

**Quantitative Multi-modal Analysis of Pediatric
Focal Epilepsy**

by

Andy Khai Siang Eow

Submitted to the Department of Electrical Engineering and Computer
Science

in partial fulfillment of the requirements for the degree of

Master of Science in Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

[June 2005]

May 2005

© Massachusetts Institute of Technology 2005. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
March 15, 2005


Certified by
W. Eric L. Grimson
Bernard Gordon Professor of Medical Engineering
Thesis Supervisor

Certified by
Simon K. Warfield
Associate Professor of Radiology, Harvard Medical School
Thesis Supervisor

Accepted by
Arthur C. Smith
Chairman, Department Committee on Graduate Students



BARKER





Room 14-0551
77 Massachusetts Avenue
Cambridge, MA 02139
Ph: 617.253.2800
Email: docs@mit.edu
<http://libraries.mit.edu/docs>

DISCLAIMER OF QUALITY

Due to the condition of the original material, there are unavoidable flaws in this reproduction. We have made every effort possible to provide you with the best copy available. If you are dissatisfied with this product and find it unusable, please contact Document Services as soon as possible.

Thank you.

The images contained in this document are of the best quality available.



Quantitative Multi-modal Analysis of Pediatric Focal Epilepsy

by

Andy Khai Siang Eow

Submitted to the Department of Electrical Engineering and Computer Science
on March 15, 2005, in partial fulfillment of the
requirements for the degree of
Master of Science in Computer Science

Abstract

For patients with medically refractive focal epilepsy, surgical intervention to remove the epileptic foci is often the last alternative for permanent cure. The success of such surgery is highly dependent on the doctor's ability to accurately locate the epileptogenic region during the pre-surgical planning and evaluation phase. Hence the goal of this project is to provide an end-to-end quantitative analysis pipeline that fuses an array of imaging modalities including magnetic resonance imaging (MRI), diffusion tensor MRI, positron emission tomography (PET), single-photon emission computerized tomography (SPECT) as well as EEG data to build patient-specific head models and to compute prior probability maps of epileptic hotspots for more accurate EEG source localization. By improving the ability to accurately locate these epileptogenic seizure sources, patients can benefit tremendously from accurate surgical resection and consequently have a better chance for complete seizure free recovery.

Thesis Supervisor: W. Eric L. Grimson
Title: Bernard Gordon Professor of Medical Engineering

Thesis Supervisor: Simon K. Warfield
Title: Associate Professor of Radiology, Harvard Medical School

Acknowledgments

Firstly, I would like to thank my thesis advisor, Eric Grimson, for all his insight, enthusiasm and support in my research. Over the past year and a half at MIT, Eric has given me significant autonomy to pursue research that excites me the most, providing valuable guidance to help me better focus on the more important issues raised along the way. His jovial and energetic nature goes a long way in building such a nurturing and ideal research environment here in the medical vision group.

I would also like to express my gratitude to my co-advisor, Simon Warfield, at the Surgical Planning Lab (SPL) in the Brigham and Women's Hospital. This project was largely made possible by a new collaboration with the doctors at the Children's Hospital. Many thanks to Simon for introducing me to the key collaborators in Children's Hospital and for his tremendous help in the collection of clinical data. His keen insight on this project has also help shape it in more ways than one.

I would also like to thank my principle clinical collaborator, Dr. Richard Robertson, at the Children's Hospital for his insightful feedback. He has been extremely supportive throughout the entire course of this project, providing us with his clinical expertise on the matter. His professional evaluation of the clinical utility of the work presented in this thesis is invaluable to the success of the project. He is simply an amazing person to work with.

This project was also largely successful due to the efforts of my fellow research partners of the Epilepsy Surgical Planning group namely Tiferet Levine who is credited with the EEG source localization part of the project and Sylvain Jaume who was involved with data pre-processing in the earlier stages of the project. These people have been instrumental to the final results as presented in this thesis.

Through the course of my research, there are also many others from MIT and the SPL who have helped me understand and use the multitude of medical imaging tools available in the field. I would like to thank Lauren O'Donnell, Mahnaz Maddah and Raul San Jose Estepar for helping me get acquainted with the advanced features of the 3D-Slicer, the Visualization Toolkit and DT-MRI data processing. I would also

like to thank Neil Weisenfeld and Killian Pohl for their advice on intensity correction and tissue segmentation. This project was also mostly implemented with the algorithms from the National Library of Medicine Insight Segmentation and Registration Toolkit(ITK) and Luis Ibanez of Kitware Inc. provided much valuable assistance in this respect.

Most importantly, I want to thank my whole family for their constant encouragement throughout my undergraduate and graduate career in the United States. I would like to especially thank my father, Alex Eow, and my mother, Yoke Yin Lee, for all their love and support over the years. It has been a wonderful journey and I look forward to even greater challenges in the years to come.

Contents

1	Introduction	12
1.1	Focal Epilepsy	12
1.2	Epilepsy Surgical Planning Project	13
2	Background	16
2.1	Accurate EEG Source Localization	16
2.2	Incorporating a Patient-Specific Realistic Head Model	20
2.3	Epileptic Hotspots Prior Probability Map	21
2.3.1	Multi-modal Analysis of Focal Epilepsy	21
2.3.2	Asymmetry Detection Framework	23
2.4	Epilepsy Surgical Planning Pipeline	25
3	Method	28
3.1	Pre-processing	28
3.1.1	Rigid Volume Correspondence for Multi-modal Analysis	28
3.1.2	Modified Watershed Tissue Segmentation	30
3.2	Plane of Symmetry Detection	36
3.3	Asymmetry Analysis	41
3.4	Building the Prior Map	47
3.5	Pipeline Summary	48
4	Results & Discussion	49
4.1	Synthetic Data Test	49

4.1.1	Data Generation	49
4.1.2	Plane of Symmetry Detection	51
4.1.3	Asymmetry Analysis	52
4.2	Patient Example 1 - Focal Cortical Dysplasia in the Right Inferior Frontal Lobe	61
4.2.1	Multi-modal Rigid Volume Correspondence	61
4.2.2	Tissue Segmentation	62
4.2.3	Plane of Symmetry Detection	66
4.2.4	Asymmetry Analysis	67
4.2.5	Prior Map Generation	74
4.3	Patient Example 2 - Focal Cortical Dysplasia in the Right Superior Temporal Lobe	75
4.3.1	Tissue Segmentation	77
4.3.2	Plane of Symmetry Detection	79
4.3.3	Asymmetry Analysis	80
4.4	Plane of Symmetry Detection Robustness Test	85
4.5	Clinical Evaluation	89
4.6	EEG Source Localization	90
5	Future Work	94
5.1	Functional Neuroimaging	94
5.2	Diffusion Tensor Imaging	95
5.2.1	Anisotropic Patient-Specific Head Model	95
5.2.2	Decreased Fiber Connectivity of FCD Sites	96
5.2.3	Neural Connectivity of Potential Surgical Sites	96
6	Conclusion	98

List of Figures

1-1	Pipeline Snapshot	14
2-1	Subdural EEG	17
2-2	32-channel Surface EEG	18
2-3	128-channel Surface EEG	19
2-4	Epilepsy Surgical Planning Pipeline	26
3-1	Segmentation Morphological Post-processing	36
3-2	Plane of Symmetry Detection	38
3-3	Skull-Stripped Cortical Volume	40
4-1	Synthetic Data	50
4-2	Synthetic Data - Plane of Symmetry	52
4-3	Synthetic Data - Deformation Field	53
4-4	Synthetic Data - Deformation Field Magnitude	54
4-5	Synthetic Data - Deformation Field Divergence	54
4-6	Synthetic Data - Asymmetry Measure	54
4-7	Synthetic Data - Hypocontrast Deformation	55
4-8	Synthetic Data - Cortical Thickening Deformation	57
4-9	Synthetic Data - Interface Blurring Deformation	58
4-10	Synthetic Data - Merged Binary Clustering Map	59
4-11	Synthetic Data - Hotspots Label Map (Closed Merged)	60
4-12	Patient Example 1 - Multi-modal Rigid Registration	62
4-13	Patient Example 1 - Intensity Inhomogeneity Correction	63

4-14 Patient Example 1 - Anisotropic Smoothing	64
4-15 Patient Example 1 - Tissue Segmentation	65
4-16 Patient Example 1 - Plane of Symmetry	66
4-17 Patient Example 1 - Deformation Field Magnitude	67
4-18 Patient Example 1 - Deformation Field Divergence	68
4-19 Patient Example 1 - Asymmetry Measure	68
4-20 Patient Example 1 - Thresholded Asymmetry Measure	69
4-21 Patient Example 1 - Merged Binary Clustering Map	70
4-22 Patient Example 1 - Cluster Size Distribution	72
4-23 Patient Example 1 - Hotspots Label Map (Closed Merged)	73
4-24 Patient Example 1 - Hotspots Prior Probability Map	76
4-25 Patient Example 2 - Intensity Inhomogeneity Correction	78
4-26 Patient Example 2 - Anisotropic Smoothing	78
4-27 Patient Example 2 - Brain Mask	79
4-28 Patient Example 2 - Plane of Symmetry	80
4-29 Patient Example 2 - Deformation Field Magnitude	81
4-30 Patient Example 2 - Deformation Field Divergence	81
4-31 Patient Example 2 - Asymmetry Measure	82
4-32 Patient Example 2 - Thresholded Asymmetry Measure	82
4-33 Patient Example 2 - Merged Binary Clustering Map	83
4-34 Patient Example 2 - Hotspots Label Map (Closed Merged)	84
4-35 Plane of Symmetry Robustness Test Results - Coronal View	86
4-36 Plane of Symmetry Robustness Test Results - Axial View	87

List of Tables

4.1 Synthetic Data - Clustering Results	60
4.2 EEG Source Localization - Noise Level Parameters and Results	92

Chapter 1

Introduction

1.1 Focal Epilepsy

Epilepsy is a chronic medical condition produced by brief electrical disturbances in the brain causing recurrent seizures which affect awareness, movement, or sensation. The Epilepsy Foundation of America reports about 1% of Americans of all ages suffer from some form of epilepsy, amounting to some \$12.5 billion in healthcare expenditure every year[18]. Of these 2.5 million people, a large percentage suffer from symptomatic partial epilepsy, more commonly known as focal epilepsy. This is one of the most common types of epilepsy, caused by localized areas of abnormality resulting from strokes, tumors, trauma, congenital brain abnormalities, sclerosis, cysts or infections.

In the etiology of focal epilepsy, focal cortical dysplasia (FCD) is perhaps one of the most epileptogenic lesions associated with early onset medically refractive focal epilepsy. It is an increasingly recognized cause of intractable epilepsy in patients of all ages presenting for epilepsy surgery[2]. Focal cortical dysplasia is an abnormal growth or development of the cortex characterized by cortical and white matter abnormalities involving one or more gyri of the cortex. This is often manifested in the form of cortical thickening i.e. encroaching of the gray-matter (GM) boundary into white-matter (WM) tissue structures, poorly defined transition between GM and WM as well as hypointense signal, i.e. appears darker relative to surrounding tissues, within the dysplastic lesion relative to normal cortex[1]. Unfortunately, 25% of focal epileptic

patients are unresponsive to anti-epileptic medication [18]. For these patients, surgical intervention to remove the epileptic foci is the last alternative for permanent cure. This is especially so for infants because studies have shown that the chances for medical control and normal neuro-development are poor[16]. Consequently, surgical treatment should be considered early in infants with FCD. However, the success of such surgery is highly dependent on the doctor's ability to accurately locate the epileptogenic region for surgical resection. Hence, accurate detection of the epileptic foci is of paramount interest for these patients. At present, about 60% of medically intractable focal epileptic patients become seizure-free after surgery[11]. We hope to improve this figure through more accurate source localization.

1.2 Epilepsy Surgical Planning Project

The Epilepsy Surgical Planning project is a collaboration with the Children's Hospital Boston. The overall goal of this project is to provide a framework for accurate detection of focal epileptic hotspots, especially for pediatric patients. The aim is to provide an end-to-end quantitative analysis pipeline that fuses a range of imaging modalities including magnetic resonance imaging (MRI), diffusion tensor MRI, positron emission tomography (PET), single-photon emission computerized tomography (SPECT) as well as EEG data to build patient-specific head models and to compute prior probability maps of epileptic hotspots for more accurate epileptic foci detection. The contribution of this thesis is to build these patient-specific head models as well as prior probability maps of cortical abnormalities based on neuroimaging techniques. Together with the work by Tiferet A. Gazit to incorporate these two improvements into the formulation of the EEG source localization problem[7], the overall contribution of the Epilepsy Surgical Planning project is to improve the accuracy of source localization through more accurate head modelling and the introduction of a spatial prior on likely hotspot location based on information from neuroimaging modalities. The multi-modal analysis of focal epilepsy will greatly improve the accuracy of source localization and allow us to perform more accurate epileptic foci detection thereby

helping surgeons better localize the target foci for resection during the surgical evaluation phase.

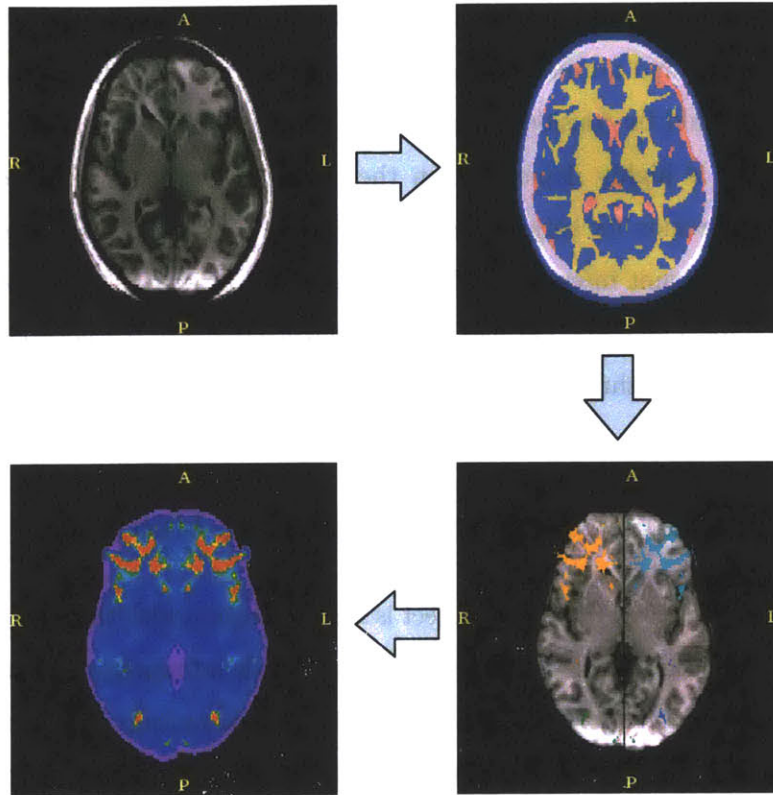


Figure 1-1: Pipeline Snapshot. Top Left: Original T1-weighted MRI. Top Right: Tissue Classification (Patient-specific Head Model). Bottom Right: Hotspots Label Map. Bottom Left: Prior Probability Map.

Figure 1-1 shows a snapshot sequence of the important outputs to expect from the pipeline proposed in this thesis. In the patient-specific head model, we delineate the imaged volume into skin, skull, cerebral-spinal-fluid, gray-matter and white-matter tissue classes. This patient-specific head model will allow us to simulate the potential distribution across the cortical volume of the patient more accurately. The prior probability map highlights areas of cortical abnormalities (red) while assigning appropriate values to other tissue classes to reflect the probability of an FCD foci being located in that voxel. This prior probability map allows us to constrain the EEG source localization search and consequently to bias the solution to these known areas

of cortical abnormalities where epileptic foci often lie.

Success of this pipeline is measured in two. Firstly, we obtain clinical validation of the prior probability map that we are indeed highlighting the relevant cortical abnormalities that may contain an epileptic foci. Secondly, we also demonstrate the improvement to EEG source localization with the incorporation of a patient-specific head model and a prior probability map. The ability to perform accurate EEG source localization even in noisy data conditions is further evidence of the success of this framework.

Chapter 2

Background

2.1 Accurate EEG Source Localization

Epileptic seizures are caused by electrical disturbances in the brain resulting in abnormal sensation, awareness or behavior. In the case of focal epilepsy, these electrical disturbances are generated from a number of abnormal focal spots within the cortex. Electroencephalogram (EEG) measurements can be used to pick up these electrical signals and EEG source localization techniques can be applied to infer the source of these electrical signals from the observed readings. The current gold standard technique, intracranial EEG monitoring, involves placing a subdural electrode grid (see Figure 2-1) on the surface of the cortex to accurately record EEG readings for optimal source localization results[21]. This method is less than ideal for obvious reasons at least from the patient's perspective. For one, placing of the electrode grid itself involves dangerous and invasive neurosurgical procedures. Furthermore, the electrode grid is left in the patient's head for durations of up to a week depending on the frequency of seizures and consequently the amount and quality of the EEG data collected. The patient is also highly susceptible to infection during this period making this entire procedure even riskier. Ultimately, if the EEG findings are non-localizing i.e. no epileptic foci can be detected from the readings, the patient would have gone through an incredible ordeal for no apparent benefit. Considerable expense, risk and possible morbidity[5] with no guaranteed results make this method of source

localization unacceptable in most cases.

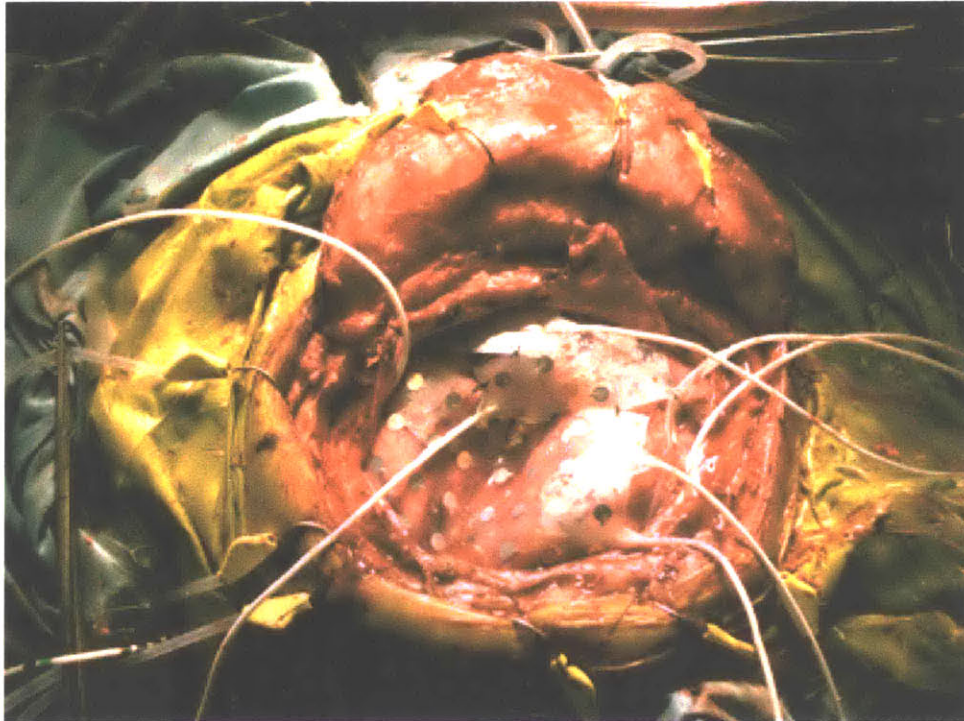


Figure 2-1: Subdural EEG. This shows the placement of multiple subdural electrode grids on the cortical surface of a patient after a craniotomy has been performed. The subdural grids can be left in the patient's head after the skull is subsequently closed for up to a week.

As such, there has been much interest in non-invasive pre-surgical techniques of source localization involving surface EEG monitoring. Unlike its subdural counterpart, surface EEG only involves placing an electrode grid on the scalp of the patient and recording the electrical signals from this grid array of surface electrodes. One huge advantage of this method is that there is no longer a need for invasive neuro-surgical procedures thus greatly reducing the amount of risk to which the patient is exposed during presurgical planning and evaluation. Unfortunately, several issues still plague this method of EEG measurement. Poor signal-to-noise ratio (SNR) and the inability of current models to accurately model the anatomy of the head exacerbates the low accuracy of epileptic foci detection through surface EEG source

localization. Besides that, source localization itself is an ill-posed inverse problem where the relationship between current distribution in the head and observed EEG potentials is not necessarily one-to-one i.e. there may exist multiple dipole locations which can generate similar surface potentials as observed on EEG measurements. As such, a common technique used nowadays relies on a doctor's visual inspection of the measured EEG signal to identify spikes corresponding to epileptic discharge within the cortical region associated with the electrode that recorded the spike. Consequently, epileptiform discharges can only be classified into 3 broad categories with this technique: regional (EEG signal spike observed over a single lobe or in 2 contiguous regions), multiregional (involving 2 or more regions in more than 1 lobe) or generalized[2].



Figure 2-2: 32-channel Surface EEG. The electrodes are often individually applied as opposed to being applied as a grid. The 10-20 system of electrode placement is often used for this setup.

In an effort to improve our ability to locate epileptic foci with greater accu-




Figure 2-3: 128-channel Surface EEG. The electrodes are applied as a grid providing for a more systematic placement of electrodes. The larger number of EEG channels also provide for higher spatial resolution.

racy, EEG equipment manufacturers are currently experimenting with 128-channel electrode grids (see Figure 2-3), up from the existing common 32-channel electrode configuration (see Figure 2-2), thereby increasing the sensitivity and resolution of EEG recordings. In this project, we also propose two optimizations in our quantitative analysis pipeline that are targeted towards improving source localization results, namely the incorporation of a realistic head model and a hotspots prior probability map. Initial studies have shown that source localization with realistic head models can achieve dipole placement errors of 11mm to 20mm depending on noise conditions of the input data[28] while a prior probability map of epileptic hotspots can help to constrain the solution set of the source localization problem. By biasing the search space of the inverse problem to suspicious regions as indicated by anatomical imaging techniques, source localization becomes more accurate and efficient by performing a targeted search on a subset of suspicious voxels as opposed to considering each imaged voxel with equal probability. Sections 2.2 and 2.3 provide further justification for these two proposed optimizations.

2.2 Incorporating a Patient-Specific Realistic Head Model

Traditional EEG analysis by visual inspection is simplistic at best and misleading at worst[5]. This highly qualitative technique usually involves manual visual inspection of raw EEG data by doctors to identify spikes in the EEG signal and to correlate these spikes to the cortical regions that are associated with the electrodes that measured these spikes[2]. Computational EEG analysis with EEG source localization provides slightly more encouraging results. Unfortunately, the current use of concentric spheres to model the human brain anatomy is highly inaccurate[5]. In reality, the physical geometry of our brain has a considerable effect on current distribution within the cortex and consequently on the EEG signals that are observed. The lack of accuracy in representing this geometry leads to a distorted interpretation of the observed EEG

signals. This problem is further exacerbated by inherently noisy surface EEG data thereby resulting in poor source localization.

Accuracy of source localization can be greatly improved if we can generate realistic anatomical models of the major tissue classes within the head i.e. skin, skull, gray-matter (GM), white-matter (WM) and cerebral spinal fluid (CSF). Segmented T1-weighted MRI and T2-weighted MRI can be used to generate isotropic anatomical models where conductivity of each tissue class is a average scalar value inferred from known literature for conductivity of that particular tissue class.

Studies have shown that skull conductivity is highly anisotropic with an anisotropy ratio of 1:10 (radially:tangentially to the skull surface) whilst white-matter exhibits anisotropy along the myelination sheaths with an anisotropy ratio of 1:9 (normal:parallel to fibers)[31]. We can easily model skull anisotropy with a simple direction dependent tensor where conductivity is 10 times higher along the skull's tangential direction as compared to its radial direction. Diffusion tensor MRI data can provide fiber-specific conductivity tensors at each white matter voxel. This allows us to accurately model the white anisotropy along the myelination sheaths for each patient individually.

2.3 Epileptic Hotspots Prior Probability Map

2.3.1 Multi-modal Analysis of Focal Epilepsy

The purpose of a hotspots prior probability map is to provide guesses as to the likely locations of epileptic foci based on anatomical magnetic resonance imaging (MRI) data as well as other functional modalities of imaging. Typical MRI findings for FCD include thickening of the cortical ribbon with abnormal gyral patterns, focal blurring of the gray-white matter interface[2] and hypointense signal within the dysplastic lesion relative to normal cortex[1]. By using a hotspots prior probability map based on these MR-visible pathologies, we can constrain the EEG source localization search space thereby providing a more accurate solution as well as improved computational efficiency.

There are strong clinical justifications for incorporating magnetic resonance imaging into source localization[6]. Advances in MRI techniques has increased the role that it plays in chronic seizure surgical planning. Consequently, MRI is now obtained at an earlier stage of the entire surgical evaluation process and it has gradually become the modality driving pre-surgical evaluation and planning[6]. When a MR-visible abnormality is identified, the work-up is primarily driven by the hypothesis that the lesion is the substrate for seizure disorder that is targeted for surgical resection. Over the last decade, it has become clear that pathologies identified on MR should be the main target of surgical resection[6].

In addition, studies have shown that 3D reconstruction and quantitative analysis of MRI can reveal additional positive information in up to 75% of patients suffering from epilepsy[24]. In recent years, higher spatial resolution in data acquisition as well as volumetric analysis of MRI have increased the proportion of epileptic cases in which lesions can be detected[24]. MRI has also been proven to be more sensitive than computerized tomography (CT) in detecting structural abnormalities because of its superior soft tissue contrast[11]. Especially in the context of an ill-posed problem such as EEG source localization, incorporation of structural MRI data can provide additional insight to constrain the search space and solution set of the inverse problem.

Other newer techniques that have become of interest in recent years include functional neuroimaging methods such as positron emission tomography (PET) and single-photon emission computerized tomography (SPECT) which provide excellent insights to the functional state of the brain[9]. Availability of such neuroimaging techniques enables us to investigate focal epilepsy not only as a structural pathology as with MRI but also as a physiological pathology reflected in abnormal blood flow and metabolism[6]. Physiological imaging investigations such as SPECT and PET provide complementary information to MRI in those patients considered for epilepsy surgery[4]. Such modalities are a relatively new but potentially very valuable addition to the suite of imaging techniques that doctors rely on to detect these elusive epileptogenic regions. Diffusion tensor MRI (DT-MRI) is a novel magnetic resonance technique and is also the only noninvasive means available today to measure molecular

diffusion in vivo. In particular, DT-MRI is able to reflect the anisotropic molecular motion of water along myelin sheaths in white matter tissue structures[3]. Incorporation of this modality in the pipeline allows us to construct realistic anisotropic head models and also to quantitatively analyze changes in subcortical fiber connectivity which is closely associated with regions of focal cortical dysplasia[15].

2.3.2 Asymmetry Detection Framework

An early method proposed for automatic detection of epileptic foci is that of inter-hemispheric volumetric quantification[24]. The goal in this method is to compare the degree of inter-hemispheric asymmetry based on the total volume of each imaged hemisphere. Quantitative analysis is performed by comparing the imaged volume of the left hemisphere to that of the right. Cortical asymmetry exists if there is a significant difference between the volume of the two hemispheres. Absolute magnitude of these volumetric differences are often small and can be easily overlooked by 2D visual inspection. Hence such a quantitative method provides for a more systematic detection of cortical asymmetry. However, while this quantitative method has better performance compared with the basic approach of 2D visual inspection, it does little beyond detecting major obvious structural abnormalities[30] and certainly makes no provisions to isolate specific epileptic foci that are associated with the detected asymmetry.

More sophisticated approaches include those discussed in [24], [30] and [29]. In [24], volumetric inter-hemispheric asymmetry comparison was extended to block analysis. The imaged brain volume was separated into smaller groups called blocks and direct comparison was made between blocks on either side of the mid-sagittal plane. If more than 1 block was significantly different, then asymmetry exists in the brain and specifically in those different blocks. This however is still a very coarse-grained approach where results are highly sensitive to how these blocks are chosen. In [30] and [29], a voxel-based approach is taken. By direct comparison with an atlas, regions of significant difference e.g. greater than 2σ from mean can indicate potential epileptogenic regions. While this method does well in isolating medium to large lesions, its

performance on smaller subtle abnormalities such as those present in focal cortical dysplasia is unclear.

Our proposed method follows the work of [25] and [13]. Specifically, we seek to detect subtle cortical changes such as abnormal thickening of the cortex or gray-matter clumps within white-matter tissue structure which are characteristic of focal cortical dysplasia (FCD). Asymmetry analysis provides a way to establish cortical thickening without reliance on a robust tissue classification method to delineate gray-matter and white-matter tissues. This is a huge advantage because segmentation is extremely reliant on good contrast between different tissue classes and this is exactly the element that is lacking in the focal pathology that we are trying to detect. Asymmetry analysis also does not rely on an anatomical atlas and this is extremely helpful because we are interested in detecting FCD in pediatric patients where there is a great degree of variability in cortical size and shape depending on their age. This makes the construction of an age-specific atlas both tedious and time-consuming. Consequently, asymmetry analysis works by detecting significant intra-patient interhemispheric differences and thus significant asymmetry across the mid-sagittal plane in the imaged brain volume strongly suggests MR-visible epileptogenic foci that should be further investigated with EEG source localization.

In our formulation, we propose an alternative method of finding the plane of symmetry as well as an efficient clustering scheme to generate a hotspots prior probability map from voxels with significant asymmetry. The ability to accurately detect this mid-sagittal plane of symmetry is of paramount importance to the pipeline because interhemispheric comparisons are only valid and relevant if this plane of symmetry is accurately defined. Contralateral structures with reference to this plane of symmetry can be effectively compared to isolate regions of significant difference. Existing methods provide the ability to model radiometric variations as a separate additive intensity inhomogeneity bias field. This however introduces too many degrees of freedom in the optimization problem of finding a plane of symmetry and might result in overfitting of the input data. Instead, our approach of finding the plane of symmetry assumes that any bias field has been previously corrected and any intensity difference

is due to structural asymmetries present in the cortex. This is also a more intuitive way of modelling structural abnormalities.

Our measure of success is also different. Unlike regular lesion detectors, the main contribution of this method is the ability to accurately detect subtle intra-patient abnormalities such as those present in focal cortical dysplasia (FCD). As such, a higher false positive rate is perfectly acceptable as long as false negative rates are low. This is reasonable because structural asymmetries exist even in normal healthy subjects.

Moreover, there is also a left / right ambiguity present in asymmetry analysis which arises from our lack of additional information to determine which hemisphere actually contains the cortical abnormality. Consequently, this also increases the false positive rates. Nonetheless, the main objective of this pipeline is not provide a definitive lesion detector. Instead, the goal of asymmetry analysis is to highlight potential spots of cortical abnormalities so as to bias the solution of EEG source localization towards these known anomalous regions that are more likely to contain epileptic foci. With this consideration in mind, it is thus more important to maintain a low false negative rate for this framework in order not to prematurely exclude any actual epileptic foci.

Note that this asymmetry detection framework is not limited to volumetric MRI data only. In fact, this framework can be also be applied to diffusion-tensor MRI (DT-MRI) data to detect reduction in sub-cortical fiber connectivity[15] as well as positron emission tomography (PET) and single-photon emission computerized tomography (SPECT) functional neuroimaging data to further identify potential epileptic hotspots which are not captured by conventional MRI imaging techniques[9].

2.4 Epilepsy Surgical Planning Pipeline

Figure 2-4 provides a summary of the pipeline stages as described in the previous sections. Grayed-out boxes represent input modalities that were not available for testing at the time of this thesis. They have nevertheless been included in this framework to

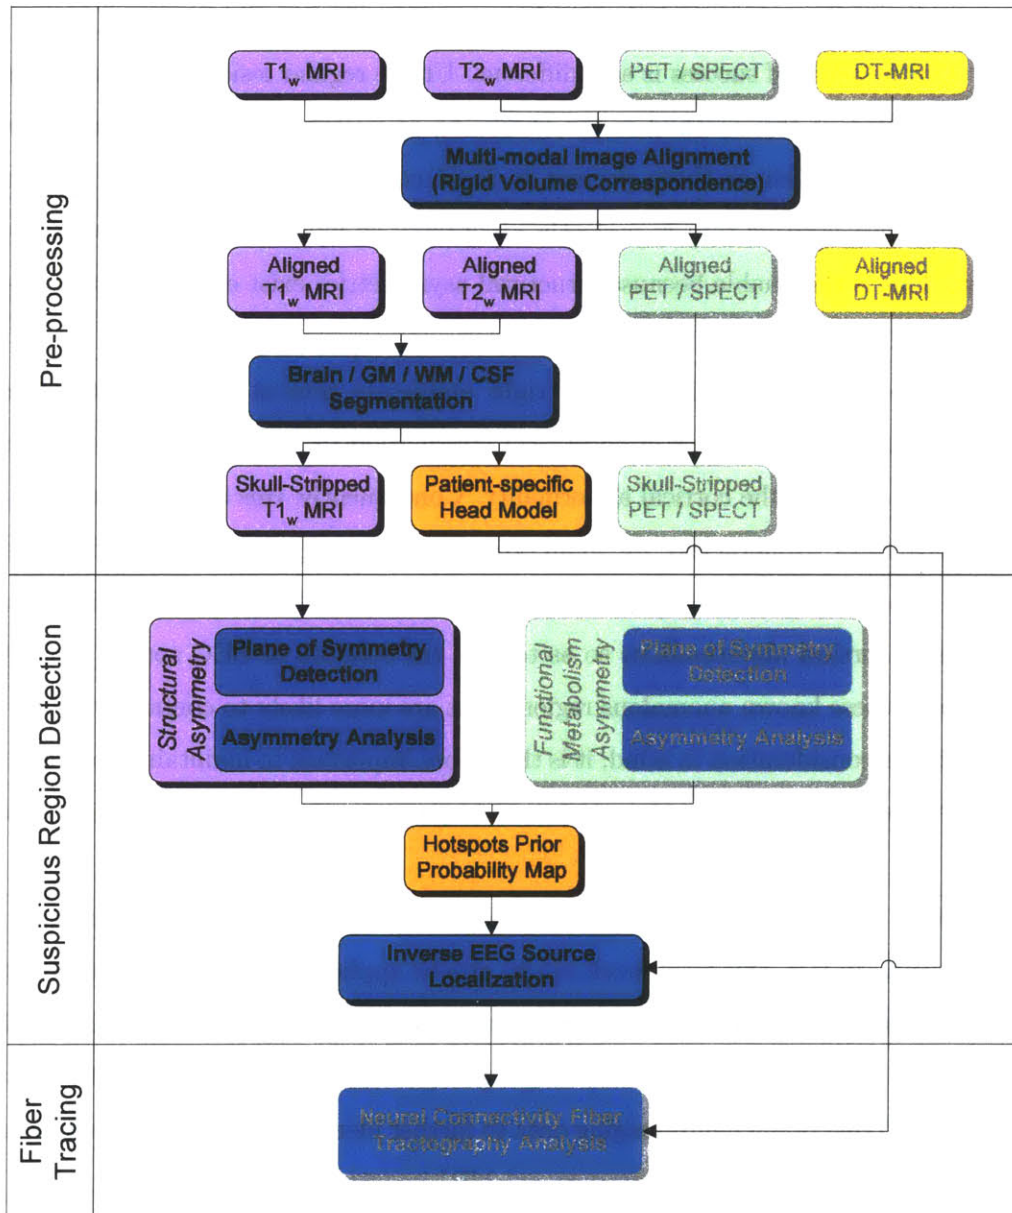


Figure 2-4: Epilepsy Surgical Planning Pipeline. The orange boxes give the desired output of this pipeline while the purple, green and yellow boxes trace the flow of different imaging modalities through the pipeline. The blue boxes represent algorithms that operate on the input volumes and grayed-out boxes represent possible future work

illustrate the potential of the pipeline to incorporate these neuroimaging modalities in the future.

There are 3 major stages in this pipeline, namely the pre-processing stage, the suspicious region detection stage and the fiber tracing stage. The pre-processing stage essentially involves aligning the different input modalities to be considered for a particular patient into the same coordinate system. Tissue classification is also carried out in this stage to construct the isotropic patient-specific head model. The suspicious region detection stage performs plane of symmetry detection and asymmetry analysis to compute the hotspots prior probability map. Note that the asymmetry analysis module can be extended for neuroimaging modalities other than structural MRI and this is further discussed in Section 5.1. The fiber tracing stage is included in this diagram for completeness. This stage was not actually implemented and tested in this thesis due to the lack of corresponding volumetric DT-MRI data for patient cases that we were analyzing. However, DT-MRI processing capabilities available in 3D-Slicer can be readily used to incorporate this modality into the quantitative analysis of focal epilepsy. Section 5.2 provides a brief overview of the potential improvements that is possible with adding the DT-MRI modality in the future. The two main contributions of this thesis are the patient-specific head model and the hotspots prior probability map. Epileptic hotspot priors as well as source localization results can be overlaid on corresponding structural and functional imaging modalities to visualize and validate detected abnormalities.

Chapter 3

Method

3.1 Pre-processing

3.1.1 Rigid Volume Correspondence for Multi-modal Analysis

Multi-modal volumetric medical image registration is an important pre-processing step to align all input modalities to the same coordinate system for subsequent volumetric analysis. In this way, each input modality can contribute meaningful and distinct characteristics of the imaged volume. For example, T1-weighted magnetic resonance imaging (MRI) provides excellent gray-matter(GM) white-matter(WM) contrast whilst cerebral-spinal-fluid(CSF) shows up best on T2-weighted MRI. This allows us to combine different imaging modalities in this quantitative analysis to obtain a more complete picture of the anatomy and pathology present in each patient.

Since each input modality is obtained using a different imaging technique, a simple voxel-based comparison (such as sum of squared difference) to determine the degree of match between any two modalities is highly ineffective. Even when perfectly registered, T1-weighted MRI of an individual is very different from a corresponding T2-weighted MRI taken from the same individual due to inherent imaging characteristics of the two techniques involved. However, it is exactly this difference that brings value in combining multiple imaging techniques.

The solution of choice to determine degree of match is the mutual information measure[27]. T1-weighted MRI and the corresponding T2-weighted MRI are not completely uncorrelated; they are after all measures of tissue properties of the same imaged volume. Intuitively, knowing the T1-weighted response of a particular tissue type can give us some information about what the T2-weighted response should be. This is the essence of the mutual information registration approach as presented in [27]. By choosing an arbitrary volume as the reference or fixed volume, all other target volumes can be registered to this fixed volume using this mutual information approach.

Several formulations of mutual information and entropy exist, as surveyed in [20]. In general, mutual information is the natural measure of dependence between random variables that takes into account the whole dependence structure of the variables involved. For our purposes, mutual information between two volumetric images A and B is defined as follows:

$$MI(A, B) = H(A) + H(B) - H(A, B) \quad (3.1)$$

where $H(X)$ is the Shannon entropy of a random variable X and $H(X, Y)$ is the joint entropy between random variables X and Y . These are in turn given by:

$$H(X) = \sum_x p(x) \ln p(x)$$

$$H(X, Y) = \sum_{x,y} p(x, y) \ln p(x, y)$$

where $p(x)$ is the probability of x and $p(x, y)$ is the joint probability of observing both x and y .

The problem of multi-modal image registration between any two volumes thus becomes one of finding the optimal transformation that maximizes their mutual information and subsequently applying this transformation to bring the target volume into the same coordinate system of the fixed volume:

$$T^* = \arg \max_T \text{MI}(A, T(B)) \quad (3.2)$$

The above optimization problem can be easily solved using any regular gradient descent algorithm optimized over the desired transformation parameters. In particular, multi-modal registration as described in this section was implemented using the registration and gradient descent framework[12] of the Insight Toolkit (ITK).

One potential problem with the optimization as described above is the overfitting of data by unnecessarily warping the target volume to fit the fixed volume. For example, the results obtained from performing the above optimization over a set of affine transformations will invariably be different from that of pure rotation and translation because of the additional degrees of freedom that affine transformations provide. Consequently, extreme skewing of the target volume may produce a local optima of good mutual information match that does not accurately register the input modalities. Given that we are only interested in performing intra-patient rigid registration, we can further constrain the set of possible transformations to that of scaling, rotation and translation only to avoid this problem.

Other neuroimaging modalities such as PET, SPECT and baseline DT-MRI volumes can be aligned to an arbitrary target volume in a similar fashion. For the purposes of this pipeline, we shall choose T1-weighted MRI as the default target volume to which all other volumes shall be aligned.

3.1.2 Modified Watershed Tissue Segmentation

The following steps are carried out in this section:

- Intensity correction to remove spatial inhomogeneities in input volumes
- Anisotropic noise smoothing to reduce image noise while maintaining important edge boundaries
- Watershed-based tissue classification

In general, magnetic resonance imaging (MRI) provides a good structural view of the brain and any abnormalities if present, due to its superior capability for good soft tissue contrast[11]. In particular, T1-weighted magnetic resonance imaging (MRI) provides excellent gray-matter (GM) white-matter (WM) contrast whilst cerebral-spinal-fluid (CSF) shows up best on T2-weighted MRI[3]. By combining these two modes of MRI, we can obtain more accurate tissue classification of the imaged cortical volume into the following tissue classes: skin, skull, cerebral-spinal-fluid, gray-matter and white-matter. Consequently, a realistic head model of the patient can be obtained from these tissue classes. This is crucial for accurate EEG source localization because each of these individual tissue classes has very different conductance properties and an accurate modeling of this can lead to better localization results. Volumetric meshing of each tissue type can give us a patient-specific isotropic head model while an anisotropic patient-specific head model can be subsequently obtained by assigning a conductivity tensor to each mesh element for more detailed modeling of the brain[10]. In addition, we can also use the above segmentation results to perform skull stripping i.e. the process of removing the skull and other extraneous tissues thereby leaving only the brain volume. As we shall see in later sections, this is important in restricting the region of interest (ROI) for asymmetry analysis because we are only concerned with asymmetries within the cortex.

In order to proceed with segmentation, the first step is to correct for intensity inhomogeneities of volumetric MRI scans. Inhomogeneities arise from numerous sources such as induced eddy currents and spatial inhomogeneity of the excitation field. In many cases, such intensity variations can reach as high as 20% in standard acquisitions and perhaps more in higher field scanners[17]. Given the intra-class intensity homogeneity assumption of most intensity-based tissue classifiers[22, 32], correction of this variation thus becomes a crucial factor for good segmentation.

The intensity inhomogeneity correction algorithm used in the pipeline follows the work of Mangin[17]. This method was chosen because it does not require a prior model of the desired tissue classes. It also uses a rapid simulated annealing schedule that converges to a global minimum, thereby avoiding potential convergence to local

minima, which is unavoidable in current EM algorithms. In this formulation, image formation is modelled as follows:

$$O(x) = S(x)F(x) + N(x)$$

where $O(x)$ is the observed signal intensity at voxel x , $S(x)$ is the intensity of the signal generated by the tissue imaged at x , $F(x)$ is an unknown smooth multiplicative bias field and $N(x)$ is white noise.

Assuming that individual tissue classes generate similar intra-class signal intensities, the task at hand is thus to find an optimal corrective field, $F_c(x)$, that corrects for the above phenomena. Note that the corrective field, $F_c(x)$, has no direct correlation with the smooth multiplicative bias field, $F(x)$; $F_c(x)$ is simply a computed field that best corrects for the image formation phenomena as shown above.

Entropy is a good measure of image quality given our intra-class homogeneity assumption. In the ideal case of perfect intra-class homogeneity, histogram binning should show clear delineation of the different tissue classes and consequently a low entropy value. As such, minimization of entropy will give us a globally optimal bias correction field that accounts for the intensity inhomogeneities as modelled above. The cost function for minimization is defined as follows:

$$C = K_I I(F_c O) + K_R R(F_c) + K_M M(F_c O) \quad (3.3)$$

where $F_c O$ is the restored image, a voxel-wise product of the corrective field F_c and the observed image O , $I(F_c O)$ is the entropy of the restored image, $R(F_c)$ is the regularization term for a smooth corrective field and $M(F_c O)$ is a quadratic measure of discrepancy between the observed image mean and restored image mean to prevent the optimal field from being uniformly null. K_I , K_R and K_M are all positive weighting constants. Minimization of the above cost function returns a globally optimal bias correction field that removes intensity inhomogeneities present in volumetric MRI acquisitions.

Another pre-segmentation optimization is that of edge-preserving noise smoothing or anisotropic diffusion. In essence, the goal is to reduce noise by performing image smoothing whilst maintaining important image features. Unlike regular Gaussian smoothing, edge-preserving noise smoothing significantly reduces noise whilst enhancing contrast between the different tissue boundaries, thereby producing better segmentation results. To this end, ITK’s Curvature Flow anisotropic diffusion filter is used to implement this feature. This filter performs edge-preserving smoothing much in the same fashion as classical anisotropic diffusion[12, 19] but uses a level set formulation where iso-intensity contours of a volumetric image form the level sets[12]. The level set function, L , is evolved under the control of the heat equation as with classical anisotropic diffusion. In this case however, speed of diffusion i.e. speed of level set evolution, L_t , is chosen to be directly proportional to curvature of the contour i.e. $L_t = \kappa|\nabla L|$ where κ is curvature. Areas with high curvature (i.e. noise) will diffuse faster and disappear whereas areas with low curvature (i.e. large-scale features) will be slow to evolve thereby preserving edges. However, care must be taken not to over-evolve the level set because even large-scale boundaries can be diffused away through continuous level set evolution.

At this point, we have a bias-corrected, noise smoothed volume that is ready for segmentation. Our method of choice for tissue classification is that of a statistical watershed transform for medical image segmentation[8]. In general, watershed classification techniques usually require seed voxels which have been already classified as a particular tissue type. Voxels are then correspondingly labeled to the respective tissue types by progressively selecting the next most similar voxel starting from the initial set of seed voxels. At present, this is a supervised classification method due to the lack of segmented healthy pediatric volumetric scans from which to construct age-specific anatomical atlases. An operator must manually seed the classification of desired tissue types to initialize the classifier. However, this manual step can be eliminated in the future with the incorporation of an age-specific atlas which can provide automatic seed selection from statistical priors. Nonetheless, watershed methods for segmentation are preferred because they are fast and parallelizable. Moreover, it

always produces complete divisions of the volume even in poor contrast thus avoiding the need for any form of contour joining. Local variations such as noise however can dramatically change final segmentation results, although anisotropic diffusion filters can correct for this high noise sensitivity of watershed methods.

The following is an outline of the method implemented in this pipeline as described in [8]. Given the intensity value I_x at each voxel x , posterior probabilities for each tissue class w_k at each voxel can be calculated using Bayes' rule as follows:

$$P_x(w_k|I_x) = \frac{P_x(I_x|w_k)P_x(w_k)}{\sum_{w_k} P_x(I_x|w_k)P_x(w_k)} \quad (3.4)$$

where k indexes over the number of desired tissue classes.

Assuming that the signal intensities of desired tissue classes follow a normal distribution, mean and variance of the distribution can be derived from the set of seed voxel labelings for each class (process of seed voxel selection is discussed in greater detail later in this section). This in essence allows us to approximate $P_x(I_x|w_k)$. In addition, especially after intensity inhomogeneity correction as mentioned above, desired tissue classes typically exhibit significant intra-class spatial homogeneity and it is thus desirable to capture this in calculation of the posterior probabilities. Markov Random Fields provide a convenient way to model local spatial correlation between neighboring voxels and can be formulated as follows:

$$P_x(w_k) = \frac{1}{Z} \exp \left(-\beta \frac{\sum_{n \in N(x)} \delta(w_k, w_{k_n})}{d(x, n)} \right) \quad (3.5)$$

where $P_x(w_k)$ is the class probability for class w_k at voxel x , $N(x)$ denotes the neighbors of voxel x , w_{k_n} is the tissue class assigned to voxel n , $d(x, n)$ is the Euclidean distance between voxels x and n and

$$\delta(w_{k_i}, w_{k_j}) = \begin{cases} -1, & \text{if } w_{k_i} = w_{k_j} \\ +1, & \text{if } w_{k_i} \neq w_{k_j} \end{cases}$$

By solving the above two equations iteratively using Iterative Conditional Modes (ICM), we can obtain the set of posterior probabilities for each class at every corresponding voxel. From these posterior probabilities, we define the similarity function for watershed transform as follows:

$$f_{w_k}(x, y) = P(w_k|I_x) - P(w_k|I_y) \quad (3.6)$$

where $P(w_k|I_x)$ and $P(w_k|I_y)$ denote the posterior probabilities of class w_k at voxels x and y .

Classical watershed transform then proceeds in the following manner. Seed voxels are initialized as particular tissue types from seed training data. Regions are labelled progressively by selecting the next most similar voxel from the set of all neighboring voxels where similarity is given by Equation 3.6.

The actual process of 3D tissue classification is two-fold. First, representative 2D slices are chosen to be manually segmented with the method as described above. Care should be taken to choose slices which contain regions that are more likely to be misclassified due to poor delineation at tissue boundaries in the imaged volume. Voxels in these slices are manually seeded according to their desired tissue class and 2D segmentation is carried out to ensure that tissue classification is satisfactory for each slice. Note that not all voxels within a slice have to be seeded in order for 2D segmentation to proceed. In practice, we have found it useful to seed only tissue boundary voxels. Given our intra-class intensity homogeneity assumption, voxels within their respective boundaries should follow the tissue class of the corresponding group of boundary voxels. After each 2D segmentation, the resulting 2D label map is checked to ensure that a majority of the voxels in the slice have been correctly classified to their desired tissue class. Further refinements can be made especially to misclassified voxels within each slice by manually seeding these voxels to their desired tissue class. Once 2D segmentation of these representative slices are all satisfactory, tissue classification is then extended to 3D with the combined set of seed voxels from all the individually segmented 2D slices. Similarly, this 3D label map can be further

refined by selecting more representative 2D slices for manual segmentation and the corresponding seed voxels that result from this 2D segmentation can be added to the existing set of seed voxels for improved 3D tissue classification.

After segmentation is complete, further morphological post-processing - erosion, largest component selection, dilation - is carried out to remove spurious foreign tissue islands that may have resulted from persistent noise in the data that was misclassified. This can particularly helpful for skull stripping because we are interested in extracting the entire brain mass within the cranial cavity and thus holes within this cavity should still be considered as part of the desired brain mass to be segmented.




Figure 3-1: Segmentation Morphological Post-processing. Left: Raw tissue classification label map before morphological post-processing. Right: Tissue classification label map after morphological post-processing.

3.2 Plane of Symmetry Detection

The methods and algorithms that have been mentioned thus far belong to the pre-processing stages of the quantitative analysis pipeline. Plane of symmetry detection marks the beginning of the asymmetry analysis portion of detecting epileptic hotspots. In order for meaningful analysis of cerebral asymmetry in the next stage, the ability to reliably find the plane of symmetry for a wide range of pathology (i.e. both small and large scale abnormalities) that can be present in the imaged volume is crucial. Given that our goal is to detect focal cortical dysplasia (FCD) in epileptic patients,

we shall focus our efforts in developing a method that works robustly for small to medium-scale abnormalities that are characteristic of FCD abnormalities.

For our purpose, we shall define the desired plane of symmetry as the plane where inter-hemispheric similarity is maximum[13]. Consequently, this plane of maximal inter-hemisphere similarity is also often the mid-sagittal plane as defined by the anterior commissure and the posterior commissure because the brain is a largely symmetrical volume about the mid-sagittal plane. This is however not always true. For example, a patient with middle cerebral artery infarction can have a significant structural abnormality amounting to more than 40% of the brain volume resulting in the shifting of mid-line structures into the contralateral hemisphere as evident the example presented in Section 4.4. In cases like this, it is hard to expect the plane of maximal inter-hemisphere similarity to coincide perfectly with the mid-sagittal plane. Nonetheless, our assumption of small to medium-scale abnormalities present in patients with FCD usually avoids this problem.

Unlike other methods which explicitly solve for the plane of symmetry, our solution is implicit by aligning the plane of symmetry with an arbitrary user-specified target plane. In this context, the chiral image of the original volume shall refer to the reflection of the original volume about this arbitrary target plane. The goal of symmetry axis detection is thus to find the optimal rigid transformation necessary to map the chiral image of the original volume back to the original volume. Once this rigid transformation is determined, this transformation is simply “halved” and applied on the chiral image to align the chiral image to the target plane of symmetry (see Figure 3-2). For completeness, the aligned chiral image of the original volume can be reflected again about the target plane to obtain the aligned original volume.

In order to effectively derive a volume’s chiral image about an arbitrary user-specified target plane, we shall use the Householder Reflection matrix formulation as follows:

$$H = I - 2 \frac{\vec{v}\vec{v}^T}{\vec{v}^T\vec{v}} \quad (3.7)$$




Figure 3-2: Plane of Symmetry Detection. H is the Householder reflection matrix that maps I to I_c . T_R^* is the optimal rigid transformation that maps I_c back to I . $T_R^{*\frac{1}{2}}$ is the final transformation that aligns I_c to the target plane of symmetry.

where H is the Householder matrix, \vec{v} is the Householder vector and I is the identity matrix. By choosing \vec{v} such that it is a vector normal to the target plane, we can construct a transformation matrix H that performs the reflection operation to derive the chiral image of the original volume. Note that the Householder matrix H is neither dependent on the sign of the direction of the Householder vector \vec{v} nor on its magnitude. H is solely determined by the absolute direction of the Householder vector \vec{v} . In the context of this pipeline where we are interested in reflection about a 2D plane, we choose a $\vec{v} \in \mathbb{R}^3$ such that H reflects every vector $\vec{x} \in \mathbb{R}^3$ about the 2D plane given by $span\{\vec{v}\}^\perp$. As such, the Householder matrix provides a very elegant way to formalize reflection about an arbitrary symmetry plane. For the purposes of this pipeline, the target symmetry plane shall simply be the mid-sagittal slice of the 3D imaged volume.

The same formulation for optimal rigid transformation estimation using ITK's mutual information registration and gradient descent framework is used here to determine the optimal rigid transformation that will map the chiral image of the original volume back to the original volume. Unlike the intensity justifications that were presented in Section 3.1.1 for using the mutual information measure in multi-modal rigid

registration, the motivation for using mutual information in this context of intrapatient alignment of an imaged volume with its chiral image is different. Mutual information is the desired similarity metric because it encapsulates the whole dependence structure of volumes to be aligned in its formulation. By having a global goodness of fit measure, localized asymmetries such as the abnormalities present with FCDs are less likely to throw the algorithm off. This provides a significant improvement over purely voxel-based approaches such as the sum-of-squared-difference (SSD) measure which is extremely sensitive to even small-scale asymmetries. Existing voxel-based methods propose the estimation of an additive intensity field to account for both bias field inhomogeneity and localized asymmetries[13] that might lead to overfitting of the data due to the excessive degrees of freedom that are inherent in the design of the algorithm. Consequently, the proposed mutual information approach does not require the explicit computation of any additive intensity field to account for these localized asymmetries, and bias field inhomogeneity can be separately corrected using the mutual information intensity inhomogeneity correction algorithm as mentioned in Section 3.1.2.

An additional enhancement for symmetry axis detection provided for by this pipeline is the skull-stripping feature in the tissue classification process. Figure 3-3 shows an example of such a skull-stripped cortical volume. The input volume to the symmetry axis detector is that of the brain mass only i.e. volume that remains after skull stripping to eliminate potential noise artifacts in unrelated tissue structures such as the skull and other extraneous tissue classes. This also restricts the region of interest for subsequent asymmetry analysis to the brain only. Intensity variations present in the skull and other extraneous tissue structures might give rise to spurious epileptic hotspots that are outside of the brain thereby unnecessarily increasing the number of false positives.




Figure 3-3: Skull-Stripped Cortical Volume. Top Row: Original T1-weighted MRI. Bottom Row: Skull-stripped cortical volume.

3.3 Asymmetry Analysis

Alignment of the mid-sagittal plane to the target symmetry plane provides the basis for further asymmetry analysis. In this section, the goal is to identify the localized areas of inter-hemispheric differences about this target plane of symmetry. Since normal healthy brains are largely symmetrical, save for the normal functional asymmetries that are present even in a healthy control, the intuition is that these differences might indicate potential epileptic sites i.e. hotspots that can be used to constrain the solution set of the EEG source localization problem. As such, the aim is to detect all “visible” epileptic foci that are captured within the imaging modality being considered. A higher false positive rate is perfectly acceptable at this point due to normal asymmetries that are present even in healthy subjects. In fact, maintaining a low false negative rate is of far greater importance so as to compute a meaningful hotspots prior probability map that does not preclude any true epileptic site.

To identify these potential epileptic hotspots, non-rigid image registration techniques can be used to find the optimal deformation field that will map the aligned original volume to its corresponding chiral volume about the target plane of symmetry. We can formalize the plane of symmetry detection and asymmetry analysis problem as follows:

- Compute the chiral volume of the original volume, $I_c = H(I)$ where H is the Householder reflection function as described in Section 3.2 and I is the original volume. H is defined by choosing a Householder vector $\vec{v} \in \mathfrak{R}^3$ such that it is a vector normal to the target symmetry plane. H then becomes a transformation matrix that reflects every vector $\vec{x} \in \mathfrak{R}^3$ about the 2D plane given by $span\{\vec{v}\}^\perp$.
- Find the optimal rigid transformation, T_R^* , that best matches I with I_c as fol-

lows:

$$\begin{aligned}
T_R(I) &= f(I; \vec{n}, \theta, \vec{t}) \\
(\vec{n}^*, \theta^*, \vec{t}^*) &= \arg \max_{\vec{n}, \theta, \vec{t}} \text{MI}(T_R(I), I_c) \\
T_R^*(I) &= f(I; \vec{n}^*, \theta^*, \vec{t}^*)
\end{aligned}$$

where f is a composite rigid transformation of rotation given by the unit quaternion parameters (\vec{n}, θ) and translation given by \vec{t} . $T_R(I)$ is the resulting volume after the rigid transformation f has been applied. MI is the mutual information operator.

- Derive the aligned original volume, I^s , and its chiral volume, I_c^s as follows:

$$T_R^{*\frac{1}{2}}(I) = f\left(I; \frac{\vec{n}^*}{\sqrt{|\vec{n}^*|^2 + \frac{\theta^{*2}}{4}}}, \frac{\frac{\theta^*}{2}}{\sqrt{|\vec{n}^*|^2 + \frac{\theta^{*2}}{4}}}, \frac{\vec{t}^*}{2}\right) \quad (3.8)$$

$$I_c^s = T_R^{*\frac{1}{2}}(I_c) \quad (3.9)$$

$$I^s = H(I_c^s) \quad (3.10)$$

- Find the optimal non-rigid deformation field, T_{NR}^* that best matches the aligned original volume, I^s , and its chiral image, I_c^s as follows:

$$T_{NR}^* = \arg \min_{T_{NR}} \text{SSD}(T_{NR}(I_c^s), I_c^s) \quad (3.11)$$

where SSD is the sum of squared difference operator.

This optimal non-rigid deformation field, T_{NR}^* , allows us to infer the inter-hemispheric differences between the left and right hemispheres of the cortical volume. Computation of T_{NR}^* is accomplished with ITK's Level Set Motion Sum-of-Squared Differences deformable registration filter which was in turn based on work by [26]. The formulation of this filter is based on the Demons algorithm in that the equations of motion are governed by instantaneous optical flow equations. However, calculation

of motion vectors differ in this level set implementation. The motion vector in this case is a function of intensity differences between the fixed and target voxel scaled by the maximum gradient over the entire field. The demons algorithm on the other hand defines its motion vectors directly as a function of both intensity differences and gradient magnitude at each voxel. Consequently, the motion vector of a voxel over a low gradient region is larger than that over a high gradient region resulting in an unstable vector field in the demons algorithm. In this level-set formulation however, motion vectors are proportional to gradients resulting in a more stable field. Moreover, since we are analyzing intra-patient asymmetry, sum-of-squared difference methods can provide a better measure of asymmetry by looking directly at the intensity values at a voxel-based level. Assuming intensity inhomogeneities have already been corrected, this direct analysis makes more intuitive sense in detecting localized regions of abnormality. Arguably an SSD formulation provides better sensitivity to localized asymmetries which is desirable at this stage of the pipeline.

Once the optimal deformation field is computed, it has to be converted to a measure of asymmetry to highlight potential epileptic hotspots for further analysis. Given the noisy nature of neuroimaging data, it is useful at this point to apply a simple magnitude threshold to remove voxels with very small displacements. These extremely small displacements are often caused by surrounding noise points and should not be included in our asymmetry measure. A typical threshold value is 2 voxels. Depending on the image resolution, this usually translates to removing deformations that are smaller than 1mm to 2mm.

In order to convert the deformation field to a measure of asymmetry, we shall consider the following vector field operator, $|F|\nabla \cdot F$ where $F = T_{NR}^*$, which was shown by [25] to be good for detecting cancerous growths or lesions. $|F|$ captures regions of high deformation whilst $\nabla \cdot F$ highlights regions of expansion or atrophy. Taken together, $|F|\nabla \cdot F$ provides more robustness in detecting abnormalities in noisy data by emphasizing the high magnitude and high divergence asymmetries.

For our purpose, increased cortical thickness, gray-white matter interface blurring and regions of hypointense signal are typical of the focal cortical dysplasia (FCD)

pathologies that we aim to detect. Clearly, the operator $|F|\nabla \cdot F$ works best in detecting hypointense regions because they are mostly nodular in nature and are characterized by high deformation and high divergence. However, for regions of cortical thickening and interface blurring, these pathologies are also sometimes characterized by high deformation and low divergence owing to the possible planar nature of these pathologies. As such, we propose the following asymmetry measure, $A_{T_{NR}^*}$:

$$A_{T_{NR}^*} = |F|(C + \nabla \cdot F) \quad \text{where} \quad C = \begin{cases} -1, & \text{if } \nabla \cdot F < 0 \\ +1, & \text{if } \nabla \cdot F \geq 0 \end{cases} \quad (3.12)$$

This new operator places an emphasis on $|F|$ and ensures that the magnitude of the second term of $A_{T_{NR}^*}$ is always at least 1. In other words, $|A_{T_{NR}^*}| \geq |F|$. Regions of high deformation and low divergence will still be able to stand out as compared with the $|F|\nabla \cdot F$ operator where such regions tend to have an insignificant asymmetry measure. With this new operator, asymmetry measure magnitude is directly proportional to the magnitude of the deformation field at every voxel. In addition, regions with high divergence will continue to stand out while regions with low divergence will not fade away as quickly as long as the deformation magnitude is high. This operator aids in the detection of cortical thickening and interface blurring as well as hypocontrast regions.

Finally, in order to generate a prior epileptic hotspot label map, we shall only consider voxels of significant asymmetry. In order to threshold our measure of asymmetry measure, we must first note that the inverse of the optimal solution that deforms the left hemisphere to the right hemisphere may not necessarily be the optimal solution that deforms the right hemisphere to the left hemisphere. Since this is an optimization problem, the deformation field derived need not be contralaterally symmetrical and for the most part, it is not. Hence, simple thresholding i.e. $\mu \pm N\sigma$ holds little significance for such an asymmetry measure with both positive and negative values that is not centered about 0. Hence we shall define a symmetric thresholding level by finding the mean, μ_{abs} and standard deviation, σ_{abs} , of the absolute asymmetry

measure, $|A_{T_{NR}^*}|$. Voxels of significant asymmetry shall be defined as follows:

$$|A_{T_{NR}^*}| > \mu_{abs} + S\sigma_{abs} \quad (3.13)$$

where S is the asymmetry sizing factor. If the asymmetries to be detected are small relative to normal healthy asymmetry, S should be small. In cases where pathology turns out to be about the same size or even smaller than normal healthy asymmetry, S can be negative. In general, S should be as large as possible so that clearer hotspot boundaries can be obtained.

Clustering is then performed on these voxels of significant asymmetry to delineate them into separate regions where each region represents the probability distribution of an epileptic foci in that area. Granted that no prior knowledge of epileptic foci count is available (nor is it reasonable to make any assumption on this number especially in an automatic framework), hierarchical clustering methods are the natural choice to construct highly descriptive cluster trees that permit convenient further manipulation i.e. the tree can be cut at a certain height to obtain clusters of a minimum specified number of member voxels. However, the single biggest disadvantage of this algorithm is its computational complexity of $O(n^2 \log(n))$. Despite being polynomial in time, this method does not scale well in practice. Initial experiments to cluster 20,000+ voxels (which is only 0.2% of the total imaged volume) did not terminate even after a day.

Hence we propose an alternative method of clustering for our pipeline using a hybrid connected component formulation in the interest of time and space efficiency. Clustering is achieved through the connected component analysis framework by considering the labeling of the neighborhood voxels of any voxel to be labeled. In this case, we shall define the neighborhood function with a structuring element similar to the construct used in image morphological operations. This allows us to directly control the size of the neighborhood region to consider for the labeling of each voxel. Hence for each voxel, 3 possible cases can arise:

- No neighborhood voxels have been classified. Assign a new label to this voxel.
- 1 neighborhood voxel has been classified. Assign the label of this neighborhood voxel to the current voxel.
- More than 1 neighborhood voxel has been classified. Assign the label of any one neighborhood voxel to the current voxel and make a note of all equivalent voxels i.e. all labels of voxels within this neighborhood is equivalent.

The time efficiency of this method is roughly $O(NS)$ where N is the total number of voxels in the input volume and S is the size of the structuring element. Space efficiency is linear in N because there can be at most N separate clusters if the structuring element is simply itself. This is a significant improvement over the computational efficiency of hierarchical methods as mentioned above. One potential drawback of this method however is its sensitivity to the size of the structuring element. Care must be taken to select a structuring element of the appropriate size in order not to mis-cluster unrelated distant voxels. By interpreting the size of the structuring element as the minimum allowable intra-cluster distance of the clustering algorithm, this provides us with a convenient intuition from which to determine an appropriate range of possible structuring element sizes. Distance between intra-cluster voxels of epileptic foci clusters should be small whilst 'noise' voxels should have larger inter-point distance. Consequently, by keeping this minimum intra-cluster distance criterion small, we are able to better cluster voxels of potential epileptic sites.

Finally, we make the observation that suspected FCD regions are usually contiguous while the detected hotspot label map often contains numerous small spurious holes within these hotspot regions due to noise. Hence, the final step in this asymmetry analysis section is to perform morphological post-processing i.e. dilation followed by erosion to close these small holes thereby smoothing the hotspot label map. This operation also decreases the false negative rate by filling in the holes in the hotspot label map.

3.4 Building the Prior Map

With the hotspots label map in place, we can now construct the prior probability map of epileptic hotspots. Note that each neuroimaging modality used in this pipeline i.e. T1-weighted MRI, PET, SPECT etc. should generate a separate hotspots label map detailing suspicious regions detected in the corresponding modality. This is achieved by co-registering each additional modality to the T1-weighted MRI reference volume and applying the methods as described in the pipeline stages above to each of these modality separately to obtain their corresponding hotspots label map. We then assign probability values as follows:

$$P_x(\text{FCD}) = \begin{cases} p_{\text{Hotspot}}, & \text{if } x \in \text{hotspot} \\ p_{\text{GM}}, & \text{if } x \in \text{GM} \\ p_{\text{WM}}, & \text{if } x \in \text{WM} \\ p_{\text{CSF}}, & \text{if } x \in \text{CSF} \\ p_{\text{Others}}, & \text{if } x \in \text{Any other tissue class} \end{cases} \quad (3.14)$$

where $P_x(\text{FCD})$ is the probability that voxel x is an epileptic FCD foci which is generating the observed surface EEG recordings used in EEG source localization. Every hotspot voxel within the detected hotspots label map is assigned the value p_{Hotspot} . Consequently, if a voxel is not lie within a hotspot region, it is then assigned a probability value according to its tissue class i.e. p_{GM} , p_{WM} , p_{CSF} and p_{Others} as obtained in the segmentation pre-processing stage.

This probability map is further smoothed by convolution with a Gaussian kernel across the entire volume to obtain the final hotspots prior probability map that is used to constrain the EEG source localization optimization problem. The purpose of this step is to produce a smooth prior probability map regardless of our initial choice of p_{Hotspot} , p_{GM} , p_{WM} , p_{CSF} , p_{Others} as well as to ensure that we do not exclude segmentation boundaries entirely from the EEG source localization search space i.e. prevent $P_x(\text{FCD}) = 0$ near boundary voxels.

3.5 Pipeline Summary

This section provides a summary of the pipeline stages that we have described in the preceding sections of this chapter. The flow of data within the pipeline can be summarized as follows:

- Rigidly align input modalities
- Intensity correction to remove spatial inhomogeneities in input volumes
- Anisotropic noise smoothing to reduce image noise while maintaining important edge boundaries.
- Perform tissue classification on structural MRI i.e. T1-weighted and T2-weighted modalities to obtain an isotropic patient-specific head model
- Obtain skull-stripped cortical volume by overlaying segmentation results and the desired input modality e.g. T1-weighted MRI, PET, SPECT etc.
- Align cortical volume with the target plane of symmetry
- Compute the deformation field to obtain a measure of inter-hemispheric asymmetry between the left and right hemisphere
- Cluster voxels of significant asymmetry to obtain the hotspots label map
- Assign probability values to voxels according to their classification based on the hotspots and segmentation label maps. Note that a hotspot voxel will always be assigned the hotspot probability value regardless of its tissue class.

Once these steps are complete, the patient-specific head model and prior probability map can be incorporated into the EEG source localization problem. The specifics of incorporating these two outputs into the formulation of EEG source localization is provided by Tiferet A. Gazit in [7]. The main thrust of this thesis is to provide a multi-modal framework to generate this patient-specific head model and a prior probability map.

Chapter 4

Results & Discussion

In this chapter, we shall present the results of running the quantitative analysis pipeline on numerous datasets of varying nature to demonstrate the clinical utility of the pipeline.

We begin by generating purely synthetic data with pathological ground truth. Section 4.1 provides a discussion of the results obtained from this dataset. Next, we test the pipeline on real clinical data involving T1-weighted MRI and T2-weighted MRI in Sections 4.2 and 4.3. In Section 4.4, we also present results that illustrate the robustness of the proposed plane of symmetry detector which underlies the entire asymmetry analysis pipeline.

To validate our results, we provide clinical evaluation in Section 4.5 from Dr. Richard Robertson, Director of Neuroradiology, Boston Children’s Hospital. Finally, a brief summary of the results from applying a prior probability map and a patient specific head model to EEG source localization is provided in Section 4.6.

4.1 Synthetic Data Test

4.1.1 Data Generation

Focal cortical dysplasia (FCD) in T1-weighted MRI acquisitions is often characterized by hypointense signals in the dysplastic lesion, thickening of the cortical ribbon with

abnormal gyral patterns or focal blurring of the gray-white matter interface. Any combination of the above 3 pathological subtypes may be exhibited in relation to a FCD lesional region. This shall be the basis of formulating the synthetic test data used in this section.

The fundamental dataset is composed of 2 simple noiseless ellipsoidal shells representing the skull-stripped cerebral volume. The outer gray shell represents cortical gray matter while the inner white ellipsoid represents cortical white matter. Hypocontrast, cortical thickening and interface blurring pathologies are subsequently added to this base dataset in the form of geometric distortions at user-specified locations in the volume. It should be noted that the ellipsoids generated in this dataset are of the form $\frac{X^2}{R_x^2} + \frac{Y^2}{R_y^2} + \frac{Z^2}{R_z^2} \leq 1$ and the 3 geometric deformations are applied in this frame. Rotation and translation of this dataset are added at a later stage before the data is fed into the plane of symmetry detector. Figure 4-1 shows this final synthetic dataset with rotation and translation applied.




Figure 4-1: Synthetic Data. Outer gray shell represents cortical gray matter while the inner white ellipsoid represents cortical white matter. Left: Hypocontrast lesion. Middle: Cortical Thickening. Right: Interface Blurring.

Hypocontrast lesional matter in T1-weighted MRI volumes often exhibits lower signal intensity relative to surrounding white matter tissue and hence appears as a darker shade of gray. From a user-specified point, we average the intensity of voxels within a certain radius and a scaling factor is applied to this average. The intensity value at these voxels are subsequently replaced with this scaled average to generate

a region of hypointense signal whose base value is determined by original local voxel intensities. Cortical thickening is generated by growing a user-specified plane along its normal. This plane should be chosen such that it follows the contour of gray-matter gyrus to generate a realistic thickening pattern. The intensity value of each “new” cortical voxel follows that of the corresponding voxel whose normal passes this “new” voxel. If multiple normals cross an individual “new” voxel, its final intensity is the average of its parent’s intensity value. Interface blurring is generated in a similar fashion. The only exception is that the final intensity value of a “new” voxel is a linearly decaying function of distance from the user-specified plane and the original intensity value of the corresponding voxel on that plane.

The following rotation and translation parameters were applied on the above dataset about the mid-sagittal plane to produce the final synthetic dataset that will be used for later analysis:

- Rotation: $\theta_{x_0} = 10^\circ$, $\theta_{y_0} = 5^\circ$, $\theta_{z_0} = 10^\circ$
- Translation: $T_{x_0} = -5$, $T_{y_0} = 3$, $T_{z_0} = 2$

No further preprocessing, i.e. multi-modal rigid registration and segmentation, is necessary because there is only 1 input modality in this example and skull-stripping is only applicable to real patient data.

4.1.2 Plane of Symmetry Detection

The plane of symmetry detector seeks to find the rigid transformation that best matches the synthetic volume to its chiral volume. For the purpose of this pipeline, we choose the y-z plane as the target plane of symmetry. The chiral volume is generated by reflecting the synthetic volume about this plane.

The following rotation and translation parameters were detected by the plane of symmetry detector:

- Quaternion Rotation: 5.15° about $(0.011, -0.93, 0.37)$




Figure 4-2: Synthetic Data - Plane of Symmetry. The gray line represents the target plane of symmetry.

- Translation: $T_x = 5.86$, $T_y = 0.46$, $T_z = 1.01$

The results of aligning the synthetic volume to the plane of symmetry is shown in Figure 4-2. Since we desire alignment about the y - z plane, we are only concerned with matching T_x to T_{x_0} . In this case, computed T_x is 5.86 which is sufficiently close to the T_{x_0} value of 5. The values of T_y and T_z have no effect on the results since they merely represent a shift along the y - z plane i.e. the target plane of symmetry. The same argument goes for rotation. The recovered quaternion axis of rotation is largely contained within the y - z plane. Hence, the quaternion angle of rotation should match that of θ_{y_0} since the chiral volume to be matched was generated by reflecting about this y - z plane. Similarly, computed angle of rotation is 5.15° which is sufficiently close to the θ_{y_0} value of 5° . Consequently, mean-squared-difference between the original synthetic volume and its chiral volume decreased significantly from 1488.8 to 71.0 after alignment with the plane of symmetry. The aligned volume shows a high degree of symmetry about the mid-sagittal plane and this can be easily seen in in Figure 4-2.

4.1.3 Asymmetry Analysis

Asymmetry analysis begins by finding the optimal deformation field that maps the symmetrically aligned volume from the previous step to its chiral volume. This de-

formation field in essence allows us to infer the inter-hemispheric differences between the left and right hemispheres of the synthetic volume i.e. the synthetic pathologies that we have added. The goal is to compute a deformation field that best warps the chiral volume back to the symmetrically aligned volume. Interpretation of the deformation field is as follows: The displacement vector of the current voxel gives the relative location of the target point in physical space. Deformation then involves “moving” the target point to the location of the current voxel. Figure 4-3 shows a 3D plot of this deformation field overlaid on the chiral image of the symmetrically aligned volume. Hence, by analyzing the resulting deformation vector at each voxel, we can infer the amount of change that is required to map the chiral volume back to the symmetrically aligned volume and consequently the degree of inter-hemispheric difference at that voxel.




Figure 4-3: Synthetic Data - Deformation Field. Displacement vector shown gives the location of the target point in physical space. Deformation involves “moving” the target point to the location of the corresponding voxel.

After the deformation field is computed, it needs to be converted to a scalar asymmetry measure in order to determine voxels of significant asymmetry to be clustered. As discussed in Section 3.3, the vector field operator $|F|(C + \nabla \cdot F)$ shall be used to convert the deformation field to a scalar asymmetry measure. Magnitude of the deformation field, $|F|$, is shown in Figure 4-4; divergence of the deformation field, $\nabla \cdot F$, is shown in Figure 4-5; asymmetry measure, $|F|(C + \nabla \cdot F)$, is shown in Figure 4-6. After the deformation field is converted a scalar asymmetry measure, a threshold is applied to extract only voxels of significant asymmetry and these voxels shall form




Figure 4-4: Synthetic Data - Deformation Field Magnitude.




Figure 4-5: Synthetic Data - Deformation Field Divergence.




Figure 4-6: Synthetic Data - Asymmetry Measure.

the set of binary input points that will be clustered in the next step.

Before we proceed further however, we shall now take a closer look at the deformation field associated with each of the three simulated target FCD pathologies to provide detailed analysis and motivation for using the resulting deformation field to quantify inter-hemispheric asymmetry in the cortical volume (see Figure 4-7, 4-8 and 4-9). Legend of the plots are as follows: Blue dots represent the source voxel to which the target point will be moved. Since these are 2D plots of 3D vectors, each vector is colored by the its z-coordinate i.e. yellow represents out of the plane, red represents into the plane and orange represents in the current plane.

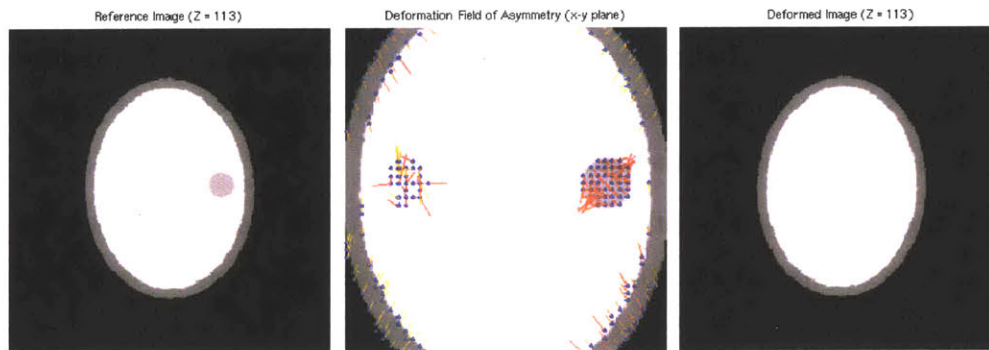


Figure 4-7: Synthetic Data - Hypocontrast Deformation. Blue dots represent the source voxel to which the target point will be moved. Since these are 2D plots of 3D vectors, each vector is colored by the its z-coordinate i.e. yellow represent out of the plane, red represents into the plane and orange represents in the current plane. Left: Chiral Image of the Symmetrically Aligned Volume. Middle: Deformation Field. Right: Final Deformed Volume.

The deformation field associated with the hypocontrast pathology is illustrated in Figure 4-7. The left reference image shows the chiral image of the symmetrically aligned volume while the right image shows the deformed volume after the deformation field has been applied. This final deformed image should ideally be a reflection of the reference image about the mid-sagittal plane i.e. the plane of symmetry for perfect non-rigid registration.

Deformation of the WM voxels surrounding the hypocontrast lesion in the right hemisphere shrinks inwards to envelope the lesion. In this case, the white-matter

(WM) target points are from out of the plane as indicated by the red diverging vectors emanating from within lesional source voxels i.e. the blue dots. From the final deformed image, we observe that the deformation field is able to “remove” the hypocontrast lesion on the right hemisphere. It is however unable to reproduce the lesion on the contralateral hemisphere as seen by the lack of the lesion on the left hemisphere of the final deformed image. A lack of hypointense seed voxels in the vicinity of the potential hypocontrast lesion in the left hemisphere makes it impossible for the non-rigid registration algorithm to obtain a deformation map that creates this hypocontrast lesion. Consequently, the deformation field on the side of the hemisphere without the hypocontrast lesion provides little information about its asymmetry with its contralateral hemisphere owing to the relatively random nature of the displacement vectors in that region. It merely shifts WM target voxels to the WM source voxels which has no impact on the deformation process.

Despite this being a futile attempt on the part of the non-rigid registration algorithm to deform the left hemisphere, fortunately the set of source voxels in the left hemisphere is a subset of the source voxels on the contralateral side. As such, while this subset of source voxels adds little asymmetry information, it does not adversely affect the deformation detected on the contralateral side. A proposal is subsequently made later in this section to completely relax the left / right ambiguity assumption to merge the set of source voxels from the opposing hemisphere to achieve lower false negative results. Since the set of source voxels in the left hemisphere is a subset of that in the right hemisphere, it does not adversely affect the output from this merge operation.

The deformation field associated with cortical thickening is illustrated in Figure 4-8. Displacement vectors for the left hemisphere show that the gray-matter (GM) patch is thickened by shifting the corresponding target point from within the original cortex towards the blue source voxels. On the right hemisphere, original thickening is reduced by shifting WM voxels from out of the plane (as indicated by the yellow displacement vectors) to the blue source voxels. The deformation field between the 2 opposing hemispheres is clearly not symmetrical i.e. displacement vectors are in




Figure 4-8: Synthetic Data - Cortical Thickening Deformation. Blue dots represent the source voxel to which the target point will be moved. Since these are 2D plots of 3D vectors, each vector is colored by the its z-coordinate i.e. yellow represent out of the plane, red represents into the plane and orange represents in the current plane. Left: Chiral Image of the Symmetrically Aligned Volume. Middle: Deformation Field. Right: Final Deformed Volume.

different directions on opposing hemispheres. The deformation field on one side shifts the GM voxels out while that on the other hemisphere shifts the WM voxels out to cover the original thickening. Nonetheless, the set of source voxels i.e. the blue dots are approximately the same in both hemispheres.

The deformation field associated with interface blurring is illustrated in Figure 4-9. In the right hemisphere where there is no blurring in the reference image, gray-matter (GM) source voxels are deformed from white-matter (WM) voxels that are into the plane while WM source voxels are deformed from the cortical GM voxels that are out of the plane to create a blurring effect in the final deformed image. In the left hemisphere where there is blurring in the reference image, blurred voxels are deformed from the corresponding GM and WM voxels to achieve a distinct interface again. Consequently, it is clear from the deformation plot that the set of source voxels (i.e. the blue dots) on the left hemisphere is different from that of the right hemisphere. The source voxels are more evenly spaced out in the left hemisphere whereas those in the right hemisphere are separated into 2 distinct bands (this becomes much clearer in Figure 4-6). In fact, the actual diseased interface blurring region is a union of the 2 opposing sets of source voxels.

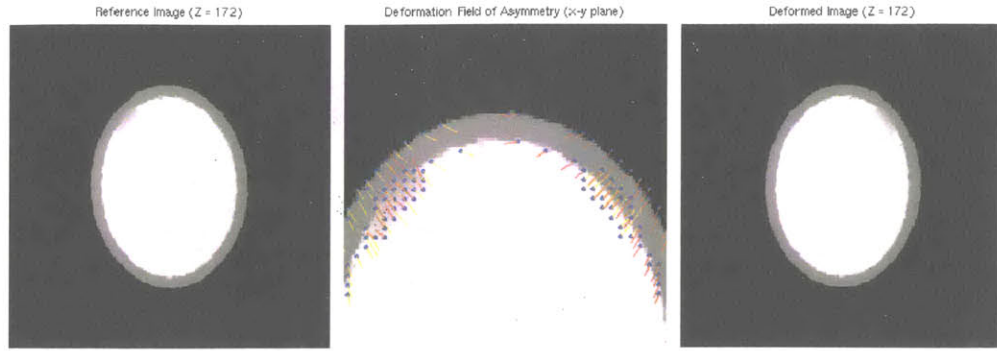


Figure 4-9: Synthetic Data - Interface Blurring Deformation. Blue dots represent the source voxel to which the target point will be moved. Since these are 2D plots of 3D vectors, each vector is colored by the its z-coordinate i.e. yellow represent out of the plane, red represents into the plane and orange represents in the current plane. Left: Chiral Image of the Symmetrically Aligned Volume. Middle: Deformation Field. Right: Final Deformed Volume.

From the above analysis, it is clear that the non-rigid registration algorithm proposed has done a great job in its goal of mapping the symmetrically aligned volume to its chiral and in so doing generating a deformation map representing the contralateral asymmetries present in the volume. If the final deformed image did not match the chiral of the reference image (i.e. the non-rigid registration algorithm has failed), the resulting deformation map would be rendered useless. In our case, except for the hypocontrast pathology, the final deformed image closely matches the chiral image of the reference image indicating successful registration and a valuable deformation map. As for the hypocontrast pathology, we observe that non-rigid registration was successful only in the hemisphere where the hypointense region was present in the reference image. At the same time, we also observe that the set of source voxels on the opposing side is contained within the bounds of the potential hypointense region despite its relatively random magnitude and direction. This is entirely justifiable because the non-rigid registration algorithm is merely attempting to modify the set of voxels within the bounds of the potential hypointense region, with little success.

Consequently, we propose to construct a binary clustering map that combines deformation from both left and right hemispheres. An initial binary map is constructed




Figure 4-10: Synthetic Data - Merged Binary Clustering Map.

from voxels of significant asymmetry. This initial binary map shall be known as the simple binary clustering map. This 3D binary map is then combined with its chiral map that is derived by reflection about the mid-sagittal symmetry plane to obtain the final binary clustering map which shall be known as the merged binary clustering map (see Figure 4-10). This merged clustering map solves the problem of non-rigid registration being successful only in one of the hemisphere as in the case of the hypocontrast lesion demonstrated in this example. In addition, this merged map also serves to reduce false negative rates by providing more comprehensive coverage of the actual diseased regions as required in the case of interface blurring. In making this left / right ambiguity assumption, the pipeline is able to pick out regions of pathology more effectively than simply considering each hemisphere in its individual confines.

With this in mind, the last step in the asymmetry analysis stage is clustering to aggregate clumps of significant asymmetry voxels and to delineate them into separate regions of potential epileptic foci hotspots. Morphological post-processing i.e. dilation followed by erosion is applied after the clustering process to close any small islands within the detected hotspot clumps thereby smoothing the resulting hotspots label map. In this section, we present 2 different binary maps for clustering, namely the original simple binary clustering map and the merged binary clustering map. In this experiment, we also evaluate the effect of performing morphological post-processing on the hotspot label map derived from the clustering process. Table 4.1 tabulates the

results of the experiment while Figure 4-11 shows the final hotspots label map which is computed based on the merged binary map with morphological post-processing.

	Binary Map	Merged Binary Map	Closed Binary Map	Closed Merged Binary Map
Ground Truth Count	5,685		6,758	
True Positive Count	3,916	5,170	6,673	6,753
True Negative Count	16,725,193	16,693,549	16,694,368	16,640,677
False Positive Count	46,338	77,982	76,090	129,781
False Negative Count	1,769	515	85	5
False Positive Rate	0.00276	0.00465	0.00454	0.00774
False Negative Rate	0.311	0.0906	0.0126	0.000740

Table 4.1: Synthetic Data - Clustering Results. Binary Map refers to the original simple binary clustering map. Merged Binary Map refers to the binary map obtain by merging the deformation from both left and right hemispheres. Closed Binary Map and Closed Merged Binary Map simply refers to the former 2 binary maps with morphological post-processing applied after clustering.




Figure 4-11: Synthetic Data - Hotspots Label Map (Closed Merged). This hotspots label map was clustered from the merged binary clustering input map and morphological post-processing was subsequently performed.

If we consider the simple binary map case, the false negative rate is relatively high at 31% mainly due to “holes” in the hotspots label map detected. By simply applying mathematical morphological operations, i.e. dilation followed by erosion of the hotspot region, we can eliminate small holes in the hotspot region thereby smoothing the hotspots label map. The false negative rate instantly drops to 1.26% indicating that there were many small holes in the simple binary map case and these

holes can be easily closed. More importantly, this low false negative rate shows that the asymmetry analysis pipeline did a good job in localizing the simulated FCD pathologies and closing these holes simply made that result more obvious. By merging the deformation from both left and right hemispheres i.e. the merged binary clustering map, false negative rates drop from 31% to 9.06% without any further processing and gets as low as 0.074% with morphological post-processing.

In summary, we have demonstrated in this example the efficacy of this pipeline to accurately detect FCD pathologies with false negative rates reaching as low as 0.074%. This provides us with an excellent basis from which to evaluate real patients with real FCD pathologies.

4.2 Patient Example 1 - Focal Cortical Dysplasia in the Right Inferior Frontal Lobe

4.2.1 Multi-modal Rigid Volume Correspondence

In this section, we evaluate the quantitative pipeline on real patient data. The patient in this case suffers from focal cortical dysplasia in the right inferior frontal lobe as given by her existing clinical workup. High resolution T1-weighted MRI and T2-weighted MRI scans are available as input into the pipeline. These inputs allow us to determine regions of structural asymmetry that might give rise to epileptic seizures. It should be noted that as additional imaging modalities are added to this patient's workup, we can process these modalities in the same manner to extract new potential regions of epileptic foci or reinforce existing structural ones that we are currently detecting. For example, PET / SPECT imaging can give us information about the functional metabolism asymmetries in this patient. Significant asymmetries present in these modalities indicate possible epileptic foci as well, thus adding valuable information to the current structural asymmetry hotspots prior map. In turn, EEG source localization will benefit greatly from this as it fuses both a structural and functional prior to guide its search.

Since we are dealing with multiple modalities, the first step in this pipeline is to bring all the modalities into alignment. For the purpose of this pipeline, we choose the T1-weighted volume as the reference volume to which every other volume will be aligned. The mutual information measure as described in Section 3.1.1 is used to optimize the degree of match between the T1-weighted reference volume and the target volume. Figure 4-12 shows the aligned T1-weighted and T2-weighted volumes.




Figure 4-12: Patient Example 1 - Multi-modal Rigid Registration. Top Row: T1-weighted slices. Bottom Row: Corresponding T2-weighted slices aligned to the T1-weight volume as shown in the top row.

4.2.2 Tissue Segmentation

Before performing tissue segmentation, two pre-processing steps can be used to enhance the segmentation results, namely intensity correction and noise smoothing. Intensity correction aims to remove any intensity inhomogeneities present in both the T1-weighted and T2-weighted volumetric MRI acquisitions. Removing these intensity

inhomogeneities allows us to assert the intra-class intensity homogeneity assumption necessary for subsequent tissue segmentation. Figure 4-13 shows the results of this intensity correction step. Alternating bright and dark bands in the original T1-weighted and T2-weighted volumes are removed except for the posterior cortex which consistently shows a significantly bright inhomogeneity artifact before and after the correction. This persistent artifact is actually characteristic of the scanner's head coil at the time of image acquisition and the removal of this effect is still an on-going effort. Fortunately for us, we can safely ignore the posterior cortex because there is no pathology in or near that section of the brain for this patient; the pathology is in the frontal lobe.




Figure 4-13: Patient Example 1 - Intensity Inhomogeneity Correction. Top Row: Intensity Corrected T1-weighted volume. Bottom Row: Intensity Corrected T2-weighted volume.

Another important step which greatly enhances segmentation is that of anisotropic smoothing. By smoothing the input volume while preserving important image edges, we can significantly reduce image noise and improve the performance of tissue segmen-

tation that follows next. Figure 4-14 shows the results of anisotropic noise smoothing. Improved edge boundaries allow the segmentation algorithm to better delineate between different tissue types and also reduces the high noise sensitivity of watershed based segmentation methods.




Figure 4-14: Patient Example 1 - Anisotropic Smoothing. Top Row: Edge-preserving Smoothed T1-weighted volume. Bottom Row: Edge-preserving Smoothed T2-weighted volume.

Once the T1-weighted and T2-weighted volumes have been pre-processed, we can proceed with tissue segmentation. Since the T1-weighted volume provides clear gray-matter (GM) white-matter (WM) delineation while the T2-weighted volume provides good cerebral-spinal-fluid (CSF) contrast, we can improve multi-class tissue segmentation by combining these 2 modalities. By considering this as a vector-valued problem, we can obtain a more accurate anatomical tissue classification with the follow tissue classes: skin, skull, CSF, GM and WM.

Tissue segmentation begins with seed voxel selection. Representative slices are chosen for supervised 2D segmentation. Seed voxels are selected as per the process




Figure 4-15: Patient Example 1 - Tissue Segmentation. Legend: Skin(Purple), Skull(White), CSF(Pink), GM(Blue), WM(Yellow).

described in Section 3.1.2. Once we have satisfactory 2D segmentation in these slices, we extend the segmentation to 3D with the combined set of seed voxels. Figure 4-15 shows the resulting tissue label map. In addition to the steps described above, we also pursued a hierarchical segmentation model. The above steps were first performed for two broad classes, namely Head or Background. This is to remove any background noise that is captured by the scanner. The head volume is extracted and this is the input volume for the next level. This level is also more commonly known as skull-stripping because the two desired classes are Brain and Non-brain. Skin, skull and other non-cortical tissues are classified as Non-brain while CSF, GM and WM are seeded as Brain. Finally, the last level involves detailed classification of the Brain volume into CSF, GM and WM tissue types. One observation is that the CSF layer in the T2-weighted volume is very faint and thin. This is usually the case in young adolescent patients where their growing brain presses against the skull bone. Consequently, the CSF layer between the cortex and the skull is very thin and this does not show up well given the spatial resolution of the MRI acquisition. Nonetheless, we are still able to obtain a very detailed tissue label map from the T1-weighted and T2-weighted MRI volumes and this can be used to construct a patient-specific head model for EEG source localization.

4.2.3 Plane of Symmetry Detection

With the segmentation label map, the entire pre-processing stage of the pipeline is complete. The next stage is suspicious region detection and given our input modalities, the goal is to detect structural asymmetries present in the patient's cortex. As such, we are only interested with the skull-stripped T1-weighted volume due to its excellent gray-matter (GM) and white-matter(WM) signal response. To obtain this volume, we simply apply the binary label map from the skull-stripping level in the segmentation process to the T1-weighted volume and the plane of symmetry detector is applied on this new cortex-only volume. Figure 4-16 shows the symmetrically aligned brain volume after the plane of symmetry detector has been applied.




Figure 4-16: Patient Example 1 - Plane of Symmetry. The gray line represents the target plane of symmetry.

The goal of the symmetry plane detector is to align the plane in this volume that maximizes inter-hemispheric similarity to a target symmetry plane by finding the rigid transformation that best matches the input brain volume to its chiral volume. In cases where pathologies are relatively subtle such as this, the plane of maximal inter-hemispheric similarity often coincides with the mid-sagittal plane. As we can see in Figure 4-16, this plane of maximal inter-hemispheric similarity is in fact the mid-sagittal plane through the anterior and posterior commissure. Consequently, the mean-squared-difference between the original volume and its chiral volume decreased significantly from 1251.1 to 432.1 after alignment with the plane of symmetry. The resulting symmetrically aligned volume shows a high degree of symmetry about the

target symmetry plane.

4.2.4 Asymmetry Analysis

Asymmetry analysis proceeds by finding the optimal deformation field that best matches the symmetrically aligned brain volume with its chiral volume. Mean-squared-difference between the brain volume and the resulting deformed chiral volume falls from 432.1 to 18.6 after deformation. This significant drop in mean-squared-difference is indicative that the resulting deformation field is an accurate representation of the inter-hemispheric asymmetry present in this patient’s brain. Figure 4-17 shows the magnitude of the deformation field while Figure 4-18 shows the divergence of the deformation field. Relatively large deformation in the frontal lobe is indicative that the pipeline is indeed accurately picking up potential FCD spots of significant asymmetry.




Figure 4-17: Patient Example 1 - Deformation Field Magnitude.

Given the noisy nature of the input MRI images, we perform a simple threshold on the deformation field magnitude to remove noise voxels characterized by relatively small displacement. A threshold of 2 voxels was chosen i.e. voxels with a displacement of less than 1.8mm. Our asymmetry measure then is computed using the vector field operator $|F|(C + \nabla \cdot F)$. Figure 4-19 shows the resulting asymmetry measure.

At this point, we have a map of asymmetries present in the patient’s brain volume, both healthy and pathological. We then make the assumption that asymmetry for




Figure 4-18: Patient Example 1 - Deformation Field Divergence.




Figure 4-19: Patient Example 1 - Asymmetry Measure.

FCD pathology is relatively larger than that of healthy asymmetry. This allows us to further threshold the volume on the asymmetry measure to isolate only voxels of significant asymmetry and classify those as epileptic hotspot regions. If we were to relax this assumption, i.e. in the worst case this would equate to no further thresholding on the asymmetry measure, we would get a final volume of asymmetry measure that details both healthy and pathological asymmetries. Given that our premise is to construct a prior probability map of epileptic hotspots, this inclusion of healthy asymmetries will not adversely hurt our final EEG source localization solution; it simply allows a larger search space. Nonetheless, this prior probability map will still improve overall EEG source localization performance because the search is still constrained only to asymmetrical areas in the brain and not the entire cortical volume. For a first pass, mean asymmetry i.e. asymmetry sizing factor, $S = 0$ is a good threshold value to use. If the clinician is expecting large FCD pathologies, we can set the asymmetry threshold higher with a larger S and if FCD pathologies are expected to be small, we can even set a negative S so that the threshold is smaller than the mean.




Figure 4-20: Patient Example 1 - Thresholded Asymmetry Measure. Thresholding significantly removes the number of asymmetrical voxels as compared with Figure 4-19.

For the purpose of this example, we shall use $S = 2$ since we know that the patient has rather extensive FCD pathologies. Figure 4-20 shows the thresholded asymmetry volume. Given the non-symmetrical nature of the deformation field, high asymmetry values exhibited in the inferior frontal region is consistent with the clinical diagnosis




Figure 4-21: Patient Example 1 - Merged Binary Clustering Map.

of FCD in the right inferior frontal lobe. As discussed in Section 4.1.3, we can use this thresholded asymmetry volume to construct a binary clustering input map that combines significant asymmetry from both the left and right hemispheres. An initial binary map is derived from the thresholded asymmetry volume. This 3D binary map is then combined with its chiral map that is derived by reflection about the mid-sagittal symmetry plane to obtain the final binary clustering map. Figure 4-21 shows the final binary clustering map.

We shall use this merged binary map as input into the last step of the asymmetry analysis stage of the pipeline i.e. clustering. The purpose of clustering is to aggregate clumps of significant asymmetry voxels and to delineate them into separate regions of potential epileptic foci hotspots. Morphological post-processing i.e. dilation followed by erosion is applied after the clustering process to close any small islands within the detected hotspot clumps thereby smoothing the resulting hotspots label map. Hence we shall use a merged binary map with morphological post-processing to obtain optimal performance as discussed in Section 4.1.3. In addition, we also threshold the final hotspots volume based on a minimum cluster size constraint. Once again, the intuition behind this is to remove relatively small noise clusters with the assumption that lesional clusters are relatively larger than these small noise clusters. Cluster size thresholding removes a significant number of small noise clusters while keeping larger lesional clusters intact. Similarly for a first pass, mean cluster size is a good threshold value to use. A more stringent threshold condition, $\mu + N\sigma$, can be used if there is prior information regarding the size of the FCD pathology in the patient. If the clinician expects the diseased region to be fairly large from preliminary workup, we can increase N to help reduce noise in the final output. On the other hand, if we are unsure of the size of the FCD pathology or have no prior knowledge, keeping a small N helps ensure that we do not exclude any small FCD pathologies and thus reduces false negative rates. Having said that, the focus of this step should not be in choosing an extremely tight bound on the threshold level to remove all noise / healthy asymmetry clusters. Instead, care must be taken to select a modest threshold level which only filters away small obvious noise clusters while leaving the larger ambiguous ones unaltered. In

this example, $N = 4$ was empirically chosen because of the exponential nature of the resulting cluster size distribution and our knowledge from previous workup that this patient suffers from relatively extensive FCD pathology. Figure 4-22 shows the cluster size distribution plot from the clustering of voxels with significant asymmetry. Notice the exponential nature of the distribution i.e. there are many small clusters but the number of larger clusters drop off very rapidly. Coupled with our prior clinical knowledge that the FCD pathology in this patient is relatively extensive, we expect the FCD pathologies to correspond to those larger clusters. Consequently, this allows us to set a larger N value that effectively prunes away many more non-pathological clusters.

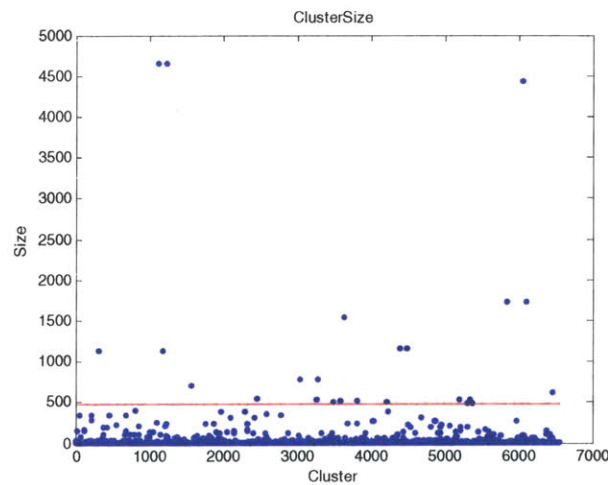


Figure 4-22: Patient Example 1 - Cluster Size Distribution. Note the exponential cluster size distribution where there are significantly more small clusters i.e. less than 100 while the large ones only number of a few. The red horizontal line gives the minimum cluster size threshold for this example.

Finally, Figure 4-23 shows the hotspots label map which is computed based on the merged binary map with morphological post-processing and a minimum cluster size constraint. Evaluation of this hotspots label map was through expert validation by Dr. Richard Robertson, Director of Neuroradiology, Boston Children’s Hospital (see Section 4.5 for a summary of his analysis). The hotspots detected by this pipeline are in agreement with the clinical diagnosis of FCD in the right inferior frontal lobe.




Figure 4-23: Patient Example 1 - Hotspots Label Map (Closed Merged). Top Row: 3D Hotspots Label Map. Middle Row: Symmetrically Aligned Cortical Volume. Bottom Row: Cortical Volume overlaid with the Hotspots Label Map.

The two particular axial slices shown in Figure 4-23 were chosen because they clearly illustrate the two main groupings of FCD pathology that were present in this patient and were accurately isolated by asymmetry analysis.

In summary, we have demonstrated in this example the effectiveness of this pipeline to accurately detect FCD pathologies in real patient data that concurs with current clinical diagnosis of the patient's condition. It should be noted that the hotspots label map generated by this pipeline will always be symmetric due to a left / right ambiguity that is inherent in the design of the system. Nonetheless, the main objective of this pipeline is not to provide a definitive lesion detector. Instead, the goal of asymmetry analysis is to highlight potential spots of cortical abnormalities so as to bias the solution of EEG source localization towards these known anomalous regions that are more likely to contain epileptic foci. Hence, it is up to the EEG source localization algorithm to determine which of the two hemispheres actually contain the epileptic foci that is generating the observed EEG measurements.

4.2.5 Prior Map Generation

For the purpose of EEG source localization, we now have to convert the hotspots label map from the previous step into a prior probability map as discussed in Section 3.4. Probability values were chosen as follows:

$$p_{Hotspot} = 0.9$$

$$p_{GM} = 0.2$$

$$p_{WM} = 0.1$$

$$p_{CSF} = 0.01$$

$$p_{Others} = 0$$

The rationale for the above choice of probabilities is to reflect the relative like-

likelihood of FCD pathologies in the corresponding cortical region. Given the hotspots label map which reflects regions of significant cortical asymmetry, we assign the highest probability of FCD pathology to voxels with this label. However, we would also like to include all gray-matter (GM) and white-matter (WM) regions because they may possibly contain epileptic foci sites that were not picked up by the asymmetry pipeline. Since GM and WM voxels are also included from a non-exclusion perspective i.e. we do not want to completely exclude non-hotspot voxels, the probability values for GM and WM should be several orders of magnitude smaller than that for hotspots. Given the nature of FCD pathologies, it is also more likely to find these diseased regions in GM tissue than WM tissue. In principle, it is not possible to find FCD pathologies in CSF but we still assign a relatively small probability to CSF voxels because of partial volume effects in MRI acquisition amongst other reasons that lead to slight segmentation inaccuracies. This is also done from a non-exclusion perspective. Finally, all non-brain, skull and background voxels are completely excluded from the solution set with probability zero. This resulting probability map is smoothed across the entire volume to obtain the final hotspots prior probability map as shown in Figure 4-24. Together with the patient-specific head model computed earlier in the pipeline, this prior probability map can be used to improve EEG source localization by biasing its solution towards cortical abnormality as highlighted by this prior probability map.

4.3 Patient Example 2 - Focal Cortical Dysplasia in the Right Superior Temporal Lobe

In this section, we shall further demonstrate the utility of the pipeline to detect potential FCD hotspots in structural neuroimaging. The nature of cortical asymmetries for the patient in this case differs from the example in Section 4.2 in that healthy normal asymmetries are large relative to pathological asymmetries. Consequently, these healthy asymmetries are more resistant to pruning and the resulting hotspots label




Figure 4-24: Patient Example 1 - Hotspots Prior Probability Map. Red regions represent areas of high FCD probability while purple regions represent areas of low FCD probability. The black background represent areas of zero probability thus excluding the background from the EEG source localization solution set because the FCD must lie within the head volume.

map contains a fair amount of non-pathological voxels with significant asymmetries.

Nonetheless, this is still indicative of success for the pipeline (see Section 4.5) as long as the pipeline is still able to pick out the FCD regions accurately i.e. maintaining a low false negative rate. As a first-pass filter to identify potential cortical abnormalities, it is perfectly reasonable to include significant healthy asymmetries at this point because these regions may be suggestive of possible subtle cortical abnormalities that are currently overlooked by visual inspection and may in fact contain additional unexpected epileptic foci. Given the overall structure of our system as well as the purpose of identifying these hotspots as a prior probability map for EEG source localization, it is in fact important to include all voxels of significant asymmetry in order to ensure better coverage of any cortical malformations present in the patient.

4.3.1 Tissue Segmentation

Since we only have T1-weighted MRI acquisition for this dataset, no multi-modal alignment is necessary. Hence, the pipeline starts off with tissue segmentation. In this example, segmentation will only be based on the T1-weighted volume since there is only one structural imaging modality available for this case. Nevertheless, T1-weighted MRI images can still provide excellent gray-matter (GM) white-matter(WM) delineation. The exception is that cerebral-spinal-fluid (CSF) does not show up well on T1-weighted MRI acquisitions. This may marginally affect segmentation performance especially when constructing the brain mask because CSF lines the cortical volume for the most part and the inability to detect this accurately may result in a brain mask that only captures CSF regions partially. For the purpose of constructing a patient-specific head model, this may prove problematic. However, we are only interested in analyzing the performance of the asymmetry analysis pipeline for this example and the ability to accurately delineate GM and WM tissues from the entire head volume is sufficient for our case. We shall not be extending this example to building a detailed segmentation label map for patient-specific head modelling or an actual prior probability map. The final output for this example is a hotspots label map identifying potential FCD regions.

The tissue segmentation process begins by pre-processing the input T1-weighted MRI volume through intensity correction and anisotropic noise smoothing. Intensity correction allows us to better assert the intra-class intensity homogeneity assumption for accurate tissue segmentation while anisotropic noise smoothing reduces the sensitivity of the segmentation algorithm to noise while maintaining important image edge features. Figure 4-25 shows the intensity corrected T1-weight MRI volume while Figure 4-26 shows the corresponding noise smoothed T1-weighted MRI volume.




Figure 4-25: Patient Example 2 - Intensity Inhomogeneity Correction.




Figure 4-26: Patient Example 2 - Anisotropic Smoothing.

As previously mentioned, tissue segmentation based on T1-weighted images alone does not isolate CSF regions well as is evident in Figure 4-27. The resulting brain mask boundary follows a tight bound with the gray-matter gyral patterns. This is because CSF in the noise-smoothed T1-weighted volume shows up as a significantly darker




Figure 4-27: Patient Example 2 - Brain Mask. Legend: Non-brain(White), Brain(Blue). The figure shows the resulting brain mask from initial segmentation overlaid on the corresponding T1-weighted MRI volume. Further processing to separate the cortical volume into gray-matter, white-matter and cerebral-spinal-fluid is not carried out in this example.

shade of gray as compared to GM tissue. In fact, this darker shade of gray is much closer to that of the surrounding skull bone as compared to regular gray-matter. As such, some CSF voxels may get misclassified as skull. Granted that building this brain mask is an intermediate step in the hierarchical segmentation framework, we would like this brain mask to incorporate CSF regions for further fine-grained segmentation and hence the value of a vector-based segmentation as shown in Section 4.2 with both T1-weighted and T2-weighted MRI input volumes. Fortunately for our case in this example, we do not require detailed CSF segmentation for asymmetry analysis since we're only concerned with building a FCD hotspots label map. The brain mask shown in Figure 4-27 adequately provides for finding significant asymmetries within the cortical volume i.e. within gray-matter and white-matter tissue structures.

4.3.2 Plane of Symmetry Detection

The brain mask from the segmentation process in the preceding section allows us to perform skull-stripping on the input T1-weighted MRI volume to obtain a new cortex-only volume which shall be the target volume of subsequent asymmetry analysis. The first step in this stage is to align the plane of maximal inter-hemispheric similarity in this new cortical volume to a target symmetry plane. Figure 4-28 shows the

resulting volume after symmetry alignment. Clearly, we can tell by visual inspection alone that the plane of maximal inter-hemispheric similarity coincides perfectly with the mid-sagittal plane of the cortical volume. Quantitatively, the mean-squared-difference between the original volume and its chiral volume only dropped from 57.9 to 32.4 after alignment. The relatively small difference in mean-squared-difference value after alignment is indicative that only minor adjustments were necessary to symmetrically align the cortical volume. A low final mean-squared-difference value shows that the left and right hemispheres are largely in agreement, save for the healthy and pathological asymmetries which we will attempt to detect in the next step.




Figure 4-28: Patient Example 2 - Plane of Symmetry. The gray line represents the target plane of symmetry.

4.3.3 Asymmetry Analysis

We begin asymmetry analysis by finding the optimal deformation field that best matches the symmetrically aligned cortical volume with its chiral volume. The mean-squared-difference value between the aligned cortical volume and the resulting deformed chiral volume drops further from 32.4 to 1.77 after the deformation field is applied. This low mean-squared-difference value is indicative that the resulting deformation field accurately accounts for most of the inter-hemispheric asymmetries present in this cortical volume. An otherwise high mean-squared-difference value would suggest that there are still significant asymmetries present in the cortical volume that the deformation field is not capturing resulting in a deformed chiral volume

that still significantly differs from the aligned cortical volume. Figure 4-29 shows the magnitude of the resulting deformation field while Figure 4-30 shows the divergence of the deformation field.




Figure 4-29: Patient Example 2 - Deformation Field Magnitude.




Figure 4-30: Patient Example 2 - Deformation Field Divergence.

Asymmetry measure is computed from the deformation field with the vector operator $|F|(C + \nabla \cdot F)$ i.e. voxel-wise multiplication of the volumes shown in Figure 4-29 and 4-30. Deformation magnitude and divergence both present with high values in the superior temporal region which is consistent with the clinical diagnosis of FCD in the right superior temporal lobe. Consequently, the resulting asymmetry measure which is a combination of deformation magnitude and divergence further accentuates the superior temporal region from the rest of the normal cortical volume as shown in Figure 4-31.




Figure 4-31: Patient Example 2 - Asymmetry Measure.




Figure 4-32: Patient Example 2 - Thresholded Asymmetry Measure. As compared with Figure 4-31, thresholding removes a significant number of voxels with relatively small asymmetries, leaving only voxels of significant asymmetries after thresholding.

Further thresholding of the asymmetry measure allows us to better isolate only voxels of significant asymmetry. Since pathological asymmetries are known to be of similar magnitude relative to healthy asymmetries in this case, we chose a conservative asymmetry sizing factor i.e. $S = -0.5$ in order to prune away the voxels of small asymmetry while keeping voxels with medium to large asymmetries intact. Figure 4-32 shows the thresholded asymmetry volume. In addition to the significant asymmetries found in the superior temporal region, the pipeline also detected regions of significant asymmetry in the superior frontal and parietal lobes as shown in the corresponding right image of Figures 4-25 to 4-34. As discussed in the beginning of this section, inclusion of these “healthy” voxels of significant asymmetry is encouraged because these regions may be suggestive of possible subtle cortical abnormalities given their considerable inter-hemispheric difference.




Figure 4-33: Patient Example 2 - Merged Binary Clustering Map.




Figure 4-34: Patient Example 2 - Hotspots Label Map (Closed Merged). Top Row: 3D Hotspots Label Map. Middle Row: Symmetrically Aligned Cortical Volume. Bottom Row: Cortical Volume overlaid with the Hotspots Label Map.

Consequently, we construct a binary clustering input map as outlined in Section 4.1.3 using the above thresholded asymmetry volume. Figure 4-33 shows the resulting binary map which will be used as the input volume for hotspots clustering. Morphological post-processing i.e. dilation followed by erosion is applied on the clustering output to close small islands within the hotspots label map thereby smoothing it. A cluster size threshold is also used to remove small noise clusters with the assumption that lesional clusters of interest are relatively larger than these noise clusters. Since regions of healthy significant asymmetry are rather sizeable for this patient, we chose the mean cluster size as the minimum threshold in order to help reduce noise in the final output while ensuring that we do not outwardly reject any true FCD clusters.

Similarly, evaluation of the final hotspots label map as shown in Figure 4-34 was through expert validation by Dr. Richard Robertson, Director of Neuroradiology, Boston Childrens Hospital (see Section 4.5 for a summary of his analysis). In summary, the hotspots detected by this pipeline are in agreement with the clinical diagnosis of FCD in the right superior temporal lobe. Furthermore, additional regions of significant inter-hemispheric cortical difference were also identified by the pipeline which indicates the sensitivity of the pipeline to potential cortical abnormalities and this is crucial if this pipeline were to be used as an effective first-pass screening tool. Overall, this example has provided further clinical validation of this pipeline's ability to accurately identify potential sites of epileptic seizures for better EEG source localization.

4.4 Plane of Symmetry Detection Robustness Test

The basis of the entire asymmetry analysis framework relies on the accurate detection of the plane of symmetry. Asymmetry analysis can only proceed once the plane of symmetry has been established. The left and right cortices are then compared against their contralateral hemispheres about this symmetry plane to detect regions of significant asymmetries. As such, the ability of the proposed plane of symmetry detector to accurately locate the mid-sagittal plane is imperative to the performance

of the asymmetry analysis components that follow after.

In this section, we shall test the robustness of the plane of symmetry detector with a chronic middle cerebral artery infarction case where a significant portion of the left cerebral hemisphere is absent, due to a stroke. Consequently, there is distinct shifting of major mid-line structures from the right hemisphere to the left hemisphere. It should be noted that this sort of drastic asymmetry is usually not associated with FCD patients. In fact, FCD asymmetry is often several orders of magnitude smaller than what is presented in this case. Nonetheless, we evaluate the plane of symmetry detector on this pathological example to illustrate the robustness of this symmetry detector.




Figure 4-35: Plane of Symmetry Robustness Test Results - Coronal View. The left image of every pair shows the original T1-weighted MRI while the right image shows the volume aligned to the plane of symmetry. The gray line represents the plane of symmetry in the slice shown.

Results of running the plane of symmetry detector on this dataset is shown in Figures 4-35 and 4-36. The plane of symmetry found best matches that which bisects




Figure 4-36: Plane of Symmetry Robustness Test Results - Axial View. The left image of every pair shows the original T1-weighted MRI while the right image shows the volume aligned to the plane of symmetry. The gray line represents the plane of symmetry in the slice shown.

the entire head volume i.e. skull, cortex and other major structures in the head. Given our plane of symmetry definition as the plane where inter-hemispheric similarity is maximum, the plane of symmetry fits the bill perfectly by any reasonable standards. In this case, the plane of symmetry is determined by the remaining cortical material in the left hemisphere that matches with its contralateral counterpart as well as other major symmetrical structural features such as the skull, ocular cavities etc.

However, due to severe shifting of mid-line structures from the right to the left hemisphere, the plane of symmetry detected in this case does not coincide perfectly with the mid-sagittal plane as would be the case in normal patients or patients with relatively small pathologies such as FCD. This is entirely expected even from a clinician's point of view because the mid-sagittal plane in this case is no longer a simple flat planar surface but one which is geometrically complex. Nonetheless, in principle, we could further improve the results of the symmetry detector even in cases of significant asymmetry by first applying the skull stripping operation as described in Section 3.1.2 to remove any non-cortical material from the volume to be aligned. With this new skull-stripped volume as input to the symmetry detector, the detected plane of symmetry will be much closer to the true mid-sagittal plane as defined by the anterior and posterior commissure. Unlike the initial full head volume input where there were other major symmetrical structures to be matched, the skull-stripped volume requires the symmetry detector to find a plane of symmetry based purely on the remaining cortical matter.

More importantly, the purpose of this exercise is to show that the plane of symmetry detector is optimal by design in locating the symmetry plane provided there is a reasonable one to be found. In cases with small i.e. FCD or no pathologies, the detected plane of symmetry often coincides with the mid-sagittal plane because the brain is largely symmetrical. Even with large pathologies as illustrated in this example, the detected plane of symmetry is still reasonably close to the actual mid-sagittal plane. With this, we can safely conclude that this formulation for plane of symmetry detection is robust for our purposes i.e. for FCD patients. We have also demonstrated that even in the presence of large deformations, this MI-based symmetry detector still

performs very well by finding the plane of symmetry that best aligns the remaining common structures in the imaged volume.

4.5 Clinical Evaluation

The following is a clinical evaluation of this quantitative analysis pipeline by Dr. Richard Robertson, Director of Neuroradiology, Boston Children's Hospital.

Currently, the evaluation of children with refractory epilepsy is a highly subjective process based on the assimilation of interpretive data from several sources. MR imaging and electroencephalography are two key elements in the assessment of these patients. Although imaging data are acquired in a digital format, the present methodology for the interpretation of brain MR is based on the visual inspection of hundreds of images each displaying tens of gyri, for subtle architectural or signal intensity asymmetries that may indicate the presence of a cortical malformation. Lesion detection is potentially affected by a number of extraneous factors such as reader experience, interpreter fatigue, work-load, room lighting, visual display characteristics and assorted distractions. The interpretation of the electroencephalography (EEG) tracings, whether obtained using either superficial scalp electrodes or an invasive subdural monitoring electrode array, is subject to similar operator dependencies. It is from this largely qualitative data that conclusions are drawn as to which region of the brain is most likely to be responsible for the generation of seizure activity and whether surgical resection of an epileptogenic focus is feasible. The development of a more quantitative means of image assessment is imperative.

The hypothesis underlying the current project is that a cortical malformation produces right-left brain asymmetry that can be detected at a voxel level by computer assisted interrogation of high-resolution volumetric MR data. The out-put of this analysis is a patient-specific head

model with a superimposed “hotspot” map indicating regions of greater than expected right-left asymmetry. This “hotspot” map is potentially useful both for directing visual inspection of the imaging as well as for correlation with EEG results and potentially more accurate and precise inverse EEG source localization.

For the purposes of testing the validity of the approach, MR data sets of patients with known cortical malformation were retrospectively analyzed. The techniques used were successful in generating patient-specific head models and in identifying as areas of high asymmetry, those regions of the brain known to contain cortical malformations. The techniques also identified other areas of normal (non-lesional) cortex as being asymmetric. This indicates that the methods are sensitive for the detection of right-left asymmetry which is important if the techniques are to be used as a first-pass screening technique for the detection of subtle cortical abnormalities.

In summary, from a clinical perspective, the methods used in this project achieved the desired goal of identifying regions of the brain containing cortical malformation using computer assisted techniques. The success of this strategy is promising for improving the diagnosis and treatment of children with epilepsy.

4.6 EEG Source Localization

Finally, this section presents a summary of the EEG source localization experiments and results from the work by Tiferet A. Gazit[7]. These experiments aim to demonstrate the improvements to source localization with the incorporation of a patient-specific head model and prior probability map as generated by the pipeline described in this thesis. Detailed implementation and results discussion on incorporating a patient-specific head model and a Bayesian prior into source localization can be found in [7] which is part of a joint research with this thesis under the Epilepsy Surgical

Planning project.

For the purpose of EEG source localization experiments as presented in [7], we shall use the patient-specific head model and prior probability map for Patient Example 1 as discussed in Section 4.2. The tissue classification label map and prior probability map generated by this quantitative analysis pipeline are the inputs into the Bayesian EEG source localization module. Note that diffusion tensor data is not available for this patient and hence the resulting patient-specific head model shall be an isotropic one i.e. generic scalar conductivity values are estimated for each tissue class from established literature. A brief discussion is provided in Section 5.2.1 on incorporating diffusion tensor imaging into the quantitative analysis pipeline and producing an anisotropic patient-specific head model that provides a conductivity tensor for each white-matter voxel in the cortical volume.

EEG data used in the source localization experiments are simulated by artificially placing a current dipole within a detected hotspot. Three types of noise variations at three different levels i.e. low, medium and high are added to this dataset in order to determine the robustness and efficacy of adding a patient-specific head model and incorporating a Bayesian prior into the formulation of the source localization problem. EEG potential noise was added in the form of an additive white Gaussian noise with zero-mean and a noise-level specific standard deviation, σ_v . Electrode placement errors in each dimension were generated from a zero-mean Gaussian distribution with a noise-level specific standard deviation, σ_e . Lastly, tissue conductivity values were perturbed by a noise-level specific operator-specified percentage, Δ_c . These perturbations represent the possible errors that may result from noisy surface EEG measurements, movement or placement errors in the 128-channel electrode grid and inaccuracies in conductivity estimation for head modelling respectively. Table 4.2 presents a summary of parameters used at each noise level and the corresponding improvements in EEG source localization results. Two sets of experiments were run on these datasets. The first set (given by the “NeuroFEM” column in Table 4.2) performs EEG source localization with NeuroFEM, a finite element software package developed at the Max Planck Institute for Human Cognitive and Brain Sciences in

Leipzig - Munich / Germany. The primary advantage that NeuroFEM provides is the ability to incorporate a patient-specific head model into EEG source localization. The second set of experiments (given by the “Prior” column in Table 4.2) performs EEG source localization with the Bayesian method developed by Gazit in [7]. In this case, the patient-specific head model and the prior probability map as computed by this pipeline are used in source localization. Additional details of these two sets of experiments and their implementations can be found in [7].

Simulated Noise	$\sigma_v / \mu V$	σ_e / mm	$\Delta_c / \%$	Location Error / mm	
				NeuroFEM	Prior
Zero Noise	NA	NA	NA	0.07	Not necessary
Low Noise	5×10^{-4}	3	5	2	Not necessary
Medium Noise	1×10^{-3}	4	7.5	4.9	0.5
High Noise	5×10^{-3}	5	10	14.6	0.5

Table 4.2: EEG Source Localization - Noise Level Parameters and Results

In the zero noise case, the EEG data is simply the original simulated set of EEG measurements with no noise added. This test serves as a sanity check on our whole experimental setup to ensure that in an ideal noiseless situation, we are able to accurately recover the simulated current dipole.

The experiment is then carried out both with and without the inclusion of a prior probability map to illustrate the improvements that this pipeline brings to source localization. In the low noise case, the addition of noise to the simulated EEG potentials results in a SNR of approximately 22dB. Standard deviation for electrode placement errors was 3mm which is consistent with the possible displacement range of a 128-channel electrode grid that would be used in a clinical setting to obtain real patient EEG data. The resulting dipole location error in this low noise case is only 2mm even without a prior map and this is considered to be adequate for most purposes. As such, incorporation of a Bayesian prior is not necessary to obtain decent localization results in a low noise situation and the results provided by NeuroFEM are sufficient. In the medium noise case, addition of EEG potential noise reduced SNR to 16dB. Coupled with a 4mm electrode location error and a 7.5% conductivity perturbation, the resulting dipole location error becomes 4.9mm using NeuroFEM.

With the Bayesian source localization method, the inclusion of a prior probability map reduces this dipole location error to 0.5mm which is well within range of good source localization. In the high noise case, improvements in source localization with the inclusion of a prior probability map is most evident. Addition of EEG potential noise causes a SNR dip to 2dB which is significantly worse than what we might expect in clinical data. Using NeuroFEM, the dipole location error is 14.6mm. This falls significantly to 0.5mm with the inclusion of a prior probability map thus indicating the success of this pipeline to improving EEG source localization.

Chapter 5

Future Work

In this chapter, we shall briefly discuss the possible future work that would be of interest to this project. These include testing of the pipeline on functional neuroimaging data, which will be discussed in Section 5.1, and the incorporation of diffusion-tensor MRI (DT-MRI) data into the pipeline, which will be discussed in Section 5.2.

5.1 Functional Neuroimaging

Functional neuroimaging modalities such as positron emission tomography (PET) and single-photon emission computerized tomography (SPECT) can provide excellent insights to the functional state of the brain[9]. Availability of such neuroimaging techniques enable us to add a physiological dimension to the quantitative analysis of focal epilepsy. For example, studies have shown that ^{18}F -fluorodeoxyglucose PET (FDG-PET) can reveal interictal areas of decreased cerebral glucose metabolism near the primary epileptic focus while ^{99m}Tc -hexamethyl propyleneamineoxime single photon emission tomography (HMPAO-SPECT) which highlights regional cerebral blood flow can often reveal interictal hypoperfusion of the whole temporal lobe in patients with medically intractable temporal lobe focal epilepsy[14].

Clearly, physiological imaging investigations can provide complementary prior information in the same way as we have demonstrated for structural MRI data. The goal for this section is to analyze functional neuroimaging modalities in parallel with

structural imaging modalities in order to ultimately construct a combined prior probability map which is composed of both structural and functional epileptic hotspots. Different weights should be given to the hotspots detected by different imaging modalities depending on the maturity of the modality in question as well as the relative confidence of the clinician in the modality for localizing the specific pathology of interest. The quantitative framework presented in this thesis provides for a systematic analysis of each of these imaging modalities and the final combined prior probability map of potential epileptic hotspots will be invaluable to improving the accuracy of EEG source localization.

5.2 Diffusion Tensor Imaging

5.2.1 Anisotropic Patient-Specific Head Model

Diffusion-tensor MRI (DT-MRI) is the technique of measuring the anisotropy of water diffusion in the cortical volume. Within white-matter (WM) tissue structures, the mobility of water is greatly restricted in directions orthogonal to the WM fiber tracts due to the densely packed multiple myelin membranes that envelop the axons in these tracts[23]. Consequently, DT-MRI captures this anisotropic water diffusion property within the cortical volume and allows us to compute a tensor at each voxel describing this diffusion in all spatial directions. By assigning the corresponding tensor to each voxel within the patient's brain volume, we can easily construct an anisotropic patient-specific head model from our existing isotropic patient-specific head model.

Given that white-matter structures exhibits anisotropy along the myelination sheaths with an anisotropy ratio of 1:9 (normal:parallel to fibers)[31], the incorporation of this diffusion tensor MRI data allows for more accurate and realistic patient-specific head modelling with fiber-specific conductivity information at each white matter voxel thereby producing better EEG source localization results.

5.2.2 Decreased Fiber Connectivity of FCD Sites

Recent studies have shown that fiber bundles adjacent to regions of FCD pathology exhibit a significant reduction in fractional anisotropy when compared to its corresponding normal contralateral region. Hence we can convert this DT-MRI data into a scalar fractional anisotropy measure and we can analyze this new data in parallel with structural and functional imaging modalities in order to construct an even more comprehensive prior probability map which is composed of structural, functional and diffusion tensor epileptic hotspots. However, it should be noted that the main role of DT-MRI, at least today, is still focused more on white-matter pathology. Granted that the main pathology of FCD includes dysplastic neurons in underlying white matter tissues, DT-MRI is still a very useful modality for detecting this anomaly.

Moreover, diffusion-tensor imaging often detects abnormalities in the brain tissue at an earlier stage when compared to conventional structural imaging modalities such as T1-weighted or T2-weighted imaging[14]. In this way, our prior probability map not only contains MR-visible hotspots, it also includes potential early stage FCD regions that might be a precursor to more severe cortical abnormality in the future.

5.2.3 Neural Connectivity of Potential Surgical Sites

Thus far in the pipeline, we have used multiple structural imaging modalities to construct a realistic isotropic head model as well as to compute a hotspots prior probability map in the interest of improving EEG source localization accuracy. The hotspots prior probability map essentially gives us the ability to hypothesize the presence of focal epileptogenic regions within the cortex to improve the convergence of EEG source localization. At the same time, based on the EEG source localization results, we can also tell if MRI-based methods are concordant with EEG-based methods in localizing epileptic seizure sources. Finally to complete the loop, extrapolating neural connectivity in tissues neighboring the detected epileptic foci can help surgeons predict potential post-surgical functional implications of removing the detected epileptogenic region.

Studies have shown that fiber tractography from diffusion tensor MRI (DT-MRI) data yields neural connectivity information between functional areas of the brain[33]. By tracing the eigenvector that corresponds to the largest eigenvalue of the diffusion tensor at each voxel i.e. the direction of largest diffusion, we can reconstruct these white-matter fiber tracts to reveal the neural connections that originates from or passes through these suspected epileptic hotspots. From the reconstructed fiber tracts, we can extrapolate the post-surgical implication of resecting the epileptogenic region. Care must be taken not to purely seed the fiber tractography within the epileptogenic region itself because FCD pathology itself experiences reduced fiber connectivity[15] and fiber tracking from these points are consequently not as meaningful and enlightening. Voxels neighboring prior hotspots as well as source-localized epileptic foci provide good seeding points for such tractography.

Chapter 6

Conclusion

This thesis describes a systematic pipeline for the quantitative analysis of pediatric focal epilepsy. In the process of designing our system, we were specifically interested in detecting focal cortical dysplasia, a type of focal epilepsy, that is known to be extremely epileptogenic and is often medically refractive. Consequently, surgical resection of the epileptic foci is often the last avenue of cure for these patients and the ability of surgeons to accurately locate these epileptic foci is of paramount importance to the eventual outcome of the surgery. Hence, the main thrust of this project is to improve our ability to accurately locate the sources of these focal epileptic seizures.

In this thesis, we proposed to accomplish this through the use of a patient-specific head model and the incorporation of a prior probability map of potential epileptic hotspots. The use of a patient-specific head model allows us to better model current conductivity and potential distribution across the cortical volume of a patient. This in turn results in more realistic simulations than what is currently available with purely generic spherical head models. It is also a vast improvement from the commonly used techniques of identifying broad epileptic regions from visual inspection of the raw EEG data and EEG source localizations with simplistic spherical head modeling.

The other major contribution of this project is the generation of a prior probability map which reflects cortical abnormalities that is picked up by various neuroimaging techniques such as MRI etc. We have demonstrated through clinical validation in this thesis the ability of the pipeline to pick out these abnormalities and to high-

light areas which are more likely to contain the epileptic foci in a prior probability map. Given the ill-posed nature of EEG source localization, this prior map helps significantly to constrain and to bias the source localization solution towards possible cortical abnormalities as detected by neuroimaging modalities. The open design of the system allows for neuroimaging modalities other than MRI to be incorporated into the pipeline in the future for more extensive detection of cortical abnormalities beyond structural MRI anomalies. This fusion of data truly allows us to perform a more directed search for the epileptic foci based on possible underlying abnormalities that might be associated with the epileptic seizure as opposed to a blind search through the entire cortical volume.

We also provided a summary of the results from EEG source localization experiments based on the patient-specific head model and prior probability map generated by the pipeline in this thesis. Initial results show that the incorporation of a patient-specific head model and a probabilistic hotspots prior is most helpful for improving source localization especially in a high noise setup. This high noise tolerance reflects the robustness of the overall system design and its potential for helping surgeons better localize epileptic foci in a clinical setting.

Moving forward, the goal is to incorporate more imaging modalities into the pipeline and to fuse these modalities which can reveal different underlying characteristics of the patient's cortical volume into a better, more encompassing prior probability map. It will also be extremely valuable to validate this pipeline with more clinical dataset especially on children with a full range of neuroimaging acquisitions i.e. MRI, PET, SPECT, DT-MRI etc., corresponding EEG measurements from a 128-channel electrode grid and more importantly surgical validation.

In summary, we have demonstrated in this thesis the utility of the quantitative analysis pipeline and how we can use the results to improve EEG source localization. We have also highlighted the potentials of this pipeline and how it can be further improved to achieve even better results.

Bibliography

- [1] S.B. Antel, A. Bernasconi, N. Bernasconi, D.L. Collins, R.E. Kearney, R. Shinghal, and D.L. Arnold. Computational models of MRI characteristics of focal cortical dysplasia improve lesion detection. *Neuroimage*, 17(4):1755-60, 2002.
- [2] J.F. Bautista, N. Foldvary-Schaefer, W.E. Bingaman, and H.O. Luders. Focal cortical dysplasia and intractable epilepsy in adults: Clinical, EEG, imaging, and surgical features. *Epilepsy Res*, 55(1-2):131-6, 2003.
- [3] R.S. Briellmann, G.S. Pell, R.M. Wellard, L.A. Mitchell, D.F. Abbot, and G.D. Jackson. Mr imaging of epilepsy: State of the art at 1.5T and potential of 3T. *Epileptic Disorders*, 5(1):3-20, 2003.
- [4] S.E. Connor and J.M. Jarosz. Magnetic resonance imaging of patients with epilepsy. *Clin Radiol*, 56(10):787-801, 2001.
- [5] J.S. Ebersole. Functional neuroimaging with eeg source models to localize epileptogenic foci noninvasively. *University of Chicago Hospitals' Clinical Comment*.
- [6] I. Fried. Magnetic resonance imaging and epilepsy: Neurosurgical decision making. *Magn Reson Imaging*, 13(8):1163-70, 1995.
- [7] T.A. Gazit. A flexible framework for non-invasive source localization in pediatric focal epilepsy. Master's thesis, Massachusetts Institute of Technology, 2005.
- [8] V. Grau, A.U.J. Mewes, M. Alcaniz, R. Kikinis, and S.K. Warfield. Improved watershed transform for medical image segmentation using prior information. *IEEE Transactions on Medical Imaging*, 23(4):447-458, 2004.
- [9] A. Hammers, M.J. Koepp, M.P. Richardson, R. Hurlemann, D.J. Brooks, and J.S. Duncan. Grey and white matter flumazenil binding in neocortical epilepsy with normal MRI. a PET study of 44 patients. *Brain*, 126(6):1300-18, 2003.
- [10] J. Haueisena, D.S. Tuch, C. Ramonc, P.H. Schimpfd, V.J. Wedeenb, J.S. Georgee, and J.W. Belliveaub. The influence of brain tissue anisotropy on human EEG and MEG. *NeuroImage*, 15(1):159-66, 2002.
- [11] A.C. Hui, J.M. Lam, Y.L. Chan, K.M. Au-Yeung, K.S. Wong, R. Kay, and W.S. Poon. Role of magnetic resonance imaging for preoperative evaluation of patients with refractory epilepsy. *Hong Kong Medical Journal*, 9(1):20-4, 2003.

- [12] L. Ibanez, W. Schroeder, L. Ng, and J. Cates. *The ITK Software Guide*. Insight Software Consortium, 2003.
- [13] S. Joshi, P. Lorenzen, G. Gerig, and E. Bullitt. Structural and radiometric asymmetry in brain images. *Medical Image Analysis*, 7(2):155-70, 2003.
- [14] S. Lamusuo, H.M. Ruottinen, J. Knuuti, R. Hrknen, U. Ruotsalainen, J. Bergman, M. Haaparanta, O. Solin, E. Mervaala, U. Nousiainen, S. Jskelinen, A. Ylinen, R. Klviinen, J.K. Rinne, M. Vapalahti, and J.O. Rinne. Comparison of [18F]FDG-PET, [99mTc]-HMPAO-SPECT, and [123I]-iomazenil-SPECT in localising the epileptogenic cortex. *J. Neurol. Neurosurg. Psychiatry*, 63(6):743-748, 1997.
- [15] S.K. Lee, D.I. Kim, S. Mori, J. Kim, H.D. Kim, K. Heo, and B.I. Lee. Diffusion tensor MRI visualizes decreased subcortical fiber connectivity in focal cortical dysplasia. *NeuroImage*, 22(4):1826-9, 2004.
- [16] A. Lortie, P. Plouin, C. Chiron, O. Delalande, and O. Dulac. Characteristics of epilepsy in focal cortical dysplasia in infancy. *Epilepsy Res*, 51(1-2):133-45, 2002.
- [17] J.F. Mangin. Entropy minimization for automatic correction of intensity non-uniformity. *Mathematical Methods in Biomedical Image Analysis*, 2000.
- [18] The Epilepsy Foundation of America. The scope of epilepsy and seizures in the united states. <http://www.epilepsyfoundation.org/answerplace/index.cfm>, 2004.
- [19] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629-39, 1990.
- [20] J.P.W. Pluim, J.B.A. Maintz, and M.A. Viergever. Mutual information based registration of medical images: A survey. *IEEE Transactions Medical Imaging*, 22(8):986-1004, 2003.
- [21] F. Rosenow and Hans Lders. Presurgical evaluation of epilepsy. *Brain*, 124(9):1683-1700, 2001.
- [22] D.W. Shattuck, S.R. Sandor-Leahy, K.A. Schaper, D.A. Rottenberg, and R.M. Leahy. Magnetic resonance image tissue classification using a partial volume model. *NeuroImage*, 13(5):856-76, 2001.
- [23] M.E. Shenton, M. Kubicki, and R.W. McCarley. Diffusion tensor imaging: Image acquisition and processing tools. *SPL Technical Report #354*, 2002.
- [24] S.M. Sisodiya, S.L. Free, D.R. Fish, and S.D. Shorvon. Increasing the yield from volumetric MRI in patients with epilepsy. *Magnetic Resonance Imaging*, 13(8):1147-52, 1995.

- [25] J.P. Thirion, S. Prima, and G. Subsol. Statistical analysis of dissymmetry in volumetric medical images. *Medical Image Analysis*, 4(2):111-21, 2000.
- [26] B.C Vemuri, J. Ye, Y. Chen, and C.M. Leonard. Image registration via level-set motion: Applications to atlas-based segmentation. *Medical Image Analysis*, 7(1):1-20, 2003.
- [27] W. Wells, P. Viola, H. Atsumi, S. Nakajima, and R. Kikinis. Multi-modal volume registration by maximization of mutual information. *Medical Image Analysis*, 1(1):35-51, 1996.
- [28] K. Whittingstall, G. Stroink, L. Gates, J.F. Connolly, and G.A. Finley. Influence of realistic head geometry differences on EEG source localization. *International Journal of Bioelectromagnetism*, 4(2):257-8, 2002.
- [29] M. Wilke, J. Kassubek, S. Ziyeh, A. Schulze-Bonhage, and H.J. Huppertz. Automated detection of gray-matter malformations using optimized voxel-based morphometry: A systematic approach. *NeuroImage*, 20(1):330-43, 2003.
- [30] F.G. Woermann, S.L. Free, M.J. Koepp, J. Ashburner, and J.S. Duncan. Voxel-by-voxel comparison of automatically segmented cerebral gray matter - a rater-independent comparison of structural MRI in patients with epilepsy. *NeuroImage*, 10:373-384, 1999.
- [31] C.H. Wolters, A. Anwander, M.A. Koch, S. Reitzinger, and M. Svendsen M. Kuhn. Influence of head tissue conductivity anisotropy on human EEG and MEG using fast high resolution finite element modelling, based on a parallel algebraic multigrid solver. 2003.
- [32] C. Zhu and T. Jiang. Multicontext fuzzy clustering for separation of brain tissues in magnetic resonance images. *NeuroImage*, 18(3):685-96, 2003.
- [33] L. Zhukov and A.H. Barr. Oriented tensor reconstruction: Tracing neural pathways from diffusion tensor MRI. *IEEE Computer Society Proceedings of the Conference on Visualization '02*:387-394, 2002.