Towards Efficient Neurosurgery: Image Analysis for Interventional MRI

Pankaj Daga

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

of

University College London.

Centre for Medical Image Computing

University College London

2014

I, Pankaj Daga, confirm that the work presented in this thesis is my own.

Where information has been derived from other sources,

I confirm that this has been indicated in the thesis.

In memory of my father.

To my mom - who loved me enough to set me free.

Abstract

Interventional magnetic resonance imaging (iMRI) is being increasingly used for performing image-guided neurosurgical procedures. Intermittent imaging through iMRI can help a neurosurgeon visualise the target and eloquent brain areas during neurosurgery and lead to better patient outcome. MRI plays an important role in planning and performing neurosurgical procedures because it can provide high-resolution anatomical images that can be used to discriminate between healthy and diseased tissue, as well as identify location and extent of functional areas. This is of significant clinical utility as it helps the surgeons maximise target resection and avoid damage to functionally important brain areas.

There is clinical interest in propagating the pre-operative surgical information to the intra-operative image space as this allows the surgeons to utilise the pre-operatively generated surgical plans during surgery. The current state of the art neuronavigation systems achieve this by performing rigid registration of pre-operative and intra-operative images. As the brain undergoes non-linear deformations after craniotomy (*brain shift*), the rigidly registered pre-operative images do not accurately align anymore with the intra-operative images acquired during surgery. This limits the accuracy of these neuronavigation systems and hampers the surgeon's ability to perform more aggressive interventions. In addition, intra-operative images are typically of lower quality with *susceptibility artefacts* inducing severe geometric and intensity distortions around areas of resection in echo planar MRI images, significantly reducing their utility in the intraoperative setting.

This thesis focuses on development of novel methods for an image processing workflow that aims to maximise the utility of iMRI in neurosurgery. I present a fast, non-rigid registration algorithm that can leverage information from both structural and diffusion weighted MRI images to localise target lesions and a critical white matter tract, the optic radiation, during surgical management of temporal lobe epilepsy. A novel method for correcting susceptibility artefacts in echo planar MRI images is also developed, which combines fieldmap and image registration based correction techniques. The work developed in this thesis has been validated and successfully integrated into the surgical workflow at the National Hospital for Neurology and Neurosurgery in London and is being clinically used to inform surgical decisions.

Acknowledgements

Over the past four years I have received support and encouragement from a great number of individuals. Prof. Sébastien Ourselin has been a mentor, colleague and friend. His guidance has made this a thoughtful and rewarding journey. My gratitude goes as well to my secondary supervisor, Prof. John Duncan, who gave me valuable guidance and support over the course of my thesis.

I am also grateful to my clinical collaborators at the Department of Clinical and Experimental Epilepsy and the National Hospital for Neurology and Neurosurgery. I would especially like to thank Dr. Gavin Winston whose collaboration was pivotal in taking this work to the clinic. Sincere thanks go to Dr. Mark White and Dr. Laura Mancini for providing me with patient imaging data and spending countless hours on testing my software.

I would like to thank all my lab colleagues and friends who made our group a fun place to work.

This thesis would not have finished without the support of my family. Whatever little good is in me is because of my mother and I would like to thank her for the sacrifices she has endured so that I could pursue my dreams. I hope this makes you proud mum. Thanks goes out to my brother and sister-in-law for their infinite support and encouragement over the years. This list would be amiss without mentioning Shweta, who constantly believes in the abilities of her brother despite evidence to the contrary. Lastly, this thesis is dedicated to my late father. Thanks for everything Dad. It's been twenty years but not a day goes by when I don't miss you.

Publication list

Peer Reviewed Journal Papers

- **Daga P.**, Pendse T., Modat M., White M., Mancini L., Winston G., McEvoy A. W., Thornton J., Yousry T., Drobnjak I., Duncan J. S., Ourselin S.: Susceptibility Artefact Correction using Dynamic Graph Cuts: Application to Neurosurgery. (2014) Medical Image Analysis.

- **Daga P.**, Winston G., Modat M., White M., Mancini L., Cardoso M. J., Symms M., Stretton J., McEvoy A. W., Thornton J., Micallef C., Yousry T., Hawkes D., Duncan J. S., Ourselin S.: Accurate Localisation of Optic Radiation during Neurosurgery in an Interventional MRI Suite. (2012) IEEE Transactions on Medical Imaging.

- Winston G., **Daga P.**, White M., Micallef C., Miserocchi A., Mancini L., Modat M., Stretton J., Sidhu M., Symms M., Lythgoe D., Thornton J., Yousry T., Ourselin S., Duncan J. S., McEvoy A. W.: Preventing Visual Field Deficits from Neurosurgery. (2014) Neurology.

- Winston G., **Daga P.**, Stretton J., Modat M., Symms M., McEvoy A. W., Ourselin S., Duncan J. S.: Optic Radiation Tractography and Vision in Anterior Temporal Lobe Resection. (2011) Annals of Neurology.

- Muhlert N., Sethi V., Schneider T., **Daga P.**, Cipolotti L., Haroon H., Parker G., Ourselin S., Wheeler-Kingshott C., Miller D., Ron M., Chard D.: Diffusion MRI-based Cortical Complexity Alterations Associated with Executive Function in Multiple Sclerosis. (2012) Journal of Magnetic Resonance Imaging.

Peer Reviewed Conference Papers

- **Daga P.**, Modat M., Winston G., White M., Mancini L., McEvoy A. W., Thornton J., Yousry T., Duncan J., Ourselin S.: Susceptibility artefact Correction by combining B_0 field maps and non-rigid registration using graph cuts, (2012) SPIE. ***Winner: Best student paper award***

- **Daga P.**, Winston G., Modat M., Cardoso M. J., White M., McEvoy A. W., Thornton J., Hawkes D., Duncan J., Ourselin S.: Improved neuronavigation through integration of intra-operative anatomical and diffusion images in an interventional MRI suite, (2011) IPCAI.

- **Daga, P.**, Winston G., Modat M., Cardoso M. J., Stretton J., Symms M., McEvoy A. W., Hawkes D., Duncan J., Ourselin S.: Integrating Structural and Diffusion MR Information for Optic Radiation Localisation in Focal Epilepsy Patients, (2011) IEEE ISBI.

- Kochan M., **Daga, P.**, Burgos N., White M., Cardoso, M. J., Mancini L., Winston G., McEvoy A. W., Thornton J., Yousry T., Duncan J., Stoyanov D., Ourselin S.: Simulated Field Maps: Toward Improved Susceptibility Artefact Correction in Interventional MRI, (2014) IPCAI.

- Modat M., Ridgway G. R., **Daga P.**, Cardoso M. J., Hawkes D., Ashburner J., Ourselin S.: Log-Euclidean Free-Form Deformation, (2011) SPIE.

- Modat M., Ridgway G. R., **Daga P.**, Cardoso M. J., Ashburner J., Ourselin S.: Parametric Non-Rigid Registration using a Stationary Velocity Field, (2012) MMBIA.

- Modat M., **Daga P.**, Cardoso M. J., Cash D., Fox N., Ourselin S.: Inverse Consistent Symmetric Free Form Deformation, (2012) WBIR.

Abstracts and Workshops

- **Daga P.**, Modat M., Micallef C., Mancini L., White M., Cardoso M. J., Kitchen N., McEvoy A. W., Thornton J., Yousry T., Hawkes D., Ourselin S.: Near Real Time Brain Shift Estimation for Interventional MRI Suite, (2010) MICCAI Workshop on High Performance Computing.

- Winston G., **Daga P.**, Stretton J., Modat M., Symms M., McEvoy A. W., Ourselin S., Duncan J.: Propagation of Probabilistic Tractography of the Optic Radiation in Epilepsy Surgery, (2011) ISMRM.

- Siow B., Cleary J., Greene N., **Daga P.**, Modat M., Ordidge R., Ourselin S., Alexander D., Lythgoe M.: Mouse Embryo Phenotyping with Contrast-Enhanced Micro-Diffusion Tensor Imaging, (2011) ISMRM.

- Roura E., Schneider T., **Daga P.**, Modat M., Muhlert N., Freixenet J., Chard D., Ourselin S., Llado X., Wheeler-Kingshott C.: Multi-channel registration of FA and T1-weighted images to standard space: patients with multiple sclerosis, (2012) ISMRM.

- Roura E., Schneider T., Modat M., **Daga P.**, Muhlert N., Chard DT., Ourselin S., Llado X., Wheeler-Kingshott C.: Evaluating a multi-channel registration approach of FA and T1w on MS patients with simulated atrophy, (2013) ISMRM.

- Kinnunen K., Ridgway G., Cash D., Leite AB., Finnegan S., **Daga P.**, Cardoso M. J., Ryan N., Espak M., Rossor M., Ourselin S., Fox N.: Abnormalities of fronto-striato-thalamic tract structure and effective connectivity in familial Alzheimer's disease, (2013) Alzheimer's & Dementia: The Journal of the Alzheimer's Association

Patents

- **Daga P.**, Ourselin S.: Apparatus and method for correcting susceptibility artefacts in a magnetic resonance image. *Under filing*.

Contents

Abstract **5**

Acknowledgements **7**

Publication list **9**

1 Introduction **25**

1.1 Clinical Background . . . 25

1.2 MRI in Neurosurgery . . . 27

1.3 Challenges in iMRI neuronavigation . . . 28

1.3.1 Brain Shift . . . 28

1.3.2 Artefacts in iMRI Images . . . 28

1.3.3 Integration Into Surgical Workflow . . . 29

1.4 Methodological Contributions . . . 29

1.5 Thesis Organization . . . 30

2 Literature Review **31**

2.1 Brain Shift Estimation . . . 31

2.1.1 Image Registration Based Brain Shift Estimation . . . 32

2.1.2 Biomechanical Model Based Brain Shift Estimation . . . 38

2.2 Susceptibility Artefacts in MRI . . . 39

2.2.1 Susceptibility Artefact Correction with Field Maps . . . 40

2.2.2 Susceptibility Artefact Correction with Image Registration . . . 43

2.3 Discussion . . . 45

3 Overview of Computational Techniques **47**

3.1 Discrete Optimisation: Graph Cuts . . . 47

3.1.1 Energy Minimisation . . . 47

3.1.2 Probabilistic Graphical Models and Random Fields . . . 48

3.1.3 Inference on MRFs: *Maximum a Posteriori* Estimation . . . 49

3.1.4 Energy Minimisation via Graph Cuts . . . 50

3.1.5 Submodular Functions . . . 51

3.1.6 Graph Construction for Submodular Functions . . . 52
3.1.7 Multi-label Optimisation with Graph Cuts . . . 54
3.2 Medical Image Registration . . . 55
3.2.1 Parametric Non-rigid Transformation Models: Cubic B-splines . . . 55
3.2.2 Similarity Measure: (Normalised) Mutual Information . . . 57
3.2.3 Optimisation: Conjugate Gradient Descent . . . 59

4 Susceptibility Artefact Correction 61
4.1 Associated Publications . . . 62
4.2 Noise in MRI Phase Images . . . 62
4.3 Phase Modelling . . . 63
4.3.1 Energy Minimization via Graph Cuts . . . 64
4.4 Uncertainty Estimation and Image Registration . . . 66
4.4.1 Uncertainty Estimation in Phase Unwrapping . . . 66
4.5 Image Registration Framework . . . 68
4.6 Validation . . . 70
4.6.1 Validation Using Simulated Data . . . 70
4.6.2 Validation Using Clinical Data . . . 72
4.7 Discussion . . . 73

5 Optic Radiation Localisation during Neurosurgery 77
5.1 Associated Publications . . . 77
5.2 Methods . . . 78
5.2.1 Fractional Anisotropy . . . 79
5.2.2 Bivariate NMI . . . 79
5.3 Validation . . . 82
5.3.1 Validation Using a Numerical Phantom . . . 82
5.3.2 Quantitative Validation on Post-Operative Clinical MRI Data . . . 84
VFD Quantificatiom . . . 84
Optic Radiation Parcellation . . . 84
5.3.3 Quantitative Validation on Interventional MRI Datasets . . . 85
Challenges With Intra-Operative Tractography . . . 88
Assessment of Registration Accuracy with Intra-Operative MRI . . . 89
5.3.4 Computational Performance Validation . . . 90

6 Clinical Integration 93

7 Clinical Findings 101
7.1 Methods . . . 101
7.1.1 Subjects . . . 101

7.1.2 Comparison Cohort . . . 101
7.1.3 Optic Radiation Tractography . . . 102
7.1.4 Surgery and Intraoperative Imaging . . . 102
7.1.5 Primary Outcome: Visual Fields . . . 103
7.1.6 Statistical Analysis . . . 103
7.2 Results . . . 103
7.2.1 Visual Field Deficits . . . 103
7.2.2 Seizure Outcome . . . 104

8 Open Software Effort **105**
8.1 NiftyReg . . . 105

9 Discussion and Conclusion **107**
9.1 Future Research Directions . . . 109

List of Figures

1.1 Human Visual System: (a) shows a schematic of the visual wiring in the brain (courtesy www.thebrain.mcgill.ca). As shown, the Meyers loop passes through the temporal lobe and hence is at risk of injury during the surgical intervention. (b) shows a dissected brain (courtesy Virtual Hospital) and the Meyers loop can be clearly identified. The blue oval highlights the area typically affected by temporal lobe resection, which can result in damage to the Meyers loop. 26

1.2 The interventional MRI surgical suite at the National Hospital for Neurology and Neurosurgery with a 1.5 tesla MR scanner and neuronavigation equipment. The surgical table interfaces with the scanner to enable the patient to be moved in and out of the scanner efficiently during surgery. 27

1.3 Illustration of brain shift. The difference between the pre-operative (*a*) and post-operative (*b*) MR images is shown by image (*c*). The input images have been registered with an affine transformation. The difference image highlights the local deformation happening to the brain due to brain shift, which cannot be captured by global image registration schemes. 28

1.4 Images showing the effect of susceptibility artefacts on interventionally acquired EPI image. (a) shows the susceptibility artefact free T1 weighted MRI image with the edges highlighted by red boundary. (b) shows the corresponding EPI image. Large deformation around the area of resection is evident. 29

2.1 A typical image registration algorithm where a similarity measure is optimised to estimate the geometric transformation that brings the target and source images into alignment. 33

2.2 A typical 3 layer multiresolution scheme used in many image registration algorithms. The images are downsampled with increasing resolution at each multiresolution level. The registration is performed at each level and the next level is initialised with the transformation estimated in the previous level. Typically, the last level performs the registration with the input images at full resolution. This helps avoid the optimiser getting trapped in local minimas due to noise and decreases computation time. 34

2.3 Various rigid transformations applied to a cube. 35

2.4 Affine transformations applied to a cube along the x-axes. Affine transformations include the rigid transformation but can also include shearing and scaling. 35

2.5 Local or non-rigid tranformations applied to a cube. More complex shape deformations can be achieved through local deformations but it comes at the cost of high computational complexity. 36

2.6 Illustration of the difference between using a simple affine versus a non-rigid registration for capturing the deformations due to brain shift. (a) shows the preoperative MRI image. (b) is the MRI image acquired intraoperatively (c) shows the result after doing an affine registration between the preoperative and intraoperative MR images. The checkerboard pattern is constructed from taking alternative square regions from the affinely registered preoperative image and the intraoperative image. It is evident that the brain structures are not aligned. (d) shows the result after performing a non-rigid registration. The checkerboard pattern reveals that the brain structures are now much better aligned. 36

2.7 (a) Timing diagram for a blipped EPI pulse sequence (Schenck, 1996). The entire k-space is acquired with a single RF pulse. The k-space collection starts in one direction sweeping continuously from one side to the next as a consequence of the oscillating frequency gradient. (b) The brief application of the phase encode gradient between echoes moves the trajectory in the k-space to a new row. Susceptibility artefacts are influenced by many factors including echo spacing and echo train length. Shorter echo spacing and echo train lengths give less time for accumulation for phase shifts and typically result in reduced distortions due to tissue susceptibility differences. Figure reprinted with permission. 40

2.8 Example phase unwrapping for MR images of the human head. (a) shows the phase image with multiple 2π wraps. (b) is the unwrapped phase image. The unwrapping was performed using the phase unwrapping software presented in chapter 4. 41

3.1 (a) A simple MRF model for image labelling. MRF is a generative model that models the joint distribution $P(X, Y)$ of the output labels of random variables X and observations Y. (b) A similar CRF model for image labelling. In contrast to the MRF model, it is a discriminative model and models the conditional posterior probability $P(X|Y)$ directly. The unary term in a CRF at a node i is a function of all of the observation data Y and the label x_i rather than just y_i and x_i only as is the case for the MRF model. In the MRF model the pairwise potentials are independent of the observations. However, in the CRF model they are also a function of the observations which allows us to include data dependent pairwise potentials. 50

3.2 A simple graph configuration on a regular grid. The squares denote the source and the sink vertices. The circles denote the other vertices. The red lines show the edges between the nodes. The dashed blue lines show edges which form a cut. (a) shows the original graph. (b) shows a graph with a valid cut. The dashed blue lines separate the graph into two sub-graphs which separate the source and the sink vertices. (c) shows a graph with an invalid cut. The cut is not valid because, if the dashed edges between nodes a and b are removed, the remaining dashed edges still form a cut. 51

3.3 Edge definitions and weights for the unary terms. (a) Graph for E_i when $E_i(1) > E_i(0)$. (b) Graph for E_i when $E_i(1) \leq E_i(0)$. The unary terms can be arbitrary as one of the terminal edges for a graph vertex is always in the cut. Hence, adding the same constant weight to both $w_{\mathbf{s},i}$ and $w_{i,\mathbf{t}}$ does not affect the choice of which edge to cut. 53

3.4 Representation of the smoothing term of the energy function of equation (3.7). (a) is the graphical representation of the edge weights as defined in equation (3.9). (b-e) shows the cut boundary and the cut cost when the vertices take the different label configurations. 53

3.5 Effect of registration on dispersion of the joint histogram. The image intensities have been normalised between 0 and 63 (a) shows the joint histogram when the images are not aligned. Corresponding image intensities do not overlap resulting in a more dispersed joint histogram. (b) shows the decrease in the dispersion of the joint histogram as the images come into alignment. Registration brings the corresponding structures into alignment and there is more overlap between corresponding intensities. This is an example when the image intensities have a linear relationship and come from the same modality. A more complex multi-modal histogram might result when performing registration between different modalities but the same principle about reduction of joint entropy applies. 58

4.1 The proposed workflow for correction of susceptibility artefacts in EPI images acquired during neurosurgery. The field map is calculated using the acquired phase images which are unwrapped using the proposed algorithm. The estimated deformation field and the uncertainty information associated with the phase unwrapping step is used to initialise the image registration step where the EPI image and the corresponding undistorted T1-weighted MRI image is used as the source and the target images respectively. The registration step is selectively driven in regions of high uncertainty to improve the results in areas where the field map might have resulted in a sub-optimal solution. 62

4.2 Graph Construction. (a) shows the construction of the elementary graph for a single pairwise term when $E^{ij}(1,0) - E^{ij}(0,0) > 0$ and $E^{ij}(1,0) - E^{ij}(1,1) > 0$. Note that the graph can only have non-negative edge weights. (b) shows the building of the graph by merging the elementary graphs together. After the graph is constructed, maximum flow algorithm can be used to find the minimum cut (denoted by the dashed line) on the graph. 66

4.3 The various inputs to POSSUM to simulate the MRI phase images. Lesions are manually drawn in the input phantom image. The B_0 inhomogeneity file describes change in magnetic field strength inside the cranium due to tissue susceptibility differences. To calculate these distortions, Maxwell's equations are solved at each voxel in an air-tissue segmentation volume using the perturbation method. Finally, the MRI pulse sequence (eg. EPI) characteristics can be specified for each simulation. 71

4.4 Results from phase unwrapping. (a) is a masked slice through a noise free wrapped image. (b) is the same image where the ground truth unwrapped image was corrupted with Gaussian noise. (c) shows the ground truth unwrapped image. (d) shows the unwrapping result from PRELUDE. Some areas with phase discontinuities are visible in the unwrapped result (highlighted in red). (e) Shows the unwrapped image using the proposed phase unwrapping algorithm where no phase discontinuities are evident. (f) shows the confidence map obtained using the proposed algorithm. 74

4.5 Images showing the result of correcting for susceptibility-induced spatial distortion using our algorithm. (a) shows the gold-standard high resolution T1-weighted image acquired during surgery. (b) shows the uncorrected B0 image with a large geometric distortion around the resected area. (c) shows the result of correcting for susceptibility artefacts using the proposed fieldmap estimation. (d) shows further improvement in the result when combined with the image registration step. 76

5.1 Numerical phantom. (a) shows the simulated cortical layer, (b) shows the 3-dimensional reconstruction of the phantom surface and (c) shows the simulated white matter tracts spanning the phantom. 83

5.2 Figure (a) illustrates the quantification of the optic radiation resection. Figure (b) shows a clinical example where the subject suffered visual deficit. The propagated pre-operative optic radiation (red) overlaps with the resected area (blue). 86

5.3 Bland Altman assessment indicates that the 95% limits of agreement between the structural only image registration and the proposed method ranged from 5.74 mm to -5.85 mm. 86

5.4 Bland Altman assessment indicates that the 95% limits of agreement between the FA only image registration and the proposed method ranged from 4.41 mm to -4.19 mm. . . 87

5.5 Regression line for the predicted damage to Meyer's loop using the proposed registration method and the observed visual field deficit in the 12 patients that suffered from visual field deficits. 87

5.6 Validation of the proposed registration scheme using the intra-operative datasets. The pre-operative images are initially non-rigidly registered to the intra-operatively acquired images using the proposed method. In a second non-rigid registration step, the intra-operative images are registered to the post-operative images. The two deformation fields acquired from the registration steps are composed together to generate the final deformation field. I show that the predicted damage to the optic radiation using this deformation field correlates strongly with the observed VFD. 90

5.7 Figure (a) shows a mesh rendering of the optic radiation obtained by directly registering the pre- to the post-operative dataset using the proposed registration scheme. I showed that it correlates best with the observed visual field deficit and is used as the ground truth for validating the intra-operative registration. Figure (b) shows the meshed optic radiation obtained by composition of the deformation fields obtained by registering the pre- to the intra-operative dataset and the intra- to the post-operative dataset. The solid colour in (a) and (b) denote the 1 voxel thick skeleton of the tracts. Figure (c) shows the close overlap of (a) and (b). 90

6.1 Workflow depicting the various inputs/outputs and the processes for the interventional MRI image analysis workflow. The various inputs and outputs are highlighted in orange while the processes are highlighted in yellow. The final output, which is the optic radiation in the intraoperative space, is highlighted in green. Some of the processes like the brain mask generation are not highlighted in this workflow for reasons of brevity. This workflow is dynamically generated for every surgery through a data interchange format called JSON. The JSON input is parsed and automatically executed within 7-8 minutes. . 94

7.1 Image analysis workflow for the two cohorts. For cohort 1, gradient non-linearity correction was applied and rigid registration was performed to propagate the Optic Radiation in the intraoperative space. For cohort 2, gradient non-linearity correction, susceptibility artefact correction and brain shift correction was applied to propagate the preoperatively segmented Optic Radiation. Image courtesy of Dr. Gavin Winston. 102

7.2 Image analysis workflow for the two cohorts. For cohort 1, gradient non-linearity correction was applied and rigid registration was performed to propagate the Optic Radiation in the intraoperative space. For cohort 2, gradient non-linearity correction, susceptibility artefact correction and brain shift correction was applied to propagate the preoperatively segmented Optic Radiation. 104

8.1 Logo for Niftyreg: a registration framework developed at University College London by Dr. Marc Modat. 105

9.1 Phase unwrapping on SWI phase image for a subject with radiation therapy induced microbleeds. (a) shows the magnitude image. (b) is the wrapped phase image and (c) is the unwrapped phase image obtained with the proposed phase unwrapping algorithm. The algorithm successfully recovers the true phase while preserving information about the microbleeds in the phase image as highlighted. 108

List of Tables

4.1 Misclassification ratio (MCR) and execution time (in seconds) for generating the fieldmap from the synthetic phase images. The MCR is defined as the ratio between the voxels that were incorrectly wrapped to the total number of voxels. For small amounts of phase noise (noted in radians), both the proposed phase unwrapping algorithm and PRELUDE perform similarly. However, for larger noise levels, the proposed algorithm results in lower MCR. The execution time of PRELUDE for high levels of phase noise does not satisfy the stringent time requirements of neurosurgery, while the proposed algorithm executes well within the time constraints. Time-1 refers to the time taken by the proposed method to do phase unwrapping without confidence map estimation. Time-2 is for phase unwrapping along with confidence map estimation. All times are reported in seconds. The mean noise variance in the standard clinical datasets produced on the iMRI was 0.71 radians (corresponding simulation result highlighted in green). 72

4.2 Mean(standard deviation) of the sum of square errors for diffusion tensor fitting in interventionally acquired diffusion weighted images for thirteen subjects. The first column (Initial) shows the initial mean error. The second column (PRELUDE) shows the fit errors after correcting for susceptibility artefacts using PRELUDE. The third column (Fieldmap only) shows the tensor fit errors after correcting for susceptibility artefacts using the fieldmap generated after unwrapping the phase maps using the proposed phase unwrapping algorithm. The fourth columns (Reg. only) shows the tensor fit errors after correcting for susceptibility artefacts using the proposed registration algorithm. The final column (Proposed) shows the tensor fit errors after combining the fieldmap and image registration methods using the proposed method. The proposed method showed statistically significant improvement over the other methods (p-value $< 10^{-3}$). The final row shows the mean tensor fit errors and standard deviation over all the cases. 75

5.1 Mean (standard deviation) Euclidean distance errors in voxels. The second column quantifies the initial misalignment. Subsequent columns correspond to the error after registration using the structural information, the WM tract information and the joint information respectively. Errors are computed within the whole phantom (All) (middle row) and the white matter (WM) (bottom row). 83

5.2 Spearman correlation coefficient (CC) of the measured visual deficit against the predicted damage to the optic radiation by using the three registration schemes (using structural images, fractional anisotropy (FA) images and using both structural and FA images through the proposed method) for the 12 subjects that suffered visual deficit. Columns 3-5 show the predicted damage (reported in mm by measuring the anteroposterior distance from the anterior part of optic radiation to the resection margin) to the optic radiation. The last row shows the CC of the A-P distance against the visual field assessment scores for subjects with VFD. 88

5.3 Mean distance and standard deviation (in mm) between the optic radiation skeleton obtained by direct registration of the pre- and post-operative datasets and the optic radiation skeleton obtained by the composition of the deformation fields from registering the pre- to an intra-operative dataset and registering the intra- to the post-operative dataset. For each case, the analysis was carried out using the two available intra-operative time points i.e. there were two intra-operative scan sessions during the surgery. Additionally, the predicted damage to the optic radiation was also measured in a similar manner to 5.3.2. Spearman correlation coefficient (CC) shows that the propagated optic radiation correlates well with the observed VFD even when using the intra-operative datasets for the intermediate registration. Case 12 was excluded from the correlation analysis as the subject did not suffer any VFD. 91

Chapter 1

Introduction

Magnetic Resonance Imaging (MRI) is an ubiquitous component of epilepsy surgery planning. The primary role of preoperative MRI is to reveal cerebral lesions that might cause epilepsy. Identification of a resectable underlying structural lesion is important, as removal of the focal abnormality improves the chances of the patient being seizure free post surgery. Around 70% of the patients who have the lesions removed through surgery enter remission (Spencer and Huh, 2008). A large body of evidence suggests that the use of multimodal images improves the localisation of epileptic lesions (Salamon et al., 2008; Vulliemoz et al., 2009; Waites et al., 2006; Wehner et al., 2007). Furthermore, MRI can also help in minimising the chance of causing new morbidity due to surgical intervention. In the case of surgical management of focal temporal lobe epilepsy, functional MRI can localise brain areas associated with language, memory and vision and can be used to predict the effects of temporal lobe resection on these brain functions. Diffusion MRI and tractography can be used to localise the main cerebral white matter tracts, like the optic radiation, thereby predicting and reducing visual field deficits due to temporal lobe resection.

In the realm of image-guided neurosurgery, interventional MRI (iMRI) is fast emerging as the preferred imaging choice. The relatively high spatial resolution, excellent soft tissue contrast and the lack of ionising radiation makes iMRI an attractive imaging option for image-guided interventions. Furthermore, along with conventional structural imaging, current commercial iMRI scanners can also perform diffusion and functional imaging which allows for imaging of functionally eloquent brain areas and critical white matter tracts along with the surgical target areas. Maximising the utility of iMRI systems requires the ability to reliably relate the preoperative multimodal imaging data and surgical planning information to the images acquired during the surgical intervention. Evidence has started to emerge that multimodal imaging during surgery can improve patient outcome. In particular, there is an interest in using diffusion weighted imaging (DWI) acquired during intervention to localise particular white matter tracts of interest (Andrea et al., 2012; Chen et al., 2009; Sun et al., 2011).

1.1 Clinical Background

Epilepsy is a common and debilitating neurological disorder. Among the various types of epilepsy, temporal lobe epilepsy is the most common focal epilepsy. Around 40% of patients with temporal

lobe epilepsy (TLE) are refractory to medication (Semah and Ryvlin, 2005), and anterior temporal lobe resection (ATLR) is an established and effective treatment for these patients (Wiebe et al., 2001). However, a careful balance has to be established between obtaining seizure control and minimising the chance of causing new morbidity. An important source of morbidity during anterior temporal lobe resection (ATLR) arises due to damage to a critical white matter tract, the Optic Radiation, during the surgical intervention. This can lead to severe visual field deficits (VFD) that can result in a significant loss of vision, even if the patient is seizure free post surgery. These deficits are typically caused by damage to Meyers loop (illustrated in figure (1.1), the most anterior part of the Optic Radiation, which shows considerable variability between patients in its location. Since the Optic Radiation cannot be identified visually during surgery, its accurate localisation and real-time display during the intervention could be crucial in improving the surgical outcome for patients undergoing anterior temporal lobe resection.

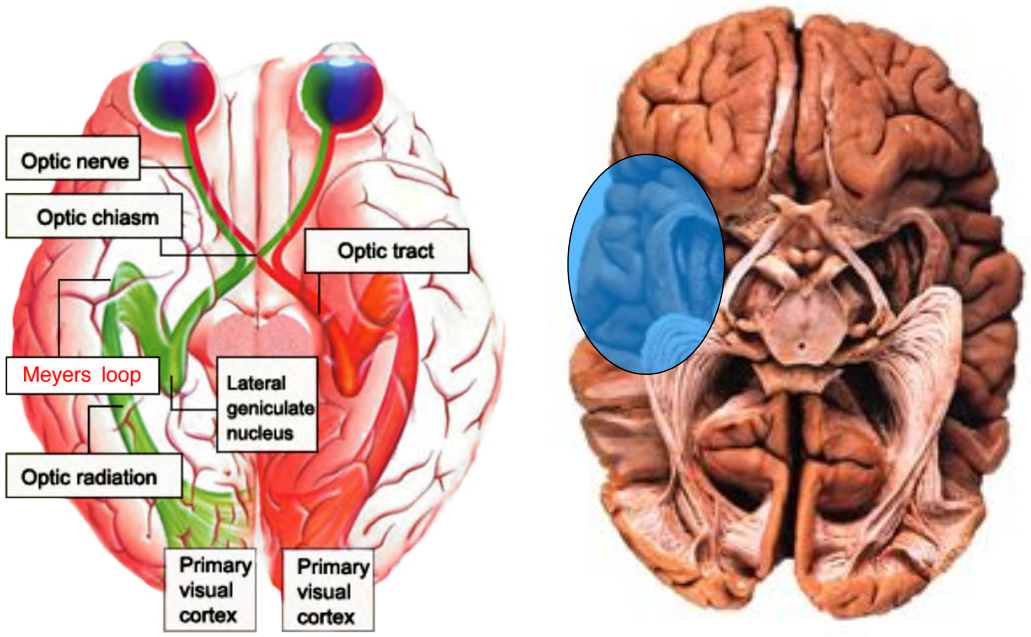

Figure 1.1: Human Visual System: (a) shows a schematic of the visual wiring in the brain (courtesy www.thebrain.mcgill.ca). As shown, the Meyers loop passes through the temporal lobe and hence is at risk of injury during the surgical intervention. (b) shows a dissected brain (courtesy Virtual Hospital) and the Meyers loop can be clearly identified. The blue oval highlights the area typically affected by temporal lobe resection, which can result in damage to the Meyers loop.

One of the key challenges facing accurate localisation of the optic radiation during surgery is the estimation of the soft tissue deformation (collectively termed *brain shift*) that occurs after craniotomy and cerebrospinal fluid drainage during a typical neurosurgical procedure. Brain shift can occur due to a variety of reasons including gravity, brain swelling, cerebrospinal fluid drainage, tumour mass effects or surgical intervention and leads to nonlinear deformation of the structures of interest, like the optic radiation. Various studies have reported significant brain shift (up to 25 mm) after craniotomy (Hall and Truwit, 2005; Nabavi et al., 2001; Nimsky et al., 2001). Brain shift was examined by Hall and Truwit (2005) and those structures found to shift significantly during surgery were located either directly over or

within a 1-cm radius of the lesion being removed. Intra-operative MRI (iMRI) provides a way to localise the structures of interest during the surgical procedure by enabling imaging of the patient intermittently during surgery.

1.2 MRI in Neurosurgery

The iMRI setup at the National Hospital for Neurology and Neurosurgery (NHNN) in London consists of a 1.5 tesla Siemens (Erlangen, Germany) Espree MRI scanner. There is a dedicated operating room 8 channel MR head coil which incorporates a surgical headrest. The operating table is fitted with an MR compatible head-holder and is placed outside the 5 Gauss line during surgery which enables the surgeons to perform the procedure using standard non MR-compatible surgical instruments. The table can interface with the MR scanner to allow the patient to be moved in and out of the scanner for intra-operative imaging. The facility is equipped with a BrainLAB VectorVision® Sky neuronavigation system which provides real-time tracking of surgical markers and tools, *global* image registration and visualisation facilities. The operating room is also equipped with an Opmi Pentero confocal surgical microscope (Carl Zeiss), supporting the injection of colour overlays from the navigation system. The location of the microscope's focal point is tracked using the navigation system and an array of four infra-red reflectors mounted on the microscope's optical head. A snapshot of the iMRI surgical room is shown in figure (1.2).

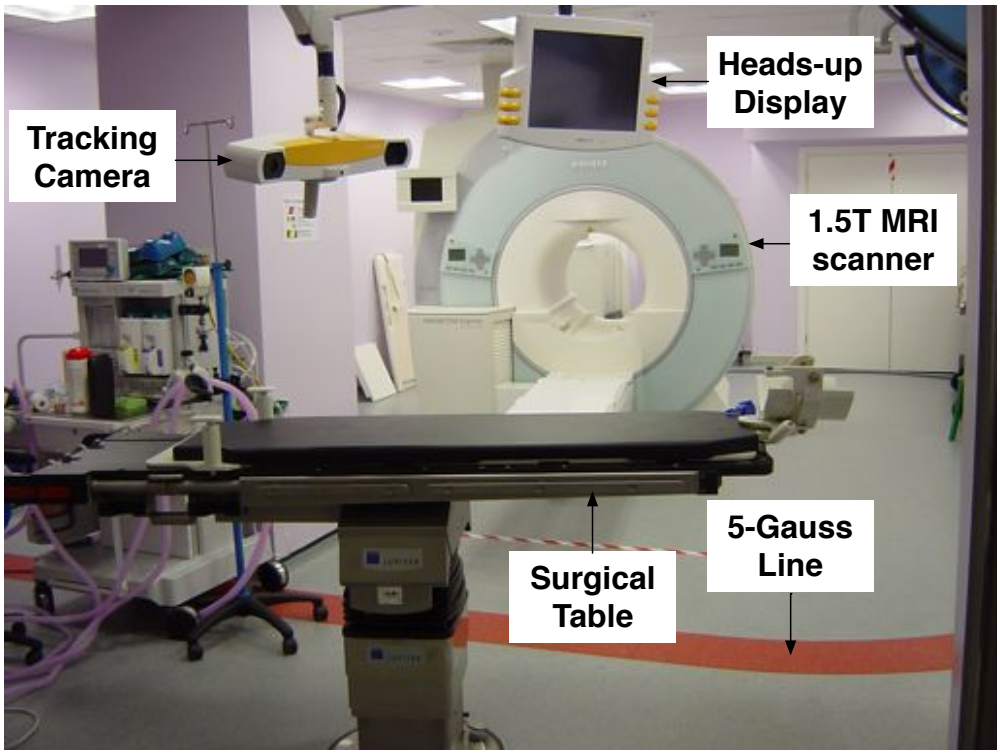

Figure 1.2: The interventional MRI surgical suite at the National Hospital for Neurology and Neurosurgery with a 1.5 tesla MR scanner and neuronavigation equipment. The surgical table interfaces with the scanner to enable the patient to be moved in and out of the scanner efficiently during surgery.

1.3 Challenges in iMRI neuronavigation

The aim of neurosurgical image guidance is to maximize the resection of target lesions while conserving healthy and important brain tissues like critical white matter tracts and functionally eloquent brain areas. There are several challenges, unique to neurosurgery, that need to be overcome for effective neuronavigation.

1.3.1 Brain Shift

As mentioned before, the main challenge to achieving effective neuronavigation is to accurately estimate the non-linear deformations in the brain arising due to brain shift. The current state of the art commercial neuronavigation systems assume a rigid body relationship between the preoperative and intra-operative images, which limits their ability to accurately estimate brain shift during neurosurgery (Shamir et al., 2009). The deformation caused by the brain shift cannot be accurately determined using a rigid or affine transformation, making it difficult to rely on pre-operative images for accurate identification of surgical targets and eloquent brain areas. Figure (1.3) presents an interventional MR case and highlights the differences between an MRI acquired before surgery and another acquired during the surgical procedure. The local deformations cannot be recovered using a global image registration and the difference image shown in figure (1.3 c) highlights the large registration errors around the area of surgical resection. A potential solution to this problem is to non-rigid image registration to estimate the deformation arising due to brain shift. However, non-rigid image registration algorithms are typically computationally expensive and are also harder to validate which hinders their use in clinical neuronavigation systems (Crum et al., 2004).

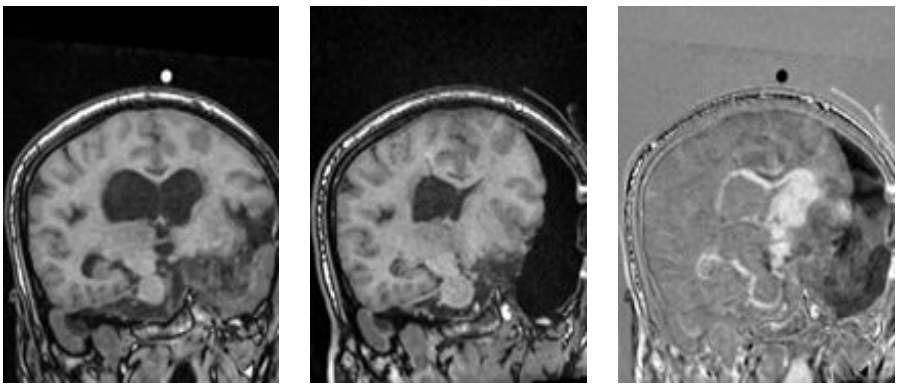

a *b* *c*

Figure 1.3: Illustration of brain shift. The difference between the pre-operative (*a*) and post-operative (*b*) MR images is shown by image (*c*). The input images have been registered with an affine transformation. The difference image highlights the local deformation happening to the brain due to brain shift, which cannot be captured by global image registration schemes.

1.3.2 Artefacts in iMRI Images

Single shot echo planar images (EPI) (Mansfield, 1977) are widely used in diffusion weighted imaging sequences. The low bandwidth in the phase encode direction makes them prone to geometric and intensity distortions arising due to susceptibility artefacts (Jezzard and Balaban, 1995). The problem becomes more severe in the neurosurgery setting due to surgical resection, which creates a tissue-air interface. As

a result, the susceptibility artefacts at the resection boundary become especially severe due to large difference in magnetic susceptibility between tissue and air, which creates large B_0 field inhomogeneities. If intraoperative EPI images are to be used for neurosurgical guidance, it is important to account for these artefacts as it is precisely around the areas of resection that the navigation system needs to be accurate for good patient outcome. Figure (1.4) shows an example of the effect of susceptibility artefacts on interventionally acquired EPI images. Large distortions around the area of resection can be seen due to the B_0 field inhomogeneity introduced by the tissue-air interface.

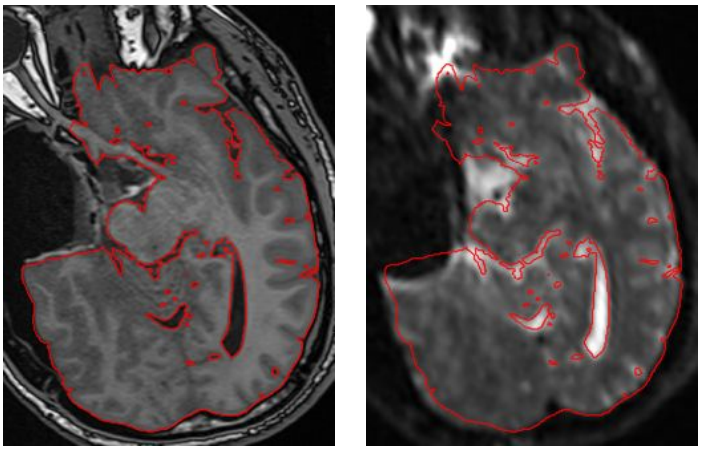

a b

Figure 1.4: Images showing the effect of susceptibility artefacts on interventionally acquired EPI image. (a) shows the susceptibility artefact free T1 weighted MRI image with the edges highlighted by red boundary. (b) shows the corresponding EPI image. Large deformation around the area of resection is evident.

1.3.3 Integration Into Surgical Workflow

The neurosurgery environment is complex and has stringent quality assurance and time constraints. Any change to the current surgical workflow must be shown to have clinically relevant benefit to patient outcome. Any proposed changes must be approved by an ethics committee and the components of the workflow must be thoroughly validated. In addition, there are strict time constraints associated with a neurosurgical procedure. Any proposed image acquisition and processing should aim to have minimal interruption to the surgical workflow. The current patient transfer time from the intra-operative scanner, after an imaging session, to the surgical bed at NHNN is between $7 - 12$ minutes. This time should be used for the processing of the images to localise the structures of interest and make the results available for neuronavigation within this time window ensuring that no additional time due to data processing is added to the surgery. Hence, the image processing algorithms designed to work in the neurosurgical environment need to be accurate, robust *and* computationally efficient.

1.4 Methodological Contributions

The primary goal of my work has been to address these challenges using medical image analysis and develop an image-guided neurosurgical platform that can be used with the iMRI and the neuronavigation setup at NHNN. To this end, I will highlight the primary methodologial contributions of my doctorate

work.

1. A susceptibility artefact correction that uses a novel phase unwrapping algorithm that can efficiently compute the B_0 field inhomogeneity map as well as the confidence associated with the estimated field map.

2. A non-rigid registration algorithm that can be used with the confidence map and the estimated field map from the phase unwrapping step and selectively refine the results in areas where the confidence in the estimated field map is low.

3. A novel, near real-time bivariate non-rigid image registration that integrates structural and diffusion MRI images in a unified similarity measure is presented. The proposed algorithm can be used within the time constraints of a neurosurgical procedure by leveraging the parallel processing capabilities of graphical processing units (GPU). I show that it can estimate brain shift and localise the optic radiation more accurately than using structural or diffusion MRI images alone.

4. A framework to integrate the methodological developments presented in this thesis into the clinical workflow at NHNN, which is being used during neurosurgical procedures to inform surgical decisions. Initial validation and evaluation of this framework in regard to the clinical outcome is also presented in this thesis.

1.5 Thesis Organization

The following chapter describes the computational techniques that are used in this thesis. In particular, I will describe the theory behind graph cuts based optimisation and medical image registration.

Chapter 2 is a literature review describing previous work in the areas of brain shift estimation and correction of susceptibility artefacts.

Chapter 3 describes the computational techniques behind graph cuts and medical image registration, which are used in this thesis.

Chapter 4 describes a novel susceptibility artefact correction algorithm that can be used in the neurosurgical setting. The proposed algorithm combines field map and image registration based correction techniques in a unified framework.

Chapter 5 describes a novel brain shift estimation technique that utilises information from structural and diffusion MRI images which is fast enough to be used in the neurosurgical setting.

Chapter 6 describes the clinical integration of the methods developed in chapters 4 and 5 at the National Hospital for Neurology and Neurosurgery in London.

Chapter 7 presents the initial clinical findings from using the system for temporal lobe resections.

Chapter 8 highlights some of the open source software contributions generated through my work.

Chapter 9 concludes the thesis by highlighting further developments that can be undertaken to carry this work forward.

Chapter 2

Literature Review

2.1 Brain Shift Estimation

One of the key challenges facing accurate image-guided surgical systems is the estimation of the soft tissue deformation (collectively termed *brain shift*) that occurs after craniotomy and tissue resection during a typical neurosurgical procedure. Brain shift can occur due to a variety of reasons including gravity, brain swelling, cerebrospinal fluid drainage, tumour mass effects or surgical intervention. An initial attempt to compensate for brain shift was developed by Kelly et al. (1986) where metal beads were implanted in the brain cortex during image guided laser resection of tumours. Brain shift caused displacement of the metal beads and their position on subsequent radiographs acquired during the intervention was then used to update the location of the tumours. Brain shift was studied quantitatively by Hill et al. (1998) who measured the deformation between the time of preoperative imaging and the start of surgical resection (i.e. after craniotomy but before any soft tissue intervention) for 21 patients. They reported mean displacements of 1.2mm, 4.4mm and 5.6 mm for the dura, first and second brain surfaces respectively. Other studies performed with the aid of intraoperative imaging, including iMRI, suggest that brain shift could be quite variable and report displacements from 1 cm upto 2.5 cm during the intervention (Hastreiter et al., 2004; Nimsky et al., 2000; Roberts et al., 1998; Winkler et al., 2005). Extensive work has been reported in the computer assisted intervention literature to compensate for brain shift using various intraoperative technologies like MRI, CT and ultrasound (Hall and Truwit, 2005; Jolesz, 2005; Kaibara et al., 2002; Lindner et al., 2006; Nabavi et al., 2001; Nagelhus et al., 2006; Nakao et al., 2003; Siewerdsen et al., 2005). There are relative merits and disadvantages associated with each of these modalities. Mobile gantry versions of CT scanners with specialised modifications to accommodate head fixation devices have been developed which make them feasible to be used in the intraoperative setting (Okudera et al., 1991). Despite these advancements, the widespread use of intraoperative CT has been hampered by concerns over radiation exposure to the subject. Ultrasound has the advantage of being portable and very cost efficient compared to CT and MRI. It also has the added advantage of not exposing the subject to any harmful ionising radiation. However, its use is limited by the low signal to noise ratio and operator dependency.

MRI has steadily been gaining ground as the imaging modality of choice for guiding interventions. iMRI offers superior soft tissue contrast without exposing the subject to harmful ionising inherent in CT.

The image quality of an iMRI scanner is contingent upon various factors including the field strength, bore size, scanner design and the requirements for patient accessibility and integration in the operating theatre. Higher quality images are obtained using a closed bore scanner whereas open bore scanners give maximal access to the patient. A low field strength (0.12T) iMRI system (Medtronic Navigation, Minneapolis, MN) allows for partial imaging of the head with the entire surgical procedure being conducted within the magnetic field using standard surgical instruments (Hall and Truwit, 2005). The 0.5T "double donut" is a mid-field strength iMRI scanner and was the first iMRI scanner developed and used specifically for interventional use (Black et al., 1997). Due to its design constraints, the magnetic field generated by these scanners is inhomogeneous, reduced signal-to-noise ratio and limited physiological and functional imaging capabilities (Martin et al., 2000).

High-field iMRI scanners (1.5T or greater) have the advantage over low- and mid-field scanners of higher image quality and availability of diverse MRI modalities like diffusion, perfusion and functional imaging. However, due to the high field strength the surgery has to be performed beyond the effective magnetic field and the patient needs to be transferred to and from the scanner when they need to be imaged during surgery. Recently, high-field (3 tesla) ceiling-mounted MRI systems have been made available. The ceiling-mounted MRI scanner can be moved in and out of the operating room as needed. With this innovation, the patient does not need to be transferred into the operating/angiography table for imaging. The Advanced Multimodality Image Guided Operating (AMIGO) suite at the Brigham and Women's Hospital in Boston, USA employs such a ceiling-mounted 3 tesla MR system. The high-field systems are closed bore systems and thus access to the patient during imaging is limited. The main disadvantage of iMRI is that large installation and setup costs are involved especially when adding in the cost of adapting or building the operating room to support it.

Advances in MRI have permitted the acquisition of rich information *preoperatively* such as structural, functional and high resolution diffusion weighted imaging. Other modalities like PET and SPECT imaging are also widely used preoperatively. These are used for surgical planning particularly to localise surgical targets (like the epilepsy focus region) and eloquent functional brain regions and critical white matter tracts that should be avoided during surgery. Surgical constraints along with iMRI limitations do not allow for acquisition of this data intraoperatively, while the preoperative images cannot be used directly for surgical guidance due to brain shift. A lot of work has been done in using the intraoperative images as means to express this rich preoperative information in the intraoperative geometry as opposed to using the lower quality and limited intraoperative images directly for guidance. This essentially means estimating the soft tissue deformations that underlie brain shift and updating the preoperative images and surgical plans to reflect the positional shift in brain structures of surgical interest.

2.1.1 Image Registration Based Brain Shift Estimation

Medical image registration is ubiquitous in medical image analysis and is the most widely used method for estimation of brain shift. Broadly speaking, image registration is the process of bringing a set of images into spatial alignment. In the current context, image registration consists of bringing the preoperative images (termed *source* or *floating* images) into alignment with the intraoperative images (termed

target or *reference* images). Various image registration algorithms have been proposed and they all follow the same general principle where the image registration task consists of finding the geometric transformation which makes the target and source images similar to each other based on some measure of similarity. Hence, the image registration task can be seen as an *optimisation* problem where we seek to find the optimum geometric transformation which will maximise the measure of similarity between the two images. Schematically, a typical registration algorithm can be visualised as figure (2.1). Broadly speaking, an image registration algorithm consists of 3 distinct modules: the transformation model, the similarity measure and the optimisation algorithm. It is an iterative process where during each iteration the source image is *warped* using the current estimated transformation. The warped and the target images are used by the similarity measure function which is maximised by the optimiser to estimate the most likely transformation between the target and source images. Non-rigid registration is typically an ill-posed problem and for such registrations the similarity measure is usually a combination of data and regularisation terms. The data term comes from the image similarity measure which is typically derived from the image intensities, image features (such as landmarks for example) or a combination of both. The regularisation terms can be viewed as a prior belief on the form of the underlying transformations and typically impose a penalty on transformation complexity. The registration is typically performed in a coarse-to-fine multiresolution pyramidal scheme where the initial alignment is performed with smoothed and downsampled input images and each successive resolution level is initialised with the transformation estimated at the previous level. This helps avoid the optimiser to find a local minima whilst decreasing computation time. This multiresolution approach is highlighted schematically in figure (2.2).

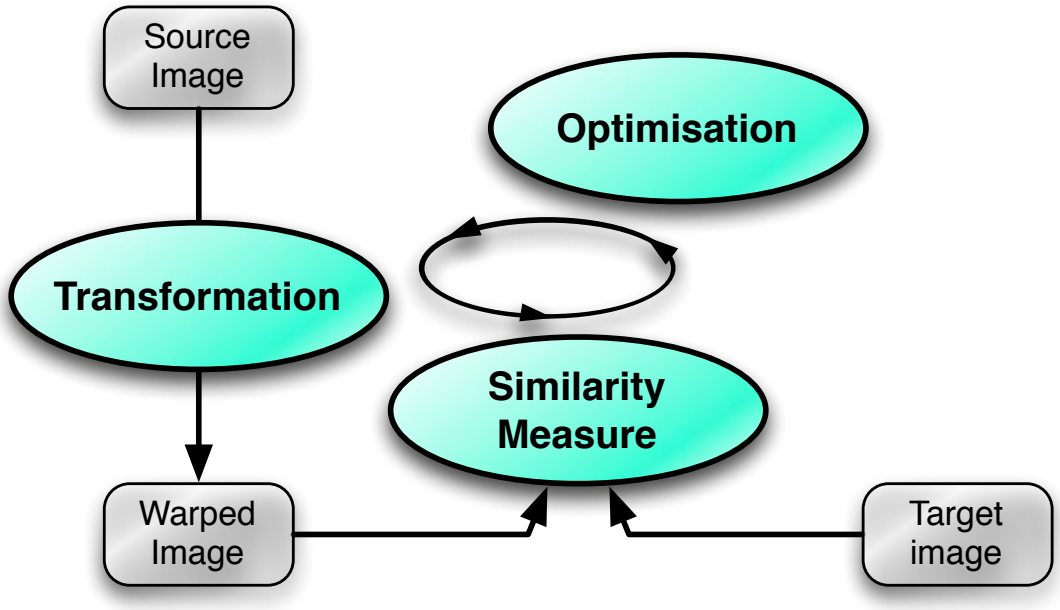

Figure 2.1: A typical image registration algorithm where a similarity measure is optimised to estimate the geometric transformation that brings the target and source images into alignment.

In the context of brain shift estimation, it is useful to divide the registration algorithms based on the employed transformation model. With this criteria in mind, the image registration algorithms broadly fall

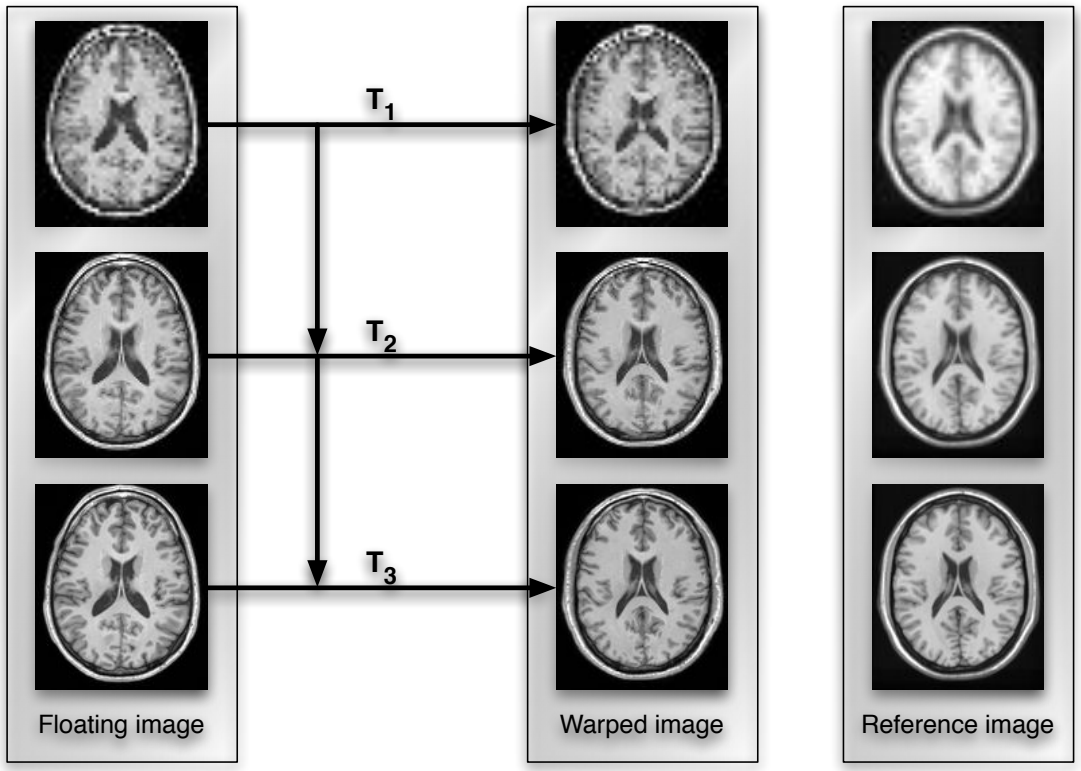

Figure 2.2: A typical 3 layer multiresolution scheme used in many image registration algorithms. The images are downsampled with increasing resolution at each multiresolution level. The registration is performed at each level and the next level is initialised with the transformation estimated in the previous level. Typically, the last level performs the registration with the input images at full resolution. This helps avoid the optimiser getting trapped in local minimas due to noise and decreases computation time.

into two categories: *global* and *local* image registration methods. Global image registration use linear transformations to relate the target image to the source image. This can be the rigid body model (used by most commercial neuronavigation systems) consisting of global translation and rotation transformations or the affine model which also includes global scaling and shearing in the transformation model. These are depicted in figures (2.3) and (2.4). The rigid model consists of 6 parameters in 3D and an object does not change shape under a rigid transformation. The affine model consists of 12 parameters and the shape changes due to scaling and shearing are global i.e. they affect the whole object equally.

In contrast, the local or non-rigid image registration algorithms typically consist of non-linear transformation models which use localised transformation models to align the target and source images. These localised non-linear transformation can capture much more complex shape deformations as highlighted in figure (2.5). However, they need to employ transformation models with many more degrees of freedom than global algorithms and are computationally much more expensive. Due to this added computational complexity, which is difficult to resolve within the time constraints inherent in a neurosurgical procedure, the current commercial neuronavigation systems use rigid body transformations to align the preoperative and intraoperative images. This, however, results in decreased ability to accurately map preoperative information onto the intraoperative scene as the non-linear deformations caused by brain shift cannot be accurately captured using a global transformation model. This is shown in figure (2.6) where the affine

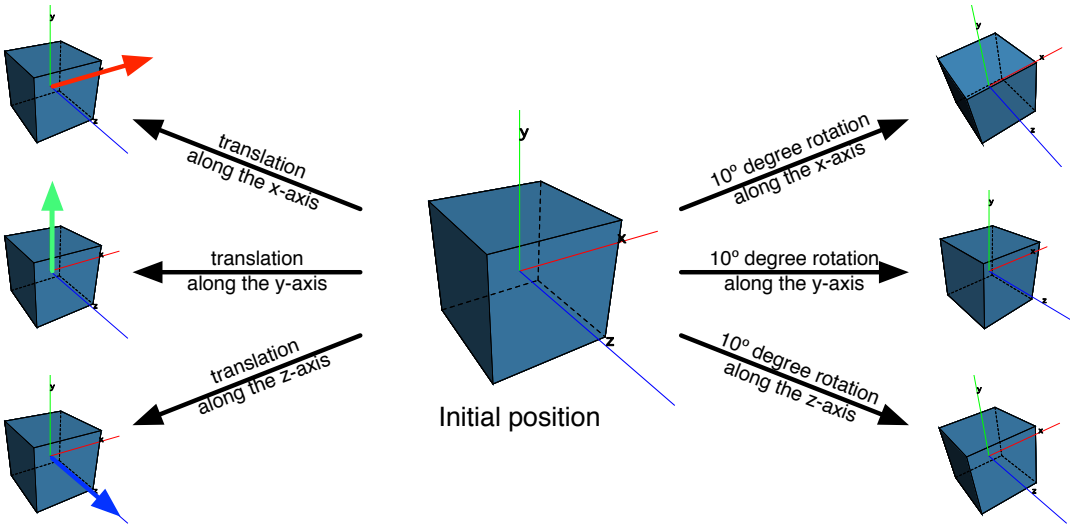

Figure 2.3: Various rigid transformations applied to a cube.

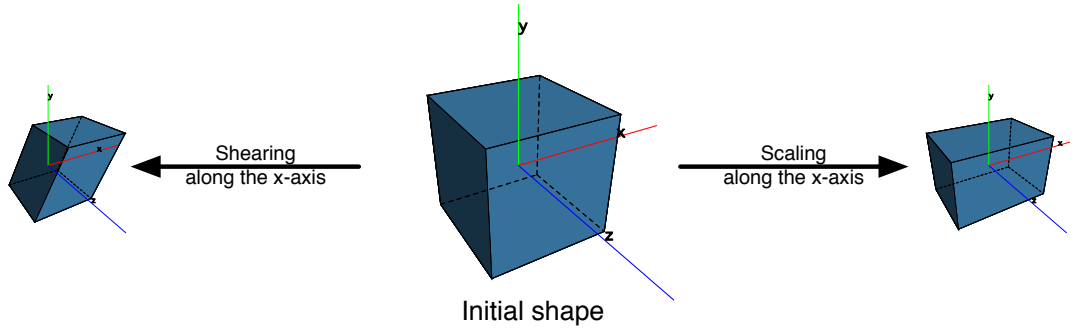

Figure 2.4: Affine transformations applied to a cube along the x-axes. Affine transformations include the rigid transformation but can also include shearing and scaling.

registration fails to capture the brain shift and the preoperative and intraoperative images are not aligned. A more accurate alignment is achieved when a non-rigid registration algorithm is used to capture the brain shift.

The current commercial neuronavigation systems assume a rigid or an affine relationship between the preoperative and postoperative images. As described, these *global* registration algorithms are not enough to accurately estimate the non-linear deformations occuring due to brain shift during neurosurgery. To address this limitation, efforts have been directed towards developing non-rigid registration algorithms for image-guided neurosurgery. The earlier algorithms proposed using a block-matching based transformation model to estimate the non-linear deformations of the brain tissue. Block-matching based transformation estimation divides the target and source images into sub volumes and searches for the optimal translation for each sub volume. Hata et al. (1998) used multimodal non-rigid registration between preoperative and intraoperative MRI using a block-matching based transformation model and mutual information as a similarity term. The algorithm used a coarse to fine multiresolution scheme and could register 3D MRI volumes (dimensions of $256 \times 256 \times 124$) in approximately 21 minutes. Another non-rigid algorithm based on block-matching and designed specifically for brain shift estimation

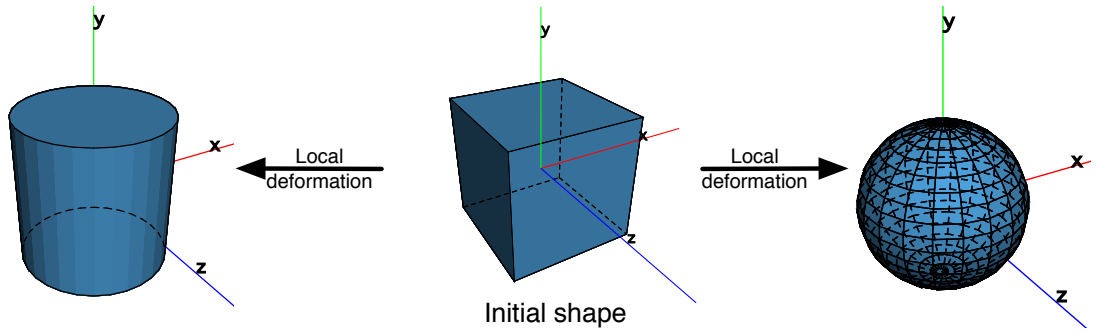

Figure 2.5: Local or non-rigid tranformations applied to a cube. More complex shape deformations can be achieved through local deformations but it comes at the cost of high computational complexity.

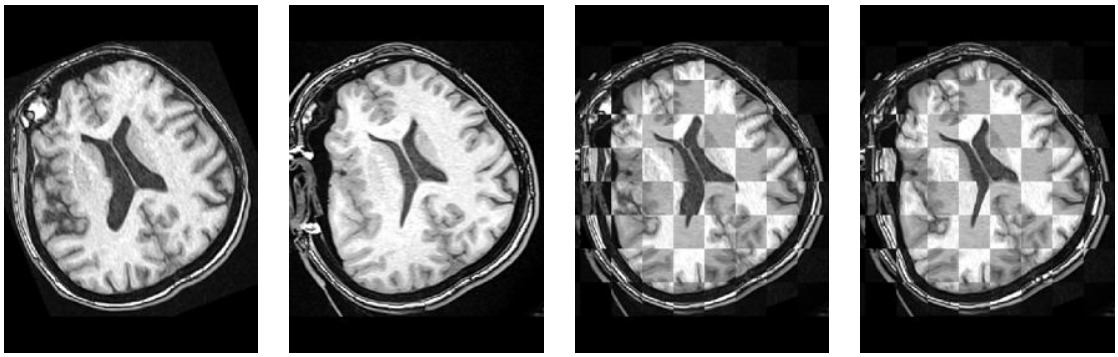

a b c d

Figure 2.6: Illustration of the difference between using a simple affine versus a non-rigid registration for capturing the deformations due to brain shift. (a) shows the preoperative MRI image. (b) is the MRI image acquired intraoperatively (c) shows the result after doing an affine registration between the preoperative and intraoperative MR images. The checkerboard pattern is constructed from taking alternative square regions from the affinely registered preoperative image and the intraoperative image. It is evident that the brain structures are not aligned. (d) shows the result after performing a non-rigid registration. The checkerboard pattern reveals that the brain structures are now much better aligned.

was proposed by Clatz et al. (2005) which used local normalised correlation coefficient as a similarity measure. They combined it with a patient specific biomechanical model of tissue deformation to ensure that the estimated brain shift is physically plausible. This work was validated on retrospective data and was subsequently extended by Archip et al. (2007) and used in the neurosurgical setting. A recent block-matching based approach was proposed by Gu and Qin (2009) where an outlier detection scheme that aimed to reduce the influence of missing features or mismatches introduced by tumour resection was used to increase robustness. Recently, a full Bayesian approach to non-rigid registration problem was adopted by Risholm et al. (2013). They characterised the full posterior distribution on the space of deformations using Markov Chain Monte Carlo sampling methods. Using this method, it was possible to also estimate the confidence associated with the estimated solution to the registration problem. They showed that the registration uncertainty increases at the area of resection and that the posterior distribution around the resection site could be multimodal. A limitation of this work is the extremely long computation times that can last from several hours to a few days, which makes this technique infeasible for use in the surgical setting. Another point to note is that the iMRI based work so far tend to use only the structural MRI information from the intraoperative scan sessions to register the preoperative

images and the recent growth in multimodal imaging capabilities of iMRI scanners has not yet been exploited in this context. Cortical surface registration has also been used in the intraoperative setting to infer volumetric brain deformation (Miga et al., 2003; Paul et al., 2009; Sinha et al., 2005; Škrinjar et al., 2002). Cortical surface data can be acquired with a wide range of imaging modalities like ultrasound and stereoscopic and laser range scanners. However, since the measured data is sparse, prior information needs to be included for accurate inference of sub-surface displacements.

There is considerable interest in adapting a wide variety of imaging modalities to the neurosurgical setting. To complement this, multimodal image registration has drawn significant interest from the medical image analysis community. Due to its low cost, real time imaging capabilities and non invasive nature, ultrasound is a popular modality in the intraoperative setting. Ultrasound imaging has been used in brain examination over the last two decades (Rubin et al., 1980) and several studies have demonstrated that ultrasound is useful in detecting tumour margins, brain shift and residual tumour tissues (Dohrmann and Rubin, 2001; Moiyadi and Shetty, 2011). Several neuronavigation systems with integrated 3D ultrasound technology have been developed and used for various procedures (Unsgaard et al., 2006). Significant work has been done in using intraoperatively acquired ultrasound images to warp the preoperative images to the intraoperative setting using registration techniques. Landmark based registration represents the majority of these approaches. Earlier works used manually identified homologous landmarks in the ultrasound image volume and the preoperative MRI were used to estimate the non-linear warp between the images (Comeau et al., 2000; Gobbi et al., 2000). The use of blood vessels as homologous landmarks in preoperative and ultrasound image have been utilised for brain shift correction (Chen et al., 2012; Lee et al., 2011; Reinertsen et al., 2007). The cerebral vasculature is a good candidate for use in image registration as they are densely distributed over the cerebral context and move with the surrounding tissue, which allow the brain shift deformations to be captured by the vasculature displacement. King et al. (2000) applied Bayesian theory and finite element modelling to estimate the brain shift. The location and shape of the object of interest are modelled as random variables and the algorithm estimates the most likely configuration of these variables given the input surface mesh generated from the preoperative image and the observed 3D ultrasound image during surgery. Intensity based registration approaches are less common primarily due to difficulty of finding a function matching ultrasound image intensities with MR image intensities. There has been some work on overcoming this problem by preprocessing the images in order to register more similar images. Arbel et al. (2001) built “pseudo” ultrasound images of objects of interest from segmented preoperative MRI images which were then used in the registration to intraoperative MRI using a cross-correlation based similarity measure. Another purely intensity based approach was proposed by Roche et al. (2001) which used the bivariate correlation ratio as a similarity measure and attempted to relate ultrasound intensities with both MR intensities and gradient information. This approach was, however, used only to perform a rigid registration.

Significant efforts have been geared towards speeding up the execution times of non-rigid registration algorithms. Hastreiter et al. (2004) exploited the 3D texture mapping capabilities of graphics hardware (GPU) to accelerate all interpolation operations during the registration. Further acceleration

was achieved with an adaptive refinement of the deformation estimate focusing only on the main deformation areas. Rohlfing and Maurer (2003) used shared-memory multiprocessor environments to speed up the free form deformation (Rueckert et al., 1999) based registration and demonstrated that it could be adapted for the brain shift problem. More recently, Modat et al. (2010) presented a refactored version of the free form deformation algorithm which also took advantage of modern graphics hardware through the use of CUDA framework (NVIDIA, 2008).

2.1.2 Biomechanical Model Based Brain Shift Estimation

Biomechanical models are becoming increasingly attractive for estimating brain shift intraoperatively because they provide whole-brain displacement fields which can be used to update the preoperative MRI images for subsequent guidance. They can be coupled with sparse intraoperative data, which can be acquired with cheaper non-tomographic imaging modalities intraoperatively, and are thus cost-effective. These models attempt to simulate the brain tissue response and predict their displacement under the particular surgical conditions. Based on different laws and assumptions, the models that are widely used can be grouped into: viscoelastic models, coupled fluid-elastic models, and porous media models (Carter et al., 2005). Viscoelastic models were one of the earliest models to be adapted for brain shift estimation and assume that brain tissue is an isotropic linear material obeying Hooke's law with a storage and loss modulus (Engin and Wang, 1970; Wang and Wineman, 1972). Coupled fluid-elastic models can model more complex behaviour and can assign different biomechanical laws to different regions of the brain. For example, Hooke's law can be used to represent the behaviour of solid brain tissue, whereas Navier-Poisson's law can be used to represent the cerebrospinal fluid in the brain (Hagemann et al., 1999). Porous media models consider brain as a spongy material where the void spaces are saturated with fluid, whose model can be represented by multi-phase consolidation theory. The tissue motion is characterised by an instantaneous deformation at the area of contact followed by additional deformation resulting from exiting pore fluid driven by a pressure gradient (Paulsen et al., 1999). These biomechanical models allow for the simulation of brain tissue motion under various surgical conditions. The displacement field computed from these simulations can be used to warp the preoperative image and update them to reflect the current state of the brain under the intraoperative setting.

To accurately simulate the deformation under a given surgical scenario, information from the current surgical setting need to be derived to simulate the deformation. The constraints are usually derived from intraoperative imaging and the model is therefore data-driven. Usually, a sparse displacement field is measured from a partial volume or partial surface of the brain at two distinct surgical stages (e.g. before and after craniotomy). Carter et al. (2005) grouped these intraoperative data measurements into two types - surface and sub-surface displacements. Various methods for measuring surface displacements have been used including contact measurements where points are acquired on the brain surface using a tracked pointer (Comeau et al., 2000; Hill et al., 1998; Roberts et al., 1998), laser range scanning (Audette et al., 2003; Miga et al., 2003), stereopsis which uses two calibrated cameras to reconstruct a three dimensional surface (Paul et al., 2009; Sun et al., 2005). The intraoperative data provided by these measurements strongly depends on the size of the craniotomy, which should be kept as small as possible. Sub-surface

displacements use 3D imaging intraoperatively to obtain dense displacement fields. Intraoperative CT was used in animal models (Miga et al., 2000) but it suffers from low soft tissue contrast and exposes the subject to harmful ionising radiation and has not been considered for human subjects. Intraoperative ultrasound has been used by a large number of studies and continues to be an area of active research. However, ultrasound cannot include the full brain volume and the low signal to noise ratio remains a significant problem. Usually only a few anatomical landmarks are visible in the ultrasound scan and they usually become less visible as the surgery progresses. Matching homologous points between ultrasound and other modalities is also a challenging problem. Additionally, ultrasound requires tissue contact, which may induce additional deformations that needs to be modelled. Recent developments are also seeing iMRI being increasingly used in the context of biomechanical modelling of brain shift (Archip et al., 2007; Warfield et al., 2002, 2005; Wittek et al., 2007). Recently, a model based approach was proposed in where an atlas of solutions that account for brain shift caused by various parameters like gravity, edema and neurosurgical drugs were computed in Chen et al. (2011) a manner similar to Dumpuri et al. (2003). This work explicitly models the dural septa and shows that this helps improve the prediction of sub-surface brain shift.

2.2 Susceptibility Artefacts in MRI

Ideally, the magnetic field in an MRI scanner would be perfectly homogeneous throughout the field of view when no external gradients are applied. However, different tissue types have dissimilar paramagnetic properties and they interact with the magnetic field in different ways. Biological tissue comprises of mostly water and air which have very dissimilar magnetic susceptibility values. Water exhibits volume magnetic susceptibility of about -9×10^{-6} in SI units whilst air has magnetic susceptibility of about 0.4×10^{-6} (Schenck, 1996). Hence, if the structure to be imaged comprises of materials with very different magnetic susceptibilities (like water and air for example), the magnetic field becomes distorted and does not stay homogeneous. These field inhomogeneities can be partially removed by *shimming*, which involves generation of corrective offset magnetic fields aiming to make the magnetic field homogeneous. However, this only partly alleviates the problem and significant field inhomogeneities usually remain even after shimming. A consequence of the magnetic field inhomogeneities is dephasing of spins and frequency shifts between the surrounding tissues which results in non-linear spatial and intensity distortions of the anatomy (Jezzard and Balaban, 1995).

Single-shot EPI provides high temporal resolution and is routinely used in diffusion weighted imaging (DWI) and functional magnetic resonance imaging (fMRI) sequences. EPI performs rapid acquisition by sampling the entire frequency space of the selected slice with one excitation pulse and fast gradient switching as shown in figure (2.7) (McRobbie et al., 2006). However, this results in very low spectral bandwidth in the phase encode direction and makes EPI extremely susceptible to magnetic field inhomogeneities. This problem is particularly severe in the context of image-guided neurosurgery as the tissue resection introduces a substantial air/tissue interface causing large geometric and intensity distortions around the area of resection. It is extremely important to correct for these distortions as it is especially around the area of resection where the need for accurate image-guidance is paramount.

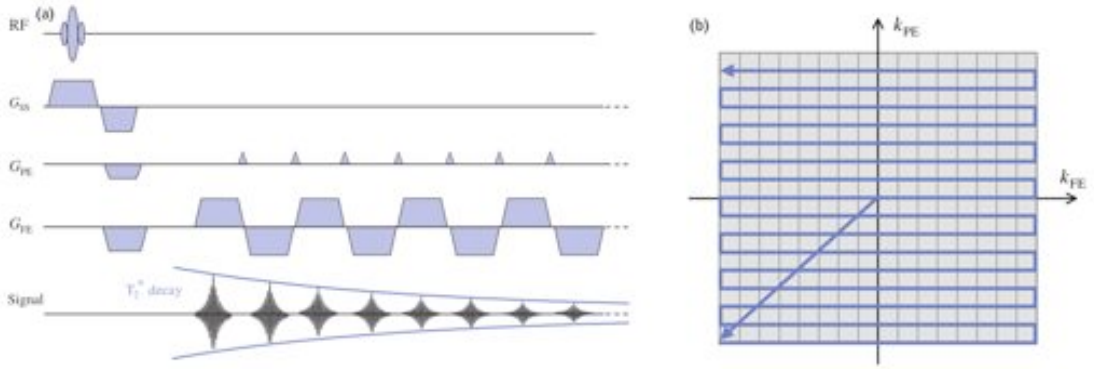

Figure 2.7: (a) Timing diagram for a blipped EPI pulse sequence (Schenck, 1996). The entire k-space is acquired with a single RF pulse. The k-space collection starts in one direction sweeping continuously from one side to the next as a consequence of the oscillating frequency gradient. (b) The brief application of the phase encode gradient between echoes moves the trajectory in the k-space to a new row. Susceptibility artefacts are influenced by many factors including echo spacing and echo train length. Shorter echo spacing and echo train lengths give less time for accumulation for phase shifts and typically result in reduced distortions due to tissue susceptibility differences. Figure reprinted with permission.

2.2.1 Susceptibility Artefact Correction with Field Maps

A popular method for correcting for susceptibility artefacts is to estimate the B_0 field inhomogeneity. This is usually done through the acquisition of dual gradient echo images which provide an estimate of magnetic field map through data acquisition at two different echo times. The field map value at each voxel is used to compute the geometric shift in the phase encode direction. The physical model of susceptibility artefact based distortion in EPI was described by Jezzard and Balaban (1995) and they presented a method for correcting for these distortions using an associated field inhomogeneity map. A B_0 field inhomogeneity map, commonly called the field map, can be calculated from a map of phase evolution of each voxel in the MR image as:

$$\Delta B_0(x, y, z) = (\gamma \Delta TE)^{-1} \Delta\Theta(x, y, z) \tag{2.1}$$

where $\Delta B_0(x, y, z)$ is the field inhomogeneity at a given voxel location, $\Theta(x, y, z)$ is the angular phase evolution measured over time ΔTE and γ is the gyromagnetic ratio. The phase evolution can be extracted from the difference of two echoes, which eliminates effects that are common to both images. Hence, in eq. (2.1), $\Theta(x, y, z)$ is the phase difference between two echoes with an echo time difference of ΔTE. The one-dimensional displacement along the phase encode direction can be computed by multiplying the field map by the acquisition time as:

$$\delta_{PE}(x, y, z) = \frac{\gamma}{2\pi} \Delta B_0(x, y, z) T_{acq} \tag{2.2}$$

where $\delta_{PE}(x, y, z)$ is the one-dimensional voxel displacements in the phase encode direction and T_{acq} is the readout time for a slice of MR data.

The robustness of the field map method depends on the ability to extract reliable phase information from measured complex data. A problem is that the phase images are uniquely defined only in the range of $(-\pi, \pi]$ and hence the phase images need to be *unwrapped* at each voxel by an unknown integer

multiple of 2π to obtain the true phase as:

$$\phi_t(i) = \phi_w(i) + 2\pi k_i \tag{2.3}$$

where $\phi_t(i)$ is the true phase at a given voxel i, $\phi_w(i)$is the wrapped phase and k_i is the unknown integer multiple of 2π that needs to be estimated.

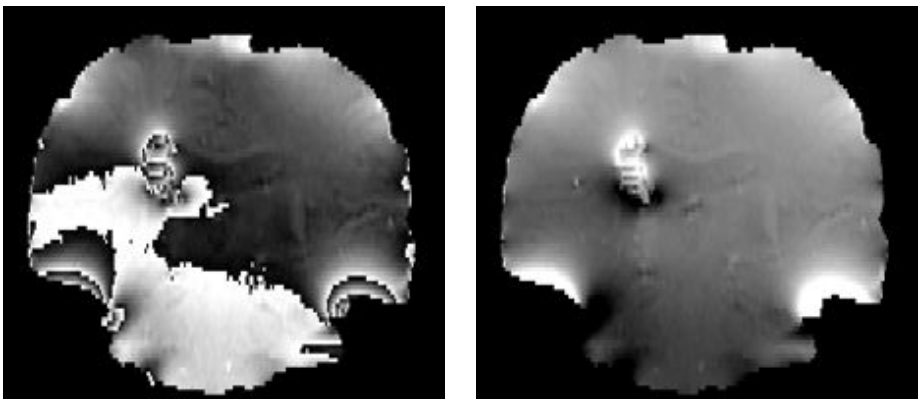

a b

Figure 2.8: Example phase unwrapping for MR images of the human head. (a) shows the phase image with multiple 2π wraps. (b) is the unwrapped phase image. The unwrapping was performed using the phase unwrapping software presented in chapter 4.

Phase unwrapping algorithms have been an active area of research since the 1980s. A majority of the early work on phase unwrapping algorithms have been designed for 2D phase data processing (An et al., 2000; Chavez et al., 2002; Liang, 1996; Moon-Ho Song et al., 1995; Ying et al., 2006; Zhou et al., 2009). Recent works have addressed the problem of unwrapping three-dimensional phase data (Abdul-Rahman et al., 2009; Cusack and Papadakis, 2002; Langley and Zhao, 2009a; Liu and Drangova, 2012; Liu et al., 2012) and Jenkinson (2003) presented a method for N-dimensional phase unwrapping. The existing phase unwrapping methods can be broadly classified into three different categories: path-following (Chavez et al., 2002; Cusack and Papadakis, 2002), cost function optimisation (Jenkinson, 2003; Moon-Ho Song et al., 1995; Nico et al., 2000; Ying et al., 2006) and parametric modelling methods (Langley and Zhao, 2009a; Liang, 1996).

The path-following phase unwrapping methods apply line integrals over a phase gradient map. In the context of one-dimensional phase unwrapping, Itoh (1982) showed that the wrapped phase gradient module 2π are the same as the corresponding true phase gradient if the latter is less than π radians everywhere (*Itoh condition*). Hence, the unwrapped phase can be obtained by integrating the wrapped phase gradient provided the Itoh condition is satisfied. However, this smoothness constraint is frequently violated in practice due to presence of noise and genuine phase discontinuities. For multidimensional phase unwrapping, the integration result not only depends in the beginning and end points but also on the chosen path of the line integral (Ghiglia et al., 1987). Most path-following methods attempt to handle this inconsistency by optimising the integration path. The branch-cut algorithm proposed by Goldstein et al. (1988) restrict the integration through the image to paths without discontinuities. These algorithms

assume that the phase discontinuities lie on the paths between the positive and negative phase residues, known as branch cuts. The phase can be unwrapped along any path that does not cross the branch cuts. Another popular variant of the path-following algorithm rely on the estimation of a quality map (Abdul-Rahman et al., 2009; Cusack and Papadakis, 2002). These algorithms aim to ascertain the noise as a function of space and unwrap the less noisy parts first. This is done to ensure that the unwrapping errors due to noise do not propagate throughout the image. These quality maps are usually derived from the wrapped phase image and include criteria such as phase derivative variance, maximum phase gradient and second phase difference (Ghiglia and Pritt, 1998). The robustness of these algorithms depend on whether reliable information about the phase noise can be extracted from the wrapped phase image. In the context of MRI phase unwrapping, the magnitude image has been used as a quality map to identify regions with high signal to noise ratio (Ying et al., 2006).

The cost-function based phase unwrapping treats the problem as a *maximum likelihood* (ML) or a *maximum a posteriori* (MAP) probability estimation problem. In this approach, phase unwrapping is formulated as an optimisation problem where a defined cost function maps unwrapped solutions to scalar costs. The optimisation routine than aims to find the unwrapping solution with the minimum associated cost. A popular automated phase-unwrapping algorithm called PRELUDE was proposed by Jenkinson (2003) and is part of the freely available FSL software package (Smith et al., 2004). PRELUDE can be used for phase unwrapping images of any dimension and has been widely used for unwrapping 3D MRI phase images. The method uses a region merging approach to optimise a cost function that penalises phase differences across boundaries between these regions. The regions are created by splitting the phase image into connected components, inside which the phase remains within a given interval. The algorithm works by iteratively merging regions until there are no more interfaces between regions. The cost function that the algorithm minimises is the sum of squared difference in phase between region interfaces.

$$C_{AB} = \sum_{j,k \in \mathcal{N}(j)} = (\phi_{Aj} - \phi_{Bk})^2$$

where C_{AB} is the cost over regions A and B. j is the index of a voxel in region A, while k is the index of a voxel in region B, such that the voxels are adjacent i.e. in the same simply connected neighbourhood: $k \in \mathcal{N}(j)$. The total cost over the whole volume is the sum of the costs over all the interfaces, which is minimised using a best-pair-first region merging approach.

Another cost-function based phase unwrapping method was proposed by Ying et al. (2006), who model the true phase function using a Conditional random field and perform inference on it by maximising the MAP probability. The work was developed for unwrapping of 2D phase images but can be extended easily to higher dimensional images. The phase is assumed to change smoothly through the image and this is encoded through the sum of square difference potential of the true phase between neighbouring voxels. The MAP configuration is found by using a dynamic programming approach (Bellman, 2003) coupled with iterated conditional mode optimisation algorithm (Besag, 1986). A quality map can be easily integrated into this algorithm to further improve its robustness by a simple modification of the

cost function to include a weighting term.

In the parametric modelling approach of (Langley and Zhao, 2009a,b), the phase map is modelled as a product of three one-dimensional Gegenbauer polynomials used as the basis functions. The unwrapped phase is modelled as $\phi(x, y.z) = Q_N(x, y, z) + r(x, y, z)$ where $r(x, y, z)$ is the residual and denotes the residual term and incorporates all expansion terms larger than N and noise. The expansion term $Q_N(x, y, z)$ is given by:

$$Q_N(x, y, z) = \sum_{n=0}^{N} \sum_{m=0}^{n} \sum_{p=0}^{m} a(n, m, p) C_{n-m}(x) C_{m-p}(y) C_p(z)$$

The expansion term is hence a linear combination of the Gegenbauer polynomials where C_a terms represent the Gegenbauer polynomials of order n and $a(n, m, p)$ are the expansion coefficients. The Gegenbauer expansion coefficients are then calculated using the gradient of the wrapped phase map. This phase modelling method can also use other complete sets of orthogonal functions as basis functions like the Legendre or Chebyshev polynomials or Fourier series with minor modifications.

Most field map based correction methods will apply phase unwrapping as a pre-processing step before the computing the field inhomogeneity as described by eq. (2.1).

2.2.2 Susceptibility Artefact Correction with Image Registration

A popular alternative to field maps is correcting for susceptibility artefacts is to use image registration techniques. Correction of susceptibility artefacts can be formulated as an image registration problem where the task is to estimate the deformation field which will bring the distorted image with a corresponding undistorted image. In this approach, EPI images are non-linearly warped to register with anatomical MR images, like the T1 or T2-weighted images. The anatomical images have a much larger spectral bandwidth and do not suffer from any significant susceptibility artefacts. The image registration process computes the deformation required to match the EPI image to the anatomical image and the resulting transformation will, in theory, provide the EPI image free of susceptibility induced distortions. A generic deformable registration based correction was first proposed by Kybic et al. (2000), which registered the baseline B0 EPI image to the undistorted T2-weighted MR image. The transformation was parameterised using B-splines and the mean squared difference between the two images was used as a similarity measure. The optimisation was done using conjugate gradient descent using analytical derivatives in a coarse-to-fine multiresolution framework. This registration algorithm did not account for the intensity distortions associated with the EPI distortions. This was addressed in the work done by Studholme et al. (2000) which added an intensity correction term, based on the Jacobian of the estimated transformation, in the registration step. This is because the change in geometry because of distortion redistributes the acquired signal over the reconstructed voxels, which is proportional to the Jacobian of the corresponding transformation. When geometric distortion occurs, there is a change in the coordinate system from intended image location (x, y, z) to displaced values (x_1, y_1, z_1), described by a transformation T. The Jacobian of the transformation is given by the following determinant:

$$J_T(x_1, y_1, z_1) = \begin{vmatrix} \frac{\partial x}{\partial x_1} & \frac{\partial x}{\partial y_1} & \frac{\partial x}{\partial z_1} \\ \frac{\partial y}{\partial x_1} & \frac{\partial y}{\partial y_1} & \frac{\partial y}{\partial z_1} \\ \frac{\partial z}{\partial x_1} & \frac{\partial z}{\partial y_1} & \frac{\partial z}{\partial z_1} \end{vmatrix}$$

Considering that the displacement due to distortion is only significant in the phase encode direction (y), then the Jacobian can be simplified as:

$$J_T(x_1, y_1, z_1) \approx \begin{vmatrix} 1 & \frac{\partial x}{\partial y_1} & 0 \\ 0 & \frac{\partial y}{\partial y_1} & 0 \\ 0 & \frac{\partial z}{\partial y_1} & 1 \end{vmatrix} \approx \frac{\partial y}{\partial y_1}$$

The registration uses the undistorted anatomical and the EPI images as the target and source image respectively and takes the intensity distortions into account by using the Jacobian corrected EPI image during the registration process. The algorithm optimises the transformation T_{AE} from the anatomical onto the EPI image. The intensity distortion are taken into account by recomputing the EPI intensities during image registration as $I_{AE} = I_E(T_{AE})J_{AE}$ where I_{AE} is the Jacobian corrected EPI image in the space of the reference anatomical image, $I_E(T_{AE})$ is the transformed EPI image where T_{AE} is the current estimate of the transformation and J_{AE} is its Jacobian. The algorithm used cubic B-splines to parameterise the deformation field and used normalised mutual information as the similarity measure, which was optimised using a simple iterative gradient descent approach. Using the Jacobian term to modify the EPI intensities as described will result in bright regions of the image being more sensitive to local changes in the transformation estimate than darker regions. To avoid this bias, the authors used the log transformation to compute the Jacobian corrected EPI image as $\hat{I}_{AE} = \log(I_E(T_{AE})) + \log(J_{AE})$. This idea was put in a variational framework by Tao et al. (2009) where they formulate the problem as a one-dimensional partial differential equation which describes the evolution of the displacement field as optimisation of EPI-structural image alignment.

Merhof et al. (2007) developed a graphics hardware accelerated version of the free form deformation algorithm which utilised the normalised mutual information as a similarity measure. The proposed method utilised simultaneous perturbation stochastic approximation (SPSA) as the optimisation routine. This optimisation method only relies on evaluation of the similarity measure and does not require the computation of the gradient of the similarity measure. The essential feature of SPSA, which provides its power and relative ease of use in difficult multivariate optimization problems, is the underlying gradient approximation that requires only two similarity measurement evaluations per iteration regardless of the dimension of the optimization problem. This also results in an algorithm which is much faster than the one utilising the classical finite difference method for approximating the gradient. However, this work does not constrain the deformation to only occur in the phase encode direction and neither did it use the intensity corrected EPI images during the registration process. Even though the work was evaluated in the context of correction of susceptibility artefacts in EPI images during neurosurgery, it should be seen as another variant on the free form deformation algorithm with a different optimiser. Furthermore, the execution time for the non-linear registration was around 50 minutes which makes it unsuitable for use

during surgical procedure. However, the computation time is likely to be reduced further on modern graphics hardware.

A registration algorithm that has received a lot of attention recently is the large deformation diffeomorphic metric mapping, commonly referred to as LDDMM (Beg et al., 2005). LDDMM provides a diffeomorphic transformation (one-to-one, invertible, smooth transformations) which preserve topology. This allows for preservation of topology even in presence of extreme distortions. Huang et al. (2008) applied LDDMM to correct for susceptibility induced deformation on 3T diffusion images. A significant drawback of this method is that it used intensity based cost functions, which cannot be used for registration of images of different modalities like EPI image with T1-weighted image, for example. Additionally, it is computationally too expensive to be used in the neurosurgical setting.

Another diffeomorphic registration algorithm to correct for susceptibility artefacts was proposed by Ruthotto et al. (2012) which formulated the problem in a variational framework. The method introduces a nonlinear regularisation functional which controls the intensity modulations and also ensures that the estimated transformation is diffeomorphic. This approach requires reversed gradient strategy i.e. acquiring two EPI volumes with inverted phase encoding gradients which results in identical images apart from their opposite directions (Andersson et al., 2003; Chang and Fitzpatrick, 1992). The registration task is then to find the transformation such that the corrected datasets are as similar as possible. Whilst this approach is highly interesting, we do not consider methods for correction of susceptibility artefacts using the reversed gradient approach in this work as the current acquisition protocol at our clinical centre is bound to the use of field maps. This is because changes in clinical protocols need to go through the ethical approval process at NHNN, which can take significantly long time to acquire.

Irfanoglu et al. (2011) combined the field map and image registration methods by first estimating a B_0 field map from an initial segmentation of a distortion free structural image and tissue susceptibility maps using the method described in Jenkinson et al. (2004). A non-uniform B-spline grid is then sampled as a function of estimated displacements. The image is densely sampled with grid knots where large distortions are expected and sparsely sampled at locations where distortions are homogeneous. This method, however, requires accurate segmentation of the undistorted T1 image and knowledge of tissue susceptibility values, which can be quite difficult to specify accurately around the resection area.

2.3 Discussion

Large efforts have been devoted towards accurately estimating brain shift during neurosurgery and it remains an active area of research. As interventional imaging capabilities continue to grow, the need for time-efficient multimodal image analysis becomes paramount. My work was motivated towards exploiting the structural and diffusion imaging capabilities of the current state of the art iMRI scanners to estimate the deformations during neurosurgery. As we are interested in localising white matter tracts like the Optic Radiation during the intervention, I propose utilising information from both structural and diffusion MR images to perform the image registration. To achieve this, I propose a modified normalised mutual information based similarity measure which combines the information from these two MR modalities in a principled manner. I also developed a GPU accelerated implementation that can

be used well-within the time constraints of a typical neurosurgical procedure. This is described in detail in chapter (5).

Presence of geometric distortions in diffusion MRI images hinder their effective use for neurosurgical guidance. Even though there is some body of work around correcting for susceptibility artefacts in EPI images, it largely remains an ignored problem in the context of neurosurgery. Hence, the susceptibility artefact correction literature does not focus on its application in the neurosurgical scenario. My work in this topic aims to combine the field map and image registration methods using a principled approach. The proposed method computes the B_0 field inhomogeneity map as well as the uncertainty associated with the estimated solution. Image registration is then used to further refine the results in regions of high uncertainty. The algorithm is fast enough to be used for neurosurgical guidance due to use of efficient graph based inference technique. The proposed method is described in detail in chapter (4).

In chapter 3, I will describe the theory behind graph cuts, which is the main computational technique behind my work on correction of susceptibility artefacts. Additionally, I will also describe the relevant components of the medical image registration algorithm which is used in the proposed method for estimation of brain shift.

Chapter 3

Overview of Computational Techniques

To address the challenges described in chapter 1, use of medical image registration and techniques based on discrete optimisation are proposed in chapters 4 and 5. This chapter will lay the theoretical groundwork for these techniques which are extensively used in the subsequent chapters of this thesis.

3.1 Discrete Optimisation: Graph Cuts

The correction of susceptibility artefacts in iMRI images has been formulated as a discrete energy minimisation problem in this thesis. Discrete optimisation techniques, in particular Graph Cuts, have received a lot of attention recently from the computer vision community particularly due to fast computation time. This section provides an overview of Graph Cuts based optimisation techniques and lays the theoretical groundwork for chapter 4. I will give an introduction to probabilistic graphical models and describe how graph cuts can be used to perform inference on classes of undirected graphical models.

3.1.1 Energy Minimisation

Many problems in medical imaging can be formulated as finding the most probable values of some hidden or unobserved variables, which can take on either discrete or continuous values. For discrete variables, these problems are referred to as *labelling* problems as the solution involves assigning the most probable label to the hidden variables. Labelling problems are ubiquitous in medical image analysis especially in the area of medical image segmentation, registration and artefact correction. These problems can be naturally formulated in terms of energy minimisation where a labelling configuration is sought that minimises some energy function.

Graph cuts have emerged as an efficient framework for solving such discrete labelling problems. In particular, efficient graph cuts based minimisation algorithms have been extremely successful for inferring the *maximum a posteriori* (MAP) solutions of Markov and Conditional Random Fields which are extensively used to model a wide range of problems in medical image analysis. These random fields belong to the class of probabilistic graphical models, which use graphs to encode the conditional dependence between the random variables. The following sections describe the background theory behind these random fields and how graph cuts can be used to do inference on them.

3.1.2 Probabilistic Graphical Models and Random Fields

Probabilistic Graphical Models (PGMs) represent the marriage between graph theory and probability theory. They are tools for dealing with two common problems that arise throughout image processing: complexity and uncertainty modelling. PGMs use a graphical structure to represent the probability distributions of random variables. Each random variable corresponds to a node in the graph and the links between the nodes encode the statistical dependence between the variables. Given a PGM, the joint or conditional probability distributions can be decomposed as a product of functions defined on the subset of the random variables. This can greatly simplify the modelling of multivariate joint distributions and efficient algorithms exist that can exploit the graph structure to compute the marginal or conditional probabilities of interest.

A PGM is represented by a graph $\mathcal{G} = \langle \mathcal{V}, \mathcal{E} \rangle$ where $\mathcal{V}$ is the set of graph nodes and $\mathcal{E}$ denotes the set of edges between the graph nodes. The nodes correspond to the random variables and the edges describe the probabilistic relationships between the random variables. In this thesis, we will only be dealing with *undirected* PGMs in which the edges $\mathcal{E}$ consist of undirected links only. Markov Random Fields (MRFs) and Conditional Random Fields (CRFs) are two closely related undirected graphical models that have been widely used for many image analysis problems including, but not limited to, image denoising, restoration and segmentation.

A Markov Random Field (MRF) is a *generative* model that models the joint probability distribution of the unknown labels X and the observations Y. MRFs model the interactions among a set of random variables through the local interactions within a selected neighbourhood system. Represented as an undirected graph, a node X_j is a neighbour of the node X_i if and only if they share an edge. The Markov property in MRFs implies that a variable X_i is conditionally independent of all other variables given its neighbours; $P(X_i|X \backslash X_i) = P(X_i|X_{N_i})$ where N_i is the set of neighbours of a random variable X_i.

As established by the Hammersley-Clifford theorem (Besag, 1974; Moussouris, 1974), the joint distribution modelled by the MRF can be specified as a *Gibbs* distribution:

$$P(X, Y) = \frac{1}{Z} \exp(-\sum_{c \in \mathcal{C}} D_c(X_c, Y_c)) \tag{3.1}$$

where $P(X, Y)$ is the joint probability distribution, $\mathcal{C}$ is the set of cliques i.e. sub-graphs in which each random variable is a neighbour of all other random variables, $D_c(X_c, Y_c)$ is an energy function that is defined on a given clique c, Z is the *partition function* and is calculated by marginalisation over all the random variables in the MRF. This is needed to ensure that $P(X, Y)$ is a valid probability distribution.

The size of the clique has a major influence on the computational complexity in an MRF model. A vast number of image processing problems have been formulated using a *first-order* MRF, where the maximum clique size is 2. In such a first-order MRF, the joint probability distribution can be written as:

$$P(X, Y) = \frac{1}{Z} \prod_{i \in \mathcal{V}} \phi(X_i, Y_i) \prod_{i,j \in \mathcal{V}, j \in N_i} \psi(X_i, X_j) \tag{3.2}$$

where $\phi(X_i, Y_i)$ is the unary potential since it is associated with only one label variable X_i. Similarly, $\psi(X_i, X_j)$ is the pairwise potential term and is defined on the clique neighbourhood. The unary

potential is typically called the *data term* and it measures how much assigning a label to an MRF node i disagrees with the observed data. The pairwise potential is usually called the *prior term* as it encodes our prior belief on labelling configurations of neighbouring MRF nodes. This typically encourages fewer label changes between neighbouring nodes and for that reason is also usually called the *smoothing term*. Figure. (3.1)(a) shows an MRF example for an image labelling problem. Given the observations $Y(y_1, y_2...y_n)$, the MRF model can be used to infer the labels X by *maximum a posteriori* (MAP) inference.

Often we would like to use pairwise potentials that are data dependent. For example, if two neighbouring voxels differ greatly in their intensities, then they are likely to belong to different classes. Hence, labelling configurations that assign them different labels should not be heavily penalised during the optimisation. This can be achieved by defining the pairwise potential $\psi(X_i, X_j)$ in such a way that it also depends on the observed data Y. A *Conditional Random Field* (CRF) is a probabilistic model which allows the use of data dependent potential functions.

In contrast to an MRF, a CRF is a *discriminative* model that directly models the posterior probability distribution of a set of random variables X, given the data Y. CRFs were first introduced in the domain of natural language processing (Lafferty et al., 2001) and have been widely used in image processing tasks (Sutton and Mccallum, 2007). Compared to the MRF model, CRFs do not model the joint probability distribution and focuses directly on the labelling problem given the observations. In addition, it also naturally considers the discontinuity of the labels since the interactions between the labels can be automatically adjusted by the observations. A CRF can be viewed as a special case of an MRF, where the MRF is globally conditioned on the data.

Similarly to the MRF model for joint distributions, the CRF model assumes that the posterior probability distribution of a set of random variables X follows the Markov property and can be represented as a Gibbs distribution and can be decomposed into a product of potential functions:

$$P(X|Y) = \frac{1}{Z(Y)} \exp(-\sum_{c \in \mathcal{C}} E_c(X_c|Y)) \tag{3.3}$$

where $E_c(X_c|Y)$ is an energy function defined on the set of random variables X_c in the clique c, conditioned on the observations Y. Now the energy function depends on the observation Y. The partition function Z is also a function of the observations Y and can be calculated by marginalising over X. Figure. (3.1)(b) shows a CRF example for an image labelling problem. In contrast to the MRF, the potential functions are related to all of the observations Y.

3.1.3 Inference on MRFs: *Maximum a Posteriori* Estimation

The most popular way to estimate an MRF is through the *maximum a posteriori* (MAP) estimation. The MAP-MRF approach was introduced by Geman and Geman (1984) in the context of image analysis. They were the first to make an analogy between image analysis and statistical mechanics systems. A typical scenario is when we wish to estimate the unobserved MRF configuration on the basis of some observed data. Then, the MAP labelling $x^\star$ of a random field is defined as equation (3.4).

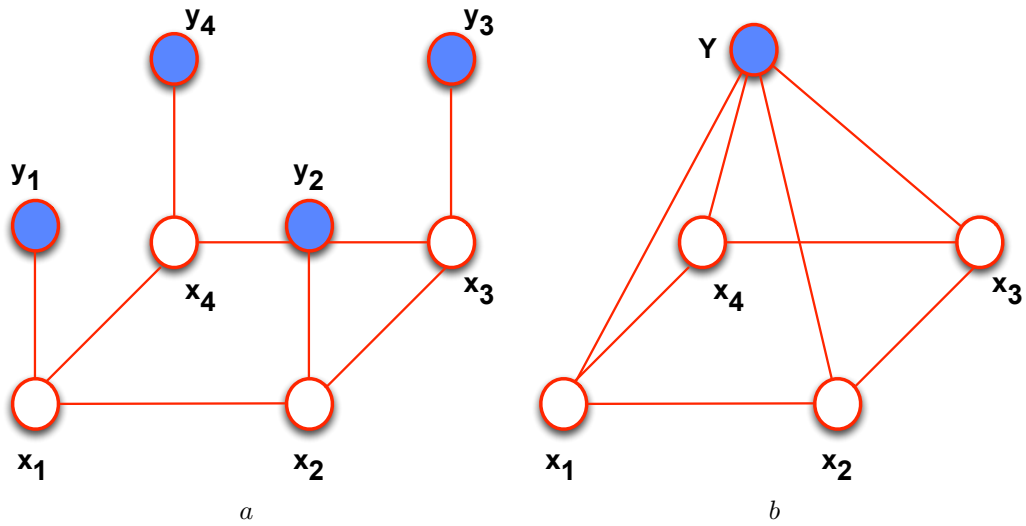

Figure 3.1: (a) A simple MRF model for image labelling. MRF is a generative model that models the joint distribution $P(X, Y)$ of the output labels of random variables X and observations Y. (b) A similar CRF model for image labelling. In contrast to the MRF model, it is a discriminative model and models the conditional posterior probability $P(X|Y)$ directly. The unary term in a CRF at a node i is a function of all of the observation data Y and the label x_i rather than just y_i and x_i only as is the case for the MRF model. In the MRF model the pairwise potentials are independent of the observations. However, in the CRF model they are also a function of the observations which allows us to include data dependent pairwise potentials.

$$x^{\star} = \underset{x \in X}{\operatorname{argmax}} P(x|Y) \tag{3.4}$$

This can be achieved by minimising the corresponding log-transfomed Gibbs energy function whose form is given by equation (3.5). Note that in the optimisation, the term involving the partition function can be treated as a constant and does not need to be considered. $E_c(x_c)$ is the potential function defined over the clique as described before for the MRF and CRF models.

$$E(x) = -\log P(x|Y) - \log Z = \sum_{c \in \mathcal{C}} E_c(x_c) \tag{3.5}$$

3.1.4 Energy Minimisation via Graph Cuts

Graph cuts have emerged as a popular framework for computing the MAP solutions for various discrete labelling problems in computer vision and have recently received much attention in the medical image analysis community. Graph cuts have achieved popularity because efficient algorithms are available for computing inference on graphs of arbitrary topology. In many cases, globally optimal solutions can be found for important classes of energy functions commonly encountered in many medical image analysis problems. Even for energy functions where the global optimal solutions cannot be guaranteed, graph cuts can usually be used to find strong local minima of the energy function. This section will introduce the basic concepts and notation associated with it.

Let $\mathcal{G} = \langle \mathcal{V}, \mathcal{E} \rangle$ be a weighted graph where $\mathcal{V}$ is the set of vertices and $\mathcal{E}$ is the set of edges. $\mathcal{V}$ has two special vertices called the terminals. Traditionally, one of them is called the *source* and the other one is called the *sink* and they are usually represented by letters **s** and **t** respectively. A *cut* $\mathcal{C} \subset \mathcal{E}$ is a

set of edges such that the terminals are separated in the induced graph. Additionally, no proper subset of $\mathcal{C}$ separates the terminals in $\mathcal{G}$. An example cut is illustrated in figure (3.2). The cost of the cut equals the sum of its edge weights. The *minimum cut* problem is to find the cut with the smallest cost. It was shown (Ford and Fulkerson, 1962) that the cut with the minimum cost can be found by computing the maximum flow between the terminal vertices in the graph.

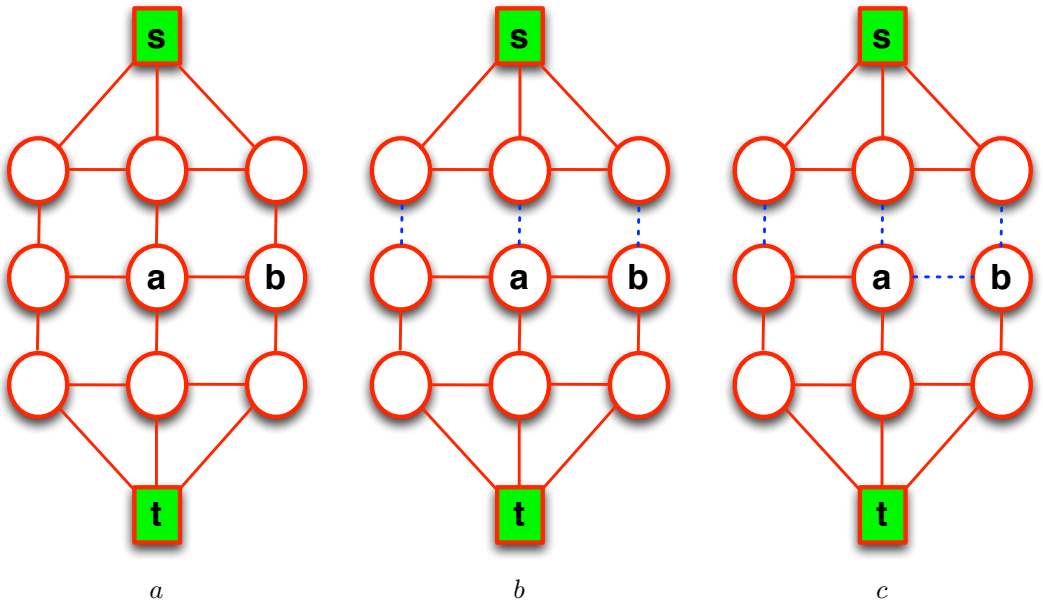

a b c

Figure 3.2: A simple graph configuration on a regular grid. The squares denote the source and the sink vertices. The circles denote the other vertices. The red lines show the edges between the nodes. The dashed blue lines show edges which form a cut. (a) shows the original graph. (b) shows a graph with a valid cut. The dashed blue lines separate the graph into two sub-graphs which separate the source and the sink vertices. (c) shows a graph with an invalid cut. The cut is not valid because, if the dashed edges between nodes a and b are removed, the remaining dashed edges still form a cut.

3.1.5 Submodular Functions

Minimising an arbitrary energy function is NP-hard in general (Kolmogorov and Zabin, 2004). However, there exist families of energy functions for which the minima can be found in polynomial time. *Submodular* set functions constitute one such family of functions which have been extensively studied. Many optimisation problems relating to submodular functions can be solved efficiently. In some respects they are similar to convex/concave functions encountered in continuous optimisation. In their seminal paper Kolmogorov and Zabin (2004) showed submodularity to be a necessary and sufficient condition for a function to be representable by a graph.

A function $E(x_1, x_2)$ of two binary variables $\{x_1, x_2\}$ is submodular if and only if the inequality of equation (3.6) is satisfied.

$$E(0,0) + E(1,1) \leq E(0,1) + E(1,0) \tag{3.6}$$

It was shown by Kolmogorov and Zabin (2004) that an energy function of n binary variables with the form of equation (3.7), where E_i is the unary energy term and E_{ij} is the pair-wise energy term, can

be represented by a graph as long as each pair-wise term E_{ij} is submodular i.e. it satisfies the inequality in equation (3.6).

$$E(x_1, x_2, ...x_n) = \sum_{i=1}^{n} E_i(x_i) + \sum_{i=1; i<j}^{n} E_{ij}(x_i, x_j) \tag{3.7}$$

3.1.6 Graph Construction for Submodular Functions

Kolmogorov and Zabin (2004) showed how to construct graphs for functions of the form of equation (3.7). All edges in the graph are assigned some weight or cost. There are two types of edges in a graph: *n*-links and *t*-links. *n*-links connect pairs of neighbouring voxels and they represent the neighbourhood interaction in a random field. Hence, the *n*-links correspond to the prior and the smoothness term and are usually used to encode the penalty term for label discontinuities between voxels. This represents the $E_{ij}(x_i, x_j)$ term in equation (3.7). The *t*-links connect voxels with the terminals (source/sink). The cost of a *t*-link connecting a voxel and a terminal corresponds to a penalty for assigning the corresponding label to the voxel. This cost is normally derived from the data term $E_i(x_i)$ in equation (3.6).

Consider a binary labelling problem i.e. the set of labels is binary: $X = \{0, 1\}$. Considering the energy function of the form as in equation (3.7), the unary function depends only on the variable x_i. Let us define the terminal edge weights as follows:

$$w_{\mathbf{s},i} = E_i(1), \qquad w_{i,\mathbf{t}} = E_i(0) \tag{3.8}$$

where $w_{s,i}$ is the weight of the edge from the source vertex to graph vertex i and $w_{i,\mathbf{t}}$ is the weight of the edge from i to the sink vertex. In the case, $x_i \in \mathbf{s}$ and takes the label 0, the edge $e_{i\to\mathbf{t}}$ is in the cut. Hence, by equation (3.8)(b), $w_{i,\mathbf{t}} = E_i(0)$ is added to the cost of the cut. Similarly, when $x_i \in \mathbf{t}$, $w_{\mathbf{s},i} = E_i(1)$ is added to the cost of the cut. To ensure non-negative weights for the unary terms, if $E_i(0) < E_i(1)$ then we add the edge $e_{\mathbf{s}\to i}$ with the weight $E_i(1) - E_i(0)$. Otherwise, the edge $e_{i\to\mathbf{t}}$ is added with the weight $E_i(0) - E_i(1)$. This is demonstrated graphically in figure (3.3).

Focussing on the pairwise smoothing term, let us define the following edge weights. This is also shown graphically in figure (3.4)(a)

$$w_{j,\mathbf{t}} = a, \qquad w_{i,i} = b, \qquad w_{\mathbf{s},i} = c \tag{3.9}$$

where i, j are nodes which are mutual neighbours. Then the cut costs for the four possible combinations of (x_i, x_j) are the following. This is also reflected in figure (3.4)(b-e)

$$E_{i,j}(0,0) = a, \qquad E_{i,j}(0,1) = b, \qquad E_{i,j}(1,0) = c + a, \qquad E_{i,j}(1,1) = c, \tag{3.10}$$

Dropping subscripts for brevity, we have equation (3.11) after simple algebraic manipulation. Since the right hand side is an edge weight, it must be non-negative for polynomial time algorithms to be applicable. Note that this constraint gives rise to the submodularity inequality condition of equation (3.6).

$$E(0,1) + E(1,0) - E(0,0) - E(1,1) = b \tag{3.11}$$

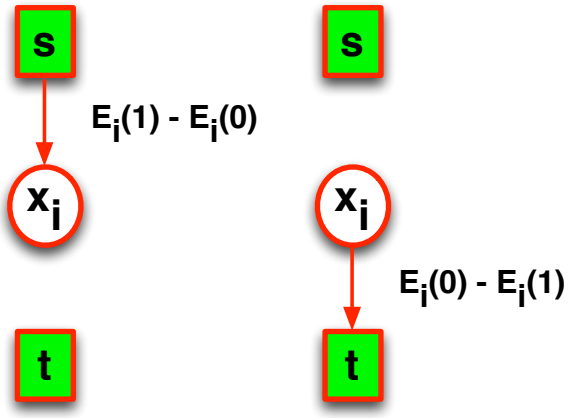

a b

Figure 3.3: Edge definitions and weights for the unary terms. (a) Graph for E_i when $E_i(1) > E_i(0)$. (b) Graph for E_i when $E_i(1) \leq E_i(0)$. The unary terms can be arbitrary as one of the terminal edges for a graph vertex is always in the cut. Hence, adding the same constant weight to both $w_{\mathbf{s},i}$ and $w_{i,\mathbf{t}}$ does not affect the choice of which edge to cut.

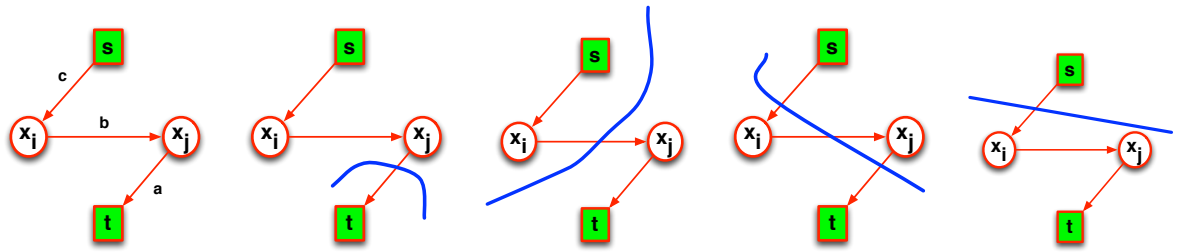

(X_i, X_j)	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
Cost	a	b	(c+a)	c

a b c d e

Figure 3.4: Representation of the smoothing term of the energy function of equation (3.7). (a) is the graphical representation of the edge weights as defined in equation (3.9). (b-e) shows the cut boundary and the cut cost when the vertices take the different label configurations.

Assuming the submodularity constraint is met, we can rewrite the cut costs for the four possible combinations as follows:

$$(x_i, x_j) = (0,0) = E(1,0) - E(1,1)$$
$$(x_i, x_j) = (0,1) = E(0,1) + E(1,0) - E(0,0) - E(1,1)$$
$$(x_i, x_j) = (1,0) = E(1,0) - E(1,1) + E(1,0) - E(0,0)$$
$$(x_i, x_j) = (1,1) = E(1,0) - E(0,0)$$

Note that if we add $E(0,0) + E(1,1) - E(1,0)$ to each of the four cost values, we see that each equals to $E_{i,j}(x_i, x_j)$. Since we are adding the same value to the four possible outcomes, it does not affect their relative costs and the minimum cut still corresponds to the labelling with the minimum energy. Combining the costs as described in equation (3.8) and equation (3.12), we have to add the following

edge weights for each neighbouring pair $(i, j) \in \mathcal{N}$ in our graph:

$$w_{\mathbf{s},i} = E_i(1) + \sum_{i,j \in N} E_{i,j}(1,0) - E_{i,j}(0,0)$$
$$w_{j,\mathbf{t}} = E_i(0) + \sum_{i,j \in N} E_{i,j}(1,0) - E_{i,j}(1,1)$$
$$w_{i,j} = E_{i,j}(0,1) + E_{i,j}(1,0) - E_{i,j}(0,0) - E_{i,j}(1,1)$$

If any of the edge weights $w_{\mathbf{s},i}$ or $w_{j,\mathbf{t}}$ is negative, it can be made non-negative by adding a constant weight to them as explained before. The minimum cut on such a graph can be found by pushing the maximum flow from the source to the sink vertices (Ford and Fulkerson, 1962).

3.1.7 Multi-label Optimisation with Graph Cuts

Graph Cuts can also be used to exactly optimise convex energy functions which involve variables taking more than two labels (Ishikawa, 2003). The graph creation proposed in Ishikawa (2003), the label of a discrete random variable is found by observing which data edge is cut. This construction is valid for a restricted class of energy functions (convex priors) and do not include energies with non-convex priors, like the Potts model (Potts, 1952). In addition, for problems with large label sets, this method is extremely memory intensive and impractical.

A popular alternative is to break the multi-way cut into a series of binary s-t cut problems. In such cases, the global optimum cannot be usually guaranteed. However, graph cuts can be used to find a solution which is a strong local mimima of the energy function (Boykov et al., 1998). These solutions for certain problems are shown to be better than the ones obtained by other methods (Boykov and Kolmogorov, 2004). Boykov *et al.* proposed two algorithms that rely on an initial labelling and an iterative application of binary graph cuts. At each iteration, an optimal *range move* is performed to either expand (α-expansion algorithm) or swap labels ($\alpha - \beta$ swap algorithm) (Boykov et al., 1998, 2001). Although convergence and error bounds are guaranteed, the initial labelling may influence the result of the algorithm. Also, it is important to note that the $\alpha - \beta$ swap algorithm can only be applied when the smoothness term is a semi-metric i.e.

$$E(\alpha, \beta) = 0 \qquad \Longleftrightarrow \ \alpha = \beta \tag{3.14a}$$
$$E(\alpha, \beta) = E(\beta, \alpha) \geq 0 \tag{3.14b}$$

The α-expansion algorithm is even more restrictive and can only be applied when the smoothness term is a metric i.e. in addition to the semi-metric conditions, the following triangle inequality must also apply.

$$E(\alpha, \beta) \geq E(\alpha, \gamma) + E(\gamma, \beta) \tag{3.15}$$

A new set of multi-label algorithms were proposed in Veksler (2009) that act on a larger set of labels than those in (Boykov et al., 2001). More recent approaches based on linear programming re-

laxation using primal-dual (Komodakis et al., 2008), message passing (Kolmogorov, 2006) and partial optimality (Kohli et al., 2008) have been proposed.

For multi-label energy minimisation, this thesis use of the α-expansion algorithm for energy minimisation in multi-label CRFs. The α-expansion algorithm belongs to the class of *move-making* algorithms, which operate by making a series of changes (also called moves) to the solution such that these changes do not lead to an increase in the solution energy. In each iteration, the algorithm searches for a lower energy solution in a pre-defined neighbourhood (also called the move space) around the current solution. It is important to highlight the distinction between *moves* and *move-spaces*. In the α-expansion algorithm we have $|X|$ possible *move-spaces* (one for every possible label in the label set X). However, we have 2^n possible *moves* within each move-space (one corresponding to each node taking on a binary value). An α-expansion *move* (where $\alpha \in X$) finds the minimum energy move within the move-space α.

The main-idea behind α-expansion algorithm is to successively segment all the nodes taking the label α from the non-α nodes where the label α is changed at every iteration. The algorithm iterates through every possible α value till the algorithm converges. The optimal α-expansion move can be performed at every iteration in polynomial time as long as the pairwise energy terms form a metric.

Chapter 4 makes extensive use of graph cuts to perform unwrapping of the MRI phase images to compute the magnetic field inhomogeneity maps. The field map based correction is used in conjunction with an image registration step, which is also formulated using graph cuts. Graph cuts provide a very fast algorithms to tackle these problems within the time constraints of a neurosurgical procedure.

The following section provides an overview of medical image registration. Medical image registration is one of the most popular ways to estimate the brain shift between the preoperative and intraoperative images. It has also been used extensively for correction of susceptibility artefacts. I will provide an introduction to image registration and explain the main components of the image registration algorithms used in this thesis in chapters 4 and 5. Image registration is an extremely active area of research in medical image processing. This section will only provide an overview of areas that are relevant in the context of the works presented in subsequent chapters.

3.2 Medical Image Registration

A broad overview of medical image registration was given in chapter 2. In the following sections, I will describe the various components of the image registration algorithm which are relevant to this thesis.

3.2.1 Parametric Non-rigid Transformation Models: Cubic B-splines

As described, accurate brain shift estimation requires the use non-rigid image registration which typically has a very high number of degrees of freedom. A popular approach is to use a parametric transformation model where the estimated deformation field is parameterised using another function. The number of degrees of freedom is usually lower than the number of voxels in the image when using a parametric transformation model. Typically, smooth transformation models are used to promote physically realistic deformations. However, regularisation still need to be employed to ensure plausible deformations. The

image registration methods described in this thesis parameterise the deformation field using cubic B-splines as used in the popular free form deformation (FFD) algorithm (Modat et al., 2010; Rueckert et al., 1999).

Cubic B-splines have the desirable property of generating deformations that are C^2 continuous i.e. their first and second derivatives are continuous. The basic idea is that a uniformly spaced cubic spline control point mesh is overlaid on the image. The spline control points control the position of certain voxels in their neighbourhood. So, by perturbing the control points, local deformations can be induced in the image. In one dimension, the new position of a point $(\vec{x})$ is given by:

$$\mathbf{T}(\vec{x}) = \sum_{l=0}^{3} B_l \left(\frac{\vec{x}}{\delta} - \lfloor \frac{\vec{x}}{\delta} \rfloor \right) \mu_{i+l} \tag{3.16}$$

where μ_i are the control points taken into account to compute the new position and δ is the spacing between the control points. To compute displacement in one dimension, 4 neighbouring control points are used, two before the indexed point and two after the indexed point. The functions B_0 to B_3 are the approximated third-order spline basis functions given by:

$$B_0(u) = \frac{(1-u)^3}{6}$$
$$B_1(u) = \frac{3u^3 - 6u^2 + 4}{6}$$
$$B_2(u) = \frac{-3u^3 + 3u^2 + 3u + 4}{6}$$
$$B_3(u) = \frac{u^3}{6}$$

In three dimensions, the new position of a point can be computed by 3-D tensor product of the one-dimensional cubic B-splines as:

$$\mathbf{T}(\vec{x}) = \sum_{l=0}^{3} \sum_{m=0}^{3} \sum_{n=0}^{3} B_l(u) B_m(v) B_n(w) \mu_{i+l,j+m,k+n} \tag{3.18}$$

where u, v, and w are the relative position of the index point along each of the axes. i, j and k are the indices of the first control points that influence the indexed point position along each of the axes. As evident from equation (3.18), the location of a point is influenced by the grid of $4 \times 4 \times 4 = 64$ surrounding control points in 3-D. This local influence of the control points is what makes cubic B-splines a very popular option to model local deformations.

Despite the parametric nature of the transformation and the C^2 continuity of the B-spline transformation model, prior information on the deformation field needs to be incorporated into the registration process to promote realistic, topology-preserving deformation. The regularisation used is based in the bending energy of the spline (equation 3.19) and is composed of the second-order derivatives of the B-spline deformation which can be computed analytically from the B-spline basis functions due to the C^2 continuity of the transformation model. Bending energy was first used in a non-rigid registration algorithm by Rueckert et al. (1999) and has the advantage of being zero for affine transformations and hence only penalises the non-affine component of the transformation.

$$\mathcal{BE} = \frac{1}{N}\sum_{\vec{x}\in\Omega}\left(\frac{\partial^2\mathbf{T}(\vec{x})}{\partial x^2}\right)^2 + \left(\frac{\partial^2\mathbf{T}(\vec{x})}{\partial y^2}\right)^2 + \left(\frac{\partial^2\mathbf{T}(\vec{x})}{\partial z^2}\right)^2 + 2\times\left[\left(\frac{\partial^2\mathbf{T}(\vec{x})}{\partial xy}\right)^2 + \left(\frac{\partial^2\mathbf{T}(\vec{x})}{\partial yz}\right)^2 + \left(\frac{\partial^2\mathbf{T}(\vec{x})}{\partial xz}\right)^2\right], \tag{3.19}$$

3.2.2 Similarity Measure: (Normalised) Mutual Information

As described before, the similarity measure is used to assess the quality of warping between the target and source images. The similarity measure, in other words, describes how similar two images are to each other after a geometric transformation. This thesis uses global and local variants of the popular mutual information (MI) as a measure of similarity during image registration. The key advantage of MI and its variants is their ability to easily handle complex relationships between the intensities in the two images. They requires no a-priori model of the relationship between the image intensities and can handle image registration between different modalities.

Information theoretic approaches for registration of medical images were introduced by Maes et al. (1997); Viola and Wells (1995) when both these groups used MI as the similarity measure. MI is a concept from information theory that measures the amount of information one image has about the other. Before introducing MI, it is important to understand the concept of *entropy*.

Entropy is the *measure of information* and the marginal and joint entropies can be defined as Shannon's entropy (Shannon, 1948) as:

$$H(A) = -\sum_i p(i)\log\, p(i)$$
$$H(A,B) = -\sum_{i,j} p_{i,j}\log\, p(i,j)$$

Shannon's entropy of an image can be estimated by computing the probability distribution of the image intensities. This can be estimated, for example, by computing a histogram of image intensities i.e. counting the number of times each grey value occurs in the image and dividing those numbers by the total number of occurrences to generate a valid probability distribution. So, an image consisting of only one intensity will correspond to having low entropy as it contains little information. On the other hand, an image with equal number of many different intensity values will have high entropy as it contains a lot of information. In terms of the shape of the probability distribution, the former corresponds with a single peaked distribution with no dispersion and the latter corresponds to a widely dispersed distribution. This can be extended to a pair of images where the joint entropy can be similarly estimated where we compute the joint probability distribution i.e. find the probability of a pair of image intensities to occur together.

Image registration can be thought of as minimising the joint entropy between two images (Studholme et al., 1995). It can be intuitively visualised as when the images are misaligned, there is no overlap between corresponding structures. Hence, the resulting joint histogram is dispersed as corresponding image intensities do not overlap. As the image comes into alignment, corresponding intensities between the two images overlap and the joint histogram becomes less dispersed. This is visualised in figure (3.5). An alternative way to think about it is in terms of uncertainty between image

intensities. When the images are not aligned, one is more uncertain about the corresponding intensities in the two images as the joint histogram is dispersed. However, this uncertainty decreases as the images come into alignment and we get a sharper and less dispersed joint histogram.

From the concept of entropy, given two images A and B, MI can be defined as:

$$MI(A, B) = H(A) + H(B) - H(A, B) \tag{3.21}$$

where $H(.)$ is the marginal *entropy* for a given image and $H(A, B)$ is their joint entropy. The expression for MI contains the term $-H(A, B)$ which implies that maximising the mutual information between two images is related to minimising the joint entropy between them. Using joint entropy alone as the similarity measure in image registration suffers from the problem that it is possible to reduce joint entropy by decreasing the information content in either image. Hence, reducing the amount of overlap between two images will decrease the joint entropy at the cost of increasing the misalignment between the images. Hence, if the misalignment between the two images is so large that they only overlap in the background areas of the image, joint entropy will be quite low. MI tries to avoid this problem by including the marginal entropies of the two images $H(A)$ and $H(B)$ in the similarity measure. When there is little overlap or overlap only between non-anatomical regions, the marginal entropy terms will also be low. Hence, they act as a penalising term and discourage transformations that decrease the information between overlapping regions of the two images.

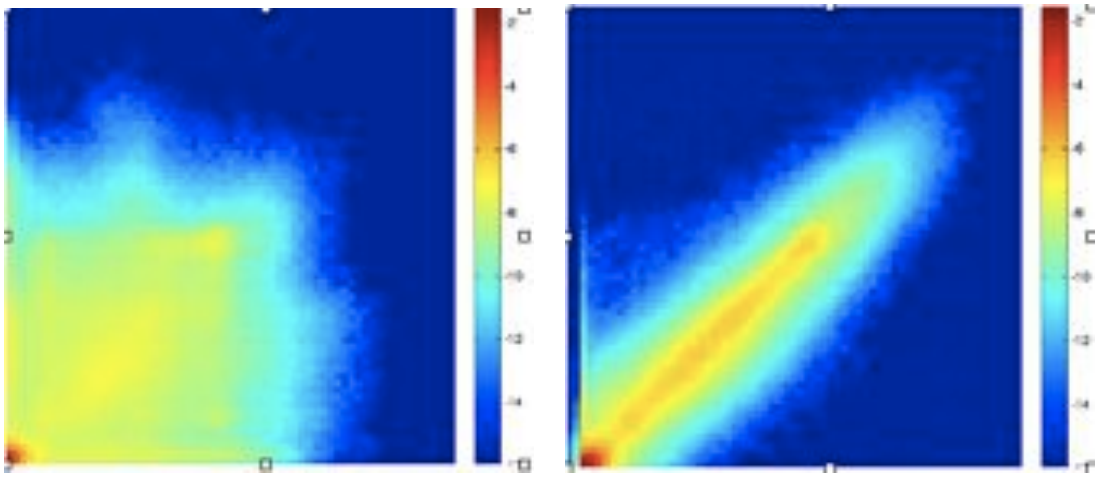

a b

Figure 3.5: Effect of registration on dispersion of the joint histogram. The image intensities have been normalised between 0 and 63 (a) shows the joint histogram when the images are not aligned. Corresponding image intensities do not overlap resulting in a more dispersed joint histogram. (b) shows the decrease in the dispersion of the joint histogram as the images come into alignment. Registration brings the corresponding structures into alignment and there is more overlap between corresponding intensities. This is an example when the image intensities have a linear relationship and come from the same modality. A more complex multi-modal histogram might result when performing registration between different modalities but the same principle about reduction of joint entropy applies.

An alternative expression for MI that encapsulates the concept that registration results in decrease of uncertainty between the two images can be written as:

$MI(A,B) = H(A) - H(A|B)$ or alternatively,

$MI(A,B) = H(B) - H(B|A)$

This expression of MI is equivalent to equation (3.21) but offers a different perspective. $H(B|A)$ is *conditional entropy* of image B given the image A. From the perspective of entropy as a measure of uncertainty, this expression of MI tells us that given two images A and B how much does the uncertainty about one image decreases *given* the other image. Maximising MI is equivalent to minimising this uncertainty as denoted by the $H(B|A)$ or $H(A|B)$ term.

Even though MI attempts to alleviate the problems related to using joint entropy alone in image registration by including the marginal entropies in its measure, it was shown by Studholme et al. (1999) that MI is not invariant to change in size of background regions. To overcome this problem, Studholme et al. (1999) proposed the *normalised mutual information* (NMI) which was empirically shown to be less sensitive to overlap size. NMI is defined as:

$$NMI(A,B) = \frac{H(A) + H(B)}{H(A,B)} \tag{3.23}$$

The registration algorithms described in chapters 4 and 5 use variants of mutual information as the similarity measure.

3.2.3 Optimisation: Conjugate Gradient Descent

Chapter 5 used a conjugate gradient ascent to find the optimal transformation between the target and source images. This approach is more efficient than a simpler steepest ascent optimisation, and is less memory intensive than Newton type algorithms. Moreover, it has the advantage to be parallel-friendly which makes it attractive for use in neurosurgical scenarios. Gradient descent based optimisation techniques require the computation of the gradient of the cost function. In chapter 5, these gradients are computed analytically which results in a significant speed-up of the registration algorithm.

This chapter introduced the computational techniques that underpin the algorithms that I have developed during my PhD. Image registration is a critical pre-processing step in most image analysis tasks. As described in chapter 2, a lot of research effort has been devoted to developing image registration algorithms. Although non-rigid image registration algorithms that use information from multiple imaging modalities exist, they are computationally expensive. In chapter 5, I propose a non-rigid registration algorithm that can register images from structural and diffusion MRI within the time constraints of a neurosurgical procedure. I propose an extension to the normalised mutual information similarity measure which allows for using multiple images of different modalities in the registration algorithm. The computational burden is overcome through the refactoring of the original free form deformation algorithm and employing GPUs to perform parallel processing.

Graph cuts have become very popular in the computer vision community for performing MAP optimisation. One of the main reasons for their widespread adoption is the reasonable low computational cost and strong guarantees on the solution obtained through their use. In the following chapter, I will apply Graph cuts to perform susceptibility artefact correction on diffusion images acquired during neu-

rosurgical intervention. I will show how we can use Graph cuts to not only obtain the MAP solution of our model parameters but also how we can get an estimate of the uncertainty associated with these parameters. This uncertainty information is then used to further improve our estimate of the distortion due to the susceptibility artefacts.

Chapter 4

Susceptibility Artefact Correction

Echo Planar Imaging (EPI) is the de-facto MRI imaging protocol of choice for diffusion weighted imaging (DWI) sequences due to its rapid acquisition time. The recent improvements in iMRI technology have made the current commercial iMRI scanners capable of performing diffusion imaging which allows for imaging of critical white matter tracts along with the surgical target areas. However, as described before, EPI images are prone to various imaging artefacts including those arising due to main magnetic field inhomogeneities. In the context of neurosurgery, this leads to severe geometric and intensity distortions around the resected brain area. I have shown that diffusion weighted MRI images along with structural images can increase the localisation accuracy of brain structures during neurosurgical procedures (Daga et al., 2012; Winston et al., 2011). There is also an interest in performing tractography on interventional DWI images to segment white matter structures of interest (Andrea et al., 2012; Cardoso et al., 2012; Chen et al., 2009; Sun et al., 2011). Hence, it becomes increasingly important to accurately compensate for susceptibility artefacts to be able to use EPI images for effective neuronavigation. There are strict time constraints associated with a neurosurgical procedure. Hence, any proposed solution must be computationally fast enough to work within these requirements. The current patient transfer time from the intra-operative scanner, after an imaging session, to the surgical bed at NHNN is between $7-9$ minutes. *All* image analysis tasks must be performed within this time window to ensure no extra time is added to the surgery.

In this chapter, I propose to meet the aforementioned challenges by combining the fieldmap and image registration based correction approach in a unified scheme. The main idea behind this work is a novel phase unwrapping algorithm that can also compute the uncertainty associated with the estimated fieldmap. The deformation field generated from the fieldmap correction step and the associated uncertainty measure are used to initialise and adaptively guide a subsequent image registration step. The overall workflow can be visualised as figure 4.1. The proposed work is also suitable to be used within the time constraints of a neurosurgical environment due to use of fast optimisation provided by graph cuts and has been successfully integrated into the surgical workflow at NHNN in London, UK.

The main contributions of this work are:

- A phase unwrapping algorithm using dynamic graph cuts that also determines the uncertainty associated with the estimated solution.

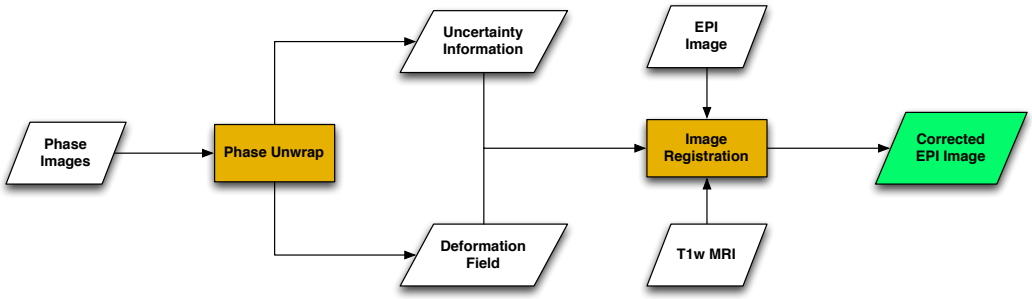

Figure 4.1: The proposed workflow for correction of susceptibility artefacts in EPI images acquired during neurosurgery. The field map is calculated using the acquired phase images which are unwrapped using the proposed algorithm. The estimated deformation field and the uncertainty information associated with the phase unwrapping step is used to initialise the image registration step where the EPI image and the corresponding undistorted T1-weighted MRI image is used as the source and the target images respectively. The registration step is selectively driven in regions of high uncertainty to improve the results in areas where the field map might have resulted in a sub-optimal solution.

- A registration algorithm that can be adaptively driven using the uncertainty information estimated from the phase unwrapping step to refine the results in areas where the fieldmap estimates are likely to be incorrect.
- Demonstrate the use of the proposed method during neurosurgery at NHNN, London on 13 patients within the time constraints of the intervention.

4.1 Associated Publications

- **Daga, P.**, Modat, M., Winston, G., White, M., Mancini, L., McEvoy, A. W., Thornton, J., Yousry, T., Duncan, J., Ourselin, S.: Susceptibility artefact Correction by combining B_0 field maps and non-rigid registration using graph cuts. (2013) Proc. SPIE, Medical Imaging. ***Winner: Best student paper award.***
- **Daga, P.**, Pendse, T., Modat, M., White, M., Mancini, L., Winston, G., McEvoy, A. W., Thornton, J., Yousry, T., Drobnjak, I., Duncan, J., Ourselin, S.: Susceptibility Artefact Correction using Dynamic Graph Cuts: Application to Neurosurgery. (2014) Medical Image Analysis.

The rest of the chapter is organised as follows: Section 4.2 describes the noise model in the MRI phase images and highlights the assumptions of our phase model. Section 4.3 describes the graph cuts based phase unwrapping method. Section 4.4 describes how uncertainty information can be computed from the phase unwrapping step and can be used with an image registration method to further improve results. Validation on synthetic and clinical datasets are descrived in section 4.6.1 and 4.6.2 respectively.

4.2 Noise in MRI Phase Images

The noise characteristics of MRI images were studied in detail by Gudbjartsson and Patz (1995). MRI phase images are reconstructed from the real and the imaginary images by calculating pixel by pixel the arctangent of their ratio. This is a nonlinear function and therefore the underlying noise distribution is not Gaussian anymore. The distribution of the phase noise, $\Delta\theta$, is given by equation (4.1).

$$p(\Delta\theta) = \frac{1}{2\pi} e^{-A^2/2\sigma^2} \left[1 + \frac{A}{\sigma}\sqrt{2\pi}\cos\Delta\theta \ \exp(A^2\cos^2\Delta\theta/2\sigma^2) \ \frac{1}{2\pi}\int_{-\infty}^{\frac{A\cos\Delta\theta}{\sigma}} \exp(-x^2/2)\ \mathrm{d}x \right] \tag{4.1}$$

where A is the noise-free phase value and σ is the standard deviation of noise in the real and imaginary channels (the noise is assumed to be identically distributed in the two channels). The underlying general distribution of the phase noise is, therefore, non-Gaussian. However, if we consider the case when $A = 0$ i.e. in background image regions where there is only noise, the distribution simplifies to $p(\Delta\theta) = 1/2\pi$ which corresponds to a uniform probability in all phase directions. Considering another case, where $A \gg \sigma$ i.e. image regions where the signal is significantly greater than noise, we also obtain a simpler distribution as:

$$p(\Delta\theta) \approx \frac{1}{2\pi(\sigma/A)^2} \ \exp(\frac{-\Delta\theta^2}{2(\sigma/A)^2}) \tag{4.2}$$

Hence, the phase noise distribution can be assumed to be additive zero mean Gaussian distributed when $A \gg \sigma$. The signal to noise ratio in iMRI images is typically lower than conventional MRI images. However, the Gaussian assumption of noise distribution is appropriate even for fairly small signal to noise ratios as was shown by Gudbjartsson and Patz (1995). The field map estimation method presented later in this chapter is formulated under this Gaussian noise distribution assumption.

4.3 Phase Modelling

As described in chapter 2, a popular method for estimating the magnetic field map is to use the phase difference between two MR images acquired at different echo times. The phase measurements at the two echo times can be used to generate the field map through equation (2.1) which can be converted to a one-dimensional voxel shift using equation (2.2). Hence, accurate correction of susceptibility artefacts is contingent upon being able to accurately measure the phase at the different echo times. However, the phase images are uniquely defined only in the range of $(-\pi, \pi]$ and the phase images need to be *unwrapped* at each voxel by an unknown integer multiple of 2π to obtain the true phase as in equation (2.3). In the absence of noise provided that the underlying field is spatially continuous, the only discontinuities that can occur in the measured phase image is due to wrapping itself. In that specific case, phase unwrapping is relatively easy to address. To unwrap, the phase difference between adjacent samples is calculated and if it is greater than π, phase wrapping has occured. In the absence of noise, the measured phase image can be correctly unwrapped provided that there are no discontinuities between adjacent voxels in the true phase image that are greater than π. While this algorithm is simple to implement, it can fail in areas with low signal to noise and these errors can propagate through the overall unwrapping process creating unwrapping failure over a large area.

To cope with this issue, I propose a Bayesian approach to the phase unwrapping problem. As already described, phase unwrapping is an ill-posed problem in the presence of noise and becomes intractable without regularisation. Similar to Ying et al. (2006), the phase is modelled as a Markov Random Field (MRF) where the true phase ϕ_t and the wrapped phase ϕ_w are treated as random variables. The

aim is to find the discrete label configuration k that gives the maximum a posteriori (MAP) estimate of the phase wraps as shown in equation (4.3). MRF is an intuitive choice for this problem as an individual voxel does not provide any information to perform the phase unwrapping and there is a need to specify spatial constraint and relationships among neighbouring voxels, which can be done conveniently through an MRF. Furthermore, there are computationally attractive options at our disposal to perform inference on such a system.

$$\phi_t = \max_k \underbrace{P(\phi_w|\phi_t)}_{\textbf{Likelihood}} \underbrace{P(\phi_t)}_{\textbf{Prior}} \tag{4.3}$$

The likelihood term in equation (4.3) is modelled as $\delta(\phi_w - W(\phi_t))$, where δ is the delta function and $W(\phi_t)$ is the wrapped true phase. This is ill-posed and additional constraints on the true phase are incorporated in terms of prior probabilities. The MR phase can be modelled as a piecewise smooth function where the smooth component is due to the inhomogeneities in the static MR field and the non-smooth component arises due to changes in the magnetic susceptibility at boundaries between tissues of different types. The spatial smoothness is enforced by modelling the true phase as a MRF and incorporating the smoothness model through a suitable potential function. In this work, I model the true phase as a six-neighbourhood pairwise MRF where the pairwise potential function used is the sum of square of difference of the true phase between adjacent neighbours. Owing to the MRF-Gibbs equivalance (Li, 1994), the phase unwrapping problem is to find the MRF labelling or configuration that minimises the energy $E(k|\phi_w)$:

$$E(k|\phi_w) = \underset{k}{\operatorname{argmin}} \sum_{i \in I} \sum_{\Omega} V(\Delta\phi_t^i) \tag{4.4}$$

where I are the image voxels, Ω is the set of neighbours for a given voxel at location i. $V(\Delta\phi_t^i)$ is the potential function defined on the difference potential between a voxel i and its neighbours in Ω. The unknown integer wraps are denoted by k. The following subsection describes how this integer constrained global optimisation problem can be efficiently solved using graph cuts.

4.3.1 Energy Minimization via Graph Cuts

As described in chapter 3, graph cuts have emerged as a popular method for optimisation of such multi-label problems (Boykov et al., 2001; Kolmogorov and Zabin, 2004). A first-order MRF of the form of equation (3.7) can be represented by a graph as long as the pairwise terms satisfy the inequality constraint of equation (3.6). If such a graph can be constructed, then fast inference algorithms are available to compute the MAP configuration of this MRF. It is easy to see that the proposed energy function of equation (4.4) has the structure of equation (3.7) with a null unary data term. The question is what pairwise energy term can we use which will satisfy the inequality constraint of equation (3.6). Unfortunately, there is no obvious way to formulate the pairwise energy term in such a fashion due to the additive term $\phi_w(i)$ in equation (2.3). However, it can be shown that as long as the pairwise energy function V of equation (4.4) is convex, this problem can still be solved through iterative graph cuts.

If the pairwise energy term V is convex and if the minima of $E(k|\phi_w)$ is not reached, a binary

image $\delta \in (0,1)$ exists such that $E(k+\delta|\phi_w) < E(k|\phi_w)$. For brevity let us consider the problem in one dimension and assume a two neighbourhood MRF system. Let $k^i_{t+1} = k^i_t + \delta^i$ be the wrap count at time $t+1$ at voxel i. Then, we have equation (4.5) where $\Delta\phi_t$ is the difference in the true phase between the MRF neighbours.

$$\Delta\phi_t = 2\pi(k^i_{t+1} - k^{i-1}_{t+1}) + (\phi^i_w - \phi^{i-1}_w) \tag{4.5}$$

After algebraic manipulation of equation (4.5), the energy function can be rewritten as equation (4.6).

$$\begin{aligned} E(k^t + \delta|\phi_w) = \underset{k}{\operatorname{argmin}} \sum_{i\in I}\sum_{N} V(2\pi(\delta^i - \delta^{i-1}) \\ +2\pi(k^i_t - k^{i-1}_t) + (\phi^i_w - \phi^{i-1}_w)) \end{aligned} \tag{4.6}$$

Now considering the terms in equation (3.6):

$$E(0,0) = V(t)$$
$$E(1,1) = V(t)$$
$$E(1,0) = V(2\pi + t)$$
$$E(0,1) = V(-2\pi + t)$$

where

$$t = 2\pi(k^i_t - k^{i-1}_t) + (\phi^i_w - \phi^{i-1}_w)$$

As V is convex, $E^{ij}(0,0)+E^{ij}(1,1) \leq E^{ij}(0,1)+E^{ij}(1,0)$ or $V(2\pi+t)+V(-2\pi+t) \geq 2\times V(t)$. Hence, the proposed energy term can be represented by a graph.

Figure 4.2(a) shows how an elementary graph between two MRF neighbours is constructed when $E^{ij}(1,0) - E^{ij}(0,0) > 0$ and $E^{ij}(1,0) - E^{ij}(1,1) > 0$. Similar constructions for other case exists as described in chapter (3). The complete graph is built by merging the elementary graphs for each node pair as illustrated in figure 4.2(b). After the complete graph is built the minimum cut on it can be found by pushing the maximum flow between the source and sink.

Hence, as long as the pairwise energy function employed is convex, we can represent the proposed MRF model with a graph. In this work, I employed the sum of the square of the L2 norm between the MRF neighbours as the pairwise energy function. However, any vector norm ≥ 1 can be used. Now, an iterative graph cut algorithm can be constructed which can efficiently find the minimum energy configuration of this MRF. This algorithm is described as pseudocode in listing (1).

Phase measurements in low signal areas tend to be less reliable and these areas can be discounted by assigning a weight to each voxel based on its magnitude. Similar to Ying et al. (2006), I use the magnitude image as a quality map and assign greater weight to voxels having large magnitude values.

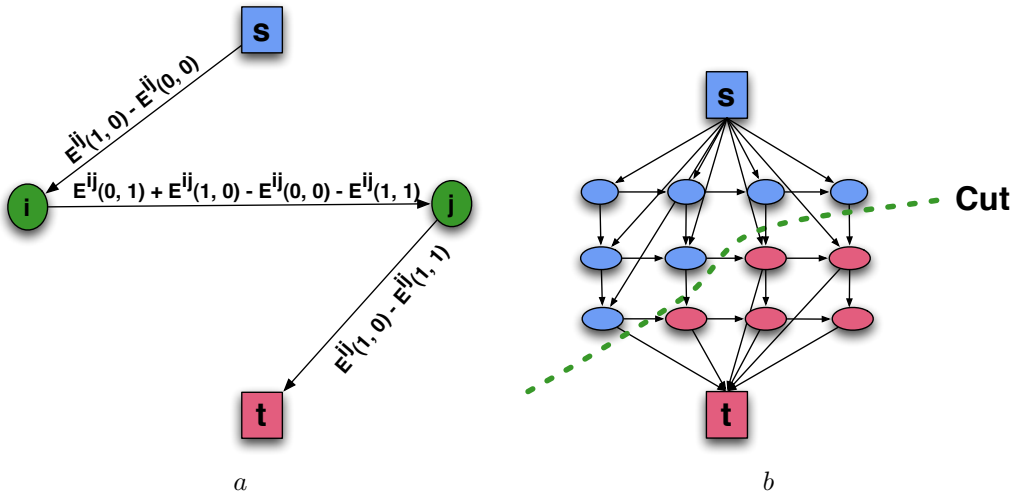

Figure 4.2: Graph Construction. (a) shows the construction of the elementary graph for a single pairwise term when $E^{ij}(1,0) - E^{ij}(0,0) > 0$ and $E^{ij}(1,0) - E^{ij}(1,1) > 0$. Note that the graph can only have non-negative edge weights. (b) shows the building of the graph by merging the elementary graphs together. After the graph is constructed, maximum flow algorithm can be used to find the minimum cut (denoted by the dashed line) on the graph.

After the phase images are unwrapped, the deformation field to correct the EPI image can be computed through equations (2.1) and (2.2). However, as previously mentioned, the estimated deformation can be inaccurate in image areas with low signal. In the following section, I will describe a way to compute the uncertainty associated with the estimated fieldmap and how this uncertainty information can be used in conjunction with an image registration step to further improve the results.

4.4 Uncertainty Estimation and Image Registration

This section explains how one can combine the fieldmap correction technique described in the previous step with image registration based techniques. The two techniques can be unified by estimating the uncertainty from the fieldmap step and using it with image registration to refine the deformation in image areas where the estimated fieldmap is likely to be inaccurate. The following sub-section describes how uncertainty information can be estimated during the phase unwrapping step.

4.4.1 Uncertainty Estimation in Phase Unwrapping

Besides fast MAP inference, another advantage of using graph cuts is its ability to be able to generate the uncertainty associated with the most likely MRF configuration. It was shown by Kohli and Torr (2008) that the uncertainty associated with the MAP solution can be estimated using graph cuts through computation of *max-marginals*. Max-marginals are a general notion and can be defined for any function as equation (4.7). Hence, the max-marginal ($\alpha_{v;j}$) is the maximum probability over all possible MRF configurations where an MRF site x_v is constrained to take the label j ($x_v = j$).

$$\alpha_{v;j} = \max_{x \in L, x_v = j} P(x|Y) \tag{4.7}$$

The max-marginals can be used to compute the confidence measure (ω) associated with any random variable labelling as equation (4.8).

Algorithm 1 The basic phase unwrapping algorithm

procedure PHASEUNWRAP($WrappedPhaseImage$)
  $k \leftarrow k' \leftarrow 0$ ▷ Set the initial wraps to 0
  $i \leftarrow 1$ ▷ Flag to keep iterating
  while $i \neq 0$ **do**
    Create $E(0,0), E(0,1), E(1,0), E(1,1)$ for every voxel ▷ Create graph
    Compute max flow ▷ Perform binary optimisation
    for all voxels x, y, z **do**
      if voxel$(x, y, z) \in T$ **then** ▷ The voxel belongs to the sink sub-graph
        $k'(x, y, z) \leftarrow k(x, y, z) + 1$ ▷ Make a 2π jump
      else
        $k'(x, y, z) \leftarrow k(x, y, z)$ ▷ Keep the current state
      end if
    end for
    if $E(k') < E(k)$ **then** ▷ Energy has decreased
      $k \leftarrow k'$
    else
      $i \leftarrow 0$ ▷ We have reached the minima
    end if
  end while
  return k ▷ k contains the number of 2π jumps at every voxel
end procedure

$$\omega_{v;j} = \frac{\max\limits_{x \in L, x_v = j} P(x|Y)}{\sum\limits_{k \in L} \max\limits_{x \in L, x_v = k} P(x|Y)} = \frac{\alpha_{v;j}}{\sum\limits_{k \in L} \alpha_{v;k}} \tag{4.8}$$

Therefore, the confidence $\omega_{v;j}$ for a random variable x_v to take the label j is given by the ratio of the max-marginal associated with assigning label j to variable x_v to the sum of max-marginals for all possible label assignments for the variable x_v.

As shown by Kohli and Torr (2008), this confidence can be expressed in terms of *min-marginal* energies. Min-marginal (ψ) is the minimum energy obtained when we constrain a random variable to take a certain label and minimise over all the remaining variables as in equation (4.9).

$$\psi_{v;j} = \operatorname*{argmin}_{x \in L, x_v = j} E(x) \tag{4.9}$$

The energy and probability of a labelling configuration are related through the expression for Gibbs energy function as:

$$E(x) = -\log P(x|Y) - \log Z \tag{4.10}$$

where Z is the partition function. Substituting the value of $P(x|Y)$ in equation (4.7) we have:

$$\begin{aligned} \alpha_{v;j} &= \max_{x \in L, x_v = j} (\exp(-E(x) - \log Z)) \\ &= \frac{1}{Z} \exp(-\operatorname*{argmin}_{x \in L, x_v = j} E(x)) \end{aligned}$$

Finally substituting equation (4.9), we have:

$$\alpha_{v;j} = \frac{1}{Z}\exp(-\psi_{v;j}) \tag{4.12}$$

Note that the knowledge of the partition function is not necessary to compute the max-marginal confidence measure. As an example, let us consider computing the max-marginal for a voxel to take a certain label 0. For the sake of simplicity, let us assume that it is a binary problem and only two configurations for this voxel are possible namely 0 and 1. The max-marginal value for this voxel to take the label 0 is given by:

$$\omega_{v;0} = \frac{\frac{1}{Z}\exp(-\psi_{v;0})}{\frac{1}{Z}\exp(-\psi_{v;0}) + \frac{1}{Z}\exp(-\psi_{v;0})} \tag{4.13}$$

Note that the Z's cancel out from the numerator and denominator.

Hence, the confidence measure ($\omega_{v;j}$) associated with any random variable x_v to take the label j can be expressed in general terms as equation (4.14), without estimating the partition function Z.

$$\omega_{v;j} = \frac{\exp(-\psi_{v;j})}{\sum\limits_{l \in L}\exp(-\psi_{v;l})} \tag{4.14}$$

Dynamic Graph Cuts can be used to compute $\omega_{v;j}$ for each voxel at every binary optimization step in a very efficient manner. A given MRF node can be constrained to belong to the source or the sink by adding an infinite capacity edge between it and the respective terminal node. No other changes need to be made to the graph and the required min-marginal can be computed by optimizing the resulting MRF. Hence, to compute the min-marginals at every binary optimisation step, one has to optimise one such MRF for every node v and each of the two labels. Usually these MRFs are very close to each other and form a slowly varying dynamic MRF system, which means that the search trees from previous computations can be efficiently reused, which greatly reduces the computation time.

This confidence map generated from the phase unwrapping step gives us a way to combine field map and image registration based susceptibility artefact correction techniques in an intuitive way. Areas of high uncertainty from the phase unwrapping step indicate where the generated field map is more likely to be unreliable. This knowledge can be used to adaptively refine the results in these areas using image registration. The following section describes how the generated deformation field and the confidence map can be used in an image registration framework to further improve the results.

4.5 Image Registration Framework

The displacement field and the confidence map generated from the phase unwrapping step are used to initialise the subsequent non-rigid registration step. As discussed in chapter 2, registration between the distorted EPI images and the undistorted T1/T2 weighted MR images is a popular alternative to using field maps for correcting for susceptibility artefacts. In this section, I will show how the two approaches can be combined using the uncertainty information derived from the phase unwrapping step.

The registration algorithm I developed follows closely from Glocker et al. (2008); So et al. (2011) and is formulated as a discrete multi-labelling problem. The deformation field is parameterised using

cubic B-splines as in Modat et al. (2010); Rueckert et al. (1999) which has the desirable property of generating smooth deformations.

A mutual information based image similarity measure was chosen for the proposed image registration algorithm. The key advantage of mutual information based measures is their ability to easily handle complex relationships between the intensities in the two images. They require no a-priori model of the relationship between the image intensities and can handle image registration between different modalities. Typically, graph cuts based optimisation algorithms cannot use such global similatity measures in the optimisation as it is difficult to adapt them directly in the data term in equation (3.7). To overcome this problem, a local variant of normalised mutual information (SEMI) as described by Zhuang et al. (2011) is used as the similarity measure. SEMI computes mutual information in a local region with respect to each of the control points. However, it uses a hierarchical weighting scheme to differentiate the contributions of different voxels to the similarity measure. The weighting scheme is chosen such that the weight given to a voxel is monotonically decreasing with respect to the distance between the voxel and the spline control point. Under this scheme, the joint histogram is computed as shown in equation (4.15) where $I_r(x)$ and $I_f(x)$ are the reference and transformed floating images. w_r and w_f are Parzen windows functions and the joint histogram is calculated for the local region Ω_s. $\Gamma_s(x)$ is a weighting function for the spatial encoding and is a Gaussian kernel centered on the control point. Hence, local joint histograms are computed for each of the control points and the corresponding data term used is generated by computing the normalised mutual information (Studholme et al., 1999) from each of these these joint histograms. The local nature of the similarity measure allows the problem to be formulated in the MRF framework which can be solved using the graph cuts framework.

$$H_s(r, f) = \sum_{x \in \Omega_s} (w_r I_r(x) w_f I_f(x)) \Gamma_s(x) \tag{4.15}$$

As registration is an ill-posed problem, priors on the estimated deformation field is usually introduced in the form of a smoothness term. A simple smoothness term would be to use the magnitude of the displacement vector difference at every registration iteration. This would result in registration scheme where incremental updates to the deformation field are penalised. This update scheme has the advantage of fulfilling the inequality constraint of equation (3.6) and can be easily accomodated into the graph cuts framework. However, it does not provide a regularisation over the full time course of the registration. In this work, I penalise the magnitude of the difference in the deformation as in Glocker et al. (2008) to perform a full regularisation as:

$$E_{ij}(x_i, x_j) = |(\mathcal{R}(i) + d_i) - (\mathcal{R}(j) + d_j)| \tag{4.16}$$

where $\mathcal{R}(.)$ projects the current displacement field to the control points and d is the displacement updates for the current iteration. It is worth noting that the inequality constraint of equation (3.6) for the pairwise term is not guaranteed to be met anymore. However, this is rarely a problem in practice as demonstrated in Glocker et al. (2008). The MRF nodes where the edge weights turn negative and the inequality constraint was violated were handled by setting those pairwise edge weights to zero. In

practice, this condition was only encountered in a handful of voxels.

The geometric distortion due to susceptibility is dominant in the phase encode (PE) direction. Hence the B-spline control points are constrained to move only in the PE direction. A discrete set of displacements is considered in the PE direction and a label assignment to a control point is associated with displacing the control point by the corresponding displacement vector. In this work, the step size is chosen to be one-third of the voxel size along the PE direction. Therefore, registration is done by solving this discrete multilabel problem modelled in the first-order MRF, where the cubic B-spline control points are the random variables and the goal is to assign individual displacement values to these nodes.

The final task that remains is to integrate the uncertainty information from the field map estimation step into the registration framework. The registration is initialised with the deformation field obtained from the field map. The goal is to adaptively drive the registration in areas where the field map results are uncertain. This is achieved by modulating the weight of the global penalty term (λ in equation 4.17) by the confidence map obtained during the phase unwrapping step. This has the effect of keeping the weight of the penalty term high in regions where the fieldmap is estimated with a low level of uncertainty thus discouraging large displacements whilst relaxing it in regions of high uncertainty to allow for more displacement. Hence, the spatially varying cost function takes the form of equation (4.17) where σ_i is the spatially varying confidence at voxel i, λ is the global penalty term weight and SEMI_i is the unary data term at control point i. The pairwise term $E_{ij}(x_i, x_j)$ is as defined in equation (4.16). The penalty term weights are initialised by projecting the confidence map on the control point grid. This cost function is optimised using an α-expansion variant of the graph cuts minimisation algorithm (Boykov et al., 2001).

$$E = -\sum_{i \in I} [(1 - \sigma_i \lambda) \times \text{SEMI}_i] + [\sigma_i \lambda \times E_{ij}(x_i, x_j)] \tag{4.17}$$

Similar to Studholme et al. (2000), the intensity distortions, due to susceptibility artefacts, are taken into account by recomputing the EPI intensities during image registration as $I_f = I_{Tf} J_T$ where I_f is the Jacobian corrected EPI image in the space of the reference anatomical image, I_{Tf} is the transformed EPI image where T is the current estimate of the transformation and J_T is its Jacobian determinant.

4.6 Validation

4.6.1 Validation Using Simulated Data

I validated the phase unwrapping algorithm using simulated phase MRI data. To conduct the simulations, an MRI simulator software package was used: POSSUM (Physics-Oriented Simulated Scanner for Understanding MRI) (Drobnjak et al., 2006, 2010). POSSUM is a simulator which generates realistic MR images. The simulator achieves this by simulating an MR scanner with various scanner input parameters operating on a physical model of the brain. The output of the simulator is the signal received from the receiver coil of the simulator scanner. The algorithm solves fundamental Bloch equations (Bloch, 1946) to model the behaviour of the magnetisation vector for each voxel of the brain and for each tissue type independently. The signal coming from one voxel is obtained by analytical integration of magnetisation over its spatial extent, and the total signal is formed by numerically summing the contributions from

all the voxels. For a given brain phantom, pulse sequence and magnetic field values, POSSUM generates realistic MR images. Magnetic field values are calculated by solving Maxwell's equations which as an input use an air-tissue segmentation of the brain, and their respective susceptibility values. These magnetic field values are fed into the Bloch equation solver in POSSUM, resulting in images with realistic susceptibility artefacts. A further, in-depth description of POSSUM is presented in Drobnjak et al. (2006).

I use a 3D digital brain phantom from the MNI BrainWeb database, which is thoroughly segmented into various tissues such as grey and white matter, cerebrospinal fluid, and has a good air-tissue segmentation (Collins et al., 1998). I assume a 1.5T scanner, and use appropriate MR parameter values for white matter ($T_1 = 833$ms, $T_2 = 83$ms, spin density $\rho = 0.86$); (grey matter $T_1 = 500$ms, $T_2 = 70$ms, $\rho = 0.77$) and CSF ($T_1 = 2569$ms, $T_2 = 329$ms, $\rho = 1$) (Rooney et al., 2007). A typical fieldmap sequence was simulated: two gradient echo images with different echo times ($TE_1 = 8$ms, $TE_2 = 10$ms). Spatial resolution was $2 \times 2 \times 2$mm and $TR = 700$ms.

In order to make the simulated images representative of images acquired during a surgical procedure, resections were introduced into the input phantom. The resections were designed to match the typical resections made during anterior temporal lobe resection for refractory epilepsy. Hence, actual T1-weighted intra-operative scans were used as reference for resection design. This modified phantom was used as an input to POSSUM and wrapped phase images and ground truth magnetic field values were simulated. The various inputs to the POSSUM simulator is shown in figure 4.3.

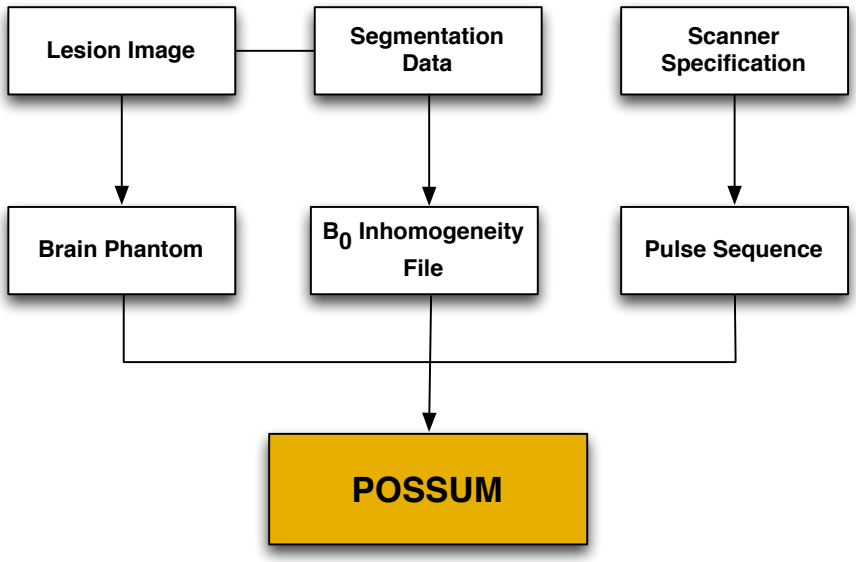

Figure 4.3: The various inputs to POSSUM to simulate the MRI phase images. Lesions are manually drawn in the input phantom image. The B_0 inhomogeneity file describes change in magnetic field strength inside the cranium due to tissue susceptibility differences. To calculate these distortions, Maxwell's equations are solved at each voxel in an air-tissue segmentation volume using the perturbation method. Finally, the MRI pulse sequence (eg. EPI) characteristics can be specified for each simulation.

For the validation, various levels of Gaussian noise were added to the ground truth unwrapped phase images. The corrupted images were then wrapped back to generate the phase images to be used as

Noise variance (rad.)	0.08	0.26	0.52	0.71	0.87	1.0	1.2
MCR (proposed)	0.01	0.05	0.11	0.13	0.17	0.19	0.23
MCR (PRELUDE)	0.01	0.06	0.14	0.19	0.23	0.24	0.31
Time-1 (proposed) sec.	5	5	6	8	8	8	9
Time-2 (proposed) sec.	25	25	28	28	27	30	32
Time (PRELUDE) sec.	4	22	161	534	1642	2532	4077

Table 4.1: Misclassification ratio (MCR) and execution time (in seconds) for generating the fieldmap from the synthetic phase images. The MCR is defined as the ratio between the voxels that were incorrectly wrapped to the total number of voxels. For small amounts of phase noise (noted in radians), both the proposed phase unwrapping algorithm and PRELUDE perform similarly. However, for larger noise levels, the proposed algorithm results in lower MCR. The execution time of PRELUDE for high levels of phase noise does not satisfy the stringent time requirements of neurosurgery, while the proposed algorithm executes well within the time constraints. Time-1 refers to the time taken by the proposed method to do phase unwrapping without confidence map estimation. Time-2 is for phase unwrapping along with confidence map estimation. All times are reported in seconds. The mean noise variance in the standard clinical datasets produced on the iMRI was 0.71 radians (corresponding simulation result highlighted in green).

input for the unwrapping algorithms. For comparison, the images were unwrapped using the proposed unwrapping algorithm as well as with PRELUDE (Jenkinson, 2003), a freely available software package available with FSL (Smith et al., 2004) and used within the neuroimaging community.

The quantitative unwrapping results for the proposed method and PRELUDE are shown in table 4.1. The results were compared with the original (ground truth) phase, and the misclassification ratio (MCR) was calculated. The MCR is the ratio of the number of voxels that were incorrectly unwrapped to the total number of voxels. Both PRELUDE and the proposed unwrapping algorithm perform comparably well under low-noise conditions. However, at higher noise levels the proposed algorithm outperforms PRELUDE both in terms of MCR and execution time. In addition, the proposed algorithm also generates the confidence associated with the unwrapping solution and can compute it within the time constraints associated with a neurosurgical procedure. A visual example is shown in figure 4.4. Some discontinuities around the lesion still exist when unwrapping with PRELUDE but not when using the proposed phase unwrapping technique. Figure 4.4(f) shows the confidence map generated along with the unwrapped image.

4.6.2 Validation Using Clinical Data

I used the proposed algorithm on 13 datasets that were acquired using interventional MRI during temporal lobe resection procedures for surgical management of temporal lobe epilepsy. The imaging was done as part of an audit to quantify the benefits of using iMRI on patient outcome for subjects having temporal lobe resections. The images were acquired using a 1.5T Espree MRI scanner (Siemens, Erlangen) designed for interventional procedures. The T1-weighted MR image, used in the registration step, had a resolution of 1.1 $\times$ 1.1 $\times$ 1.3mm using a 3D FLASH sequence with TR = 5.25ms, TE = 2.5ms and flip angle = 15°. The EPI images used a single shot scheme with GRAPPA parallel imaging (acceleration factor of 2) and had a spatial resolution of 2.5mm $\times 2.5$mm $\times 2.7$ mm. The phase encoding was applied in the anterior-posterior direction and the total read-out time was 35.52 ms. The noise variance in these datasets was measured in manually selected region of interest known to only contain air. The mean noise

variance was 0.71 radians.

Validation of the proposed susceptibility correction in the absence of ground truth deformation is not trivial. A popular approach has been to identify landmarks on the EPI and T1-weighted or T2-weighted MR images (obtained with conventional spin or gradient echo sequences with negligible spatial distortion) and measure the distance between the landmarks before and after performing the correction. However, this method tends to bias the results towards image registration based schemes. This is because intensity based registration algorithms tend to perform better in regions with high contrast which is precisely where landmarks can be reliably identified. Secondly, it is very difficult to reliably pick landmarks on interventionally acquired EPI images due to increased levels of noise, low spatial resolution and presence of deformation. Since I was interested in achieving accurate artefact correction in the white matter areas, I focused on looking at the effect of susceptibility correction on residual tensor fit errors. One significant source of tensor fit errors is the geometric distortions arising from susceptibility artefacts. Hence, accurate correction of susceptibility artefacts should reduce residual errors after performing tensor fitting. A previous study also demonstrated that nonlinear correction of susceptibility artefacts resulted in smaller tensor fit errors (Kim et al., 2006).

The normalised sums of square of diffusion tensor fit errors (χ^2) is given by equation 4.18 where N signals are fitted and S_m and S_f are the measured and fitted signals respectively (Papadakis et al., 2002).

$$\chi^2 = \frac{\sum_{i=1}^{N}(S_m - S_f)^2}{\sum_{i=1}^{N} S_m^2} \tag{4.18}$$

The diffusion tensors were reconstructed using *dtifit* (Smith et al., 2004) and sum of square residual errors for the diffusion tensor fits were obtained for the 13 subjects. For the validation, the initial sums of square residual tensor fit errors were computed for all subjects. Correction was performed after unwrapping the phase maps using PRELUDE and the proposed phase unwrapping algorithm. I also performed the correction using the registration algorithm described in section 4.5 and finally using the proposed method combining the fieldmap and image registration algorithm. The quantitative results are described in table 4.2. A paired t-test showed that the proposed method showed a statistically significant reduction (p-value $< 10^{-3}$) in residual tensor fit errors when compared to fieldmap and image registration based techniques alone. Figure 4.5 shows a representative slice where the corrected B0 image using the proposed method shows good visual correspondence with the undistorted T1-weighted image.

4.7 Discussion

I have presented a novel susceptibility correction algorithm that can be used with the time constraints of a typical neurosurgical procedure. While distortion correction is routinely done in the diffusion and fMRI imaging community, it is not the case in the interventional MRI community. This is because the use of diffusion MRI imaging for guiding interventions is not a common practice. In this work, I have taken two of the most commonly used methods for correction of susceptibility artefacts and unified them in a principled manner. Initial validation results indicate that combining information from field maps and

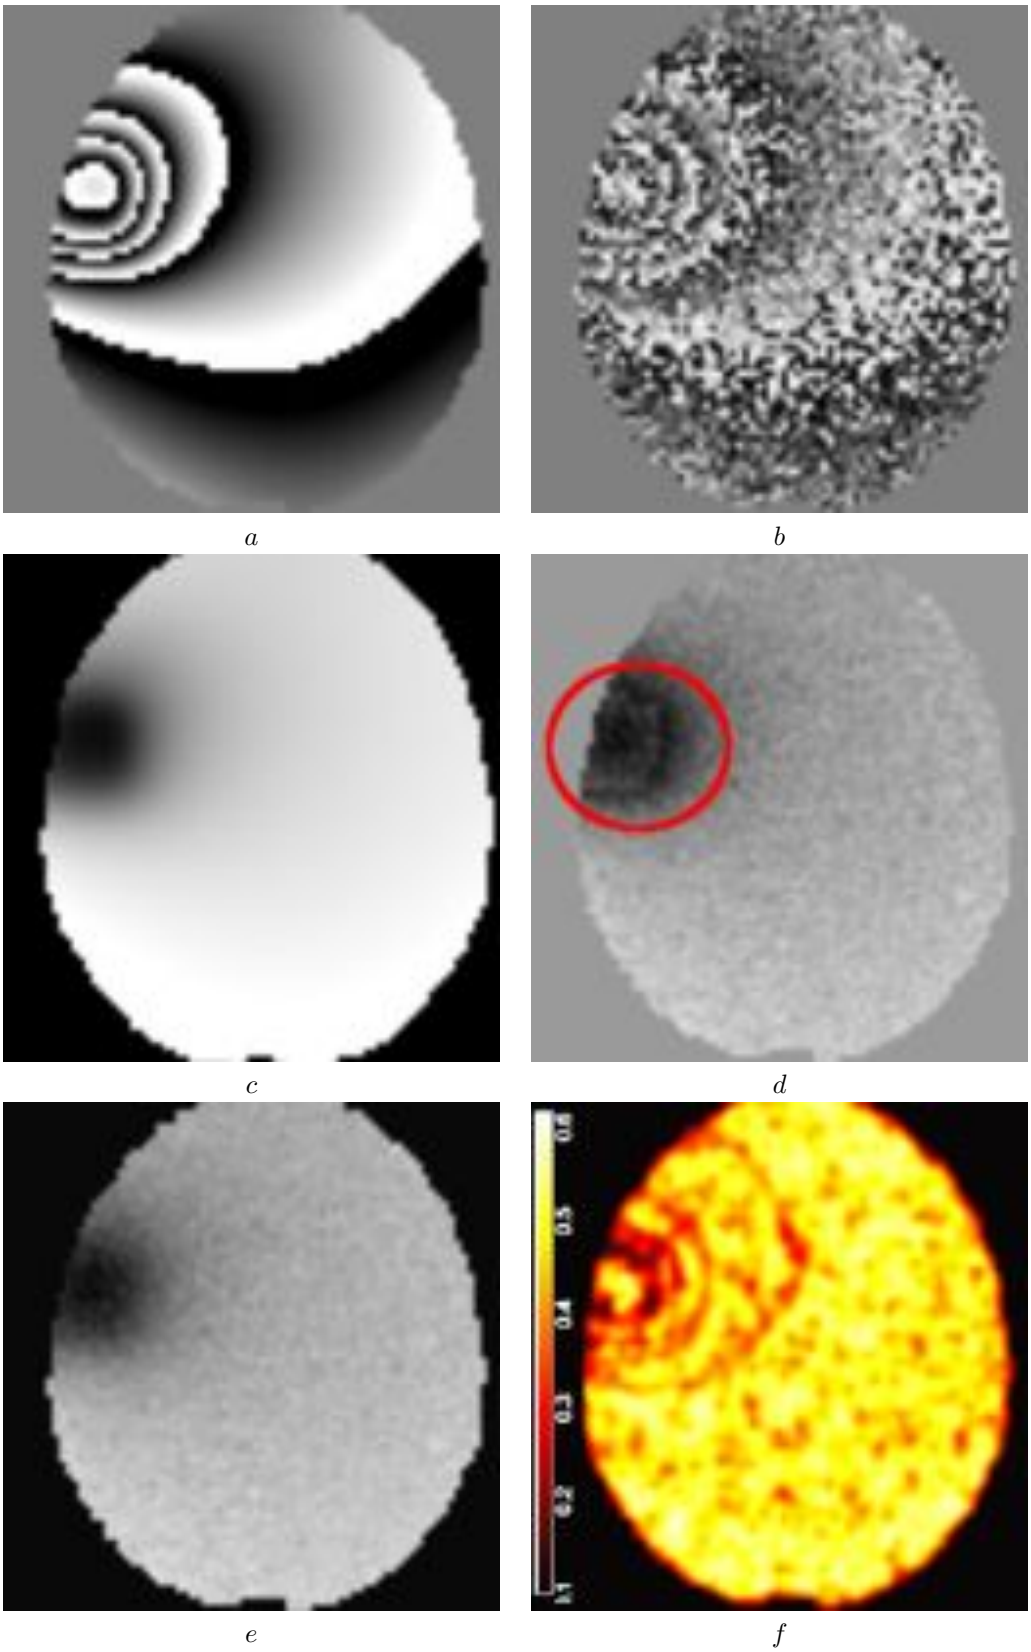

Figure 4.4: Results from phase unwrapping. (a) is a masked slice through a noise free wrapped image. (b) is the same image where the ground truth unwrapped image was corrupted with Gaussian noise. (c) shows the ground truth unwrapped image. (d) shows the unwrapping result from PRELUDE. Some areas with phase discontinuities are visible in the unwrapped result (highlighted in red). (e) Shows the unwrapped image using the proposed phase unwrapping algorithm where no phase discontinuities are evident. (f) shows the confidence map obtained using the proposed algorithm.

Initial	PRELUDE	Fieldmap only	Reg. only	Proposed
3.08(1.94)	1.92(1.14)	1.51(1.23)	1.31(0.97)	1.23(0.84)
2.94(1.89)	1.51(1.42)	1.48(1.35)	1.14(0.87)	1.12(0.73)
2.97(3.56)	1.94(2.26)	1.94(2.26)	1.98(1.56)	1.03(1.21)
3.40(3.17)	2.71(1.54)	2.42(1.06)	2.51(1.99)	2.34(0.98)
1.76(1.42)	1.43(1.36)	1.38(1.11)	1.12(0.76)	1.08(0.81)
2.27(2.30)	1.23(1.08)	1.36(1.02)	1.68(1.54)	1.22(1.12)
3.85(3.91)	2.78(2.51)	2.42(2.02)	2.53(1.91)	2.39(1.62)
2.70(2.37)	2.12(1.43)	2.04(1.51)	2.04(1.63)	1.72(1.43)
3.60(3.51)	2.53(2.01)	2.19(1.84)	2.61(2.30)	1.81(0.93)
2.32(1.85)	1.32(1.01)	1.45(1.04)	1.76(1.34)	1.41(0.96)
2.17(2.11)	1.16(0.86)	1.12(0.72)	1.47(1.14)	1.07(0.93)
2.81(2.62)	1.93(1.62)	1.59(1.22)	2.12(1.69)	1.41(1.04)
2.02(2.17)	1.16(0.91)	1.07(0.86)	1.41(1.38)	1.01(0.92)
2.76(0.63)	**1.82(0.58)**	**1.69(0.46)**	**1.82(0.52)**	**1.44(0.47)**

Table 4.2: Mean(standard deviation) of the sum of square errors for diffusion tensor fitting in interventionally acquired diffusion weighted images for thirteen subjects. The first column (Initial) shows the initial mean error. The second column (PRELUDE) shows the fit errors after correcting for susceptibility artefacts using PRELUDE. The third column (Fieldmap only) shows the tensor fit errors after correcting for susceptibility artefacts using the fieldmap generated after unwrapping the phase maps using the proposed phase unwrapping algorithm. The fourth columns (Reg. only) shows the tensor fit errors after correcting for susceptibility artefacts using the proposed registration algorithm. The final column (Proposed) shows the tensor fit errors after combining the fieldmap and image registration methods using the proposed method. The proposed method showed statistically significant improvement over the other methods (p-value $< 10^{-3}$). The final row shows the mean tensor fit errors and standard deviation over all the cases.

image registration yield a better estimate of the underlying distortion. However, further validation needs to be done to ascertain whether the proposed method offers any tangible benefits in terms of patient outcome. Typically, the accuracy with which a surgeon can perform a particular neurosurgical procedure is limited and this imposes a limit on how much additional benefit we can confer through interventional imaging.

This susceptibility correction algorithm is used as a preprocessing step before the brain shift estimation to correct the interventionally acquired diffusion MRI images. The following chapter will describe an intensity based image registration algorithm, which combines information from boh structural and diffusion MRI, and can be used for accurate estimate of brain shift during neurosurgery.

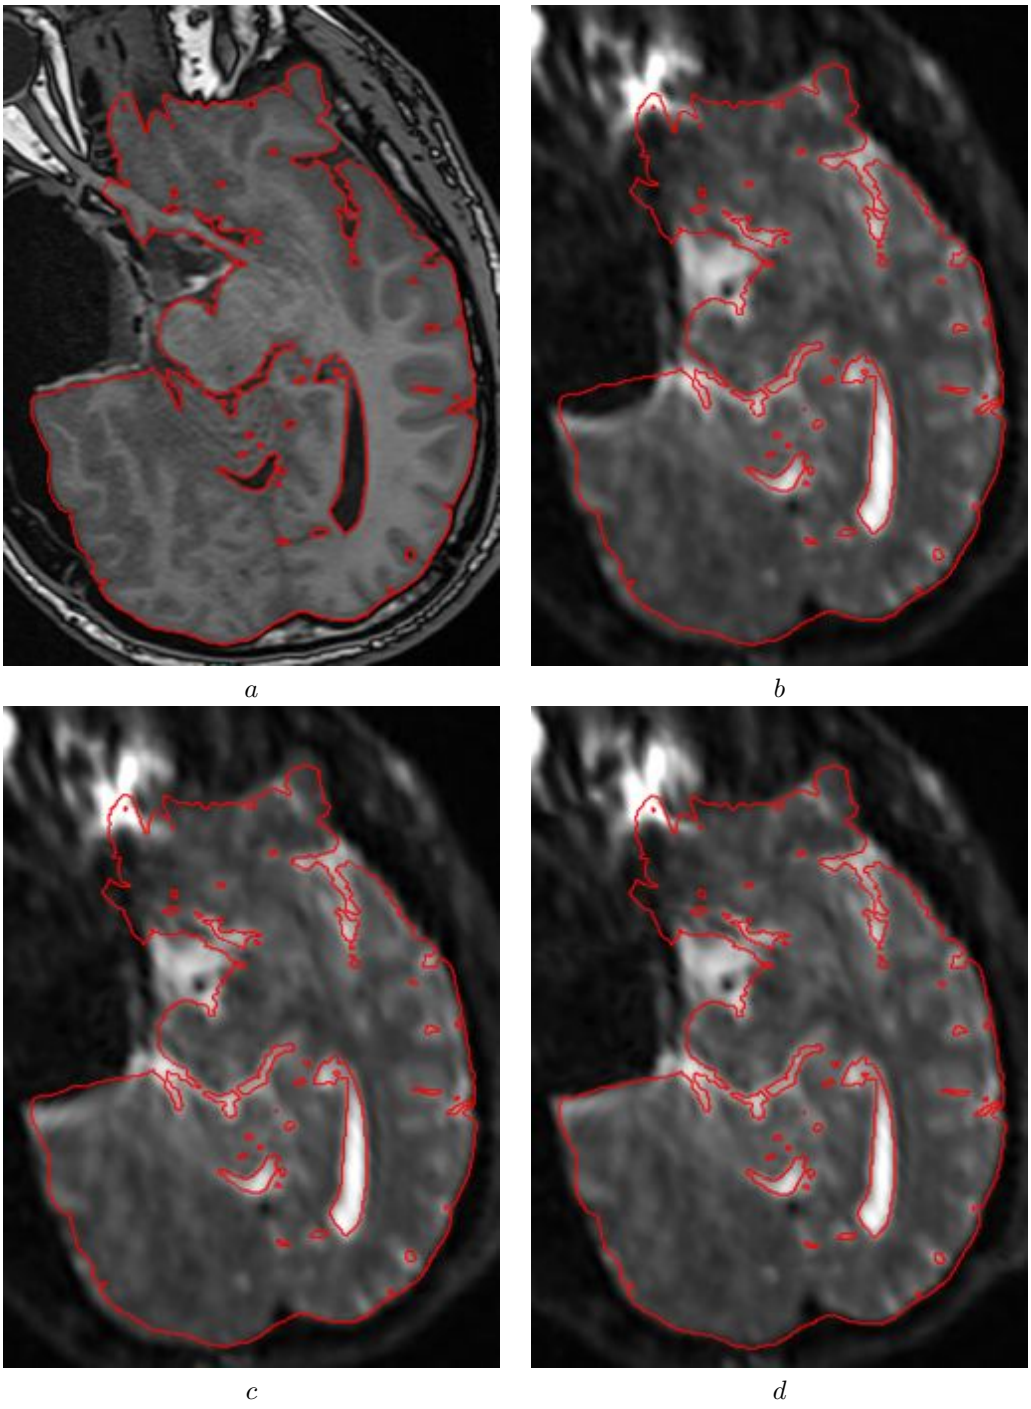

Figure 4.5: Images showing the result of correcting for susceptibility-induced spatial distortion using our algorithm. (a) shows the gold-standard high resolution T1-weighted image acquired during surgery. (b) shows the uncorrected B0 image with a large geometric distortion around the resected area. (c) shows the result of correcting for susceptibility artefacts using the proposed fieldmap estimation. (d) shows further improvement in the result when combined with the image registration step.

Chapter 5

Optic Radiation Localisation during Neurosurgery

As stated in the chapter 1, accurate localisation of the target lesions as well as functionally eloquent brain areas is needed to minimise chances of new morbidity to the patient and improve surgical outcome. The current commercial iMRI neuronavigation systems use rigid registration between the preoperative and intraoperative images and as described in chapter 2, rigid registration cannot capture the deformations caused by brain shift as these deformations are highly non-linear. Furthermore, the current iMRI neuronavigation systems also do not use the multimodal imaging capabilities of the iMRI scanners to estimate the mapping between preoperative and intraoperative images.

In this chapter, I propose a new image registration method designed specifically for estimation of brain shift during neurosurgery. I will demonstrate that the proposed algorithm can be executed within the time constraints of a typical neurosurgery procedure. Furthermore, the proposed algorithm makes use of diffusion weighted imaging along with traditional structural MRI to estimate the brain shift. I will show that this results in a more accurate estimate of the brain shift and leads to better localisation of optic radiation during surgery. This work has been integrated into the surgical workflow at NHNN in London, UK. The main contributions of this work are:

- An image registration algorithm designed for estimation of brain shift during neurosurgery.
- A similarity measure for use in image registration that utilises information from both structural and diffusion MRI to estimate brain shift and localise the optic radiation during surgery.
- Validation that shows improved localisation of optic radiation when using the proposed similarity measure.

5.1 Associated Publications

- **Daga P.**, Winston G., Modat M., White M., Mancini L., Cardoso M. J., Symms M., Stretton J., McEvoy A. W., Thornton J., Micallef C., Yousry T., Hawkes D., Duncan J. S., Ourselin S.: Accurate Localisation of Optic Radiation during Neurosurgery in an Interventional MRI Suite. (2012) IEEE Transactions on Medical Imaging.

- Winston G., **Daga P.**, Stretton J., Modat M., Symms M., McEvoy A. W., Ourselin S., Duncan J. S.: Optic Radiation Tractography and Vision in Anterior Temporal Lobe Resection. (2011) Annals of Neurology.
- **Daga, P.**, Winston, G., Modat, M., Cardoso, M. J., White, M., McEvoy, A. W., Thornton, J., Hawkes, D., Duncan, J., Ourselin, S.: Improved neuronavigation through integration of intra-operative anatomical and diffusion images in an interventional MRI suite, (2011), IPCAI.
- **Daga, P.**, Winston, G., Modat, M., Cardoso, M. J., Stretton, J., Symms, M., McEvoy, A. W., Hawkes, D., Duncan, J., Ourselin, S.: Integrating Structural and Diffusion MR Information for Optic Radiation Localisation in Focal Epilepsy Patients, (2011), IEEE ISBI.
- Winston, G., **Daga, P.**, Stretton, J., Modat, M., Symms, M., McEvoy, A. W., Ourselin, S., Duncan, J.: Propagation of Probabilistic Tractography of the Optic Radiation in Epilepsy Surgery, (2011) ISMRM.

5.2 Methods

An intuitive solution to the challenge of accurate localisation of surgical targets and functionally eloquent brain regions is to incorporate information from structural and diffusion MR images into a non-rigid image registration scheme. To the best of my knowledge, there is relatively little body of work around such multichannel image registration schemes. This is especially true when we look at registration algorithms that can be used within a neurosurgical environment. iMRI systems have steadily grown in their capabilities to provide multimodal imaging. However, the image analysis algorithms used within the iMRI environment have failed to keep up with these developments. There are a few general purpose multi-channel image registration schemes developed by the medical image analysis community but none of them are designed specifically for the purposes of intraoperative brain shift estimation. A multi-channel variation of the demons algorithm was proposed by Park et al. (2003) to register DTI datasets and create a group diffusion tensor atlas. Similarly, Avants et al. (2007) presented a multivariate approach using fused structural and diffusion data. Even though both these works used images from various modalities, there is no explicit formulation to utilise the shared information present in the various images and they ignore any influence that one image modality may have in explaining the structure of the other. More recently, Li and Verma (2011) proposed a multichannel registration scheme that fuses information from multiple modalities using feature analysis through Gabor wavelets transform. A novel multivariate mutual information (MI) based similarity measure called diffusion paired MI which uses structural MRI and DTI datasets in a unified similarity measure was presented by Studholme (2008). This method exploits the shared information between the structural MRI image and the diffusion tensor components. However, it requires the computation of 7 four-dimensional joint histograms for computation of the joint entropy. This computational complexity renders it currently unsuitable for use in a neurosurgical setting due to the time constraints.

In this work, I propose using a *bivariate* normalised mutual information as the image similarity measure in order to incorporate information from both structural and diffusion MRI imaging modalities

in image registration. An advantage of using this measure is also that it utilises the *shared* information within the images of these two modalities. To incorporate the information from diffusion MRI images, I use fractional anisotropy (FA), which is widely used scalar index derived from eigenvalues of a diffusion tensor.

5.2.1 Fractional Anisotropy

A commonly used model to infer the tissue microstructure from diffusion weighted imaging is the diffusion tensor imaging (DTI) model. DTI models the diffusion process with a Gaussian distribution and estimates a symmetric, positive-definite 3×3 matrix called the *diffusion tensor* (DT). FA is a scalar value derived from the DT that describes the degree of anisotropy of the diffusion process. The expression for FA is given as:

$$FA = \sqrt{\frac{3}{2}} \frac{\sqrt{(\lambda_1 - \hat{\lambda})^2 + (\lambda_2 - \hat{\lambda})^2 + (\lambda_3 - \hat{\lambda})^2}}{\sqrt{\lambda_1 + \lambda_2 + \lambda_3}} \tag{5.1}$$

where λ_1, λ_2 and λ_3 are the eigenvalues of the DT and $\hat{\lambda} = (\lambda_1 + \lambda_2 + \lambda_3)/3$ also called *trace* of the DT. FA is a scalar value between 0 and 1. FA is low in regions where the diffusion tends to be isotropic (e.g. the cerebrospinal fluid) and high where there is preferred diffusion along one direction due to highly ordered white matter tracts (corpus callosum, for instance). Hence, FA allows inference on the underlying tissue microstructure environment.

5.2.2 Bivariate NMI

The bivariate NMI (Daga et al., 2011a,b, 2012) that I propose that combines structural MRI and FA images in a unified image registration similarity measure. Instead of a single target and source image, we now consider a pair of target and source images. In this case, the target and source image pairs consist of intra-operative and pre-operatively acquired structural and fractional anisotropy (FA) images respectively. The bivariate NMI similarity measure ($\mathcal{S}$) between the target and source images $\{R_1, R_2\}$ and $\{F_1, F_2\}$ respectively is given by extending the conventional NMI definition as follows:

$$\mathcal{S}(R_1, R_2, F_1(\mathbf{T}), F_2(\mathbf{T})) = \frac{H(R_1, R_2) + H(F_1(\mathbf{T}), F_2(\mathbf{T}))}{H(R_1, R_2, F_1(\mathbf{T}), F_2(\mathbf{T}))}, \tag{5.2}$$

where $\mathbf{T}$ is the current deformation between the target and source images. $H(R_1, R_2)$ and $H(F_1(\mathbf{T}), F_2(\mathbf{T}))$ represent the joint entropy between the two target images and the two deformed source images respectively. $H(R_1, R_2, F_1(\mathbf{T}), F_2(\mathbf{T}))$ is the joint entropy between the four input images and is computed using Shannon's formula for entropy as:

$$H(R_1, R_2, F_1(\mathbf{T}), F_2(\mathbf{T})) = - \sum_{r_1, r_2, f_1, f_2} p(r_1, r_2, f_1, f_2) \times \log(p(r_1, r_2, f_1, f_2)),$$

where r_1, r_2, f_1 and f_2 are the voxel intensities of images R_1, R_2, $F_1(\mathbf{T})$ and $F_2(\mathbf{T})$ respectively.

Each probability is computed using a joint histogram $\mathcal{H}$ as:

$$p(r_1, r_2, f_1, f_2) = \frac{\mathcal{H}(r_1, r_2, f_1, f_2)}{\sum_{r_1, r_2, f_1, f_2} \mathcal{H}(r_1, r_2, f_1, f_2)},$$

where a Parzen Window (Mattes et al., 2003; Thevenaz and Unser, 2000; Viola and Wells, 1995) is used to estimate the joint histogram:

$$\begin{aligned}\mathcal{H}(r_1, r_2, f_1, f_2) = \sum_{\vec{x} \in R} \big[& \beta^3(R_1(\vec{x}), r_1) \times \beta^3(R_2(\vec{x}), r_2) \\ & \times \beta^3(F_1(\mathbf{T}(\vec{x})), f_1) \times \beta^3(F_2(\mathbf{T}(\vec{x})), f_2) \big]\end{aligned}$$

where β^3 is a cubic B-Spline kernel which is used as the Parzen window kernel. The Parzen window technique essentially involves adding weight in the vicinity of the voxel intensities rather than doing a simple increment of the joint histogram bin. This is shown to be more robust in presence of noise and intensity non uniformities.

The transformation model used in the proposed image registration is the parametric cubic spline model. The spline transformation model is described in more detail in chapter (4) and is omitted here.

In order to ensure plausibility of the estimated transformation, a penalty term is usually added to the similarity measure as a smoothness constraint. Hence, there is a balance between unconstraint optimization of the similarity measure and the smoothness of the estimated transformation. The bending energy (BE) of the transforming spline is used as a penalty term to constrain the solution as equation 3.19.

The objective function for the registration is given by:

$$\Omega(R, F(\mathbf{T})) = \alpha \times \mathcal{S} - (1 - \alpha) \times \mathrm{BE}$$

This objective function is optimised using a conjugate gradient ascent scheme. Gradient ascent is a first-order optimisation scheme and requires computation of the first derivative of the similarity measure with respect to the control point position. In order to compute the derivative of the NMI, one must compute the derivative of the joint entropy terms. This can be achieved by computing the derivative of the probability of each group of intensities. The derivative of each probability can be calculated by computing the derivative of the joint histogram according to each degree of freedom which is given by:

$$\begin{aligned}\frac{\partial \mathcal{H}(r_1, r_2, f_1, f_2)}{\partial \mu^{\xi}_{ijk}} = & \sum_{\vec{x} \in R} \beta^3\left(R_1\left(\vec{x}\right), r_1\right) \times \beta^3\left(R_2\left(\vec{x}\right), r_2\right) \\ & \times \Bigg(\left.\frac{\partial \beta^3\left(u, f_1\right)}{\partial u}\right|_{u = F_1(\mathbf{T}(\vec{x}))} \left.\frac{\partial F_1\left(p\right)}{\partial p}\right|_{p = \mathbf{T}(\vec{x})} \frac{\partial \mathbf{T}(\vec{x})}{\partial \mu^{\xi}_{ijk}} \\ & \quad \times \beta^3\left(F_2\left(\mathbf{T}\left(\vec{x}\right)\right), f_2\right) + \beta^3\left(F_1\left(\mathbf{T}\left(\vec{x}\right)\right), f_1\right) \\ & \times \left.\frac{\partial \beta^3\left(u, f_2\right)}{\partial u}\right|_{u = F_2(\mathbf{T}(\vec{x}))} \left.\frac{\partial F_2\left(p\right)}{\partial p}\right|_{p = \mathbf{T}(\vec{x})} \frac{\partial \mathbf{T}(\vec{x})}{\partial \mu^{\xi}_{ijk}} \Bigg)\end{aligned} \tag{5.3}$$

where ξ are the x, y and z components of the control point μ_{ijk}. $\frac{\partial F(p)}{\partial p}$ is the gradient of the deformed floating image with respect to the current transformation parameters. $\frac{\partial \beta^3(u,f)}{\partial u}$ are the first

order derivatives of the cubic B-spline given by:

$$dB_0(u)/du = (-u^2 + 2u - 1)/2$$
$$dB_1(u)/du = (3u^2 - 4u)/2$$
$$dB_2(u)/du = (-3u^2 + 2u + 1)/2$$
$$dB_3(u)/du = u^2/2$$

where the input parameter u is the input to the B-spline basis function and has a support width of [-2, 2] for cubic B-splines. For computational efficiency, the proposed gradient of the bivariate NMI is initially computed at each voxel position and then convolved with the appropriate B-spline kernel to produce the gradient value at each B-spline node position. The analytical derivative of the BE term is also needed and it was computed as defined in Modat et al. (2010).

The registration is performed using a multi-resolution approach where three levels of pyramidal downsampling are used to perform the registration. The registration is performed at the coarsest levels first and the deformation field is propagated to initialise the next finer level. In the experiments, the control point spacing was 5 voxels. For example, an image with dimensions of $282 \times 352 \times 154$ voxels result in a B-spline grid of $61 \times 75 \times 35$ control points.

As I have previously mentioned, non-rigid registration is a computationally expensive process and a fast implementation is needed to use this technology in the neurosurgical setting. In this work, I leverage the parallel computing capabilities of modern GPUs and implement the proposed algorithm in a parallel friendly manner to satisfy the stringent time constraints of the surgical procedure. Current GPU-based image registration implementations have been reviewed in Fluck et al. (2011). However, most of these implementations use sum of squared differences (SSD) as the similarity measure, which is not suitable for multi-modal image registration. Mutual Information based implementations have also been reported but none of them currently provide a multichannel similarity measure.

The CUDA framework (NVIDIA, 2008) was used for the implementation of the proposed algorithm on the GPU. CUDA utilises the many-core architecture of the modern GPUs for data-parallel computation processes. The majority of the GPU functions (termed *kernels*) for the registration algorithm are previously described in Modat et al. (2010). The implementation for the proposed multivariate scheme extends the registration algorithm to include GPU accelerated computation of the joint histogram and the analytical gradient of the proposed similarity measure as defined previously.

Computation of the marginal and joint entropies as described in equation (5.2) is computationally expensive when done serially on a CPU. The core of the computational complexity is shared between the Parzen Window smoothing of the joint histogram and the marginalisation along the reference and resampled floating image axes to compute the marginal entropies. Considering the use of 64 bins per image, a serial implementation has to perform 4×64^4 iterations in order to smooth the joint histogram using the Parzen window approach. In my parallel implementation, the smoothing of the joint histogram is done on the GPU by using four serial CUDA kernels; one for each dimension of the joint histogram. For a bin size of 64 along each dimension, each of these kernels launch 64^3 CUDA threads where each

concurrent thread effectively smoothes one line of the joint histogram along the given direction.

The marginalisation of the joint histogram along the two reference and the two resampled floating image axes is split into four CUDA kernels corresponding to each joint histogram axis. Each concurrent thread sums a line along a given direction. To manage loss of accuracy due to the use of single precision floating point in GPUs, *compensated summation* (Kahan, 1965) is used for accumulation in these kernels.

One of the major contributors to the speed improvement of the proposed algorithm is the use of the analytical objective function gradient as described in equation (5.3). This is significant as the usual symmetric difference based computation of the gradient is computationally extremely expensive. However, the implementation of equation (5.3) as it stands involves significant computational redundancy, since each voxel is included in the neighbourhood of many control points. In addition, it is also memory intensive as each spline node requires one joint histogram per degree of freedom. To alleviate these problems, a voxel-centric rather than a node-centric approach is used. The gradient of the proposed similarity measure is initially computed at the voxel position as in Modat et al. (2010). This allows each concurrent CUDA thread to process each voxel independently. A convolution with the cubic B-spline curve which corresponds to the basis functions in the deformation model is applied to the voxel-based gradient field to obtain the required gradient values at the spline control point positions.

The source code for the registration algorithm can be freely downloaded under an open source licence[1].

5.3 Validation

The validation of the proposed algorithm focuses on the two main criteria: registration accuracy and the computation time. As previously mentioned, the registration needs to accurately localise the structures of interest within the time constraints of neurosurgery. For assessing the accuracy of the image registration, I use a numerical phantom and also pre- and post-operative clinical datasets from a set of 20 patients who underwent anterior temporal lobe resection for treatment of refractory focal epilepsy. To validate the applicability of the algorithm in the interventional setting, I apply the algorithm retrospectively to 10 interventional MRI datasets acquired from 5 subjects and assess its accuracy by correlating the predicted outcome with the observed post-operative VFD.

5.3.1 Validation Using a Numerical Phantom

A numerical phantom was constructed in order to assess the accuracy of the proposed registration algorithm. For the structural image phantom (see figure 5.1), a very high resolution digital phantom containing finger and sheet like collapsed sulci and gyri was created, simulating the structure of the cortex. The phantom was created on a 0.25 mm equivalent isotropic image with a size of $180 \times 180 \times 120$ voxels. Gaussian noise was added in the Fourier domain to create the Rician noise corrupted phantom. The fibre tracts were created to span the white matter region of the phantom as shown in figure 5.1(c). This phantom allowed for comparison of the proposed registration algorithm against univariate registration schemes that use anatomical or diffusion only images.

[1] http://sourceforge.net/projects/niftyreg

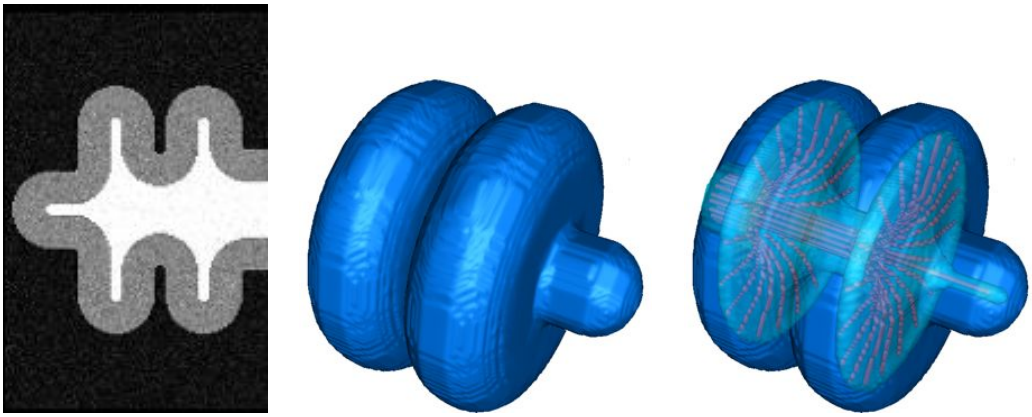

a *b* *c*

Figure 5.1: Numerical phantom. (a) shows the simulated cortical layer, (b) shows the 3-dimensional reconstruction of the phantom surface and (c) shows the simulated white matter tracts spanning the phantom.

Known random deformations were applied to the phantom and three different registration schemes were used to recover the deformation. First, the structural phantom was registered to the deformed structural phantom to simulate registration between structural images. Secondly, the WM phantom, which was also a scalar image and designed to simulate fractional anisotropy images, was registered to the deformed WM phantom to simulate registration using the diffusion imaging modality. Finally, I registered the images using the proposed registration algorithm using the multivariate NMI as the similarity measure. For structural and WM only registrations, univariate NMI was used as a similarity measure and the same registration algorithm in terms of the transformation model and optimisation scheme was used in all the registrations. I performed the analysis by repeating the experiment with 100 different deformations. The results over the whole phantom and the white matter area are illustrated in table 5.1. The proposed method improved registration accuracy over the whole phantom and also over the simulated white matter regions. Even though the designed numerical phantom is a simple simulation of the clinical environment of the temporal lobe, it demonstrates that the use of complementary information in a registration scheme can indeed improve registration accuracy.

	Initial	Structural	Tract	Joint
All	1.69(0.22)	1.15(0.17)	1.50(0.22)	1.05(0.17)
WM	1.70(0.09)	1.08(0.07)	0.92(0.08)	0.87(0.06)

Table 5.1: Mean (standard deviation) Euclidean distance errors in voxels. The second column quantifies the initial misalignment. Subsequent columns correspond to the error after registration using the structural information, the WM tract information and the joint information respectively. Errors are computed within the whole phantom (All) (middle row) and the white matter (WM) (bottom row).

5.3.2 Quantitative Validation on Post-Operative Clinical MRI Data

I used data from 20 subjects who had undergone temporal lobe resection for treatment of refractory focal epilepsy. Structural MRI scans, DTI and visual field measurements were acquired before surgery and $3-5$ months following surgery. Significant VFD can be caused by damage to the optic radiation during the intervention. I analysed the pre- and post-operative MRI scans and correlated the visual field deficit, which was determined by a visual field assessment, with the optic radiation resection as predicted by the different registration schemes. Standard clinical sequences were performed on a 3T GE Excite II scanner (General Electric, Milwaukee, WI, USA) including a coronal T1-weighted volumetric fast spoiled gradient echo (SPGR) acquisition with 170 contiguous 1.1 mm thick slices. The field of view was 24 cm, the acquisition matrix size was 256 $\times$ 256 and the reconstructed image resolution was 0.9 mm $\times$ 0.9 mm $\times$ 1.1 mm. DTI data were acquired using a cardiac-triggered single-shot spin-echo planar imaging (EPI) sequence with TE $=$ 73 ms. Sets of 60 contiguous 2.4-mm thick axial slices were obtained covering the whole brain, with diffusion sensitizing gradients applied in each of 52 non-collinear directions [b value of 1200 mm^2 s^{-1} (sigma $=$ 21 ms, delta $=$ 29 ms, using full gradient strength of 40 mT m^{-1})] along with six non-diffusion weighted scans. The field of view was 24 cm, and the acquisition matrix size was 96 $\times$ 96, zero filled to 128 $\times$ 128 during reconstruction giving a reconstructed voxel size of 1.875 mm $\times$ 1.875 mm $\times$ 2.4 mm.

VFD Quantificatiom

Pre- and post-operative visual fields were assessed by Goldmann perimetry and the V/4e isoptre was used for analysis. Due to the high variability observed between Goldmann perimetry sessions (Parrish et al., 1984) and the lack of pre-operative data in some patients, visual field loss was calculated using the areas of the upper quadrants as follows:

$$\textbf{VFD} = 1 - \frac{\text{[area of upper quadrants contralateral to resection]}}{\text{[area of upper quadrants ipsilateral to resection]}}$$

The use of the unaffected upper quadrant ipsilateral to the side of surgery as the reference for each patient allowed for the use of post-operative data alone and eliminated inter-session variability. No significant asymmetry in the upper quadrants was observed on pre-operative Goldmann perimetry, and no deficits within the ipsilateral upper quadrants were observed in post-operative fields.

Optic Radiation Parcellation

The optic radiation parecellation was done according to the established clinical protocol at the Chalfont Centre for Epilepsy in London. All the parcellations were performed by my clinical collaborator Dr. Gavin Winston. The optic radiation was identified in the pre-operative diffusion images by conducting multi-tensor Probabilistic Index of Connectivity (PICo) as implemented in Camino (Cook et al., 2006). A 15 voxel seed region across the base of Meyer's loop was defined with a way point in the lateral wall of the occipital horn of the lateral ventricle and a midline exclusion mask. Tracking from the seed was performed using 50000 Monte Carlo iterations, an angular threshold of 180° and a fractional anisotropy

threshold of 0.1, in order to ensure that the paths detected would not erroneously enter areas of cerebrospinal fluid, and yet had sufficient angular flexibility to allow tracking of Meyer's loop. Finally, a coronal exclusion mask was used to remove artefactual connections to adjacent white matter tracts, such as the fronto-occipital fasciculus, anterior commissure and uncinate fasciculus (Yogarajah et al., 2009). An objective, iterative process was performed to determine the optimum location for this mask whereby the exclusion mask was moved posteriorly until it began to coincide with Meyer's loop, identified by a visible thinning of the estimated trajectory of the optic radiation, typically associated with a reduction in tract volume greater than 10%. A connectivity distribution was generated from each voxel in the seed region and combined into an overall connectivity map representing the maximum observed connection probability to each voxel within the brain from all the voxels within the seed region. For display purposes, the connectivity distributions were thresholded at 5%, representing a compromise between retaining anatomically valid tracts and removing obviously artefactual connections.

I registered the pre-operative dataset to the post-operative dataset using only anatomical images, only diffusion images and using the proposed method. Automatic skull-stripping (Smith, 2002) was performed on the images prior to registration to ensure that all non-brain related tissues were removed from the image. Optic radiation was propagated using the deformation field generated by the three registration schemes. The damage to the optic radiation was quantified by measuring the anteroposterior distance from the anterior part of Meyer's loop in an axial plane to the resection margin measured in millimetres in the T1-weighted MRI image. This is illustrated in figure 5.2. The validation results are shown in table 5.2. Spearman rank order correlation coefficient was used to measure the relationship between the measured visual field deficit and the predicted damage to the optic radiation. It is worth noting that the Spearman correlation coefficient does not assume that both datasets are normally distributed. The predicted damage to the optic radiation when using the proposed registration scheme correlates better with the measured visual field deficit (Spearman correlation coefficient: 0.79, $p = 0.002$) and there is a trend towards higher correlation for the proposed method. Further clinical details about this study can be found in Winston et al. (2011).

Figures 5.3 and 5.4 show the Bland-Altman plots generated when comparing the structural only and FA only image registration schemes with the proposed method. The scatterplot shows the average of the damage to the optic radiation as predicted by the methods under comparison on the horizontal axes and their difference on the vertical axes. The plots show that the proposed method differs from both the structural and FA only methods as a few of the observations are close or outside the range of agreement which was defined as mean bias ± 1.96 standard deviations, which is the 95% limits of agreement assuming that the differences are normally distributed. However, the question of whether this disagreement between the methods is *clinically important* needs to be investigated and is currently being undertaken through a clinical study at NHNN.

5.3.3 Quantitative Validation on Interventional MRI Datasets

Evidence for improvement of patient outcome must be demonstrated before changes to a clinical workflow can be made. For this purpose intra-operative DTI datasets were acquired from twelve subjects

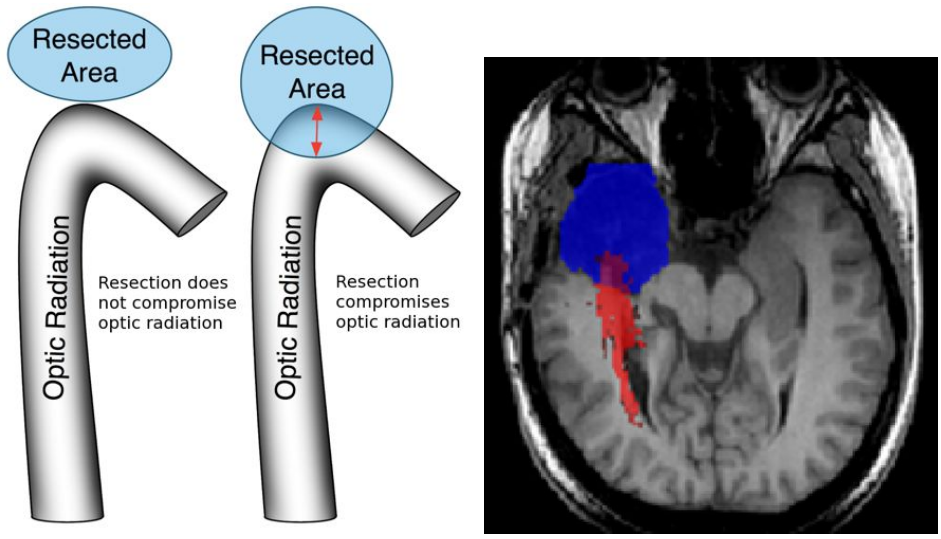

(a) (b)

Figure 5.2: Figure (a) illustrates the quantification of the optic radiation resection. Figure (b) shows a clinical example where the subject suffered visual deficit. The propagated pre-operative optic radiation (red) overlaps with the resected area (blue).

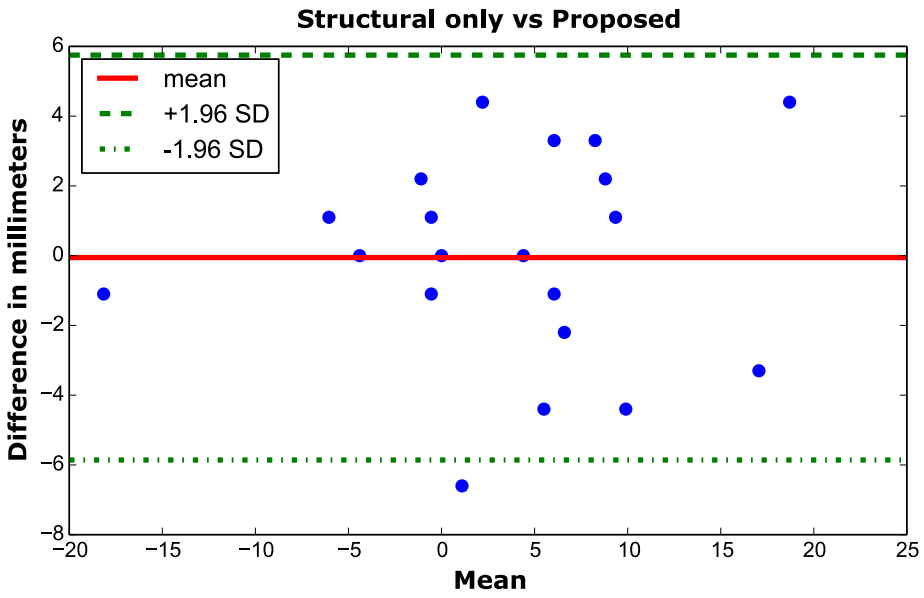

Figure 5.3: Bland Altman assessment indicates that the 95% limits of agreement between the structural only image registration and the proposed method ranged from 5.74 mm to -5.85 mm.

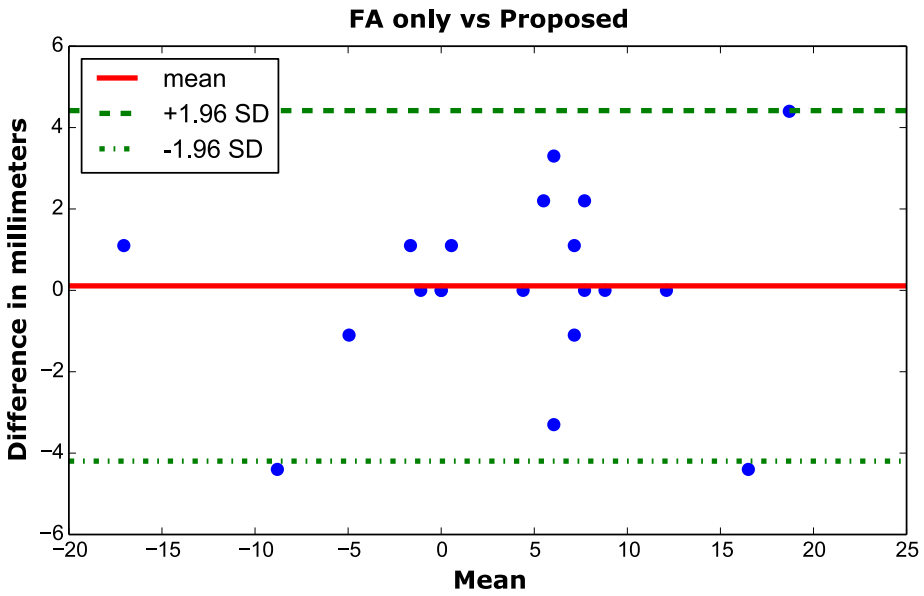

Figure 5.4: Bland Altman assessment indicates that the 95% limits of agreement between the FA only image registration and the proposed method ranged from 4.41 mm to -4.19 mm.

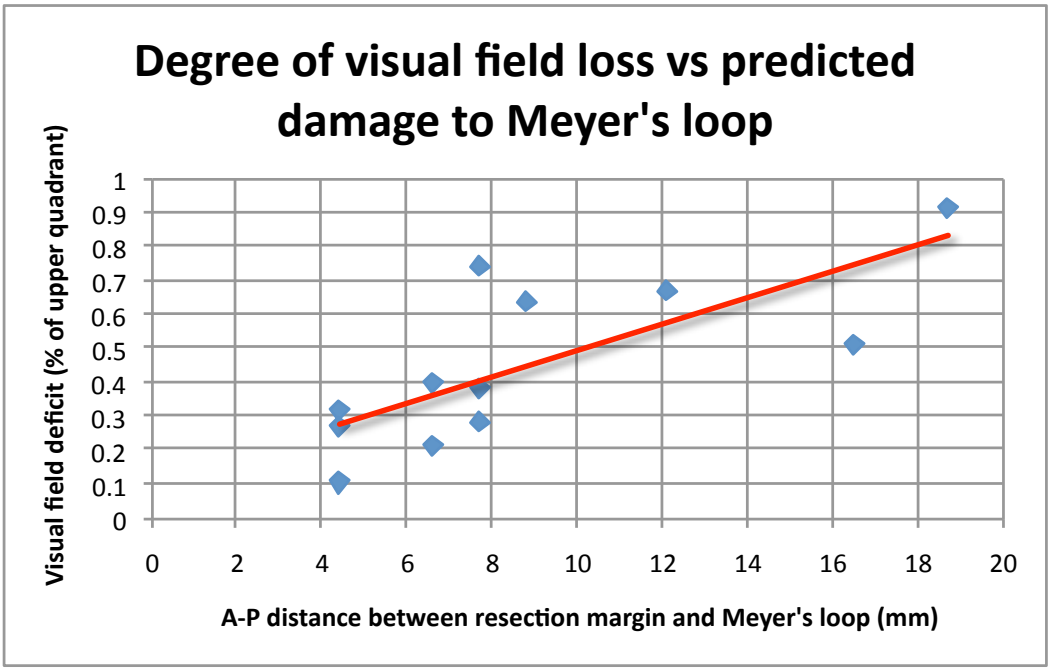

Figure 5.5: Regression line for the predicted damage to Meyer's loop using the proposed registration method and the observed visual field deficit in the 12 patients that suffered from visual field deficits.

Case	Visual Deficit	Structural	FA	Proposed
1	0	−1.1	0	0
2	0	0	1.1	0
3	0	−18.7	−16.5	−17.6
4	0	−5.5	−11	−6.6
5	0	−4.4	−5.5	−4.4
6	0	4.4	0	0
7	0	0	−1.1	−2.2
8	0	0	−1.1	−1.1
9	0.103	−2.2	7.7	4.4
10	0.208	9.9	7.7	6.6
11	0.265	7.7	6.6	4.4
12	0.276	3.3	6.6	7.7
13	0.315	4.4	4.4	4.4
14	0.380	9.9	7.7	7.7
15	0.393	5.5	8.8	6.6
16	0.508	20.9	20.9	16.5
17	0.634	9.9	8.8	8.8
18	0.667	7.7	12.1	12.1
19	0.740	5.5	4.4	7.7
20	0.915	15.4	14.3	18.7
CC		**0.44**	**0.42**	**0.79**

Table 5.2: Spearman correlation coefficient (CC) of the measured visual deficit against the predicted damage to the optic radiation by using the three registration schemes (using structural images, fractional anisotropy (FA) images and using both structural and FA images through the proposed method) for the 12 subjects that suffered visual deficit. Columns 3-5 show the predicted damage (reported in mm by measuring the anteroposterior distance from the anterior part of optic radiation to the resection margin) to the optic radiation. The last row shows the CC of the A-P distance against the visual field assessment scores for subjects with VFD.

(two interventional time-points for each subject for a total of twenty four datasets) undergoing temporal lobe resection for treatment of refractory focal epilepsy. These data sets were assessed retrospectively as part of a formal clinical audit exercise according to the data governance protocols at NHNN. In all the twelve cases pre-operative T1-weighted MR and diffusion weighted MR data were acquired. Pre-operative MR scans were acquired on a 3T MR GE Excite II scanner (General Electric, Waukesha, Milwaukee, WI, USA) and included a T1-weighted coronal volumetric acquisition with a spatial resolution of $0.9 \times 0.9 \times 1.1$ mm. DTI data using 52 gradient directions was acquired using a single-shot spin-echo planar imaging (EPI) sequence with a spatial resolution of $1.9 \times 1.9 \times 2.4$ mm. For all 12 subjects, MRI data acquired during the intervention after the temporal pole resection. The intra-operative protocol included a T1-weighted 3D FLASH sequence with TR = 5.25ms, TE = 2.5ms, flip angle = 15° and have a spatial resolution of $1.1 \times 1.1 \times 1.3$mm. DTI data using 30 gradient directions and a spatial resolution of $2.5 \times 2.5 \times 2.7$ mm were also acquired. The FA images generated from the DTI data were corrected for susceptibility artefacts using the algorithm described in chapter 4.

Challenges With Intra-Operative Tractography

The validation of the pre- to intra-operative image registration results is challenging due to the absence of a suitable gold standard. Tractography techniques use diffusion MRI data acquired on the millimetre scale to infer underlying axonal connectivity on the micrometer scale. An assumption inherent in the dif-

fusion tensor model is that the fitted principle eigenvector represents the orientation of a coherent axonal bundle within each voxel. It has become increasingly apparent that the majority of voxels contain multiple fibre orientations (Jeurissen et al., 2010) and that models that take account of this can better depict tracts in regions of crossing fibres (Behrens et al., 2007). Moreover, the use of deterministic tractography techniques which provide a single estimate of the path at each point can lead to erroneous tracts in the presence of noise. The structure of interest, Meyer's loop of the optic radiation, is a tightly curving structure which lies in close proximity to another white matter bundle, the uncinate fasciculus, and thus poses a particular challenge for tractography algorithms. A multitensor probabilistic tractography algorithm in which up to two independent fibre populations are modelled per voxel was thus employed for the pre- and post-operative data within this paper. This uses a previously validated technique (Yogarajah et al., 2009) and probabilistic tractography has been shown to be superior in its depiction of Meyer's loop to deterministic algorithms (Nilsson, 2010).

In intra-operative datasets, the lower signal-to-noise ratio and fewer gradient directions do not allow the fitting of a multitensor model (Behrens et al., 2007). The lower SNR leads to greater uncertainty within the tractography and consequently greater spread with distance in the tractography results. The poorer spatial resolution also hinders depiction of tightly curved structures as the direction of a tract may change within a voxel violating the assumptions of the tensor model. For these reasons and the fact that tractography algorithms are inherently highly sensitive to the seed region selected (Ciccarelli et al., 2003), a direct comparison of pre-operative and intra-operative tractography results is not appropriate.

Assessment of Registration Accuracy with Intra-Operative MRI

Despite the fact that tractography is not reliable on intra-operative datasets, these datasets still provide information to drive image registration. FA is a diffusion anisotropy measure that is known to be robust to noise (Armitage and Bastin, 2000; Hasan et al., 2004) and, therefore, is a good candidate for use in image registration algorithms. I validated the registration accuracy in the intra-operative setting by performing the registration with the proposed method between the pre- and *intra-operative* datasets and between the *intra-* and post-operative datasets. The deformation fields obtained from these two registrations were composed to produce the deformation field from the pre- to post-operative datasets. This is schematically illustrated in figure 5.6. The pre-operatively delineated tract was propagated using this deformation field and compared with the tract obtained from registration of the pre- to post-operative dataset as in section 5.3.2 for similarity. The propagated tract obtained by direct registration of the pre- and post-operative datasets using the proposed method was shown to correlate best with the observed VFD in section 5.3.2 and is used as the ground truth in this experiment.

In order to assess the consistency of tract propagation, the average distance between the skeleton of the ground truth tract and the closest voxel in the skeleton of the propagated tract was calculated. The skeletons were generated first by placing an initial bounding box around the tract. This bounding box was then skeletonised, with the probability of belonging to the tract as a priority function. Thus, as the skeleton will not be medial in a euclidean sense but rather in a probabilistic sense, thin 1D isthmuses were used as locking points resulting in a 1 voxel thick 6 connected skeleton. As the average distance is

calculated for every point in the skeleton and one skeleton might be longer than the other, the reported values will slightly overestimate the true distance between the skeletons. This overestimation effect can also occur due to the resolution of the image and the fact that the skeleton is not sub-voxel accurate. The true distance will thus be slightly smaller than the reported values.

A visual example is shown in figure 5.7. The mean distance measures and the associated standard deviation are shown in table 5.3. In addition, a correlation analysis as in 5.3.2 was performed by calculating the correlation coefficient between the VFD and the predicted damage to the optic radiation by using the deformation field obtained by the registration scheme described in figure 5.6. They showed a strong correlation (Spearman correlation coefficient: 0.76, $p = 0.003$) for this small cohort.

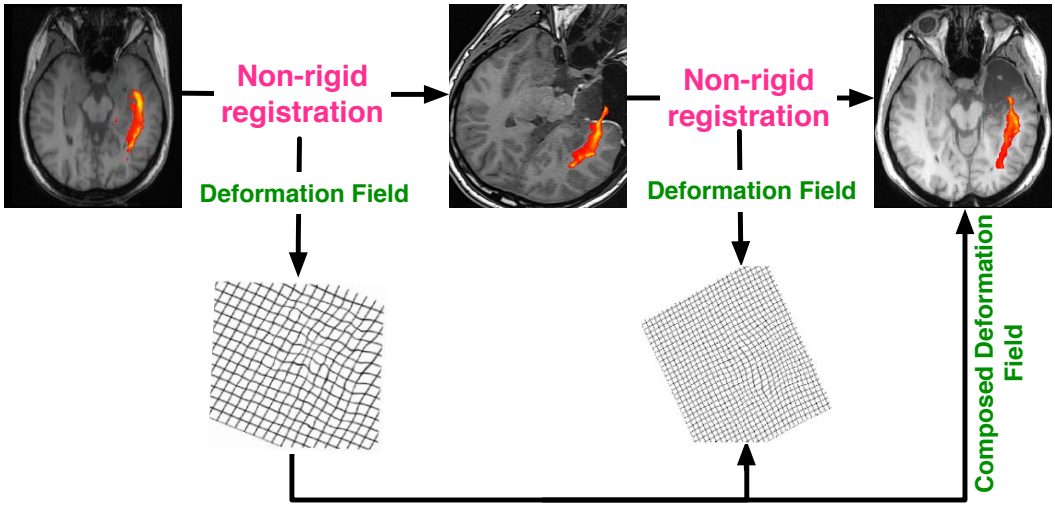

Figure 5.6: Validation of the proposed registration scheme using the intra-operative datasets. The pre-operative images are initially non-rigidly registered to the intra-operatively acquired images using the proposed method. In a second non-rigid registration step, the intra-operative images are registered to the post-operative images. The two deformation fields acquired from the registration steps are composed together to generate the final deformation field. I show that the predicted damage to the optic radiation using this deformation field correlates strongly with the observed VFD.

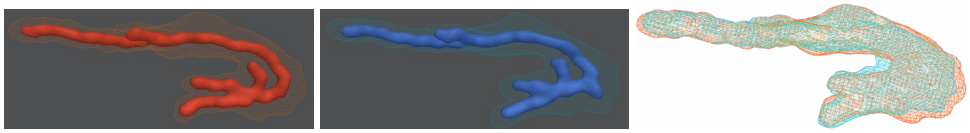

(a) (b) (c)

Figure 5.7: Figure (a) shows a mesh rendering of the optic radiation obtained by directly registering the pre- to the post-operative dataset using the proposed registration scheme. I showed that it correlates best with the observed visual field deficit and is used as the ground truth for validating the intra-operative registration. Figure (b) shows the meshed optic radiation obtained by composition of the deformation fields obtained by registering the pre- to the intra-operative dataset and the intra- to the post-operative dataset. The solid colour in (a) and (b) denote the 1 voxel thick skeleton of the tracts. Figure (c) shows the close overlap of (a) and (b).

5.3.4 Computational Performance Validation

Through the use of the parallel processing capabilities of the GPUs, significant reduction to the computation time was achieved. We used NVIDIA's C2050 Tesla processors for the benchmarks. The mean

Case	Intra-operative time point	VFD	Mean distance(std.)
1	1	10.3	2.37(1.05)
	2		2.43(1.81)
2	1	20.8	1.57(1.96)
	2		1.06(1.31)
3	1	26.5	3.31(2.02)
	2		3.01(1.61)
4	1	27.6	1.53(1.73)
	2		1.06(1.52)
5	1	31.5	1.12(0.92)
	2		1.43(0.98)
6	1	38.0	3.42(2.15)
	2		3.57(2.07)
7	1	50.8	2.82(2.21)
	2		2.93(2.77)
8	1	59.7	1.65(1.29)
	2		1.72(1.18)
9	1	66.7	2.86(2.71)
	2		2.57(2.13)
10	1	73.2	2.16(2.21)
	2		2.11(2.72)
11	1	91.5	1.54(1.05)
	2		1.77(1.09)
12	1	0	1.89(1.86)
	2		2.50(2.56)
CC			**0.76**

Table 5.3: Mean distance and standard deviation (in mm) between the optic radiation skeleton obtained by direct registration of the pre- and post-operative datasets and the optic radiation skeleton obtained by the composition of the deformation fields from registering the pre- to an intra-operative dataset and registering the intra- to the post-operative dataset. For each case, the analysis was carried out using the two available intra-operative time points i.e. there were two intra-operative scan sessions during the surgery. Additionally, the predicted damage to the optic radiation was also measured in a similar manner to 5.3.2. Spearman correlation coefficient (CC) shows that the propagated optic radiation correlates well with the observed VFD even when using the intra-operative datasets for the intermediate registration. Case 12 was excluded from the correlation analysis as the subject did not suffer any VFD.

time for affine registration of the target and source images for the five interventional subjects was 18 seconds. The average time for non-rigid registration using the proposed method is 2 minutes and 55 seconds. In comparison, the mean time for CPU based affine registration is 37 seconds and for the non-rigid registration it is 25 minutes and 54 seconds. In addition, there is an overhead of doing the skull-stripping and generating the FA images for the interventional scans. However, these times are not significant. The current transfer time of the patient from the scanner to the operating table at NHNN is between $7 - 12$ minutes and the proposed registration algorithm is fast enough to cope with this time constraint. GPUs are constantly evolving in design and we expect further increase in computation time through the use of updated hardware.

The following chapter will describe the clinical integration of methods developed in this thesis into the surgical workflow at NHNN, London.

Chapter 6

Clinical Integration

One of my main motivations for pursuing this PhD was to be able to work on a project that might eventually be used in clinical practice to improve patient care. As part of my PhD, I integrated the algorithms described in chapter (4-5) into the surgical workflow at NHNN, London. This work was done in close collaboration with my collaborators at NHNN: Dr. Mark White and Dr. Laura Mancini. This chapter provides an overview on the underlying workflow architecture. In order to ensure that the workflow is usable, it is automatic and requires minimal user interaction. The workflow is schematically depicted in figure (6.1). Certain tasks like extraction of the brain mask from the intraoperative T1-weighted image are not highlighted for reasons of brevity. For every surgical procedure, this workflow is transformed into a format that can be easily parsed by a software program through the use of Javascript Object Notation (JSON). JSON[1] is a lightweight, human readable, data interchange format and has wide support with software libraries available for parsing JSON input for all major computer languages. The JSON specification file for the neurosurgical workflow is automatically generated and the only user interaction step is to specify the patient identifier. The requisite pre and intraoperative images are retrieved from the hospital PACS system made avalable for processing. Each JSON block corresponds to a image analysis process that needs to be executed. For processes where protocol specific information is required, for example the table offset position for the gradient non-linearity correction, they are extracted from the image metadata.

[1] http://www.json.org

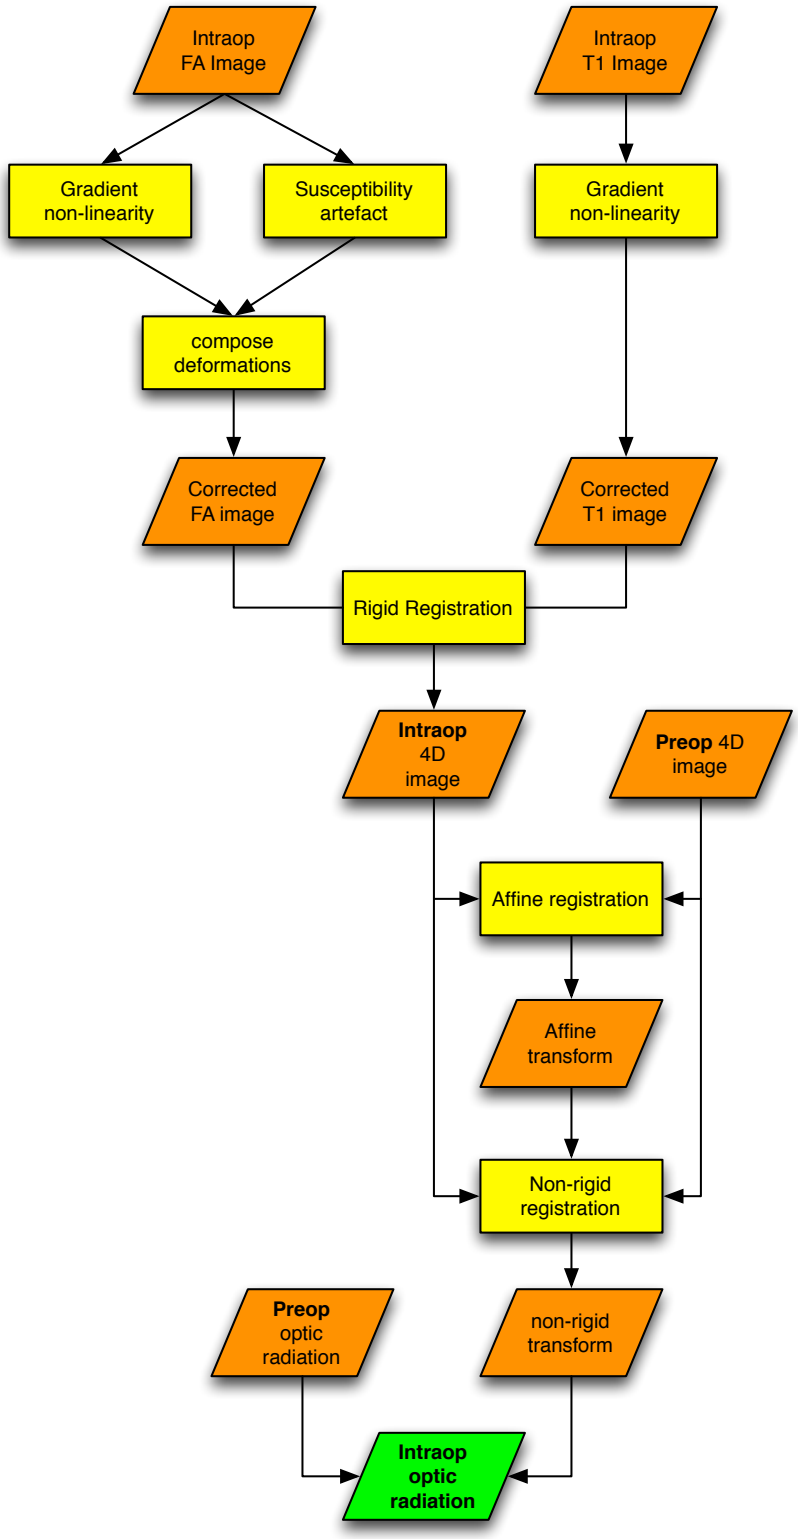

Figure 6.1: Workflow depicting the various inputs/outputs and the processes for the interventional MRI image analysis workflow. The various inputs and outputs are highlighted in orange while the processes are highlighted in yellow. The final output, which is the optic radiation in the intraoperative space, is highlighted in green. Some of the processes like the brain mask generation are not highlighted in this workflow for reasons of brevity. This workflow is dynamically generated for every surgery through a data interchange format called JSON. The JSON input is parsed and automatically executed within 7-8 minutes.

The following shows an example JSON files that represent the flowchart of figure (6.1) describing the full intraoperative image analysis pipeline for a specific subject. The “comments” section in each JSON block documents the task that will be performed making documentation part of the data generation process. This dynamically generated JSON file is parsed by a Python[2] based software program which executes the workflow serially. The whole workflow can be executed in 7-8 minutes. The resulting propagation of the optic radiation in the intraoperative space is examined by a radiologist before injecting it into the neuronavigation system for surgical guidance.

```
{
"correct_non_linearity": [
    {
        "input": "/input/i1_t1.nii",
        "output_def": "output/gradunwarp_t1_def.nii",
        "coeffs": "input/coeff.grad",
        "off_x": "0",
        "off_y": "0",
        "off_z": "24",
        "comment": "Correct the T1 image for gradient non-linearities."
    },

    {
        "input": "/input/i1_fa.nii",
        "output_def": "/output/gradunwarp_fa_def.nii",
        "coeffs": "/input/coeff.grad",
        "off_x": "0",
        "off_y": "0",
        "off_z": "44",
        "comment": "Correct the FA image for gradient non-linearities."
    },

    {
        "input": "/input/i1_fm_pha.nii",
        "output_def": "/output/gradunwarp_phase_def.nii",
        "coeffs": "/input/coeff.grad",
        "off_x": "0",
        "off_y": "0",
        "off_z": "44",
        "comment": "Correct the phase image for gradient non-linearities."
```

[2] http://www.python.org

```
    }
],

"resample_after_gradwarp": [
    {
        "input": "/input/i1_t1.nii",
        "output": "/output/i1_t1_gradwarp_corrected.nii",
        "def": "/output/gradunwarp_t1_def.nii",
        "comment": "Apply the gradwarp correction to the T1 image."
    },

    {
        "input": "/input/i1_fm_pha.nii",
        "output": "/output/i1_fm_pha_gradwarp_corrected.nii",
        "def": "/output/gradunwarp_phase_def.nii",
        "comment": "Apply the gradwarp correction to the phase image."
    }
],

"susceptibility": [
    {
        "input": "/input/i1_fa.nii",
        "output_def": "/output/susceptibility_def.nii",
        "phase_diff_image": "/input/i1_fm_pha.nii",
        "mag": "/input/i1_fm_mag1.nii",
        "phase_encode_direction": "+y",
        "read_out_time": "35.52",
        "echo_time_difference": "4.76",
        "comment": "Correct the FA image for susceptibility artefacts."
    }
],

"compose_grad_warp_susceptibility_deformations": [
    {
        "input": "/input/i1_fa.nii",
        "gradwarp_def": "/output/gradunwarp_fa_def.nii",
        "susceptibility_def": "/output/susceptibility_def.nii",
        "output_def": "/output/gradwarp_susceptibility_correction_def.nii",
```

```
        "comment": "Compose the deformation fields from correction steps".
    }
],


"resample_diffusion_to_anatomical": [
    {
        "target": "/output/i1_t1_gradwarp_corrected.nii",
        "source": "/input/i1_fa.nii",
        "output": "/output/i1_fa_to_t1.nii",
        "nr_def": "/output/gradwarp_susceptibility_correction_def.nii",
        "comment": "Resample the FA image into the T1 space".
    }
],

"reorient_before_reg": [
    {
        "input": "/output/i1_t1_gradwarp_corrected.nii",
        "surgery_side": "r",
        "output": "/output/i1_t1_gradwarp_corrected_reoriented.nii",
        "comment": "Reorient the T1 image to a standard axes."
    },

    {
        "input": "/output/i1_fa_to_t1.nii",
        "surgery_side": "r",
        "output": "/output/i1_fa_to_t1_reoriented.nii",
        "comment": "Reorient the FA image to a standard axes."
    }
],

"brain_mask": [
    {
        "input": "/output/i1_t1_gradwarp_corrected_reoriented.nii",
        "output": "/output/i1_t1_gradwarp_corrected_reoriented_mask.nii",
        "comment": "Generate intraop T1 brain mask."
    },
```

```
    {
        "input": "/input/preop_t1.nii",
        "output": "/output/preop_t1_mask.nii",
        "comment": "Generate preop T1 brain mask."
    }
],

"rigid_registration": [
    {
        "target": "/input/preop_t1.nii",
        "source": "/input/preop_b0.nii",
        "mask": "/output/preop_t1_mask.nii",
        "output": "/output/pre_b0_to_t1.nii",
        "transform": "/output/pre_b0_to_t1.txt",
        "comment": "Rigidly register preop T1 to the B0 image."
    }
],

"resample_rigid": [
    {
        "target": "/input/preop_t1.nii",
        "source": "/input/preop_fa.nii",
        "transform": "/output/pre_b0_to_t1.txt",
        "output": "/output/preop_fa_to_t1.nii",
        "comment": "Resample the preop FA to t1 space using the rigid transformation."
    }
],

"affine_registration": [
    {
        "target": "/output/i1_t1_gradwarp_corrected_reoriented.nii",
        "source": "/input/preop_t1.nii",
        "mask": "/output/i1_t1_gradwarp_corrected_reoriented_mask.nii",
        "output": "/output/pre_t1_to_i1_t1_affine.nii",
        "transform": "/output/pre_t1_to_i1_t1_affine.txt",
        "comment": "Affinely register the preop T1 to intraop T1."
    }
],
```

```
"merge": [
    {
        "image_1": "/output/i1_t1_gradwarp_corrected_reoriented.nii",
        "image_2": "/output/i1_fa_to_t1_reoriented.nii",
        "output": "/output/i1_4D.nii",
        "comment": "Generate a 4D image to be used for non-rigid registration."
    },

    {
        "image_1": "/input/preop_t1.nii",
        "image_2": "/output/preop_fa_to_t1.nii",
        "output": "/output/preop_4D.nii",
        "comment": "Generate a 4D image to be used for non-rigid registration."
    }
],

"mc_nonrigid_registration": [
    {
        "target": "/output/i1_4D.nii",
        "source": "/output/preop_4D.nii",
        "mask": "/output/i1_t1_gradwarp_corrected_reoriented_mask.nii",
        "aff": "/output/pre_t1_to_i1_t1_affine.txt",
        "output": "/output/pre_t1_to_i1_t1_nrr.nii",
        "transform": "/output/pre_t1_to_i1_t1_cpp.nii",
        "comment": "Perform the multichannel non-rigid registration."
    }
],

"propagate": [
    {
        "target": "/output/i1_t1_gradwarp_corrected_reoriented.nii",
        "source": "/input/pre_or.nii",
        "transform": "/output/pre_t1_to_i1_t1_cpp.nii",
        "output": "/output/preop_or_to_i1_t1.nii",
        "comment": "Propagate the preoperative OR in the intraop space."
    }
],
```

```
"unorient_after_reg": [
    {
        "input": "/output/preop_or_to_i1_t1.nii",
        "surgery_side": "r",
        "output": "/output/preop_or_to_i1_t1_unoriented.nii",
        "comment": "Orient the image into the original intraop T1 image."
    }
]
}
```

Chapter 7

Clinical Findings

In the previous chapters, I described the development of computational techniques that estimates the brain shift during a neurosurgical procedure. The estimated brain shift was used to propagate the pre-operative tractography onto the intraoperative images and localise it during surgery. I showed that this technique could accurately predict the degree of VFD and it could be used within the time constraints of a neurosurgical procedure. I suggested that display in a neuronavigation suite of the location of the optic radiation would be useful in avoiding surgical damage.

In this chapter, assessment of whether the display of preoperative tractography during ATLR can reduce the severity of VFD and increase the proportion of patients that can drive is performed. The patients are also followed up post surgery to assess whether it has an affect on post-operative seizure outcome. Secondly, assessment of whether the correction of brain shift during surgery using iMRI provides additional benefit is also performed. This work was lead by my clinical collaborator Dr. Gavin Winston.

7.1 Methods

7.1.1 Subjects

21 patients (age range, 23-63 years; median, 36 years; 8 male) with medically refractory TLE undergoing ATLR at NHNN, London were studied. All patients had structural MRI scans performed at 3T, video electroencephalographic (EEG) telemetry, neuropsychology, neuropsychiatry, and if necessary intracranial EEG recordings prior to surgery. Structural MRI scans, diffusion tensor imaging (DTI) and visual fields were acquired before surgery and 3 months following surgery (range 70-145 days). The study was approved by the National Hospital for Neurology and Neurosurgery and the Institute of Neurology Joint Research Ethics Committee, and informed written consent was obtained from all subjects.

7.1.2 Comparison Cohort

For comparison to previous clinical practice, a cohort of patients who underwent the same assessment and ATLR by the same neurosurgeon in a conventional operating theatre without tractography-based image guidance between 2009 and 2012 was selected, comprising 44 patients (age range, 17-68 years; median, 39 years; 17 male; 21 left, 23 right ATLR).

7.1.3 Optic Radiation Tractography

Preoperative and postoperative MRI studies were performed on a 3T GE Signa HDx scanner (General Electric, Waukesha, Milwaukee, WI) as previously described. Tractography of the optic radiation was performed using the multi-tensor probabilistic index of connectivity model (19) in the Camino toolkit (20). Tractography data were corrected for image distortion due to gradient non-linearities and magnetic susceptibility artefacts as described in chapter 4.

7.1.4 Surgery and Intraoperative Imaging

All patients underwent ATLR in the iMRI suite at NHNN. During surgery, the neuronavigation system provides real-time tracking of surgical markers and tools and visualization facilities. The operating room is equipped with a confocal surgical microscope that supports the injection of colour overlays from the neuronavigation system. The location of the microscopes focal point is tracked using the navigation system and an array of four infra-red reflectors mounted on the microscopes optical head. Before surgery, anatomical scans were performed for use with the neuronavigation system. Repeat anatomical and diffusion scans were acquired after the craniotomy and dura opening (timepoint 1) to provide guidance in entering the ventricle which was manually delineated for display by a radiologist and at the end of the surgery (timepoint 2) to confirm adequate resection. This is highlighted schematically in figure (7.1).

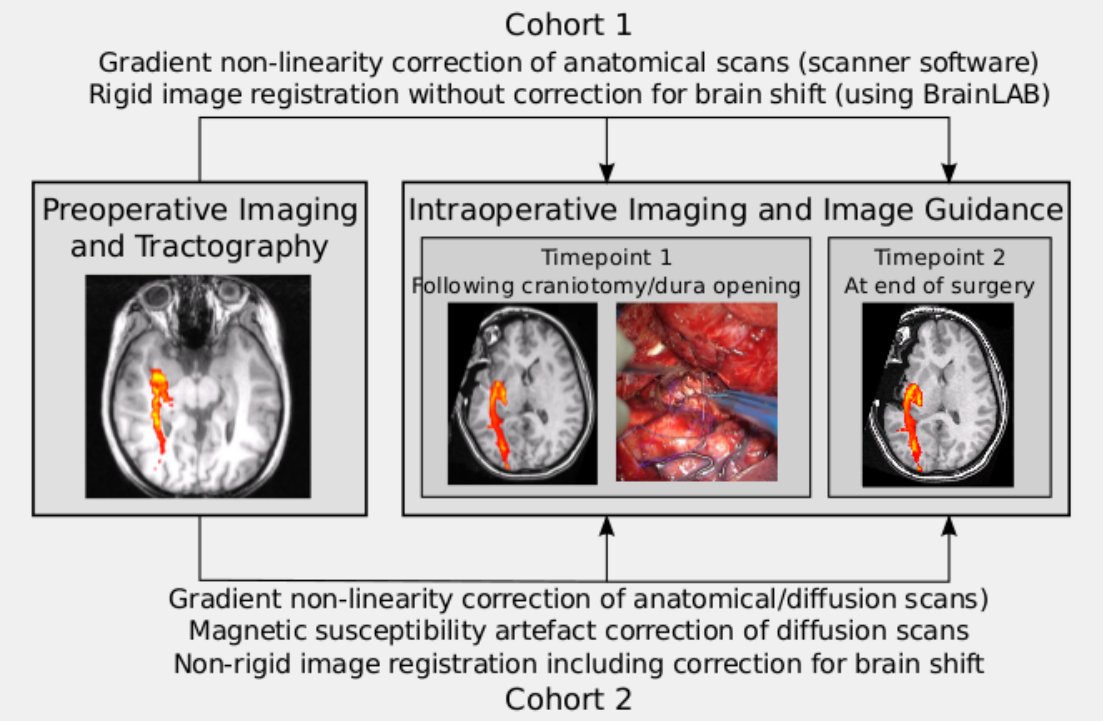

Figure 7.1: Image analysis workflow for the two cohorts. For cohort 1, gradient non-linearity correction was applied and rigid registration was performed to propagate the Optic Radiation in the intraoperative space. For cohort 2, gradient non-linearity correction, susceptibility artefact correction and brain shift correction was applied to propagate the preoperatively segmented Optic Radiation. Image courtesy of Dr. Gavin Winston.

In the first cohort of patients (9 subjects), preoperative imaging including tractography of the optic radiation was transferred to the neuronavigation system and registered to intraoperative images using registration provided by the BrainLAB neuronavigation software. This performs only a rigid transformation which does not correct for brain shift. In the second cohort of patients (12 subjects), preoperative and intraoperative images were processed with the workflow that included gradient non-linearity and susceptibility correction followed by full non-linear registration of the combined preoperative T1-weighted image and FA map to the intra-operative imaging as described in chapter (6). The corrected images were transferred to the neuronavigation system for display. Processing was performed using graphical processing units to ensure the entire procedure could be performed quickly enough not to delay surgery.

The outline of the optic radiation was projected onto the navigation display and the operating microscope display. In cohort 1, additional error margins of 1.5mm in the anatomical antero-posterior direction and 1.5mm isotropic were added to account for the lack of compensation for susceptibility artefacts and potential brain shift respectively.

7.1.5 Primary Outcome: Visual Fields

Pre- and postoperative visual fields were assessed using Goldmann perimetry. To quantify the VFD, postoperative visual fields were scanned and the areas enclosed by the V4e and I4e isopters in each upper quadrant (UQ) were determined. Visual field loss for each isopters was calculated as described in chapter 5 and the mean of the two figures was taken. The use of a single timepoint eliminates the high variability observed between Goldmann perimetry sessions (Parrish et al., 1984).

The number of patients not permitted to drive due to the VFD was determined in accordance with UK Driver and Vehicle Licensing Agency regulations (25) with additional binocular Esterman perimetry if necessary. UK regulations are based on EU Directive 2009/112/EC that requires a horizontal visual field of at least 120 degrees (at least 50 degrees left and right) and 20 degrees up and down with no deficits in the central 20 degrees.

7.1.6 Statistical Analysis

Performing the Shapiro-Wilks normality test on the VFD and degree of hippocampal resection showed that they do not come from a normally distributed population. Hence, non-parameteric Mann-Whitney U or independent-samples Kruskal-Wallis tests were used to detect any difference in the distribution between groups. In contrast, the Shapiro-Wilks test showed that the observed brain shifts were normally distributed.

7.2 Results

7.2.1 Visual Field Deficits

None of the 21 patients undergoing surgery with iMRI guidance developed a VFD that precluded driving. The VFD were 0-41.7% of the contralateral superior quadrant (median 17.9%, IQR 28.0%) in cohort 1, 0-49.2% (median 9.2%, IQR 30.5%) in cohort 2 and 0-49.2% (median 14.5%, IQR 27.5%) overall.

Five patients in the historical cohort had equivocal Goldmann perimetry but declined Estermann as

they did not wish to drive. Of the remaining patients, 5/39 (12.8%) failed to meet DVLA criteria as a result of surgery. The VFD were 0-90.9% of the contralateral superior quadrant (median 24.0%, IQR 32.6%).

The distribution of VFD from those with iMRI guidance (cohorts 1 and 2 combined) was significantly different from those without iMRI guidance (independent-samples Mann-Whitney U test p=0.043) as shown in figure (7.2). The difference was not significant between the historical controls and each iMRI guided cohort individually. In cohort 2, two patients had previous surgery with one having a pre-existing minor VFD that did not preclude driving. Exclusion of these patients did not affect the significant difference but the median VFD fell to 3.4% (IQR 36.0%) in cohort 2 and to 11.0% (IQR 32.3%) in the iMRI-guided cohort overall.

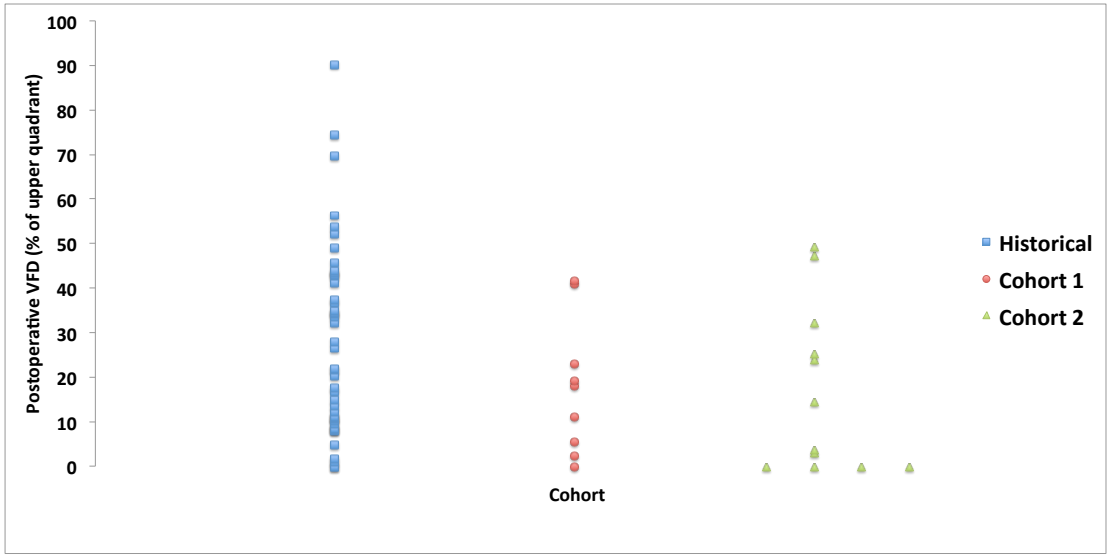

Figure 7.2: Image analysis workflow for the two cohorts. For cohort 1, gradient non-linearity correction was applied and rigid registration was performed to propagate the Optic Radiation in the intraoperative space. For cohort 2, gradient non-linearity correction, susceptibility artefact correction and brain shift correction was applied to propagate the preoperatively segmented Optic Radiation.

7.2.2 Seizure Outcome

At 3 months, 89% of patients in cohort 1, 92% in cohort 2 and 91% in the historical cohort had a good outcome. At 12 months, 80% in cohort 1 and 83% in the historical cohort had a good outcome. The seizure outcome results for cohort 2 were not available at the time of writing.

Chapter 8

Open Software Effort

Open source is a development approach that promotes transparency and promises more quality, reliability and flexibility in the production of software (Wheeler, 2005). Due to this open nature, most licenses allow anyone to contribute, understand, refactor and reuse the code with no restrictions. As a supporter of this approach, the code developed during my PhD is available under a Berkeley Software Distribution (BSD) license.

With a BSD license, redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code or binaries must retain all the copyright notices, the list of conditions and a disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the organization nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

8.1 NiftyReg

NiftyReg, part of the NifTK suite of software developed at University College London, is an image registration framework developed by Dr. Marc Modat. The software implements the fast free form deformation algorithm as described in Modat (2012). The software implements an efficient implementation of the free form deformation algorithm Rueckert et al. (1999) and has been widely used by the medical image analysis community. The code can be freely downloaded and used under the BSD licence from the website `http://sourceforge.net/projects/niftyreg/`

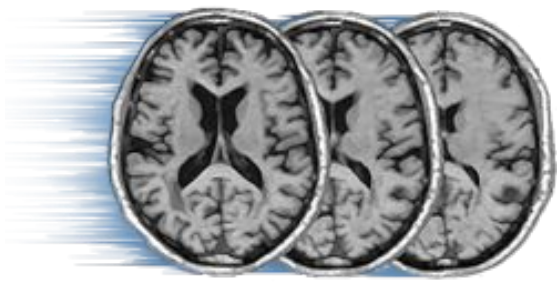

Figure 8.1: Logo for Niftyreg: a registration framework developed at University College London by Dr. Marc Modat.

The image registration algorithm described in chapter (5) is implemented in the NiftyReg framework. The input target and source images can be specified as 4-dimensional nifti images. The software also allows the specification of the number of bins to use for the 4-dimensional joint histogram. An example command line looks as follows:

```
reg_f3d -target <target_4D_image> -source <source_4D_image>
        -tmask <target_mask_image> -aff <affine_transformation_file>
        -result <output_4D_Image> -cpp <output_cpp_file>
        -tbn 0 64 -tbn 1 32 -sbn 0 64 -sbn 1 32 -be 0.01 -gpu
```

where the description of the parameters are as follows:

- reg_f3d: The NiftyReg executable to perform non-rigid registration.
- target: The **input** parameter to specify the *target* image. Specified as a 4dimensional image (target_4D_Image).
- source: The **input** parameter to specify the *source* image. Specified as a 4dimensional image (source_4D_image).
- tmask: The **input** binary mask in the space of the target image (target_mask_image). Specifies the area of overlap between the images.
- aff: The **input** affine transformation. Specified as a text file containing the 4×4 transformation matrix (affine_transformation_file).
- result: The **output** registered source image in the space of the target image (output_4D_image).
- cpp: The **output** control point position file which specifies the estimated deformation (output_cpp_file).
- tbn: The **input** parameter to specify the bin size for computation of joint histogram for the target images (can be specified for both the images).
- sbn: The **input** parameter to specify the bin size for computation of joint histogram for the source images (can be specified for both the images).
- be: The **input** parameter to specify the weight of the bending energy term.
- gpu: The **input** parameter to specify to use GPUs to do the registration.

Chapter 9

Discussion and Conclusion

This thesis presented a clinical workflow and novel algorithms to accurately estimate brain shift and propagate preoperatively segmented optic radiation onto intraoperative images. The algorithms presented in this thesis were specifically developed to work in the intraoperative setting and execute well within the time constraints of a typical neurosurgical procedure. This work was applied on a cohort of patients undergoing ATLR for surgical management of TLE. However, the techniques presented in this thesis can be applied to a much broader spectrum of neurosurgical procedures than ATLR.

In chapter 4, I tackled the problem of correction of susceptibility artefacts in EPI images acquired during a neurosurgical procedure. The susceptibility artefact correction method was based around a novel phase unwrapping framework which formulated the unwrapping as a global optimisation problem. I presented a graph cuts based algorithm which could solve this in a very efficient manner. This algorithm also has the added advantage of being able to calculate the uncertainty associated with the estimated field map without adding significantly to the computation time. The uncertainty information was exploited to selectively improve deformation estimation in regions with high uncertainty using a non-rigid registration algorithm.

The phase unwrapping algorithm developed in chapter 4 has the potential to be used in applications outside the realm of neurosurgery. Accurate measurement of phase is critical in various other contexts in MRI like flow imaging and susceptibility weighted imaging (SWI). SWI is an exciting MRI imaging modality that use the phase images to generate tissue contrast which is inherently different from conventional structural MR imaging. SWI exploits the magnetic susceptibility differences between various tissues and the phase images generated from SWI are useful for detection of cerebral microbleeds in patients with traumatic brain injuries (Haacke et al., 2009). SWI requires long echo times and suffers from severe phase wraps especially in regions of sharp tissue susceptibility differences. Hence, unwrapping of the SWI phase images is essential for their meaningful analysis. Figure 9.1 shows an example where the proposed phase unwrapping algorithm successfully unwrapped an SWI phase image, while maintaining detail about the microbleed locations as highlighted. The unwrapping was performed in less than 15 minutes on a standard computer for a 3D image of size $512 \times 512 \times 320$. The same image took around 2 days to unwrap using PRELUDE.

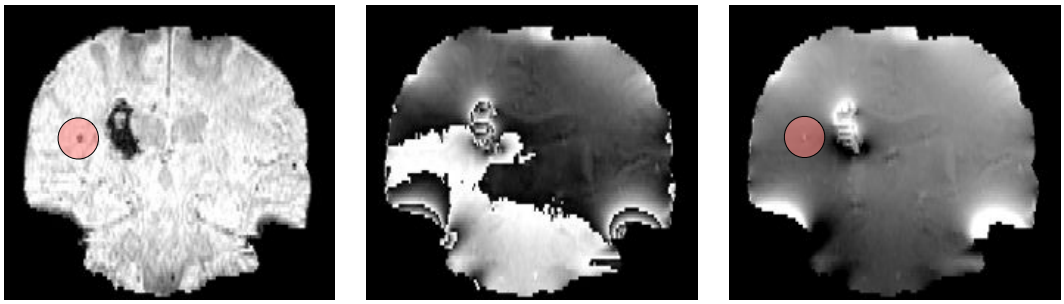

a *b* *c*

Figure 9.1: Phase unwrapping on SWI phase image for a subject with radiation therapy induced microbleeds. (a) shows the magnitude image. (b) is the wrapped phase image and (c) is the unwrapped phase image obtained with the proposed phase unwrapping algorithm. The algorithm successfully recovers the true phase while preserving information about the microbleeds in the phase image as highlighted.

The susceptibility correction method presented in this thesis is, however, not without limitations. A limitation is that I am restricted to using convex pairwise energy function during the graph creation in the phase unwrapping step. Ideally, I would like to employ functions which can handle the discontinuities at sharp tissue boundaries better than the sum of squared difference energy function used here. Another limitation of the proposed work is the use of min-marginals to estimate the confidence in the phase unwrapping solution. However, min-marginals are not true marginal probabilities and can diverge quite significantly from the true probabilities in the worst case scenario (Tarlow and Adams, 2012).

In chapter 5, I presented a brain shift estimation algorithm which utilises information from both structural and diffusion MR images. Diffusion MR images have not been previously used in estimation of brain shift. However, accurate localisation of critical white matter tracts is of surgical interest and with the growing capabilities of the iMRI scanners, multimodal MR images can now be acquired during neurosurgery. Its use on a limited number of clinical data showed that integrating complimentary information from both these modalities in the proposed brain shift estimation algorithm enabled more accurate localisation of the optic radiation. The algorithm was accelerated using GPUs and this enabled it to be used during neurosurgery. However, we cannot draw very firm conclusions from the work presented in this thesis. The initial validation has been done on a small number of subjects and further clinical studies are needed to quantify any potential gains that can arise from the proposed methods. Another issue to note is that we did not perform any intra- and inter-rater variability analysis for clinical tasks like the parcellation of the optic radiation or manually computing the overlap between the resection boundary and the optic radiation.

In chapter 6, I presented a clinical workflow that integrates the algorithms developed in this thesis into the surgical workflow at NHNN, London. This workflow is now used routinely during surgical procedures and will enable further insight into some of the issues mentioned already.

In chapter 7, the clinical findings were presented where the effects of using iMRI (with and without brain shift correction) on VFD and seizure outcome was compared against a historical cohort which did not use iMRI. The initial study indicates significant reduction of VFD in the iMRI cohort. However,

brain shift correction did not confer additional benefit. This could be because the observed brain shift of the anterior tip of the lateral ventricle after craniotomy was not severe (maximum 2.8mm, mean 2.0mm). However, the proposed methods can be applied to a diverse range of neurosurgical procedures like surgical management of low-grade glioblastomas and future studies with larger cohort sizes and diverse procedures could reveal potential benefits arising out of its application.

9.1 Future Research Directions

Notwithstanding the improvements shown in this thesis, there are still many derived problems and associated conceptual ideas that are worth pursuing, either by myself or anyone in the medical imaging community. This section describes my thoughts on a few of the many possible future developments in this field.

Majority of the image registration algorithms, including the ones proposed in this thesis are based on computing the most likely deformation field which is a point estimate and does not provide any confidence measure associated with it. This confidence estimate along with the most likely deformation could be useful in informing surgical decisions. This problem was tackled by Risholm et al. (2013) where they characterised the full posterior distribution over the space of deformations using Markov Chain Monte Carlo (MCMC) sampling methods. However, this method is computationally very expensive and completely infeasible for use in the neurosurgical setting. To reduce the number of parameters to estimate, the deformation was parameterised using a very coarse grid, which would increase the registration uncertainty naturally due to partial volume effects. I feel that significant strides can be made in this area using approximate inference method like expectation propagation (Minka, 2001), for example. Exciting developments in the field of MCMC sampling methods like the Hamiltonian Monte Carlo methods are also worth investing in the interest of speeding up computation time. Algorithms using parallel samplers using GPUs or large multi-core systems could also further reduce computation times.

Another related improvement can be made in the area of uncertainty estimation in the phase unwrapping algorithm. As mentioned before, max-marginals are not true marginal probabilities. Approximate inference methods could also have a part to play here to generate a confidence map using more principled methods than what is proposed in this thesis, while still being computationally efficient.

One of the findings coming out of this work is that surgical guidance using iMRI results in a decrease in VFD without adversely affecting seizure outcome. However, a major disadvantage of iMRI is its limited availability due to the high costs associated with it. However, the information generated by the iMRI system could be used to learn a model of the brain shift for particular kinds of surgery. This model can then be potentially used in conjunction with sparse and low cost imaging techniques like ultrasound, for example, to estimate brain shift in the absense of iMRI. This, I believe, would take the utility of this work beyond the limited iMRI centres and truly make a difference.

References

Abdul-Rahman, H., Arevalillo-Herráez, M., Gdeisat, M., Burton, D., Lalor, M., Lilley, F., Moore, C., Sheltraw, D., Qudeisat, M., Aug 2009. Robust three-dimensional best-path phase-unwrapping algorithm that avoids singularity loops. Applied Optics 48 (23), 4582–4596.

An, L., Xiang, Q.-S., Chavez, S., 2000. A fast implementation of the minimum spanning tree method for phase unwrapping. IEEE Transactions on Medical Imaging 19 (8), 805–808.

Andersson, J. L. R., Skare, S., Ashburner, J., Oct. 2003. How to correct susceptibility distortions in spin-echo echo-planar images: Application to diffusion tensor imaging. NeuroImage 20 (2), 870–888.

Andrea, G., Angelini, A., Romano, A., Di Lauro, A., Sessa, G., Bozzao, A., Ferrante, L., 2012. Intraoperative DTI and brain mapping for surgery of neoplasm of the motor cortex and the corticospinal tract: our protocol and series in brainsuite. Neurosurgical Review 35, 401–412.

Arbel, T., Morandi, X., Comeau, R., Louis Collins, D., 2001. Automatic non-linear MRI-ultrasound registration for the correction of intra-operative brain deformations. In: Medical Image Computing and Computer-Assisted Intervention MICCAI 2001. Vol. 2208. pp. 913–922.

Archip, N., Clatz, O., Whalen, S., Kacher, D., Fedorov, A., Kot, A., Chrisochoides, N., Jolesz, F., Golby, A., Black, P. M., Warfield, S. K., 2007. Non-rigid alignment of pre-operative MRI, fMRI, and DT-MRI with intra-operative MRI for enhanced visualization and navigation in image-guided neurosurgery. NeuroImage 35 (2), 609 – 624.

Armitage, P., Bastin, M., 2000. Selecting an appropriate anisotropy index for displaying diffusion tensor imaging data with improved contrast and sensitivity. Magnetic Resonance in Medicine 44 (1), 117–121.

Audette, M. A., Siddiqi, K., Ferrie, F. P., Peters, T. M., 2003. An integrated range-sensing, segmentation and registration framework for the characterization of intra-surgical brain deformations in image-guided surgery. Computer Vision and Image Understanding 89 (2-3), 226–251.

Avants, B., Duda, J., Zhang, H., Gee, J., 2007. Multivariate normalization with symmetric diffeomorphisms for multivariate studies. In: Medical Image Computing and Computer-Assisted Intervention MICCAI 2007. pp. 359–366.

Beg, M. F., Miller, M. I., Trouv, A., Younes, L., 2005. Computing Large Deformation Metric Mappings via Geodesic Flows of Diffeomorphisms. International Journal of Computer Vision 61, 139–157.

Behrens, T. E., Johansen-Berg, H., Jbabdi, S., Rushworth, M. F., Woolrich, M. W., 2007. Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? Neuroimage 34 (1), 144–55.

Bellman, R. E., 2003. Dynamic Programming.

Besag, J., 1974. Spatial Interaction and the Statistical Analysis of Lattice Systems. Journal of the Royal Statistical Society. Series B (Methodological) 36 (2), 192–236.

Besag, J., 1986. On the Statistical Analysis of Dirty Pictures. Journal of the Royal Statistical Society. Series B (Methodological) 48 (3), 259–302.

Black, P., Moriarty, T., Alexander, E., Stieg, P., Woodard, E., Gleason, P., Martin, C., Kikinis, R., Schwartz, R., Jolesz, F., Oct. 1997. Development and implementation of intraoperative magnetic resonance imaging and its neurosurgical applications. Neurosurgery 41 (4), 831–842.

Bloch, F., Oct 1946. Nuclear induction. Phys. Rev. 70, 460–474.

Boykov, Y., Kolmogorov, V., 2004. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. IEEE Transactions on Pattern Analysis and Machine Intelligence 26 (9), 1124–1137.

Boykov, Y., Veksler, O., Zabih, R., 1998. Markov random fields with efficient approximations. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition.

Boykov, Y., Veksler, O., Zabih, R., 2001. Fast approximate energy minimization via graph cuts. IEEE Transactions on Pattern Analysis and Machine Intelligence 23 (11), 1222–1239.

Cardoso, M., Winston, G., Modat, M., Keihaninejad, S., Duncan, J., Ourselin, S., 2012. Geodesic shape-based averaging. In: Medical Image Computing and Computer-Assisted Intervention MICCAI 2012. pp. 26–33.

Carter, T., Sermesant, M., Cash, D., Barratt, D., Tanner, C., Hawkes, D., 2005. Application of soft tissue modelling to image-guided surgery. Medical Engineering & Physics 27 (10), 893–909.

Chang, H., Fitzpatrick, J. M., 1992. A technique for accurate magnetic resonance imaging in the presence of field inhomogeneities. IEEE Transactions on Medical Imaging 11 (3), 319–329.

Chavez, S., Xiang, Q. S., An, L., 2002. Understanding phase maps in MRI: a new cutline phase unwrapping method. IEEE Transactions on Medical Imaging 21 (8), 966–977.

Chen, I., Coffey, A., Ding, S., Dumpuri, P., Dawant, B., Thompson, R., Miga, M., 2011. Intraoperative brain shift compensation: Accounting for dural septa. Biomedical Engineering, IEEE Transactions on 58 (3), 499–508.

Chen, S.-S., Reinertsen, I., Coup, P., Yan, C., Mercier, L., Del Maestro, D., Collins, D., 2012. Validation of a hybrid doppler ultrasound vessel-based registration algorithm for neurosurgery. International Journal of Computer Assisted Radiology and Surgery 7 (5), 667–685.

Chen, X., Weigel, D., Ganslandt, O., Buchfelder, M., Nimsky, C., 2009. Prediction of visual field deficits by diffusion tensor imaging in temporal lobe epilepsy surgery. NeuroImage 45 (2), 286 – 297.

Ciccarelli, O., Parker, G. J. M., Toosy, A. T., Wheeler-Kingshott, C. A. M., Barker, G. J., Boulby, P. A., Miller, D. H., Thompson, A. J., 2003. From diffusion tractography to quantitative white matter tract measures: a reproducibility study. NeuroImage 18 (2), 348 – 359.

Clatz, O., Delingette, H., Talos, I.-F., Golby, A., Kikinis, R., Jolesz, F., Ayache, N., Warfield, S., 2005. Robust nonrigid registration to capture brain shift from intraoperative MRI. IEEE Transactions on Medical Imaging 24 (11), 1417–1427.

Collins, D., Zijdenbos, A., Kollokian, V., Sled, J., Kabani, N., Holmes, C., Evans, A., 1998. Design and construction of a realistic digital brain phantom. IEEE Transactions on Medical Imaging 17 (3), 463 –468.

Comeau, R. M., Sadikot, A. F., Fenster, A., Peters, T. M., 2000. Intraoperative ultrasound for guidance and tissue shift correction in image-guided neurosurgery. Medical Physics 27 (4), 787–800.

Cook, P., Bai, Y., Nedjati-Gilani, S., Seunarine, K., Hall, M., Parker, G., Alexander, D., 2006. Camino: Open-source diffusion-MRI reconstruction and processing. In: International Society for Magnetic Resonance in Medicine.

Crum, W., Hartkens, T., Hill, D., 2004. Non-rigid image registration: theory and practice. British Journal of Radiology 77 (suppl 2), S140–S153.

Cusack, R., Papadakis, N., 2002. New robust 3-D phase unwrapping algorithms: Application to magnetic field mapping and undistorting echoplanar images. NeuroImage 16 (3, Part A), 754 – 764.

Daga, P., Winston, G., Modat, M., Cardoso, M., Stretton, J., Symms, M., McEvoy, A., Hawkes, D., Duncan, J., Ourselin, S., 2011a. Integrating structural and diffusion MR information for optic radiation localisation in focal epilepsy patients. In: Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on. pp. 353 –356.

Daga, P., Winston, G., Modat, M., Cardoso, M., White, M., Mancini, L., McEvoy, A., Thornton, J., Micallef, C., Yousry, T., Hawkes, D., Duncan, J., Ourselin, S., 2011b. Improved neuronavigation through integration of intraoperative anatomical and diffusion images in an interventional MRI suite. In: Information Processing in Computer-Assisted Interventions. pp. 168–178.

Daga, P., Winston, G., Modat, M., White, M., Mancini, L., Cardoso, M., Symms, M., Stretton, J., McEvoy, A., Thornton, J., Micallef, C., Yousry, T., Hawkes, D., Duncan, J., Ourselin, S., 2012. Accurate localization of optic radiation during neurosurgery in an interventional MRI suite. IEEE Transactions on Medical Imaging 31 (4), 882 –891.

Dohrmann, G., Rubin, J., 2001. History of intraoperative ultrasound in neurosurgery. Neurosurgery Clinics of North America 12 (1), 155–166.

Drobnjak, I., Gavaghan, D., Sli, E., Pitt-Francis, J., Jenkinson, M., 2006. Development of a functional magnetic resonance imaging simulator for modeling realistic rigid-body motion artifacts. Magnetic Resonance in Medicine 56 (2), 364–80.

Drobnjak, I., Pell, G., Jenkinson, M., 2010. Simulating the effects of time-varying magnetic fields with a realistic simulated scanner. Magnetic Resonance Imaging 28 (7), 1014–21.

Dumpuri, P., Chen, C.-C., Miga, M., 2003. Model-updated image guidance: A statistical approach to gravity-induced brain shift. In: Medical Image Computing and Computer-Assisted Intervention - MICCAI 2003. Vol. 2878. pp. 375–382.

Engin, A. E., Wang, H.-C., 1970. A mathematical model to determine viscoelastic behavior of in vivo primate brain. Journal of Biomechanics 3 (3), 283 – 296.

Fluck, O., Vetter, C., Wein, W., Kamen, A., Preim, B., Westermann, R., Dec. 2011. A survey of medical image registration on graphics hardware. Computer Methods and Programs in Biomedicine 104 (3), 45–57.

Ford, L. R., Fulkerson, D. R., 1962. Flows in Networks. Princeton University Press.

Geman, S., Geman, D., 1984. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. IEEE Transactions on Pattern Analysis and Machine Intelligence 6 (6), 721–741.

Ghiglia, D. C., Mastin, G. A., Romero, L. A., 1987. Cellular-automata method for phase unwrapping. Applied Optics 4 (1), 267–280.

Ghiglia, D. C., Pritt, M. D., 1998. Two-Dimensional Phase Unwrapping: Theory, Algorithms, and Software. Wiley.

Glocker, B., Komodakis, N., Tziritas, G., Navab, N., Paragios, N., 2008. Dense image registration through MRFs and efficient linear programming. Medical Image Analysis 12 (6), 731 – 741.

Gobbi, D. G., Comeau, R. M., Peters, T. M., 2000. Ultrasound/MRI overlay with image warping for neurosurgery. In: Proceedings of the Third International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 106–114.

Goldstein, R. M., Zebker, H. A., Werner, C. L., 1988. Satellite radar interferometry: Two-dimensional phase unwrapping. Radio Science 23 (4), 713–720.

Gu, Z., Qin, B., 2009. Nonrigid registration of brain tumor resection MR images based on joint saliency map and keypoint clustering. Sensors 9 (12), 10270–10290.

Gudbjartsson, H., Patz, S., Dec. 1995. The Rician distribution of noisy MRI data. Magnetic resonance in medicine 34 (6), 910–914.

Haacke, E., Mittal, S., Wu, Z., Neelavalli, J., Cheng, Y.-C., January 2009. Susceptibility-weighted imaging: Technical aspects and clinical applications, part 1. American Journal of Neuroradiology 30 (1), 19–30.

Hagemann, A., Rohr, K., Stiehl, H., Spetzger, U., Gilsbach, J., 1999. Biomechanical modeling of the human head for physically based, nonrigid image registration. IEEE Transactions on Medical Imaging 18 (10), 875–884.

Hall, W., Truwit, C., 2005. Intraoperative MR imaging. Magnetic Resonance Imaging Clinics of North America 13 (3), 533 – 543.

Hasan, K. M., Alexander, A. L., Narayana, P. A., 2004. Does fractional anisotropy have better noise immunity characteristics than relative anisotropy in diffusion tensor MRI? An analytical approach. Magnetic Resonance in Medicine 51 (2), 413–417.

Hastreiter, P., Rezk-Salama, C., Soza, G., Bauer, M., Greiner, G., Fahlbusch, R., Ganslandt, O., Nimsky, C., 2004. Strategies for brain shift evaluation. Medical Image Analysis 8 (4), 447 – 464.

Hata, N., Dohi, T., Warfield, S., III, W. W., Kikinis, R., Jolesz, F., 1998. Multimodality deformable registration of pre- and intraoperative images for MRI-guided brain surgery. In: Medical Image Computing and Computer-Assisted Intervention MICCAI 2000. Vol. 1. pp. 1067–1074.

Hill, D. L. G., Maurer, C. R., Maciunas, R. J., Barwise, J. A., Fitzpatrick, M. J., Wang, M. Y., 1998. Measurement of Intraoperative Brain Surface Deformation under a Craniotomy. Neurosurgery 43 (3), 514–526.

Huang, H., Ceritoglu, C., Li, X., Qiu, A., Miller, M. I., van Zijl, P. C., Mori, S., 2008. Correction of B0 susceptibility induced distortion in diffusion-weighted images using large-deformation diffeomorphic metric mapping. Magnetic Resonance Imaging 26 (9), 1294 – 1302.

Irfanoglu, M. O., Walker, L., Sammet, S., Pierpaoli, C., Machiraju, R., 2011. Susceptibility distortion correction for echo planar images with non-uniform b-spline grid sampling: A diffusion tensor image study. In: MICCAI. pp. 174–181.

Ishikawa, H., Oct. 2003. Exact optimization for markov random fields with convex priors. IEEE Transactions on Pattern Analysis and Machine Intelligence 25 (10), 1333–1336.

Itoh, K., 1982. Analysis of the phase unwrapping algorithm. Applied Optics 21, 2470–2470.

Jenkinson, M., 2003. Fast, automated, n-dimensional phase-unwrapping algorithm. Magnetic Resonance in Medicine 49 (1), 193–197.

Jenkinson, M., Wilson, J. L., Jezzard, P., 2004. Perturbation method for magnetic field calculations of nonconductive objects. Magnetic Resonance in Medicine 52 (3), 471–477.

Jeurissen, B., Leemans, A., Tournier, J.-D., Jones, D., Sijbers, J., 2010. Estimating the number of fiber orientations in diffusion MRI voxels: A constrained spherical deconvolution study. In: International Society Magnetic Resonance in Medicine. p. 573.

Jezzard, P., Balaban, R., 1995. Correction for geometric distortion in echo planar images from B0 field variations. Magnetic Resonance in Medicine 34 (1), 65–73.

Jolesz, F. A., 2005. Future perspectives for intraoperative MRI. Neurosurgery Clinics of North America 16 (1), 201 – 213.

Kahan, W., 1965. Pracniques: further remarks on reducing truncation errors. Communications of the ACM 8, 783–799.

Kaibara, T., Myles, S. T., Lee, M. A., Sutherland, G. R., 2002. Optimizing epilepsy surgery with intraoperative MR imaging. Epilepsia 43 (4), 425–429.

Kelly, P. J., Kall, B. A., Goerss, S., Earnest, F., 1986. Computer-assisted stereotaxic laser resection of intra-axial brain neoplasms. Journal of Neurosurgery 64 (3), 427–439.

Kim, D.-J., Park, H.-J., Kang, K.-W., Shin, Y.-W., Kim, J.-J., Moon, W.-J., Chung, E.-C., Kim, I. Y., Kwon, J. S., Kim, S. I., 2006. How does distortion correction correlate with anisotropic indices? A diffusion tensor imaging study. Magnetic Resonance Imaging 24 (10), 1369 – 1376.

King, A., Blackall, J., Penney, G., Edwards, P., Hill, D., Hawkes, D., 2000. Bayesian estimation of intraoperative deformation for image-guided surgery using 3-D ultrasound. In: Medical Image Computing and Computer-Assisted Intervention MICCAI 2000. Vol. 1935. pp. 588–597.

Kohli, P., Shekhovtsov, A., Rother, C., Kolmogorov, V., Torr, P., 2008. On partial optimality in multilabel mrfs. In: Proceedings of the 25th international conference on Machine learning. pp. 480–487.

Kohli, P., Torr, P. H. S., 2008. Measuring uncertainty in graph cut solutions. Computer Vision and Image Understanding 112 (1), 30–38.

Kolmogorov, V., 2006. Convergent tree-reweighted message passing for energy minimization. IEEE Transactions on Pattern Analysis and Machine Intelligence 28 (10), 1568–1583.

Kolmogorov, V., Zabin, R., 2004. What energy functions can be minimized via graph cuts? IEEE Transactions on Pattern Analysis and Machine Intelligence 26 (2), 147 –159.

Komodakis, N., Tziritas, G., Paragios, N., 2008. Performance vs computational efficiency for optimizing single and dynamic MRFs: Setting the state of the art with primal-dual strategies. Computer Vision and Image Understanding 112 (1), 14–29.

Kybic, J., Thevenaz, P., Nirkko, A., Unser, M., 2000. Unwarping of unidirectionally distorted EPI images. IEEE Transactions on Medical Imaging 19 (2), 80 –93.

Lafferty, J. D., McCallum, A., Pereira, F. C. N., 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In: Proceedings of the Eighteenth International Conference on Machine Learning. pp. 282–289.

Langley, J., Zhao, Q., 2009a. A model-based 3D phase unwrapping algorithm using gegenbauer polynomials. Physics in Medicine and Biology 54 (17), 5237–5252.

Langley, J., Zhao, Q., 2009b. Unwrapping magnetic resonance phase maps with chebyshev polynomials. Magnetic Resonance Imaging 27 (9), 1293 – 1301.

Lee, D., Nam, W. H., Lee, J. Y., Ra, J. B., 2011. Non-rigid registration between 3D ultrasound and CT images of the liver based on intensity and gradient information. Physics in Medicine and Biology 56 (1), 117–137.

Li, S. Z., 1994. Markov random field models in computer vision. In: Proceedings of the third European conference on Computer Vision (Vol. II). pp. 361–370.

Li, Y., Verma, R., Mar. 2011. Multichannel image registration by feature-based information fusion. IEEE Transactions on Medical Imaging 30 (3), 707–720.

Liang, Z.-P., 1996. A model-based method for phase unwrapping. IEEE Transactions on Medical Imaging 15 (6), 893–897.

Lindner, D., Trantakis, C., Renner, C., Arnold, S., Schmitgen, A., Schneider, J., Meixensberger, J., 2006. Application of intraoperative 3D ultrasound during navigated tumor resection. Minimally Invasive Neurosurgery 49 (4), 197–202.

Liu, J., Drangova, M., 2012. Intervention-based multidimensional phase unwrapping using recursive orthogonal referring. Magnetic Resonance in Medicine 68 (4), 1303–1316.

Liu, W., Tang, X., Ma, Y., Gao, J.-H., 2012. 3D phase unwrapping using global expected phase as a reference: Application to MRI global shimming. Magnetic Resonance in Medicine 70 (1), 160–168.

Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., Suetens, P., 1997. Multimodality image registration by maximization of mutual information. IEEE Transactions on Medical Imaging 16 (2), 187–198.

Mansfield, P., 1977. Multi-planar image formation using nmr spin echoes. Journal of Physics C: Solid State Physics 10 (3), L55.

Martin, A., Hall, W., Liu, H., Pozza, C., Michel, E., Casey, S., Maxwell, R., Truwit, C., 2000. Brain tumor resection: Intraoperative monitoring with high-field-strength MR imaginginitial results. Radiology 215 (1), 221–228.

Mattes, D., Haynor, D., Vesselle, H., Lewellen, T., Eubank, W., 2003. PET-CT image registration in the chest using free-form deformations. IEEE Transactions on Medical Imaging 22 (1), 120–128.

McRobbie, D. W., Moore, E. A., Graves, M. J., Prince, M. R., 2006. MRI from picture to proton. Cambridge University Press.

Merhof, D., Soza, G., Stadlbauer, A., Greiner, G., Nimsky, C., 2007. Correction of susceptibility artifacts in diffusion tensor data using non-linear registration. Medical Image Analysis 11 (6), 588 – 603.

Miga, M., Paulsen, K., Hoopes, P., Kennedy, F.E., J., Hartov, A., Roberts, D., 2000. In vivo quantification of a homogeneous brain deformation model for updating preoperative images during surgery. IEEE Transactions on Biomedical Engineering 47 (2), 266–273.

Miga, M., Sinha, T., Cash, D., Galloway, R., Weil, R., 2003. Cortical surface registration for image-guided neurosurgery using laser-range scanning. IEEE Transactions on Medical Imaging 22 (8), 973–985.

Minka, T. P., 2001. Expectation propagation for approximate bayesian inference. In: Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence. pp. 362–369.

Modat, M., 2012. Efficient Dense Non-Rigid Registration using the Free-Form Deformation Framework. Ph.D. thesis, University College London.

Modat, M., Ridgway, G., Taylor, Z., Lehmann, M., Barnes, J., Hawkes, D., Fox, N., Ourselin, S., 2010. Fast free-form deformation using graphics processing units. Computer methods and Programs in Biomedicine 98 (3), 278–284.

Moiyadi, A., Shetty, P., 2011. Objective assessment of utility of intraoperative ultrasound in resection of central nervous system tumors: A cost-effective tool for intraoperative navigation in neurosurgery. Journal of Neurosciences in Rural Practice 2 (1), 4–11.

Moon-Ho Song, S., Napel, S., Pelc, N., Glover, G., 1995. Phase unwrapping of MR phase images using poisson equation. IEEE Transactions on Image Processing 4 (5), 667–676.

Moussouris, J., 1974. Gibbs and markov random systems with constraints. Journal of Statistical Physics 10 (1), 11–33.

Nabavi, A., Black, P. M., Gering, D. T., Westin, C. F., Mehta, V., Pergolizzi, R. S., Ferrant, M., Warfield, S. K., Hata, N., Schwartz, R. B., Wells, W. M., Kikinis, R., Jolesz, F. A., 2001. Serial intraoperative magnetic resonance imaging of brain shift. Neurosurgery 48 (4), 787–97; discussion 797–8.

Nagelhus, H., Lindseth, F., Selbekk, T., Wollf, A., Solberg, O. V., Harg, E., Rygh, O. M., Tangen, G. A., Rasmussen, I., Augdal, S., Couweleers, F., Unsgaard, G., 2006. Computer-assisted 3D ultrasound-guided neurosurgery: Technological contributions, including multimodal registration and advanced display, demonstrating future perspectives. The International Journal of Medical Robotics and Computer Assisted Surgery 2 (1), 45–59.

Nakao, N., Nakai, K., Itakura, T., 2003. Updating of neuronavigation based on images intraoperatively acquired with a mobile computerized tomographic scanner: technical note. Minimally Invasive Neurosurgery 46 (2), 117–20.

Nico, G., Palubinskas, G., Datcu, M., 2000. Bayesian approaches to phase unwrapping: Theoretical study. IEEE Transactions on Signal Processing 48 (9), 2545–2556.

Nilsson, D., 2010. Imaging visual pathways in temporal lobe epilepsy surgery. In: European Conference in Epileptology.

Nimsky, C., Ganslandt, O., Cerny, S., Hastreiter, P., Greiner, G., Fahlbusch, R., 2000. Quantification of, visualization of, and compensation for brain shift using intraoperative magnetic resonance imaging. Neurosurgery 47 (5), 1070–1079.

Nimsky, C., Ganslandt, O., Hastreiter, P., Fahlbusch, R., 12 2001. Intraoperative compensation for brain shift. Surgical neurology 56 (6), 357–364.

NVIDIA, 2008. NVIDIA CUDA Programming Guide 2.0.

Okudera, H., Kobayashi, S., Kyoshima, K., Gibo, H., Takemae, T., Sugita, K., 1991. Development of the operating computerized tomographic scanner system for neurosurgery. Acta Neurochirurgica 111 (1-2), 61–63.

Papadakis, N., Martin, K., Mustafa, M., Wilkinson, I., Griffiths, P., Huang, C.-H., Woodruff, P., 2002. Study of the effect of CSF suppression on white matter diffusion anisotropy mapping of healthy human brain. Magnetic Resonance in Medicine 48 (2), 394–398.

Park, H., Kubicki, M., Shenton, M., Guimond, A., McCarley, R., Maier, S., Kikinis, R., Jolesz, F., Westin, C., 2003. Spatial normalization of diffusion tensor MRI using multiple channels. Neuroimage 20 (4), 1995–2009.

Parrish, R., Schiffman, J., Anderson, D., 1984. Static and kinetic visual field testing: Reproducibility in normal volunteers. Arch Ophthalmol 102 (10), 1497–1502.

Paul, P., Morandi, X., Jannin, P., 2009. A surface registration method for quantification of intraoperative brain deformations in image-guided neurosurgery. Transactions on Information Technology in Biomedicine 13 (6), 976–983.

Paulsen, K. D., Miga, M. I., Kennedy, F. E., Hoopes, P. J., Hartov, A., Roberts, D. W., 1999. A computational model for tracking subsurface tissue deformation during stereotactic neurosurgery. IEEE Trans Biomed Eng 46 (2), 213–225.

Potts, R. B., 1952. Some generalized order-disorder transformations. Mathematical Proceedings of the Cambridge Philosophical Society 48 (01), 106–109.

Reinertsen, I., Lindseth, F., Unsgaard, G., Collins, D. L., 2007. Clinical validation of vessel-based registration for correction of brain-shift. Medical image analysis 11 (6), 673–684.

Risholm, P., Janoos, F., Norton, I., Golby, A. J., III, W. M. W., 2013. Bayesian characterization of uncertainty in intra-subject non-rigid registration. Medical Image Analysis 17 (5), 538 – 555.

Roberts, D. W., Hartov, A., Kennedy, F. E., Miga, M. I., Paulsen, K. D., Oct. 1998. Intraoperative brain shift and deformation: a quantitative analysis of cortical displacement in 28 cases. Neurosurgery 43 (4), 749–758.

Roche, A., Pennec, X., Malandain, G., Ayache, N., 2001. Rigid registration of 3D ultrasound with MR images: A new approach combining intensity and gradient information. IEEE Transactions on Medical Imaging 20 (10), 1038–1049.

Rohlfing, T., Maurer, Jr., C. R., 2003. Nonrigid image registration in shared-memory multiprocessor environments with application to brains, breasts, and bees. IEEE Transactions on Information Technology in Biomedicine 7 (1), 16–25.

Rooney, W. D., Johnson, G., Li, X., Cohen, E. R., Kim, S.-G., Ugurbil, K., Springer, C. S., 2007. Magnetic field and tissue dependencies of human brain longitudinal 1h2o relaxation in vivo. Magnetic Resonance in Medicine 57 (2), 308–318.

Rubin, J. M., Mirfakhraee, M., Duda, E. E., Dohrmann, G. J., Brown, F., 1980. Intraoperative ultrasound examination of the brain. Radiology 137 (3), 831–832.

Rueckert, D., Sonoda, L., Hayes, C., Hill, D., Leach, M., Hawkes, D., 1999. Nonrigid registration using free-form deformations: Application to breast MR images. IEEE Transactions on Medical Imaging 18 (8), 712–721.

Ruthotto, L., Kugel, H., Olesch, J., Fischer, B., Modersitzki, J., Burger, M., Wolters, C. H., 2012. Diffeomorphic susceptibility artifact correction of diffusion-weighted magnetic resonance images. Physics in Medicine and Biology 57 (18), 5715.

Salamon, N., Kung, J., Shaw, S. J., Koo, J., Koh, S., Wu, J. Y., Lerner, J. T., Sankar, R., Shields, W. D., Engel, J., Fried, I., Miyata, H., Yong, W. H., Vinters, H. V., Mathern, G. W., 2008. FDG-PET/MRI coregistration improves detection of cortical dysplasia in patients with epilepsy. Neurology 71 (20), 1594–1601.

Schenck, J. F., 1996. The role of magnetic susceptibility in magnetic resonance imaging: Mri magnetic compatibility of the first and second kinds. Medical Physics 23 (6), 815–850.

Semah, F., Ryvlin, P., 2005. Can we predict refractory epilepsy at the time of diagnosis? Epileptic Disorders 7 (1 Suppl), S10–3.

Shamir, R., Joskowicz, L., Spektor, S., Shoshan, Y., 2009. Localization and registration accuracy in image guided neurosurgery: a clinical study. International Journal of Computer Assisted Radiology and Surgery 4 (1), 45–52.

Shannon, C., 1948. A mathematical theory of communication. Bell system technical journal 27.

Siewerdsen, J. H., Moseley, D. J., Burch, S., Bisland, S. K., Bogaards, A., Wilson, B. C., Jaffray, D. A., 2005. Volume CT with a flat-panel detector on a mobile, isocentric C-arm: Pre-clinical investigation in guidance of minimally invasive surgery. Medical Physics 32 (1), 241–254.

Sinha, T., Dawant, B., Duay, V., Cash, D., Weil, R., Thompson, R., Weaver, K., Miga, M., 2005. A method to track cortical surface deformations using a laser range scanner. IEEE Transactions on Medical Imaging 24 (6), 767–781.

Škrinjar, O., Nabavi, A., Duncan, J., 2002. Model-driven brain shift compensation. Medical Image Analysis 6 (4), 361 – 373.

Smith, S., 2002. Fast robust automated brain extraction. Human brain mapping 17 (3), 143–155.

Smith, S. M., Jenkinson, M., Woolrich, M. W., Beckmann, C. F., Behrens, T. E., Johansen-Berg, H., Bannister, P. R., Luca, M. D., Drobnjak, I., Flitney, D. E., Niazy, R. K., Saunders, J., Vickers, J., Zhang, Y., Stefano, N. D., Brady, J. M., Matthews, P. M., 2004. Advances in functional and structural MR image analysis and implementation as FSL. NeuroImage 23, Supplement 1 (0), S208 – S219.

So, R. W. K., Tang, T. W. H., Chung, A. C. S., 2011. Non-rigid image registration of brain magnetic resonance images using graph-cuts. Pattern Recognition 44 (10-11), 2450–2467.

Spencer, S., Huh, L., 2008. Outcomes of epilepsy surgery in adults and children. The Lancet Neurology 7 (6), 525 – 537.

Studholme, C., 2008. Dense feature deformation morphometry: Incorporating DTI data into conventional MRI morphometry. Medical Image Analysis 12 (6), 742–751.

Studholme, C., Constable, R., Duncan, J., 2000. Accurate alignment of functional EPI data to anatomical MRI using a physics-based distortion model. IEEE Transactions on Medical Imaging 19 (11), 1115–1127.

Studholme, C., Hill, D., Hawkes, D., Jan. 1999. An overlap invariant entropy measure of 3D medical image alignment. Pattern Recognition 32 (1), 71–86.

Studholme, C., Hill, D. L. G., Hawkes, D. J., 1995. Multiresolution voxel similarity measures for MR-PET registration. In: Information processing in Medical Imaging (IPMI).

Sun, G., lei Chen, X., Zhao, Y., Wang, F., Song, Z.-J., bo Wang, Y., Wang, D., nan Xu, B., 2011. Intraoperative MRI with integrated functional neuronavigation-guided resection of supratentorial cavernous malformations in eloquent brain areas. Journal of Clinical Neuroscience 18 (10), 1350 – 1354.

Sun, H., Lunn, K. E., Farid, H., Wu, Z., Roberts, D. W., Hartov, A., Paulsen, K. D., 2005. Stereopsis-guided brain shift compensation. IEEE Transactions on Medical Imaging 24 (8), 1039–1052.

Sutton, C., Mccallum, A., 2007. An introduction to conditional random fields for relational learning. In: Introduction to Statistical Relational Learning.

Tao, R., Fletcher, P., Gerber, S., Whitaker, T., 2009. A variational image-based approach to the correction of susceptibility artifacts in the alignment of diffusion weighted and structural MRI. In: IPMI. pp. 664–675.

Tarlow, D., Adams, R. P., 2012. Revisiting uncertainty in graph cut solutions. In: IEEE Conference on Computer Vision and Pattern Recognition.

Thevenaz, P., Unser, M., 2000. Optimization of mutual information for multiresolution image registration. IEEE Transactions on Image Processing 9 (12), 2083–2099.

Unsgaard, G., Rygh, O. M., Selbekk, T., Mller, T. B., Kolstad, F., Lindseth, F., Hernes, T. A. N., 2006. Intra-operative 3D ultrasound in neurosurgery. Acta Neurochirurgica 148 (3), 235–253.

Veksler, O., 2009. Multi-label moves for MRFs with truncated convex priors. In: Energy Minimization Methods in Computer Vision and Pattern Recognition. Vol. 5681. pp. 1–13.

Viola, P., Wells, W., Jun. 1995. Alignment by maximization of mutual information. In: Computer Vision, 1995. Proceedings., Fifth International Conference on. pp. 16–23.

Vulliemoz, S., Thornton, R., Rodionov, R., Carmichael, D., Guye, M., Lhatoo, S., McEvoy, A., Spinelli, L., Michel, C., Duncan, J., Lemieux, L., 2009. The spatio-temporal mapping of epileptic networks: Combination of EEG–fMRI and EEG source imaging. NeuroImage 46 (3), 834 – 843.

Waites, A. B., Briellmann, R. S., Saling, M. M., Abbott, D. F., Jackson, G. D., 2006. Functional connectivity networks are disrupted in left temporal lobe epilepsy. Annals of Neurology 59 (2), 335–343.

Wang, H. C., Wineman, A. S., 1972. A mathematical model for the determination of viscoelastic behavior of brain in vivoi oscillatory response. Journal of Biomechanics 5 (5), 431 – 446.

Warfield, S., Talos, I.-F., Tei, A., Bharatha, A., Nabavi, A., Ferrant, M., Black, P., Jolesz, F., Kikinis, R., 2002. Real-time registration of volumetric brain MRI by biomechanical simulation of deformation during image guided neurosurgery 5, 3–11.

Warfield, S. K., Haker, S. J., Talos, I.-F., Kemper, C. A., Weisenfeld, N., Mewes, A. U., Goldberg-Zimring, D., Zou, K. H., Westin, C.-F., Wells, W. M., Tempany, C. M., Golby, A., Black, P. M., Jolesz, F. A., Kikinis, R., 2005. Capturing intraoperative deformations: Research experience at brigham and women's hospital. Medical Image Analysis 9 (2), 145 – 162.

Wehner, T., LaPresto, E., Tkach, J., Liu, P., Bingaman, W., Prayson, R. A., Ruggieri, P., Diehl, B., 2007. The value of interictal diffusion-weighted imaging in lateralizing temporal lobe epilepsy. Neurology 68 (2), 122–127.

Wheeler, D. A., 2005. Why open source software / free software (oss/fs, floss, or foss)? look at the numbers!

Wiebe, S., Blume, W., Girvin, J., Eliasziw, M., 2001. A randomized, controlled trial of surgery for temporal-lobe epilepsy. New England Journal of Medicine 345 (5), 311–318.

Winkler, D., Tittgemeyer, M., Schwarz, J., Preul, C., Strecker, K., Meixensberger, J., 2005. The first evaluation of brain shift during functional neurosurgery by deformation field analysis. Journal of Neurology, Neurosurgery and Psychiatry 76 (8), 1161–1163.

Winston, G. P., Daga, P., Stretton, J., Modat, M., Symms, M. R., McEvoy, A. W., Ourselin, S., Duncan, J. S., 2011. Optic radiation tractography and vision in anterior temporal lobe resection. Annals of Neurology 71 (3), 334–341.

Wittek, A., Miller, K., Kikinis, R., Warfield, S. K., 2007. Patient-specific model of brain deformation: Application to medical image registration. Journal of Biomechanics 40 (4), 919 – 929.

Ying, L., Liang, Z.-P., Munson, D.C., J., Koetter, R., Frey, B., 2006. Unwrapping of MR phase images using a markov random field model. IEEE Transactions on Medical Imaging 25 (1), 128 –136.

Yogarajah, M., Focke, N., Bonelli, S., Cercignani, M., Acheson, J., Parker, G., Alexander, D., McEvoy, A. W., Symms, M., Koepp, M., Duncan, J., 2009. Defining Meyer's loop-temporal lobe resections, visual field deficits and diffusion tensor tractography. Brain 132 (6), 1656–1668.

Zhou, K., Zaitsev, M., Bao, S., 2009. Reliable two-dimensional phase unwrapping method using region growing and local linear estimation. Magnetic Resonance in Medicine 62 (4), 1085–1090.

Zhuang, X., Arridge, S., Hawkes, D., Ourselin, S., 2011. A nonrigid registration framework using spatially encoded mutual information and free-form deformations. IEEE Transactions on Medical Imaging, 1819 –1828.