions of interest. This motivated an initial investiga-

tion of possible extra insights we might get from higher-order registration experiments using the tensor difference, δ_3, similarity measure. Order-two, three and four polynomial registration with either FS or PPD performed similar or better than no reorientation for both prolate regions, except for order-two registration with PPD reorientation, where a local minima problem may explain the unexpectedly low registration performance in the corpus callosum region. However, the poor alignment in the pons region for all registrations gives rise to the higher than expected angular separations of the minor eigenvector. Interestingly, the best alignment for the pons region was found with affine registration and no reorientation, which can be attributed to poor registration in the other experiments.

9.2 Future Outlook

The development of DT-MRI registration methods is at a relatively early stage and there are various avenues for future research. The results in this thesis highlight the need for fast sophisticated optimisation algorithms to overcome the increased local minima in tensor registration objective functions over scalar registration. We can only expect minor improvements in orientation matching with low dimensional transformations. Higher dimensional transformations will show the true potential of this technique. Although we reject the hypothesis that similarity measures derived from the full tensor matrix find a more accurate registration than those based on the derived scalar indices, we note that other full-tensor similarity measures may still be devised that perform better that those we have tested. Batchelor et al. [18] use Riemannian geometry that takes into account the positive definitiveness of diffusion tensors to define a distance between diffusion tensors. This could be used as a similarity measure for DT registration. Other possible areas of future research include the extension of current DT-MRI registration techniques to use the additional information in the output of multiple-fibre reconstructions [3]. Further progress in performance evaluation is required before these techniques are used clinically. Other areas of future research are discussed below.

The results in the thesis have shown that a more sophisticated registration al-

gorithm is necessary to achieve accurate alignment of the tensor images in order to assess whether orientation matching improves registration over registration algorithms that ignore orientation. An alternative to using polynomials is to break the image up into regions that can move independently of each other. Following an approximate alignment of the source and target images using a rigid or affine transformation, the images can be split up into an array of nodes. These nodes may be the control points of a spline such as a B-spline [114], or may be regarded as the centre of a local region of interest (ROI). In [114] Rueckert et al. model the global motion of the breast by an affine transformation while the local breast motion is described by a free-form deformation (FFD) based on B-splines. Each node in the source image is iteratively translated, or a local affine transformation is iteratively determined for each ROI, in order to optimise the objective function.

To improve the efficiency of the registration algorithm we could adopt a multi-resolution or scale-space framework. A multi-resolution framework starts at a coarse resolution of n mm (where the volume is sub-sampled such that the voxel side-length becomes n mm). We would then find a solution with the optimisation method and then refine the solution by progressively increasing the resolution (decreasing the value of n). For instance, a method may start with n = 8mm and then progress to n = 4mm and then finally to n = 1mm. An advantage of this approach is that the initial optimisation, at large n, has a dramatically reduced computational load, since the number of sample points is substantially less. In addition, for large sub-samplings the gross features of the image should dominate and so the overall alignment should be easier to find. The idea of this hierarchical framework is that at the highest resolution, the global minimum is the nearest minimum to the starting point and can be found by local optimisation. In practice, this requires that the starting estimate of the registration transformation is reasonably close to the correct solution. However, sub-sampling to lower resolutions may not reduce the number of local minima sufficiently. Vaarkamp et al. [131] adopt a hierarchical multi-resolution approach to match CT images for radiotherapy treatment planning and find that this approach dramatically increases both the speed

of their algorithm and the capture range, i.e. portion of the parameter space in which the algorithm is more likely to converge to the correct minimum. Jenkinson et al. [75] find that the use of Powell's optimisation [107] method together with a multi-resolution approach is not sufficient to find the global minimum reliably. Pluim et al. [104] perform some work on interpolation [104] and show how optimising in a multi-resolution framework can actually create additional minima.

To assess the performance of registration using diffusion tensor images, we could compute fibre tracts from corresponding voxels in the target and warped-source images and use the similarity of the trajectories to obtain a measure of registration quality, as proposed by Park et al. [100].

Appendix A

Annealing Temperature Range

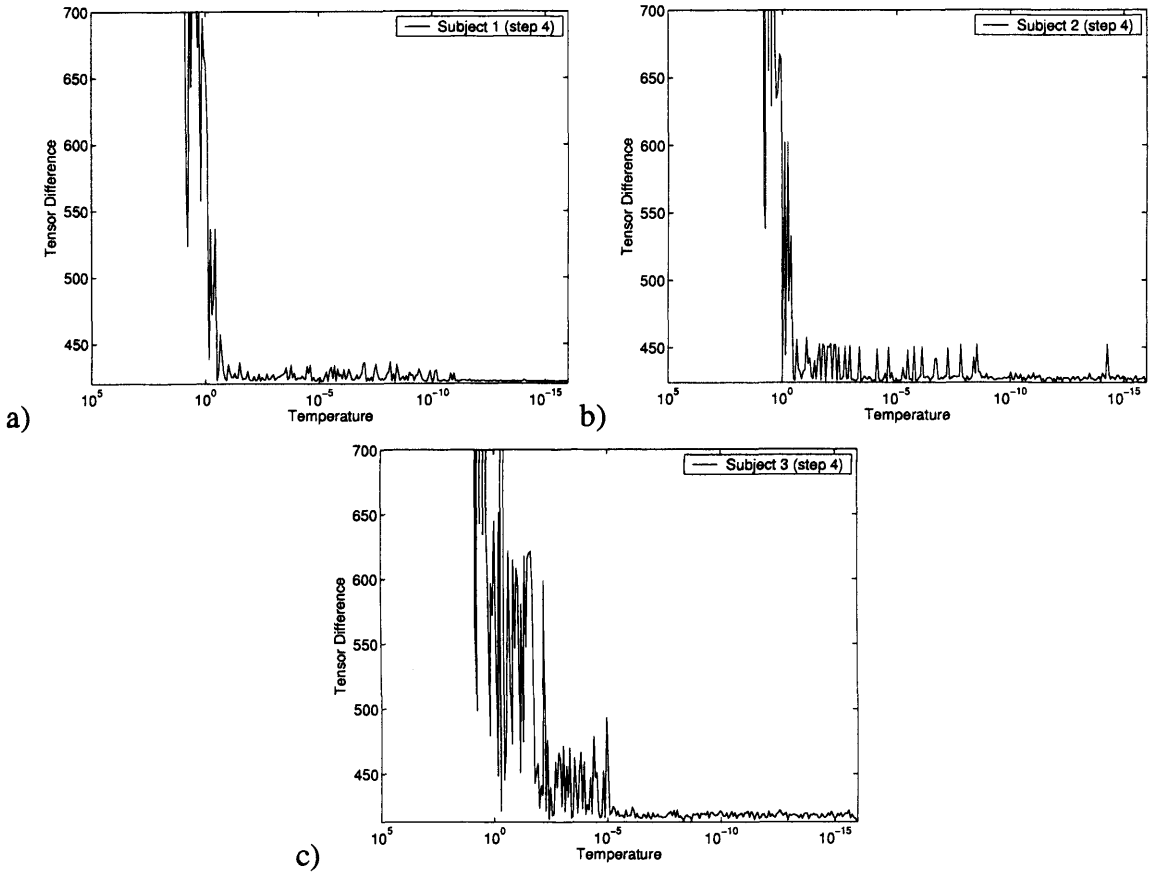

Figure A.1: Plot of the tensor difference similarity measure against temperature, where the temperature has been initialised to 10^5 and the freezing point is 10^{-15}. Figures (a)-(c) come from experiments with a step size of 4 and are the results for 3 data sets.

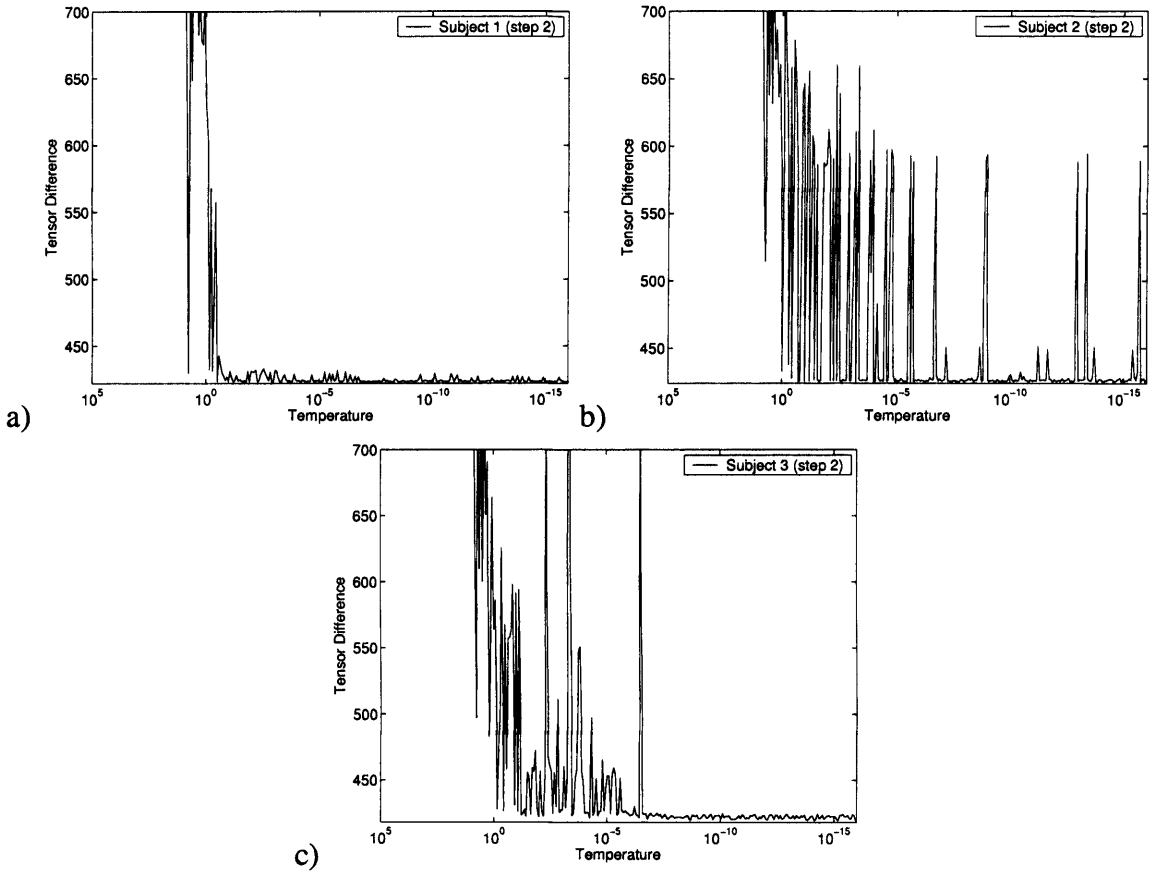

Figure A.2: Plot of the tensor difference similarity measure against temperature, where the temperature has been initialised to 10^5 and the freezing point is 10^{-15}. Figures (a)-(c) consider a step size of 2 and are the results for 3 data sets.

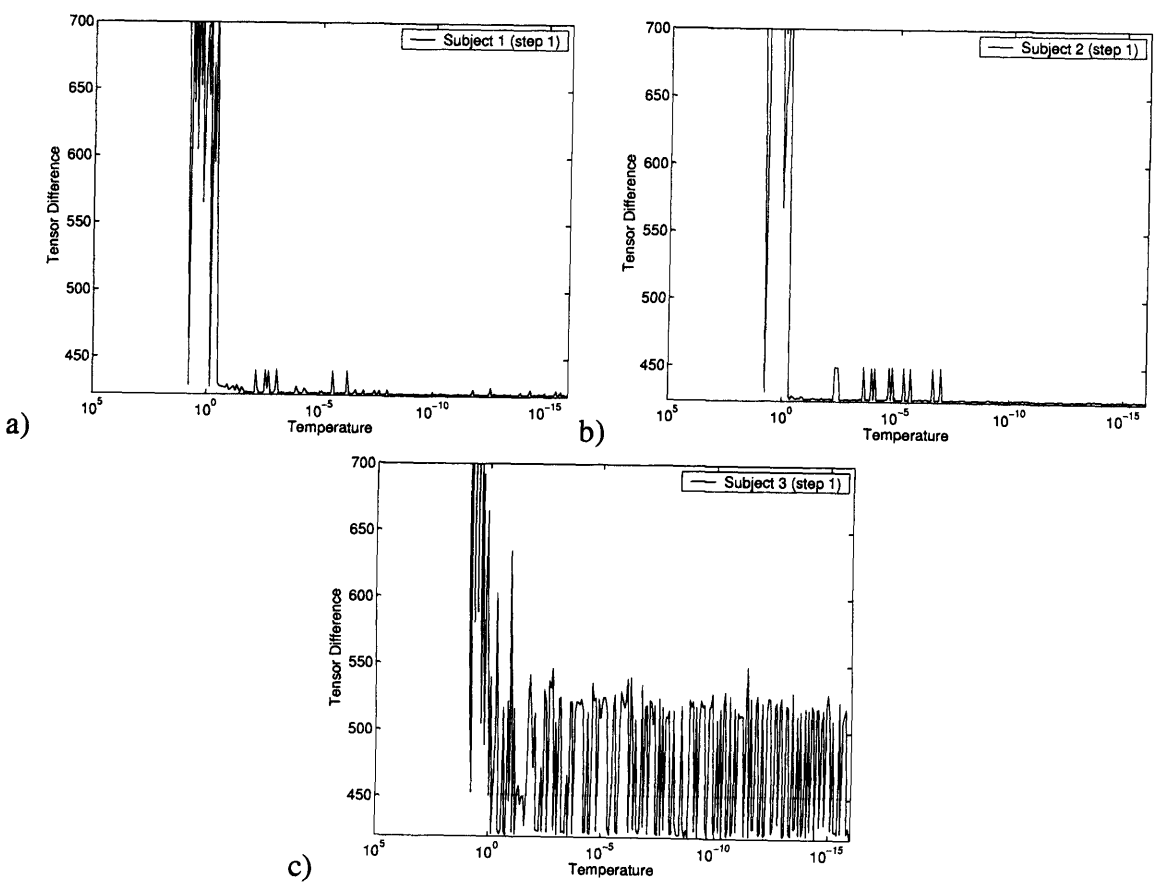

Figure A.3: Plot of the tensor difference similarity measure against temperature, where the temperature has been initialised to 10^5 and the freezing point is 10^{-15}. Figures (a)-(c) consider a step size of 1 and are the results for 3 data sets.

Appendix B

Gradient Annealing Results

Table B.1: The summed relative-anisotropy difference, δ_1, modulus difference, δ_2, tensor difference, δ_3, normalised tensor-difference, δ_4 and principal-direction difference, δ_7, for DT-MR brain registration using Powell's method and gradient annealing, for each of the three reorientation strategies.

Subject	None	FS	PPD	None	FS	PPD
		Powell (δ_1)			**Annealing** (δ_1)	
1	0.063	0.063	0.063	0.061	0.061	0.061
2	0.065	0.065	0.065	0.063	0.062	0.062
3	0.066	0.066	0.065	0.062	0.062	0.062
4	0.079	0.079	0.079	0.066	0.065	0.065
5	0.069	0.066	0.065	0.064	0.064	0.064
		Powell (δ_2)			**Annealing** (δ_2)	
1	233	233	231	230	230	231
2	194	194	194	186	193	192
3	197	202	202	194	198	198
4	218	218	218	201	198	202
5	189	189	189	182	182	181
		Powell (δ_3)			**Annealing** (δ_3)	
1	555	500	566	555	448	447
2	436	438	437	420	421	420
3	455	449	448	428	429	433
4	457	443	440	428	421	424
5	415	408	458	402	399	399
		Powell (δ_4)			**Annealing** (δ_4)	
1	0.155	0.165	0.165	0.153	0.153	0.154
2	0.159	0.158	0.156	0.149	0.148	0.149
3	0.158	0.158	0.158	0.156	0.156	0.157
4	0.199	0.199	0.199	0.199	0.195	0.195
5	0.163	0.153	0.145	0.162	0.143	0.143
		Powell (δ_6)			**Annealing** (δ_6)	
1	-0.584	-0.584	-0.584	-0.585	-0.585	-0.585
2	-0.584	-0.584	-0.584	-0.584	-0.584	-0.584
3	-0.584	-0.584	-0.584	-0.584	-0.584	-0.584
4	-0.584	-0.584	-0.584	-0.584	-0.584	-0.584
5	-0.584	-0.584	-0.584	-0.584	-0.584	-0.584
		Powell (δ_7)			**Annealing** (δ_7)	
1	0.571	0.600	0.572	0.568	0.568	0.569
2	0.616	0.612	0.613	0.591	0.590	0.588
3	0.665	0.664	0.665	0.641	0.611	0.619
4	0.680	0.682	0.644	0.610	0.608	0.617
5	0.657	0.648	0.618	0.602	0.595	0.590

Table B.2: The orientation statistics for the DT-MR brain registrations with no reorientation (NR), finite strain (FS) and preservation of principal direction (PPD) reorientation.

	CC (E_1)			IC (E_1)			Pons (E_3)		
Subject	**None**	**FS**	**PPD**	**None**	**FS**	**PPD**	**None**	**FS**	**PPD**
		(δ_1)			(δ_1)			(δ_1)	
1	0.15	0.16	0.13	0.13	0.10	0.09	0.57	0.60	0.62
2	0.17	0.20	0.23	0.12	0.13	0.13	0.26	0.30	0.28
3	0.29	0.26	0.24	0.14	0.14	0.15	0.38	0.39	0.40
4	0.21	0.22	0.24	0.40	0.42	0.25	1.24	0.53	0.74
5	0.21	0.15	0.19	0.21	0.08	0.08	0.40	0.35	0.41
		(δ_2)			(δ_2)			(δ_2)	
1	0.28	0.30	0.19	0.14	0.13	0.16	0.97	0.85	0.92
2	0.14	0.17	0.14	0.13	0.15	0.22	0.27	0.63	0.50
3	0.18	0.18	0.31	0.09	0.14	0.15	0.37	0.44	0.36
4	1.10	0.98	1.22	0.49	0.58	1.22	1.20	0.89	1.30
5	0.15	0.16	0.20	0.15	0.13	0.12	0.54	0.32	0.44
		(δ_3)			(δ_3)			(δ_3)	
1	0.58	0.14	0.29	0.51	0.18	0.18	0.97	0.61	0.68
2	0.17	0.20	0.19	0.12	0.12	0.19	0.25	0.33	0.63
3	0.24	0.19	0.19	0.11	0.15	0.16	0.39	0.43	0.40
4	0.48	0.30	0.29	0.39	0.39	0.44	1.23	0.62	0.96
5	0.18	0.15	0.18	0.11	0.08	0.07	0.38	0.45	0.36
		(δ_4)			(δ_4)			(δ_4)	
1	0.19	0.14	0.17	0.15	0.15	0.16	0.74	0.65	0.75
2	0.16	0.21	0.28	0.12	0.13	0.13	0.28	0.25	0.23
3	0.19	0.20	0.23	0.15	0.16	0.16	0.47	0.48	0.48
4	0.57	0.56	0.56	1.07	1.09	1.09	1.23	1.25	1.20
5	0.18	0.18	0.28	0.23	0.08	0.09	0.34	0.57	0.35
		(δ_6)			(δ_6)			(δ_6)	
1	0.19	0.17	0.15	0.08	0.20	0.09	0.59	0.50	0.51
2	0.19	0.20	0.24	0.14	0.14	0.12	0.38	0.49	0.33
3	0.13	0.13	0.15	0.13	0.11	0.13	0.35	0.37	0.29
4	0.23	0.25	0.20	0.23	0.28	0.30	0.50	0.55	0.25
5	0.12	0.15	0.19	0.19	0.13	0.14	0.68	0.62	0.48
		(δ_7)			(δ_7)			(δ_7)	
1	0.10	0.08	0.09	0.18	0.18	0.18	0.94	0.95	0.99
2	0.14	0.22	0.28	0.18	0.12	0.11	0.28	0.32	0.32
3	0.22	0.14	0.50	0.13	0.14	0.14	0.75	0.43	0.64
4	0.20	0.16	0.21	0.28	0.36	0.46	0.65	1.11	0.77
5	0.11	0.17	0.21	0.22	0.12	0.09	0.55	0.77	0.60

Bibliography

[1] I. Agartz, J.L. Anderson, and S. Skare. Abnormal brain white matter in schizophrenia: a DT-MRI imaging study. *Neuroreport*, 12(10):2251–2254, 2001.

[2] D.C. Alexander. Note on indices of shape and similarity for diffusion tensors. Technical Report RN/01/14, UCL (University College London), Department of Computer ScienceResearch Note, 2001.

[3] D.C. Alexander. *Visualization and Image Processing of Tensor Fields.*, chapter An introduction to diffusion MRI. Springer, 2005.

[4] D.C. Alexander and J.C. Gee. Elastic matching of diffusion tensor images. *Computer Vision and Image Understanding: Special Issue on Analysis of Volumetric Images*, 7(2):233–250, 2000.

[5] D.C. Alexander, C. Pierpaoli, P.J. Basser, and J.C. Gee. Spatial transformations of diffusion tensor images. *IEEE Transactions on Medical Imaging*, 20(11):1131–1139, 2001.

[6] Y. Amit. A non-linear variational problem for image matching. *SIAM Journal on Scientific Computing*, 15(1):207–224, 1994.

[7] J.L.R. Anderson, S. Skare, and J. Ashburner. How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. *NeuroImage*, 20:870–888, 2003.

[8] N. Arad, N. Dyan, D. Reisfield, and Y. Yeshurun. Image warping by radial basis functions: Application to facial expressions. *Computer Vision, Graphics and*

Image Processing: Graphical Models and Image Processing, 56(2):161–172, 1994.

[9] J. Ashburner and K.J. Friston. Spatial normalization. In A.W. Toga, editor, *Brain Warping*, chapter 2, pages 254–266. San Diego: Academic Press, 1999.

[10] M.A. Audette, F.P. Ferrie, and T.M. Peters. An algorithmic overview of surface registration techniques for medical imaging. *Med Image Anal.*, 4:201–17, 2000.

[11] R. Bajcsy and S. Kovacic. Multiresolution elastic matching. *Comput. Vision Graph. Image Processing*, 46:1–21, 1989.

[12] D.C. Barber, W.B. Tindale, E. Hunt, A. Mayes, and H.J. Sagar. Automatic registration of SPECT images as an alternative to immobilization in neuroactivation studies. *Phys. Med. Biol.*, 40:449–463, 1995.

[13] P.J. Basser. New histological and physiological stains derived from diffusion-tensor MR images. *Annals of the New York Academy of Sciences*, 820:123–138, 1997.

[14] P.J. Basser and D.K. Jones. Diffusion-tensor MRI: Theory, experimental design and data analysis - a technical review. *NMR Biomed*, 15:456–67, 2002.

[15] P.J. Basser, J. Matiello, and D. Le Bihan. MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66:259–267, 1994.

[16] P.J. Basser, S. Pajevic, and C Pierpaoli. In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*, 44:625–632, 2000.

[17] P.J. Basser and C. Pierpaoli. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *J. Magnetic Resonance*, Series B 111:209–219, 1996.

[18] P.G. Batchelor, F. Calamante, M. Moaker, D. Atkinson, A. Connelly, and D.G. Hill. A rigorous method for calculus on diffusion tensors: Taking into account

positivity of eigenvalues. In *Proceedings 12th Scientific Meeting*, page 1205, Kyoto, 2004. ISMRM.

[19] T.E.J. Behrens, H. Johansen-Berg, M.W. Woolrich, S.M Smith, C.A.M. Wheeler-Kingshott, P.A. Boulby, G.J. Barker, E.L. Sillery, K. Sheehan, O. Ciccarelli, A.J. Thompson, J.M. Brady, and P.M. Matthews. Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging. *Nature*, 6(7):750–757, 2003.

[20] P.J. Besl and N.D. McKay. A method for registration of 3D shapes. *IEEE Transactions in Medical Imaging*, 14:239–256, 1992.

[21] F. Bookstein. *Morphometric Tools for Landmark Data: Geometry and Biology.* Cambridge University Press, 1997.

[22] F.L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *Transactions on Pattern Analysis and Machine Intelligence*, 11(6):567–585, 1989.

[23] F.L. Bookstein. Thin-plate splines and the atlas problem for biomedical images. In A. C. F. Colchester and D. J. Hawkes, editors, *Information Processing in Medical Imaging, Proc.12th International Conference*, pages 326–342, Wye, UK, 7-12 July 1991.

[24] F.L. Bookstein and W.D.K. Green. A feature space for edgels in images with landmarks. *Journal of Mathematical Imaging and Vision*, 2:231–261, 1993.

[25] C. Broit. *Optimal Registration of Deformed Images.* Ph.D. Thesis, Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA, 1981.

[26] L.G. Brown. A survey of image registration techniques. *ACM Computing Surveys*, 24(4):325–376, 1992.

[27] H.J. Buiten and B.V. Putten. Quality assessment of remote sensing image registration: Analysis and testing of control point residuals. *ISPRS J. Photogrammetry and Remote Sensing.*, 52(2):57–73, 1997.

[28] T. Butz and J.P. Thiran. Affine registration with feature space mutual information. In W. Niessenand and M. Viergever, editors, *Proceedings of medical image computing and computer-assisted intervention.*, pages 549–556. MICCAI, 2001.

[29] G.E. Christensen. *Deformable Shape Models for Anatomy.* Doctoral dissertation in electrical engineering, Washington University, St. Louis, MO, August 1994.

[30] G.E. Christensen, R.D. Rabbitt, and M.I. Miller. Deformable templates using large deformation kinematics. *IEEE Transactions on Image Processing*, 5(9):399–412, September 1996.

[31] O. Ciccarelli, A.T. Toosy, G.J. Parker, C.A. Wheeler-Kingshott, G.J. Barker, D.H. Miller, and A.J. Thompson. Diffusion tractography based group mapping of major white-matter pathways in the human brain. *NeuroImage*, 19:1545–1555, 2003.

[32] D.L. Collins, P. Neelin, T.M. Peters, and A.C. Evans. Automatic 3D intersubject registration of MR volumetric data in standardized Talairach space. *J. Comput. Assist. Tomography*, 18:192–205, 1994.

[33] T.E. Conturo, N.F. Lori, T.S. Cull, E. Akbudak, A.Z. Snyder, J.S. Shimony, R.C. McKinstry, H. Burton, and M.E. Raichle. Tracking neuronal fiber pathways in the living human brain. In *Proceedings National Acad. Sci. USA*, volume 96, pages 10422–10427, August 1999.

[34] P. Cook and D.C. Alexander. Inter-subject comparison of brain connectivity using diffusion-tensor magnetic resonance imaging. *Proceedings of Medical Image Understanding and Analysis (MIUA)*, pages 117–120, 2003.

[35] J. Crank. *The Mathematics of Diffusion.* Oxford University Press, Oxford, UK, 1975.

[36] W.R. Crum, T. Hartkens, and D.L.G. Hill. Non-rigid image registration: theory and practice. *British Journal of Radiology*, 77:S140–S153, 2004.

[37] W.R. Crum, R.I. Scahill, and N.C. Fox. Automated hippocampal segmentation by regional fluid registration of serial MRI: Validation and application in alzheimer's disease. *NeuroImage*, 13:847–845, 2001.

[38] E. Cuchet, J. Knoplioch, D. Dormont, and C. Marsault. Registration in neurosurgery and neuroradiotherapy applications. *Medical Robotics and Computer Assisted Surgery*, pages 31–38, 1995.

[39] O. Cuisnaire, J. Thiran, B. Macq, C. Michel, A. deVolder, and F. Marqués. Automatic registration of 3D MR images with a computerised brain atlas. In *SPIE Medical Imaging*, pages 438–448, California, 1996.

[40] K.M. Curran and D.C. Alexander. Diffusion tensor orientation matching for image registration. In *Proc. SPIE Medical Imaging: Image Processing.*, volume 5032, pages 149–156, 2003.

[41] K.M. Curran and D.C. Alexander. Orientation coherence optimisation in tensor image registration. In D. Rueckert, J. Hajnal, and G. Z. Yang, editors, *Proc. Medical Image Understanding and Analysis (MIUA)*, pages 259–272, 2004.

[42] K.M. Curran and D.C. Alexander. Comparison of similarity measures for driving tensor registration. In *Proceedings 13th Scientific Meeting*. ISMRM, 2005. accepted.

[43] C. Davatzikos. Spatial transformation and registration of brain images using elastically deformable models. *Comp. Vision and Image Understanding*, 66(2), May 1997.

[44] C. A. Davatzikos, J.L. Prince, and R.N. Bryan. Image registration based on boundary mapping. *IEEE Transactions on Medical Imaging*, 15(1):112–115, 1995.

[45] M.H. Davis, A. Khotanzad, D.P. Flamig, and S.E. Harms. A physics-based coordinate transformation for 3-D image matching. *IEEE Transactions on Medical Imaging*, 16(3):317–328, 1997.

[46] J. Declerck, J. Feldmar, M.L. Goris, and F. Betting. Automatic registration and alignment on a template of cardiac stress and rest reoriented SPECT images. *IEEE Transactions on Medical Imaging*, 6:727–737, 1997.

[47] Q. Dong, R.C. Welch, T.L. Chenevert, R.C. Carlos, P. Maly-Sundgren, D.M. Gomez-Hasan, and S.K. Mukherji. Clinical applications of diffusion tensor imaging. *Journal of Magnetic Resonance Imaging*, 6:6–18, 2002.

[48] K.A. Dowsland. *Simulated Annealing. In Modern Heuristic Techniques for Combinatorial Problems.* McGraw-Hill, New York, 1995.

[49] J.T. Duda, M. Rivera, D.C. Alexander, and J.C. Gee. A method for non-rigid registration of diffusion tensor magnetic resonance images. In *Proc. SPIE Medical Imaging 2003: Image Processing.*, volume 5032, pages 1186–1196, 2003.

[50] P.J. Edwards, D.L.G. Hill, J.A. Little, and D.J. Hawkes. A three-component deformation model for image-guided surgery. *Medical Image Analysis*, 2(4):355–367, 1998.

[51] M. Ferrant, S.K. Warfield, C.R. Guttman, R.V. Mulkern, F.A. Jolesz, and R. Kikinis. 3D image matching using a finite element based elastic deformation model. In *Proc. of MICCAI*, pages 202–209, 1998.

[52] J. Fitzparick, J. West, and C. Maurer Jr. Predicting error in rigid-body, point-based registration. *IEEE Transactions on Medical Imaging*, 17:694–702, October 1998.

[53] J. M. Fitzparick and J.B. West. The distribution of target registration error in rigid-body point-based registration. *IEEE Transactions on Medical Imaging*, 20:917–27, 2001.

[54] J.M. Fitzpatrick, D.L. Hill, and C.R. Maurer. Image registration. In M Sonka and JM Fitzpatrick, editors, *Handbook of Medical Imaging, Vol. 2: Medical Image Processing and Analysis.*, volume 2, chapter 8. SPIE Press, Bellingham, WA, 2000.

[55] M. Fornefett, K. Rohr, and H.S. Stiehl. Elastic registration of medical images using radial basis functions with compact support. *Proc. Computer Vision and Pattern Recognition*, pages 402–407, 1999.

[56] S. Forrest. Genetic algorithms: Principles of natural selection applied to computation. *Science*, 261:872–878, 1993.

[57] P.A. Freeborough and N.C. Fox. Modeling brain deformations in alzheimer's disease by fluid registration of serial MR images. *J. Comp Assisted Tomography*, 22:838–843, 1998.

[58] K.J. Friston, J. Ashburner, C.D. Frith, J.B. Poline, J.D. Heather, and R.S. Frackowiak. Spatial registration and normalisation of images. *Human Brain Mapping*, 2:165–189, 1995.

[59] K.S. Fu, R.C. Gonzalez, and C.S.G. Lee. *Robotics: Control, Sensing, Vision and Intelligence.* McGraw-Hill, New York, 1987.

[60] J.C. Gee and D.C. Alexander. *Visualization and Image Processing of Tensor Fields.*, chapter DT-MRI Registration. Springer, 2005.

[61] J.C. Gee and R.K. Bajcsy. Elastic matching: Continuum mechanical and probabilistic analysis. In A.W. Toga, editor, *Brain Warping*, chapter 11. Academic Press, 1999.

[62] J.C. Gee, C. Barillot, L. Le Briquer, D.R. Haynor, and R. Bajcsy. Matching structural images of the human brain using statistical and geometrical features. In *Proc. 3rd International Conference Visualization in Biomedical Computing (VBC)*, pages 191–204, 1994.

[63] J.C. Gee, M. Reivicj, and R. Bajcsy. Elastically deforming a three-dimensional atlas to match anatomical brain images. *Journal of Computed Assisted Tomography*, 17:225–236, 1993.

[64] S. Geman and D. Geman. Stochastic relaxation, gibbs distribution and the bayesian restoration in images. *IEEE Trans. Patt. Anal. Mac. Int.*, 6(6):721–741, 1984.

[65] G.H. Golub and C.F. Van Loan. *Matrix Computations*. The John Hopkins University Press., Baltimore, Maryland, 1983.

[66] A. Goshtasby. Image registration by local approximation methods. *Image and Vision Computing*, 6(4):255–261, 1988.

[67] A Guimond, CR Guttmann, SK Warfield, and CF Westin. Deformable registration of DT-MRI data based on transformation invariant tensor characteristics. In *Proceedings of the IEEE International Symposium on Biomedical Imaging, Washington, DC.*, 2002.

[68] A. Guimond, A. Roche, N. Ayache, and J. Meunier. Three-dimensional multimodal brain warping using the demons algorithm and adaptive intensity corrections. *IEEE Transactions on Medical Imaging*, 20(1):58–69, 2001.

[69] J.V. Hajnal, D.L.G. Hill, and D.J. Hawkes. *Medical image registration.* CRC Press, Boca Raton, 2001.

[70] D.L.G. Hill, P.G. Batchelor, M. Holden, and D.J. Hawkes. Medical image registration. *Phys. Med. Biol.*, 46:R1–R45, 2001.

[71] D.L.G. Hill, C. Studholme, and D.J. Hawkes. Voxel similarity measures for automated image registration. In *Proc. Visualization in Biomedical Computing*, volume 2359, pages 205–216. SPIE, 1994.

[72] M. Holden, D.L.G. Hill, E.R.E. Denton, J.M. Jarosz, T.C.S. Cox, T. Rohlfing, J. Goodey, and D.J. Hawkes. Voxel similarity measures for 3-D serial MR brain image registration. *IEEE Transactions in Medical Imaging*, 19(2):94–102, February 2000.

[73] L. Ingber. Simulated annealing: Practice versus theory. *Journal of Mathematical and Computational Modelling*, 18:29–57, 1993.

[74] K.M. Jansons and D.C. Alexander. Persistant angular structure: new insights from diffusion magnetic resonance imaging data. *Inverse Problems*, 19:1031–1046, 2003.

[75] M. Jenkinson and S.M. Smith. A global optimisation method for robust affine registration of brain images. *Medical Image Analysis*, 5(2):143–156, 2001.

[76] D.K. Jones, L.D. Griffin, D.C. Alexander, M. Catani, M.A. Horsfield, R. Howard, and S.C.R. Williams. Spatial normalisation and averaging of diffusion tensor MRI data sets. *NeuroImaging*, 17(2):592–617, 2002.

[77] D.K. Jones, M.A. Horsefield, and A. Simmons. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. *Magnetic Resonance in Medicine*, 42:515–525, 1999.

[78] D.K. Jones, A. Simmons, S.C.R. Williams, and M.A. Horsefield. Non-invasive assessment of axonal fiber connextivity in the human brain via diffusion tensor MRI. *Magnetic Resonance in Medicine*, 42:37–41, 1999.

[79] K.A. De Jong. Genetic algorithms are not function optimizers. In *Proceedings In: Foundations of Genetic Algorithms*, July 1992.

[80] S. Kiebel, J. Ashburner, J. Poline, and K. Friston. MRI and PET coregistration - a cross validation of statistical parametric mapping and automated image registration. *NeuroImage*, 5, 1997.

[81] S. Kirkpatrick, C. D. Gellatt, and M. P. Vecchi. Simulated annealing. *Science*, 220(671), 1983.

[82] D. LeBihan, J.F. Mangin, C. Poupon, C. Clark, S. Pappata, N. Molko, and H. Chabriat. Diffusion tensor imaging: Concepts and applications. *Journal of Magnetic Resonance Imaging*, 13:534–546, 2001.

[83] L. Lemieux, U.C. Wieshmann, N.F. Moran, D.R. Fish, and S.D. Shoovon. The detection and significance of subtle changes in mixed-signal brain lesions by serial MRI scan matching and spatial normalisation. *Med. Image Anal.*, 2(3):227–242, 1998.

[84] H. Lester. *Non-linear registration of medical images.* Ph.D. Thesis, University College London, Department of Computer Science, London, UK, 1998.

[85] H. Lester and S.R. Arridge. A survey of hierarchical non-linear medical image registration. *Pattern Recognition*, 32:129–149, 1999.

[86] J.A. Little, D.L.G. Hill, and D.J. Hawkes. Deformations incorporating rigid structures. *Computer Vision and Image Understanding*, 66(2):223–232, 1997.

[87] J.B.A. Maintz and M.A. Viergever. A survey of medical image registration. *Medical Image Analysis*, 2:129–149, 1998.

[88] L.E. Malvern. *Introduction to the mechanics of a continuous medium.* Prentice Hall, Englewood Cliffs, N.J., 1969.

[89] J.F. Mangin, C. Poupon, C. Clark, D. LeBihan, and I. Bloch. Distortion correction and robust tensor estimation for mr diffusion imaging. *Medical Image Analysis*, 6:191–8, 2002.

[90] B. S. Manjunath and W. Y. Ma. Texture features for browsing and retrieval of image data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(8):837–842, 1996.

[91] K.E. Mardia and P.E. Jupp. *Directional statistics*. Wiley, Chichester, 2000.

[92] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equations of state calculations by fast computing machine. *J. Chem. Phys.*, 21:1087–1091, 1953.

[93] O. Monga and S. Benayoun. Using partial derivatives of 3D images to extract typical surface features. *Comput. Vision Image Understanding*, 61:171–189, 1995.

[94] S. Mori, BJ.. Crain, V.P. Chacko, and P.C.M. van Zijl. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Ann. Neurol.*, 45:265–269, 1999.

[95] S. Mori, M. Kraut, P.C. van Zijl, M. Solaiyappan, H. Brem, and M. Pomper. Diffusion tensor imaging and fiber tracking of brain tumours. In *Proc. International Society Magnetic Resonance in Medicine.*, 2001.

[96] J.F. Nielson, N.R. Ghugre, and A. Panigrahy. Affine and polynomial mutual information coregistration for artifact elimination in diffusion tensor imaging of newborns. *Magnetic Resonance Medicine.*, 22:1319–1323, 2004.

[97] J.L. Ostuni and A.K. Santha. Analysis of interpolation effects in the reslicing of functional MR images. *Journal of computer assisted tomography*, 21(5):803–810, 1997.

[98] M. O'Sullivan, D.K. Jones, P.E. Summers, R.G. Morris, S.C. Williams, and A.S. Markus. Evidence for cortical disconnection as a mechanism of age-related cognitive decline. *Neurology*, 57:632–638, 2001.

[99] A. Papoulis. *Probability, Random variables and Stochastic Processes.* McGraw-Hill, New York, 1984.

[100] H. J. Park, M. Kubicki, M. E. Shenton, A. Guimond, R. W. McCarley, S. E. Maier, R. Kikinis, F. A. Jolesz, and C.-F. Westin. Spatial normalization of diffusion tensor MRI using multiple channels. *NeuroImage*, 20(4):1995–2009, 2003.

[101] G. J. M. Parker, C. A. M. Wheeler-Kingshott, and G. J. Barker. Estimating distributed anatomical connectivity using fast marching methods and diffusion tensor imaging. *IEEE Transactions in Medical Imaging*, 21:505–512, 2002.

[102] X. Pennec and J.P. Thirion. A framework for uncertainty and validation of 3D registration methods based on points and frames. *Int. J. Comput. Vision*, 25:203–229, 1997.

[103] C. Pierpaoli and P.J. Basser. Toward a quantitative assessment of diffusion anisotropy. *Magnetic Resonance in Medicine*, 36:893–906, 1996.

[104] J.P.W. Pluim, J.B.A. Maintz, and M.A. Viergever. Interpolation artefacts in mutual information based image registration. *Computer Vision and Image Understanding*, 77(2):211–232, 2000.

[105] J.P.W. Pluim, J.B.A. Maintz, and M.A. Viergever. Mutual-information-based registration of medical images: a survey. *IEEE Transactions on Medical Imaging*, 22(8):986–1004, 2003.

[106] C. Poupon, C.A. Clark, and V. Froulin. Regularisation of diffusion-based direction maps for the tracking of brain white matter fasciles. *NeuroImage*, 12:184–195, 2000.

[107] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery. *Numerical recipes in C. The art of scientific computing.* Cambridge University Press, 2nd edition, 1992.

[108] V.J. Rayward-Smith, I.H. Osman, C.R. Reeves, and G.D. Smith. *Modern Heuristic Search Methods*. John Wiley and Sons, 1996.

[109] G.K. Rhode, A. Aldroubi, and B.M. Dawant. The adaptive bases algorithm for intensity-based non-rigid image registration. *IEEE Transactions on Medical Imaging*, 22:1470–1479, 2003.

[110] G.K. Rhode, A.S. Barnett, P.J. Basser, S. Marenco, and C. Pierpaoli. Comprehensive approach for correction of motion and distortion in diffusion-weighted MRI. *Magnetic Resonance in Medicine*, 51:103–14, 2004.

[111] G.K. Rhode, S. Pajevic, and C. Pierpaoli. Multi-channel registration of diffusion tensor images using directional information. In *Proc. IEEE International Symposium on Biomedical Imaging (Arlington)*, pages 712–715, 2004.

[112] G.K. Rhode, C. Pierpaoli, S. Pajevic, and P.J. Basser. A comprehensive approach for multi-channel image registration. In *Proceedings of 2nd International Workshop on Biomedical Image Registration.*, pages 214–223, 2003.

[113] A. Roche, G. Malaindain, X. Pennec, and N. Ayache. The correlation ratio as a new similarity measure for multimodal image registration. In *Proc. MICCAI 1998, Lecture Notes in Computer Science*, volume 1496, pages 1115–24, 1998.

[114] D. Rueckert, L.I. Sonoda, C. Hayes, D.L.G. Hill, M.O. Leach, and D.J. Hawkes. Non-rigid registration using free-form deformations: Application to breast mr images. *IEEE Transactions on Medical Imaging*, 18(8):712–721, 1999.

[115] J. Ruiz-Alzola, C.F. Westin, S.K. Warfield, C. Alberola, S. Maier, and R. Kikinis. Nonrigid registration of 3D tensor medical data. *Medical Image Analysis*, 2002.

[116] L.R. Schad, R. Boesecke, W. Schlegel, G.H. Hartmann, V. Sturm, L.G. Strauss, and W.J. Lorenz. Three dimensional image correlation of CT, MR and PET studies in radiotherapy treatment planning of brain tumors. *Journal of Computed Assisted Tomography*, 11:948–954, 1987.

[117] J.A. Schnabel, C. Tanner, A.D. Castellano-Smith, A. Degenhard, and M.O. Leach et al. Validation of volume-preserving non-rigid registration using finite element methods: application to breast MR images. *IEEE Transactions on Medical Imaging*, 22:238–47, 2003.

[118] C. E. Shannon. A mathematical theory of communication. *Bell Syst. Te ch. J.*, 27:379–423/623–656, 1948.

[119] D.G. Shen, Y. Zhan, and C. Davatzikos. Segmentation of prostrate boundaries from ultrasound images using statistical shape model. *IEEE Transactions in Medical Imaging*, 4:539–551, 2003.

[120] N. Sicotte, R. Woods, and J. Mazziotta. Automated image registration using a 105 parameter non-linear model. *NeuroImage*, 3(3), 1996.

[121] S.M. Smith. Fast robust automated brain extraction. *Human Brain Mapping*, 17(3):143–155, 2002.

[122] S.C. Strother, J.R. Anderson, X.L. Xu, J.S. Liow, D.C. Bonar, and D.A. Rottenberg. Quantitative comparisons of image registration techniques based on high-resolution MRI of the brain. *Journal of Computer Assisted Tomography*, 18:954–62., 1994.

[123] C. Studholme, D. L. G. Hill, and D. J. Hawkes. An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognition*, 32:71–86, 1999.

[124] M. R. Symms, G.J. Barker, F. Franconi, and C.A. Clark. Correction of eddy-current distortions in diffusion weighted echo-pla nar images with a two-dimensional registration technique. In *Proc. 5th Annual Meeting of the ISMRM*, page 1723, Vancouver, 1997. Berkeley, USA: ISMRM.

[125] R. Szeliski and S. Lavallee. Matching 3D anatomical surfaces with non-rigid deformations using octree-splines. *IEEE Workshop on Biomedical Image Image Analysis*, 11:144–153, 1994.

[126] J. Talairach and P. Tournoux. *Co-planar stereotactic atlas of the human brain.* Stuttgart: George Thieme, Verlag, 1988.

[127] N.A. Thacker, A. Jackson, and D. Moriarty. Improved quality of re-sliced MR images using re-normalized sinc interpolation. *Journal of Magnetic Resonance Imaging*, 10:582–588, 1999.

[128] P. Thompson and A.W. Toga. A surface-based technique for warping three-dimensional images of the brain. *IEEE Transactions on Medical Imaging*, 15(4):402–417, 1996.

[129] J.D. Tournier, F. Calamante, D.Gadian, and A. Connelly. Direct estimation of the fiber orientation density function from diffusion-weighted mri data using spherical deconvolution. *NeuroImage*, 23(3):1176–1185, 2004.

[130] D.S. Tuch. Q-ball imaging. *Magnetic Resonance in Medicine*, 52(6):1358–1372, 2004.

[131] J. Vaarkamp, D. C. Barber, J. Conway, and M. H. Robinson. Automatic image registration of diagnostic and radiotherapy treatment planning ct head images. *Int. J. Radiation Oncology Biol. Phys.*, 47(1):225–230, 2000.

[132] R. Verma and C. Davatzikos. Matching of diffusion tensor images using gabor features. *Proceedings IEEE International Symposium on Biomedical Imaging (Arlington)*, pages 396–399, 2004.

[133] P. Vermersch, J. Roche, M. Hamon, C. Daems-Monpeurt, P.J. Pruvo, P. Dewailly, and H. Petit. White matter magnetic resonance imaging hyperintensity in alzheimer's disease: Correlations with corpus callosum atrophy. *Neurology*, 243(2):231–234, 1996.

[134] U.C. Weishmann, G.J. Parker, M.R. Symms, G.J. Barker, and S.D. Shorvon. In vivo tractography using diffusion tensor imaging and spatial normalisation: Rotation of eigenvectors into talairach space. In *Proc. International Society*

for Magnetic Resonance in Medicine , 7th Scientific Meeting and Exhibition, Philadelphia, PA, 24-28 May 1999.

[135] J. West, J.M. Fitzpatrick, M.Y. Wang, B.M. Dawant, C.R. Maurer, R.M. Kessler, R.J. Maciunas, C. Barillot, D. Lemoine, A. Collignon, F. Maes, P. Suetens, D. Vandermeulen, P. van den Elsen, S. Napel, T.S. Sumanaweera, B. Harkness, P.F. Hemler, D.L.G. Hill, D.J. Hawkes, C. Studholme, J.B.A. Maintz, M.A. Viergever, G. Malandain, X. Pennec, M.E. Noz, G.Q. Maguire, M. Pollack, C.A. Pelizzari, R.A. Robb, D. Hanson, and R.P. Woods. Comparison and evaluation of retrospective intermodality brain image registration techniques. *Journal of Computer Assisted Tomography*, 21:554–566, 1997.

[136] C.F. Westin, S.E. Maier, H. Mamata, A. Nabavi, F.A. Jolesz, and R. Kikinis. Processing and visualization for diffusion tensor MRI. *Medical Image Analysis*, 6:93–108, 2002.

[137] C.A.M. Wheeler-Kingshott, P.A. Boulby, M. Symms, and G.J. Barker. Optimised cardiac gating for high-resolution whole brain dti on a standard scanner. In *Proc. International Society for Magnetic Resonance in Medicine.*, 2002.

[138] PW Woodruff, GD Pearlson, MJ Geer, PE Barta, and HD Chilcoat. A computerised magnetic resonance imaging study of corpus callosum morphology in schizophrenia. *Psychological Medicine*, 23:45–56, 1993.

[139] R.P. Woods, S.T. Grafton, J.D.G. Watson, N.L. Sicotte, and J.C. Mazziotta. Automated image registration: I. general methods and intrasubject, intramodality validation. *Journal of Computer Assisted Tomography*, 22:141–154, 1998.

[140] R.P. Woods, S.T. Grafton, J.D.G. Watson, N.L. Sicotte, and J.C. Mazziotta. Automated image registration: Ii. intersubject validation of linear and nonlinear models. *Journal of Computer Assisted Tomography*, 22:155–165., 1998.

[141] RP Woods, JC Mazziotta, and SR Cherry. MRI-PET registration with automated algorithm. *Journal of Computer Assisted Tomography*, 17:536–546, 1993.

[142] Y.T. Wu. Image registration using wavelet-based motion model. *Intl. J. Comp. Vision.*, 38(2):129–152, July 2000.

[143] D. Xu, S. Mori, D. Shen, and C. Davatzikos. Statistically-based reorientation of diffusion tensor field. In *IEEE International Symposium on Biomedical Imaging: Macro to Nano*, pages 757–760, Washington, D.C., 7-10 July 2002.

[144] D. Xu, S. Mori, D. Shen, P. C. M. van Zijl, and C. Davatzikos. Spatial normalization of diffusion tensor fields. *Magnetic Resonance in Medicine*, 50:175–182, December 2003.

[145] D. Xu, S. Mori, M. Solaiyappan, P. C. M. van Zijl, and C. Davatzikos. A framework for callosal fiber distribution analysis. *NeuroImage*, 14:1361–1369, December 2001.

[146] H. Zhang, P.A. Yushkevich, and J.C. Gee. Registration of diffusion tensor images. In *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 842–847, Washington, D.C., 2004.

[147] A. Žilinskas. A review of statistical models for global optimization. *J. Global Optim.*, 2(2):145–153, 1992.

[148] B. Zitova and J. Flusser. Image registration methods: a survey. *IVC*, 21(11):977–1000, October 2003.