Enhancing Registration for Image-Guided Neurosurgery

Martin Kochan

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

I, Martin Kochan, 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.

To my parents and Janka with Slavo.

Abstract

Pharmacologically refractive temporal lobe epilepsy and malignant glioma brain tumours are examples of pathologies that are clinically managed through neurosurgical intervention. The aims of neurosurgery are, where possible, to perform a resection of the surgical target while minimising morbidity to critical structures in the vicinity of the resected brain area. Image-guidance technology aims to assist this task by displaying a model of brain anatomy to the surgical team, which may include an overlay of surgical planning information derived from preoperative scanning such as the segmented resection target and nearby critical brain structures. Accurate neuronavigation is hindered by brain shift, the complex and non-rigid deformation of the brain that arises during surgery, which invalidates assumed rigid geometric correspondence between the neuronavigation model and the true shifted positions of relevant brain areas. Imaging using an interventional MRI (iMRI) scanner in a next-generation operating room can serve as a reference for intraoperative updates of the neuronavigation. An established clinical image processing workflow for iMRI-based guidance involves the correction of relevant imaging artefacts and the estimation of deformation due to brain shift based on non-rigid registration. The present thesis introduces two refinements aimed at enhancing the accuracy and reliability of iMRI-based guidance. A method is presented for the correction of magnetic susceptibility artefacts, which affect diffusion and functional MRI datasets, based on simulating magnetic field variation in the head from structural iMRI scans. Next, a method is presented for estimating brain shift using discrete non-rigid registration and a novel local similarity measure equipped with an edge-preserving property which is shown to improve the accuracy of the estimated deformation in the vicinity of the resected area for a number of cases of surgery performed for the management of temporal lobe epilepsy and glioma.

Impact Statement

Pharmacologically refractive temporal lobe epilepsy (TLE) and malignant glioma brain tumours are examples of pathologies managed by resective neurosurgery in which the imperatives are to remove the surgical target and to minimise morbidity to nearby critical brain areas. Image-guidance technology can assist this task by displaying to the surgeon a model of brain anatomy that may include surgical planning information such as critical brain areas identified from preoperative MRI. Image-guidance assumes a rigid transformation between pre/intraoperative positions of brain areas, which is invalidated by severe non-rigid deformation (brain shift). The advanced operating suite at the National Hospital for Neurology and Neurosurgery (NHNN) is equipped with an interventional MRI (iMRI) scanner that enables intraoperative imaging. A previous collaboration between the UCL Centre for Medical Image Computing, the UCL Institute of Neurology and the NHNN established novel fast, automated software methods to enable an intraoperative update of a surgical plan. These include the correction of iMRI-specific imaging artefacts, and the estimation of brain shift based on non-rigid registration whereby the best and most plausible alignment is sought between preoperative MRI and intraoperative iMRI images. In this thesis, the above steps were investigated with regard to accuracy near the resection boundary where accuracy is required.

A significant imaging artefact in iMRI is the susceptibility artefact, which is the geometric distortion of diffusion MRI (dMRI) and functional MRI (fMRI) data that arises near the airtissue transition at the resection boundary. This work confirmed that it is feasible to partially correct for the artefact by simulating the magnetic field from iMRI structural scans acquired at an advanced stage of anterior temporal lobe resection (ATLR) performed for the management of TLE in 12 surgical cases. A refinement of this approach may complement more established artefact correction methods such as MRI-acquired field maps that exhibit reduced reliability near the resection boundary due to the low signal-to-noise ratio of iMRI. The proposed method can in principle correct non-interventional dMRI and fMRI datasets from retrospective studies, as many older studies did not acquire field maps and supplementary artefact correction may be needed to enable their seamless inclusion within new studies.

Next, an investigation was performed into non-rigid registration for brain shift estimation. An important aspect of registration schemes is how similarity between the image pair is measured for local neighbourhoods. An investigation of the effect of introducing an edge-awareness property into local similarity measurement revealed improvement in registration accuracy near the resection boundary for 12 iMRI ATLR cases and 8 iMRI cases of glioma tumour resection. This is beneficial because the extent of retained vs. resected tissue is challenging to determine automatically due to liquid in the resection cavity and contrast changes. The proposed similarity measurement can be used for non-interventional MRI datasets for registration-based image analysis where effect of interest occurs near intensity edges such as for evaluation of dementia-related disease progression where volume change tends to occur near the ventricles or for lung inhale/exhale pairs where the rigid ribcage borders the lungs that undergo sliding motion.

Acknowledgements

Firsly, I would like to thank my supervisory team for their support during my time at UCL. I genuinely thank Dr Dan Stoyanov, for his patience, encouragement through times of self-doubt, key insights on the state of the art in image-guided interventions, a drive to pass on to all of us the intuitions on the general give-and-take of the modern research environment and for constructive feedback throughout the last few years, all of which were immensely valuable. I am also very grateful to Prof Sébastien Ourselin, for his guidance at key moments during the programme, access to his team and for giving me the opportunity to work with the fascinating iMRI data and to visit neurosurgeries, which was a real privilege.

I would also like to thank my PhD viva examiners, Dr Ben Glocker and Dr Eleftheria Panagiotaki for the insight and the constructive approach that they brought to the table on the day and for their very thorough feedback regarding the thesis.

I am grateful to all my colleages with whom I collaborated for being approachable, meticulous and really professional. I would especially like to thank Dr Pankaj Daga for introducing me to the field and just teaching a passion for solving the right problems in a way both principled and pragmatic. I have much to learn in this regard. I also thank Dr Marc Modat for his time and feedback at a key stage during the drafting of my thesis.

I met a few great people at CMIC along the way who were a great inspiration and I hope that we will remain friends. Gergely, Miklos, Pankaj, Rachel — it was a pleasure to know you in the first office. To the good folks at SRV — Max, Xiaofei, Geoff, Rene and the whole growing team — it has been amazing to watch everyone pull together and support each other at key moments despite everybody working so hard.

Finally, I cannot thank enough my family — I am quite lucky to have such supportive and caring parents and sister and I grow to appreciate this more and more.

Publication List

Peer Reviewed Journal Papers

Kochan, M., Daga, P., Burgos, N., White, M., Cardoso, M. J., Mancini, L., Winston, G. P., McEvoy, A. W., Thornton, J., Yousry, T., Duncan, J. S., Stoyanov, D., and Ourselin, S. (2015). Simulated field maps for susceptibility artefact correction in interventional MRI. *International Journal of Computer Assisted Radiology and Surgery*, 10(9):14051416.

Peer Reviewed Conference Papers

Kochan, M., Daga, P., Burgos, N., White, M., Cardoso, M. J., Mancini, L., Winston, G. P., McEvoy, A. W., Thornton, J., Yousry, T., Duncan, J. S., Stoyanov, D., and Ourselin, S. (2014). Simulated Field Maps: Toward Improved Susceptibility Artefact Correction in Interventional MRI. In *Information Processing in Computer-Assisted Interventions*, pages 226235. Springer.

Kochan, M., Modat, M., Vercauteren, T., White, M., Mancini, L., Winston, G. P., McEvoy, A. W., Thornton, J. S., Yousry, T., Duncan, J. S., Ourselin, S., and Stoyanov, D. (2016). Bilateral weighted adaptive local similarity measure for registration in neurosurgery. In *Medical Image Computing and Computer-Assisted Intervention-MICCAI 2016*, volume 9902, page 81. Springer.

Contents

1	Intr	roduction	1
	1.1	Pathologies Requiring Neurosurgical Management	4
	1.2	Use of iMRI Guidance in an Operating Room	6
	1.3	Key Challenges in iMRI-Guided Neurosurgery	9
	1.4	Aims of the Thesis	14
	1.5	Methodological Contributions of the Thesis	14
2	Rev	iew of Image Guided Neurosurgery	16
	2.1	Intraoperative Imaging of Brain Shift	16
	2.2	Biomechanical Modelling Based Brain Shift Estimation	18
	2.3	Registration Based Brain Shift Estimation	20
		2.3.1 Geometric Transformations	26
		2.3.2 Similarity Metrics	30
		2.3.3 Regularisation	33
	2.4	Advanced Registration Algorithms	36
		2.4.1 Modelling Missing Correspondences	37
		2.4.2 Fulfilling Time Constraints of Neurosurgery	39
		2.4.3 Incorporating Registration Uncertainty	40
		2.4.4 Evaluating Estimated Brain Shift	42
	2.5	Summary	44
3	Fiel	d Map Simulation for Distortion Correction	45

	moroduction		4
	3.1.1 The Origin of the	he Susceptibility Artefact	
	3.1.2 Susceptibility A	rtefact Correction	
	3.1.3 Proposed Corre	ction using Simulated Field Maps	
3.2	Methods		
	3.2.1 Field Map in Te	erms of Voxel Displacement	
	3.2.2 Air-Tissue Segn	nentation	
	3.2.3 Field Map Estir	nation	
	3.2.4 Data Acquisitio	n	
3.3	Experiments with Clini	ical Data and Results	
3.4	Discussion		
	3.4.1 Accuracy of Air	-Tissue Segmentation	
	3.4.2 Comparison of .	Acquired and Simulation Field Maps	
	3.4.3 Running Time	of the Pipeline	
3.5	${\rm Conclusion} . \ . \ . \ .$		
ł Re	istration near Resect	ion Boundary	
4.1	Introduction		
4.2	Adaptively Weighted L	ocal Similarity Measure	
4.3	Registration using a Di	screte Optimisation Framework	
	Experiments		
4.4			
4.4	4.4.1 Patch Matching	on 2D Synthetic Phantom	
4.4	4.4.1 Patch Matching4.4.2 Recovery of a 32	on 2D Synthetic Phantom	
4.4	4.4.1 Patch Matching4.4.2 Recovery of a 324.4.3 Evaluation on if	g on 2D Synthetic Phantom	
4.4	4.4.1 Patch Matching4.4.2 Recovery of a 34.4.3 Evaluation on if4.4.4 Segmentation P	g on 2D Synthetic Phantom	
4.4 4.5	 4.4.1 Patch Matching 4.4.2 Recovery of a 3 4.4.3 Evaluation on if 4.4.4 Segmentation P Discussion and Conclust 	g on 2D Synthetic Phantom	
4.44.55 Dis	 4.4.1 Patch Matching 4.4.2 Recovery of a 31 4.4.3 Evaluation on if 4.4.4 Segmentation P Discussion and Conclusion 	g on 2D Synthetic Phantom	

List of Figures

1.1	The visual system in the human brain. (a) The paths of the fiber tracts to the	
	visual cortex (courtesy http://www.thebrain.mcgill.ca). (b) The morphology of the	
	human brain (courtesy Virtual Hospital) with the part resection during the ATLR	
	procedure highlighted.	5
1.2	The appearance of a malignant glioma in (a) a T1-weighted structural scan and	
	(b) a T1-weighted FLAIR scan. The images are taken from preoperative imaging	
	performed prior to the surgery for the patient $#3$ from the glioma registration	
	dataset described in Section 4.4.3.	6
1.3	The INGS suite at the NHNN in London with an integrated iMRI scanner. (Image	
	courtesy of Dr Pankaj Daga.)	7
1.4	Brain shift for two representative cases of ATLR. Top row, patient #1. Bottom	
	row, patient #9. The ATLR dataset is described in Section 3.2.4. Left, preopera-	
	tive structural T1-weighted scan acquired about one month prior to surgery, shown	
	affinely pre-registered to the intraoperative scan. Right, intraoperative scan ac-	
	quired using an closed-bore iMRI scanner in the operating room with the patient's	
	head in the intraoperative orientation. Highlights A–E are explained in the main text.	10

1.5	Missing correspondences. Left and middle: Example pre/intraoperative image slice	
	(patient $\#9$ from the ATLR dataset, Section 3.2.4). Right, top row: A schematic	
	phantom with landmarks. Right, bottom row: arrow ends mark each landmark's	
	corresponding location in the other image. Preoperative landmarks E and F have	
	no correspondence in the intraoperative space. The other landmarks do not move	
	and so map to the same location.	11
1.6	Volume folding due to partial collapse of the ventricles. Left and middle: Ex-	
	ample pre/intraoperative slice (patient $\#9$ from the ATLR dataset, Section 3.2.4).	
	Right, top row: a schematic of volume between landmarks D, E, F in the preoperative	
	space collapsing into a line segment in the intraoperative space. Right, bottom row:	
	arrow ends mark each landmark's corresponding location in the other image. $\ . \ .$	12
1.7	Volume folding due to tissue retraction. A schematic phantom of the segment con-	
	necting B–D being retracted to the right during surgery. Top row: The volume	
	between F, D, B in the intraop. image folds into a line segment in the preop. image.	
	Bottom row: Arrow ends mark each landmark's corresponding location in the other	
	image	13
1.8	A schematic of the image registration part of the software pipeline used clinically	
	at the NHNN during the study of [Winston et al., 2014] for brain shift estimation	
	based on multi-channel registration. The susceptibility artefact correction step is	
	not shown here	15
2.1	Registration for an ATLR dataset. The input volumes are: (a) the preoperative T1w	
	with the head in the standard orientation and (f) the intraoperative T1w volume	
	with the head in the intraoperative position. The results are: (b) preop. registered to	
	intraop, affinely i.e. using a global transformation. (c) preop, registered to intraop.	
	non-rigidly. (d) intraop. registered to preop. non-rigidly. (e) intraop. registered to	
	preop. affinely. Note: all images are 2D slices through 3D volumes.	21

2.2	For the convenience of a simpler resampling of the warped floating image, the reg-	
	is tration problem can be posed so that the backward transformation ${\cal T}$ is sought	
	which maps each reference image voxel grid point \vec{x} to a location $T(\vec{x})$ in the geo-	
	metric space of the floating image that is allowed to lie between the grid points of	
	the floating image. The sought intensity of the point in the warped floating image	
	can be found by interpolating the intensities of the floating image grid points that	
	neighbour $T(\vec{x})$. (Image courtesy of Dr Marc Modat.)	23
2.3	Rigid transformations of a cube involve translations and rotations. (Image courtesy	
	of Dr Marc Modat.)	24
2.4	Affine transformations of a cube involve translations, rotations, scaling and shearing.	
	(Image courtesy of Dr Marc Modat.)	24
2.5	Two examples of non-rigid transformation of a cube using distinct deformations.	
	(Image courtesy of Dr Marc Modat.)	25
3.3	The processing pipeline for correcting the susceptibility artefact in a DW-MRI im-	
	age. The dashed lines represent the step into which either a phase-unwrapped	
	acquired field map or a simulated field map can be fed. The entry points of the	
	simulation are non-distorted preoperative T1w and intraoperative T1w and T2w	
	images	55
3.4	Air-tissue segmentation. (a) Intraoperative $T1w+T2w$ sum image. The section runs	
	through a plane close to the anatomical coronal plane with the head at an angle due	
	to intraoperative orientation. (b) An accompanying pseudo-CT (range -1000–1000 $$	
	HU). (c) The result of the proposed segmentation within the intraoperative FOV	
	(red for air, green for soft-tissue, blue for bone). (d) The final air-tissue segmenta-	
	tion (black for air, white for tissue). The appended inferior volume outside of the	
	intraoperative FOV is an approximate air-tissue segmentation based on preoperative	
	T1w MRI	56
3.5	Histogram (cropped for clarity of the figure) of the $T1w+T2w$ image in Fig. 3.4(a)	
	and the fit of the GMM model for the segmentation in Fig. 3.4(a). The data come	
	from intraoperative scans for subject #3. \ldots	59

- 3.7 Field maps expressed as mm of displacement along the phase encode direction. The view is centered at anterior temporal lobe resection cavity. The brain surface outlined using a surface extractor is shown for reference (red outline). (a–c) A phase-wrapped acquired field map for subject #3, showing a step-change in phase value close to the resection margin. (d–f) The acquired field map after phase-unwrapping; only the volume inside the brain mask is shown, because the phase-unwrapping was restricted to the brain only. (g–i) The proposed simulated field map. (j–l) A voxel-wise absolute difference between the simulated and the phase-unwrapped acquired field maps, only considered within the brain. Left to right: coronal (a,d,g,j), sagittal (b,e,h,k) and axial sections (c,f,i,l), respectively. Slice orientations are close to the standard anatomical planes. I used a brain surface extractor included in NiftyView (http://cmic.cs.ucl.ac.uk/home/software).

61

62

4.1	Brain masks extracted using the FSL-BET fast automatic tool (yellow with blue	
	border). (a) Mask for ATLR case $\#1$ which has an air-filled cavity. The base of the	
	cavity is incorrectly included in the brain mask. (b) Mask for ATLR case $\#9$ whose	
	resection cavity is filled with a saline solution. (The ATLR dataset is described in	
	Section 3.2.4.)	76
4.2	(a) T1-weighted reference image and (b) intensity-based weights for a point (blue	
	cross). (c) T1-weighted floating image and (d) intensity-based weights for the point.	
	(e) Final weights based on (b), (d) and distance from the point. \ldots	80
4.3	The employed range kernel is more long-tailed than the Gaussian kernel. \ldots .	81
4.4	2D numerical phantoms. (a–d) A contrast-enhanced lesion near a resection. (e–h)	
	The medial longitudinal fissure. (a,e) The reference image and the fixed patch.	
	The cross shows the location of the centre of the fixed patch. The outline shows the	
	size of the fixed patch. (b,f) The floating image and the moving patch. The cross	
	shows the location of the centre of the moving patch at zero displacement. The inner	
	outline shows the size of the moving patch at zero displacement. The outer outline	
	shows the extent of the area within which the patch is allowed to move (up until	
	the edge of the patch touches the outline at maximum allowed displacement). (c,g)	
	The similarity profile of LNCC as a function of the displacement of the moving	
	patch. Neg. (pos.) x displacement is left (right) and neg. (pos.) y displacement is	
	down (up). (d,h) The same for LNCC-AW.	84
4.5	Axial view of 3D BrainWeb based phantom. (a) Reference image (inserted resec-	
	tion). (b) Floating image (synthetic deformation). (c) Target registration error. (d)	
	Map of log Jacobian determinant for ground truth deformation (forward field). (e)	
	Same map for fields recovered using LNCC, (f) LNCC-AW with $\alpha = 0.30$ and (g)	
	LNCC-AW with $\alpha = 0.10.$	85

- 4.7 Registration of glioma resection case #3. Views are through the (a—d) axial, (e—h) coronal and (i—l) saggital plane. (a, e, i) Reference volume (intraoperative T1w iMRI with patient's head in intraoperative orientation). Landmarks annotated within the slice are shown as yellow crosses (smaller if out-of-slice). (b, f, j) Floating volume (pre-craniotomy T1w iMRI). Homologous landmarks are shown. Reference brain surface is outlined in white for comparison. (c, g, k) Warped volume for LNCC. (d, h, l) Warped volume for LNCC-AW. The last row shows the estimated deformation in saggital view. (m, n) Magnitude of displacement for LNCC and LNCC-AW, respectively [0 10 mm]. (o, p) Absolute log Jacobian determinant for LNCC and LNCC-AW, respectively [-1 1].

88

List of Tables

3.2	Misregistration (in mm) between the ground truth landmarks and the landmarks	
	in the uncorrected b0-DW-MRI images (U) and those corrected using acquired (A)	
	and simulated field maps (S). The results are reported for 4 ROIs: (1) near resection	
	margin, (2) brain stem, (3) head-holder fixation pins, (4) rest of the brain. \ldots	66
3.3	Results of one-tailed Student's t-test of landmark misregistration reduction between	
	uncorrected (Unc.) images and those corrected (Corr.) using acquired (AFM) and	
	simulated (SFM) field maps (p -value and lower bound of the CI at 1.25% significance	
	level in mm).	67
3.4	Mean (std) of normalised sum of square diffusion tensor fit errors summed up across	
	the whole brain volume. The results are reported for the uncorrected (Unc.) DW-	
	MRI datasets and those corrected using the acquired (AFM) and the simulated field	
	maps (SFM). The summary line contains the column averages.	69
		~ ~
4.1	Mean (std) mandmark misregistration for the epilepsy (ATLR) iMRI dataset	89
4.2	Mean (std) of abs. log Jacobian determinant map in vicinity of resection for the	
	ATLR iMRI dataset.	90
4.3	Mean (std) mandmark misregistration for the glioma resection iMRI dataset. (The	
	box-plot is illustrative only due to small number of samples.) $\ldots \ldots \ldots \ldots$	91
4.4	Mean (std) of abs. log Jacobian determinant map in vicinity of resection for the	
	glioma iMRI dataset. (The box-plot is illustrative only due to small number of	
	samples.)	93

Chapter 1

Introduction

Pharmacologically refractive temporal lobe epilepsy and malignant glioma brain tumours are examples of pathologies that are clinically managed through neurosurgical intervention. The aims of the surgeons throughout a neurosurgical procedure are, where possible, to perform a resection in such a manner that the the area of pathological tissue is completely resected while choosing a surgical approach trajectory that stays clear of critical brain structures. Obviously, these tasks may get in conflict with each other during surgery and the surgeon may need to make a decision regarding which goal takes precedence.

Neurosurgical procedures may follow a patient-specific surgical plan that is designed and agreed by the clinical team prior to surgery and which outlines the surgical target, the approach trajectory and the areas to be avoided. These data are derived from preoperative imaging performed within the weeks before the surgery. Magnetic resonance imaging (MRI) has emerged the imaging modality of choice for both diagnosis and for preoperative imaging due to its high resolution, low signal to noise (SNR) ratio and flexible imaging capabilities. At its most basic, MRI enables high-resolution high-contrast structural imaging that can yield a rich tomographic image of anatomical structures in the brain stem, as well as the cortical folding pattern with substantial detail. Diffusion-weighted MRI (DW-MRI) imaging followed by a probabilistic tractography post-processing step can reveal locations of axonal nerve fibre tracts that cross the white matter which are neither visible in the

CHAPTER 1. INTRODUCTION

structural scans nor recognisable by the naked eye during surgery and which exhibit intra-patient variability [Essayed et al., 2017]. Functional MRI (fMRI) can reveal eloquent (activity-specific) cortical regions such as the motor cortex or the brain speech centres, which again exhibit intrapatient variability. The mentioned modalities are useful for mapping the brain areas to be avoided as part of the surgical planning. MRI can also partly reveal the location and extent of pathological areas. This information is crucial for surgical planning and will be outlined further in Section 1.1.

A practical task that arises during neurosurgery is how to display preoperative planning information to the surgical team in an effective manner. In conventional surgery the surgeons would simply keep a "mental picture" of the surgical plan as they progress through the procedure and navigate using the anatomical landmarks which they reach and recognise as points of reference in the neuroanatomy. However, this is inaccurate due to the difficulty of identifying landmarks, due to the linear/planar nature of landmarks and due to uncertainty of localisation in areas between the sparse landmarks that are visually present to the surgeon as homogeneous tissue.

This problem attempts to be solved in a practical manner by the image-guided neurosurgery (IGNS) technology (also image-based guidance, neuronavigation, frameless stereotaxy), which was first introduced by [Roberts et al., 1986]. In the setting of IGNS, a reference frame is rigidly attached to the patient's head and its position and orientation in the 3D physical space of the operating theatre is tracked in real time, often using optical tracking technology. Further, a rigid transformation is established by means of features that can be sampled on the patient's skull. Corresponding features can be identified in the preoperative scan. Options to achieve this include the use of anatomical landmarks, fiducial markers or skull surface. The estimation of this rigid transformation is referred to as rigid registration. Finally, surgical instruments can be tracked in real-time and thus any point in the physical space can be located in the virtual space of the preoperative image. For instance, the most basic instrument is a tracked pointer (stylus) with whose distal end the surgeon can touch a surface anatomical landmarks (e.g. a cortical blood vessel bifurcation) and thus locate the corresponding point in the preoperative image and, by extension, the surgical plan.

There are three sources of error in the IGNS technology that may contribute to a given point

CHAPTER 1. INTRODUCTION

of interest in the virtual space to be displaced from its true corresponding location in the physical space [Gerard et al., 2017]. Firstly, there is an error in the tracking of the reference frame compounded by error in the rigid transformation between the physical landmarks and their positions in the preoperative image. Secondly, imaging artefacts may displace points of interest in the virtual space of the preoperative image. Thirdly, there is an error in the assumption that the brain forms a rigid system together with the skull to which the reference frame is attached. The latter assumption is violated by the complex and difficult to characterise non-rigid deformation of the soft tissue of the brain, also termed *brain shift*, which arises after craniotomy and develops with the progress of the surgery.

The phenomenon of brain shift can be mitigated by intraoperative imaging. Intraoperative IMRI (iMRI) has emerged as the gold-standard modality for intraoperative imaging of the brain whereby the MRI scanner is located directly in the advanced operating room. iMRI enables accurate structural imaging of the brain affected by brain shift [Nimsky et al., 2000, Nabavi et al., 2001] and resection control [Nimsky et al., 2001]. In the context of IGNS the intraoperative image can serve as an update to the neuronavigation. Indeed, the surgical plan and the delineations of brain areas of interest can potentially be geometrically transformed or *warped* to reflect the intraoperative state of the neuronavigit the non-rigid deformation of the brain is accurately estimated in a principled manner. There exist several approaches for estimating brain shift that involve the preoperative image and intraoperative scanning using an interventional imaging modality.

Recently an experimental clinical study performed at UCL and the National Hospital for Neurology and Neurosurgery (NHNN) in London [Winston et al., 2014] demonstrated a clinical benefit of using neuronavigation updated based on reference scans acquired on an iMRI scanner during 21 surgeries performed for management of pharmacologically refractive temporal lobe epilepsy (TLE). Technologically their work was centered around a brain shift estimation step that took additional advantage of a limited multi-modal imaging capability of a high-field 1.5 T iMRI scanner [Daga et al., 2012]. Their pipeline also involved a correction of MRI magnetic susceptibility artefact [Daga et al., 2014], which is severe for iMRI data and which requires correction in order to facilitate reliable brain shift estimation.

In the present thesis, I investigate whether accuracy of neuronavigation based on iMRI imaging

can be improved in the vicinity of the resection boundary, as this is the most relevant brain area in terms of clinical importance. I explore this point with focus on the MRI susceptibility artefact, which can only be corrected based on data which may itself suffer from noise near the resection boundary. I also investigate the influence of the resection cavity in the context of brain shift estimation, as it may be challenging to delineate the resection cavity accurately in intraoperative iMRI scans. The rest of this chapter provides a more detailed introduction and an outline of the contributions of my work.

1.1 Pathologies Requiring Neurosurgical Management

Epilepsy is a common and debilitating neurological disorder that develops in 50 in 100,000 people per year. In approximately third of these people, antiepileptic drugs do not control seizures. About half of these latter people have focal epilepsy, in which seizures spread from a spatially limited epileptogenic zone (an epileptic focus) in the brain. The epileptic seizures in these individuals can potentially be resolved by neurosurgical treatment if the patient is eligible for the surgery based on clinical and safety criteria and if the patient accepts the risks associated with the surgery [Duncan et al., 2016].

Anterior temporal lobe resection (ATLR) is an established treatment for temporal lobe epilepsy, in which the epileptogenic zone is located in the temporal lobe of the brain. A complication of ATLR commonly observed post-surgery is the visual field deficit (VFD) in the contra-lateral visual field caused by damage to the *optic radiation* nerve fibre tract whose location exhibits large interpatient variability and whose part, called the Meyer loop, shown in Fig. 1.1(a), can pass through the area affected by ATLR for some patients [Daga et al., 2012, Duncan et al., 2016], as shown in Fig. 1.1(b). The VFD can reduce the quality of life for the seizure-free patient that limits their return to full activity and therefore, it is highly desirable to limit any damage to the optic radiation tract.

Malignant brain tumours are cancers with extremely poor prognosis, which are further feared due to their adverse effects on quality of life and neurological functions. **Malignant gliomas** form 80% of the malignant brain tumours and had an annual incidence of 5.26 per 100,000 of



Fig. 1.1: The visual system in the human brain. (a) The paths of the fiber tracts to the visual cortex (courtesy http://www.thebrain.mcgill.ca). (b) The morphology of the human brain (courtesy Virtual Hospital) with the part resection during the ATLR procedure highlighted.

population (in the US in 2005–2009) [Dolecek et al., 2012]. The more common malignant glioma is glioblastoma (World Health Organisation (WHO) histological grade IV) and the less common is anaplastic astrocytoma (WHO grade III). Fig. 1.2 shows the appearance of a malignant glioma in diagnostic MRI and demonstrates a possible size of a growing tumour. The diagnostic protocol usually consists multi-modal morphological MRI imaging composed of T1-weighted imaging enhanced with a contrast agent such as gadolinium, T1-weighted imaging with the FLAIR pulse sequence (FLluid Attenuation Inversion Recovery) that nulls fluids, susceptibility-weighted imaging that is sensitive to potential hemorrhagic components of the tumour and perfusion imaging using the dynamic contrast enhancement imaging approach [Omuro and DeAngelis, 2013].

The recommended therapy consists of symptomatic treatment followed by neurosurgical management, further followed by a combination of chemotherapy and radiotherapy. The median survival time is only 15 months for glioblastoma and 2–3 years for anaplastic astrocytoma [Omuro and DeAngelis, 2013]. The aims of neurosurgery are relief of the mass effect, cytoreduction (the removal or reduction of the tumour mass) and for histological sampling. Due to the aggressively infiltrative nature of glioma, resection is not expected to prevent secondary tumours. A growing body of evidence suggests that a maximum extent of tumour-volume resection



Fig. 1.2: The appearance of a malignant glioma in (a) a T1-weighted structural scan and (b) a T1-weighted FLAIR scan. The images are taken from preoperative imaging performed prior to the surgery for the patient #3 from the glioma registration dataset described in Section 4.4.3.

is a positive predictor of patient outcome [Kuhnt et al., 2011] but this has to be balanced with the preservation of cognitive and neurological functions. The eloquent (functionally-specific) cortical areas that are removed as part of the surgical access to the tumour site are the most direct source of functional loss, which depends on the location of the tumour and the function associated with the resected cortical region. As in the case of ATLR, resection near eloquent cortical areas may involve disruption to nerve fibre tracts. For instance, the corticospinal tract, which connects the spinal cord with the motor cortex and the somatosensory cortex, has been studied extensively in the context of neurosurgical planning in patients with gliomas [Nimsky et al., 2005, Essayed et al., 2017].

1.2 Use of iMRI Guidance in an Operating Room

The iMRI suite at the National Hospital for Neurology and Neurosurgery (NHNN) is an example of a next-generation OR, in which the iMRI 1.5 Tesla MAGNETOM[®] Espree scanner (Siemens, Erlangen, Germany) is installed directly in the OR, and is fully integrated with the Brainlab VectorVision[®] Sky IGNS system [Gumprecht et al., 1999]. The setup of the OR is shown in Fig. 1.3. The operating table is located beyond the 5 Gauss field line of the scanner, which enables the surgery to be performed using standard surgical instrumentation and equipment. The operating table can be swivelled and docked with the gantry of the iMRI scanner and subsequently



Fig. 1.3: The INGS suite at the NHNN in London with an integrated iMRI scanner. (Image courtesy of Dr Pankaj Daga.)

the patient can be moved into the scanner for imaging. In a recent study, the scanning protocol was kept to approx. 30–40 minutes and the transport of the patient from the operating table to the scanner and back took 7 to 10 minutes either way [Winston et al., 2014].

The closed-bore high-field 1.5T iMRI design is capable of structural scanning at approx. $1.1 \times 1.1 \times 1.25$ mm resolution with relatively low signal to noise ratio (SNR). The Brainlab data processing workflow can download image data from the iMRI scanner and also facilitates real-time instrument tracking and visualisation. The tracking is facilitated by an infrared optical stereo camera, which establishes a 3D coordinate system of the operating theater. Prior to craniotomy, the surgical team fixates a head holder onto the patient's head using titanium pins at several points in the skull. The head holder doubles as an MRI receive/transmit coil but also contains an arrangement of spheres reflective in the infrared spectrum, which form a reference for the tracking system. Since these spheres are also visible as white outlines in the iMRI scan, a rigid geometric transformation between the virtual image space and the head holder (the skull) can be established in a fully automated fashion with high accuracy.

The most recent iMRI structural scan is visualised using a three-plane view on the heads-

CHAPTER 1. INTRODUCTION

up display whose viewing plane intersection point or the "current view" can be adjusted using a tracked pointer. Crucially, heads-up display view also includes outlines of segmented areas of interest established as part of surgical planning. For instance, in ATLR procedures performed at the NHNN, the outlines of the optic radiation tract (Section 1.1) are displayed (an in-plane outline and a projection of maximum extent outline). An alternative means of visualisation is provided by the surgical microscope (OPMI Pentero, Carl Zeiss, Germany) which the surgical team use during advanced stages of surgery. The optical head of the microscope is tracked and the binocular input is digitally recorded and re-projected to the eyepiece, and the BrainLab system augments the video with outlines of segmented areas in a manner similar to the heads-up display.

The 1.5T field of the iMRI scanner also in principle enables multi-modal imaging in the form of DW-MRI data and fMRI. However, in practice, the use of these two modalities intraoperatively has so far remained limited. In the case of DW-MRI, the relevant information is a reliable outline of the tract expected to pass through or near the resected area. Determining an accurate tractography of DW-MRI data requires the interaction of an expert radiologist and also demands data processing times that are prohibitive for the intraoperative setting. While DW-MRI data acquired intraoperatively using iMRI were used successfully to reveal a coarse mapping of voxels with a high assumed content of nerve fibres, this relied on simple modelling assumptions and was thus only use in order to complement the use of precomputed tractography that was itself derived from preoperative imaging [Daga et al., 2012]. Further, the feasibility and practicality of performing intraoperative fMRI are still very open research areas and to my knowledge, this modality is not being used in the iMRI suite. Should this technique be developed in the future, the main assessment criteria of its benefit would be whether it could be used to locate brain activation areas in their intraoperative locations with a sufficient accuracy and reliability. However, clearly any imaging of brain activation areas that would require awaking the patient from anaesthesia would be bound to considerably increase the complexity of the surgical procedure. Thus, in summary, the locations/extent of relevant tracts and/or task-specific brain activation areas can be considered as information that is part of preoperative surgical planning only.

The boundaries of the critical anatomical structures are presented to the surgeon as outlines rendered by the IGNS system. While an iMRI scan acquired during surgery implicitly presents

CHAPTER 1. INTRODUCTION

an up-to-date structural model of the neuroanatomy, the assumption by the BrainLab system of a rigid relationship between the patient's head and the virtual image space does not change. Therefore, the boundaries of critical structures in the IGNS system retain locations as they were in the rigidly registered preoperative image.

At the time of the writing there are two competitors that dominate the IGNS market: Brainlab with its IGNS portfolio (which contains more recent models) and the Medtronic (Louisville, CO, USA) StealthStation[®] range (with StealthStation[®] S8 the current flagship system). Both competitors primarily offer stand-alone rack-based display systems that can be installed in any OR. To the best of my knowledge, their offerings contain rigid registration of preoperative images only.

1.3 Key Challenges in iMRI-Guided Neurosurgery

In the following I will outline in more detail the sources of error to neuronavigation in the context of iMRI image guidance. While some of the following reiterates what has already been outlined above, it is now possible to describe the phenomena involved in more detail.

Brain shift is an effect associated with both keyhole neurosurgery (limited surgery through burr holes in the skull) and open neurosurgery (extensive surgery following craniotomy) and is a collective term for the biomechanical deformation of the brain due to cerebrospinal fluid (CSF) drainage, gravity, edema and tissue resection and retraction. Fig. 1.4 illustrates the situation for two representative cases of ATLR. In a classical iMRI study by [Nimsky et al., 2000] the authors report cortical surface displacement of up to 24 mm for 62 cases of open neurosurgery for tumour removal. While brain shift occurs primarily in the direction of gravity, complex and irregular displacements also occur perpendicular to the direction of gravity.

As mentioned earlier, brain shift presents a significant complication even in the presence of neuronavigation as surgical planning information is presented based on rigid registration only and thus becomes offset. Therefore a dense estimate of the deformation due to brain shift is needed in order to propagate the surgical planning information. However, this estimation problem is challenging due to a number of issues. There are highlighted on the T1-weighted MRI image pair



Fig. 1.4: Brain shift for two representative cases of ATLR. Top row, patient #1. Bottom row, patient #9. The ATLR dataset is described in Section 3.2.4. Left, preoperative structural T1-weighted scan acquired about one month prior to surgery, shown affinely pre-registered to the intraoperative scan. Right, intraoperative scan acquired using an closed-bore iMRI scanner in the operating room with the patient's head in the intraoperative orientation. Highlights A–E are explained in the main text.

in Fig. 1.4 and can be summarised as follows:

- The most obvious presentation of brain shift is the **severe non-rigid deformation** (see highlights C and D in Fig. 1.4) which occurs as the brain tissue is deforming due to the absence of the pressure of the CSF. Moreover, the deformation is most prominent near the resection cavity as the soft tissue sags due to the missing tissue.
- There is a complex **intensity relationship** between the image pair. The global relationship is affected by the increased noise in the iMRI dataset and also by the fact that different scanners and pulse sequences (even within sub-modalities such as T1-weighted MRI) are used. Further, local contrast changes occur due to CSF leakage (Fig. 1.4, highlight C); bleeding near the resection margin (Fig. 1.4, highlight B); and the bias field artefact (intensity varying



Fig. 1.5: Missing correspondences. Left and middle: Example pre/intraoperative image slice (patient #9 from the ATLR dataset, Section 3.2.4). Right, top row: A schematic phantom with landmarks. Right, bottom row: arrow ends mark each landmark's corresponding location in the other image. Preoperative landmarks E and F have no correspondence in the intraoperative space. The other landmarks do not move and so map to the same location.

across the image due to variations in power of received/transmitted MRI signal).

- The tissue missing in the intraoperative image due to resection (Fig. 1.4, highlight A) gives rise to **missing correspondences** in the mapping from the preoperative to the intraoperative image space. In a strict sense, the voxel locations in the preoperative image inside the tissue that is later resected do not map to any voxel in the intraoperative image; this is shown schematically in Fig. 1.5. In a looser sense, the said voxels map to some arbitrary locations in the resection cavity whose exact coordinates are irrelevant.
- CSF leakage can lead to a collapse of the ventricles and the consequent touching of the inner surfaces of ventricles (Fig. 1.4, highlight D), which introduces **folding** as multiple points in the preoperative image map to the same point in the intraoperative image at the interface of the touching surfaces and the volume inside the ventricle in the preoperative space folds into a surface in the intraoperative space (as in a paper-folding analogy implying differing space continuity of either space). This is schematically illustrated in Fig. 1.6. CSF leakage can also lead to touching surfaces of adjacent sulci (Fig. 1.4, highlight C).
- Neurosurgery potentially involves the use of the "retractor", which is a firm planar instrument (a blade applied along its face) that can be fixed to keep part of brain tissue pulled aside



Fig. 1.6: Volume folding due to partial collapse of the ventricles. Left and middle: Example pre/intraoperative slice (patient #9 from the ATLR dataset, Section 3.2.4). Right, top row: a schematic of volume between landmarks D, E, F in the preoperative space collapsing into a line segment in the intraoperative space. Right, bottom row: arrow ends mark each landmark's corresponding location in the other image.

to allow better access to deep brain structures. The gap, as seen in the intraoperative space, maps onto an surface/interface in the preoperative image. This is schematically illustrated in Fig. 1.7. Retraction also alters the stresses in the tissue: while at the preoperative time point the stresses are at equilibrium, at the intraoperative time point there will be a compression on the pushing side of the retractor and a relaxation on the other side [Miga et al., 2001].

• Datasets that include tumours may exhibit **tumour mass effect** i.e. the expansive pressure of the tumour on the surrounding structures that can potentially induce complex and variable tissue displacement within and outside of the tumour margin dependent on the specific shape, composition and location of the tumour. The mass effect in the intraoperative time-step is likely to be significantly altered for an exposed or partially/fully resected tumour.

Imaging Artefacts in iMRI. The high-field of the closed-bore iMRI scanners enables good quality structural imaging at $1.1 \times 1.1 \times 1.25$ mm spatial resolution. However, the design constraints of iMRI imply a lower image quality than that of conventional MRI, which results in lower SNR and the presence of artefacts:

• The gradient-nonlinearity artefact is a geometric distortion, which arises due to the wide



Fig. 1.7: Volume folding due to tissue retraction. A schematic phantom of the segment connecting B–D being retracted to the right during surgery. Top row: The volume between F, D, B in the intraop. image folds into a line segment in the preop. image. Bottom row: Arrow ends mark each landmark's corresponding location in the other image.

diameter of the bore of the main magnet in the iMRI design. As a consequence, the field of the magnet is not as uniform as for conventional MRI scanners. This artefact can be completely corrected in a lossless manner using spherical harmonics and a set of parameters for the iMRI scanner which have been established as part of the design and/or commissioning of the scanner [Janke et al., 2004, Glover and Pelc, 1986].

• An advantage of high-field iMRI is the availability of DW-MRI and functional MRI (fMRI). However, these two techniques rely on the *echo planar imaging* (EPI) pulse sequence, which is affected by the *susceptibility artefact*. This artefact arises because EPI is sensitive to small inhomogeneities in the magnetic field of the main superconducting magnet (the B_0 field) that arise near air-tissue boundaries in the head. In intraoperative iMRI, the boundary between the soft tissue of the brain and the air-filled resection cavity leads to severe distortion that coincides with the anatomical area relevant to the surgery.

Time Constraints of Surgery. The scanning protocol at the NHNN is kept to approx. 30 minutes while the transfer of the patient from the operating table to the scanner and back takes approximately 7 to 10 minutes [Winston et al., 2014]. As it is essential not to add to the duration of the surgery, the iMRI processing pipeline must finish during the transfer of the patient to the surgical table. This creates a strict time requirement for all the processing steps. In particular,

the brain shift estimation step should ideally be kept to well under the patient transfer time.

1.4 Aims of the Thesis

The aims of the present thesis are to improve surgical navigation for neurosurgery in the advanced operating room. The direction adopted in this thesis follows from encouraging outcomes of a recent clinical experimental study performed at UCL and the NHNN, whose findings were reported in [Winston et al., 2014]. In this study, 21 patients underwent ATLR under iMRI guidance in which a T1-weighted structural iMRI scan was overlaid with a tractography of the ORT derived from preoperative data. For the first 9 of these patients, the tractography was shown without brain shift compensation. For the 12 subsequent patients, a novel brain shift estimation method based on fast multimodal non-rigid registration by [Daga et al., 2012] was applied to geometrically shift (propagate) a parcellation of the ORT, as illustrated in Fig. 1.8. An important part of the brain shift correction pipeline was also a MRI susceptibility artefact correction step [Daga et al., 2014]. The iMRI guidance led to a significant reduction of post-operative VFD relative to a control group of 44 patients who underwent ATLR without iMRI guidance. The post-operative VFD was not significantly different between the iMRI cases with or without brain shift compensation, possibly due to small number of patients or because the surgeons could guess the shifted position of the ORT without brain shift compensation. However, this work opens the field of intraoperative brain shift estimation based on iMRI to further investigation, in particular with a view toward procedures in which the approach trajectory to the surgical target is less repeatable than in ATLR, or procedures in which the brain shift is especially severe such as on gliomas. Therefore, the aims of this thesis are to enhance the steps in the above pipeline in order to improve the accuracy of iMRI guidance.

1.5 Methodological Contributions of the Thesis

This thesis presents two main methodological contributions:

• A scheme for the correction of the susceptibility artefact that arises in DW-MRI datasets



Fig. 1.8: A schematic of the image registration part of the software pipeline used clinically at the NHNN during the study of [Winston et al., 2014] for brain shift estimation based on multi-channel registration. The susceptibility artefact correction step is not shown here.

acquired intraoperatively during ATLR surgeries. The scheme is based on simulating the B_0 field map from an intraoperative iMRI structural scan. I demonstrate that it is feasible to avoid the need for field map acquisition by the iMRI scanner, which has a reduced reliability near the resection boundary. The scheme is presented in Chapter 3.

• A non-rigid registration scheme for the estimation of brain shift from an intraoperative iMRI structural scan and an preoperative MRI structural scan. The scheme employs a local similarity measure with an edge-preservation property. I demostrate on clinical data from ATLR and glioma tumour resection surgeries that the proposed method increases registration accuracy near the resection boundary and provides an invariance to the intensity variations that arise between pre/intraoperative time-steps. The scheme is presented in Chapter 4.

The chapters dedicated to these contributions are followed by a discussion, in Chapter 5, of research directions that could be pursued to build on the presented work.

Chapter 2

Review of Image Guided Neurosurgery

In this chapter I describe the work on brain shift estimation in the literature. Section 2.1 describes the imaging technology that is used to image the intraoperative brain deformation. Section 2.2 describes approaches to estimating brain shift based on biomechanical modelling (physical simulation) constrained by sparse imaging data and a discussion of the basic limitation of these models. Section 2.3 describes estimation based on non-rigid registration.

2.1 Intraoperative Imaging of Brain Shift

Brain shift has been recognised in neurosurgery since its inception as a clinical discipline. An early attempt to assess brain shift was by [Kelly et al., 1986] who attached small metal spheres to cortical landmarks and observed their displacement. [Hill et al., 1998] sampled cortical landmarks with a tracked pointer and measured their motion, but could not observe movement in deep structures.

Interventional MRI is an established technology in clinical use at a small number of facilities around the world since the mid-1990s. A good review of iMRI technology is given

CHAPTER 2. REVIEW OF IMAGE GUIDED NEUROSURGERY

in [Hall and Truwit, 2008]. In the early 0.5T Signa double-coil design, which was in operation at the Brigham and Women's Hospital (Boston, Massachusetts), used a pair of coils separated by a 56-cm gap, which enabled access to the patient [Black et al., 1997]. However, long imaging times were required which offset the advantage of easy access to the patient, and the images had a poor SNR and were distorted due to field-inhomogeneity. Closed-bore high-field (1.5T and 3T) systems enable imaging with favourable SNR and minor (fully corrigible) fieldinhomogeneity and enable the acquisition of DTI and fMRI. The high-field mandates that surgery be performed beyond the 5-Gauss line. In one design of the advanced operating room, the surgical table can be operated by the surgical team and can dock with the gantry of the scanner, thus allowing the transport of the patient from beyond the 5-Gauss line into the bore of the scanner [Hall and Truwit, 2008, Winston et al., 2014]. According to [Hall and Truwit, 2008], the 5-Gauss line requirement can potentially be relaxed, as they reported that advances in MRcompatibility of instrumentation and equipment enabled them to perform minimally-invasive and open cranial procedures in the immediate vicinity of the bore of a 3T scanner. An alternative design of the operating room, which is implemented for instance in the new Advanced Multimodality Image Guided Operating suite at the Brigham and Women's Hospital [Jolesz, 2011], involves an iMRI scanner suspended on a rail that can be moved over the patient for scanning and also out of the theatre when not needed, which increases the utilisation of the scanner. The duration of patient transport and of iMRI scanning have to be considered when planning the surgery. As mentioned in Section 1.2, a combined structural and diffusion scanning protocol can take almost an hour (including patient transport). However, this disadvantage has to be weighted against the strong benefits of iMRI imaging.

Interventional MRI is regarded as the gold-standard for visualising and quantifying brain shift. [Nimsky et al., 2000] performed the first large-scale iMRI study using a 0.2T open-bore scanner (n = 64; 32 were gliomas), which revealed cortical shifts of up to 24 mm and shifts of deep tumour margins in excess of 3 mm for 66% of cases. The complex nature of brain shift-induced deformations became clear. This was corroborated by [Nabavi et al., 2001] who performed serial scanning (n = 25, 4 or more time-points per case) at 0.5T and noticed complex motion. A follow-up study [Nimsky et al., 2001] demostrated the benefit of iMRI for resection control. Considerable
brain shift was observed in deep brain stimulation (DBS) electrode implantations through frontal burr holes by [Ivan et al., 2014] at 1.5T (n = 44) who reported shifts of 0.0–10.1 mm for landmarks throughout the brain and shifts in excess of 2 mm in deep brain structures for 9% of cases.

In addition to iMRI, several sparse imaging modalities have emerged in the context of neurosurgery, namely intraoperative ultrasound (iUS), stereovision (iSV) and laser range scanning (iLRS). These modalities are minimally distruptive to the surgical procedure. 3D iUS imaging using a probe touching the exposed brain surface provides imaging of structures below the surface such as ventricles [Comeau et al., 2000] and tumour margins [Unsgaard et al., 2005]. iUS in the Doppler mode was used to image vasculature which can be used opposite to MRI angiography [Reinertsen et al., 2007]. While iUS image quality is poor, suffers from heavy artefacts and the origin of contrast is different from MRI, [Wein et al., 2013] and [Rivaz et al., 2015] aimed to locally match intensities of iUS and structural MRI volumes with reasonable results. iSV is based on reconstruction of brain surface using a pair of calibrated cameras [Skrinjar et al., 2002, Sun et al., 2005, Paul et al., 2009, Ji et al., 2014] and is ideal for workflow integration with the surgical microscope. iLRS [Miga et al., 2003, Ding et al., 2011, Miga et al., 2016] probes acquire intensity-encoded point clouds of the exposed surface. Surface-based methods require registration with brain surface from preoperative imaging [Miga et al., 2003, Sun et al., 2005]. Some authors relied on tracking surface features in video to maintain accuracy [Paul et al., 2009, Ding et al., 2011]. Recently, [Ji et al., 2014] presented a method for reconstruction of surface using microscope even after its focus has been readjusted. Since iSV and LRS provide only cortical surface, they are mostly used to constrain biomechanical models which require fewer and less dispersed data-points than volumetric 3D registration.

2.2 Biomechanical Modelling Based Brain Shift Estimation

Biomechanical models attempt to physically model the brain. The preoperative image is meshed into tetragonal finite elements and physical modelling is performed with each element treated as interacting with its neighbouring elements. The models are run in iterative fashion until an equilibrium condition is met. Intraoperative imaging is used to provide boundary conditions or otherwise inform the behaviour of the model.

Basic models treat the tissues in the brain as elastic material [Ferrant et al., 2001, Škrinjar et al., 2002] that follows Hooke's laws. [Paulsen et al., 1999] presented a biphasic model in which each element contains an elastic phase, which represents the solid/cellular tissue that can deform instantaneously, and a viscous fluid phase, which represents a transient flow of the extracellular CSF based on pressure gradient. This original model was extended to account for gravity sag and resection and retraction. Follow-up work introduced increasingly complex models that included the effect of gravity sag, resection and retraction, the brain swelling effects of edema, influence of mannitol that is often administered to manage swelling.

An important part of the model is how the sparse intraoperative data is introduced. For instance, [Ferrant et al., 2001] and [Škrinjar et al., 2002] introduce the brain surface as a given boundary condition, which has the disadvantage that an inaccurately extracted surface may confuse the model. [Lunn et al., 2005, Ji et al., 2009] attempted to allow for inaccurate data by using relaxed models that better match sparse data. For similar purposes, [Dumpuri et al., 2007] presented a method that pre-computes an atlas of brain deformations for a set of common parameters (head angle, degree of CSF evacuation) and then attempts to estimate a deformation that best fits the observed data points. An advanced degree of integration with sparse modalities has been achieved. [Ji et al., 2009] used sparse iUS data, [Ji et al., 2014] used iSV data and [Miga et al., 2016] used iLRS data to constrain their models. These models have been able to provide estimates within the time-constraints of brain shift.

Clearly, biomechanical modelling offers compelling methods for brain shift estimation. However, these models suffer from the reliance on many steps and parameters which may potentially not be accurate but which influence each other. For instance, any tissue segmentation in the presence of pathology or brain surface extraction steps have limited accuracy; the mechanical parameters of tissue, CSF, tumour mass and edema are only known approximately; the motion is constrained by the falx cerebri and the tentorium cerebelli, which require more complex models; amount of CSF drainage and/or ventricular collapse are highly unpredictable; the extent of edema or response to administred mannitol may be difficult to predict. Thus, while modelling using complex models is certainly valuable and accurate in a range of situations, there are clearly many unknown and coupled parameters.

In the next section, I will explore the technique of non-rigid registration in general and in particular its use for estimating brain shift. This methodology is suitable for situations when there is a good coverage of data to drive the process, which is primarily the case for iMRI and, to a lesser extent, for iUS. Potentially in such circumstances, running a simpler method which is nevertheless regularised in a principled manner can produce results that are more robust to unpredictable physical behaviour. However, I envisage that in the future there will be more fusion between biomechanical modelling and non-rigid registration based-methods.

2.3 Registration Based Brain Shift Estimation

In registration, a geometric transformation is sought from the space of a *floating image* into the space of a *reference image* such that the transformation takes into account the physical correspondences between the image pair: a geometric transformation is sought such that the locations in the floating image that are in physical correspondence (also, are *homologous*) with locations in the reference image are moved until they become aligned with the locations in the reference image. The result of the registration is the geometrically transformed (also, *warped*) floating image, or alternatively, the found geometric transformation. Image registration is a very active research topic in the medical image analysis community. A comprehensive recent review of the topic is given in [Sotiras et al., 2013]. In the following, I present a limited review aimed at registration-based brain shift estimation.

In the context of neurosurgery, one has the preoperative (further, preop.) and the intraoperative (further, intraop.) MRI volumes of the same patient. The choice of reference and floating image is ambiguous as illustrated in Fig. 2.1: both preop. volume and intraop. volume can in theory be chosen as the registration reference while the other volume is warped into alignment with the reference. However, a choice of the intraop. image as the reference has the advantage that it maps every voxel of the brain in the preop. image to a location in the intraop. image that can be either inside the brain or inside the resected area. This allows for the portion of the brain to be delineated in the preop. image which is resected in the intraop. image. Consider, for instance,



Fig. 2.1: Registration for an ATLR dataset. The input volumes are: (a) the preoperative T1w with the head in the standard orientation and (f) the intraoperative T1w volume with the head in the intraoperative position. The results are: (b) preop. registered to intraop. affinely i.e. using a global transformation, (c) preop. registered to intraop. non-rigidly, (d) intraop. registered to preop. non-rigidly, (e) intraop. registered to preop. affinely. Note: all images are 2D slices through 3D volumes.

the correspondence between the points marked by the crosshairs in the image pair in Fig. 2.1(c,f)and the fact that the locations in the warped preop. image (c) that are aligned with the resection cavity in the intraop. image (f) represent the portion of the temporal lobe that has been resected during the surgery.

The utilised geometric transformation establishes a mapping from the reference to the floating image space. Let R be the reference image $R : \Omega_R \to \mathbb{R}^M$ where the discrete domain $\Omega_R \subset \mathbb{R}^3$ consists of a uniformly-spaced grid of points and where M is the number of image channels. Similarly, let $F : \Omega_F \to \mathbb{R}^M$ be the floating image. The discrete nature of Ω_R and Ω_F implies that one can regard either image as a discrete sampling of an underlying intensity function defined on a continuous space \mathbb{R}^3 , whereby the intensity of a voxel is the intensity at the voxel centre. The goal of registration is to find a mapping $t : \Omega_F \to \mathbb{R}^3$ that ensures that for every image grid point in the discrete domain Ω_F , its corresponding location in the reference image is found. The range \mathbb{R}^3 of the mapping t is not restricted to grid points Ω_R but also includes the space between the grid points in order to enable the estimation of sub-voxel motion.

Once the transformation is found, it remains to find the intensities of the warped floating image sampled at the discrete grid points Ω_R of the reference image (also, to resample the floating image into the reference image grid); in the case of a multi-channel floating image (M > 1) the resampling step is performed for each channel individually as the channels are in implicit alignment with each other. However, the resampling step is not straightforward to perform using the transformation t because of the aforementioned permitted sub-pixel motion. Namely, t does not map to points restricted to be on the image grid $\vec{x} \in \Omega_R$ of the reference image but to locations that are potentially between the grid points. Therefore, there is a need to invert the transformation t, which cannot in general be done in an unambiguous way. On the other hand, if the registration problem is formulated as seeking the backward transformation $T : \Omega_R \to \mathbb{R}^3$, then for every grid location $\vec{x} \in \Omega_R$ in the reference image, a general location $T(\vec{x})$ is naturally assigned in the geometric space of the floating image. Subsequently, the intensity of the warped point is found by interpolating the intensities of the floating image grid points that neighbour the mapped point $T(\vec{x})$. In the following, the latter formulation using the backward mapping T will be used to remain consistent with a widespread convention in the registration literature.

The transformation T models possible ways of aligning images. The simple models are global, such as rigid transformation, illustrated in Fig. 2.3, which models translation and rotation of the floating image. Another global model is the affine transformation, illustrated in Fig. 2.4, which models translation, rotation, scaling and shearing. Global transformations can model alignment between rigid objects (e.g. the skull, ribs) and/or are useful for initialisation: e.g. brain shift is relatively mild away from the resection as illustrated in Fig. 2.1(b,e). As mentioned in Section 1.2, a limitation of IGNS systems is that they only perform affine registration. However, soft tissue deformations e.g. brain shift in the vicinity of resection, can clearly only be modeled as local non-linear transformations or *non-rigid deformation*, as illustrated in Fig. 2.5 and Fig. 2.1(c,d). Models for representing transformations are reviewed in Section 2.3.1 and all involve a vector of parameters, its length determining the degree of freedom (DOF).

Correspondences and/or similarity between the reference and floating image can only be iden-



Fig. 2.2: For the convenience of a simpler resampling of the warped floating image, the registration problem can be posed so that the backward transformation T is sought which maps each reference image voxel grid point \vec{x} to a location $T(\vec{x})$ in the geometric space of the floating image that is allowed to lie between the grid points of the floating image. The sought intensity of the point in the warped floating image can be found by interpolating the intensities of the floating image grid points that neighbour $T(\vec{x})$. (Image courtesy of Dr Marc Modat.)

tified in areas with relatively clear and unique (also, *salient*) features e.g. the cortical folds in Fig. 2.1(a,f). Conversely, in homogeneous areas such as the white matter in structural images, the correspondences are not identifiable or are ambiguous in terms of location. This implies that multiple transformations satisfy the correspondence criterion, which makes registration an ill-posed problem.

Registration can be posed as an optimisation problem in which the aim is to find a geometric transformation that is optimum in the sense that it achieves the best similarity between the reference and the transformed floating image while being in some way biologically plausible. There are broadly two algorithmic approaches to this optimisation problem described in the registration literature [Sotiras et al., 2013]. In *feature-based* registration, the registration begins with the detection of highly salient areas (*features*) in either image; these sparse features inform subsequent steps. In *intensity-based* registration, all locations in either image are considered as relevant or "feature-rich" in principle and it is considered necessary to evaluate the similarity between both images in such a way that takes into account all the image voxels. Thus intensity-registration reg-



Fig. 2.3: Rigid transformations of a cube involve translations and rotations. (Image courtesy of Dr Marc Modat.)



Fig. 2.4: Affine transformations of a cube involve translations, rotations, scaling and shearing. (Image courtesy of Dr Marc Modat.)

istration scheme address the optimisation problem by combining the following components: Hence every *intensity-based* registration scheme consists of:

- The **geometric transformation** as outlined above.
- The **similarity measure** that reveals the degree of similarity between the reference and the transformed floating image. Every point in either image is assumed to be equally relevant.
- The **regularisation** of the transformation that ensures that the transformation behaves in a biologically plausible manner.

Due to the volume of the literature devoted to image registration, I only review intensity-based registration methods and omit their qualification as such for brevity.

Formally, registration can formulated as an optimisation problem whereby a transformation is



Fig. 2.5: Two examples of non-rigid transformation of a cube using distinct deformations. (Image courtesy of Dr Marc Modat.)

sought \hat{T} that maximises a functional composed of two energies (costs),

$$\hat{T} = \operatorname*{argmax}_{T} \lambda_{S} \cdot E_{S}(R, F(T)) - \lambda_{R} \cdot E_{R}(T), \qquad (2.1)$$

where $E_S(R, F(T))$ is the similarity measure energy, $E_R(T)$ is the regularisation energy and λ_S and λ_R are the respective weights of these terms. The transformation T is in practice represented using a vector of parameters.

The optimisation of Eq. 2.1 is typically performed using gradient-based continuous optimisation. This is possible if the gradient of the terms $E_S(R, F(T))$ and $E_R(T)$ with respect to the transformation parameters can either be derived analytically or is available in approximate form using finite derivatives. Non-linear optimisation methods such as gradient ascent, Newton's method or conjugate gradient are applicable. The are two major disadvantages of continuous optimisation: firstly, deriving the analytical gradient of the similarity term $E_S(R, F(T))$ can be challenging for more complex similarity measures, and secondly, non-linear optimisation search can only find a local optimum and relies on a good initialisation. In the case of non-rigid registration such initialisation can come from affine registration or from feature-based registration. The latter is a conceptually different approach to registration compared to the general approach which relies on matching potentially quite distant and typically sparse features (highly salient regions) and subsequently interpolating or approximating the deformation field; an overview is included in [Sotiras et al., 2013].

An alternative approach to solving the non-linear optimisation problem is to employ discrete optimisation, in which the parameters that specify the geometric transformation are quantised and thereby rendered discrete-valued. The use of this approach was first proposed in the context of non-rigid registration by [Glocker et al., 2008a] and it relies on methods such as Markov Random Field (MRF) optimisation [Glocker et al., 2008a, Heinrich et al., 2013]. The gradient of the similarity measure is not required, which enables the use of a wider range of similarity measures. Discrete registration in theory aims to find the global optimum but even for a moderate number of parameters (in the hundreds or thousands), an exact solution is intractable due to combinatorial explosion. Thus authors used various heuristics; this is outlined at greater length in Section 2.4.2.

2.3.1 Geometric Transformations

Global (rigid and affine) registration was introduced above. [Ourselin et al., 2000] performed rigid registration by calculating a local displacement map using block-matching. First an exhaustive search is performed at every relevant reference image position for the displacement, out of an allowed set, with the best similarity between the block pair. Then the translation and rotation parameters are estimated from the displacement map. While [Ourselin et al., 2000] used this scheme to register histological sections, it generalises trivially to 3D and affine transformation. [Modat et al., 2014] presented a symmetrical extension of this scheme.

As per above, non-rigid transformation is required for brain shift estimation. The simplest representation is the dense *deformation* field,

$$\vec{U}(\vec{x}) = \vec{x} + \vec{u}(\vec{x}),$$
 (2.2)

where the term "deformation" refers to absolute coordinate mapping, the adjective "dense" (often omitted) refers to the transformation being precisely given at every voxel (centre) location $x \in \Omega_R$. Note that the transformation values $\vec{T}(\vec{x})$ at locations x here arbitrarily inside the convex hull of Ω_R , i.e. evaluated in *between* the voxel centres, the transformation parameterised using the dense field $\vec{U}(\vec{x})$ can be interpolated similarly as the image intensities themselves. Analogously to the deformation field, the mapping $\vec{u} : \Omega_R \to \mathbb{R}^3$ is referred to as the (dense) *displacement* field as matter of convention. The dense formulation has $3 |\Omega_R|$ degrees of freedom, which leads to increased computational complexity and potentially (in the absence of further regularisation) to over-fitting.

To mitigate this problem, approaches for parametrising the deformation based on a limited number of parameters have been proposed as follows. [Ashburner et al., 1999] used a linear combination of cosine basis functions of multiple frequencies; however, due to the periodic nature of cosine, only few low-frequency basis functions can be used, which only enables coarse registration.

A widely-used and flexible approach to parametrise the deformation is the free-form deformation (FFD), which was proposed by [Sederberg and Parry, 1986] in the field of computer graphics, was introduced into the medical image analysis field by [Bardinet et al., 1996] who used it to model the deformation of the heart, and was first employed for intensity-based registration in the seminal work of [Rueckert et al., 1999]. A grid of control points $\{\vec{\mu}_{i,j,k}\}$ is overlaid onto Ω_R . Each control point acts as the center of a cubic B-Spline whose local influence (i.e. it being non-zero) reaches to within two grid spacings in x, y, z directions. Each control point has a 3×1 vector of coefficients $\vec{\mu}_{i,j,k}$ associated with it, which determines its contribution the component fields $T_x(\vec{x}), T_y(\vec{x}), T_z(\vec{x})$ of the interpolated deformation $T(\vec{x})$.

$$T(\vec{x}) = \vec{x} + \sum_{i,j,k} \beta^3 \left(\frac{x}{\delta_x} - i\right) \beta^3 \left(\frac{y}{\delta_y} - j\right) \beta^3 \left(\frac{z}{\delta_z} - k\right) \vec{\mu}_{i,j,k},\tag{2.3}$$

where i, j, k are the grid indices of the control points, $\delta_x, \delta_y, \delta_z$ are the control point grid spacings along image dimensions, and $\beta^3(\xi)$ is a cubic B-spline kernel with a limited support $-2 \leq \xi < 2$. Note that the B-spline coefficients of the control points are different from the displacements at the control points, as the latter are a result of a linear combination, but an inverse relation exists. The kernel $\beta^3(\xi)$ is piece-wise composed of polynomial basis functions, which are chosen so that the kernel has continuous first and second order derivatives. A further advantage of FFD is that exact coefficients/displacements exist for a grid with twice the resolution (half the spacing).

Both non-parametric and parametric representations of deformation result in a relatively large number of parameters to optimise, which may lead to over-fitting and long running times. For this reason, often a multi-resolution scheme is employed, in which a multi-resolution pyramid is created by downsampling the original images several times; subsequently, it is possible to perform a registration for the lowest resolution initially and use the results as initialisations for subsequent higher resolutions e.g. as in [Heinrich et al., 2014c]. Alternatively, for parametric representations it is possible to use a multi-level scheme, in which registration is performed initially with the transformation represented by a small number of parameters, e.g. with a coarse FFD grid, and the results are used to initialise registrations performed with a higher number of parameters, e.g. with a denser FFD grid; this approach is used e.g. in [Modat et al., 2010, Heinrich et al., 2013] whereby in [Modat et al., 2010] the multi-level and multi-resolution are combined.

Diffeomorphic Registration

For data with no missing correspondences and no topological changes (folding) there exists a true mapping that is bijective for which it holds that $T^{-1}(T(\vec{x})) = \vec{x}$. Hence it is highly desirable to constrain registration to find a mapping that is *diffeomorphic*, which is one that is bijective and also smooth with a smooth inverse. This is a more representative model of the deformation and also ensures that no folding arises in the transformation.

A useful property is that a composition of diffeomorphisms is also a diffeomorphism. [Rueckert et al., 2006] also presented a "hard-constraint" approach, in which they ensured the final transformation of an iterative registration scheme to be diffeomorphic, by maintaining a composition chain of transformations known to be diffeomorphic, i.e. $T_n(\cdots((T_2(T_1(\vec{x}))))))$. Each time an iteration failed a specific diffeomorphism test, the result of the preceding iteration was added to the chain. The relevant test had been presented for the case of cubic B-spline FFD by [Choi and Lee, 2000] who found that, assuming a control point grid spacing δ , iff all displacements at control point positions are below 0.4δ , the transformation is bijective and C^2 continuous. However, in theory, if a diffeomorphic transformation cannot be found immediately after extending the chain, the final transformation will not be diffeomorphic.

An elegant approach to obtain a diffeomorphism is to integrate a velocity field over continuous time. [Beg et al., 2005] presented a method called Large Deformation Diffeomorphic Metric Mapping (LDDMM) in which a velocity field that is non-stationary in continuous time is optimised so as to minimise the SSD similarity cost along with the time integral of the norm of the velocity field. The result is a gradual morph from the floating to the reference image. The continuous time is represented by a large number of discrete short time steps. However, the number of parameters to optimise is extremely large, because a dense velocity field needs to be estimated for every time step.

[Arsigny et al., 2006] observed that integration of a velocity field that is stationary in time is equivalent to an exponentiation operation within Lie group theory and, conversely, that a logarithm of a diffeomorphism constructed in this manner is the original stationary velocity field. They remark that a log-Euclidean metric can be defined for these diffeomorphisms which is the difference between the vector field norms of the logarithms of a pair of diffeomorphism. They also present an efficient "scaling and squaring" method for velocity field exponentiation that requires only N compositions when using 2^N discrete time steps. [Ashburner, 2007] presented a registration algorithm called Diffeomorphic Anatomical Registration using Exponentiated Lie Algebra (DARTEL), which uses the stationary velocity field and the scaling and squaring approach. While the optimisation of a stationary velocity field is less flexible than that of non-stationary field, the optimisation of a single velocity field is much faster.

Symmetric Registration

Symmetry means that the algorithm should produce the same result no matter which of the image pair is chosen as the reference image or the floating image, respectively. This can be stated more precisely. Let A and B be images and let T_{AB} be the result of running a registration algorithm with A (B) as the reference (floating) image, and let T_{BA} be the result of running the same algorithm with B (A) as the reference (floating) image, respectively. Then the algorithm is symmetric if it supports diffeomorphism (i.e. T_{AB} and T_{BA} are both diffeomorphic) and $T_{AB} = T_{BA}^{-1}$ for all $\vec{x} \in \Omega_A$ and $T_{AB}^{-1} = T_{BA}$ for all $\vec{y} \in \Omega_B$.

[Avants et al., 2008] presented a method that relies on maintaining a virtual mid-point image, say, denoted D. Therefore mappings T_{AD} (with associated inverse T_{AD}^{-1}) and T_{BD} (with associated inverse T_{BD}^{-1}) are optimised together; both A and B are resampled to the space of D to enable the use of a similarity measure. This approach can be regarded as a special case of group-wise registration, in which multiple images are collectively registered to a common virtual mid-point

2.3.2 Similarity Metrics

Global similarity measures are the most basic similarity measures and return a *scalar* value for a given pair of the reference image R (or $R(\vec{x})$ in an expanded form) and the (resampled) warped floating image F(T) (or $F(T(\vec{x}))$ in an expanded form). The simplest similarity measure, the sum of absolute differences (SAD), is based on the voxel-wise difference image R - F(T)) and is defined as

$$SAD(R, F(T)) = \sum_{\vec{x} \in \Omega_R} \|R(\vec{x}) - F(T(\vec{x}))\|.$$
(2.4)

A related measure is the sum of squared differences (SSD) defined as

$$SSD(R, F(T)) = \sum_{\vec{x} \in \Omega_R} (R(\vec{x}) - F(T(\vec{x})))^2.$$
(2.5)

However, the assumption that an image pair must be identical to depict the same structure is incorrect. Firstly, every image includes acquisition noise, whereby in MRI the noise is Rician, which is often approximated by Gaussian noise. Secondly, imaging settings may vary between acquisitions and scanners for the same modality such as different echo (TE) and relaxation time (TR) in T1w images, different noise levels etc., which leads to changes in contrast. Thirdly, both SSD and SAD do not have an upper bound that would be easily reproducible for various images, which precludes a meaningful choice of relative weights of the similarity and regularisation terms in Eq. 2.1.

The correlation coefficient (CC) is defined as

$$CC(R, F(T))^{2} = \frac{\langle R, F(T) \rangle^{2}}{\langle R, R \rangle \cdot \langle F(T), F(T) \rangle},$$
(2.6)

where $\langle R, R \rangle$ and $\langle F(T), F(T) \rangle$ are the variances and $\langle R, F(T) \rangle = \overline{R \cdot F(T)} - \overline{R} \cdot \overline{F(T)}$ is the covariance expressed using intensity means \overline{R} and $\overline{F(T)}$, and the variances in Eq. 2.6 are defined

analogously. CC reveals the affine correlation between the images, which makes it invariant to intensity offset and scaling such as can be observed in images of the same structural modality but acquired on different scanners, is well-behaved for moderate noise levels and constrained to the interval [0, 1]. However, the affine relationship is not sufficient to capture the complex intensity relationship of multi-modal image pairs such as T1w/T2w MRI or MRI/CT (X-ray computed tomography) pairs.

[Viola and Wells III, 1997] and [Collignon et al., 1995, Maes et al., 1997] independently introduced mutual information, which builds on the intuition that a tissue gives rise to a limited range of intensities. Let $\mathcal{H}(r, w)$ be a joint histogram that is built by partitioning the allowed intensity levels in the images R and F(T) into discrete bins and counting the number of occurrences when the voxel intensities at a given location jointly belong in the bins r and w in the images R and F(T), respectively. Since the same tissue gives rise to a limited range of intensities, for co-registered images the joint histogram tends to exhibit a reduced dispersion. Let mutual information be defined using the Shannon-Wiener entropy [Shannon, 1948] as

$$I(R; F(T)) = H(R) + H(F(T)) - H(R, F(T)),$$
(2.7)

where H(R) and H(F(T)) are marginal entropies obtained by marginalising over either other image and H(R, F(T)) is mutual entropy based on the joint histogram: H(R, F(T)) = $-\sum_{r} \sum_{w} p(r, w) \log p(r, w)$. The probabilities are calculated from the joint histogram as per $p(r, w) = \frac{\mathcal{H}(r, w)}{\sum r, w \mathcal{H}(r, w)}$. [Viola and Wells III, 1997] proposed to fill the joint histogram using a so-called Parzen Window to account for discrete intensities typical for most acquisitions and [Thévenaz and Unser, 2000] proposed to use a cubic B-spline kernel this purpose, whereby the joint histogram becomes

$$\mathcal{H}(r,w) = \sum_{\vec{x}\in\Omega_R} \beta^3(r, R(\vec{x})) \cdot \beta^3(w, F(T(\vec{x})))$$
(2.8)

where $\beta^3(\cdot, \cdot)$ is the cubic B-spline kernel. [Studholme et al., 1999] observed that mutual information alone is not invariant to a partial lack of image overlap, which can arise frequently in affine registration or when registering images with differing field-of-view: they proposed the normalised

CHAPTER 2. REVIEW OF IMAGE GUIDED NEUROSURGERY

mutual information (NMI), which they defined as

$$NMI(R, F(T)) = \frac{H(R) \cdot H(F(T))}{H(R, F(T))}.$$
(2.9)

The measures mentioned so far (global measures) assume that there exists a constant relationship between tissue class and tissue intensity in the whole image, which is generally not true for MRI datasets because of the so-called "bias field" artefact i.e. local fluctuations in the intensity level of the same tissue across the MRI image based on the physical proximity of any given location to the MRI receive coils. A more fundamental problem with global measures is that they are not ideally suited for images where non-rigid deformation is highly localised, such as due to brain shift. The contribution of a given voxel location \vec{x} to the similarity measure is simply the pair of intensities exactly at the voxel pair at the given location, which do not fully represent the morphological similarity of the neighbourhood centered at the considered location. By contrast, local similarity measures capture some relationship between the voxels in the neighbourhood of a given location in the images R and F(T). [Cachier et al., 2003] introduced a local version of the CC, termed the local normalised correlation coefficient (LNCC), which measures the Pearson correlation coefficient between the voxel intensities and therefore captures an affine intensity relationship. The measure is thus invariant to local intensity level variations such as those that arise due to the MRI bias field artefact. The measure also captures local morphological similarity unless the local deformations are severe. Since the result of this computation is a map $LNCC(\vec{x})$, [Cachier et al., 2003] simply considered its sum over the image domain to derive the gradient with respect to the dense deformation field.

Another limitation of the above measures is that they only reveal similarity between volumes with one intensity value (*channel*) per voxel. The multi-modal capability of MRI gives rise to multi-channel data so that, for instance, the merging of a T1-weighted MRI structural scan and an MRI scan acquired using the FLAIR pulse sequence (Section 1.1, Fig. 1.2) into a multi-channel (4D) volume enables the appreciation of the details of the anatomy alongside the extent of the tumour. A trivial way to employ multiple channels is to apply the similarity measure on each channel separately and then sum up the similarity values, which can be performed for both global and local similarity measures. However, this simple approach has several limitations. Firstly, the additive mixing of similarity values introduces ambiguity about the contributions of individual channels. Secondly, the correlations that exist between different channels within the respective images R and F(T) are ignored, which makes it impossible to evaluate the degree of redundancy between the channels in either image. Thirdly, the cross-correlations between the channels in Rand those channels in F(T) that are not their direct counterparts, are also ignored.

Various authors aimed to find similarity measures that can include some of the aforementioned information. [Daga et al., 2012] extended NMI by adding two dimensions to the joint histogram from the second channel of the reference and the floating image; however, the joint histogram grows prohibitively sparse for more channels. Correlation-based measures are less susceptible to this dimensionality problem. [Wein et al., 2008, Wein et al., 2013] proposed a measure (LCLC, linear combination of local correlation) that produces a similarity value for each voxel location by considering patches i.e. windows surrounding the voxel in the image pair; the returned value is the voxel-wise correlation of a single-channel patch with a least squares-fitted linear combination of the channels of a multi-channel patch. The latter paper showed that this can be useful for registering preoperative MRI and iUS images for neurosurgery. [Heinrich et al., 2014a] used local canonical correlation analysis (LCCA) to find an optimum mapping of floating multi-channel patch into the space of the reference multi-channel patch space where they maximally correlate. [Rivaz et al., 2015] presented a local weight approach to MI based on self-similarity which they demonstrated on an MRI to iUS registration use in neurosurgery.

2.3.3 Regularisation

The final component of registration algorithms is the regularisation part whose purpose is to constrain the resulting transformation so that it is more biologically plausible, especially in homogeneous areas. Note that regularisation naturally arises as a feature of some transformation models: for example, the FFD imposes smoothness proportional to the coarseness of the grid [Rueckert et al., 1999] and a diffeomorphism naturally prevents folding [Beg et al., 2005]. However, virtually all registration algorithms also contain a way of imposing a further soft con-

straint on the transformation to make it less likely that an implausible solution is found. This is achieved by an inclusion of a regularisation term $E_R(T)$ as in Eq. 2.1.

However, separate similarity E(R, F(T)) and regularisation $E_R(T)$ terms do not arise in some algorithms; I will refer to such algorithms to having an "implicit" regularisation scheme. Conversely, I will regard the algorithms in which $E_S(R, F(T))$ and $E_R(T)$ can be written as independent terms as having an "explicit" regularisation scheme.

Implicit Regularisation

Some classical optimisation strategies that are particularly suitable for algorithms operating with the transformation parametrised using the dense deformation field $\vec{U}(\vec{u})$, rely on an intermediate *unconstrained* deformation field, which can be obtained by relying solely on similarity matching, for instance using continuous methods such as the gradient of SSD and/or related methods (e.g. the classical "optical flow" of [Horn and Schunck, 1981]) or by evaluating a local similarity measure for image pair shifted by a discrete set of translations (as e.g. in [Heinrich et al., 2014c]). Regardless of the method employed to find the unconstained field, the regularisation is achieved by processing the field. The two steps are then repeated until the processed field result does not vary anymore significantly between iterations.

A well-known example is the demons algorithm [Thirion, 1998]. Let $\vec{U^n}(\vec{x})$ be the current dense deformation estimate at the start of *n*-th iteration and let $\vec{v^n}$ be the unconstrained deformation field that has been obtained by applying a similarity measure between for the reference image and the floating image warped using $\vec{U^n}(\vec{x})$. Then the regularisation of the deformation is performed by applying an isotropic smoothing filter S (e.g. by convolution with a Gaussian kernel) on the field to yield a "fluid-like" deformation

$$U^{\vec{n}+1}(\vec{x}) = U^{\vec{n}} + S(v^{\vec{n}}), \tag{2.10}$$

or alternatively, by applying the filter on the updated field to yield an "elastic-like" deformation:

$$U^{\vec{n}+1}(\vec{x}) = S(\vec{U^n} + \vec{v^n}).$$
(2.11)

The fluid-like regularisation is more suited to large deformations while elastic-like regularisation is more likely to enforce joint tissue movement in homogeneous tissues. The authors noted an analogy between the smoothing filter and heat-diffusion of the unconstrained field. The algorithm is fast and easy to implement. However, while the smoothing step is inspired by diffusion, no real physical modeling of the deformation field takes place, which renders the working principle relatively tenuous.

[Christensen et al., 1996] modelled the deformation field as a viscuous fluid (with specific material/viscosity constants) that reacts to body forces generated by the image pair at the given iteration; in this case the body forces are simply given by the gradient of the SSD. The registration corresponds to finding the equilibrium of such a model physical system. The solution of the fluid model for a given iteration is based on solving a PDE, which presents a significant computational hurdle in the original work and also in later work that uses a faster Fast Fourier Transform (FFT) based solver [Cahill et al., 2007]. The approach is sensitive to the choice of the viscosity constants and the physical analogy remains tenous as no tissue modeling is performed.

A general limitation of the above methods is that the regularisation of the transformation cannot be decoupled from the application of similarity measure. In particular, it is not possible to write independent similarity and regularisation terms that could be optimised jointly and neither can the regularisation term be extended easily.

Explicit Regularisation

I now review commonly used penalty terms: linear elastic energy, bending energy and the Jacobian determinant. Since these are independent terms, they can also be used in tandem.

Let $J(\vec{x})$ denote the Jacobian matrix of T, sampled at image grid locations:

$$\boldsymbol{J}(\vec{x}) = \begin{bmatrix} \frac{\partial \vec{T}(\vec{x})}{\partial x} & \frac{\partial \vec{T}(\vec{x})}{\partial y} & \frac{\partial \vec{T}(\vec{x})}{\partial z} \end{bmatrix}.$$
 (2.12)

Then the linear-elastic energy is employed as a penalty e.g. in [Ashburner and Ridgway, 2013] as:

$$E_{\rm LE} = \sum_{\vec{x} \in \Omega_R} \| \boldsymbol{J}(\vec{x}) + (\boldsymbol{J}(\vec{x}))^{\rm T}) \|_{\rm F}^2, \qquad (2.13)$$

where $\| \bullet \|_{\mathrm{F}}^2$ is the Frobenius norm. By analogy to elastic-like demons, the linear-elastic energy may enforce joint motion in homogeneous areas but also may hamper the estimation of large deformations.

The bending energy as used [Rueckert et al., 1999] is motivated by the bending energy of a thin (metal) and expresses the overall degree of deformation:

$$E_{\rm BE} = \sum_{\vec{x}\in\Omega_R} \left(\left(\frac{\partial^2 \vec{T}(\vec{x})}{\partial x^2} \right)^2 + \left(\frac{\partial^2 \vec{T}(\vec{x})}{\partial y^2} \right)^2 + \left(\frac{\partial^2 \vec{T}(\vec{x})}{\partial z^2} \right)^2 + 2 \left[\left(\frac{\partial^2 \vec{T}(\vec{x})}{\partial xy} \right)^2 + \left(\frac{\partial^2 \vec{T}(\vec{x})}{\partial yz} \right)^2 + \left(\frac{\partial^2 \vec{T}(\vec{x})}{\partial xz} \right)^2 \right] \right)^2$$
(2.14)

The bending energy tends to prevent the optimisation from clearly implausible deformations. However, it does not prevent local volume expansion (compression).

[Rohlfing et al., 2003] proposed a penalty term based on the map of the determinant of the Jacobian $|J(\vec{x})|$, which describes the local volume change at \vec{x} . A clear consequence is that this tends to limit small volume changes in anatomically homogeneous areas where there typically is no good reason why such changes should arise. Furthermore, since a negative value of the Jacobian determinant implies volume loss, [Rueckert et al., 2006] also propose to use a log-transformed Jacobian determinant term as a soft-constraint to prevent folding.

2.4 Advanced Registration Algorithms

In this section I review more advanced techniques with focus on brain shift estimation using registration and related work.

2.4.1 Modelling Missing Correspondences

As outlined in Section 1.3 and illustrated in Fig. 1.5, the resection gives rise to missing correspondences due to resected tissue in the intraoperative image. Clearly the registration can be improved by modelling the extent of the resected tissue. On the other hand, the problems of brain deformation estimation and of detecting missing correspondences are coupled as either problem complicates the solution of the other.

One approach is to use a pre-segmented binary inclusion mask on the reference image and to ensure that only voxels from the mask influence the similarity measure. For instance [Brett et al., 2001] masked out focal lesions in this manner. More recently, [Daga et al., 2012] included the brain and excluded the resected area/cavity in iMRI structural scan in this way. However, for the pre/intraoperative MRI/iMRI volume pair, binary masking has two limitations. Firstly, an accurate pre-segmentation of the brain in iMRI that excludes the resection is complicated by the presence of the saline solution and surgical gauze in the cavity (Fig. 1.4). Secondly, the knowledge of the extent of the missing correspondence in the preoperative time-step does not follow from knowing the extent of the resection in the intraoperative time-step.

A simple approach was presented in [Ou et al., 2011] who observed that some local areas exhibit high *mutual saliency* between the reference and floating image, while by contrast, areas with homogeneous features or missing correspondences have low mutual saliency. Here, the mutual saliency can be thought of as a location-specific value that determines how much the value of a chosen local similarity measure depends on whether the reference and the floating image are exactly aligned as opposed to when the floating image is artifically slightly displaced. More specifically, an overall ratio is calculated between a set of smaller and a set of larger considered displacements. A high ratio at a given location implies a greater local sensitivity to artificial displacements, and vice-versa. Such mutual saliency values can also be used to weight the influence of voxels when aggregating local similarity measure values into a single global value. The authors used such a weighting in a registration scheme and reported improved registration accuracy near missing correspondence regions. However, the approach cannot distinguish between homogeneous and missing correspondence areas.

CHAPTER 2. REVIEW OF IMAGE GUIDED NEUROSURGERY

A more principled approach is to model the missing correspondences explicitly. In machine vision the analogical problem of optical flow estimation in the presence of occlusions has attracted attention. [Alvarez et al., 2007] proposed a symmetrical estimation scheme that can accomodate occlusions and discontinuities. Following from this work, [Risholm et al., 2009] proposed a symmetrical demons scheme that aims to account for resection (i.e. occlusions) and retraction (i.e. discontinuities). They detect resection by evolving a level set driven by image intensity and gradient and detect retraction by thresholding the strain/stress in the last deformation estimate. They next use an anisotropic diffusion step for the demons update whereby the diffusion is restricted pointing outward from the detected resected area and also across the detected retraction boundary. The paper thus treats multiple phenomena inherent in brain shift. The authors test the method on a synthetic dataset with an inserted resection cavity and a retraction channel, and on a clinical iMRI dataset, in both cases with a visually convincing result. This work has several limitations. Firstly, the authors report the detection methods as sensitive to a choice of parameters. Secondly, the anisotropic diffusion step is reported as very slow and the demons scheme is not extensible (Section 2.3.3). Thirdly, while the level set successfully enforced the spatial consistency of the correspondence labelling, it is dependent on the evolution of the level set and so may potentially not behave robustly for complicated iMRI cases; a similar argument applies to a recent work on MRI/iMRI registration [Chen et al., 2015] that also relies on evolving a level set for the same purpose.

Perhaps the most flexible and principled treatment of the joint problem is afforded by probabilistic inference models. [Periaswamy and Farid, 2006] presented an algorithm that classifies each voxel of the reference image into having a valid/missing correspondence in the floating image, and calculates the similarity measure as the joint log-likelihood of all the voxel pairs, whereby for a valid correspondence voxel pair the error between intensities is assumed to follow a Gaussian distribution, while for a missing correspondence voxel pair the error is assumed to follow a uniform distribution (as an outlier). The authors assume the correspondence labels to be independent. The authors make use of the Expectation Maximisation (EM) algorithm [Dempster et al., 1977] to alternately estimate the labels and the deformation. This work is limited by the use of a poly-affine transformation model, the SSD as the similarity measure, and the fact that no priors are imposed directly on the correspondence labelling.

The latter problem is best solved by using a full maximum a posteriori (MAP) framework in which both the deformation and the labelling parameters are modeled probabilistically. Chitphakdithai *et al.* studied the joint correspondence segmentation and registration problem for pre- and post-resection MRI datasets using the EM algorithm: they imposed an intensity prior [Chitphakdithai and Duncan, 2010] and later a location prior and an MRF-based spatial consistency regularisation of the segmentation [Chitphakdithai et al., 2011]. More recently, [Parisot et al., 2012] presented a similar joint scheme for the problem of glioma segmentation in the presence of heavy mass effect, which builds on the discrete framework of [Glocker et al., 2008a] (Section 2.3.1); the discrete formulation arguably makes it easier to model the joint registration and segmentation problem, as the latter is discrete in nature. Recent work on the analogous problem of joint occlusion detection and optical flow estimation is outlined in [Fortun et al., 2015].

2.4.2 Fulfilling Time Constraints of Neurosurgery

As explained in Section 1.3, the transport of the patient from the iMRI scanner to the surgical table takes 7–15 minutes after the conclusion of the intraoperative scanning. This imposes a limit on the running time for the image registration, which should ideally finish within several minutes.

One way to achieve a speed-up is to take advantage of GPU-based acceleration if particular steps of a registration scheme lend themselves to parallelisation. An example of a parallelisable algorithm is the popular FFD parametrisation of the non-rigid transformation (Section 2.3.1). FFD itself provides a considerable optimisation by limiting the number of parameters that need to be optimised. However, the calculation of the FFD interpolation can become a time bottleneck, which can become especially problematic if the interpolation step is used as part of a non-linear optimisation scheme in which each iteration involves an FFD interpolation. This problem was resolved by [Modat et al., 2010] who implemented the FFD interpolation using the CUDA[®] programming interface to the GPU cards made by Nvidia (Santa Clara, Calif.) [Sanders and Kandrot, 2010]. They also parallelised the NMI similarity measure (Section 2.3.2) in the same fashion and introduced some additional optimisations unrelated to parallelisation. The resulting registration scheme could execute within a few minutes for a pair of 3D volumes with 1 mm voxel resolution. In a subsequent work, [Daga et al., 2012] extended this scheme with a two-channel extension of the NMI similarity measure (Section 2.3.2), for which they also employed CUDA acceleration. They applied the modified scheme to estimate brain shift from MRI/iMRI image pairs and confirmed that the running time satisfied the time constraints of neurosurgery. This paved the way to the deployment of the scheme within an experimental clinical pipeline [Winston et al., 2014]. The clinical value of this pipeline was already outlined in Section 1.4.

Recently, discrete registration has attracted attention due to lack of need for the gradient of the similarity measure, and the fact that it aims to find a global optimum. Due to these advantages, I apply discrete registration to brain shift estimation later in Chapter 4. An exact solution is not tractable due to a combinatorial explosion and various heuristics are employed in the literature to circumvent this problem. [Glocker et al., 2008a] limited the set of displacement labels that were allowed at MRF nodes to those along the normal axes. [Heinrich et al., 2013] used a full 3D displacement label set at MRF nodes, but limited optimised MRF edges to a minimum spanning tree (MST), which enabled a running time of several minutes. Both [Glocker et al., 2008a] and [Heinrich et al., 2013] used a grid of MRF nodes and interpolated the dense deformation using FFD, in order to achieve speed-up. Recently, [Heinrich et al., 2016] used supervoxels as MST nodes, and their algorithm also has a running time of several minutes.

2.4.3 Incorporating Registration Uncertainty

The majority of registration algorithms yield a point estimate of the deformation parameters i.e. the optimum deformation as per Eq. 2.1. However, the ill-posed nature of registration implies inherent uncertainty in the estimated deformation as multiple solutions may exist that are close in terms of total energy, as the similarity measure varies mildly in homogeneous areas and may have low specificity for complex/noisy data, and the regularisation energy of the transformation may vary mildly for possible parameter values near the resection margin. This uncertainty is clearly relevant to surgical guidance.

One approach to uncertainty estimation requires the registration problem to be reformu-

CHAPTER 2. REVIEW OF IMAGE GUIDED NEUROSURGERY

lated in the Bayesian probabilistic modelling and inference framework, which is a generalisation of the MAP framework outlined in Section 2.4.1 [Chitphakdithai and Duncan, 2010, Chitphakdithai et al., 2011], and Markov Chain Monte Carlo (MCMC) sampling from the joint posterior distribution of the transformation parameters. In this vein, [Risholm et al., 2013] presented a registration scheme with an SSD-based Boltzmann distribution likelihood term and a Boltzmann distribution prior on the elasticity parameters of a finite element (FE) elastic biomechanical model. They applied this method on neurosurgical MRI/iMRI structural scans and found that the posterior distribution of deformation was multi-modal (i.e. had multiple local optima) in the vicinity of the resection cavity. A key limitation of the sampling approach is an extremely long running time, on the order of days, required to cover the high-dimensional parameter space, which is exacerbated by the fact that individual sample draws involve long Markov chains.

A considerably faster approach also based on the Bayesian formulation of the problem is approximate inference, where the true (unknown) joint distribution is approximated by a simpler analytical distribution. For instance, [Simpson et al., 2012] applied Variational Bayes to continuous non-rigid registration, whereby they approximated the distribution of FFD coefficients by the multivariate Gaussian distribution and minimised an information-theoretic that decreases along with an improving approximation. The limitation of this approach is that it does not easily allow for modelling multi-modal distributions.

In discrete registration, where transformation parameters are optimised using an MRF, approximate marginal distributions of the parameters can be obtained. The marginal distributions effectively denote the uncertainty in the assignment of particular labels (3D displacement) to individual MRF nodes. For instance, [Glocker et al., 2008b] estimated the uncertainty of discrete displacement labelling for each FFD control point by evaluating so-called min-marginal energies, from which approximate marginal distributions can be derived within only several multiples of the running time that is needed to find just the global optimum [Kohli and Torr, 2008]. Similarly, [Heinrich et al., 2016] used approximate marginal distributions for a discrete registration scheme defined on an MRF with edges limited to a spanning tree (Section 2.4.2), for which approximate marginal distributions could be obtained with favourable approximation bounds; they generated many instances of perturbed spanning trees in order not to bias the uncertainty estimates toward a single spanning tree and, finally, fused the uncertainty estimates to obtain marginal distributions for a dense deformation field.

2.4.4 Evaluating Estimated Brain Shift

The evaluation of estimated brain shift is a challenging task as outlined below.

An incomplete approach to validate registration accuracy is to annotate point anatomical landmarks in the structural pre/intraoperative image pair and to compare the *target registration error* (TRE) for the affinely and non-rigidly registered datasets [Shamir et al., 2009]. However, it should be noted that it is not straightforward to annotate point landmarks. First, the number of landmarks that can be reliably identified in the image pair is limited. Secondly, it may be challenging to identify the movement of landmarks in the presence of complex deformations as medical image viewing tools are based on the three anatomical planes (coronal, sagittal, axial) but the deformations due to brain shift are not restricted along these planes (out-of-plane deformations) [Nimsky et al., 2000, Nabavi et al., 2001]. Despite these limitations, manual annotation is a valuable validation method. A good example of an open dataset for evaluating brain shift estimation algorithms that includes manual annotations was provided by [Mercier et al., 2012] although the annotations were made for iUS volumes only.

To offer а reliable landmark-based means of validating brain deformations, [Chen et al., 2012] presented a physical phantom based on polyvinyl alcohol cryogel, which is cast from a mold based on a standard brain atlas and contains two baloons that can be filled with water using a catheter and emptied to simulate brain deformation loosely akin to brain shift, while ball-shaped MRI/US-contrast landmarks are inserted in the cortical part of the phantom; the authors released T1w and T2w MRI and US volumes for the phantom that are usable to validate registration assuming the inserted landmarks are inpainted to prevent introducing bias. The disadvantage of this dataset is that the cortical model does not contain a white/gray matter separation.

Paper	Similarity	Transfor-	Regular-	Optim-	Symm-	Uncer-	Speed
		mation	isation	isation	etric	tainty	
[Modat et al., 2010]	IMI	FFD	Exp.	Cont.	No	No	< 3m.†
[Glocker et al., 2008a]	Multiple	FFD	Exp.	Disc.	No	N_{O}	10-50 m.
[Heinrich et al., 2013]	MIND	FFD	Exp.	Disc.	Yes	No	< 5m.
[Simpson et al., 2012]	SSD	FFD	Exp.	Cont.	N_{O}	Yes	N/A
[Parisot et al., 2012]	Multiple	FFD	Exp.	Disc.	N_{O}	Yes	1—8 m.
[Heinrich et al., 2016]	SADG	FFD	Exp.	Disc.	Yes	\mathbf{Yes}	90 s.
[Daga et al., 2012]	IMI	FFD	Exp.	Cont.	No	No	$< 3m.\dagger$
[Risholm et al., 2009]	SSD	Dense	Imp.	Cont.	Yes	No	$20m.\ddagger$
[Ou et al., 2011]	GAM	FFD	Exp.	Disc.	N_{O}	N_{O}	115 m.
[Periaswamy and Farid, 2006]	Gauss.	Dense^*	Exp.	Cont.	N_{O}	No	30 m.
[Chitphakdithai and Duncan, 2010]	Gauss.	FFD	Exp.	Cont.	N_{O}	N_{O}	N/A
[Chitphakdithai et al., 2011]	Gauss.	FFD	Exp.	Cont.	No	N_{O}	N/A
[Risholm et al., 2013]	SSD	FEM	Exp.	Cont.	N_{O}	Yes	46 h.
[Wassermann et al., 2014]	LME	D/GP	Exp.	Cont.	Yes	Yes	N/A
"able 2.1: Selected mublications relevant to regis	stration-based	brain shift est	imation. k	cev to abbi	reviations:	NMI (Nor	malised Mutual

(explicit), Imp. (implicit); Cont. (continuous), Disc. (discrete). Running times are as given in the paper or maxima giving my estimate for resolution $256 \times 256 \times 256$ on a CPU (e.g. Intel Core i5 2.6 GHz), \dagger denotes time for a GPU (medium-spec GePorce GTX 680), \ddagger Information), MIND (Modality-Independent Neighborhood Descriptor), SSD (Sum of Squared Differences of Intensities), SADG (Sum of Absolute Differences of Intensities and Gradients), GAM (Gabor attribute matching, bespoke measure), Gauss. (Gaussian probability of intensity difference), LME (landmark misregistration error, but framework is extensible to intensity-based); FFD (Free Form Deformation), Dense* (dense and locally affine), FE (finite element model), D/GP (diffeomorphic flow given by a Gaussian Process); Exp. denotes 2D only.

CHAPTER 2. REVIEW OF IMAGE GUIDED NEUROSURGERY

2.5 Summary

In this chapter, I reviewed the literature on image guided neurosurgery, namely the relevant interventional modalities including iMRI as well as brain shift estimation techniques based on biomechanical modelling and image registration methods. In Section 2.2, I argued that the use of biomechanical models is complicated by the need for a highly accurate tissue segmentation and the need to know a large number of parameters, many of which are only known with a large degree of uncertainty. On the other hand, the iMRI data considered in this thesis are relatively high-quality and high-resolution volumetric images with a relatively dense coverage of features such as cortical folds and deep structures. These data are thus sufficient to drive image registration with a relatively high degree of accuracy. For these reasons, in the remainder of the review, I focused on the registration based approach to brain shift estimation; the most revelant work is listed in Table 2.1. The registration based approach also provides the context for subsequent chapters.

In Chapter 3, I will explore the potential of DW-MRI data acquired using iMRI that are a powerful source of data for multimodal registration [Modat et al., 2010, Daga et al., 2012]; in particular I will build on previous work [Daga et al., 2014] and focus on the need to compensate for distortions that arise in these data and affect the accuracy of guidance. In Chapter 4, I will focus directly on the image registration problem, with regard to the need to account for specific appearance differences that exist between preoperative and iMRI structural images.

Chapter 3

Field Map Simulation for Distortion Correction

In Chapter 1 an outline was given of the utility of multi-modal iMRI imaging in the context of brain shift estimation. Structural scans such as T1w MRI volumes can reveal anatomical changes such as deformed cortical sulci. Other modalities can reveal tissue characteristics that are not present in structural scans. Such characteristics can be especially useful if they are not discernible to the surgeon by the naked eye. An important category are the nerve fibre tracts such as the optic radiation tract (Section 1.1). DW-MRI has emerged as a modality that can reveal the locations of nerve fibre tracts and can be acquired by high-field iMRI scanners (Section 1.2).

In DW-MRI, a series of diffusion-weighted images are acquired with pulsed diffusion-encoding magnetic field gradients [McRobbie et al., 2017]. In each diffusion-weighted image, the MRI signal from water molecules that diffuse along the gradient direction becomes attenuated [Stejskal and Tanner, 1965]. A separate image is acquired with no diffusion-encoding gradient applied in order to serve as a signal intensity reference. Subsequently, the signals acquired for a given voxel are fitted to a voxel-based model of water diffusion. A widely used model, Diffusion Tensor Imaging (DTI), assumes that the water molecules have reached a symmetric Gaussian distribution in space after being allowed to diffuse for a short period. This model is agnostic of tissue microstructure in the voxel and can be parametrised in a simple manner by a diffusion tensor [Basser et al., 1994]. More complex models aim to explicitly represent tissue microstructure and its influence on the measured signals. For instance, fibre bundle geometry or presence of extracellular space can be factors in such models [Essayed et al., 2017]. DW-MRI data are the source of data for probabilistic tractography techniques that aim to identify nerve fibre tracts spanning many voxels such as the optic radiation and the corticospinal tract (Section 1.1). Tractography based on preoperative DW-MRI imaging has emerged as a key step in neurosurgical planning [Nimsky et al., 2005, Duncan et al., 2016, Essayed et al., 2017].

High-field (1.5T) iMRI systems opened another compelling application area for DW-MRI [Hall and Truwit, 2008]. As I mentioned in Section 1.2, tractography from DW-MRI data acquired using iMRI has remained elusive. However, [Daga et al., 2012] attempted to overcome this problem by instead propagating tractography derived from preoperative DW-MRI using nonrigid registration with iMRI images. In order to utilise DW-MRI data acquired using iMRI, they built an intensity channel from a fractional anisotropy (FA) map, which they computed using the DTI model. FA is a value that can be calculated for each voxel from the diffusion tensor to provide a normalised measure of diffusion anisotropy such that zero (unity) may indicate the lack (presence) of a crossing nerve fibre tract. The FA map was used alongside the structural scan and provided an outline of tracts in white matter areas where the structural scan lacks features. This demonstrates that DW-MRI data acquired on an iMRI scanner can be used to improve the accuracy of neuronavigation.

However, a question arises regarding how favourably the DW-MRI datasets compare to structural iMRI scans in terms of localisation accuracy. An important factor is that in DW-MRI, many diffusion-weighted images need to be acquired per dataset. For instance, the DTI model requires at least 30 diffusion-weighted images in order for the diffusion tensor to be fitted robustly [Jones, 2004]. More complex diffusion models require further diffusion-weighting directions acquired at multiple strengths of diffusion-encoding gradients in order to fit the models [Tuch et al., 2002]. A conventional acquisition of an MRI volume (using a conventional pulse sequence) takes up to several minutes to acquire a volume at high resolution (e.g. 1 mm isotropic voxels). DW-MRI volumes have a lower resolution (e.g. 2.5 mm isotropic voxels) but it would



Fig. 3.1: Susceptibility artefact correction. (a) An intraoperative T1-weighted image which is not itself affected by the distortion. The brain surface outlined using a surface extraction tool is shown for reference (red outline). (b) An image from the DW-MRI dataset affected by the artefact. Arrows point at the region of severe distortion which coincides with the resection margin. (c,d) The same image corrected using (c) an acquired and (d) a simulated field map.

still be prohibitively slow to acquire volumes for all the diffusion directions using a conventional pulse sequence. Instead, DW-MRI volumes are acquired using the *echo planar imaging* (EPI, [Mansfield, 1977]) sequence, which is considerably faster. However, EPI suffers from a so-called *susceptibility artefact*, which is described further below.

The susceptibility artefact distorts features revealed though DW-MRI imaging such as the outlines of putative tracts in FA maps. In the case of DW-MRI imaging using iMRI, the distortion is particularly severe near the air-tissue boundary presented by the resection cavity. Several methods for correcting the artefact exist and are outlined below. However, the existing methods tend to rely on data which have a reduced reliability near the resection margin, which is the same region for which an accurate correction is sought.

In this chapter I propose an alternative method which relies on the conventional T1-weighted structural scan and on simulating the physical phenomena that give rise to the artefact. My main motivation is that the structural scan is routinely acquired as part of iMRI guidance at high-resolution and does not itself suffer from the same type of imaging artefact. I perform an extensive evaluation of the proposed method and discuss its usability within and beyond the context of image-guidance for neurosurgery.

3.1 Introduction

The susceptibility artefact in EPI is a geometric and intensity distortion. The artefact manifests itself in the affected areas in three ways: first, voxels get locally translated relative to where they would be placed in a scan performed with a conventional pulse sequence; second, the signal from several voxels may get compressed into fewer voxels (local non-linear translation); third, the signal strength severely drops, giving rise to areas of low intensity and poor signal-to-noise ratio (SNR), which is also referred to as signal dropout. The image areas affected by the susceptibility artefact are located near air-tissue boundaries. In non-interventional MRI images, the affected brain areas are adjacent to the air-gaps in the paranasal sinuses and the petrous part of the temporal bone, and the resulting translations can be up to 1—1.5 cm for 3 T scanners [Jezzard and Balaban, 1995].

In iMRI data of open neurosurgery, an obvious and large air-tissue boundary is presented by the resection and is implicitly located in the region where as accurate DW-MRI data as possible are desirable. To illustrate the problem, Fig. 3.1(a) shows a slice through an intraoperative T1w scan acquired using a conventional pulse sequence and hence without the artefact, while Fig. 3.1(b) shows a slice through *one* of the DW-MRI volumes of the same surgical case, in which the susceptibility artefact is apparent and is marked with arrows.

Note that the volume used in Fig. 3.1(b) is the so-called b0 volume (further, b0-DW-MRI), which is actually the only volume from the DW-MRI dataset for which the diffusion is *not* constrained along any physical direction; because DW-MRI is based on T2w imaging, the b0-DW-MRI volume is simply a T2w scan acquired using EPI. Thus it is straightforward to visually compare the b0-DW-MRI volume with the conventional T1w scan in order to appreciate the susceptibility artefact. Note that *all* the volumes from the DW-MRI dataset are equally affected by the same distortion. Therefore the results of any post-processing, such DTI fitting and/or tractography, will be geometrically displaced. Furthermore, the signal dropout and the additive compression of the MRI signal, both of which are irrecoverable, can complicate DW-MRI post-processing in the areas affected by the artefact [Kim et al., 2006].

3.1.1 The Origin of the Susceptibility Artefact

The artefact arises during EPI imaging as follows. In MRI, the spatial origin of signal along the three imaging dimensions is always determined (or *encoded*) by linear gradient fields that are sequentially superimposed onto the main field, or the B_0 field, of the superconducting magnet [McRobbie et al., 2017]. The spatial encoding implies that for image acquisition, each of the image dimensions has to be chosen to serve as one of the nominal "slice select", "phase encode" (PE) or "frequency encode" (FE) dimensions. The spatial encoding implicitly assumes that the B_0 field is spatially homogeneous throughout the imaged volume and conversely, any inhomogeneity will necessarily result in the perceived spatial origin of the MRI signal being virtually shifted i.e. the corresponding voxels are displaced in the image. The degree of the displacement along a given imaging dimension is inversely proportional to the signal bandwidth in that imaging dimension. A key disadvantage of EPI is that it has a limited bandwidth in the PE imaging dimension. In practice, this causes even small B_0 field inhomogeneities on the order of several parts per million (ppm) to cause perceptible linear displacements along the PE dimension, while shifts along the other two nominal dimensions of the image are negligible.

Some B_0 inhomogeneities arise due to design constraints on the magnet; they are reduced (or *shimmed*) to several ppm by means of dedicated superconducting shim coils [Clare et al., 2006]. Further B_0 inhomogeneities arise due to the perturbations of the field by non-uniform geometric distribution of magnetic susceptibility in the imaged volume. Such perturbations are largest near air-tissue boundaries such as the paranasal sinuses and the petrous part of the temporal bone [Jezzard and Balaban, 1995]; in the case of open neurosurgery, an obvious and large air-tissue boundary is presented by the resection and is implicitly located in the region where as accurate DW-MRI data as possible are desirable. The perturbation due to tissue susceptibility distribution is shimmed using another set of dedicated room-temperature shim coils whose currents are optimised dynamically by the scanner during scanning.

However, some B_0 inhomogeneities due to higher-order field perturbations due to magnetic susceptibility cannot be shimmed. The effect of these residual inhomogeneities is what is referred to as the susceptibility artefact characterised above.



Fig. 3.2: B_0 field maps for subject #3 from the below study. (a) A raw acquired field map which contains modulo- 2π phase-wraps. (b) The acquired field map phase-unwrapped. (c) The proposed simulated field map. An advanced view is shown in Fig. 3.7.

3.1.2 Susceptibility Artefact Correction

I introduce two methods for susceptibility artefact correction for DW-MRI and remark on scenarios where correction is not possible due to omissions in the imaging protocol.

B₀ Field Map Based Correction

One approach to correct the susceptibility artefact is to *acquire* a residual B_0 inhomogeneity map ("field map"), for which several minutes of additional acquisition time are required. The EPI image can then be corrected by considering the field map as defining the linear displacement along the PE dimension. This approach was recently used in a clinical iMRI study and facilitated a reduction of the intraoperative susceptibility artefact [Daga et al., 2014].

The acquisition of the field map is based on the gradient-recalled echo pulse sequence [Jezzard and Balaban, 1995], which measures the (nuclear spin) phase evolution between a pair of consecutive MRI signal echoes. The resulting phase-difference map ("acquired field map") is modulo- 2π wrapped due to the unknown true number of elapsed phase revolutions. An example phase-wrapped acquired field map is shown in Fig. 3.2(a). Techniques for the necessary phase-unwrapping to recover the true field map include the widely used deterministic algorithm of [Jenkinson, 2003] and a recent graphical probabilistic inference model of [Daga et al., 2014]. An example field map phase-unwrapped using their model is shown in Fig. 3.2(b) and an example b0-DW-MRI image undistorted using the same phase-unwrapping method is shown in Fig. 3.1(c). The phase-unwrapping of B_0 field maps is ill-posed and challenging due to low SNR of raw field maps as can be appreciated in Fig. 3.2(a). The noise is most problematic in areas where phasewraps are tightly spaced as this causes ambiguity about the true number of phase-wraps. Such problem tends to occur in areas of highest distortion i.e. near the resection. Thus the estimated field map may not reflect the true field in the same area which is most relevant clinically. However, the phase-unwrapping approach could potentially benefit from a spatial prior that was based on a physically meaningful simulation of the distribution of the magnetic field.

Registration Based Correction

Another approach to correcting the artefact is to use registration to estimate the 1D deformation constrained along the PE dimension. For instance the b0-DW-MRI volume (being a lowresolution distorted T2w volume) can be registered to a high-resolution non-distorted structural volume (a T1w or a T2w volume), for instance using a multi-modal similarity measure such as NMI [Studholme et al., 1999]. All the diffusion-weighted volumes can then be corrected by applying the same 1D deformation, as they are affected by the artefact in the same way as the b0-DW-MRI volume [Pujol et al., 2015]. This registration approach has been shown to reduce tensor fitting errors in DW-MRI datasets [Kim et al., 2006]. However, using this relatively simple method in an iMRI setting is not ideal as accurate registration near the resection cavity is very challenging and likely inaccurate due to the heavy susceptibility artefact itself as can be appreciated in Fig. 3.1(b).

A more advanced registration approach involves the acquisition of the b0-DW-MRI volume twice using opposite direction for the PE dimension (*blip-up blip-down* EPI), which causes the distortion to manifest itself in opposite directions in both volumes. The registration then aims to jointly estimate the mid-point deformation field while explicitly compensating between the heavy compression and the heavy expansion the two volumes [Andersson et al., 2003]. The estimated deformation can then again be applied to the diffusion gradient volumes. This approach is appealing as it avoids the intermediate step of field map acquisition and operates directly on the distorted images. The time required to acquire an opposite PE direction volume for the b0-DW-MRI volume is less than the time required for field map acquisition. However, this method has to my knowledge not been applied in an iMRI data setting yet and so it is not clear whether it would be sufficiently robust to the heavy signal dropout near the resection cavity.

Registration approaches tend to find an accurate alignment between images in areas with sufficiently dense and clear salient features to drive the registration (Section 2.3). Conversely, in areas with homogeneous appearance (such as the white matter) or with low-quality features (such as areas with image blurring/fading), registration results depend primarily on the employed regularisation term(s) and the corresponding parameter settings (Section 2.3.3). However, in the context of distortion correction, it is not immediately clear what these settings should be. For these reasons, for instance [Daga et al., 2014] employed a phase-unwrapped field map as a *prior* that informed a subsequent registration step; the inclusion of the prior was enabled by a graphical model formalism employed by their scheme. Note that in principle any meaningful prior could be used for this purpose.

Scenarios with Missing Data Needed for Correction

The susceptible artefact is not severe only for intraoperative iMRI data but is also significant for non-interventional data. Despite that, the correction of the artefact has in practice been neglected by the neuroimaging community [Glasser et al., 2013]. Thus in many DW-MRI studies the B_0 field maps were not collected alongside EPI datasets and neither has blip-up blip-down EPI acquisition been used. A mean to crudely correct DW-MRI is to non-rigidly register the b0-DW-MRI volume(s) to the structural scan, but this is not optimal as per Section 3.1.2. Thus it will not be possible to retroactively correct DW-MRI data to assess reliability of results reported in previous studies and/or to include these data in new studies alongside data for which correction is included in the imaging protocol.

While the focus of this chapter is on DW-MRI, functional MRI (fMRI) also relies on EPI [Hutton et al., 2002] and a similar logic applies to datasets with missing field maps. In fact registration based correction (registration against structural MRI, blip-up blip-down EPI) may be unfeasible for fMRI as brain activation maps lack sufficient texture, which underlines the problem with missing field map data.

3.1.3 Proposed Correction using Simulated Field Maps

In the previous section I outlined situations in which it would be useful to obtain a physically reliable model of the geometric distribution of the magnetic field in the head at an intraoperative time-step. I propose to simulate a field map from air-tissue segmentations generated using non-EPI T1-weighted (T1w) and T2w iMRI images acquired as part of a clinical iMRI scanning protocol. The structural scans are appealing for this purpose as they are high-resolution and are not affected by the susceptibility artefact.

Previously, Jenkinson *et al.* [Jenkinson et al., 2004] developed a perturbation method to calculate a B_0 inhomogeneity field from air-tissue segmentation derived from computed tomography (CT) images. Poynton *et al.* [Poynton et al., 2009] demonstrated that non-interventional T1w images can be segmented into air and tissue classes using a probabilistic CT atlas and reported that a subsequent application of the method [Jenkinson et al., 2004] results in a close overall agreement between the acquired and simulated field maps. Their results are heavily motivated by the problem of datasets with missing field maps as per Section 3.1.2 and based on the observation that the structural scan reflects the anatomy and the position of the head in the scanner and is naturally always included in the imaging protocol.

In the present problem with iMRI the main issue instead is that the phase-unwrapping of field maps can potentially have low confidence near the resection cavity due to the combination of noise and the tight packing of phase-wraps in this area. A specific problem posed by the iMRI data is that, since a resection cavity and an air-filled craniotomy area are present in the intraoperative image, thus a probabilistic atlas derived from *non*-interventional datasets is not well-suited for the air-tissue segmentation.

I propose to instead inform the air-tissue segmentation using priors derived from a synthetic CT image, which I propose to be computed from the preoperative T1w image based on a database of MRI/CT pair templates. I also propose to subsequently feed the air-tissue segmentation into the simulation step due to [Jenkinson et al., 2004]. To illustrate this approach, an example simulated field map is shown alongside the acquired one in Fig. 3.1(c). The correction that can be obtained using the proposed simulated field map approach is illustrated in Fig. 3.1(d).
I evaluate the proposed field map simulation method by comparison with field maps acquired during iMRI guided ATLR neurosurgery for 12 cases and show that the proposed method generates field maps in close agreement with the acquired field maps; this result was already presented in [Kochan et al., 2014]. I further interpret the differences between the acquired and simulated field maps based on a quantitative evaluation of the effects of susceptibility correction using manually annotated anatomical landmarks and tensor fit error maps; this result was already published in [Kochan et al., 2015].

3.2 Methods

The methods are described as follows. Section 3.2.1 introduces the method used to correct EPI images based on field maps and Section 3.2.4 outlines the preprocessing step used for raw acquired field maps in order to enable the comparison of acquired and simulated field maps. Section 3.2.2 describes an air-tissue segmentation scheme used to bootstrap a field map simulation outlined in Section 3.2.3. An overview of the processing pipeline, needed to correct a DW-MRI image based on either an acquired or a simulated field map, is shown in Fig. 3.3.

3.2.1 Field Map in Terms of Voxel Displacement

Let the magnetic field at point \vec{x} be $B_0 + \Delta B_0(\vec{x})$ [T] where B_0 is the homogeneous field and $\Delta B_0(\vec{x})$ is the inhomogeneity field map, which can be equivalently expressed as $\gamma \Delta B_0(\vec{x})$ [rad/s] or $\frac{\gamma \Delta B_0(\vec{x})}{2\pi}$ [Hz]. For the purposes of image correction, one is interested in the millimeter displacement along the PE direction that the inhomogeneity causes to an EPI image. The displacement can be calculated based on theory in [Jezzard and Balaban, 1995, Hutton et al., 2002]. Consider the acquisition of a single EPI slice with matrix size $N \times N$ and voxel dimensions $r_{\rm FE}$ in the FE direction and $r_{\rm PE}$ in the PE direction, respectively. The EPI slice is reconstructed by the inverse Fourier transform of the MR signal. In the PE direction, the MR signal sampling rate is $\frac{N}{T_{\rm RO}}$ [Hz], where $T_{\rm RO}$ is the readout time. The image bandwidth in the PE direction is $\frac{N}{NT_{\rm RO}} = \frac{1}{T_{\rm RO}}$ [Hz/pixel] or $T_{\rm RO}$ [pixel/Hz]. Since the PE gradient is used to encode position along the PE direction, the



Fig. 3.3: The processing pipeline for correcting the susceptibility artefact in a DW-MRI image. The dashed lines represent the step into which either a phase-unwrapped acquired field map or a simulated field map can be fed. The entry points of the simulation are non-distorted preoperative T1w and intraoperative T1w and T2w images.

above offset corresponds to a distortion along the PE direction of size

$$d_{\rm PE}(\vec{x}) = \frac{\gamma \Delta B_0(\vec{x})}{2\pi} T_{\rm RO} r_{\rm PE}$$
(3.1)

in millimeters. The magnetic field $\Delta B_0(\vec{x})$ is either found by simulation or, in case of the acquired field map, can be derived from the phase evolution $\Delta \phi(\vec{x})$ [rad] of phase during the echo time difference $T_{\rm ED}$:

$$\Delta B_0(\vec{x}) = \frac{\Delta \phi(\vec{x})}{\gamma T_{\rm ED}} \tag{3.2}$$

A scalar displacement map was calculated using Equation 3.1 and was converted into a dense displacement field along the PE direction [Daga et al., 2014]. Identical displacement was applied to all the EPI images in each DW-MRI dataset by resampling the images with cubic spline interpolation using the resampling utility from the NiftyReg package [Modat et al., 2010].



Fig. 3.4: Air-tissue segmentation. (a) Intraoperative T1w+T2w sum image. The section runs through a plane close to the anatomical coronal plane with the head at an angle due to intraoperative orientation. (b) An accompanying pseudo-CT (range -1000–1000 HU). (c) The result of the proposed segmentation within the intraoperative FOV (red for air, green for soft-tissue, blue for bone). (d) The final air-tissue segmentation (black for air, white for tissue). The appended inferior volume outside of the intraoperative FOV is an approximate air-tissue segmentation based on preoperative T1w MRI.

3.2.2 Air-Tissue Segmentation

The magnetic susceptibility values for soft-tissue ($\approx -9.1 \times 10^{-6}$) and bone ($\approx -11.4 \times 10^{-6}$) are similar, but both are significantly different from that of air ($\approx 0.4 \times 10^{-6}$) [Poynton et al., 2009]. Therefore, a binary labeling of the head was needed, into tissue (which combines soft tissue and bone) and air.

For each subject, a segmentation was performed on the sum (T1w+T2w) of the intraoperative T1-weighted (T1w) and T2-weighted (T2w) MRI images, which were already co-registered, as the head was fixed during acquisitions. In T1w images, soft tissue with low free water content (e.g. the brain, muscles, fat) appears bright, while in T2w images, the voxels with high free water content (such as CSF and the eyes) appear bright. Therefore, an accurate delineation of the soft-tissues can be obtained by summing the T1w and T2w MRI images. However, since bone and air have similar low intensities in both T1w and T2w images, they are indistinguishable based on intensity alone. In order to distinguish between air and bone, a spatial prior probability map (spatial prior) was used for each class, which enabled the segmentation of the T1w+T2w image into three partial volume classes (air, soft-tissue and bone).

Air and bone have distinct attenuation properties, which makes CT images a good candidate for the spatial priors, as their intensities expressed in Hounsfield units (HU), are a linear transformation of the linear attenuation coefficients. However, CT images were not part of the preoperative planning imaging protocol. The attenuation information was derived from a pseudo-CT, synthesised from a preoperative T1w image, as shown in Fig. 3.4(b).

CT Synthesis

The CT synthesis method, developed by Burgos *et al.* [Burgos et al., 2013, Burgos et al., 2014] relies on a pre-acquired set of aligned T1w MRI/CT image pairs from multiple subjects forming an MRI-CT database. Two gender-specific databases were used, both comprising 35 subjects part of an epilepsy study. To generate the CT from the preoperative T1w MRI image, each MRI image from the database was non-rigidly registered to the preoperative T1w MRI image. The CT images were then mapped using the same transformation to the preoperative T1w MRI image. A local image similarity measure between the preoperative T1w MRI and the set of registered MRIs from the database was used as a surrogate of the underlying morphological similarity, under the assumption that if two MRIs are similar at a certain spatial location, the two CTs will also be similar at this location. Finally, the set of registered CTs was fused using a spatially varying weighted averaging, with weights derived from the similarity measure, generating a pseudo-CT.

Constrained Spatial Priors

To constrain the spatial priors to the probability range 0–1, the sigmoid intensity transformation $\operatorname{sig}(I(\vec{x})) = 1/(1 + \exp(\phi_0 + \phi_1 I(\vec{x})))$ was used, where $I(\vec{x})$ is the image intensity at voxel \vec{x} and the constants ϕ_0 and ϕ_1 are chosen so as to transform a particular intensity range to the transitional interval 0.001–0.999. The soft-tissue prior was established by intensity-transforming the T1w+T2w image intensity range 50–90 (mild gray).

The pseudo-CT, synthesised from the preoperative T1w MRI, was rigidly propagated to the intraoperative space. The advantage of using CT images is that the Hounsfield scale is quantitative, whereby air corresponds to -1000 HU, soft-tissue ranges from 0 to 100 HU and bone from 700 to

3000 HU. Using the sigmoid function, the air prior was allowed to vary from 0.999 to 0.001 for -1000–0 HU and the bone prior was allowed to vary from 0.001 to 0.999 for 100–700 HU.

The spatial priors had no biological basis in the areas altered by the surgery. The air prior was linearly reduced into the range 0.1–0.9 to enforce the resection cavity to become segmented as air. The spatial priors were smoothed with a Gaussian kernel ($\sigma = 2$ mm) to enforce the skull to become segmented as bone, and the part of the skull that is removed by the surgeon during the craniotomy to become segmented as air.

Segmentation Scheme

The segmentation was obtained using the *seg_EM* tool from the NiftySeg package [Cardoso et al., 2009], which provides an intensity-based segmentation scheme based on the Gaussian Mixture Model (GMM) implemented using the Expectation-Maximisation (EM) algorithm and a Markov Random Field (MRF) spatial regularisation to provide a Maximum a Posteriori solution, as per Fig. 3.4(c). Fig. 3.5 shows the corresponding fitted GMM. The MRF prior strenth was set to an increased value ($\beta = 1.0$) to further enable a correct air segmentation in the craniotomy area.

iMRI and pseudo-CT fields of view (FOV) are largely limited to the cranial part of the head. The subsequent field map estimation step (Section 3.2.3) assumes that the voxels outside of the segmentation FOV are air-filled, which I observed to result in unrealistic field maps due to the virtual tissue cut-off below the cranium. I utilised the wide FOV of the rigidly registered preoperative T1w and appended a 100 mm-deep volume inferiorly to the intraoperative FOV containing an approximate air-tissue segmentation obtained by thresholding the preoperative T1w (Fig. 3.4(d)).

3.2.3 Field Map Estimation

The field map estimation was performed using the tool b0calc from the FSL software package [Smith et al., 2004, Jenkinson et al., 2004], which models the first order perturbations of the main magnetic field. The susceptibility χ can be expanded as $\chi = \chi_0 + \delta \chi_1$, where χ_0 is the magnetic susceptibility of air, δ is the susceptibility difference between air and brain tissue and χ_1



Fig. 3.5: Histogram (cropped for clarity of the figure) of the T1w+T2w image in Fig. 3.4(a) and the fit of the GMM model for the segmentation in Fig. 3.4(a). The data come from intraoperative scans for subject #3.

is a binary variable describing the tissue type. The first order perturbations of the z-component of the main magnetic field (B_z^1) can be written in terms of the main magnetic field (B_z^0) :

$$B_{z}^{1} = \frac{\chi_{1}}{3 + \chi_{0}} B_{z}^{0} - \frac{1}{1 + \chi_{0}} \left(\left(\frac{\partial^{2} G}{\partial z^{2}} \right) * (\chi_{1} B_{z}^{0}) \right)$$
(3.3)

where G is the Green's function $G(x) = (4\pi r)^{-1}$ and $r = \sqrt{x^2 + y^2 + z^2}$. Note that the expression is simplified considerably due to the fact that there only is a non-zero component B_z^0 in the longitudinal axis (z-direction) of the main magnetic field.

Equation (3.3) cannot be evaluated analytically for an arbitrary susceptibility distribution χ_1 . However, an analytical solution $H(\vec{x})$ exists for a single voxel of size (a, b, c) centered at origin (with $\chi_1 = 1$ inside a voxel and $\chi_1 = 0$ outside of the voxel) for a constant field along the z-axis and is given by:

$$H(\vec{x}) = \left(\frac{\partial^2 G}{\partial z^2}\right) * (\chi_1 B_z^0) = \sum_{i,j,k \in (-1,1)} (ijk) F\left(x + \frac{ia}{2}, y + \frac{jb}{2}, z + \frac{kc}{2}\right)$$
(3.4)

where $F(\vec{x}) = \frac{1}{4\pi} \arctan(\frac{xy}{zr})$.

Due to the linearity of Equation (3.3), the shifted single voxel solutions $H(\vec{x} - \vec{x'})$ can be added together using a discrete convolution to compute the total field:

$$B_z^1(\vec{x}) = \sum_{\vec{x'}} \chi_1(\vec{x'}) H(\vec{x} - \vec{x'})$$
(3.5)

where $\vec{x'}$ are the voxel locations, $\chi_1(\vec{x'})$ is the voxel-based susceptibility map and \vec{x} is the point where the field is evaluated. The discrete convolution can be implemented using the 3D Fast Fourier Transform (FFT).

Although this approach simulates the field distribution due to the main coil, it does not account for the decrease in inhomogeneity as effected by the room-temperature shim coils, which are wound to form fields that follow first and second order spherical harmonics $S(\vec{x}) = [x, y, z, z^2 - (x^2 + y^2)/2, xz, yz, x^2 - y^2, 2xy](\vec{x})$, where $\vec{x} = \vec{0}$ at the magnet isocentre [Gruetter and Boesch, 1992]. The field in the scanner becomes $B_z^1(\vec{x}) - S\vec{\theta}$, where the coefficients $\vec{\theta} = [\theta_1, \theta_2, \dots, \theta_8]^T$ are proportional to the currents in the shim coils, which are dynamically optimised by the scanner during image acquisition based on the field in the imaged volume [Gruetter and Boesch, 1992]. In this simulation, I approximate the shim currents as a linear combination that minimises the inhomogeneity field across the field of view, as used in [Poynton et al., 2009]. I perform a leastsquares fit of the spherical harmonics to determine $\hat{\vec{\theta}} = \arg\min_{\vec{\theta}} \left(\left(B_z^1(\vec{x}) - S\vec{\theta} \right)^T \left(B_z^1(\vec{x}) - S\vec{\theta} \right) \right)$.

3.2.4 Data Acquisition

The proposed algorithm was validated on 12 datasets that were acquired using interventional MRI during ATLR procedures. The study was approved by the National Hospital for Neurology and Neurosurgery and the UCL Institute of Neurology Joint Research Ethics Committee. Written informed consent was obtained from all participants. The images were acquired using a 1.5T Espree MRI scanner (Siemens, Erlangen, Germany) designed for interventional procedures. The intraoperative protocol included a conventional T1-weighted FLASH image (TR = 5.25 ms, TE = 2.5 ms, flip angle = 15° , $0.547 \times 0.547 \times 1.25$ mm grid of $512 \times 512 \times 176$ voxels) a conventional T2-



Fig. 3.6: Scatter plot of the acquired and the simulated field map in corresponding voxels inside the brain for subject #3.

weighted turbo spin echo image (TR = 3200 ms, TE = 510 ms, flip angle = 120° , $1.0 \times 1.0 \times 1.0 \times 1.0$ mm grid of $256 \times 256 \times 176$ voxels), and a DW-MRI dataset acquired using a single shot EPI sequence with GRAPPA-based parallel imaging (acceleration factor of 2, $2.5 \times 2.5 \times 2.5$ mm grid ($r_{\rm PE}$ = 2.5mm) of $84 \times 84 \times 49$ voxels, signal readout time $T_{\rm RO}$ = 35.52 ms) and a field map (see further below). The DW-MRI dataset consisted of set of 64 diffusion-weighted EPI images with diffusion weighting strength b = 1000 s \cdot mm⁻² accompanied by a b0-DW-MRI image. The GRAPPA imaging resulted in a halving of EPI echoes acquired in each readout but the readout time above is given nominally with the reduction factored in.

Field Map Acquisition

The field maps were acquired using a gradient-recalled echo pulse sequence [Jezzard and Balaban, 1995] (2.91667 × 2.91667 × 2.9 mm grid of $72 \times 72 \times 43$ voxels, echo time difference of $T_{\rm ED} = 4.76$ ms). The phase-unwrapping was performed using a method detailed in [Daga et al., 2014], which uses a Markov random field (MRF) formulation to enforce spatial smoothness in the estimated true field map. Since the recovered phase difference necessarily had



Fig. 3.7: Field maps expressed as mm of displacement along the phase encode direction. The view is centered at anterior temporal lobe resection cavity. The brain surface outlined using a surface extractor is shown for reference (red outline). (a–c) A phase-wrapped acquired field map for subject #3, showing a step-change in phase value close to the resection margin. (d–f) The acquired field map after phase-unwrapping; only the volume inside the brain mask is shown, because the phase-unwrapping was restricted to the brain only. (g–i) The proposed simulated field map. (j–l) A voxel-wise absolute difference between the simulated and the phase-unwrapped acquired field maps, only considered within the brain. Left to right: coronal (a,d,g,j), sagittal (b,e,h,k) and axial sections (c,f,i,l), respectively. Slice orientations are close to the standard anatomical planes. I used a brain surface extractor included in NiftyView (http://cmic.cs.ucl.ac.uk/home/software).



Fig. 3.8: Several axial slices through absolute difference between the simulated and phaseunwrapped acquired field maps, expressed as mm of displacement along the phase encode direction shown for subject #3. (a) Contralateral temporal lobe level. (b) Eye level. (c) Superior frontal and parietal lobe level.

an arbitrary constant component, the map was de-meaned based on the voxels inside the brain mask. The resulting map $(\Delta \phi(\vec{x}))$ was converted using Eqs. 3.1 and 3.2 into displacement units.

3.3 Experiments with Clinical Data and Results

Since there is no direct way of measuring the true field maps in vivo, I compared the simulated field maps to the acquired field maps. I expressed the field maps in terms of mm of displacement along the PE direction (as per Equation 3.1). In all except one of the given DW-MRI datasets, the PE direction coincided with the anatomical anterior-posterior axis i.e. a positive field map value corresponded to displacement toward the posterior end of the image, while in one DW-MRI dataset, the PE direction was inverted. A scatter plot of the two types of maps for subject #3 is shown in Fig. 3.6. I also calculated an absolute difference between the two types of maps (Fig. 3.7 and Fig. 3.8). I calculated statistics for the difference between the simulated and acquired field maps within the brain mask volume. The results for the 12 subjects are reported in Table 3.1. For most of the brain, there is a close agreement, whereby the maps differ by 1.15 ± 1.49 mm on average (3.43 mm in the 95th percentile) for all the subjects within the brain volume. However, in some areas, there are larger disagreements and the differences follow a long-tailed distribution, whereby the the maps differ by 4.51 ± 2.76 mm on average for all the subjects (9.99 mm in the 95th percentile) in the voxels in which the absolute difference is above the voxel size (2.5 mm). I note that a small discrepancy exists regarding the differences between the acquired and simulated field

Subj.	Whole bra	in	Disp. $> 2.5 \text{ mm}$			
	Mean (std)	P95	Mean (std)	P95		
1	1.18(1.53)	3.49	4.58 (2.86)	10.46		
2	1.06 (1.24)	3.17	4.10 (2.01)	7.91		
3	1.05(1.37)	3.10	4.46(2.65)	9.67		
4	1.10(1.53)	3.48	4.71 (3.04)	10.42		
5	1.22(1.54)	3.28	4.55(3.08)	10.54		
6	1.02(1.36)	3.00	4.57(2.76)	10.26		
7	1.08 (1.47)	3.03	4.77(3.24)	11.70		
8	$0.89\ (\ 1.35\)$	2.99	4.53(2.74)	10.05		
9	1.28(1.59)	3.51	4.44(3.05)	10.44		
10	1.34(1.72)	4.15	4.60 (2.73)	10.23		
11	1.36(1.69)	4.32	4.40 (2.45)	9.13		
12	1.19 (1.46)	3.62	4.35(2.49)	9.09		
Avg.	1.15 (1.49)	3.43	4.51 (2.76)	9.99		

Table 3.1: Absolute difference between the displacement (in mm) in the phase encode dimension as per the proposed simulated field map and the acquired field map. The mean (std) and 95th percentile values are reported for all the voxels in the brain and for those with the abs. difference above the voxel size (2.5 mm), respectively. The summary line lists the column averages.

maps reported in Table 3.1 as given in publications [Kochan et al., 2014] and [Kochan et al., 2015] in which these results were reported previously. Namely, in [Kochan et al., 2014] an incorrect value of 2.9 mm was used for the scaling factor $r_{\rm PE}$ in Eq. 3.1 while in [Kochan et al., 2015] the correct value of 2.5 mm was used and the correct results are also reproduced here.

The acquired field maps do not constitute a ground truth and therefore the accuracy of the simulated field maps had to be assessed by observing the corrected EPI images in several regions of interest in the brain. The DW-MRI dataset for each subject was corrected as per Section 3.2.1 using the acquired field map and the proposed simulated field map, respectively. Examples for two subjects are shown in Fig. 3.9. I chose the b0-DW-MRI image from each DW-MRI dataset as the representative EPI image, because it contains visually identifiable landmarks. Since the conventional T2-weighted image has a similar tissue contrast profile as the b0-DW-MRI image, but has a higher resolution and is not affected by the susceptibility artefact, I used it as the ground truth image. I manually identified anatomical landmarks in the ground truth and susceptibility-distorted b0-DW-MRI images and measured the landmark misregistration in millimeters. I then propagated the landmarks from the distorted b0-DW-MRI images using the acquired and simulated field maps, respectively, and measured the misregistration between the ground truth and the

propagated landmarks. I only considered landmarks with reliable correspondences. Although in the T2-weighted images it was possible to identify features and localise them at image-resolution accuracy (1 mm³), conversely in the b0-DW-MRI images in areas of signal dropout or severe susceptibility distortion it was not possible to find reliable corresponding landmarks. In such cases, I collected nearest reliably identifiable landmarks.

I collected landmarks in 4 brain regions of interest (ROIs): (1) near the resection margin (posterior and superior from the resection cavity, 7–8 landmarks/subject, see Fig. 3.9(a)); (2) brain stem (due to clearly distinguishable margin and proximity to the resection cavity, 5 landmarks/subject, see Fig. 3.9(a)); (3) areas of high distortion in the proximity of the sites of the head-holder fixation pins (4–5 landmarks/subject, see Fig. 3.9(b)); (4) landmarks evenly distributed in the rest of the brain (6 landmarks/subject). Table 3.2 summarises the landmark distances for the four ROIs. The titanium head fixation pins (ROI 3) are part of the head-holder (Noras, Hoechberg, Germany) and in their vicinity DW-MRI exhibits increased distortion, while T1w and T2w images exhibit signal dropout in the scalp.

I performed 3 one-tailed paired Student's *t*-tests for each ROI with the following alternative hypotheses: (a) landmark misregistration is smaller in images corrected using acquired field maps than in uncorrected images; (b) the same when correcting using simulated field maps; (c) landmark misregistration is smaller in images corrected using acquired field maps than in images corrected using simulated field maps. I set the significance level for a test regarding one ROI at 1.25% as per Bonferroni correction (5/4 = 1.25%) in order to keep the family-wise error rate (FWER) below 5%, where FWER is the probability of at least one null hypothesis being rejected due to chance alone. The results are reported in Table 3.3.

However, landmark based validation is complicated due to non-repeatability of manual visual landmark localisation, as well as by the low resolution and poor quality of the b0-DW-MRI image and limited density of unambiguous landmarks. I corroborated the validation with an additional approach as in [Daga et al., 2014]. Since the DW-MRI datasets are ultimately used to estimate the presence of white matter areas in the brain, I evaluated the effect of susceptibility artefact correction on residual tensor fitting errors from the DTI model. The rationale for this is as follows. In the DTI model, the MRI signal intensities from the diffusion-weighted images and the b0-DW-



Fig. 3.9: Landmark based evaluation of susceptibility artefact correction. Manually located anatomical landmarks are shown color-coded based on distance from ground truth position (large cross: in-slice, small cross: projection of landmark onto the shown slice). (a,e) Intraoperative T2-weighted image unaffected by the distortion (ground truth image for landmarks). (b,f) An uncorrected susceptibility-distorted b0-DW-MRI image. (c,g) The b0-DW-MRI image corrected using the acquired field map. (d,h) The b0-DW-MRI image corrected using the proposed simulated field map. The top row (a–d) shows the vicinity of the resection cavity (ROI 1 in Table 3.2) and brain stem (ROI 2 in Table 3.2) for subject #3 with resection in the right hemisphere. The bottom row (e–h) shows the area of strong brain surface distortion (ROI 3 in Table 3.2) for subject #4.

Subj.	. ROI 1		ROI 2		ROI 3			ROI 4				
	U	Α	S	U	Α	S	U	Α	S	U	A	S
1	2.80	0.40	1.80	4.86	3.26	4.26	10.75	4.00	8.75	2.80	1.00	2.20
2	2.33	1.17	1.33	4.20	2.62	4.80	4.00	1.45	3.00	1.67	1.00	2.01
3	3.50	0.88	2.37	3.00	1.52	2.18	6.00	2.40	5.60	4.75	3.86	4.89
4	3.51	1.77	1.90	3.40	1.60	1.80	6.07	2.65	6.07	1.94	0.88	1.21
5	2.04	0.79	0.92	3.05	0.68	1.20	8.00	4.75	7.75	2.17	0.67	2.00
6	2.56	1.51	1.81	2.40	1.45	1.10	8.00	4.25	6.75	2.00	1.00	1.50
7	2.83	0.67	1.50	2.95	1.80	1.55	8.75	5.25	7.50	2.00	0.50	1.50
8	3.43	1.28	2.00	2.78	1.20	1.60	8.75	5.00	5.75	0.83	0.84	1.00
9	6.10	4.02	5.56	1.59	0.64	1.00	3.60	1.50	3.40	9.02	7.99	8.26
10	2.48	1.56	2.27	3.60	1.20	2.00	5.60	2.80	5.00	2.67	0.67	1.83
11	6.88	2.20	4.89	2.20	0.40	0.60	2.21	0.75	2.75	2.33	1.33	2.33
12	3.07	1.21	2.25	4.00	1.74	1.60	6.35	3.81	6.75	1.50	0.84	1.17
Avg.	3.46	1.46	2.38	3.17	1.51	1.97	6.51	3.22	5.76	2.81	1.72	2.49

Table 3.2: Misregistration (in mm) between the ground truth landmarks and the landmarks in the uncorrected b0-DW-MRI images (U) and those corrected using acquired (A) and simulated field maps (S). The results are reported for 4 ROIs: (1) near resection margin, (2) brain stem, (3) head-holder fixation pins, (4) rest of the brain.

Misregistration	ROI 1	ROI 2	ROI 3	ROI 4
(a) Corr $/\Lambda EM < Unc$	1.34×10^{-5}	8.91×10^{-8}	1.37×10^{-6}	$\begin{array}{c} \textbf{ROI 4} \\ \hline 1.26 \times 10^{-5} \\ (0.68) \\ \hline 1.02 \times 10^{-2} \\ (0.13) \\ \hline 2.58 \times 10^{-5} \end{array}$
(a) $\operatorname{Cont}/\operatorname{AFM} \leq \operatorname{Chc}$.	(1.25) (1.29) (2.31)	(2.31)	(0.68)	
(b) Corr /SEM < Una	4.34×10^{-6}	1.10×10^{-4}	1.34×10^{-2}	1.02×10^{-2}
(b) $\operatorname{Cont./SFM} < \operatorname{Chc.}$	(0.72)	(0.62)	(-0.012)	(0.13)
(a) Corr $/$ AFM $<$ Corr $/$ SFM	7.16×10^{-4}	1.88×10^{-2}	1.66×10^{-6}	$\begin{array}{c} \textbf{ROI 4} \\ \hline 1.26 \times 10^{-5} \\ (0.68) \\ \hline 1.02 \times 10^{-2} \\ (0.13) \\ \hline 2.58 \times 10^{-5} \\ (0.46) \end{array}$
(c) $\operatorname{Cont}_{\operatorname{AFM}} < \operatorname{Cont}_{\operatorname{AFM}}$	(0.36)	(-0.045)	(1.77)	(0.46)

Table 3.3: Results of one-tailed Student's t-test of landmark misregistration reduction between uncorrected (Unc.) images and those corrected (Corr.) using acquired (AFM) and simulated (SFM) field maps (*p*-value and lower bound of the CI at 1.25% significance level in mm).

MRI image are fitted to the model by regression [Basser et al., 1994]. The fitting is not perfect because of the modeling inadequacy of the DTI model, the limited voxel resolution, and the poor SNR. [Papadakis et al., 2002] proposed the normalised sum of square residual diffusion tensor fit errors (χ^2) as a measure of goodness of fit of the DTI model for a given voxel location. Here, the measure χ^2 is defined as:

$$\chi^{2} = \frac{\sum_{i=1}^{N} \left(S_{m}i - S_{f}i\right)^{2}}{\sum_{i=1}^{N} S_{mi}^{2}}$$
(3.6)

where N signals are fitted, and $S_m i$ and $S_f i$ are the measured signals and the fitted signals, respectively.

The susceptibility artefact causes non-linear (spatially varying) distortion. [Kim et al., 2006] showed that correcting non-linear distortions in a DW-MRI dataset reduced tensor fit errors. This suggests that a reduction in tensor fit errors might be indicative of a better susceptibility artefact correction. To support this assertion further, I note an observation made by [Leemans and Jones, 2009] in a connection with a step for the correction of subject motion during DW-MRI acquisition (this step is clearly not needed for iMRI). In the subject motion correction step, each diffusion-weighted image is registered affinely to the b0-DW-MRI image reference. [Leemans and Jones, 2009] showed that during subsequent tensor fitting, it is important to rotate the direction of each diffusion-encoding gradient. Their reason is that the affine transformation rotates and skews the voxels. Accordingly, the gradient direction of each diffusion-weighted image tensor is that the affine transformation rotates and skews the voxels. Accordingly, the gradient direction of each diffusion-weighted image tensor is that the affine transformation rotates and skews the voxels. Accordingly, the gradient direction of each diffusion-weighted image tensor is that the affine transformation rotates and skews the voxels. Accordingly, the gradient direction of each diffusion-weighted image tensor is that the affine transformation rotates and skews the voxels. Accordingly, the gradient direction of each diffusion-weighted image tensor is that the affine transformation rotates and skews the voxels. Accordingly, the gradient direction of each diffusion-weighted image tensor is that the affine transformation rotates and skews the voxels. Accordingly, the gradient direction of each diffusion-weighted image tensor is that the affine transformation rotates and skews the voxels. Accordingly the gradient direction of each diffusion-weighted image tensor is that the affine transformation rotates and skews the voxels displacement field, which will have a lo

effective direction of the diffusion-encoding gradient will become rotated from the nominal direction. For such voxels the tensor fit error will contain a component *solely* due to the non-linearity. Therefore, a reduction in tensor fit error after the correction for non-linear distortions, such as those due to the susceptibility artefact, should be proportional to the accuracy of the correction.

The diffusion tensors were reconstructed using *dtifit* from the FSL package [Smith et al., 2004], which also calculated the normalised tensor fit error maps. For the validation, tensor fitting was performed on the uncorrected DW-MRI datasets for the 12 subjects, on DW-MRI datasets with every of the 65 diffusion direction images individually corrected using the acquired field maps, and on DW-MRI datasets likewise corrected using the proposed simulated field maps, as shown in Fig. 3.10. The maps of normalised sum of square diffusion tensor fit errors were added up across the whole brain, as reported in Table 3.4. One-tailed paired Student's t-tests showed a statistically significant reduction for DW-MRI datasets corrected using both the acquired field maps ($p = 3.43 \times 10^{-8}$, confidence interval lower bound at 5% significance level: 0.72 mm) and the simulated field maps ($p = 3.56 \times 10^{-8}$, CI: 0.41 mm). The said reduction was significantly larger for DW-MRI datasets corrected using the acquired field maps ($p = 2.32 \times 10^{-6}$, CI: 0.28 mm).



Fig. 3.10: Map of normalised sum of square diffusion tensor fit errors shown for subject #3 calculated from (a) an uncorrected example DW-MRI dataset, (b) the DW-MRI dataset corrected using an acquired field map and (c) a simulated field map.

Subj.	Unc.	AFM	SFM
1	5.16(7.61)	4.07(5.61)	4.60 (6.61)
2	$3.43\ (\ 5.59\)$	2.79(4.21)	3.10(5.02)
3	$3.92\ (\ 6.79\)$	$3.09\ (\ 5.19\)$	$3.29\ (\ 5.57\)$
4	4.96(7.58)	4.16 (6.17)	4.47 (6.70)
5	5.06(6.98)	4.01 (5.30)	4.50 (6.15)
6	$3.48\ (\ 6.57\)$	2.72 (4.81)	3.02(5.42)
7	3.86(6.86)	3.06(5.43)	3.39(6.08)
8	4.27 (6.40)	3.60(5.16)	3.82(5.55)
9	5.39(7.18)	4.37(5.68)	4.72 (6.20)
10	3.33(6.00)	2.65(4.67)	3.07(5.60)
11	5.86 (7.74)	4.60(5.89)	5.27(6.85)
12	$3.28\ (\ 5.55\)$	2.84(4.83)	$2.99\ (\ 5.03\)$
Avg.	4.33 (6.74)	3.50(5.25)	3.85(5.90)

Table 3.4: Mean (std) of normalised sum of square diffusion tensor fit errors summed up across the whole brain volume. The results are reported for the uncorrected (Unc.) DW-MRI datasets and those corrected using the acquired (AFM) and the simulated field maps (SFM). The summary line contains the column averages.

3.4 Discussion

I first briefly discuss the accuracy of the air-tissue segmentation. I then compare the original acquired field maps and the proposed simulated field maps directly. However, since the acquired maps are not in themselves a reliable gold standard, I next compare the effects of susceptibility artefact correction based on either type of maps in several regions of interest.

3.4.1 Accuracy of Air-Tissue Segmentation

I found that the GMM model classified T1w+T2w voxels representing the CSF and the eyes as softtissue, which is correct for susceptibility modelling. The main challenge for an accurate air-tissue segmentation was the ambiguity between the air and bone voxels. The combination of the use of a strong MRF neighbourhood consistency parameter and the smoothing of the air prior ensured that skull voxels were generally correctly classified as bone while resection cavity voxels were correctly classified as air. The strong MRF also forced skull bone marrow, bright in T1w+T2w, to be included in the bone class, which is acceptable due to their virtually identical susceptibility. The distribution for the bone class in the GMM model received a low weight due to the sparsity of the bone voxels in the prior (Fig. 3.5). In future work, the T1w+T2w image could be replaced by multimodal segmentation.

The air cavities (e.g. paranasal sinuses) appear to be slightly oversegmented with respect to their appearance in the synthetic CT, although this is not straightforward to judge without a reliable ground truth. This oversegmentation likely arose due to the compromises regarding parameters (MRF consistency, air prior smoothing) which in turn arose primarily due to the lack of a robust automated method for segmentation of the resection cavity and the craniotomy area. This limitation could be overcome in the future using a shape model such as one based on active contours [Xu and Prince, 1998].

3.4.2 Comparison of Acquired and Simulation Field Maps

The average difference found between the acquired and simulated field maps of 1.15 ± 1.49 mm for all subjects within the brain volume (in terms of displacement along the PE direction) is less than the voxel size of the DW-MRI dataset (2.5 mm, typical for DW-MRI datasets) and has to be considered with respect to the desired resection accuracy, which is patient and surgeon specific and difficult to define. I believe that the observed near 1 mm resection accuracy in areas of low field map difference is clinically useful. However, there are areas of difference between the two types of field maps.

In the absence of ground truth field maps, I used landmark misregistration and sum of diffusion tensor fit errors as surrogate measures for statistical evaluation of susceptibility distortion correction using acquired and simulated field maps. The diffusion tensor fit errors based results indicate that both types of field maps facilitate DW-MRI dataset correction, although the acquired field maps facilitate a better correction compared to the simulated ones (Table 3.4). The landmark based results indicate that the acquired field maps facilitate landmark correction in all ROIs, while the simulated field maps facilitate significant correction in all ROIs except in the regions of high DW-MRI distortion near the head-holder fixation pins (second row, ROI 3, in Table 3.3). Also, in all ROIs except near the brain stem the results due to acquired field maps are better than due the simulated field maps (third row, ROI 2, in Table 3.3).

I attempt to interpret the differences between the field maps in specific areas. The simulated

field maps are more positive in the immediate vicinity (i.e. up to 1 mm) from the resection margin, as shown in Fig. 3.7(j–l), i.e. the simulated field maps denote a further displacement along the PE direction than the acquired field maps. This area is challenging to evaluate using both landmark based and tensor fit error based approach due to the heavy noise in the DW-MRI data. The landmark based evaluation in the area up to approx. 10mm from the resection margin (ROI 1) and in the nearby brain stem (ROI 2) (Fig. 3.9, a–d) is more conclusive and indicates that the acquired field maps provide superior DW-MRI correction in those regions. The simulated field maps tend to co-vary with the acquired field maps in this region (Fig. 3.7, a–i). Also, under the head-holder fixation pins (ROI 3), the simulated field maps are more positive than the acquired field maps. In these regions (Fig. 3.9, e–h) I found during landmark-based evaluation that only the acquired field maps provided significant EPI correction, and hence the simulated field maps are overestimated. This is possibly due to the perturbation of the magnetic field due to the magnetic susceptibility of the titanium material from which the pins are manufactured. I suggest that a separate study of the influence of the titanium pins, for instance based on simple digital phantoms inserted into an existing air-tissue segmentation, be undertaken to investigate the influence of the titanium pins.

I finally observe that near the air-filled petrous part of the temporal bone in both hemispheres and anteriorly in the frontal lobe (superior to the paranasal sinuses) the simulated field maps tend to form regions that are up to 10 mm more positive than the acquired field maps, as shown in Fig. 3.8. This is likely caused by the overestimation of the air-filled cavities as described in Section 3.4.1.

I also note that the present basic method to estimate the room-temperature shim coil currents (Section 3.2.3) likely does not faithfully reproduce the dynamic optimisation implemented in the real iMRI scanner as the latter is potentially dependent on physical constraints such as maximum allowable shim coil current [Gruetter and Boesch, 1992]. Such detailed entries are not part of the image data produced by the iMRI scanner. These data are in the form of files in the DICOM format, which is the standard data format for sharing medical images including but not limited to MRI [Mildenberger et al., 2002]. The DICOM header entries in the format are evidently not specialised enough to include information about shim coil currents. Such detailed entries may potentially be present in the manufacturer specific sub-header.

3.4.3 Running Time of the Pipeline

The time to apply the field maps to a DW-MRI dataset is several seconds. The computational time to phase-unwrap the raw acquired field map is below 1 minute. However, the computational time required to simulate the field map at the full T1-weighted iMRI resolution, in the current unaccelerated implementation, is approx. 1 hour (Intel Core i5 @ 3.30 GHz), out of which the air-tissue segmentation takes approx. 2 minutes. The time given assumes that the pseudo-CT is precomputed or a real CT is imaged prior to the surgery and hence can be discounted from the intraoperative time. Although the proposed simulation method clearly takes too long in its present form to be applicable in an interventional setting, most of it is dominated by forward and inverse FFT (Section 3.2.3) which lends itself relatively easily to GPU acceleration due to widely available libraries.

3.5 Conclusion

I performed a feasibility study on an iMRI dataset of 12 cases of ALTR surgery to investigate whether B_0 field maps can be simulated from intraoperative structural scans. I found a close agreement between the simulated and the original acquired field maps and that the simulated field maps facilitated susceptibility artefact correction in all brain regions of interest, which is an encouraging result. On the other hand I found that the original acquired field maps facilitated statistically more accurate susceptibility artefact correction in all brain regions of interest and attribute this to sensitivity to inaccuracies in air-tissue segmentation and the lack of explicit accounting for the titanium head-holder fixation pins in this study, both of which should be investigated in a further study.

This study also corroborates the usefulness of field map simulation for *non-interventional* datasets as reported in [Poynton et al., 2009]. I found that the inclusion of the spatial prior derived from the synthetic CT yields reliable simulated field maps in most of the brain volume. At the same time, while I found localised inaccuracies, I attribute these primarily due to challenges specific to the intraoperative time-step. Since the widespread neglect of susceptibility artefact correction in

DW-MRI and fMRI studies in general is a serious problem (Section 3.1.2), I suggest a separate study into the use of the proposed approach for non-interventional datasets.

Chapter 4

Registration near Resection Boundary

In the previous chapter the focus was on preparing the data to be located in the respective preoperative and intraoperative image spaces as accurately as possible in accordance with their physical locations. In this chapter I proceed to estimate the non-rigid deformation between the preoperative and the intraoperative image. Previously, GPU-enabled non-rigid registration was successfully applied to the brain shift estimation problem within the time-constraints of neuro-surgery [Modat et al., 2010, Daga et al., 2012]. Likewise, I choose to approach the problem using the framework of non-rigid registration.

Registration schemes are composed of an optimisation strategy, a regularisation method and a similarity measure (Section 2.3). The quality of the registration estimate will thus depend on how well these components account for the issues arising from iMRI imaging and the complex motion due to brain shift. For instance, continuous registration schemes are likely to get stuck in a local optimum for datasets with complex motion, which have many degrees of freedom that need to be recovered: continuous registration schemes perform gradient-based search in the convex basin of the initial (identity) transformation and for data with many degrees of freedom it is relatively unlikely that the search for the transformation will have begun in the basin of the global optimum.

Further, the regularisation scheme may need to take into account specific constraints such as volume preservation within contiguous structures. In this chapter, I focus on the last component i.e. similarity measure, because it drives the registration and thus significantly influences the accuracy and plausibility of the registration.

4.1 Introduction

The complex motion due to brain shift and iMRI imaging present challenges to conventional similarity measures described in Section 2.3.2. The most obvious challenge in MRI/iMRI T1w structural scan pairs is the missing correspondences arising from the presence of the resection in the intraoperative scan. This resection is typically in the form of a resection cavity that is mostly air-filled, however the appearance of the cavity is more complex for many surgery cases and can be filled by surgical gauge or partly filled with a saline solution to control bleeding; for these reasons it may be challenging to produce an accurate binary segmentation delineating the extent of the brain (brain mask) in an fast, automated manner, which is necessary due to the time-contraints of neurosurgery. For instance, the authors in [Daga et al., 2012] employed the widely-used FSL-BET tool to extract the brain mask automatically, however this tool is designed for skull-stripping of structural scans of non-surgical data in the standard anatomical orientation and even with reasonable effort at parameter tuning it tends to oversegment the brain in the intraoperative scan to contain a substantial (20%+) part of the resection cavity near the resection boundary as shown in Fig. 4.1. This poses a problem to conventional similarity measures. For instance, [Daga et al., 2012] evaluate NMI (Eq.2.9) such that only voxels that lie within the brain mask in the intraoperative scan are included in the joint histogram. However, because the brain mask is often oversegmented, this approach may lower the accuracy of motion estimation near the base of the resection cavity where accuracy is especially desirable. In order to avoid reliance on a potentially incorrect externally supplied binary mask, it would be useful to introduce a similarity measure with its own mechanism for segmenting structures.

Another challenge is the complex and heavy deformation and the highly localised nature of brain shift whereby the most severe deformation tends to arise within a few cm from the re-



Fig. 4.1: Brain masks extracted using the FSL-BET fast automatic tool (yellow with blue border). (a) Mask for ATLR case #1 which has an air-filled cavity. The base of the cavity is incorrectly included in the brain mask. (b) Mask for ATLR case #9 whose resection cavity is filled with a saline solution. (The ATLR dataset is described in Section 3.2.4.)

section. In Section 2.3.2, local similarity measures were introduced as better suited to capture local morphological similarity than global measures, because they attempt to measure some local relationship. Typically these measures calculate some statistic from the intensities of the voxel pairs co-located in the reference and floating image, respectively, which are restricted to a local neighbourhood (Section 2.3.2). The chosen statistic is then assigned to the voxel at the centre of the neighbourhood. The calculation of the local statistic typically relies on image smoothing with a static kernel for intermediate steps. However, filtering with a static kernel is not edge-aware: when calculating the local statistics for a given location near a tissue or resection boundary, the voxels located *beyond* the boundary are included in the local statistics, even though they are not representative of the location within the boundary. A more meaningful measure of similarity would arise if the influence of the voxels beyond tissue/resection boundary could be suppressed provided the boundary could be identified in a meaningful way.

A further challenge posed by the combination of brain shift and iMRI imaging is the contrast changes between the preoperative image (which is typically acquired on a conventional MRI scanner), which arise due to CSF drainage, bleeding, edema, bias field and low signal to noise ratio (SNR) of iMRI. These factors induce changes to the appearance of the same anatomical area and may invalidate the local relationship that is assumed by the chosen similarity measure to exist between brain structures such as the cortical white and gray matter and deep structures such as brain stem and hippocampus. It may be beneficial to only assume the measured relationship to exist in the *texture* of tissues by suppressing the contribution of the voxels beyond the borders of tissues. In this regard, it is worth to observe that brain structures tend to have a relatively stable intensity level beyond minor variation due to texture. Therefore, most boundaries of contiguous structures can be assumed to coincide with sudden changes in local intensity.

In this chapter, I aim to address the challenges outlined above. To this end, I propose to use a suitable local similarity measure and introduce an *adaptive weighting* of the local voxels, which endows the measure with an edge-awareness property. I embed this measure in a discrete registration scheme and perform extensive evaluation in several experiments on numerical phantoms, on 12 clinical iMRI cases of ATLR surgery and in a parcellation propagation experiment. These results were first published in a conference paper [Kochan et al., 2016]. Specifically in this chapter, I also add a further validation on 8 iMRI datasets of resection for clinical management of glioma and provide extended explanations and additional figures.

4.2 Adaptively Weighted Local Similarity Measure

In line with the reasoning presented in the previous section, I propose to use a local similarity measure in order to account for tissue deformations and contrast changes between the pre/intraoperative structural T1w scans. Since the image pair is unimodal, the LNCC (local normalised correlation coefficient) measure (Section 2.3.2), which measures the local affine intensity relationship between a pair of images, is sufficient and can serve as a simple local similarity measure [Cachier et al., 2003]. I note that for multi-channel and/or multi-modal images, further suitable local similarity measures are also outlined in the same section, and so LNCC is considered in this chapter as an example of a local similarity measure without a loss of generality.

Further, in order to account for the presence of the resection cavity in the intraoperative image and to also account for more severe contrast changes that arise between the image pair, I propose to introduce an adaptive weighting scheme into the local similarity measure. The basic tenet of the proposed adaptive weighting is that the contribution of a given voxel in the local neighbourhood to the evaluation of the measure should be weighted based on whether it is likely to be from within or from beyond a resection boundary (or, alternatively, a tissue boundary). Such a weighting can also be regarded as one that introduces an edge-awareness property into the similarity measure. In order to choose a suitable adaptive weighting scheme, I draw from the concept of *bilateral filtering*, in which a mean of intensities for a local neighbourhood is calculated by weighting the influence of individual voxels based on their spatial distance to the central voxel of the neighbourhood *and* also based on the difference between their intensity and the intensity of the central voxel [Tomasi and Manduchi, 1998]. Thus, in such adaptive weighting, voxels beyond an intensity edge become down-weighted.

An adaptive weighting approach inspired by bilateral filtering was used by [Yoon and Kweon, 2006] in the context of a stereo reconstruction problem, which is a 1D registration problem whereby 1D displacements along scanlines (disparities) are sought between a "left" and a "right" 2D image. More specifically, [Yoon and Kweon, 2006] presented a patch-based SSD (Section 2.3.2) similarity measure that was tailored to the stereo reconstruction problem by weighting of contribution of a given pixel location in a patch to the total SSD value based on the colour-space distance of the "left" and the "right" pixels from the central pixel of the "left" and "right" patch, respectively. This enabled a more accurate disparity search, because pixel that lie on the same surface in the 3D scene tend to have the same colour (or slightly varying shades of the same colour). It is clear that this reasoning can be trivially translated to the registration of gray-scale (intensity-based) images by simply replacing the colour-space distances with intensity differences.

In the following I define LNCC more formally and proceed to introduce the proposed bilateral adaptive weighting into LNCC, following closely from the approach of [Yoon and Kweon, 2006], in more concrete terms. Let R be the reference volume and F the floating volume defined on the same voxel grid and in the same coordinate space, then LNCC for the local neighborhood of a point \vec{v} is defined as

$$\mathrm{LNCC}_{\vec{v}}(R,F)^2 = \frac{\langle R,F\rangle_{\vec{v}}^2}{\langle R,R\rangle_{\vec{v}} \cdot \langle F,F\rangle_{\vec{v}}},\tag{4.1}$$

where the $\langle R, R \rangle_{\vec{v}}$ and $\langle F, F \rangle_{\vec{v}}$ are the local variances and $\langle R, F \rangle_{\vec{v}}$ is the local covariance. The

latter is defined as

$$\langle R, F \rangle_{\vec{v}} = \overline{R \cdot F}_{\vec{v}} - \overline{R}_{\vec{v}} \cdot \overline{F}_{\vec{v}}, \tag{4.2}$$

where $\overline{R}_{\vec{v}}$ and $\overline{F}_{\vec{v}}$ are the respective local means for the point \vec{v} . The local variances are defined analogously. The local mean for R is defined as

$$\overline{R}_{\vec{v}} = \frac{1}{N} \sum_{\vec{x}} R(\vec{v} - \vec{x}) w_{\vec{v}}(\vec{x}), \qquad (4.3)$$

where N is the number of voxels in the neighborhood of \vec{v} , \vec{x} is the offset relative to \vec{v} and $w_{\vec{v}}(\vec{x})$ are the weights, here given by a generic term that depends on \vec{v} . The local mean for F is defined analogously. The local weights can follow any kernel, with the box filter arguably being the simplest but anisotropic (biased toward the x, y and z axes). In the following the 3D Gaussian kernel

$$w_{\vec{v}}(\vec{x}) = G_{\beta}(\vec{x}) = \frac{1}{\sqrt{2\pi\beta}} \exp\left(-\frac{|\vec{x}|^2}{2\beta^2}\right),\tag{4.4}$$

is used, where β controls the neighborhood's size (negligible for $|\vec{x}| > 3\beta$). The 3D Gaussian kernel is isotropic and also separable into three 1D Gaussian filters i.e. $G_{\beta}(\vec{x}) = 2\beta G_{\beta}(x)G_{\beta}(y)G_{\beta}(z)$ where $\vec{x} = [x, y, z]^{\mathrm{T}}$. If the identical kernel is applied for all neighbourhoods independently of \vec{v} then the local mean can be implemented using convolution

$$\overline{I}_{\vec{v}} = (G_{\beta} * I)(\vec{v}). \tag{4.5}$$

The convolution operation can be implemented fast thanks to the separability of the Gaussian filter.

I introduce bilateral adaptive weighting whereby a unique kernel is used for every neigbourhood centre location \vec{v} . The modified measure is referred to as LNCC-AW. A bilateral filtered smoothing of an arbitrary image I is

$$\overline{I}_{\vec{v}}^{\text{bilat.}} = \frac{1}{N} \sum_{\vec{x}} I(\vec{v} - \vec{x}) \cdot G_{\beta}(\vec{x}) \cdot G_{\alpha} \left(I(\vec{v} - \vec{x}) - I(\vec{v}) \right), \tag{4.6}$$

where $G_{\alpha}(d)$ is a range kernel i.e. a kernel for the intensity difference $d = I(\vec{v} - \vec{x}) - I(\vec{v})$.



Fig. 4.2: (a) T1-weighted reference image and (b) intensity-based weights for a point (blue cross). (c) T1-weighted floating image and (d) intensity-based weights for the point. (e) Final weights based on (b), (d) and distance from the point.

The edge-preserving property arises as the voxels beyond an intensity rise/drop are excluded. Given images R and F to register, the adaptive weighting is guided by both the images as in [Yoon and Kweon, 2006] by using the composite term

$$w_{\vec{v}}(\vec{x}) = G_{\beta}(\vec{x}) \cdot G_{\alpha} \left(R(\vec{v} - \vec{x}) - R(\vec{v}) \right) \cdot G_{\alpha} \left(F(\vec{v} - \vec{x}) - F(\vec{v}) \right)$$
(4.7)

as illustrated in Fig. 4.2. However, since the weights $w_{\vec{v}}(\vec{x})$ now vary spatially, it is no longer possible to implement the local mean using the convolution operation. Similarly, it is no more possible to take advantage of kernel separability. The local mean is implemented directly to assess the benefit of the weighting.

The range kernel $G_{\alpha}(d)$ determines the adaptive weighting. A desired property is to facilitate boundary-preservation, which implies a need for a limited width of the kernel. On the other hand, the kernel should preferably not eliminate the whole intensity range, for two reasons. Firstly, most brain structures are homogeneous (have little or no texture) and meaningful correlation can only



Fig. 4.3: The employed range kernel is more long-tailed than the Gaussian kernel.

be measured if all the structures are included to some degree. Secondly, iMRI is relatively noisy and a narrow instensity range would give rise to inconsistent weighting across homogeneous areas as the intensity of the central voxel of each neighbourhood varied with the noise. The original paper [Yoon and Kweon, 2006] used the Gaussian kernel $G_{\alpha}(d) = (1/\sqrt{2\pi\alpha}) \exp(d^2/2\alpha^2)$ for colour-space distance range. The Gaussian kernel can include a wide intensity range if a large enough α is chosen, but this involves sacrificing the strength of achievable boundary-preservation. I instead propose to use a range kernel shaped like the Student's *t*-distribution, which has a more gradual drop-off and a longer tail, thus allowing a lower value of α which does not involve compromising on edge-preservation:

$$G_{\alpha}(d) = \frac{\Gamma(\frac{\nu+1}{2})}{\sqrt{\nu\pi\alpha^2}\Gamma(\frac{\nu}{2})} \left(1 + \frac{d^2}{\nu\alpha^2}\right)^{-\frac{\nu+1}{2}}.$$
(4.8)

I choose $\nu = 2$ degrees of freedom, which corresponds to the most long-tailed one in the family of Student's *t*-distributions. See Fig. 4.3 for a comparison of the kernels. For $\alpha = \infty$, the weighting reduces to locally non-varying, and correspondingly, LNCC-AW reduces to LNCC.

4.3 Registration using a Discrete Optimisation Framework

The derivation of analytical gradient of the similarity measure, for instance with respect to a voxel-based deformation field, for use in gradient-based non-rigid registration schemes becomes complicated when using adaptive weighting, because the gradient depends on the local weights which in turn depend on the deformation. However, [Glocker et al., 2008a] reformulated non-rigid registration as a discrete Markov Random Field (MRF) optimisation problem, for which the similarity measure gradient is not needed.

The proposed measure is employed in a related discrete optimisation scheme of [Heinrich et al., 2013]. A grid \mathcal{P} of B-spline transformation control points $p \in \mathcal{P}$ with positions \vec{c}_p is overlaid onto the reference image. The control point displacements in the floating image are $\vec{u}_p = [u_p, v_p, w_p]$ with discrete valued components. For efficiency, a minimum spanning tree \mathcal{N} of the most relevant edges $(p, q) \in \mathcal{N}$ is optimised rather than a full MRF. Displacements are sought minimising the energy

$$\sum_{p \in \mathcal{P}} \left(1 - \| \text{LNCC}_{\vec{c}_p}(R(\vec{x}), F(\vec{x} + \vec{u}_p)) \| \right) + \alpha \sum_{(p,q) \in \mathcal{N}} \frac{\| \vec{u}_p - \vec{u}_q \|^2}{\| \vec{x}_p - \vec{x}_q \|}.$$
(4.9)

A multi-level scheme is used: the reference and floating image are used at full resolution at all levels, while a B-spline control point grid is used with increasing spacing of 7, 6, 5, 4 and 3 voxels at each level; the control point displacements estimated at a given level serve to initialise the subsequent level. 3D displacement labels allowed for each level l = 1...5 are enumerated from the set $s_l \cdot \{-r_l, ..., +r_l\}^3$ where $s_l \in [5, 4, 3, 2, 1]$ are the step-sizes and $r_l \in [6, 5, 4, 3, 2]$ are the ranges.

The above scheme does not guarantee symmetry of the transformation and is biased toward the reference image. However, the choice of the preoperative and/or intraoperative image as the reference is in principle ambivalent. Moreover, the true physical deformation is invertible in the areas where valid correspondences exist between the image pair. Therefore the recovered deformation should be close to invertible in order to be physically plausible. The symmetry of the final transformation is enforced as follows [Heinrich et al., 2013]. Both forward *and* backward control point dis-

placements are found independently using the above scheme. Then, bespoke diffeomorphic forward and backward transformations are calculated by recasting the B-spline fields as stationary velocity fields and integrating them using the scaling-and-squaring approach of [Arsigny et al., 2006]. Diffeomorphisms have a unique inverse transformation and a unique mid-point transformation (i.e. a mapping to the mid-point image space). Thus the final transformation which is symmetric can be found by composing the mid-point of the forward diffeomorphism with the inverse of the mid-point of the backward diffeomorphism.

4.4 Experiments

4.4.1 Patch Matching on 2D Synthetic Phantom

The matching accuracy is compared for two 2D synthetic phantoms. A fixed patch, representing a local neighborhood, is placed in the reference image centred at a point of interest. A moving patch, representing an independently moving local neighborhood, is placed in the floating image and is centred initially at the same grid location but allowed to move with respect to this location. A similarity profile is plotted of LNCC and LNCC-AW, respectively, which is the value of the similarity measure plotted as a function of the displacement between the fixed and moving patch centres. Two phantoms are assessed. A contrast-enhanced lesion near a resection phantom is shown in Fig. 4.4(a-d). The similarity profile for LNCC has a mild maximum at the true zero displacement due to voxels included from the resected area. The similarity profile for LNCC-AW has a clear maximum due voxels down-weighted in the resected area. A phantom of the *medial* longitudinal fissure is shown in Fig. 4.4(e-h). The reference and floating image are the same axial slice from the BrainWeb database. The patch is centered next to the medial longitudinal fissure that contains dark voxels in the CSF and the falx cerebri. The similarity profile for LNCC has a band of false matches due to voxels included in the fissure. The similarity profile for LNCC-AW has a unique maximum at the true zero displacement due to these voxels being down-weighted. The similarity profiles for both phantoms indicate that the contribution of the adaptive weighting is to remove some false-positives in matching where at least one of the patches is near an intensity edge.



Fig. 4.4: 2D numerical phantoms. (a–d) A contrast-enhanced lesion near a resection. (e–h) The medial longitudinal fissure. (a,e) The reference image and the fixed patch. The cross shows the location of the centre of the fixed patch. The outline shows the size of the fixed patch. (b,f) The floating image and the moving patch. The cross shows the location of the centre of the moving patch. The inner outline shows the size of the moving patch at zero displacement. The inner outline shows the size of the moving patch at zero displacement. The inner outline shows the size of the moving patch at zero displacement. The inner outline at maximum allowed displacement). (c,g) The similarity profile of LNCC as a function of the displacement of the moving patch. Neg. (pos.) x displacement is left (right) and neg. (pos.) y displacement is down (up). (d,h) The same for LNCC-AW.

4.4.2 Recovery of a 3D Synthetic Deformation

I perform a registration experiment on a BrainWeb dataset. The reference image is made by inserting a synthetic resection cavity in the right temporal lobe. The floating image is resampled using B-spline interpolation from the BrainWeb image using a synthetic sinusoidal deforming field (period 100 mm in all directions, displacement amplitude 4 mm). The voxel intensities are normalised to 0–1 range. The multi-level discrete symmetric registration scheme was used as described in Section 4.3 and was run for LNCC ($\beta = 5$ mm) and twice for LNCC-AW ($\beta = 5$ mm, $\alpha = 0.30$ and $\alpha = 0.10$).

The registration accuracy is quantified using landmarks found in the reference using 3D-SIFT [Toews and Wells, 2013]. I include 43 landmarks from a 2 cm region from the resection margin. The landmarks are propagated using the true and recovered deformations. The target registration error (TRE) is shown in Fig. 4.5(c). TRE for LNCC-AW is significantly lower (for both



Fig. 4.5: Axial view of 3D BrainWeb based phantom. (a) Reference image (inserted resection). (b) Floating image (synthetic deformation). (c) Target registration error. (d) Map of log Jacobian determinant for ground truth deformation (forward field). (e) Same map for fields recovered using LNCC, (f) LNCC-AW with $\alpha = 0.30$ and (g) LNCC-AW with $\alpha = 0.10$.

 $\alpha = 0.30$ and $\alpha = 0.10$) than for LNCC (paired *t*-tests, p < 0.001). The degree of the reduction is apparent from Fig. 4.5(c) and could be partly owing to the fact that local appearance is relatively conserved between the image pair and the ground truth is exact.

I also visually assess the smoothness of the recovered deformation by calculating the absolute log Jacobian determinant map. The Jacobian determinant measures local volume contraction and expansion assuming a linearisation of the deformation at a voxel, while the absolute log transformation enables contraction and expansion to be treated in the same way. The log of Jacobian determinant maps for the true and the recovered deformation are shown in Fig. 4.5(d–g). The deformations recovered using LNCC-AW follow the true deformation closer than using LNCC.

4.4.3 Evaluation on iMRI Surgical Datasets

Anterior Temporal Lobe Resection Dataset

I evaluate the measure on 12 cases of anterior temporal lobe resection (ATLR, Section 1.1). The dataset is described in Section 3.2.4. For each case, I perform the following pre-processing and

registration steps:

- The preoperative and the intraoperative images are skull-stripped using the FSL-BET tool.
- Both images are normalised by mapping the 1st–99th intensity percentile linearly onto the intensity range 0–1.
- The intraoperative image is resampled to a resolution of $1.1 \times 1.1 \times 1.1$ mm.
- The preoperative image is selected as the floating image and is registered affinely to the intraoperative image that is selected as the reference image. The affine registration is performed using a method described in [Modat et al., 2014], which also resamples the registered image onto the same image grid as the reference image. The resulting image can be considered as the floating image for the purposes of non-rigid registration and will be simply referred to as such in the subsequent steps.
- Auxiliary guidance images are created for the reference and the floating image, respectively, by applying the bilateral filter defined in Eq. 4.6 onto either image with the filtering parameters $\beta = 2.2$ mm and $\alpha = 0.03$. I empirically found these parameters to yield a mild smoothing in homogeneous areas whilst preserving the edges in the clinical images.
- The floating image is registered non-rigidly to the reference image using the proposed non-rigid registration scheme. LNCC with the parameter value β = 5.5 mm is used as the similarity measure. The discrete optimisation parameters of are identical (in voxel units) to those used in Section 4.4.2. The guidance image pair is used to provide the image intensities for the adaptive weighting in Eq. 4.7, in order to limit the influence of the signal noise in the iMRI image on the weighting.
- Finally, non-rigid registration is performed using the same procedure as in the previous step but with LNCC-AW chosen as the similarity measure. The registration is performed once for $\alpha = 0.30$ and again for $\alpha = 0.10$, with $\beta = 5.5$ mm in both cases.

The non-rigid registration takes approx. 10 hours per subject using 4 threads on a computing cluster node. The outcome of the registration for the surgical case #3 with LNCC and LNCC-AW ($\alpha = 0.10$) is shown in Fig. 4.6. The left column of the figure shows the intraoperative reference

in three nominal anatomical planes of the scan when the patient is in the intraoperative position. The second column (Fig. 4.6(b,f,j)) shows the affinely registered preoperative structural scan in the same planar views. Overlaid on the views is the extent of the optic radiation tract (ORT), which has been parcellated by a clinical expert on a statistical map resulting from the probabilistic tractography inferred from the preoperative DW-MRI dataset; the procedure to perform the ORT parcellation is described in more detail in [Winston et al., 2012]. The third and the last columns in the figure show the preoperative scan warped using the estimated non-rigid deformation due to LNCC (Fig. 4.6(c,g,k)) and LNCC-AW (Fig. 4.6(d,h,l)), respectively. In the latter two columns, the ORT parcellation is warped using the same deformation fields used for the structural scan. Therefore, the parcellation "moves" with the structural scan, which makes it easier to appreciate the qualitative difference between non-rigid registration due to LNCC and LNCC-AW, respectively. For instance, the brain shift estimate due to LNCC, as seen in the plane Fig. 4.6(g), would map the ORT just at the resection margin. By contrast, the estimate due to LNCC-AW in the same plane, as seen in Fig. 4.6(h), would place the tract mostly *outside* of the resection margin. This suggests that the application of the proposed adaptive weighting could have a clinical significance. Unfortunately, tractography on the *intra*operative DW-MRI dataset is considerably complicated due to several factors, such as the low signal to noise ratio of iMRI, the presence of the susceptibility artefact due to the air-tissue boundary (Chapter 3) and also due to a lack of reliable choice of seed points in the optic tract for the tractography algorithm due to the resection. Therefore, I assess the quality of registration without relying on DW-MRI data from iMRI.

For each case, I annotated 50–60 point-based landmarks pairs in the pre/intraoperative image a few cm from the resection margin. The landmarks were propagated using the recovered deformations. The target registration error (TRE) of landmarks is summarised in Table 4.1 and is significantly lower for LNCC-AW with $\alpha = 0.30$ (paired *t*-test, p = 0.0236) and LNCC-AW with $\alpha = 0.10$ (p = 0.0054), respectively, than for LNCC. I take the confidence interval (CI) at 5% false-positive rate as the effect size. The CI is 0.440 - 0.200 mm for $\alpha = 0.30$ and 0.010 - 0.126 mm for $\alpha = 0.10$, respectively, well below the image resolution of 1.10 mm. One reason for the detected effect being small may be that the landmark pairs that are reliably identifiable as homologous pairs are located away from the area of severe brain shift, which is highly localised, and are thus



Fig. 4.6: Registration of ATLR resection case #3. Views are through the (a—d) axial, (e—h) coronal and (i—l) saggital plane. (a, e, i) Reference volume (intraoperative T1w iMRI with patient's head in intraoperative orientation). Landmarks annotated within the slice are shown as yellow crosses (smaller if out-of-slice). (b, f, j) Floating volume (pre-craniotomy T1w iMRI). Homologous landmarks are shown. Reference brain surface is outlined in white for comparison. The optic radiation tract (ORT) parcellated in the preoperative scan space is shown. (c, g, k) Warped volume (and location of ORT) for deformation obtained using LNCC with $\beta = 5.5$ mm. (d, h, l) Warped volume (and locatiod of ORT) for deformation obtained using LNCC-AW with $\beta = 5.5$ mm and $\alpha = 0.10$. The last row shows the estimated deformation in axial view. (m, n) Magnitude of displacement for LNCC and LNCC-AW, respectively [0 15 mm]. (o, p) Absolute log Jacobian determinant for LNCC and LNCC-AW, respectively [-1 1].

Case	Affine	LNCC	$\mathbf{LNCC-AW}$ $(\alpha = 0.30)$	$\mathbf{LNCC-AW}$ $(\alpha = 0.10)$	9	r		
1	6.55(2.53)	1.79 (1.14)	1.64(0.98)	1.51(0.88)				
2	2.79 (4.24)	2.54 (4.05)	2.38 (4.04)	2.30 (4.06)				-
3	2.90 (1.86)	1.73 (1.41)	1.73(1.36)	1.70(1.37)	6 -			-
4	3.25 (1.90)	1.82 (1.28)	1.69(1.27)	1.76 (1.21)	5 -			-
5	4.54 (1.98)	1.99 (1.02)	1.97(0.97)	1.89(0.91)	4			
6	4.69 (2.70)	2.72(2.23)	2.63(2.19)	2.51(2.05)				
7	3.29 (1.61)	1.68 (1.00)	1.58(0.99)	1.50(1.01)	3	T	<u> </u>	1
8	3.89 (2.00)	1.38 (0.97)	1.17(0.64)	1.14 (0.60)	2 -			
9	7.08 (2.62)	3.04 (3.34)	3.15(3.26)	3.20 (3.31)	1	1	Ĺ.	<u> </u>
10	4.82 (1.84)	1.87 (1.00)	1.78 (1.00)	1.76 (1.05)				
11	7.72 (2.81)	2.12 (1.20)	2.10 (1.20)	2.01(1.21)	Affine	LNCC	LNCC-AW	LNCC-AW
12	6.39 (2.02)	1.91 (1.16)	1.94 (1.12)	1.86(1.03)			(a=0.30)	(a=0.10)

Table 4.1: Mean (std) mandmark misregistration for the epilepsy (ATLR) iMRI dataset.

away from the intensity-edge of the resection boundary where the adaptive weighting would be reasonably expected to have effect. Another reason for the small detected effect may be ambiguity in the localisation of manually annotated landmarks as these are linear/planar in nature.

The manual annotation was performed by the author. One reason why the annotation was not performed by an expert radiologist is that to sample a representative number of landmark pairs (point-picking) is a very time-consuming process that needs to be repeated for all the datasets. A similar problem was faced by [Mercier et al., 2012] who annotated landmarks in iUS image pairs manually over the course of several months. Further, the presence of brain shift complicates the annotation as the available viewing software to which radiologists are accustomed uses viewing through the three standard anatomical planes (coronal, sagittal, axial). However, the motion due to brain shift has been described as complex and difficult to characterise [Nimsky et al., 2000, Nabavi et al., 2001, Ji et al., 2014] and can exhibit considerable out-of-plane movements. A less severe but comparable problem arises for MRI datasets capturing inhale and exhale time-steps of lung motion due to the sliding motion of the lungs. The research work on the development of reliable non-rigid registration algorithms for lung motion recovery has been partly motivated by a recognition that a visual comparison of the lung inhale/exhale datasets can be error-prone for expert radiologists [Murphy et al., 2011].

I assess the smoothness of the estimated deformation using the absolute log Jacobian determinant map as per above. I limit the region of interest (ROI) to within the brain mask less than 2 cm from the base of the resection cavity, which I locate manually in the iMRI volume. The
Case	LNCC	$LNCC-AW (\alpha = 0.30)$	$LNCC-AW (\alpha = 0.10)$	0.3		
1	0.219 (0.181)	0.151 (0.113)	0.219 (0.196)	0.25		-
2	0.151 (0.118)	0.113 (0.080)	0.154 (0.120)		t	
3	0.127(0.113)	$0.091 (\ 0.065 \)$	0.128(0.099)	0.2		-
4	0.180 (0.188)	0.119 (0.091)	0.158(0.168)			
5	0.168 (0.128)	0.133 (0.091)	0.132 (0.104)	0.15		
6	0.161 (0.181)	0.143(0.142)	0.152(0.168)			
7	0.178(0.160)	0.083(0.085)	0.132(0.127)	0.1		
8	0.120 (0.142)	0.099(0.078)	0.108 (0.113)			
9	0.120 (0.100)	0.133(0.099)	0.132(0.111)	0.05		-
10	0.151 (0.181)	0.111 (0.104)	0.137(0.148)			
11	0.222 (0.148)	0.153(0.098)	0.209 (0.141)	LNCC	LNCC-AW	LNCC-AW
12	0.128 (0.126)	0.078(0.059)	0.110 (0.090)		(a=0.30)	(a=0.10)

Table 4.2: Mean (std) of abs. log Jacobian determinant map in vicinity of resection for the ATLRiMRI dataset.

means within this ROI are shown in Table 4.2 and are significantly lower for LNCC-AW with $\alpha = 0.30$ (paired *t*-test, p = 0.0133) and for LNCC-AW with $\alpha = 0.10$ (p < 0.001), respectively, than for LNCC. In the context of the clinical dataset, the local smoothness of the estimated deformation can be regarded as a surrogate measure of biomechanical plausibility of the recovered deformation: when the mechanical forces are at equilibrium, the strain is likely to vary smoothly within contiguous structures i.e. ones with homogenous physical properties and it step-changes tend to arise at boundaries of structures such as the ventricles, the tentorium, the falx cerebri, the cortical surface, the walls of limbic system structures, the tumour boundary. Therefore a smoother estimated deformation may indicate (although not imply) a more plausible estimate of motion due to brain shift.

Glioma Resection Dataset

I also evaluate the measure on 8 cases of resective neurosurgery for the management of glioma brain tumour. The iMRI datasets were acquired during surgeries performed in the iMRI operating suite at the NHNN. The tumour location and the degree of tumour mass effect varied between cases but in general the severity of brain shift was considerably stronger than for the ATLR dataset. In all cases a T1w intraoperative iMRI scan from an advanced stage of surgery was used as the reference while a T1w iMRI scan taken immediately prior to craniotomy was used as the floating volume, in this time-step the patient's head (skull) is at the same angle as later during the procedure

Case	Affine	LNCC	$ \begin{array}{c} \mathbf{LNCC-AW} \\ (\alpha=0.10) \end{array} $	8 - - 7 - 1		
1	3.14(1.65)	1.84 (1.27)	1.18 (0.52)	6		
2	8.01 (3.37)	2.28(1.25)	2.26(1.31)	5		
3	3.66(1.29)	2.24(0.97)	1.98(0.82)	4		
4	3.61(1.96)	1.58(0.88)	$1.35\ (\ 0.83\)$			
5	5.97(1.47)	1.56(0.76)	$1.33\ (\ 0.76\)$	3-		
6	2.22(0.83)	1.51(0.89)	1.27 (0.84)	2		
7	5.69(2.99)	1.87(1.22)	1.84(1.20)	1	1	
8	2.53(1.01)	1.20(0.87)	1.04(0.55)			
				Affine	LNCC	L

9 -

Table 4.3: Mean (std) mandmark misregistration for the glioma resection iMRI dataset. (The box-plot is illustrative only due to small number of samples.)

and hence no affine registration was needed. The iMRI volumes were acquired using the same imaging protocol as the one for T1w volumes described in Section 3.2.4. Written informed consent was obtained from all participants. I apply the same image preprocessing, the same registration scheme and the same choice of parameters as for ATLR registration above. I only consider LNCC-AW parameter value $\alpha = 0.10$. I use reference and floating volumes for guidance directly without intermediate bilateral filtering. The results for glioma case #3 are shown in Fig. 4.7.

I annotated 30–50 representative landmarks near the resection cavity to get a representative sample of motion due to brain shift. The number was lower than for ATLR because the there are typically fewer reliably identifiable landmark pairs. The target registration error (mean misregistration) for propagated landmarks (Table 4.3) is significantly lower for LNCC-AW than LNCC (paired *t*-test, p = 0.0098, 0.090–0.48 mm confidence interval at 5% significance level). The detected effect is below voxel size but is approx. twice as strong as for ATLR, which is in line with the brain shift being more severe. The same limitations for use of manually annotated landmarks apply as discussed for ATLR.

I evaluate the absolute log Jacobian determinant map in a spherical ROI within the brain within 2 cm from the resection cavity base. The mean values are shown in Table 4.4 and are significantly lower for LNCC-AW with $\alpha = 0.10$ than for LNCC (paired *t*-test, p = 0.0309). There is a marked improvement in the smoothness of the estimated deformation, which indicates (although does not imply) a more plausible registration.



Fig. 4.7: Registration of glioma resection case #3. Views are through the $(\mathbf{a}-\mathbf{d})$ axial, $(\mathbf{e}-\mathbf{h})$ coronal and $(\mathbf{i}-\mathbf{l})$ saggital plane. $(\mathbf{a}, \mathbf{e}, \mathbf{i})$ Reference volume (intraoperative T1w iMRI with patient's head in intraoperative orientation). Landmarks annotated within the slice are shown as yellow crosses (smaller if out-of-slice). $(\mathbf{b}, \mathbf{f}, \mathbf{j})$ Floating volume (pre-craniotomy T1w iMRI). Homologous landmarks are shown. Reference brain surface is outlined in white for comparison. $(\mathbf{c}, \mathbf{g}, \mathbf{k})$ Warped volume for LNCC. $(\mathbf{d}, \mathbf{h}, \mathbf{l})$ Warped volume for LNCC-AW. The last row shows the estimated deformation in saggital view. (\mathbf{m}, \mathbf{n}) Magnitude of displacement for LNCC and LNCC-AW, respectively [0 - 10 mm]. (\mathbf{o}, \mathbf{p}) Absolute log Jacobian determinant for LNCC and LNCC-AW, respectively [-1 - 1].

			0.0		
Case	LNCC	$LNCC-AW$ ($\alpha = 0.10$)	0.25 -	 	-
1	0.044 (0.062)	0.065 (0.081)	0.2		-
2	$0.206\ (\ 0.155\)$	$0.179\ (\ 0.132\)$			T
3	0.134 (0.111)	$0.085\ (\ 0.091\)$	0.15		
4	0.195 (0.201)	0.058(0.047)	0.1		
5	0.138(0.102)	0.078(0.051)	0.1		
6 7	0.078(0.090)	0.073(0.087)	0.05		
1	0.234(0.109)	0.109(0.117)		1	-
0	0.004 (0.000)	0.033 (0.040)	₀└	LNCC	LNCC-AW
				2.100	(0.10)

0.3

Table 4.4: Mean (std) of abs. log Jacobian determinant map in vicinity of resection for the glioma iMRI dataset. (The box-plot is illustrative only due to small number of samples.)

4.4.4 Segmentation Propagation Experiment

The results of the experiments on clinical iMRI data indicate reduced target localisation error and smoother recovered deformation for LNCC-AW with respect to LNCC. However, the accuracy of the annotated landmarks is limited and smooth recovered field is only a surrogate measure of accurate registration. I therefore add another experiment to corroborate the above results, based on inter-subject registration of non-surgical structural scans.

I use a database of 35 T1w scans with parcellations of 140 key structures provided by Neuromorphometrics for the MICCAI 2012 Grand Challenge and Workshop on Multi-Atlas Labeling ¹. The image intensities are normalised and each image is used as a reference image and the remaining images as floating images. For each of the 1190 image pairs, I perform affine registration and non-rigid registrations using LNCC ($\beta = 5 \text{ mm}$) and LNCC-AW ($\beta = 5 \text{ mm}, \alpha = 0.10 \text{ only}$) using discrete registration parameters as above. The floating image segmentations are propagated using nearest-neighbor interpolation and calculate Dice score for each label. The average Dice score for 1190 affine registration image pairs is 0.422 ± 0.00187 , for LNCC based non-rigid registrations it is 0.517 ± 0.0101 and for LNCC-AW based registrations it is 0.526 ± 0.00947 . Average Dice score is significantly higher when using LNCC-AW than LNCC ($p < 10^{-6}$).

The measurement of mapping based on overlap of parcellations enables a principled assessment of mapping as it takes into account the region based nature of the human cortex, which is potentially

¹https://masi.vuse.vanderbilt.edu/workshop2012/index.php/Challenge_Details

missing prostical mass

more meaningful than to use approximate point landmarks if there is a missing practical means to account for their ambiguity. Further, as the parcellations are relatively small and densely packed, a higher degree of overlap may indicate a more plausible estimated dense deformation field. The use of all image pairs increases the statistical power of the analysis.

4.5 Discussion and Conclusion

I introduced bilateral adaptive weighting into a local similarity measure (LNCC). The modification facilitated a more accurate landmark localisation in several T1w registration experiments. In a study on clinical iMRI data, a smoother deformation was recovered near the resection margin, which is biomechanically more plausible and potentially enables a more accurate surgical guidance near the resection margin. The brain shift I assessed in the experiments on clinical iMRI data arose from CSF leakage and postural drainage in the ATLR cases, and also from mass effect at the margins of tumours. However, in principle the proposed approach can improve accuracy near distinct intensity edges at collapsed cysts or haematomas from bleeding into the brain, which could be confirmed separately if additional iMRI data became available.

The unoptimised bilateral weighting introduced a time bottleneck which resulted in registration time of around 10 hours per volume pair, which prevents an application of the technology based on the proposed scheme in an intraoperative setting. I envisage that this limitation should be first addressed through the use of an algorithm with a better time-complexity before additional avenues are considered. The discrete optimisation and MRF-based regularisation only take approx. one minute of the running time of the scheme. Powerful options are open toward optimising bilateral weighting, such as guided image filtering [He et al., 2013].

The B-spline control point grid spacings (7, 6, 5, 4 and 3 voxels at consecutive levels) enabled a gradual refinement of the estimated deformation from a coarser to a finer level. This was motivated by the need to provide a high chance of reaching the appropriate global optimum in each step, owing to the high number of degrees of freedom inherent in deformations due to brain shift. However, the results of using an exponential spacing such as 32, 16, 8, 4, 2 voxels at consecutive level, or some intermediate spacing, could also be evaluated in a separate work.

CHAPTER 4. REGISTRATION NEAR RESECTION BOUNDARY

In the present chapter, I considered the problem of unimodal registration of the MRI/iMRI T1w structural scan pair. Therefore, LNCC was appropriate as a testing case for adaptive weighting. However, the single-channel nature of LNCC prevents its application to multi-channel registration if additional intensity channels become available. Examples of additional iMRI channels could include a FLAIR T1w structural scan in which fluids are nulled in order to reveal the metabolically active part of residual tumour (Section 1.1) or voxel-based maps derived from DW-MRI data using the DTI model, such as FA (Chapter 3), mean diffusivity (MD, [Basser et al., 1994]) or simple diffusion direction "colour" maps (e.g. the projections of the DTI-fitted eigenvector with the highest eigenvalue into the x, y and z axes). A multi-channel generalisation of LNCC was recently presented by [Heinrich et al., 2014b] and the bilateral weighting could be extended to improve the said local measure.

Further, while the proposed approach improves the specificity of the similarity measure, the estimated deformation field could also be regularised using bilateral filtering. This approach is applicable with implicit regularisation methods (Section 2.3.3) in which a smoothing step is used to filter an unconstrained deformation field. This presents a natural entry point for using the bilateral filter in order to ensure that the enforced uniformity of motion for neighbouring locations does not extend beyond structural boundaries. Indeed, such a regularisation was used by [Papież et al., 2014] in a demons-like registration scheme aimed at lung MRI inhale/exhale datasets, where it helped to enforce a uniform sliding motion of the lungs while marking the immediately adjacent ribs as static. I also note that a related family of medical image registration algorithms, exemplified by [Heinrich et al., 2014c], use a smoothing step to regularise a map of costs associated with translating the floating image by a specific displacement; here, bilateral filter could also be employed in a natural manner. Indeed, this was already done in the machine vision field in [Hosni et al., 2013]. Finally, an interesting direction to investigate could be a unified scheme where the proposed bilateral adaptive weighted similarity measure is used alongside the bilateral regularisation step.

The clinical evaluation on ATLR and glioma datasets was complicated by the fact that I found it challenging to reliably identify homologous pairs of anatomical landmarks between preoperative MRI and intraoperative iMRI structural scans, in the vicinity of resection. The reasons for this were severe deformation, appearance changes and linear/planar nature of the landmarks. I believe that the lack of veritable ground truth for evaluation of brain shift estimation algorithms is a significant problem in the field of brain shift estimation and I return to this point in the final discussion chapter in Section 5.1.

Chapter 5

Discussion

In my thesis I focussed on the question of whether it possible to improve the accuracy of iMRI based image guidance in the vicinity of the resection boundary where the resection cavity itself poses challenges to iMRI image acquisition and registration. My main motivation was that information about the structures within a few cm from the resection boundary is the most relevant to the outcome of surgery. I investigated the constituent steps of an interventional image processing pipeline which had been established in a clinical experimental study described in [Winston et al., 2014]. These are essentially two constituent steps: firstly, a correction of image artefacts, in particular the susceptibility artefact in diffusion MRI, which was described in [Daga et al., 2014], and a non-rigid registration step for the estimation of intraoperative brain shift, which was described in [Daga et al., 2012]. I identified specific limitations of these methods and attempted to address them.

In Chapter 3 I focussed on the correction of the susceptibility artefact in EPI that manifests itself near the air-tissue interface of the resection boundary. The EPI modalities affected by the artefact, namely diffusion MRI and fMRI, are potentially key to the accuracy of neuronavigation in the vicinity of resection. The work presented in [Daga et al., 2014] introduced a fast algorithm for correcting the susceptibility artefact, based on acquisition of field maps and their unwrapping. Their algorithm is implemented as an MRF optimization problem and is relatively robust to noise

CHAPTER 5. DISCUSSION

in raw phase maps but the phase-unwrapped maps tend to be estimated with reduced confidence near the resection boundary. For this reason, I investigated whether it is possible to simulate field maps from structural scans. The resulting simulated field maps were in close agreement with the acquired ones and enabled the correction of the artefact. On the other hand, the simulated maps did not facilitate a correction statistically significantly more accurate than the correction due to the acquired maps. This was mostly due to strong dependence on accuracy of air-tissue segmentation, which could potentially be improved with a methodological refinement. Further, the simulations took about one hour per case, which is too slow for intraoperative use, although this problem could likely be overcome by GPU acceleration. In conclusion, the simulation approach is a feasible one and has a real potential to be useful for the interventional use-case when perfected.

I envisage that simulated field maps could serve as patient-specific voxel-based spatial priors to inform phase-unwrapping in areas of low-confidence. This would enable the acquired field maps, which were used in a clinical setting as per [Winston et al., 2014], to be retained as the primary data source. The limitations pertaining to accuracy were mostly due to the intraoperative nature of the images and the usefuless of the method for non-interventional datasets could be studied as the susceptibility artefact has been ignored in many studies that relied on EPI imaging [Glasser et al., 2013]. A correction approach that is becoming popular is blip-up blip-down EPI imaging due to [Andersson et al., 2003], which is essentially a registration-based approach as it relies on recovering a mid-step between a pair of images that are affected by the same artefact in mutually opposite directions. In areas of homogeneous intensity the mid-point may be challenging to identify uniquely, which makes the problem ill-posed. For this reason, a meaningful regularisation is necessary and a spatial prior based on the simulated field map could potentially be well-suited for this application. This may be true especially for fMRI datasets which are naturally sparse in the anatomical space. Finally, recent new developments in fast MRI acquisition have opened a path for replacing EPI with pulse sequences less sensitive to B_0 field inhomogeneities [Ben-Eliezer and Frydman, 2011] but imaging protocols in clinical use will not change rapidly and there will likely remain a strong use-case for the presented approach for the foreseeable future.

In Chapter 4 I explored the non-rigid registration step for brain shift estimation. The previous

CHAPTER 5. DISCUSSION

work in this topic of greatest relevance is the registration scheme presented in [Daga et al., 2012] which provides multi-channel deformable registration with GPU acceleration which enables the registration to run within the time-constraints of the intraoperative iMRI imaging workflow. I considered the said scheme for a new iMRI dataset of 8 cases of glioma resection surgery performed at NHNN. I noted that the scheme tended to get stuck in local optima due to its reliance on gradientbased optimisation strategy, and also due to its use of a global image similarity measure which is not ideally suited for non-rigid registration. In order to overcome the said limitations, I proposed a registration scheme which combines a discrete registration, which seeks the global optimum of deformation parameters, with a local similarity measure equipped with an edge-preservation property based on bilateral adaptive weighting. The proposed adaptive weighting improved registration accuracy near tissue and resection boundaries in a number of experiments including an evaluation on NHNN ATLR and glioma iMRI datasets.

The proposed scheme had several limitations, and in particular, the edge-awareness property of the similarity measure introduces a speed bottleneck that renders the scheme too slow for intraoperative application, although powerful options are open for achieving a significant speedup (Section 4.5). The proposed scheme also only supports uni-modal and single-channel volumes and therefore is primarily intended to be used as an initialisation of a subsequent multi-channel registration step. Additionally, a multi-channel extension is possible (Section 4.5).

To summarise, I found that it is possible to improve the accuracy of surgical guidance in the area near of the resection boundary. I believe the presented methods are worth of being developed further so as to enable their experimental verification with the long-term goal of clinical translation. The main limitation of the presented method is its long running time. However, I outlined pathways to accelerating its implementation.

5.1 Open Research Problems

In the following I will attempt to discuss wider open research problems in the field of iMRI guidance for neurosurgery. It is clear that this is very much a new field and there exist multiple research directions. Propagation of surgical plans based on registration remains the most promising method to intraoperative update of neuronavigation, for two reasons: surgical planning based on preoperative imaging is an established and active area thereby enabling meaningful delination of critical areas; at the same time, there have been promising results in fast and relatively accurate non-rigid registration as exemplified by the work of [Daga et al., 2012] and [Heinrich et al., 2013]. However, the fact that most severe brain shift arises within several cm from resection boundary, which is also most relevant to outcome of the surgery, remains a major challenge. For instance, the registration presented in Chapter 4 achieves a convincing degree of overlap of the cortical gyri and sulci near the resection cavity but there the alignment becomes imperfect nearer the resection, which signifies inaccurately recovered deformation. There are several ways in which registration for brain shift estimation could be refined. These are briefly discussed as follows.

Firstly, conventional local similarity measures (Section 2.3.2) assume locally affine geometric transformation and fail to capture local similarity in severely deformed regions. I propose to use patch-based feature descriptors tailored to locally non-rigid transformation. Traditionally, such descriptors were restricted to feature-based registration as their gradient with respect to the transformation parameters is unknown. Recently, discrete registration, which does not require the gradient, expanded their use to intensity-based schemes as a similarity value between reference and floating image patches can be calculated as the norm of the difference of their normalised descriptors. An example hand-crafted descriptor that is robust to local deformations and contrast changes has been presented in the field of machine vision [Simo-Serra et al., 2015]. Recently, great interest has arisen in "deep learning" approaches that involve the training of convolutional neural networks which are apt at capturing high-dimensional features with very high sensitivity and specificity and which are being adopted vigorously by the medical imaging community [Greenspan et al., 2016, Litjens et al., 2017]. A most fruitful path to learning robust local image descriptors is likely to follow this route.

Secondly, the complex deformations inherent in brain shift present a challenge to gradientbased registration schemes which only search for local optima. Discrete optimisation methods search for an approximation to the global optimum and, as corroborated by Chapter 4, have become applicable within the time constraints of neurosurgery with state-of-the-art accuracy [Heinrich et al., 2013]. Discrete registration methods necessarily employ heuristics in order to make the solution computationally tractable. For instance, [Glocker et al., 2008a] optimised an MRF using a full set of node edges but used a sparse set of possible labels. In order to mitigate the latter limitation, they performed the discrete registration step repeatedly and composed the intermediate estimated transformations. However, by doing so they in effect abandoned the search for the global optimum. More recently, [Heinrich et al., 2013, Heinrich et al., 2016] optimised an MRF using a tree subset of node edges, which does not provide a solution to the full problem. Potentially a better approximation to the global optimum could be found by a method that would perform an optimisation of the full set of edges while allowing a large set of labels. It would be interesting to explore whether avenues toward such methodology have been explored in the wider literature perhaps borrowing from the optical flow field [Besse et al., 2014].

Thirdly, missing correspondences due to resection cannot be decoupled from the registration near the resection. In Section 2.4.1 I outlined previous work on the joint registration and valid/missing correspondence labelling problem. This should be pursued further, potentially within a discrete optimisation framework similar to [Parisot et al., 2012]. A point of note is that the valid/missing correspondence relationship is symmetrical such that resected volume (in the intraoperative space) has corresponding would-be resected tissue (in the preoperative space) whereby both do not have valid correspondences in each other image. The detection of volume that is folded in the intraoperative space due to partial collapse of ventricles could also be considered as a labelling problem involving a cost for opening a contiguous folding region similarly to strain thresholding in [Risholm et al., 2009].

Finally with regards to registration, most registration algorithms only return a point estimate of the deformation, which ignores the ill-posed nature of registration. The implication in the context of image guidance is that the surgeon cannot make a fully informed decision about whether to proceed with tissue resection. Uncertainty quantification for registration is an active field as outlined in Section 2.4.3. The time constraints of neurosurgery restrict the methods to approximate inference. Here, related work is that of [Simpson et al., 2012] and [Heinrich et al., 2016] whereby the latter work presents a method to obtain a dense uncertainty estimate within several minutes of running time.

The above refinements to registration are meaningful research problems. However, a practical question remains about how to actually validate results of registration. While iMRI has been regarded as the gold standard imaging modality to visualise brain shift since it became available [Nimsky et al., 2000], validating results of brain shift estimation is not trivial. As observed in Sections 4.4.3 and 4.4.3, an accurate annotation of preoperative MRI and intraoperative iMRI structural scans is in fact complicated by a lack of reliably *identifiable* homologous pairs of anatomical landmarks in the vicinity of the resection due to severe deformation, appearance changes, and the linear/planar nature of the landmarks. A way to mitigate this problem would be to prepare a physical phantom simulating brain shift. [Reinertsen and Collins, 2006] manufactured a phantom from polyvinyl alcohol cryogel (PVA-C) formed inside a brain surface mold, with a membrane which could be variably filled with water via a catheter to introduce deformation. MRI-compatible spherical markers were embedded in the phantom. The authors imaged the phantom and released the MRI volumes into the public domain. I assessed whether the data could be used for additional validation on the registration study in Chapter 4 but rejected them due to the following reason. I considered it necessary to inpaint the markers in order to prevent introducing a bias into the registration. At the same time, I found only some of the images had enough contrast to be suitable for registration. However, the same volumes had hyper-intense halos around the markers, which prevented inpainting in a veritable manner. This was unfortunate, as the phantom otherwise appears ideally-suited to verifying any brain shift estimation algorithm. I believe it would be useful to prepare a similar phantom and image it with a better MRI protocol. Alternatively, a realistic phantom could be obtained by ex-vivo MRI imaging of an animal (porcine) brain, which to the best of my knowledge has not been done before. It is debatable how MRI-visible markers could be embedded in the tissue such that they remain fixed in place. Potentially markers based on glass microspheres could be considered in this context, as [Li et al., 2005] imaged chicken breast tissue using MRI with such markers inserted via a cannula; this could potentially be done under ultrasound guidance for deeper structures.

Prior to concluding this discussion, it remains to comment on the obvious value of the existing iMRI datasets. The Spencer ATLR procedure performed at NHNN appears to be quite repetitive in terms of the location of the resection target and also in terms of the steps of the surgical procedure.

CHAPTER 5. DISCUSSION

These consist, following the opening of the dura, of tracing the floor of middle cranial fossa via the collateral sulcus before entering the lateral ventricle and making a cut through the temporal stem to the amygdala [Winston et al., 2014]. In the surgeries involved in the cited study, the results of the registration of the intraoperative iMRI data taken at a time-point immediately prior to the commencement of resection in deeper brain areas were used to make an informed decision on the subsequent path of resection. Conversely, due to the repetitive nature of the procedure up to this time-point, these iMRI data are representative of brain shift that arose before a point of intervention. It may be valuable to investigate whether a statistical model of the involved deformations can be built, perhaps using a group-wise average of the preoperative MRI images as a normalised image space. If the fitted model proved to report acceptably low uncertainties in the extended neighbourhood of the resection target, it could serve as a voxel-wise deformation prior. One possible application of such a prior would be to enable an assessment, given a preoperative MRI taken at the surgical planning stage that can be registered to the average space, of possible future shifts of key structures (e.g. the ORT tractography) as they could be expected to arise by the particular intraoperative time-point. Another possible application of the deformation prior could be to improve the accuracy of registration of intraoperative iMRI data. Looking beyond iMRI, an example of an interventional modality whose usefulness could likely be considerably improved by using a representative deformation prior is iUS (Section 2.1), which enables the imaging of deep-set structures at low cost and with little disruption to neurosurgery. A key disadvantage of iUS is that the acquired images only contain sparsely distributed features, which introduces uncertainty into iUS based guidance (Section 2.3.2). In this context, the incorporation of a deformation prior based on iMRI datasets into the registration of a iUS image could potentially reduce the uncertainty in the areas with few features. This could indirectly extend the benefits of iMRI to situations where its use is deemed too disruptive to the flow of surgery. Last but not least, since the high procurement and operational costs of iMRI limit its adoption to a small number of clinics, a publication in the future of iMRI-based deformation priors for certain repetitive and frequent procedures such as ATLR might be welcomed by facilities that embrace iUS as an affordable alternative to iMRI.

To conclude, in this thesis I presented techniques that have a potential to improve the accuracy of iMRI based image guidance for neurosurgery. In Chapter 3 I demonstrated that it is possible to correct DW-MRI and fMRI artefacts, which greatly expand the multimodal capabilities of iMRI imaging, in a way that is in principle robust to iMRI acquisition noise. In Chapter 4 I demonstrated that choosing a suitable local similarity measure and endowing it with an edge-awareness property leads to improved registration and guidance accuracy near the resection cavity. In the present discussion, I identified ways in which these findings can be pursued further to translate them into the clinic. Image guidance for neurosurgery remains an open area of research. Therefore, I identified a number of relevant open problems which, I believe, can be pursued by a new PhD student or other researchers, and which could lead to concrete advancements in this field.

Bibliography

- [Alvarez et al., 2007] Alvarez, L., Deriche, R., Papadopoulo, T., and Sánchez, J. (2007). Symmetrical dense optical flow estimation with occlusions detection. *International Journal of Computer Vision*, 75(3):371–385.
- [Andersson et al., 2003] Andersson, J. L., Skare, S., and Ashburner, J. (2003). How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. *NeuroImage*, 20(2):870–888.
- [Arsigny et al., 2006] Arsigny, V., Commowick, O., Pennec, X., and Ayache, N. (2006). A logeuclidean framework for statistics on diffeomorphisms. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 924–931. Springer.
- [Ashburner, 2007] Ashburner, J. (2007). A fast diffeomorphic image registration algorithm. NeuroImage, 38(1):95–113.
- [Ashburner et al., 1999] Ashburner, J., Friston, K. J., et al. (1999). Nonlinear spatial normalization using basis functions. *Human Brain Mapping*, 7(4):254–266.
- [Ashburner and Ridgway, 2013] Ashburner, J. and Ridgway, G. R. (2013). Symmetric diffeomorphic modeling of longitudinal structural MRI. Frontiers in Neuroscience, 6:197.
- [Avants et al., 2008] Avants, B. B., Epstein, C. L., Grossman, M., and Gee, J. C. (2008). Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Medical Image Analysis*, 12(1):26–41.

- [Bardinet et al., 1996] Bardinet, E., Cohen, L. D., and Ayache, N. (1996). Tracking and motion analysis of the left ventricle with deformable superquadrics. *Medical Image Analysis*, 1(2):129– 149.
- [Basser et al., 1994] Basser, P. J., Mattiello, J., and LeBihan, D. (1994). MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66(1):259–267.
- [Beg et al., 2005] Beg, M. F., Miller, M. I., Trouvé, A., and Younes, L. (2005). Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International Journal of Computer Vision*, 61(2):139–157.
- [Ben-Eliezer and Frydman, 2011] Ben-Eliezer, N. and Frydman, L. (2011). Spatiotemporal encoding as a robust basis for fast three-dimensional in vivo MRI. *NMR in Biomedicine*, 24(10):1191– 1201.
- [Besse et al., 2014] Besse, F., Rother, C., Fitzgibbon, A., and Kautz, J. (2014). PMPB: Patch-Match belief propagation for correspondence field estimation. *International Journal of Computer* Vision, 110(1):2–13.
- [Black et al., 1997] Black, P. M., Moriarty, T., Alexander, E., Stieg, P., Woodard, E. J., Gleason, P. L., Martin, C. H., Kikinis, R., Schwartz, R. B., and Jolesz, F. A. (1997). Development and implementation of intraoperative magnetic resonance imaging and its neurosurgical applications. *Neurosurgery*, 41(4):831–845.
- [Brett et al., 2001] Brett, M., Leff, A. P., Rorden, C., and Ashburner, J. (2001). Spatial normalization of brain images with focal lesions using cost function masking. *NeuroImage*, 14(2):486–500.
- [Burgos et al., 2013] Burgos, N., Cardoso, M. J., Modat, M., Pedemonte, S., Dickson, J., Barnes, A., Duncan, J. S., Atkinson, D., Arridge, S. R., Hutton, B. F., and Ourselin, S. (2013). Attenuation correction synthesis for hybrid PET-MR scanners. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 147–154. Springer.
- [Burgos et al., 2014] Burgos, N., Cardoso, M. J., Thielemans, K., Modat, M., Pedemonte, S., Dickson, J., Barnes, A., Ahmed, R., Mahoney, C. J., Schott, J. M., et al. (2014). Attenu-

ation correction synthesis for hybrid PET-MR scanners: application to brain studies. *IEEE Transactions on Medical Imaging*, 33(12):2332–2341.

- [Cachier et al., 2003] Cachier, P., Bardinet, E., Dormont, D., Pennec, X., and Ayache, N. (2003). Iconic feature based nonrigid registration: the PASHA algorithm. *Computer Vision and Image Understanding*, 89(2):272–298.
- [Cahill et al., 2007] Cahill, N. D., Noble, J. A., and Hawkes, D. J. (2007). Fourier methods for nonparametric image registration. In *Computer Vision and Pattern Recognition*, pages 1–8. IEEE.
- [Cardoso et al., 2009] Cardoso, M. J., Clarkson, M. J., Ridgway, G. R., Modat, M., Fox, N. C., and Ourselin, S. (2009). Improved maximum a posteriori cortical segmentation by iterative relaxation of priors. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 441–449. Springer.
- [Chen et al., 2015] Chen, K., Derksen, A., Heldmann, S., Hallmann, M., and Berkels, B. (2015). Deformable image registration with automatic non-correspondence detection. In Scale Space and Variational Methods in Computer Vision, pages 360–371. Springer.
- [Chen et al., 2012] Chen, S. J.-S., Hellier, P., Marchal, M., Gauvrit, J.-Y., Carpentier, R., Morandi, X., and Collins, D. L. (2012). An anthropomorphic polyvinyl alcohol brain phantom based on Colin27 for use in multimodal imaging. *Medical Physics*, 39(1):554–561.
- [Chitphakdithai and Duncan, 2010] Chitphakdithai, N. and Duncan, J. S. (2010). Non-rigid registration with missing correspondences in preoperative and postresection brain images. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 367–374. Springer.
- [Chitphakdithai et al., 2011] Chitphakdithai, N., Vives, K. P., and Duncan, J. S. (2011). Registration of brain resection MRI with intensity and location priors. In *IEEE International Symposium* on Biomedical Imaging: From Nano to Macro, pages 1520–1523. IEEE.
- [Choi and Lee, 2000] Choi, Y. and Lee, S. (2000). Injectivity conditions of 2D and 3D uniform cubic B-spline functions. *Graphical Models*, 62(6):411–427.

- [Christensen et al., 1996] Christensen, G. E., Rabbitt, R. D., and Miller, M. I. (1996). Deformable templates using large deformation kinematics. *IEEE Transactions on Image Processing*, 5(10):1435–1447.
- [Clare et al., 2006] Clare, S., Evans, J., and Jezzard, P. (2006). Requirements for room temperature shimming of the human brain. *Magnetic Resonance in Medicine*, 55(1):210–214.
- [Collignon et al., 1995] Collignon, A., Maes, F., Delaere, D., Vandermeulen, D., Suetens, P., and Marchal, G. (1995). Automated multi-modality image registration based on information theory. In International Conference on Information Processing in Medical Imaging, volume 3, pages 263–274.
- [Comeau et al., 2000] Comeau, R. M., Sadikot, A. F., Fenster, A., and Peters, T. M. (2000). Intraoperative ultrasound for guidance and tissue shift correction in image-guided neurosurgery. *Medical Physics*, 27(4):787–800.
- [Daga et al., 2014] Daga, P., Pendse, T., Modat, M., White, M., Mancini, L., Winston, G. P., McEvoy, A. W., Thornton, J., Yousry, T., Drobnjak, I., et al. (2014). Susceptibility artefact correction using dynamic graph cuts: application to neurosurgery. *Medical Image Analysis*, 18(7):1132–1142.
- [Daga et al., 2012] 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. J., Duncan, J. S., and Ourselin, S. (2012). Accurate localization of optic radiation during neurosurgery in an interventional MRI suite. *IEEE Transactions on Medical Imaging*, 31(4):882– 891.
- [Dempster et al., 1977] Dempster, A., Laird, N., and Rubin, D. (1977). Maximum likelihood from incomplete data via the EM algorithm. Journal of the Royal Statistical Society. Series B (Methodological), 39(1):1–38.
- [Ding et al., 2011] Ding, S., Miga, M. I., Pheiffer, T. S., Simpson, A. L., Thompson, R. C., and Dawant, B. M. (2011). Tracking of vessels in intra-operative microscope video sequences for

cortical displacement estimation. *IEEE Transactions on Biomedical Engineering*, 58(7):1985–1993.

- [Dolecek et al., 2012] Dolecek, T. A., Propp, J. M., Stroup, N. E., and Kruchko, C. (2012). CB-TRUS statistical report: primary brain and central nervous system tumors diagnosed in the United States in 2005–2009. *Neuro-oncology*, 15:646–647.
- [Dumpuri et al., 2007] Dumpuri, P., Thompson, R. C., Dawant, B. M., Cao, A., and Miga, M. I. (2007). An atlas-based method to compensate for brain shift: preliminary results. *Medical Image Analysis*, 11(2):128–145.
- [Duncan et al., 2016] Duncan, J. S., Winston, G. P., Koepp, M. J., and Ourselin, S. (2016). Brain imaging in the assessment for epilepsy surgery. *The Lancet Neurology*, 15(4):420–433.
- [Essayed et al., 2017] Essayed, W. I., Zhang, F., Unadkat, P., Cosgrove, G. R., Golby, A. J., and O'Donnell, L. J. (2017). White matter tractography for neurosurgical planning: A topographybased review of the current state of the art. *NeuroImage: Clinical*, 15:659–672.
- [Ferrant et al., 2001] Ferrant, M., Nabavi, A., Macq, B., Jolesz, F. A., Kikinis, R., and Warfield, S. K. (2001). Registration of 3-D intraoperative MR images of the brain using a finite-element biomechanical model. *IEEE Transactions on Medical Imaging*, 20(12):1384–1397.
- [Fortun et al., 2015] Fortun, D., Bouthemy, P., and Kervrann, C. (2015). Optical flow modeling and computation: a survey. *Computer Vision and Image Understanding*, 134:1–21.
- [Gerard et al., 2017] Gerard, I. J., Kersten-Oertel, M., Petrecca, K., Sirhan, D., Hall, J. A., and Collins, D. L. (2017). Brain shift in neuronavigation of brain tumors: A review. *Medical Image Analysis*, 35:403–420.
- [Glasser et al., 2013] Glasser, M. F., Sotiropoulos, S. N., Wilson, J. A., Coalson, T. S., Fischl, B., Andersson, J. L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J. R., et al. (2013). The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage*, 80:105–124.
- [Glocker et al., 2008a] Glocker, B., Komodakis, N., Tziritas, G., Navab, N., and Paragios, N. (2008a). Dense image registration through MRFs and efficient linear programming. *Medical Image Analysis*, 12(6):731–741.

- [Glocker et al., 2008b] Glocker, B., Paragios, N., Komodakis, N., Tziritas, G., and Navab, N. (2008b). Optical flow estimation with uncertainties through dynamic MRFs. In *Computer Vision and Pattern Recognition*, 2008. CVPR 2008. IEEE Conference on, pages 1–8. IEEE.
- [Glover and Pelc, 1986] Glover, G. H. and Pelc, N. J. (1986). Method for correcting image distortion due to gradient nonuniformity. US Patent 4,591,789.
- [Greenspan et al., 2016] Greenspan, H., van Ginneken, B., and Summers, R. M. (2016). Guest editorial deep learning in medical imaging: Overview and future promise of an exciting new technique. *IEEE Transactions on Medical Imaging*, 35(5):1153–1159.
- [Gruetter and Boesch, 1992] Gruetter, R. and Boesch, C. (1992). Fast, noniterative shimming of spatially localized signals. In vivo analysis of the nagnetic field along axes. *Journal of Magnetic Resonance*, 96(2):323–334.
- [Gumprecht et al., 1999] Gumprecht, H. K., Widenka, D. C., and Lumenta, C. B. (1999). Brain-Lab VectorVision Neuronavigation System: technology and clinical experiences in 131 cases. *Neurosurgery*, 44(1):97–104.
- [Hall and Truwit, 2008] Hall, W. A. and Truwit, C. L. (2008). Intraoperative MR-guided neurosurgery. Journal of Magnetic Resonance Imaging, 27(2):368–375.
- [He et al., 2013] He, K., Sun, J., and Tang, X. (2013). Guided image filtering. IEEE Transactions on Pattern Analysis and Machine Intelligence, 35(6):1397–1409.
- [Heinrich et al., 2013] Heinrich, H., Jenkinson, M., Brady, M., and Schnabel, J. (2013). MRFbased deformable registration and ventilation estimation of lung CT. *IEEE Transactions on Medical Imaging*, 32(7):1239–1248.
- [Heinrich et al., 2014a] Heinrich, M. P., Papież, B. W., Schnabel, J. A., and Handels, H. (2014a). Multispectral image registration based on local canonical correlation analysis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 202–209. Springer.
- [Heinrich et al., 2014b] Heinrich, M. P., Papież, B. W., Schnabel, J. A., and Handels, H. (2014b). Multispectral image registration based on local canonical correlation analysis. In Golland, P.,

Hata, N., Barillot, C., Hornegger, J., and Howe, R., editors, International Conference on Medical Image Computing and Computer-Assisted Intervention, volume 8673 of LNCS, pages 202–209. Springer, Heidelberg.

- [Heinrich et al., 2014c] Heinrich, M. P., Papież, B. W., Schnabel, J. A., and Handels, H. (2014c). Non-parametric discrete registration with convex optimisation. In *International Workshop on Biomedical Image Registration*, pages 51–61. Springer.
- [Heinrich et al., 2016] Heinrich, M. P., Simpson, I. J., Papież, B. W., Brady, M., and Schnabel, J. A. (2016). Deformable image registration by combining uncertainty estimates from supervoxel belief propagation. *Medical Image Analysis*, 27:57–71.
- [Hill et al., 1998] Hill, D. L., Maurer, C. R., Maciunas, R. J., Maciunas, R. J., Barwise, J. A., Fitzpatrick, J. M., and Wang, M. Y. (1998). Measurement of intraoperative brain surface deformation under a craniotomy. *Neurosurgery*, 43(3):514–526.
- [Horn and Schunck, 1981] Horn, B. K. and Schunck, B. G. (1981). Determining optical flow. Artificial Intelligence, 17(1-3):185–203.
- [Hosni et al., 2013] Hosni, A., Rhemann, C., Bleyer, M., Rother, C., and Gelautz, M. (2013). Fast cost-volume filtering for visual correspondence and beyond. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(2):504–511.
- [Hutton et al., 2002] Hutton, C., Bork, A., Josephs, O., Deichmann, R., Ashburner, J., and Turner, R. (2002). Image distortion correction in fMRI: a quantitative evaluation. *NeuroImage*, 16(1):217–240.
- [Ivan et al., 2014] Ivan, M. E., Yarlagadda, J., Saxena, A. P., Martin, A. J., Starr, P. A., Sootsman, W. K., and Larson, P. S. (2014). Brain shift during bur hole-based procedures using interventional MRI. *Journal of Neurosurgery*, 121(1):149–160.
- [Janke et al., 2004] Janke, A., Zhao, H., Cowin, G. J., Galloway, G. J., and Doddrell, D. M. (2004). Use of spherical harmonic deconvolution methods to compensate for nonlinear gradient effects on MRI images. *Magnetic Resonance in Medicine*, 52(1):115–122.

- [Jenkinson, 2003] Jenkinson, M. (2003). Fast, automated, N-dimensional phase-unwrapping algorithm. Magnetic Resonance in Medicine, 49(1):193–197.
- [Jenkinson et al., 2004] Jenkinson, M., Wilson, J. L., and Jezzard, P. (2004). Perturbation method for magnetic field calculations of nonconductive objects. *Magnetic Resonance in Medicine*, 52(3):471–477.
- [Jezzard and Balaban, 1995] Jezzard, P. and Balaban, R. S. (1995). Correction for geometric distortion in echo planar images from B0 field variations. *Magnetic Resonance in Medicine*, 34(1):65–73.
- [Ji et al., 2014] Ji, S., Fan, X., Roberts, D. W., Hartov, A., and Paulsen, K. D. (2014). Cortical surface shift estimation using stereovision and optical flow motion tracking via projection image registration. *Medical image analysis*, 18(7):1169–1183.
- [Ji et al., 2009] Ji, S., Hartov, A., Roberts, D., and Paulsen, K. (2009). Data assimilation using a gradient descent method for estimation of intraoperative brain deformation. *Medical Image Analysis*, 13(5):744–756.
- [Jolesz, 2011] Jolesz, F. A. (2011). Intraoperative Imaging in Neurosurgery: Where Will the Future Take Us?, pages 21–25. Springer Vienna, Vienna.
- [Jones, 2004] Jones, D. K. (2004). The effect of gradient sampling schemes on measures derived from diffusion tensor MRI: a Monte Carlo study. *Magnetic Resonance in Medicine*, 51(4):807– 815.
- [Kelly et al., 1986] Kelly, P. J., Kall, B. A., Goerss, S., and Earnest IV, F. (1986). Computerassisted stereotaxic laser resection of intra-axial brain neoplasms. *Journal of Neurosurgery*, 64(3):427–439.
- [Kim et al., 2006] 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., and Kim, S. I. (2006). How does distortion correction correlate with anisotropic indices? A diffusion tensor imaging study. *Magnetic Resonance Imaging*, 24(10):1369–1376.

- [Kochan et al., 2014] Kochan, M., Daga, P., Burgos, N., White, M., Cardoso, M. J., Mancini, L., Winston, G. P., McEvoy, A. W., Thornton, J., Yousry, T., Duncan, J. S., Stoyanov, D., and Ourselin, S. (2014). Simulated field maps: toward improved susceptibility artefact correction in interventional MRI. In *International Conference on Information Processing in Computer-Assisted Interventions*, pages 226–235. Springer.
- [Kochan et al., 2015] Kochan, M., Daga, P., Burgos, N., White, M., Cardoso, M. J., Mancini, L., Winston, G. P., McEvoy, A. W., Thornton, J., Yousry, T., Duncan, J. S., Stoyanov, D., and Ourselin, S. (2015). Simulated field maps for susceptibility artefact correction in interventional MRI. International Journal of Computer Assisted Radiology and Surgery, 10(9):1405–1416.
- [Kochan et al., 2016] Kochan, M., Modat, M., Vercauteren, T., White, M., Mancini, L., Winston, G. P., McEvoy, A. W., Thornton, J. S., Yousry, T., Duncan, J. S., Ourselin, S., and Stoyanov, D. (2016). Bilateral weighted adaptive local similarity measure for registration in neurosurgery. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, volume 9902, page 81. Springer.
- [Kohli and Torr, 2008] Kohli, P. and Torr, P. H. (2008). Measuring uncertainty in graph cut solutions. Computer Vision and Image Understanding, 112(1):30–38.
- [Kuhnt et al., 2011] Kuhnt, D., Ganslandt, O., Schlaffer, S.-M., Buchfelder, M., and Nimsky, C. (2011). Quantification of glioma removal by intraoperative high-field magnetic resonance imaging: an update. *Neurosurgery*, 69(4):852–863.
- [Leemans and Jones, 2009] Leemans, A. and Jones, D. K. (2009). The B-matrix must be rotated when correcting for subject motion in DTI data. *Magnetic resonance in medicine*, 61(6):1336– 1349.
- [Li et al., 2005] Li, Y., Wang, J., Holloway, C., and Plewes, D. B. (2005). Development of an MRI/x-ray/ultrasound compatible marker for pre-operative breast tumour localization. *Physics* in Medicine and Biology, 50(14):3349.
- [Litjens et al., 2017] Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A. A., Ciompi, F., Ghafoorian, M., van der Laak, J. A., van Ginneken, B., and Sánchez, C. I. (2017). A survey on deep learning

in medical image analysis. arXiv preprint arXiv:1702.05747.

- [Lunn et al., 2005] Lunn, K. E., Paulsen, K. D., Lynch, D. R., Roberts, D. W., Kennedy, F. E., and Hartov, A. (2005). Assimilating intraoperative data with brain shift modeling using the adjoint equations. *Medical Image Analysis*, 9(3):281–293.
- [Maes et al., 1997] Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., and Suetens, P. (1997). Multimodality image registration by maximization of mutual information. *IEEE Trans*actions on Medical Imaging, 16(2):187–198.
- [Mansfield, 1977] Mansfield, P. (1977). Multi-planar image formation using NMR spin echoes. Journal of Physics C: Solid State Physics, 10(3):L55.
- [McRobbie et al., 2017] McRobbie, D. W., Moore, E. A., and Graves, M. J. (2017). MRI from Picture to Proton. Cambridge university press.
- [Mercier et al., 2012] Mercier, L., Del Maestro, R. F., Petrecca, K., Araujo, D., Haegelen, C., and Collins, D. L. (2012). Online database of clinical MR and ultrasound images of brain tumors. *Medical Physics*, 39(6):3253–3261.
- [Miga et al., 2001] Miga, M. I., Roberts, D. W., Kennedy, F. E., Platenik, L. A., Hartov, A., Lunn, K. E., and Paulsen, K. D. (2001). Modeling of retraction and resection for intraoperative updating of images. *Neurosurgery*, 49(1):75–85.
- [Miga et al., 2003] Miga, M. I., Sinha, T. K., Cash, D. M., Galloway, R. L., and Weil, R. J. (2003). Cortical surface registration for image-guided neurosurgery using laser-range scanning. *IEEE Transactions on Medical Imaging*, 22(8):973–985.
- [Miga et al., 2016] Miga, M. I., Sun, K., Chen, I., Clements, L. W., Pheiffer, T. S., Simpson, A. L., and Thompson, R. C. (2016). Clinical evaluation of a model-updated image-guidance approach to brain shift compensation: experience in 16 cases. *International journal of computer assisted* radiology and surgery, 11(8):1467–1474.
- [Mildenberger et al., 2002] Mildenberger, P., Eichelberg, M., and Martin, E. (2002). Introduction to the DICOM standard. *European radiology*, 12(4):920–927.

- [Modat et al., 2014] Modat, M., Cash, D. M., Daga, P., Winston, G. P., Duncan, J. S., and Ourselin, S. (2014). Global image registration using a symmetric block-matching approach. *Journal of Medical Imaging*, 1(2).
- [Modat et al., 2010] Modat, M., Ridgway, G. R., Taylor, Z. A., Lehmann, M., Barnes, J., Hawkes, D. J., Fox, N. C., and Ourselin, S. (2010). Fast free-form deformation using graphics processing units. *Computer Methods and Programs in Biomedicine*, 98(3):278–284.
- [Murphy et al., 2011] Murphy, K., van Ginneken, B., Reinhardt, J. M., Kabus, S., Ding, K., Deng, X., Cao, K., Du, K., Christensen, G. E., Garcia, V., Vercauteren, T., Ayache, N., Commowick, O., Malandain, G., Glocker, B., Paragios, N., Navab, N., Gorbunova, V., Sporring, J., de Bruijne, M., Han, X., Heinrich, M. P., Schnabel, J. A., Jenkinson, M., Lorenz, C., Modat, M., McClelland, J. R., Ourselin, S., Muenzing, S. E. A., Viergever, M. A., Nigris, D. D., Collins, D. L., Arbel, T., Peroni, M., Li, R., Sharp, G. C., Schmidt-Richberg, A., Ehrhardt, J., Werner, R., Smeets, D., Loeckx, D., Song, G., Tustison, N., Avants, B., Gee, J. C., Staring, M., Klein, S., Stoel, B. C., Urschler, M., Werlberger, M., Vandemeulebroucke, J., Rit, S., Sarrut, D., and Pluim, J. P. W. (2011). Evaluation of Registration Methods on Thoracic CT: The EMPIRE10 Challenge. *IEEE Transactions on Medical Imaging*, 30(11):1901–1920.
- [Nabavi et al., 2001] Nabavi, A., McL. Black, P., Gering, D. T., Westin, C.-F., Mehta, V., Pergolizzi Jr, R. S., Ferrant, M., Warfield, S. K., Hata, N., Schwartz, R. B., et al. (2001). Serial intraoperative magnetic resonance imaging of brain shift. *Neurosurgery*, 48(4):787–798.
- [Nimsky et al., 2000] Nimsky, C., Ganslandt, O., Cerny, S., Hastreiter, P., Greiner, G., and Fahlbusch, R. (2000). Quantification of, visualization of, and compensation for brain shift using intraoperative magnetic resonance imaging. *Neurosurgery*, 47(5):1070–1080.
- [Nimsky et al., 2001] Nimsky, C., Ganslandt, O., Hastreiter, P., and Fahlbusch, R. (2001). Intraoperative compensation for brain shift. *Surgical Neurology*, 56(6):357–364.
- [Nimsky et al., 2005] Nimsky, C., Ganslandt, O., Hastreiter, P., Wang, R., Benner, T., Sorensen, A. G., and Fahlbusch, R. (2005). Preoperative and intraoperative diffusion tensor imaging-based fiber tracking in glioma surgery. *Neurosurgery*, 56(1):130–138.

- [Omuro and DeAngelis, 2013] Omuro, A. and DeAngelis, L. M. (2013). Glioblastoma and other malignant gliomas: a clinical review. Journal of the American Medical Association, 310(17):1842–1850.
- [Ou et al., 2011] Ou, Y., Sotiras, A., Paragios, N., and Davatzikos, C. (2011). DRAMMS: Deformable registration via attribute matching and mutual-saliency weighting. *Medical Image Analysis*, 15(4):622–639.
- [Ourselin et al., 2000] Ourselin, S., Roche, A., Subsol, G., Pennec, X., and Ayache, N. (2000). Reconstructing a 3D structure from serial histological sections. *Image and Vision Computing*, 19:25–31.
- [Papadakis et al., 2002] Papadakis, N. G., Martin, K. M., Mustafa, M. H., Wilkinson, I. D., Griffiths, P. D., Huang, C. L.-H., and Woodruff, P. W. (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.
- [Papież et al., 2014] Papież, B. W., Heinrich, M. P., Fehrenbach, J., Risser, L., and Schnabel, J. A. (2014). An implicit sliding-motion preserving regularisation via bilateral filtering for deformable image registration. *Medical Image Analysis*, 18(8):1299–1311.
- [Parisot et al., 2012] Parisot, S., Duffau, H., Chemouny, S., and Paragios, N. (2012). Joint tumor segmentation and dense deformable registration of brain MR images. In *International Conference* on Medical Image Computing and Computer-Assisted Intervention, pages 651–658. Springer.
- [Paul et al., 2009] Paul, P., Morandi, X., and Jannin, P. (2009). A surface registration method for quantification of intraoperative brain deformations in image-guided neurosurgery. *IEEE Transactions on Information Technology in Biomedicine*, 13(6):976–983.
- [Paulsen et al., 1999] Paulsen, K. D., Miga, M. I., Kennedy, F. E., Hoopens, P., Hartov, A., and Roberts, D. W. (1999). A computational model for tracking subsurface tissue deformation during stereotactic neurosurgery. *IEEE Transactions on Biomedical Engineering*, 46(2):213–225.
- [Periaswamy and Farid, 2006] Periaswamy, S. and Farid, H. (2006). Medical image registration with partial data. *Medical Image Analysis*, 10(3):452–464.

- [Poynton et al., 2009] Poynton, C., Jenkinson, M., and Wells III, W. (2009). Atlas-based improved prediction of magnetic field inhomogeneity for distortion correction of EPI data. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 951–959. Springer.
- [Pujol et al., 2015] Pujol, S., Wells, W., Pierpaoli, C., Brun, C., Gee, J., Cheng, G., Vemuri, B., Commowick, O., Prima, S., Stamm, A., et al. (2015). The DTI challenge: toward standardized evaluation of diffusion tensor imaging tractography for neurosurgery. *Journal of Neuroimaging*, 25(6):875–882.
- [Reinertsen and Collins, 2006] Reinertsen, I. and Collins, D. (2006). A realistic phantom for brainshift simulations. *Medical Physics*, 33(9):3234–3240.
- [Reinertsen et al., 2007] Reinertsen, I., Descoteaux, M., Siddiqi, K., and Collins, D. L. (2007). Validation of vessel-based registration for correction of brain shift. *Medical image analysis*, 11(4):374–388.
- [Risholm et al., 2013] Risholm, P., Janoos, F., Norton, I., Golby, A. J., and Wells, W. M. (2013). Bayesian characterization of uncertainty in intra-subject non-rigid registration. *Medical Image Analysis*, 17(5):538–555.
- [Risholm et al., 2009] Risholm, P., Samset, E., Talos, I.-F., and Wells, W. (2009). A non-rigid registration framework that accommodates resection and retraction. In *International Conference* on Information Processing in Medical Imaging, pages 447–458. Springer Berlin Heidelberg.
- [Rivaz et al., 2015] Rivaz, H., Chen, S. J.-S., and Collins, D. L. (2015). Automatic deformable MRultrasound registration for image-guided neurosurgery. *IEEE Transactions on Medical Imaging*, 34(2):366–380.
- [Roberts et al., 1986] Roberts, D. W., Strohbehn, J. W., Hatch, J. F., Murray, W., and Kettenberger, H. (1986). A frameless stereotaxic integration of computerized tomographic imaging and the operating microscope. *Journal of Neurosurgery*, 65(4):545–549.
- [Rohlfing et al., 2003] Rohlfing, T., Maurer, C. R., Bluemke, D. A., and Jacobs, M. A. (2003). Volume-preserving nonrigid registration of MR breast images using free-form deformation with

an incompressibility constraint. IEEE Transactions on Medical Imaging, 22(6):730-741.

- [Rueckert et al., 2006] Rueckert, D., Aljabar, P., Heckemann, R. A., Hajnal, J. V., and Hammers,
 A. (2006). Diffeomorphic registration using B-splines. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 702–709. Springer.
- [Rueckert et al., 1999] Rueckert, D., Sonoda, L. I., Hayes, C., Hill, D. L., Leach, M. O., and Hawkes, D. J. (1999). Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Transactions on Medical Imaging*, 18(8):712–721.
- [Sanders and Kandrot, 2010] Sanders, J. and Kandrot, E. (2010). CUDA by example: an introduction to general-purpose GPU programming. Addison-Wesley Professional, 1st edition.
- [Sederberg and Parry, 1986] Sederberg, T. W. and Parry, S. R. (1986). Free-form deformation of solid geometric models. ACM SIGGRAPH Computer Graphics, 20(4):151–160.
- [Shamir et al., 2009] Shamir, R. R., Joskowicz, L., Spektor, S., and 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, 1948] Shannon, C. E. (1948). A mathematical theory of communication. The Bell System Technical Journal, 27(3):379–423.
- [Simo-Serra et al., 2015] Simo-Serra, E., Torras, C., and Moreno-Noguer, F. (2015). DaLI: deformation and light invariant descriptor. International Journal of Computer Vision, 115(2):136– 154.
- [Simpson et al., 2012] Simpson, I. J., Schnabel, J. A., Groves, A. R., Andersson, J. L., and Woolrich, M. W. (2012). Probabilistic inference of regularisation in non-rigid registration. *NeuroIm*age, 59(3):2438–2451.
- [Škrinjar et al., 2002] Škrinjar, O., Nabavi, A., and Duncan, J. (2002). Model-driven brain shift compensation. *Medical Image Analysis*, 6(4):361–373.
- [Smith et al., 2004] Smith, S. M., Jenkinson, M., Woolrich, M. W., Beckmann, C. F., Behrens, T. E., Johansen-Berg, H., Bannister, P. R., De Luca, M., Drobnjak, I., Flitney, D. E., Niazya, R. K., Saundersa, J., Vickersa, J., Zhanga, Y., De Stefano, N., Brady, J. M., and Matthews,

P. M. (2004). Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage*, 23(Supplement 1):S208–S219.

- [Sotiras et al., 2013] Sotiras, A., Davatzikos, C., and Paragios, N. (2013). Deformable medical image registration: a survey. *IEEE Transactions on Medical Imaging*, 32(7):1153–1190.
- [Stejskal and Tanner, 1965] Stejskal, E. O. and Tanner, J. E. (1965). Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient. *The Journal of Chemical Physics*, 42(1):288–292.
- [Studholme et al., 1999] Studholme, C., Hill, D. L., and Hawkes, D. J. (1999). An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognition*, 32(1):71–86.
- [Sun et al., 2005] Sun, H., Lunn, K. E., Farid, H., Wu, Z., Roberts, D. W., Hartov, A., and Paulsen, K. D. (2005). Stereopsis-guided brain shift compensation. *IEEE Transactions on Medical Imaging*, 24(8):1039–1052.
- [Thévenaz and Unser, 2000] Thévenaz, P. and Unser, M. (2000). Optimization of mutual information for multiresolution image registration. *IEEE Transactions on Image Processing*, 9(12):2083– 2099.
- [Thirion, 1998] Thirion, J.-P. (1998). Image matching as a diffusion process: an analogy with Maxwell's demons. *Medical Image Analysis*, 2(3):243–260.
- [Toews and Wells, 2013] Toews, M. and Wells, W. M. (2013). Efficient and robust model-to-image alignment using 3D scale-invariant features. *Medical Image Analysis*, 17(3):271–282.
- [Tomasi and Manduchi, 1998] Tomasi, C. and Manduchi, R. (1998). Bilateral filtering for gray and color images. In *International Conference on Computer Vision*, pages 839–846. IEEE.
- [Tuch et al., 2002] Tuch, D. S., Reese, T. G., Wiegell, M. R., Makris, N., Belliveau, J. W., and Wedeen, V. J. (2002). High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magnetic Resonance in Medicine*, 48(4):577–582.
- [Unsgaard et al., 2005] Unsgaard, G., Selbekk, T., Müller, T. B., Ommedal, S., Torp, S., Myhr, G., Bang, J., and Hernes, T. N. (2005). Ability of navigated 3D ultrasound to delineate gliomas

and metastases–comparison of image interpretations with histopathology. *Acta Neurochirurgica*, 147(12):1259–1269.

- [Viola and Wells III, 1997] Viola, P. and Wells III, W. M. (1997). Alignment by maximization of mutual information. International Journal of Computer Vision, 24(2):137–154.
- [Wassermann et al., 2014] Wassermann, D., Toews, M., Niethammer, M., and Wells III, W. (2014). Probabilistic diffeomorphic registration: representing uncertainty. In *International Workshop* on *Biomedical Image Registration*, pages 72–82. Springer.
- [Wein et al., 2008] Wein, W., Brunke, S., Khamene, A., Callstrom, M. R., and Navab, N. (2008). Automatic CT-ultrasound registration for diagnostic imaging and image-guided intervention. *Medical Image Analysis*, 12(5):577–585.
- [Wein et al., 2013] Wein, W., Ladikos, A., Fuerst, B., Shah, A., Sharma, K., and Navab, N. (2013). Global registration of ultrasound to mri using the LC2 metric for enabling neurosurgical guidance. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 34–41. Springer Berlin Heidelberg.
- [Winston et al., 2012] Winston, G. P., Daga, P., Stretton, J., Modat, M., Symms, M. R., McEvoy, A. W., Ourselin, S., and Duncan, J. S. (2012). Optic radiation tractography and vision in anterior temporal lobe resection. *Annals of Neurology*, 71(3):334–341.
- [Winston et al., 2014] Winston, G. P., Micallef, C., Symms, M. R., Alexander, D. C., Duncan, J. S., and Zhang, H. (2014). Advanced diffusion imaging sequences could aid assessing patients with focal cortical dysplasia and epilepsy. *Epilepsy Research*, 108(2):336–339.
- [Xu and Prince, 1998] Xu, C. and Prince, J. L. (1998). Snakes, shapes, and gradient vector flow. IEEE Transactions on Image Processing, 7(3):359–369.
- [Yoon and Kweon, 2006] Yoon, K.-J. and Kweon, I. S. (2006). Adaptive support-weight approach for correspondence search. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(4):650–656.