

LOCALISATION AND CHARACTERISATION OF BRAIN PATHOLOGY FROM STRUCTURAL MRI DATA

vorgelegt von

Maddalena Strumia

Freiburg, 2016

Dissertation zur Erlangung des Doktorgrades der Technischen Fakultät
der Albert-Ludwigs-Universität Freiburg im Breisgau



Dekan:

Prof. Dr. Georg Lausen

Prüfungskommission:

Prof. Dr. Matthias Teschner

(Vorsitz)

Prof. Dr. Rolf Backofen

(Beisitz)

Prof. Dr.-Ing. Thomas Brox

(1. Gutachter)

Prof. Dr. Dr. h. c. Jürgen Hennig

(2. Gutachter)

Tag des Kolloquiums:

27.04.2016

To my family and Marius

Abstract

The segmentation of the brain in medical images is a challenging problem owing to the high complexity of the brain structure. In the presence of tissue alterations, such as tumours and diffuse lesions, the localisation and segmentation of a given tissue, or region, of interest becomes even more challenging due to large anatomic deformations.

This work addresses several applications of automatic segmentation of brain structures in magnetic resonance (MR) images, with the aim of improving clinical diagnosis. These include the identification of brain tissues, the segmentation of pathological brain regions, and the quantification of abnormal brain vasculature. In particular, three brain diseases have been considered: focal cortical dysplasia, multiple sclerosis and glioblastoma tumours. Because of the high anatomical variability of these diseases, a specific method has been developed for each of them.

Focal cortical dysplasia is characterised by an abnormal stratification of the brain grey matter. Its identification in MR images is a difficult task due to the limited contrast of the lesions. Current methods rely solely on the information extracted from T_1 -weighted MR images and, to a lesser extent, also from FLAIR images. In this work the delineation of this kind of lesion is performed through the extraction of several features from both T_1 -weighted and FLAIR images, which represent intensity, shape, and form, and by creating a baseline for normal-appearing tissue based on the MR images of several healthy subjects. These features are extracted for both the patient and the healthy volunteer datasets and are subsequently combined in a probabilistic framework. The proposed method has been tested on 11 patient datasets and has been compared with a state-of-the-art method. For all patient data the focal cortical dysplasia has been detected with a high precision and with improved results with respect to the compared method.

Multiple sclerosis is a neurodegenerative disease of the central nervous system, whose main manifestations are lesions in the white brain matter. In T_1 -weighted and FLAIR images, multiple sclerosis lesions have similar intensities as the grey matter. Many methods rely strongly on an atlas-driven pre-segmentation, which can be inaccurate at the boundary between grey and white matter. As a result, the detected white matter can contain areas of grey matter that might be segmented as lesion. Here, a geometric brain model is proposed to simultaneously segment all tissue types as well as the lesions without using an atlas. The segmentation problem is formulated as an energy minimisation approach and a local solution is found via the graph-cuts method. The energy consists of three terms, namely the appearance, local and global energies. The appearance energy models the intensity of the different tissues, whereas the local energy models the geometric constraints. The global energy is composed of a connectivity prior for the grey matter and a minimum size constraint for the lesions. The proposed connectivity prior achieves better results compared to state-of-the-art methods and recovers a number of false neg-

atives for the lesions. On the other hand, the minimum size constraint forces the lesions to have a minimum size and therefore avoids possible noisy lesion segmentations. The proposed method has been tested on two different databases, and for both competitive results in comparison to several state-of-the-art methods have been achieved.

Glioblastoma tumours are one of most aggressive types of primary brain tumours. To maintain their aggressive proliferation they initiate the formation of new blood vessels which present irregular shapes and immature organisation patterns, and they have a high permeability. In this work, novel quantification measures for these abnormalities are proposed. These measures are based on the local incoherence of vessel orientations and are computed for both tumour and normal-appearing vasculature. The method has been tested on 12 patients, and in every dataset all the proposed measures are able to discriminate very accurately between tumour and normal vessels. Additionally, one of our measures has shown correlation with two histological markers, *i.e.* MIB and endothelial proliferation. A weak but positive correlation with MIB has been found, while the endothelial proliferation has a negative correlation with our measure.

The segmentation methods proposed in this work open the way for a more precise diagnostic evaluation in particular in diseases where lesion identification is currently challenging. In addition the novel quantification measures proposed for the glioblastoma vessel analysis in combination with histologic correlation can help to connect macroscopic vessel appearance with cellular abnormalities in blood vessels and might therefore be suited to identify prognostic subgroups in glioblastoma.

Zusammenfassung

Die Segmentierung des Gehirns in medizinischen Bildern ist ein schwieriges Problem aufgrund der hohen Komplexität der Gehirnstruktur. In Gegenwart von Gewebeveränderungen, wie beispielsweise Tumoren oder diffusen Läsionen, wird die Lokalisierung und Segmentierung eines vorgegebenen Gewebes oder einer Region noch schwieriger wegen der großen anatomischen Deformationen.

Diese Arbeit befasst sich mit einigen Anwendungen der automatischen Segmentierung von Hirnstrukturen in Magnetresonanztomographie-(MRT)-Bildern, mit dem Ziel, die klinische Diagnostik zu verbessern. Dazu gehören die Identifizierung von Hirngeweben, die Segmentierung der pathologischen Hirnregionen, und die Quantifizierung von abnormen Hirngefäßsystemen. Insbesondere wurden drei Hirnerkrankungen untersucht: die fokale kortikale Dysplasie, die Multiple Sklerose und Glioblastome. Wegen der hohen anatomischen Variabilität dieser Erkrankungen wurde eine angepasste Methode für jede von ihnen entwickelt.

Die fokale kortikale Dysplasie ist durch eine abnorme Stratifizierung der grauen Hirnsubstanz gekennzeichnet. Die Identifizierung in MR-Bildern ist eine schwierige Aufgabe aufgrund des begrenzten Kontrasts der Läsionen. Derzeitige Verfahren verlassen sich ausschließlich auf die Informationen die aus T_1 -gewichteten MR-Bilder extrahiert werden, und in geringerem Umfang auch auf FLAIR Bilder. In dieser Arbeit wurde die Abgrenzung der Läsionen durch die Extraktion mehrerer Eigenschaften aus den T_1 -gewichteten und den FLAIR Bildern durchgeführt: die Intensität, die Form, und durch die Erstellung eines Basisdatensatzes für Normalgewebe basierend auf MR-Bildern von gesunden Probanden. Diese Eigenschaften wurden sowohl in Datensätzen von Patienten als auch von gesunden Probanden bestimmt, und anschließend mit einem probabilistischen Ansatz kombiniert. Das vorgeschlagene Verfahren wurde an 11 Patientendatensätzen evaluiert und mit etablierten Verfahren verglichen. In allen Patientendaten konnte die fokale kortikale Dysplasie mit hoher Präzision und mit verbesserten Ergebnissen in Bezug auf die etablierten Verfahren nachgewiesen werden.

Multiple Sklerose ist eine neurodegenerative Erkrankung des zentralen Nervensystems, deren Hauptmanifestationen Läsionen in der weißen Hirnsubstanz sind. In T_1 -gewichteten Bildern und in FLAIR Bildern haben MS-Läsionen ähnliche Signalintensitäten wie die graue Hirnsubstanz. Viele Verfahren beruhen auf einer Atlas-basierten Vorsegmentierung, die an der Grenze zwischen grauer und weißer Substanz ungenau sein kann. Infolgedessen kann die detektierte weiße Substanz Bereiche der grauen Substanz enthalten, die als Läsionen segmentiert werden. In dieser Arbeit wird ein geometrisches Hirnmodell ohne Verwendung eines Atlas vorgeschlagen um gleichzeitig alle Gewebetypen sowie die Läsionen zu segmentieren. Das Segmentationsproblem wird als Energieminimierung formuliert, und eine lokale Lösung wird über ein Graph-Cut-Verfahren

bestimmt. Die Energie besteht aus drei Termen: der Form, der lokalen Energie, und der globalen Energie. Die Formenergie gibt die Signalintensität der verschiedenen Gewebe wider, während die lokale Energie die geometrischen Randbedingungen modelliert. Die globale Energie besteht aus einer Konnektivitäts-Annahme (connectivity prior) für die graue Substanz, einer Nebenbedingung einer Mindestgröße (minimum size constraint) für die Läsionen. Die Konnektivitäts-Annahme erzielt bessere Ergebnisse als konventionelle Methoden, insbesondere bei der Identifikation von falsch negativen Ereignissen. Andererseits erzwingt die Größeneinschränkung eine Mindestgröße der Läsionen und vermeidet so mögliche Fehlsegmentierungen im Rauschen. Das vorgeschlagene Verfahren wurde auf zwei unterschiedlichen Datenbanken getestet, und für beide konnten im Vergleich zu konventionellen Methoden vergleichbare Ergebnisse erreicht werden.

Glioblastome stellen eine der aggressivsten Arten von primären Hirntumoren dar. Zur Aufrechterhaltung ihrer aggressiven Proliferation initiieren sie die Bildung neuer Blutgefäße, die an ihrer unregelmäßigen Form und dem unreifen Muster erkennbar sind, und die eine hohe Permeabilität aufweisen. In dieser Arbeit werden neue Quantifizierungsmethoden für diese Gefäßanomalien vorgeschlagen. Diese Methoden basieren auf der lokalen Inkohärenz der Gefäßorientierungen und werden sowohl für Tumorgefäße wie auch für normal erscheinende Gefäßsysteme berechnet. Die Methode wurde an 12 Patienten getestet, und in jedem Datensatz konnte das Verfahren sehr genau zwischen Tumor- und Normalgefäßen unterscheiden. Darüber hinaus korrelierte ein Messparameter mit zwei histologischen Markern, dem MIB und der Endothelproliferation. Es zeigte sich eine schwach positive Korrelation mit dem MIB, während die Endothelproliferation eine negative Korrelation aufwies. Die in dieser Arbeit vorgeschlagenen Segmentierungsverfahren öffnen den Weg für eine präzisere Diagnostik insbesondere in Krankheiten, bei denen die Läsionsidentifizierung momentan noch schwierig ist. Außerdem können die neuen Messparameter für die Glioblastom-Gefäßanalyse in Kombination mit histologischen Analysen dazu beitragen, die makroskopische Gefäßstruktur mit zellulären Abnormalitäten der Blutgefäße zu verbinden, so dass möglicherweise prognostische Untergruppen von Glioblastompatienten identifiziert werden können.

Acknowledgements

I want to thank all those people who made this thesis possible and an unforgettable experience.

Fist of all I would like to thank Prof. Jürgen Hennig, Prof. Michael Bock and Prof. Thomas Brox for giving me the opportunity to do the PhD, and for the interesting and very helpful discussions about all the projects I worked on. Thank you Michael for teaching me how to write in a clear and effective way and for your advices and guidance through the several topics I encountered.

I would like to thank my supervisor Dr. Frank Schmidt for introducing me to the world of research in computer vision, especially to graph-cuts. Thank you for always bring up new great ideas and for the time you dedicated to me answering all my questions. Thank you also for your support especially during the long evenings before the deadlines, and for pushing me to get always better results in such a short time. I also thank Dr. Stathis Hadjidemetriou for his support and advices in the first few months of my PhD.

I would also like to thank the Medical Physics group for their warm welcome and the great time during the three and a half years I spent in Freiburg. Thank you for the nice Christmas parties and events we had, and also for the great opportunity to have an insight into the magnetic resonance world.

A special thank goes to my colleagues and friends from the “Spelunke group”: Denis, Elmar, Ben, Seb, Michael, Iulius, Marius, David and Robert. Thank you all for the great evenings we had at our favorite pub discussing interesting and intriguing topics in front of a beer. Thank you also for the support you gave me during the several hard periods I had. I will never forget those moments. Thank you David for the inestimable time we spent together, for our coffee breaks where we discussed work and everyday life problems. Thank you for listening, supporting me, proof-reading my papers, and for always having a nice word that changed my day. I had a great time with you and independently on where we will be, we will always be great friends.

I would also like to thank you, José, especially for the time you dedicated me in the last period of my PhD, for reading and correcting my thesis several times. Thank you also for making the lunch breaks such enjoyable moments full of great discussions.

Finally, I take the opportunity to express my profound gratitude to my parents, my brothers Guglielmo and Jacopo, my sister Carola and to my life partner Marius for their love and precious support. A special thank goes to my mum, a great woman to whom I owe everything. Thank you mum for always being there for me, for listening and advising me in the best way and for your great mind always able to understand things in such a peculiar way. Without your encouragements I wouldn't have been able to publish this thesis. I also would like to thank you, Marius for your love. You always supported me in so many uncountable ways, being there whenever I needed you. Thank you for

Acknowledgements

your patience and for enduring the very stressful times when I worked long nights and weekends. Thank you for your invaluable help with small and big things, and for your great ability in bringing new prospectives into my work. You always believed in me, and gave me the confidence to face every difficulty I encountered. With you I feel being able to manage everything. I love you.

Contents

Abstract	iii
Zusammenfassung	v
Acknowledgements	vii
List of Publications	xiii
1 Introduction	1
1.1 Outline of the Thesis	4
2 Basic Concepts of MRI	5
2.1 Larmor Frequency and Net Magnetisation	5
2.2 Relaxation	7
2.3 Image Creation	8
2.4 Pulse Sequences	11
3 Image Segmentation	13
3.1 Bayesian Formulation	14
3.2 Segmentation Methods	15
3.2.1 Thresholding	15
3.2.2 Feature-based Techniques	16
3.2.3 Classifiers	18
3.2.4 Clustering	20
3.2.5 Markov Random Fields and Graph-based Techniques	23
3.2.6 Atlas-based Approaches	25
3.2.7 Active Contours and Level-set Methods	28
3.3 Relation between Denoising and Segmentation	33
3.3.1 Total Variation	34
3.3.2 Other Methods	35
4 Introduction to Graph-cuts	39
4.1 General Formulation	39
4.2 Graph Construction	40
4.3 Submodularity Condition	41
4.4 Multi-labelling Segmentation	43
4.5 Local Optimisation	45

4.5.1	Swap Move	46
4.5.2	Expansion Move	47
4.6	Further Developments	48
4.6.1	Objects Interaction	48
4.6.2	Connectivity Constraint	50
5	Localisation of Focal Cortical Dysplasia	51
5.1	Related Work	52
5.2	Pre-processing	53
5.3	Extraction of Features	54
5.3.1	Intensity-based Features	55
5.3.2	Texture-based Features	55
5.3.3	Form-based Features	56
5.3.4	Spatial Probability Map for FCD Lesions	57
5.4	Results	59
5.5	Summary and Discussion	63
6	Segmentation of Multiple Sclerosis Lesions	65
6.1	Related Work	66
6.2	Notation	68
6.3	Pre-processing	69
6.3.1	Image Registration	69
6.3.2	Inhomogeneity Correction	69
6.3.3	Partial Segmentation	71
6.4	Geometric Brain Model	72
6.4.1	Appearance Energy	73
6.4.2	Local Geometric Model	74
6.4.3	Global Geometric Model	76
6.5	Energy Optimisation	77
6.5.1	Input Images	78
6.5.2	Computation of the Initial Labelling	79
6.5.3	Local Energy Minimisation	83
6.5.4	Hyper-parameter Learning	85
6.6	Results	87
6.6.1	Measures	87
6.6.2	MICCAI Dataset	90
6.6.3	In-House Database	90
6.6.4	Robustness against Noise	94
6.7	Summary and Discussion	95
7	Quantification of Tumour Vessel Abnormality	99
7.1	Related Work	100
7.2	Method	101
7.2.1	Pre-processing and Vessel Segmentation	101

7.2.2	Measures of Vessel Abnormality	104
7.3	Comparison to Histology	106
7.4	Results	107
7.5	Summary and Discussion	109
8	Conclusion	113
8.1	Summary of Contributions	113
8.2	Limitations and Future Work	114

List of Publications

Journal Papers

- Strumia M., Schmidt F. R., Anastasopoulos C., Granziera C., Krueger G., Brox T. *White Matter MS-Lesion Segmentation Using a Geometric Brain Model*. IEEE Transactions on Medical Imaging, (2016).
- Strumia M., Reichardt W, Staszewski O., Heiland H., Weyerbrock A., Mader I., Bock M. *Glioblastoma Vessels Abnormality Quantification using Time-of-flight MR Angiography*. Magnetic Resonance Materials in Physics, Biology and Medicine, (2016).
- Pilutti D., Strumia M., Büchert M., Hadjidemetriou S. *Non-Parametric Bayesian Registration (NParBR) of Body Tumors in DCE-MRI Data*. IEEE Transactions on Medical Imaging, (2015).
- Pilutti D., Strumia M., Hadjidemetriou S. *Bi-modal Non-rigid Registration of Brain MRI Data with Deconvolution of Joint Statistics*. IEEE Transactions on Image Processing, vol. 23(9), pp. 3999-4009, (2014).

Conference Papers

- Pilutti D., Strumia M., Hadjidemetriou S. *Non-Parametric Bayesian Registration (NParBR) on CT Lungs Data-EMPIRE10 Challenge*. url:<http://empire10.isi.uu.nl/> (2015).
- Strumia M., Ramantani G., Mader I., Henning J., Bai L., Hadjidemetriou S. *Analysis of Structural MRI Data for the Localisation of Focal Cortical Dysplasia in Epilepsy*. MICCAI-Workshop Clinical Image-Based Procedures. From Planning to Intervention, pp. 25-32, (2013).
- Pilutti D., Strumia M., Hadjidemetriou S. *Bi-modal Non-rigid Registration of Brain MRI Data Based on Deconvolution of Joint Statistics*. MICCAI-workshop Multimodal Brain Image Analysis, pp. 40-50, (2013).
- Strumia M., Anastasopoulos C., Mader I., Henning J., Bai L., Hadjidemetriou S. *Comparative characterisation of susceptibility weighted MRI for brain white matter lesions in MS*. MICCAI-Workshop Multimodal Brain Image Analysis, pp.157-166, (2012).

- Strumia M., Feltell D., Evangelou N., Gowland P., Tench C., Bai L. *Grey matter segmentation of 7T MR images*. IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), pp. 3710-3714, (2011).

Abstracts

- Strumia M., Reichardt W, Staszewski O., Heiland H., Weyerbrock A., Mader I., Bock M. *Correlation of Vessel Abnormality Quantification from TOF MRA with Histologic and Gene Markers*. International Society for Magnetic Resonance in Medicine, (2016).
- Strumia M., Reichardt W., Mader I., Hennig J., Bock M. *Glioblastoma Vessels Abnormality Quantification in TOF-MRA*. International Society for Magnetic Resonance in Medicine, (2014).
- Vicari M., Izadpanah K., Serra A., Dragonu I., Li G., Strumia M., Hennig J. *Accelerated 3D RARE for Positional Weight-Bearing MRI of ACJ Bone Fixation with Metal Implants*. International Society for Magnetic Resonance in Medicine, (2014).

Chapter 1

Introduction

Medical image segmentation is one of the primary post-processing techniques to support diagnostic image analysis in clinical practice. In fact, it is commonly used in a wide range of applications, including localisation of pathology, treatment planning, computer-integrated surgery, and to monitor the evolution of degenerative diseases. Due to the complexity of the human body, the segmentation of medical images is still a challenging task, and each organ, or tissue of interest requires in general specific segmentation techniques that may not be appropriate for other organs or tissues.

The human visual system is able to discriminate between several objects in an image on account of different shape, texture and/or local intensity. These commonly used features are also exploited by segmentation algorithms that automatically extract individual objects/organs or discriminate among multiple ones. In many pathologic conditions, and especially in the case of cancer, the affected organ(s) often show large anatomic deformations, as well as noticeable changes in their fat/water content. Such alterations may affect the shape, the size, and the relative anatomical location of a given organ. Therefore, segmentation techniques developed for healthy subjects can fail, in the presence of disease, to provide equally satisfactory, quantitative results. More advanced algorithms that take into account organ-specific changes caused by disease are therefore needed.

In particular, the accurate segmentation of diseased brain tissue remains an open problem, due both to the inherently high structural complexity of the brain and to tissue deformations caused by disease. For this reason, in this thesis novel algorithms have been developed for the segmentation of focal cortical dysplasia and multiple sclerosis lesions, and for the quantification of vessel abnormality in glioblastoma brain tumours. The remainder of this chapter gives a brief introduction to brain anatomy, pathology and imaging methods, as relevant to this work, and closes with a description of the structure of the thesis.

The human nervous system is divided into the central nervous system (CNS) and the peripheral nervous system (PNS). The CNS, in turn, is divided into the brain and the spinal cord. The brain is responsible for integrating sensory information and for providing higher motor and cognitive functions, including language, planning, reasoning, and abstract thought. The brain is therefore one of the most extensively studied organs. For this purpose, various imaging techniques, combined with post-processing methods, have proved extremely useful. In this work, magnetic resonance (MR) images were used. In clinical practice, magnetic resonance imaging (MRI) has established itself as a technique of choice for imaging not only for the brain, but many other organs as




Figure 1.1: In (a) an axial, in (b) a sagittal, in (c) a coronal slice of a MRI healthy brain. Starting from the outer border there are the skull, the corticospinal fluid (CSF), the cortical grey matter (GM), the white matter (WM), the subcortical nuclei/grey matter (SGM), and the ventricles (V).

well. Due to the absence of ionising radiation, MRI is less invasive than alternative imaging techniques, such as computed tomography (CT). Moreover, it provides soft tissue contrast as well as variable image contrasts. As an example, three anatomical MR images, at different orientations, of a healthy brain are shown in Figure 1.1. A plethora of segmentation algorithms have been developed to delineate the different tissues and substructures of the brain, providing useful insight in the anatomical variation of the brain and its pathology.

The brain is separated from the enclosing skull by membranous tissue. Between these membrane layers the corticospinal fluid (CSF) is circulating, which also fills the ventricular system (V) in the centre of the brain. The brain itself is classified in two tissue types, namely the white matter (WM) and the grey matter (GM). The GM on the surface of the brain is called cortex and is mainly formed by the neuron cell bodies and their respective nuclei. The internal GM (SGM) contains several subcortical nuclei (e.g., the globus pallidus, the caudate nucleus, and the putamen). The WM is mainly formed by neuron axons. In Figure 1.1 a typical MRI of the human head is depicted with T_1 -weighted (T_1 -w) contrast. The intensity of the brain tissues is higher compared to the skull and the CSF, which appear dark, and the intensity of the WM is higher than the GM.

Concerning brain pathology, both brain tumours and neurodegenerative diseases have received particular research attention. The precise identification and delineation of a tumour boundary is crucial to the successful surgical removal of the cancerous tissue. This can prolong patient life expectancy and, in certain cases, lead to nearly normal living conditions. As regards neurodegenerative diseases, it is of foremost importance to monitor their progress over time, as this will help pharmaceutical research to develop new drugs able to slow the degeneration.

Focal cortical dysplasia (FCD) is a malformation of cortical development, and one of the major causes of pharmacologically intractable epilepsy. This disease is caused by an abnormal development of the cortical grey matter leading to an atypical stratification

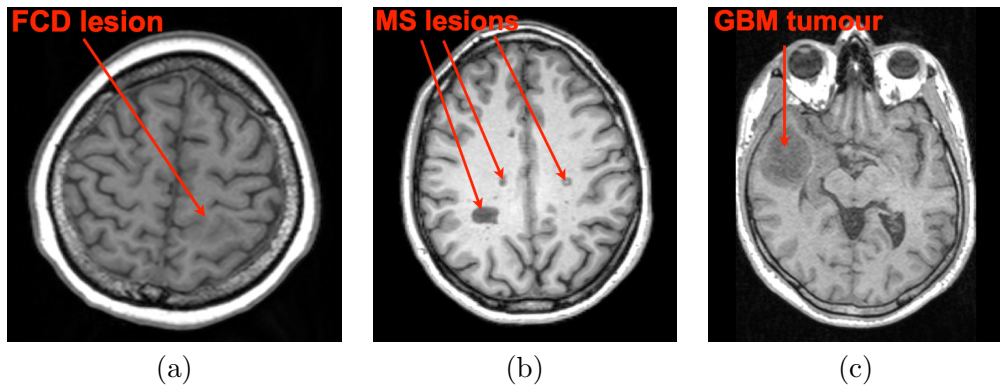


Figure 1.2: Axial MR slices illustrating the three pathologies considered in this work. In (a) an FCD patient, in (b) an MS patient, and in (c) a patient with a glioblastoma tumour.

of the tissue. Patients affected by FCD experience recurrent seizures, and can presently achieve seizure freedom only after successful surgery. Therefore, the precise pre-surgical localisation and complete surgical removal of the dysplastic tissue is pivotal for a good outcome and for the avoidance of any additional interventions.

Multiple Sclerosis (MS) is the most frequent chronic inflammatory disease of the CNS caused by inflammatory demyelination (destruction of the myelin sheath covering the nerve axons). The pathogenesis of MS lesions is still unclear and one of the most challenging aspects of the disease that remains to be understood is the nature of, and mechanisms leading to, tissue injury. The main clinical manifestations of this disease are caused by lesions affecting the brain white matter and grey matter. Accordingly, lesion identification and volumetry are crucial steps in quantifying the burden of the disease. In the clinical routine the lesions are still manually annotated, which is time-consuming, leads to subjective variability, and lacks reproducibility.

Glioblastoma multiforme is one of the most aggressive primary brain tumours. Due to its rapid proliferation pattern, glioblastoma depends on the formation of new supplying blood vessels (neo-angiogenesis). Such newly formed vessels are often characterised by an increased number and size, an irregular shape, immature vessel organisation, and, in particular, by a high degree of vascular permeability. Many anti-angiogenic therapeutic approaches attempt to inhibit the formation and maturation of new blood vessels, and thus to disconnect the growing tumour from its vascular supply with nutrients. The therapeutic success of anti-angiogenic therapies depends on many factors, and the early detection and quantification of the changes of the vascular architecture under therapy is highly desirable in order to improve and stratify anti-angiogenic therapies.

Figure 1.2 shows a representative axial slice for each of the three diseases considered in this work.

1.1 Outline of the Thesis

Chapter 2 explains the basic principles of nuclear magnetic resonance imaging. Starting from the level of the nucleus, the principle for generating MR images of given tissues and organs is described. Because MR imaging can provide a variety of image contrasts in order to differentiate the tissues of interest, the relevant brain imaging contrasts used in this work are introduced.

Chapter 3 introduces the background theory on image segmentation with a brief description of the most interesting and relevant techniques. In this connection, an initial grey matter segmentation method is presented. This method combines the well-known atlas based and level-set image segmentation techniques to improve the segmentation over pure atlas based methods.

Chapter 4 further introduces the background theory of graph-cuts based methods. Starting from the binary segmentation method, the extension to a multi-labelling approach is described. More recent developments, such as object interactions are also described. The theory presented in this chapter is fundamental for the understanding of our contribution to the problem of lesion segmentation presented in Chapter 6.

The next three chapters present the main contributions developed in this thesis:

Chapter 5 presents our work on focal cortical dysplasia lesion segmentation. These lesions are, in many cases, not visible with conventional MR imaging. Therefore, we propose a segmentation method which combines several features extracted from multiple brain MR images in order to identify and segment this type of lesions. Due to the high complexity of the problem, we rely on information extracted from healthy brains which provides a baseline for ‘normal’ brain tissue. Both the extracted features and the healthy baseline are subsequently combined into a probabilistic framework to highlight abnormal cortical regions.

Chapter 6 presents our method for the segmentation of multiple sclerosis lesions in white matter. The method is based on a geometric model of the brain tissues, which constraints the lesions to the white matter. The method improves over other techniques described in the literature by introducing a connectivity prior and a minimum lesion size constraint, both of which help to reduce the number of incorrectly identified lesions.

Chapter 7 presents a clinically oriented study concerning glioblastoma tumours. Following an initial segmentation of the brain vasculature, tumour vessel abnormality is quantified by considering the local vessel orientation both in tumours and in healthy vasculature. Interestingly, one of our measures correlates well with existing histological markers of vessel abnormality in tumours.

Chapter 8 summarises the contributions of this thesis, presents the main conclusions and suggests possible improvements to the techniques developed in this work.

Chapter 2

Basic Concepts of MRI

In this chapter a brief explanation of the basic concepts of Magnetic Resonance-Imaging (MRI) is provided, ranging from how an image is generated from atomic nuclei within the body, to the contrast mechanisms in clinical MRI. These concepts provide an important basis for the methods presented in the subsequent chapters of this thesis. For a more detailed description of the theoretical aspects, an interested reader may be referred to textbooks such as [1, 79].

2.1 Larmor Frequency and Net Magnetisation

Nuclear Magnetic Resonance (NMR) is a phenomenon based on the interaction between a magnetic field and an atomic nucleus. Nuclei possess an intrinsic quantum mechanical property, known as spin.

The spin value is quantised in integer multiples of $1/2$, and it is dependent on the atomic number, *i.e.*, spin $I = 0, 1/2, 1, 3/2, 2$. MRI uses mainly hydrogen ^1H nuclei ($I = 1/2$) due to their high concentration in fat and water of which the human body mainly is composed. The ^1H nucleus can be considered classically, as a vector (spin vector) with a defined orientation and magnitude, Figure 2.1(a).

If a macroscopic ensemble of hydrogen nuclei, *i.e.* arbitrary tissue or pure water, is considered, the spins axes are randomly orientated in all directions. The net magnetisation (sum over all spin vectors) is then 0 (Figure 2.1(b)). Instead, if the tissue is placed inside a static magnetic field \vec{B}_0 , the spin angular momentum of the individual nuclei




Figure 2.1: Example of nucleus rotating around its axes in (a). Arbitrary positions of nuclei inside a tissue in (b).




Figure 2.2: Precession of the magnetisation in the presence of the magnetic field \vec{B}_0 . The axis of rotation is parallel to \vec{B}_0 .

interacts with this static field causing a precession motion of the nuclei, where the axis of rotation is in the direction of the static magnetic field \vec{B}_0 .

The motion of each nucleus can be described by a coordinate system where x and y are perpendicular to, and z parallel to \vec{B}_0 , respectively (Figure 2.2). The perpendicular coordinates are non-zero and vary with time, while the z coordinate remains constant.

The frequency of precession is proportional to the strength of the magnetic field and is described by the Larmor equation:

$$w_0 = \gamma B_0 \quad (2.1)$$

where w_0 is the Larmor frequency (MHz), and γ is known as the gyromagnetic ratio and is constant for each nucleus. For a hydrogen nuclei, γ has the value of $2\pi \cdot 42.577$ MHz/Tesla.

In the magnetic field the hydrogen nuclei can occupy one of two energy states. In the lower energy state the spin vectors are oriented in the same direction of \vec{B}_0 (parallel), while in the higher energy state they are antiparallel to \vec{B}_0 . Therefore, the energy difference between the two energy levels ΔE is proportional to the magnetic field strength B_0 :

$$\Delta E = hw_0 = \frac{h\gamma B_0}{2\pi} \quad (2.2)$$

where h is the Planck's constant.

The spin population in thermal equilibrium is given by the Boltzmann distribution. The ratio between population on the higher (n_+) and the lower energy level (n_-) can be described by:

$$\frac{n_-}{n_+} = \exp\left(-\frac{\Delta E}{kT}\right) = \exp\left(-\frac{hw_0}{kT}\right) \quad (2.3)$$

where T is the temperature and k is the Boltzmann constant. At 1.5 Tesla (T) and room temperature $T = 300K$, $hw_0 \ll kT$, $n_-/n_+ = 0.999995$, *i.e.* the lower energy state has only additional five spins for every million.




Figure 2.3: Equilibrium orientation of \vec{M}_0 in (a); a pulse flips \vec{M}_0 by an angle α (b). The magnetisation \vec{M}_0 consists of two components (b)-(c): M_z or longitudinal component, and M_{xy} or transverse component. During relaxation to equilibrium M_z increases while M_{xy} decreases (c).

Given the number of excess spins, *i.e.* $n = (n_- - n_+)$, the total magnetisation M_0 is given by:

$$M_0 = \gamma hn = \frac{N\gamma^2 h^2 B_0}{kT} \quad (2.4)$$

where $N = n_- + n_+$ is the total number of spins in the ensemble. At equilibrium, the orientation of \vec{M}_0 is parallel to \vec{B}_0 .

2.2 Relaxation

To create an MR signal, the patient or sample is exposed to an electro-magnetic *Radiofrequency* (RF) pulse at a frequency matching the Larmor Frequency ω_0 . Since the resonance condition $\Delta E = h\omega_0$ is fulfilled, the energy contained in the RF pulse is absorbed by the nuclei, exciting them from the lower to the higher energy state.

In a classical description of the RF absorption, the RF pulse deflects \vec{M}_0 from its equilibrium (orientation parallel to \vec{B}_0) by an angle (flip angle α) depending on its duration and magnitude. The magnetisation now consists of two components: a longitudinal component M_z parallel to the magnetic field \vec{B}_0 and a transverse component M_{xy} perpendicular to the magnetic field and rotating with the Larmor frequency around the \vec{B}_0 axis. When the transmitter is turned off, the nuclei immediately start to return to the original equilibrium state, Figure 2.3. The precessing transverse magnetisation induces an MR signal $S(t)$, where t is the time, as a voltage in a nearby receiver coil (Faraday induction):

$$S(t) \propto \vec{M}_{xy}(t) = |\vec{M}_{xy}| \exp(i\omega_0 t) = M_x(t) + iM_y(t). \quad (2.5)$$

The process of recovery to equilibrium state is called relaxation, and is tissue-dependent. In MR, two of the commonly measured main relaxation processes are T_1 and T_2 . The spin lattice relaxation is described by the T_1 time required by the M_z to return to

Tissue type	T_1 [ms]	T_2 [ms]
Grey matter	1820	99
White matter	1084	56
Heart	1420	32
Liver	812	42
Blood	1932	275

Table 2.1: Relaxation times at 3T for different tissue types.

63% of its equilibrium value. On the other hand, the spin-spin relaxation process is described by time T_2 required by M_{xy} to decay to 37% of its initial value. Both effects are due to interaction processes between spins and a perturbation of the static magnetic field. Mathematically, the temporal evolution of the magnetisation components can be expressed as exponential decays:

$$M_z(t) = |M_0| \cdot (1 - \exp^{-t/T_1}) - M_z(0) \cdot \exp^{-t/T_1} \quad (2.6)$$

$$M_{xy}(t) = M_{xy}(0) \cdot \exp^{-t/T_2}. \quad (2.7)$$

By altering the time between excitation with an RF pulse and signal acquisition, T_1 and T_2 influence the contrast on an MR image. The relaxation times vary across tissue types due to biochemical differences, Table 2.1.

2.3 Image Creation

In order to generate an image, the MR signal is made spatially dependent through the application of magnetic field gradients. These gradients are small perturbations superimposed on the main magnetic field \vec{B}_0 . In the presence of a gradient, the magnetic field $B(\vec{r})$ is linearly dependent on the position $r = (x, y, z)$ inside the magnet:

$$B(\vec{r}) = B_0 + \vec{G} \cdot \vec{r} = B_0 + (G_x x + G_y y + G_z z) \quad (2.8)$$

where \vec{G} is a linear gradient along a direction. Three physical gradients are used, one in each of the x , y , and z directions. Each of the physical gradients, $G_x(t)$, $G_y(t)$ and $G_z(t)$, can be combined to produce one or more of the three functional gradients required to obtain an image. These functional gradients are known as: slice selection, readout or frequency encoding, and phase encoding. The combination of the gradients, RF pulses, data sampling periods, and the timing between each of them is known as a pulse sequence and used to acquire an image, Section 2.4. The presence of the gradients changes the nuclei precession frequency, *i.e.* $w_0(\vec{r})$, Equation (2.1):

$$w_0(\vec{r}) = \gamma(B_0 + \vec{G} \cdot \vec{r}). \quad (2.9)$$




Figure 2.4: Example of two slice selection gradients, G_1 and G_2 with a frequency $\omega_{\text{RF}} \pm BW$ of the RF pulse. With the stronger gradient G_1 the slice thickness (Δz_1) is narrower than the weaker gradient (Δz_2).

Equation (2.9) states that in the presence of a gradient field, a nucleus will precess at a unique frequency that depends on its exact position inside the gradient field.

The localisation of a 2D slice in space is accomplished using frequency selective excitations in conjunction with a slice selection gradient \vec{G}_{SS} . The gradient direction defines the slice orientation, while the gradient amplitude determines the slice thickness, and the frequency of the RF pulse determines the slice position along a vector normal to the slice. A frequency selective RF pulse has two parts: a carrier frequency and a narrow range of bandwidth (BW) determined by the RF pulse duration, Figure 2.4. The application of \vec{G}_{SS} has the effect of causing the nuclei in the sample to precess at different frequencies, depending on spatial location along the direction of the gradient. When an RF pulse is applied in the presence of such a gradient only the nuclei exhibiting a precession frequency matching that of the RF pulse carrier frequency $\omega_{\text{RF}} \pm BW$ satisfy the resonance condition.

After the selection of a single slice, the two in-plane directions x and y are spatially encoded by applying gradients along these directions. The measured signal is manipulated, as an example, by a gradient applied in the x direction, *i.e.* G_x , as follows:

$$S(t) = \int \int M_{x,y}(x, y) \cdot e^{-i\gamma x G_x t} dx dy. \quad (2.10)$$

The variable t can be expressed as the k -space k_x , defined as:

$$k_x = \gamma G_x t. \quad (2.11)$$

Substituting Equation 2.11 into Equation 2.10 yields:

$$S(k_x) = \int \int M_{x,y}(x, y) \cdot e^{-ik_x x} dx dy. \quad (2.12)$$




Figure 2.5: Acquired k -space (a), and the corresponding 2D Fourier transformed magnitude image (b).

Equation (2.12) shows the relationship between the measured signal $S(k_x)$ and the transverse magnetisation, *i.e.* $S(k_x)$ is the Fourier transformation of M_{xy} . In MR the signal acquisition, Equation (2.10), is acquired in the presence of a frequency encoding gradient \vec{G}_{RO} , which consists out of two individual gradients known as pre-phasing and readout gradient (perpendicular to \vec{G}_{SS}). The pre-phasing gradient is applied first. The readout gradient is of opposite polarity and has an integral of twice the size of the pre-phasing gradient. When the frequency encoding gradient is applied, the pre-phasing gradient causes the nuclei inside the sample to precess at different frequencies that depends on the position of the nuclei with respect to the gradient. The application of the readout gradient causes a reversal of the frequency change, thus making the precession of the nuclei synchronous. This is at a maximum when the integral of the readout gradient is equal to the negative integral of the pre-phasing gradient. This point in time is known as an ‘echo’. The application of \vec{G}_{RO} corresponds to a projection along the frequency encoding direction of the signal originating from the entire slice. This corresponds to a line in the k -space at $(k_x, 0)$.

To encode the second dimension, the so-called phase encoding gradient \vec{G}_{PE} , is used. This gradient is oriented perpendicular to the \vec{G}_{SS} and \vec{G}_{RO} and is played out in between them. It adds a linear phase, proportional to gradient strength, to the signal, which corresponds to a shift in k_y direction. To sample the complete slice, different phase encoding gradient strength are applied before playing out \vec{G}_{RO} . For a resolution of N pixels in phase encoding direction, signal acquisition has to be performed N times by gradually alternating \vec{G}_{PE} .

Measurement techniques can be divided into 2D and 3D categories based on the excited volume that it is used to generate the signal. For 3D acquisition an additional phase encoding gradient perpendicular to \vec{G}_{PE} and \vec{G}_{RO} is necessary to sample the third direction.




Figure 2.6: 2D slices of a representative brain. In (a) an MPRAGE image in (b) a FLAIR image, in (c) a T_2 weighted image, and in (d) TOF image contrast.

The MR signals are collected by a receive coil, digitalised as a function of time and converted into a complex format. These data are known as raw data or k -space data. Each signal point has a real and imaginary value. The image or the displayed matrix that it is commonly seen, is obtained by the Fourier transform of the complex raw data, Equation 2.12 (Figure 2.5). The maximum signal contained in k -space is located at the central point of the raw data matrix, where the echo has occurred and is primarily responsible for the image contrast. The outer portions of the raw data have relative low signal value and mainly provide edge definition to the resultant image.

2.4 Pulse Sequences

In order to acquire an image a combination of RF pulses and gradients are required to localise the position of the nuclei in the tissue of interest. The combination of all these components is known as *pulse sequence*. The pulse sequence contains the hardware instructions, such as RF pulses and gradient pulses, necessary to acquire the MR image in a specific manner. By changing sequence parameters, such as timing or RF pulses, different contrasts between the different tissue types can be achieved.

The repetition time (TR) is the time between successive RF excitations and determines the amount of T_1 weighting contributing to the image contrast. The echo time (TE) is the time between the excitation pulse and the signal echo. It determines the amount of T_2 weighting. For inversion recovery sequences, where the imaging sequence is preceded by an inversion pulse, the inversion time (TI) is defined as the time between the 180° inversion pulse and the image excitation pulse and also contribute T_1 weighting (T_1 -w) of an image. Short TI have a low minimum T_1 relaxation, while long TI times allow significant T_1 relaxation prior to the imaging excitation pulse. A correct choice of TI enables signal suppression of tissues based on their T_1 relaxation times. Flip angle defines the angle of rotation away from the equilibrium axis that \vec{M}_0 undergoes through RF absorption. The flip angle together with TR and TI values determines the amount of T_1 weighting present in the image.

Examples of pulse sequences for tissue identification, used in this thesis, are the Magnetisation-Prepared Rapid Acquisition Gradient Echo (MPRAGE), the Fluid-Attenuated Inversion Recovery (FLAIR), T_2 -weighted (T_2 -w) Turbo Spin Echo (TSE) and Time Of Flight (TOF), which are routinely used to image the brain, Figure 2.6.

The MPRAGE is related to the T_1 relaxation value and captures high tissue contrast while providing high spatial resolution.

The FLAIR sequence allows to suppress signal from fluid by adjusting the TI. Therefore, the signal from the CSF can be eliminated for brain imaging and the detection of subtle changes at the periphery of the hemispheres and in the periventricular region close to CSF can be determined. The usefulness of FLAIR sequences has been evaluated in diseases of the central nervous system such as multiple sclerosis and infarction [36]. An example of a T_2 weighted image is shown in Figure 2.6(c) where the CSF is hyperintense due to long T_2 time, which enables the identification of brain tumours and oedema.

Another example of a pulse sequence, which is very useful to analyse vasculature in the brain, is called Time Of Flight (TOF) MR angiography. This type of sequence saturates the signal of static tissue of a slice or volume using high flip angles and short TR. Unsaturated inflowing blood creates a bright vascular image without use of contrast media. Placing an additional saturation slab on one side parallel to the slice can selectively destroy the MR signal from the in-flowing blood from this side of the slice. This allows the technique to be flow direction sensitive and to differ between arteries and veins. In Figure 2.1(d) the TOF image shows the bright arterial blood, whereas static tissue is almost completely saturated.

Chapter 3

Image Segmentation

The segmentation of images is one of the central problems in image understanding both for computer vision and medical imaging applications. The goal of image segmentation is to separate or detect an object of interest, or foreground, from the background. In multi-object segmentation, the image is divided into different components, where the number of components depends on the problem to be solved. In this case the segmentation can also be regarded as a process of assigning a label to each object. Given the image $I : \Omega \rightarrow \mathbb{R}^c$, where $c \geq 1$ is the number of input images, $\Omega \subset \mathbb{R}^3$ is the domain, and the disjoint regions $\{R_1, R_2, \dots, R_m\}$, the problem of image segmentation can be formulated as:

$$\Omega = \bigcup_{i=1}^m R_i, \quad R_i \cap R_j = \emptyset, \forall i \neq j$$

where each region can represent an object. A segmentation method can be based on intensity information, texture, motion, shape, spatial relationship or other information extracted from the image. For example an object may be detected due to its high contrast against the surrounding objects, or in the case of video analysis the object of interest could be detected by its motion in consecutive frames. In medical images, typically, the object to be segmented is an organ, such as the liver, or the heart ventricles, or a lesion, or also a specific tissue, like brain grey or white matter. For such images, the most discriminative information is in general the intensity difference between neighbouring tissue types. Organ shape can also be a very helpful feature in guiding the segmentation process, especially when the object is surrounded by structures with similar intensities (Figure 3.1).

A very important application of medical image segmentation is to identify or to monitor pathologies such as tumour growth or degenerative diseases. For a large number of diseases, the boundaries between healthy and pathological tissue are rather smooth, therefore the precise identification of e.g., tumours and other lesions is a difficult task which requires highly experienced radiologists. Often, signal intensities are not discriminative enough and additional information needs to be considered. Typically, a comparison with healthy subjects can help in seeing the difference between the pathological organ/tissue and normal tissue. In addition, tumour tissue can have a very different texture as compared to the healthy surrounding tissues. In particularly difficult cases, specific information on the location of the pathology can be very helpful to exclude potential errors and in targeting the segmentation to a limited region of interest.




Figure 3.1: Examples of medical images: (a) brain image where the intensity contrast already discriminates between grey (GM) and white matter (WM), (b) heart image where the shape prior (yellow ring) can help to segment the left ventricle, (c) brain tumour shows a different texture from the normal appearing tissues, and (d) brain lesions localised in the white matter.

In the following sections, a Bayesian formulation of the segmentation problem is presented, followed by an overview of the most common segmentation methods with a particular attention to the algorithms that have been applied to brain MR images. Lastly, methods which are formulated as denoising approaches, but which can also be used for image segmentation, are described.

3.1 Bayesian Formulation

Many approaches model the image segmentation problem using a probabilistic framework. Given an input image $I : \Omega \rightarrow \mathbb{R}^c$ and the labelling $f : \Omega \rightarrow \mathcal{L}$, where \mathcal{L} is the label space, one wishes to maximise the posterior probability with respect to f :

$$P(f|I) = \frac{P(I|f)P(f)}{P(I)} \quad (3.1)$$

where $P(I|f)$ represents the likelihood that an image I originates from the labelling f , while $P(f)$ is the prior probability of the labelling. The maximisation of Equation (3.1), is equivalent to the minimisation of its negative logarithm:

$$-\log P(f|I) = -\log(P(I|f)) - \log(P(f)) + \log(P(I)). \quad (3.2)$$

The $\log(P(I))$ is treated as constant value since it denotes the probability that the observed image is possible. Let $E_{\text{App}}(f) = -\log(P(I|f))$ represent an ‘appearance energy’ term corresponding to the intensity information of the image, and let $E_{\text{Knowledge}} = -\log P(f)$ denote the energy associated with the *a priori* information about the object to be segmented e.g., information concerning the relative tissue location. Then Equation (3.2) can be rewritten as an energy minimisation problem:

$$E(f) = E_{\text{App}}(f) + \beta E_{\text{Knowledge}}(f) \quad (3.3)$$

where the parameter β weights the two energies. With this formulation the objective is to find a labelling f which minimises Equation (3.3). The above definitions show that the appearance energy is related to the probability of the grey values given the labelling f , while the knowledge or internal energy is described by the negative log-likelihood of the *a priori* probability of a given f . Note that if the label space \mathcal{L} contains only two labels, then the labelling f represents a binary segmentation, or also called mask. On the other hand, in a multi-labelling framework, the label space contains more than two labels and the region of the image identified by a label is defined as a region $R_\ell = \{x \in \Omega | f(x) = \ell\}$, $\forall \ell \in \mathcal{L}$.

There is a variety of global and local optimisation techniques that can be applied to maximise Equation (3.1) or to minimise Equation (3.3) [11, 52]. In Chapter 6 it is described how it is possible to achieve the segmentation of multiple sclerosis brain lesions using Equation (3.3), and how a local optimal solution is computed.

3.2 Segmentation Methods

A number of automated and semi-automated segmentation algorithms have been proposed in the literature [63, 100], with a large interest on medical images [107, 114]. In this section the main focus is on some of these methods and on the description of their main characteristics, advantages and limitations.

The application of these methods to MR medical images can be challenging due to several artefacts present in these images: noise, inhomogeneity, and partial volume effects. The inhomogeneity, also called bias-field, is the presence of smoothly varying intensities inside tissues, and the partial volume effect denote the contribution of multiple tissue to the intensity of a single voxel. These artefacts are generally modelled in the segmentation algorithms.

3.2.1 Thresholding

This class of methods encompasses some of the simplest approaches to perform segmentation, which regard pixel intensity as the main feature to discriminate between two or multiple regions. Thresholding methods attempt to find an intensity value, called threshold, which separates the desired region or tissue from neighbouring ones. The segmentation is then achieved by grouping all the pixels with an intensity greater than the threshold into one class, and the rest into a second class:

$$f(x) = \begin{cases} 1, & \text{if } I(x) > T \\ 0, & \text{if } I(x) \leq T \end{cases}$$

where T is the threshold value, 0 and 1 represent two different classes. The selection of two or more thresholds is known as multithresholding and separates the image into mul-

multiple regions. The formulation of a thresholding method can be seen as the minimisation of Equation (3.3) where the internal energy is not considered, therefore $\beta = 0$:

$$E(f) = E_{\text{App}}(f).$$

This technique was applied, for example, to cardiac MR segmentation [119] and it is commonly used as a pre-selection step for more advanced segmentation methods, such as for brain extraction [103, 121]. A main limitation of thresholding approaches is the lack of spatial information related to the segmented region, which causes these methods to be particularly sensitive to noise and bias field. Both of these artefacts distort the intensity histogram and therefore, hinder the segmentation. Evolved thresholding methods have been proposed, which consider local intensity differences [84] as well as connectivity priors [76].

3.2.2 Feature-based Techniques

The most intuitive feature that can be extracted from an image is the pixel/voxel intensity, as for the thresholding methods (Section 3.2.1). It is common to model the distribution of such intensities to discriminate pixels which do not fit into the chosen model and might therefore be identified as outliers or disease pixels. One of the commonly used models to represent image intensities is the Gaussian distribution:

$$\mathcal{N}(I(x)|\mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(I(x) - \mu)^2}{2\sigma^2}\right) \quad (3.4)$$

where $I(x)$ represents the intensity of the pixel x , μ is the mean value and σ is the standard deviation of the pixel intensities. In case there are multiple input images the 1D-Gaussian model becomes multi-dimensional, where the number of input images is represented by c . The intensity distributions in MR images, usually, cannot be described by a single Gaussian distribution due to the presence of multiple tissues. Therefore a more flexible way to model signal intensities is by means of a weighted sum, or linear superposition of Gaussians, *i.e.* a Gaussian Mixture Model (GMM):

$$p(x) = \sum_{k=1}^K \pi_k \mathcal{N}(x|\mu_k, \sigma_k) \quad (3.5)$$

where K is the number of Gaussians and π are the mixture parameters weighting the influence of each Gaussian, subject to $\sum_{k=1}^K \pi_k = 1$. This last constraint forces $p(x)$ in Equation (3.5) to have a maximum value of 1 and therefore to be a valid probability. Every Gaussian density $\mathcal{N}(x|\mu_k, \sigma_k)$ has its own mean μ_k , and variance σ_k . When the number of input images is greater than one, *i.e.* $c > 1$, then the Gaussian density becomes $\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$ where the bold fonts represent vectors and $\boldsymbol{\Sigma}_k$ is the covariance matrix which, intuitively, is a generalisation of the variance to multiple dimensions. Note




Figure 3.2: Examples of intensity in (a), texture in (b) and motion segmentation in (c).

that for the MR image contrasts used in this thesis each tissue type can be modelled by a Gaussian.

There are other ways to describe and to characterise an image. For example it is possible to analyse the texture, the motion or the shape of the object/tissues that are contained in the image, Figure 3.2. Features can be classified in three different types: intensity, texture or shape features. Intensity features are the most commonly used. They are based on the pixel intensity values and are therefore derived from the histogram of pixel intensities. Examples of intensity features are the mean, variance, skewness, kurtosis, and the entropy. The mean is the average value of the intensities, and the variance represents the intensity variation around the mean. Skewness represents the symmetry of the intensity histogram around the mean value; the kurtosis is the flatness of the histogram; and the entropy represents the histogram uniformity.

Texture features provide a higher-order description of the image and contain information about the spatial distribution of the variation of the grey level. In other words they represent the pattern of the image. These features can be scalar quantities, discrete histograms, or any type of characteristic which gives texture properties of the image, such as spatial structure, contrast, orientation, and brightness. Texture features have been widely used to discriminate between tumours and healthy tissues in the brain [59, 65]. An interesting characteristic of certain tumours, such as glioblastoma or astrocytoma, is that they lack texture, as compared to normal appearing tissues, whereas within the tumour it is hard to discriminate between cysts, necrotic and solid parts.

Shape features provide geometrical information about an object which is part of a given image, therefore they represent a prior knowledge about the object to be segmented. Examples of these features include the centroid, area, perimeter, and orientation of the object of interest.

In order to perform a feature-based segmentation, the first step is to extract one or multiple features from the input image. Then a classification method is required to separate the features into one or multiple classes. A simple way to achieve the classification can be by applying a thresholding technique (Section 3.2.1); more advance methods will be presented in Section 3.2.3. An example of this procedure will be presented in




Figure 3.3: Example of a brain image with two patches. The top one is in the white matter and the bottom one lies between the ventricles and the white matter.

Chapter 5 where we propose a feature-based approach to discriminate between epileptic lesions and normal-appearing tissue in the brain.

Recently a feature called ‘patch’ has been successfully used in medical imaging applications [149]. A patch is a part of the image with a specified size (Figure 3.3). Patches represent the local structure, shape and image intensity of the object they are taken from. Images, especially medical ones, reveal similarities across multiple locations, therefore, a subset of the total amount of patches can be used to represent the image. The process of learning the relevant patches is called ‘dictionary learning’ [138].

3.2.3 Classifiers

Classifier methods are pattern recognition techniques that are capable of separating a feature space, not necessarily intensity features, derived from an image into two regions, e.g. tumour and non-tumour voxels. A one-dimensional feature space may simply consist in the pixel intensities of an image, and a two-dimensional feature space can be created using, as an example, the intensities from two different image contrasts. See Section 3.2.2 for more details on image features. Classifiers are supervised method that require a training set, or rather a subset of input data for which the classification, or label, is known. This set is then used to ‘train’, or learn, the parameters of the classifier. After the training phase, the rest of the input data are classified. This classification process can be binary (two labels) or can assign multiple labels.

One of the simplest classifiers used in medical imaging is the K-Nearest Neighbours (K-NN) classifier. This is a non parametric classifier which discriminates the pixels based on the most common class among its K nearest neighbours. It is very easy to implement and requires only an integer value K, a set of labelled examples and a metric to measure the distance between two input data sets. Some disadvantages of this classification method are the poor performances when the training set is small since it tends to over-fit the input data, its sensitivity to irrelevant or redundant features as all features contribute

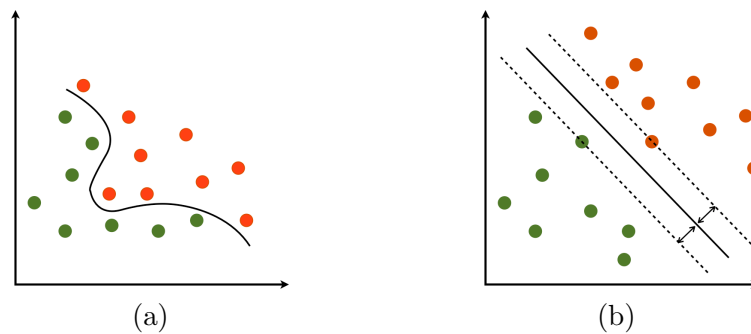


Figure 3.4: Example of input data classified by the black curve in (a). Mapping of the input data in (a) into a higher dimensional space (b) where the data are separated by a hyperplane (solid black line). The input data lying on the dashed line are the support vectors, and represent points with equal distance to the hyperplane. The distance between the hyperplane and the dashed lines is called *margin*.

equally to the similarity, and all the training samples must be stored. This type of classifier was applied in combination with more advanced techniques to the segmentation of gliomas [68].

Support Vector Machines (SVM) are more advanced type of classifiers, which perform a binary linear classification of the input data. In order to achieve a non-linear classification, a mapping of the input data into a high-dimensional space is performed with a kernel function. In the high dimensional space a separating hyperplane is identified as the one having the largest distance to the nearest training-data point of any class, Figure 3.4. The SVM technique is widely used in object detection and recognition as well as text recognition and biometrics [125]. Additionally, it has been applied in medical imaging for brain tumour segmentation [158], lesion segmentation [45, 75], and disease identification [89, 160]. On the other hand, SVM assumes that the data are independently and identically distributed, which is not the case for medical images, since both intensity inhomogeneity and noise can be spatially dependent. Therefore, SVM is usually combined with other methods to consider spatial information [77]. Note that when a Gaussian kernel is used for the mapping, SVM is similar to the K-NN classifier, therefore it has similar properties.

A commonly used parametric classifier is the maximum-likelihood, or Bayesian, classifier. It assumes that the pixel intensities are independent samples taken from a mixture of probability distributions, usually Gaussian (Equation (3.5)). The training phase is obtained by taking representative samples for each density function and estimating the corresponding Gaussians parameters μ_k, Σ_k accordingly. The classification of new data is performed assigning to each pixel the class with the higher posterior probability.

In general classifiers are not iterative, therefore once the classifier is trained the new input data can be directly assigned to a class. A disadvantage is that they generally do not perform any spatial modelling. This weakness has been addressed by using classifiers in conjunction with other methods which take the spatial relationship into account. In

addition they require a manually labelled training set which can be time consuming. On the other hand, in the case of application to brain images, if the same training set is used for different subjects, this can lead to biased results as anatomical differences between subjects are not taken into account.

3.2.4 Clustering

Clustering methods perform a segmentation grouping a set of pixels/voxels in such a way that the pixels/voxels in the same group, called cluster, are ‘closer’ to each other than the pixels/voxels in other clusters. This class of methods consider spatial consistency and can be seen as an improvement over the thresholding technique, Section 3.2.1, where the pixels are clustered based only on a threshold value and no spatial information is included. Clustering algorithms also belong to the class of machine learning techniques but, unlike the classifier methods, they are unsupervised methods. Therefore they do not require a training set to learn the classification parameters, as in the case for classifiers (Section 3.2.3), but instead they iterate between segmentation/labelling of the input data and estimation of the classification parameters. Especially, clustering methods are training themselves with the available data.

Commonly used clustering algorithms are K -means, fuzzy c -mean and expectation maximisation [40]. The K -means clustering method aims to partition the data into K clusters by iteratively computing the mean for each cluster and subsequently assigning the pixels to the cluster with the nearest mean. The objective function can be written as:

$$J = \sum_{n=1}^N \sum_{k=1}^K r_{nk} \|\mathbf{x}_n - \boldsymbol{\mu}_k\|^2 \quad (3.6)$$

where N is the number of input data points, $\boldsymbol{\mu}_k$ represents the mean of the k th cluster, and $r_{nk} \in \{0, 1\}$ is a binary variable which indicates if input data n belongs to the cluster k . Note that \mathbf{x} can represent, as for the classifier methods, a vector with multiple features. The objective is to find the $\{r_{nk}\}$ and $\{\boldsymbol{\mu}_k\}$ for which J is minimised.

The fuzzy c -means algorithm is a generalisation of the K -means clustering and allows for a ‘soft’ segmentation, *i.e.* the coefficients r_{nk} can assume any value in the interval $[0, 1]$.

The Expectation-Maximisation (EM) algorithm for image segmentation is generally applied to Gaussian Mixture Models (GMM), see Section 3.2.2. The EM is a method that seeks to find the parameters which maximise the likelihood of the model. The image intensities are assumed to be generated from Gaussian distributions, and the EM provides a segmentation/partition of the image based solely on these intensities. Assuming that the data are identically and independently distributed, the log-likelihood of Equation (3.5), for multi-dimensional data, is given by:

$$\ln p(\mathbf{X}|\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \sum_{n=1}^N \ln \left\{ \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}_n | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \right\} \quad (3.7)$$

where $\mathbf{X} = \mathbf{x}_1, \dots, \mathbf{x}_N$ are the input data, $\mathcal{N}(\mathbf{x}_n | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$ is a Gaussian random variable with mean $\boldsymbol{\mu}_k$ and covariance matrix $\boldsymbol{\Sigma}_k$, and the mixing coefficients $\boldsymbol{\pi}$ represent the respective weights of the Gaussian distribution in the mixture. Note that the maximum likelihood approach, as applied to GMM models, has singularities. Let consider the case where $\mathcal{N}(\mathbf{x}_n | \boldsymbol{\mu}_k, \sigma_k^2 I)$, where I is the identity matrix, then in the limiting case $\sigma_k \rightarrow 0$ this term goes to infinity. Therefore the maximisation of the log-likelihood function is not a well-posed problem because such singularities can always happen when a Gaussian collapses into one point data and σ tends to zero. A possible solution to this problem is to employ certain heuristics, such as detecting when a Gaussian is collapsing and therefore resetting its mean to a random value and the variance to a large value. Another possibility is to remove from the model the collapsed Gaussian.

In order to maximise Equation (3.7), the derivatives with respect to each parameter are computed. Computing the derivative with respect to the mean $\boldsymbol{\mu}_k$ and setting it equal to zero, the following equation is obtained:

$$\sum_{n=1}^N \frac{\pi_k \mathcal{N}(\mathbf{x}_n | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)}{\underbrace{\sum_j \pi_j \mathcal{N}(\mathbf{x}_n | \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)}_{\gamma(z_{nk})}} \boldsymbol{\Sigma}_k^{-1} (\mathbf{x}_n - \boldsymbol{\mu}_k) = 0 \quad (3.8)$$

where $\gamma(z_{nk})$, known as the responsibilities, represent the posterior probability that \mathbf{x}_n was generated by the k th component. Multiplying both sides of the above equation by $\boldsymbol{\Sigma}_k$ yields (after rearrangement):

$$\boldsymbol{\mu}_k = \frac{1}{N_k} \sum_{n=1}^N \gamma(z_{nk}) \mathbf{x}_n, \quad \text{where } N_k = \sum_{n=1}^N \gamma(z_{nk}). \quad (3.9)$$

The variable N_k is related to the actual input points which do contribute to cluster k . The derivative of Equation (3.7) with respect to $\boldsymbol{\Sigma}_k$ results in:

$$\boldsymbol{\Sigma}_k = \frac{1}{N_k} \sum_{n=1}^N \gamma(z_{nk}) (\mathbf{x}_n - \boldsymbol{\mu}_k) (\mathbf{x}_n - \boldsymbol{\mu}_k)^T. \quad (3.10)$$

The maximisation with respect to the mixing coefficients π_k is subject to the constraint $\sum_k \pi_k = 1$. This can be achieved with the use of the Lagrange multipliers, λ , and maximising the following expression:

$$\ln p(\mathbf{X} | \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) + \lambda \left(\sum_{k=1}^K \pi_k - 1 \right). \quad (3.11)$$

Computing the derivative with respect to π_k and setting it equal to zero results in:

$$\sum_{n=1}^N \frac{\mathcal{N}(\mathbf{x}_n | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)}{\sum_j \pi_j \mathcal{N}(\mathbf{x}_n | \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)} + \lambda = 0. \quad (3.12)$$

Multiplying both sides by π_k , and taking the second Equation (3.9) into account, gives $N_k + \pi_k \lambda = 0$. Next, summing over k , $\lambda = -N$ is obtained, which gives:

$$\pi_k = \frac{N_k}{N}. \quad (3.13)$$

The results in Equations (3.9), (3.10), and (3.13) show that it is not possible to directly compute the model parameters, since they depend on the responsibilities $\gamma(z_{nk})$, and the responsibilities themselves depend on the model parameters. Therefore a possible way to estimate the maximum likelihood parameters is in an iterative fashion, *i.e.* the Expectation-Maximisation (EM) algorithm.

The EM algorithm starts with some random values for means, covariance matrices and mixing coefficients. Then two steps are iterated in succession: the Expectation, or E-step, and the Maximisation, or M-step. During the E-step, the current model parameters are used to estimate the posterior probabilities, or responsibilities, $\gamma(z_{nk})$, while in the M-step the model variables are re-estimated using Equations (3.9), (3.10), and (3.13). The general EM steps are summarised in Algorithm 1.

The final classification is performed considering the estimated GMMs: every input data will belong to the class with the highest log-likelihood.

A drawback of the EM algorithm is that it requires many iterations to reach convergence. Therefore it is common practise to use the K -means algorithm in order to find a suitable initialisation for the GMM model parameters. Additionally when the log-likelihood function has multiple local maxima, the EM algorithm does not guarantee that the largest one is actually found. The combination of K -means with the EM

Algorithm 1

- 1: Choose an initial value for $\boldsymbol{\mu}_k$, $\boldsymbol{\Sigma}_k$ and π_k . Evaluate the initial log likelihood value, Equation (3.7)
- 2: **repeat**
- 3: E-step: expectation of the posterior probabilities:

$$\gamma(z_{nk}) = \frac{\pi_k \mathcal{N}(\mathbf{x}_n | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)}{\sum_{j=1}^K \pi_j \mathcal{N}(\mathbf{x}_n | \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)}$$

- 4: M-step: maximise the expectation of the log-likelihood in the model parameters, Equations (3.9), (3.10), and (3.13).
 - 5: **until** Convergence of either the parameters or of the log-likelihood.
-




Figure 3.5: Examples of MRF graphs. Vertices (circles) are connected by the edges (black lines). In (a) a 4-connected graph. In (b) an 8-connected graph which can be used to obtain better geometrical details.

algorithm is also used in the proposed method that will be presented in Chapter 6 to model the different brain tissue intensities.

The EM clustering method is widely used in combination with atlas-based approaches, Section 3.2.6, to get a soft segmentation of the brain tissues [4].

3.2.5 Markov Random Fields and Graph-based Techniques

A Markov random field (MRF), is a statistical model represented by a set of random variables $\mathbf{X} = (X_1, X_2, \dots, X_N)$ that can be used with segmentation methods. A MRF models spatial relationships between neighbouring pixels. Such local correlations are particularly useful for modelling image properties, and in medical imaging they are used to represent the fact that a pixel has a high probability of belonging in the tissue/class of its neighbours. Therefore an anatomical structure composed by a single pixel has a very low probability of occurrence in a MRF model.

Markov random fields can be modelled using graphs. A graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ consists of a set of vertices $\mathcal{V} = (v_1, v_2, \dots, v_N)$, which can correspond to the image pixels/voxels, and a set of edges \mathcal{E} which represent the connection between two vertices: $(v_i, v_j) \in \mathcal{E}, v_i, v_j \in \mathcal{V}$. In MRF the edges are considered undirected, so (v_i, v_j) and (v_j, v_i) refer to the same edge which is typically qualified as a weight $w(v_i, v_j)$. A simple representation of a graph can be seen in Figure 3.5.

By definition, the variables in the MRF need to follow the *Markov Property*, which states that in a stochastic process the conditional probability distribution of future states for a process depends only on the present state and not on the precedent sequence of state transitions [85]. When applied to an image, this property may be interpreted as restricting the local dependence of any given pixel intensity to the intensities of just close neighbouring pixels. Therefore the conditional probability of a random variable can be expressed as:

$$P(X_i | X_k, k \neq i) = P(X_i | X_k, k \in N_i) \quad (3.14)$$

where N_i represents the neighbourhood of X_i .

Markov random fields are often incorporated into brain tissue clustering segmentation algorithms such as K -means by using a Bayesian prior model for brain tissues [58, 159]. The segmentation is obtained by maximising the posterior probability of the segmentation given the image data. Additionally, MRFs have been applied to brain tumour segmentation in combination, for example, with support vector machines [77] and atlas models [9].

A difficulty associated with the use of MRF models is the proper selection of the parameters controlling the strength of spatial interactions [85]. Too high values can result in an excessively smooth segmentation, hence in the loss of important structural details. In addition, MRF methods usually require computationally intensive algorithms. Despite these disadvantages, MRFs are widely used not only to model segmentation classes, but also to model intensity inhomogeneities that can occur in MR images [159] as well as texture properties.

In the remainder of this subsection the graph theory presented above is used to perform image segmentation, *i.e.* the graph-based approaches. The inclusion of MRF into graph-based segmentation techniques will be described in Chapter 4.

Given an image $I : \Omega \rightarrow \mathbb{R}$, the pixels $x \in \Omega$, which are needed to be separated, are inserted into the graph as vertices, *i.e.* there is a one-to-one correspondence between the vertices and the pixels. Computing an image segmentation based on a graph means to find mutually exclusive components, such that each component A is a connected graph $G' = (\mathcal{V}', \mathcal{E}')$, where $\mathcal{V}' \subseteq \mathcal{V}$, $\mathcal{E}' \subseteq \mathcal{E}$ and \mathcal{E}' contains only edges connecting vertices in \mathcal{V}' . Therefore all the components (A_1, A_2, \dots, A_k) in which the graph is partitioned must satisfy $A_i \cap A_j = \emptyset, \forall i \neq j$ and $A_1 \cup \dots \cup A_k = G$. Note that a component is ‘a part’ of the image I , hence a subset of the pixels.

In graph theory a cut of the graph is related to the edges which are cut to identify, in binary segmentations, two disjoint sets A_i, A_j . The cost of the cut is defined as:

$$\text{cut}(A_i, A_j) = \sum_{v \in A_i, u \in A_j} w(v, u). \quad (3.15)$$

In order to find the best cut the problem is usually formulated as an energy optimisation [104]. Wu and Leahy [152] were the first to apply the concept of graph cuts to image segmentation. They formulated Equation (3.15) as a minimisation of a cost function, which gives the so-called minimum cut. The method seeks to find the minimum cut which separates existing segments in the graph. The minimum cut criterion is intuitive but it is biased towards small components. This can be observed from Equation (3.15) where the cost of the cut increases with the number of edges that need cutting. Therefore the algorithm favours the separation of the image into small sets of points.

This problem was addressed by the normalised cut (*Ncut*) algorithm proposed by Shi and Malik [117]. Instead of looking at the total edge weights connecting two partitions,

they propose to compute the cost based on the edge weights of all the vertices in a set, *i.e.* edges between A_i to all the vertices in the graph:

$$Ncut = \frac{cut(A_i, A_j)}{assoc(A_i, \mathcal{V})} + \frac{cut(A_i, A_j)}{assoc(A_j, \mathcal{V})} \quad (3.16)$$

where $assoc(A_i, \mathcal{V}) = \sum_{u \in A_i, v \in \mathcal{V}} w(u, v)$ and $assoc(A_j, \mathcal{V})$ is similarly defined. With this formulation the $Ncut$ value is no longer small for partitions with few vertices, therefore the problem of minimum cut bias is avoided.

Cox et al. incorporated the interior region and the boundary information into the minimisation [34]. The minimisation considers the ratio between the cost of the exterior boundary and the one of the interior part:

$$RegionCut(A_i, A_j) = \frac{length(A_i)}{weight(A_i)} \quad (3.17)$$

where $length(A_i)$ represents the length of the boundary and $weight(A_i)$ is the cost of the segmented area. This formulation favours large regions and the denominator term is balanced by the cost of the contour. For example, if the edge costs are relatively low, then a larger region enclosing low intensity edges may be preferred over a small region with high contrast.

The minimum mean cut proposed by Wang and Siskind addresses this problem by defining an edge-based function [147]:

$$MeanCut(A_i, A_j) = \frac{cut(A_i, A_j | w(u, v))}{cut(A_i, A_j | 1)} \quad (3.18)$$

where the numerator represents the cut cost given the edge weights $w(u, v)$, and the denominator is similarly defined with all edge weights set equal to 1. This approach finds the best cuts which minimises the average edge weight on the boundary.

The methods presented above have been used in different image segmentation applications. The main common drawback of all these approaches is their lack of spatial information. Studies in psychology demonstrated that contextual knowledge is crucial for understanding visual information. The MRF theory provides a useful way of modelling contextual information such as the relation between neighbouring pixels. One of the most prominent algorithm which includes MRF, is the max-flow/min-cut presented by Boykov and Jolly [17], usually referred to as graph-cuts. One of the main contributions on this work, presented in Chapter 6, is based on this theory and its evolutions. Therefore, a complete description of the necessary theory to understand the method we propose in Chapter 6, is presented in Chapter 4.

3.2.6 Atlas-based Approaches

An atlas, or template, is a reference image which contains information about the anatomic region to be segmented. A number of atlases have been created during the last

years, in particular for the brain tissues. First a brief description on how to generate an atlas is presented, and subsequently its application to image segmentation.

Here the focus is on how to generate a brain atlas, but the same technique can be applied to any other anatomical organs. There are two types of atlases: topological and probabilistic [25]. Topological atlases are based on a single subject, which is considered to be representative of the object to be segmented, in terms of size, shape or intensity. One of the first topological atlas is the Talairach atlas which was proposed to identify deep grey matter structures [135]. Atlases based on a single subject are not representative for anatomical variability. Therefore a variety of statistical, or probabilistic, atlases have been proposed and are based on a population. This kind of atlas is in continuous evolution. They are relatively flexible since a new subject can be easily incorporated or the input images can be subdivided based on specific criteria such as age or gender. Examples of this kind of atlas can be found in [43, 74]. The creation of a probabilistic atlas is a process which includes segmentation and registration. Here, only a brief introduction to registration is given, for a detailed review on registration techniques see [38, 163].

The first problem that occurs when analysing multiple image contrasts is to assure that the images are registered. Two images are assumed to be registered when they can be perfectly superimposed, so that any structure present in a given image appears in exactly the same location in the other image. When analysing complex structure, such as brain tissues, this constraint becomes very important since a shift of just a few voxels can result in different tissues being wrongly aligned, hence leading to an incorrect segmentation.

The various types of registration approaches can be classified according to their degrees of freedom. The most common general transformations are rigid, affine, and non-linear. Rigid transformations allow for translation, rotation along the three Cartesian axes and a global scaling factor. Affine registration use more complicated distortions such as shearing. Finally, non-linear registration allows for more complex transformations, such as stretching and dilating, with an increased number of degrees of freedom. The latter transformations are very useful for the alignment of abdominal organs such as liver and lungs [108], whereas rigid or affine transformations are often sufficient for the brain.

The first step of the atlas generation, performs a registration, usually affine, of all the subject data into a common space. Following a standardisation of image intensities, the set of registered images is averaged voxel-wise and the resulting image is a probabilistic map. This procedure can be applied to every kind of image contrast available, which will produce different atlases. A very useful, hence widely used, type of atlas is one that describes tissue probabilities. These are computed in the same way as described above and require that for each subject data the tissue segmentation is available. Usually a manual annotation of the different tissues is performed. An example of a T_1 -weighted intensity atlas including tissue probability maps of grey matter, white matter and cerebrospinal fluid is shown in Figure 3.6.




Figure 3.6: Example of a T_1 -weighted intensity atlas in (a), and the tissue probability maps for grey, white matter and cerebrospinal fluid in (b), (c) and (d) respectively.

The process of generating an atlas is very time consuming. A well-known problem during the registration is that, depending on the method of the registration employed, certain brain anatomical structures may appear consistently deformed. Additionally a substantial number of healthy subjects are needed in order to generate statistically significant probability maps, and for specific anatomical regions this can be a challenging task.

When atlases are used to perform segmentation the initial step is to find a geometrical transformation $\tau \in T$, where T is the space of all possible transformations, which optimally aligns the atlas with the input image as closely as possible. Given the input image $I : \Omega \rightarrow \mathbb{R}$ and an intensity atlas $\pi^I : \Omega \rightarrow \mathbb{R}$, the process of finding τ , can be formulated as:

$$\arg \max_{\tau} m(I, \tau(\pi^I)) \quad (3.19)$$

where m is the metric used to evaluate the similarity between the input image and the transformed atlas. This initial registration step is performed using an intensity atlas, and the same transformation is applied to the tissue probability maps. As a result, these registered maps are the segmentation. Since the brain has a large anatomical variability, this single step is not guaranteed to provide alone a correct segmentation of the region of interest. Therefore, a possible way to improve the tissue segmentation is to integrate the atlas information into the formulation. Specifically, with reference to Equation (3.3) the atlas can be encoded into the $E_{\text{Knowledge}}$ as it gives spatial information to the voxels.

Most brain segmentation methods rely on a good initial atlas-based segmentation step and, subsequently, they implement different techniques to produce the final tissue segmentation [92, 144]. Van Leemput et al. [144] propose to initialise the segmentation using an atlas of tissue probabilities and to constrain the classification process by weighting these probabilities for each voxel. This method also accounts for neighbourhood information with a Markov Random Field (MRF) model (Section 3.2.5).

Other methods consider the registration and the segmentation as two closely related and mutually dependent steps. The registration with an atlas can benefit from an

initial segmentation step, and the brain segmentation requires prior information obtained from the registration to the atlas. Ashburner and Friston [5] proposed an algorithm in which they iteratively perform the registration and the segmentation steps. The energy that they propose to minimise contains the registration parameters, the intensity model formulation (GMM) and the bias field:

$$E = - \sum_{i=1}^N \ln \left(\frac{\rho_i(\beta)}{\sum_{k=1}^K \gamma_k b_{ik}(\alpha)} \sum_{k=1}^K \frac{\gamma_k b_{ik}(\alpha)}{2\pi\sigma_k^2} \exp \left(-\frac{(\rho_i(\beta)y_i - \mu_k)^2}{2\sigma_k^2} \right) \right). \quad (3.20)$$

In Equation (3.20), the total number of pixels in the image I is represented by N and the intensity of the pixels is denoted by y_i . The GMM model depends on the K Gaussian distributions, the means μ_k , the standard deviations σ_k and the mixing parameters γ_k . The bias field is encoded in $\rho(\beta)$ where β are the unknown parameters, while the registration parameters are represented as the α variable. Since the minimisation involves mutually dependent parameters, the authors proposed to use an iterated conditional model. The minimisation starts with an initial estimate and then iterates until a local optimal solution is found. Each iteration involves alternating between estimating different groups of parameters while keeping the others fixed. The GMM parameters are updated using an EM algorithm, while keeping the bias field and the registration parameters constant. The deformation of the probability maps is re-estimated while the GMMs and the bias field are kept fixed. A similar approach is used to estimate the bias field. For further implementation details, the reader is referred to [5, 105].

The algorithm proposed by Ashburner and Friston [5] is freely available as part of the Statistical Parametric Mapping (SPM) tool [124]. In Chapter 5 it is used as a pre-processing step for our proposed method for the detection of epileptic lesions. A modification, instead, has been developed for the segmentation of the brain grey matter, Section 3.2.7.

3.2.7 Active Contours and Level-set Methods

Image segmentation can be performed using active contours (or snakes) models. For a comprehensive review, the reader is referred to [98]. The basic idea of contour models is to evolve a curve based on constraints from a given image I , to segment an object in I . Ideally the curve starts somewhere near the object to be detected, and iteratively moves in the direction normal to the curve itself until reaches the boundary of the object. The goal of the final segmentation is to preserve important features of the object, such as edges and corners, while smoothing irrelevant details. The segmentation is achieved by performing a local optimisation and as for the MRF (Section 3.2.5) spatial consistency is included. The classical snake model was proposed by Kass et al. [67] and depends on the gradient of the image I , so it is able to stop the curve evolution at the boundary of the segmented object:

$$F(C) = \alpha \int |C'(s)|^2 ds + \beta \int |C''(s)| ds - \lambda \int |\nabla I(C(s))|^2 ds. \quad (3.21)$$

Here, C is the parametrised curve and α , β and λ are positive parameters. The first two terms control the smoothness of the curve, while the last term attracts the contour towards the object that presents high intensity gradients. Equation (3.21) is non-convex and therefore one cannot compute a global solution. Straightforward implementations of this energy model are not capable of handling changes in the topology of the evolving contour. Therefore the topology of the final curve will be as that for the initial curve C_0 (curve C at time zero) unless special procedures, many times heuristic, are implemented for detecting possible splitting and merging. This is problematic if an unknown number of objects must be simultaneously detected.

In implicit curve representation, contours are represented as the zero level curve of some embedding function $\phi : \Omega \rightarrow \mathbb{R}$:

$$C = \{x \in \Omega | \phi(x) = 0\}. \quad (3.22)$$

There are different methods to evolve implicit contours. The most popular ones are known as level-set methods, initially proposed in [99], where C is propagated by evolving a time dependent function $\phi(x, t)$ according to a Partial Differential Equation (PDE).

The first formulation considers a velocity field $\vec{V}(x)$ which represents the evolution at each point x of the implicit surface $\phi(x)$. If one wants to move all the points on the surface with respect to \vec{V} , then the ordinary differential equation

$$\frac{dx}{dt} = \vec{V}(x) \quad (3.23)$$

must be solved for all the points x which are on the front at time t , *i.e.* $\phi(x) = 0$. Practically, this means that the front must be discretised into a finite number of pieces, which can be a challenging task especially if the connectivity of the surface changes and if the surface elements have become too distorted. In order to avoid the aforementioned problems the implicit function ϕ can be used to represent both the interface and its own evolution. This can be represented by the convection, or level set, equation:

$$\frac{\partial \phi}{\partial t} + \vec{V} \cdot \nabla \phi = 0. \quad (3.24)$$

This PDE represents the motion of the interface where $\phi(x) = 0$. At every new time point, the curve evolves and the new location of the surface needs to be calculated. A well know discretisation method to compute this evolution is the upwind differencing described in [98].

The first application of the level-set formulation to image segmentation was proposed in [27], and a first application to synthetic medical images was proposed in [91]. The

curve evolution or level-set formulation for the snake energy in Equation (3.21) was proposed in [28, 69]:

$$\begin{aligned}\frac{\partial\phi}{\partial t} &= |\nabla\phi|\operatorname{div}\left(g(I)\frac{\nabla\phi}{|\nabla\phi|}\right) \\ &= g(I)|\nabla\phi|\operatorname{div}\left(\frac{\nabla\phi}{|\nabla\phi|}\right) + \nabla g(I) \cdot \nabla\phi\end{aligned}\tag{3.25}$$

where div is the divergence operator, and the modulus of the gradient $|\nabla I|$, in Equation (3.21) is replaced by a more general edge based function $g(I)$. Ideally this function should stop the evolution of the curve whenever an edge is encountered. This level-set formulation allows topological changes and geometric flexibility, and has been successfully applied to problems in two and three spatial dimensions.

An evolution of the level-set methods was formulated by Chen and Vese [31]. These authors propose an alternative active contours method whose stopping criterion does not involve edge functions. The stopping term of this approach is based on the Mumford-Shah segmentation technique [95]. The model developed by these authors can detect objects boundaries having a variety of gradient values, for instance, objects that are very smooth, or even have discontinuous boundaries. In addition, the model and its level-set formulation are such that interior contours are automatically detected, and the initial curve can be selected to be anywhere in the image. The basic idea of the model considers an evolving curve C and the two regions that this curve identifies: an interior region with an average intensity of c_0 , and an exterior one with an average intensity of c_1 . The model is formulated as follows:

$$\begin{aligned}F(c_1, c_2, C) &= \mu \cdot \operatorname{Length}(C) + \nu \cdot \operatorname{Area}(\operatorname{inside}(C)) \\ &\quad + \lambda_1 \int_{\operatorname{inside}(C)} |I - c_0|^2 + \lambda_2 \int_{\operatorname{outside}(C)} |I - c_1|^2\end{aligned}\tag{3.26}$$

where $\mu, \nu \geq 0$, $\lambda_1, \lambda_2 > 0$ are scalars. The first and the second terms are the regularising terms. From this formulation it is possible to observe that the scalars μ and ν are controlling the smoothness of the curve, particularly useful in the presence of noise. Using the variational level-set formulation presented in [161], the unknown variable C can be replaced by the unknown variable ϕ . Defining the Heaviside step function H , and the one dimensional Dirac measure δ_0 as:

$$H(z) = \begin{cases} 1, & \text{if } z \geq 0 \\ 0, & \text{if } z < 0 \end{cases} \quad \delta_0(z) = \frac{d}{dz}H(z)\tag{3.27}$$

the regularisation terms in Equation (3.26) can be formulated as:

$$\begin{aligned} \text{Length}\{\phi = 0\} &= \int_{\Omega} |\nabla H(\phi)| = \int_{\Omega} \delta_0(\phi) |\nabla \phi|, \\ \text{Area}\{\phi \geq 0\} &= \int_{\Omega} H(\phi) \end{aligned} \quad (3.28)$$

and the third and fourth terms as:

$$\begin{aligned} \int_{\phi > 0} |I - c_0|^2 &= \int_{\Omega} |I - c_0|^2 H(\phi), \\ \int_{\phi > 0} |I - c_1|^2 &= \int_{\Omega} |I - c_0|^2 (1 - H(\phi)). \end{aligned} \quad (3.29)$$

Using the above equalities, the energy $F(c_1, c_2, \phi)$ can be written as:

$$\begin{aligned} F(c_1, c_2, \phi) &= \mu \int_{\Omega} \delta(\phi) |\nabla \phi| + \nu \int_{\Omega} H(\phi) \\ &\quad + \lambda_1 \int_{\Omega} |I - c_1|^2 H(\phi) + \lambda_2 \int_{\Omega} |I - c_2|^2 (1 - H(\phi)). \end{aligned} \quad (3.30)$$

Keeping c_1 and c_2 fixed and minimising F with respect to ϕ , the Euler-Lagrange equation for ϕ is:

$$\frac{\partial \phi}{\partial t} = \delta_{\epsilon}(\phi) \left[\mu \operatorname{div} \left(\frac{\nabla \phi}{|\phi|} \right) - \nu - \lambda_1 (I - c_1)^2 + \lambda_2 (I - c_2)^2 \right] \quad (3.31)$$

where δ_{ϵ} is a globally positive approximation of the δ function.

In general, there is a large number of level-set methods which can be classified into two major classes: edge-based models [28, 82], and region-based models [31, 102, 112]. Edge-based methods use local edge information to attract the contour towards the object boundaries. Region based methods identify a region descriptor, or features, of the object to be segmented, in order to guide the evolution of the curve. Most of these methods are following the same generic steps. At the initial step, the contour ϕ needs to be placed either outside or inside the region of interest, and is then allowed to evolve according to a partial differential equation. The positioning of the initial contour is still a challenging task for many level-set methods.

The method we proposed in [126] is an initial attempt to perform a segmentation of the grey matter tissue in T_1 -w images acquired at 7 Tesla. Our algorithm is based on an atlas probability map segmentation which corrects the high inhomogeneity present in 7T images and provides an initial segmentation of the main tissues: Grey Matter (GM), White Matter (WM), and cerebrospinal Fluid (CSF). The atlas based method is similar to the one proposed by [5] described in Section 3.2.6. Therefore registration to a T_1 -weighted (T_1 -w) atlas is interleaved with the segmentation into the three main tissues and bias estimation. Since the input images have a very high resolution these steps are performed using a multi resolution approach. In the lowest resolution image the




Figure 3.7: 2D slice of a T_1 -w image acquired at 7T in (a), and the bias corrected image in (b). Tissue segmentation obtained with an atlas based approach into GM (c), WM (d), and CSF (e).

bias field parameters are estimated (Section 3.2.6), and subsequently applied, without re-estimation, to the higher resolution images. When convergence is reached, we have T_1 -w image corrected for bias field and a segmentation into grey matter, white matter and cerebrospinal fluid, Figure 3.7. Even if bias correction improves the image contrast, high frequency noise is still present in the image. Therefore the atlas-based grey matter segmentation may contain several regions which are not correctly detected.

Therefore a second step is introduced where the segmentation of the grey matter is refined using a level-set approach. In addition to the original input image, a standardised image S is introduced, in which every tissue type is assigned a constant pixel intensity, Figure 3.8(a). The computation of S is based on the mean values of the tissues types computed with the atlas based approach previously described:

$$S_x = \text{GM}_x \cdot \mu_{\text{GM}} + \text{WM}_x \cdot \mu_{\text{WM}} + \text{CSF}_x \cdot \mu_{\text{CSF}}$$

where x denotes a voxel in the image. The seed points, which define the initial contour, are randomly sampled on the atlas based GM segmentation output as those voxels which have a high probability value of being GM. Given the standard level-set Equation (3.24) the speed function \vec{V} is modelled as being dependent on both the original and standardised images:

$$\vec{V}(x) = \beta_1 [(1 - \alpha)D_1(x) + \alpha\kappa(x)] + \beta_2 [(1 - \alpha)D_2(x) + \alpha\kappa(x)]. \quad (3.32)$$

The variable κ represents the mean curvature of the contour and it is defined as $\kappa = \nabla \cdot \frac{\nabla\phi}{|\nabla\phi|}$, α it is a constant value, β_1 and β_2 are constants weighting the two terms of the speed function, while the terms D_1 and D_2 encode the intensity information from the original and standardised image, respectively:

$$D_1(x) = 2 \cdot \frac{|\mu_{\text{GM}} - I(x)|^n}{\epsilon + |\mu_{\text{GM}} - I(x)|^n} - 1 \quad D_2(x) = 2 \cdot \frac{|\mu_{\text{GM}} - S(x)|^n}{\epsilon + |\mu_{\text{GM}} - S(x)|^n} - 1 \quad (3.33)$$

where ϵ and n are constant values, and empirically fixed at $\epsilon = 0.01$ and $n = 1.5$. The proposed method has been tested on several 7T T_1 -w images, an example is shown in Figure 3.8(b)-(c). For all the images an improved GM segmentation has been obtained




Figure 3.8: 2D slice of the standardised image S in (a), final GM segmentation obtained with the proposed method in (b), and a 3D visualisation of the GM in (c).

compared to the results achieved only with an atlas based approach. The results demonstrated that level-set approaches provide good results but they need to be combined with other methods to overcome the inherent difficulties of the initialisation phase.

Similar approaches have been proposed where the initial contour has been identified with thresholding approaches [81, 83].

The presented class of methods can be used independently or in combination to perform the segmentation of object(s)/organ(s) of interest. The choice of the appropriate algorithm depends on the application and on the number and type of images used. As already mentioned, a feature-based approach to detect focal cortical dysplasia is proposed in Chapter 5, while a graph-based method is proposed to segment brain white matter lesions (Chapter 6).

3.3 Relation between Denoising and Segmentation

Image denoising and segmentation are two fundamental, and closely related problems in image processing and computer vision. The goal of denoising is to remove noise and/or spurious details from a given, possibly corrupted, digital image while maintaining essential features such as edges. The commonly used models for noise are, for example, Gaussian random noise or salt and pepper. As described earlier, the goal of segmentation is to divide a given image into regions that belong to distinct objects. The quality of the segmentation is inversely related to the level of noise present in the image. Therefore, any segmentation algorithm requires an initial step which either removes as much noise as possible or incorporates it into the model formulation. Denoising and segmentation can be formulated using variational principles expressed in terms of PDEs.

The description starts with the presentation of the total variation denoising paradigm, which is related to the contributions of this thesis presented in Chapter 6. The active contour model presented in Section 3.2.7 and used in the method we proposed in [126], can also be seen as a denoising algorithm. The curve, that models the object to be segmented, has to preserve important features, such as corners, while removing irrelevant

details that may be caused by noise. Then, for completeness, a brief description of two additional methods is provided.

3.3.1 Total Variation

The Total Variation (TV) is defined as the integral of absolute gradient of a signal. The signal in images can be seen as the pixel/voxel intensity. TV denoising is based on the principle that images affected by noise have a high total variation. Therefore, reducing the TV of the image subject to it being as close as possible to the original image, removes the unwanted noise.

One of the best known methods is the total variation based model published by Rudin, Osher, and Fatemi (ROF) [115]. Let $u(x) : \Omega \rightarrow \mathbb{R}$, where $x \in \Omega$ and $\Omega \in \mathbb{R}^3$ be the image without noise. In the ROF model, the corrupted image $I(x) : \Omega \rightarrow \mathbb{R}$ can be described as:

$$I(x) = u(x) + n(x) \quad (3.34)$$

where n is a Gaussian white noise of variance σ^2 . The ROF model minimises the following energy:

$$\min_u \int_{\Omega} |\nabla u| d\Omega \quad \text{s.t.} \quad \int_{\Omega} (u - I)^2 d\Omega = \sigma^2. \quad (3.35)$$

Chambolle showed that the non-convex ROF model, Equation (3.35) can be turned into a convex problem by replacing the equality constraint into an inequality constraint [30]. Therefore the unconstrained model becomes:

$$\min_u \left[\int_{\Omega} |\nabla u| d\Omega + \lambda \int_{\Omega} (u - I)^2 d\Omega \right] \quad (3.36)$$

where λ is a scalar parameter that controls the tradeoff between smoothness and fidelity to the original data. If $\lambda = \infty$ the minimisation will result in the original image, while if $\lambda = 0$ the result is an image where the intensity values are the mean value of the original image. The Euler-Lagrange equation for the unconstrained ROF model is:

$$-\nabla \cdot \left(\frac{\nabla u}{|\nabla u|} \right) + 2\lambda(u - I) + \frac{\partial u}{\partial N} = 0 \quad (3.37)$$

where $\partial u / \partial N$ is the derivative of u normal to the boundary. This equation is not differentiable at points where $|\nabla u| = 0$, therefore different techniques have been developed to minimise it. A simple approach to overcome this problem is to regularise the gradient as $|\nabla u|_{\epsilon} = \sqrt{|\nabla u|^2 + \epsilon}$. The parameter ϵ influences the performance of the algorithm and must therefore be chosen with care: a too low value of ϵ will slow down the algorithm in regions with low intensity gradients, whereas high values will cause edge blurring. Chambolle proposed to use the dual formulation for the TV norm [29]:

$$|\nabla u| = \max_{\mathbf{p}} \mathbf{p} \cdot \nabla u \quad \text{with} \quad |\mathbf{p}| \leq 1 \quad (3.38)$$

where \mathbf{p} is a vector field. Substituting this equation into Equation (3.36) gives the primal-dual formulation:

$$J = \min_u \max_{\mathbf{p}} \int_{\Omega} \mathbf{p} \cdot \nabla u \, d\Omega + \lambda \int_{\Omega} (u - I)^2 \, d\Omega. \quad (3.39)$$

The Euler-Lagrange formulation for the primal-dual results in:

$$\begin{aligned} \frac{\partial J}{\partial u} &= 2\lambda(u - I) - \operatorname{div} \mathbf{p} \\ \frac{\partial J}{\partial \mathbf{p}} &= \nabla u \end{aligned} \quad (3.40)$$

where div is the divergence operator.

An efficient algorithm to minimise the primal-dual is to use a gradient descent proposed in [29]. The algorithm steps are briefly described in Algorithm 2, where the variable t represents the iterations.

Algorithm 2

- 1: Initialise $u_t = I$, and p_t with random values *s.t.* $\|p_t\| \leq 1$.
 - 2: **repeat**
 - 3: $u_{t+1} = u_t - \tau[2(u - I) - \lambda \operatorname{div} p]$
 - 4: $p_{t+1} = p_t + \sigma(\lambda \nabla u)$
 - 5: $p_{t+1} = \frac{p_{t+1}}{\max\{1, \|p_{t+1}\|\}}$
 - 6: $u_t = u_{t+1}$, $p_t = p_{t+1}$
 - 7: **until** Convergence of the vector field p and of u .
-

The variables λ , τ , and σ in Algorithm 2, are scalar values which need to follow the inequality $\sigma \cdot \tau \leq \frac{1}{8\lambda^2}$.

There are additional methods proposed to solve the unconstrained ROF model, as for example the method of Total Generalised Variation proposed in [18] or the first-order scheme for convex optimisation proposed in [150]. The total variation model has been applied to various vision problems, including segmentation and registration. The method proposed in [141] performs image segmentation by minimising an energy which considers a weighted total variation, as described in [20], and additionally incorporated user constraints.

In Chapter 6 we will see how the denoising algorithm proposed by Chambolle [29] can help to detect and segment lesions in the brain.

3.3.2 Other Methods

Anisotropic Diffusion

In an image, the presence of noise can cause spurious local variations in pixel intensities, *i.e.*, locally unphysical pixel intensity gradients. Here, the theory of diffusion processes

provides a sensible approach to image denoising. Physically, diffusion consists in the motion of particles (atoms or molecules) from a region of higher concentration to a region of lower concentration. Under certain experimental conditions, this is expressed quantitatively by Ficks law as:

$$J = -D\nabla u. \quad (3.41)$$

Equation (3.41) states that a concentration gradient, ∇u , gives rise to a particle flux, J , that tends to equalise the concentration. Mathematically, these quantities are related through the diffusion tensor D . If the flux is parallel to the concentration gradient, the diffusion is said to be isotropic, otherwise anisotropic. In diffusion processes, mass is neither created nor destroyed. The conservation of matter is expressed by the continuity equation:

$$\frac{\partial u}{\partial t} = -\operatorname{div} J \quad (3.42)$$

where t represents the time and div is the divergence operator. Introducing Fick's law, Equation (3.41), into Equation (3.42) the diffusion equation is obtained:

$$\frac{\partial u}{\partial t} = \partial_t u = \operatorname{div}(D\nabla u). \quad (3.43)$$

In image processing the diffusion equation can be applied to an image I , by considering pixel intensities, or grey values, to represent the local concentration gradient. Thus, the local gradient of image grey values will tend to be equalised, resulting in a smoothed image. The diffusion equation applied to an image I becomes:

$$I_t = \operatorname{div}(c(x, y, t)\nabla I) \quad (3.44)$$

where I_t describes the evolution of the smooth image over time t , and c is the diffusion coefficient. A constant value for c leads to a linear diffusion equation with homogeneous diffusivity. With this formulation the filter will smooth the noise in the image I but also important structures as for example the edges. To avoid this undesirable effect an inhomogeneous diffusivity can be chosen as $c = g(|\nabla I|)$, as in the anisotropic diffusion method of Perona and Malik [106]. The underlying idea is to smooth out the image while preserving the image edges:

$$I_t = \operatorname{div}(g(|\nabla I|)\nabla I) \quad (3.45)$$

where g is a non-negative, monotonically decreasing function with $g(0) = 1$. This definition of the function g guarantees that in regions with large variations in pixel intensity, the diffusion will be comparatively low, hence the location of the edges will be preserved. On the other hand, in regions with small variations in pixel intensity, the image will be further smoothed out due to the relatively high local diffusion coefficient. The definition that Perona and Malik gave for the class of functions g is ill-posed [44, 97]. As discussed by Catté et al. [44] there are some admissible functions, as for example $g(s) = e^{-s}$, for which the solution of Equation (3.45) is unstable. In fact, in order to obtain both

existence and uniqueness of the solution the function g must satisfy that $s \cdot g(s)$ is non decreasing [60]. Catté et al. [44] proposed a new model which overcomes the problem mentioned above and which is robust in the presence of noise:

$$I_t = \operatorname{div}(g(|\nabla G_\sigma * I|)\nabla I) \quad (3.46)$$

where G_σ is the Gaussian function. Thus the modification consists in replacing the gradient $|\nabla I|$ by its estimate $|\nabla G_\sigma * I|$. The main drawback of this regularisation is that it introduces a blurring effect which depends on the choice of the parameter σ . The optimal choice of σ is noise dependent, thus requires *a priori* knowledge of the noise level.

Weickert [148] proposed to regularise Equation (3.45) using a truly anisotropic diffusion:

$$I_t = \operatorname{div}(D(|\nabla G_\sigma * I|)\nabla I). \quad (3.47)$$

The parameter σ again needs to be estimated but the blurring effect is reduced by the diffusion tensor D .

The remarkable results of anisotropic diffusion techniques are mainly associated with the introduction of anisotropic smoothing. This technique is able to simultaneously eliminate noise and preserve, or even enhance edges. This is in contrast to conventional spatial filtering techniques that do not preserve region boundaries or small structures. For this reason this is still an active field of research and many different methods have been proposed to regularise and/or relax Perona and Malik's diffusion method [54, 139].

Anisotropic diffusion methods have been used in different applications such as noise reduction and edge detection [7, 155, 156]. These methods have also been applied to enhance structures in medical images such as X-ray angiograms, microscopy images of kidney [66] as well as ultrasound images [156].

Non-Local Means

The main idea of Non-Local (NL)-means filtering is that the restoration of a noisy texture patch (Section 3.2.2) is improved by using similar texture patches in other parts of the image.

This method was developed by Buades et al. and proposes a new way to denoise images [22]. Instead of looking at every pixel independently it considers also their neighbourhood. Given the noisy image $I : \Omega \rightarrow \mathbb{R}$, the NL-means for pixel i is defined as:

$$NL[I](i) = \sum_j w(i, j)I(j) \quad (3.48)$$

where the weight w depends on the similarity between pixel i and j . This similarity depends on the intensity values of $I(N_i)$ and $I(N_j)$, where N is the neighbourhood, or ‘patch’, of a pixel, and is defined as:

$$w(i, j) = \frac{1}{Z(i)} \exp\left(-\frac{\|I(N_i) - I(N_j)\|_2^2}{h^2}\right) \quad (3.49)$$

where h is a parameter which controls the degree of filtering and $Z(i)$ is a normalisation constant defined as:

$$Z(i) = \sum_j \exp\left(-\frac{\|I(N_i) - I(N_j)\|_2^2}{h^2}\right). \quad (3.50)$$

Therefore, image denoising is achieved by averaging out the intensities of the most similar pixels in the image. Since the similarity measure considers the neighbourhood of each pixel, this method is able to preserve texture.

The NL-means method has been widely used for texture denoising [21], video processing [90] and surface smoothing [154]. More recently, it also has been applied to brain image segmentation [33]. In this paper the authors use a set of manually segmented images as a prior for the segmentation. When a new image needs to be processed they compare the local patches of the image to the patches of the segmented images.

Chapter 4

Introduction to Graph-cuts

In this chapter the relevant background for our contribution, that will be proposed in Chapter 6, is explained. The description starts with the graph-based segmentation algorithm, *i.e.* the max-flow/min-cut, proposed by Boykov and Jolly [17]. The algorithm is usually referred as graph-cuts and is able to find the global optimum for binary segmentation problems. Then, the extension to multi labelling problems is provided as well as to local optimisation techniques. Finally, further developments, such as object interactions, are described.

4.1 General Formulation

Graph-cuts is one of the most prominent algorithm which includes MRF (Section 3.2.5), and here the same graph notation is used. Given the unknown labelling $f : \Omega \rightarrow \mathcal{L}$, where \mathcal{L} is the labelling space, and an image I , the energy function that Boykov and Jolly [17] propose to minimise is:

$$E(f) = \sum_{i \in I} D_i(f_i) + \lambda \sum_{(i,j) \in N} V_{i,j}(I_i, I_j) \quad (4.1)$$

where N represents the neighbourhood of each vertex $i \in \Omega$, and λ is a weighting parameter. The first term represents the cost of assigning the label f_i to the pixel i , while the second term represents the penalty of neighbour interactions and can therefore be seen as a smoothing term. The parameter λ weights the relative importance between the data and the regularisation constraint: if λ is small, then the data term is more important while if λ is large, then the regularising term is more relevant. In the extreme case of $\lambda = 0$, the minimisation relies only on the intensity information of the image, hence the resulting segmentation is identical to a thresholding approach. In the other extreme case of $\lambda = \infty$, the optimal solution consists in all the pixels being assigned to the same label. Therefore, a correct value for λ needs to be chosen and this choice depends on the specific application.

Note the formal similarity between Equation (4.1) and Equation (3.3): any energy function of the type given by Equation (4.1) can be derived as the maximum a posteriori estimate, and hence is justified by Bayesian statistics. The data term $D(\cdot)$ or the E_{App} energy both represent the intrinsic information contained in the image. As an example, this energy can encode pixel-wise intensity information, which can be modelled using

a Gaussian or GMM distribution. On the other hand, the knowledge term $V(\cdot)$ or the energy $E_{\text{Knowledge}}$ both encode spatial information. In [17], Boykov and Jolly propose to use the following function:

$$V_{i,j}(I_i, I_j) = \exp\left(-\frac{(I_i - I_j)^2}{2\sigma^2}\right) \cdot \frac{1}{\text{dist}(i, j)} \quad (4.2)$$

where I_i represents the intensity of the pixel i , and dist measures the distance between i and j . This function highly penalises discontinuities between pixels of similar intensities, *i.e.* $|I_i - I_j| < \sigma$, while it has a small penalty value if $|I_i - I_j| > \sigma$. Note that if $f_i = f_j$ there is no additional cost to the cut, while if $f_i \neq f_j$ then the $V_{i,j}(I_i, I_j)$ will contribute to the cost.

4.2 Graph Construction

The minimisation of Equation (4.1) needs that all the voxels in the image are inserted in the graph as vertices, so there is a one-to-one correspondence between the vertices and the pixels. Moreover, two additional vertices, namely the *source* (s) and the *sink* (t) are required. These two vertices must be in the graph and they represent the background and the foreground in a binary segmentation framework. Therefore a cut separates the *source* from the *sink*: the vertices which, after the cut, are connected to the *source* represent the background of the image, while those connected to the *sink* are the segmented object.

In order to model the two energy terms in Equation (4.1), the graph contains two types directed edges: the t -link and the n -link. The t -links are the terminal edges: prior to the cut, they connect each vertex that takes parts in the segmentation to both the *source* and the *sink*; after the cut, a segmentation is obtained where each vertex is connected via a t -link either to the *source* (if the vertex is part of the background) or to the *sink* (if it is part of the foreground). The n -links connect neighbouring vertices/pixels and hence encode spatial information. A graph representation of a 3×3 image is shown in Figure 4.1.

Now it is shown how to construct this directed graph, with the assumption that the *source* represents the intensity level 0 and the *sink* represents the intensity level 1. In Equation (4.1) the data term $D_i(f_i)$ represents the cost of assigning the label $f_i \in \{0, 1\}$ to the pixel i . If the pixel i is assigned to the *source* then $D_i(0) < D_i(1)$. To encode this information into the graph two t -links for each pixel node are created: the first t -link connects pixel i to the *source*, as denoted by (s, i) , and has the weight $D_i(1)$. Similarly the second t -link connects i to the *sink*, denoted by (i, t) , and has the weight $D_i(0)$. The weights in the graph are non-negative, *i.e.* $D_i \geq 0$.

Next, the regularisation term is introduced, where it is assumed that the original image is piece-wise constant. This assumption can be encoded with the use of n -links, which connect neighbouring pixels. The size of the neighbourhood depends on the specific application, as an example one can use 4-connected or 8-connected neighbourhood. The weights of these n -links depend on the function V (see Equation (4.1)) and on λ .




Figure 4.1: Graphs model for a 3×3 image in (a). The orange and green vertices are, respectively, the *source* and the *sink*. Orange and green edges connecting the vertices to *s* or *t* are the *t-links*, while the black edges connecting the vertices are the *n-links*. In (b) a cut of the graph (dashed line) separates between nodes connected to the *source* and to the *sink* vertices.

To see how to perform a cut on such a graph, the reader is referred to [14]. The cost of any cut C of the graph depends on the weight values of the *n-links* and *t-links* which are cut. First let consider the contribution of the *t-links* to the cut. For every node in the graph it must be decided where the edge with lowest weight lies. If the lowest weight is that of the edge (i, t) , then $D_i(0)$ will contribute to the cut cost and the pixel i will be assigned to the *source* (label 0). If the cut is on the edge (s, i) , then $D_i(1)$ will contribute to the cost and in this case the pixel i will be assigned to the *sink* (label 1). For the contribution of the *n-links* let assume that the pixels i and j are neighbours and are connected by the edges (i, j) and (j, i) . If i and j are on the same side of the cut, then the respective weights of (i, j) and (j, i) do not contribute to the cost of the cut. However if i and j are on different side of the cut then these *n-links* are cut and contribute to its cost.

4.3 Submodularity Condition

In general the graph-cuts method by its construction, computes binary segmentations. An important work, published by Kolmogorov and Zabih [73], formalises the class of function that can be minimised by graph cuts. The author showed that the function E in Equation (4.1) is graph representable if it is a regular/submodular function. This condition states that the cost of assigning two neighbouring pixels to the same label has to be less than or equal to the cost of assigning different labels to them:

$$V_{i,j}(0, 0) + V_{i,j}(1, 1) \leq V_{i,j}(0, 1) + V_{i,j}(1, 0). \quad (4.3)$$




Figure 4.2: Graph representation for the regularity term.

Submodularity is ensured when all edge weights are non-negative. The $V_{i,j}$, second term of Equation (4.1), represents the potential cost of the two neighbouring vertices i, j in four different situations: $V_{i,j}(0,0)$ represents the cost that both vertices, or nodes, are assigned to the *source*, $V_{i,j}(1,1)$ is the cost when they are both assigned to the *sink* and $V_{i,j}(1,0), V_{i,j}(0,1)$ are the costs when the cut separates the two nodes. The possible interactions between the nodes i and j can be schematically represented with the following tables

$$V_{i,j} = \begin{array}{|c|c|} \hline V_{i,j}(0,0) & V_{i,j}(0,1) \\ \hline V_{i,j}(1,0) & V_{i,j}(1,1) \\ \hline \end{array} = \begin{array}{|c|c|} \hline A & B \\ \hline C & D \\ \hline \end{array}$$

$$\begin{array}{|c|c|} \hline A & B \\ \hline C & D \\ \hline \end{array} = A + \begin{array}{|c|c|} \hline 0 & 0 \\ \hline C - A & C - A \\ \hline \end{array} + \begin{array}{|c|c|} \hline 0 & D - C \\ \hline 0 & D - C \\ \hline \end{array} + \begin{array}{|c|c|} \hline 0 & B + C - A - D \\ \hline 0 & 0 \\ \hline \end{array}$$

Consider the second expression. The first term A is a constant value which does not need to be inserted in the graph; the second and third terms depend only on the vertices i and j , respectively; and the last term is the regularity term (see Equation (4.3)). The graph construction for the overall energy contribution due to pixels i and j is illustrated in Figure 4.2, where the weights $C - A > 0$ and $D - C > 0$ are placed on the t -links while the regularity term weights the n -link between i and j .

If the graph construction follows the regularity condition, the solution to the minimisation problem is guaranteed to be optimal. The proof can be seen in [73].




Figure 4.3: Ishikawa construction for 3 labels. In (a) the t -links connecting every node in the layers to s and t . From (b) to (d) the three possible configurations, where the cut is represented by the dashed line. (b) represents $\mathcal{L} = 1$, or rather $i_0 = i_1 = 0$, (c) represents $\mathcal{L} = 3$ where $i_0 = i_1 = 1$, and in (d) the additional ∞ -edge is added between the layers to force $i_0 = 0, i_1 = 1$ which represents $\mathcal{L} = 2$.

4.4 Multi-labelling Segmentation

Graph-cuts are devised for binary segmentation problems and have found many applications, but for multi-labelling problems the technique must be extended. Given the labelling $f : \Omega \rightarrow \mathcal{L}$, where \mathcal{L} is the labelling space which contains more than two labels, Equation 4.1 can be extended in:

$$E(f) = \sum_{i \in I} D_i(f_i) + \lambda_{\ell_1, \ell_2} \sum_{(i, j) \in N} V_{i, j}(I_i, I_j) \quad (4.4)$$

where the parameter λ is now dependent on two labels $\ell_1, \ell_2 \in \mathcal{L}$. A multi-labelling problem is in general NP-hard, but becomes polynomial if a total ordering of \mathcal{L} , *i.e.* $\leq_{\mathcal{L}}$, is given such that λ_{ℓ_1, ℓ_2} , is convex with respect to this ordering. Let us start to observe that λ_{ℓ_1, ℓ_2} , in Equation (4.4), is a scalar value, but can also be interpreted as a function: $\lambda : \mathcal{L}^2 \rightarrow \mathbb{R}$ [64]. Ishikawa studied functions that can be written as $\lambda_{\ell_1, \ell_2} = F(|\ell_1 - \ell_2|)$ for a convex function F [64]. Then, the polynomial run time was achieved by Ishikawa, by first allowing an independent binary decision for each label ℓ which requires that the label space \mathcal{L} has to be expanded to the set of its subsets, *i.e.* $2^{\mathcal{L}}$. Thus, the multi-labelling $f : \Omega \rightarrow \mathcal{L}$ becomes the extended multi-labelling $\hat{f} : \Omega \rightarrow 2^{\mathcal{L}}$. Secondly, he introduced the inclusion constraints

$$\ell \in \hat{f}(x) \Rightarrow \ell' \in \hat{f}(x) \quad \text{for all } \ell' \leq_{\mathcal{L}} \ell,$$




Figure 4.4: Two moves introduced in [16]. In (a) the initial configuration, in (b) an α/β -swap move and in (c) an α -expansion.

which ensure that $\hat{f}(x)$ has the form $\{\ell' \in \mathcal{L} | \ell' \leq_{\mathcal{L}} \ell\}$ for some $\ell \in \mathcal{L}$. He showed that the data terms can be chosen in a way that an optimum of the original multi-labelling $f: \Omega \rightarrow \mathcal{L}$ is derived from the optimal $\hat{f}: \Omega \rightarrow 2^{\mathcal{L}}$ via $f(x) := \max \hat{f}(x)$.

Given the label space $\mathcal{L} = \{1, \dots, k\}$, the graph is to be constructed by creating $k - 1$ copies i_0, i_1, \dots, i_{k-2} , of each vertex/pixel i . Every copy of a given vertex is itself considered as a vertex and can be seen as being part of a separate ‘layer’ in the graph, where it encodes a specific label. All copies are then connected to the *source* and to the *sink*, $(s, i_0), (s, i_1), \dots, (s, i_{k-2})$ and $(i_0, t), (i_1, t), \dots, (i_{k-2}, t)$, as was previously done for the binary case, Figure 4.3(a). These n -links represent the relationship between neighbouring pixels, which can be within a ‘layer’ or across ‘layers’. Intra-layer n -links have the same meaning as for the binary graph so they represent the smoothing term for each label. Through-layer n -links represent the ordering of the labels. In addition they can be used to encode, for example, topological constraints or minimum distance constraints between two or multiple labels. This kind of graph is usually referred as ‘Ishikawa construction’.

It is worth seeing how to encode different labels. Let us suppose that $k = 3$ labels, without loss of generality denoted as $\mathcal{L} = \{1, 2, 3\}$. According to the Ishikawa construction, two graph layers, hence also two copies i_0, i_1 , of each vertex/pixel, are required. In the graph, these nodes (copies) are connected to both s and t , as shown in Figure 4.3(a). Thus, there are four different possible configurations: if both i_0 and i_1 remain connected to the *source*, then they are both assigned the label 0 (configuration 0 0, Figure 4.3(b)); if i_0 is attached to the *source* while i_1 is attached to the *sink*, their respective labels are 0 1 (configuration 0 1); interchanging connections give rise to a configuration 1 0; finally in the configuration 1 1 both nodes remain connected to the *sink*, Figure 4.3(c). It is seen that both configurations, 0 0 and 1 1, are ordered; they represent label 1 and 3, respectively. On the other hand, for the configurations 0 1 and 1 0 there is no obvious ordering between the two. Therefore, Ishikawa introduced an ∞ -edge between the different layers as shown in Figure 4.3(d). This additional edge enforces that if the vertex in layer 0 is assigned to the *sink*, then so is also for the vertex in layer 1 and both vertices are assigned label 1. With this additional constraint there are only three possible configurations, namely 0 0, 0 1, 1 1.

4.5 Local Optimisation

The pairwise term $V_{i,j}$ in Equation (4.1), is not suitable for many computer vision applications if the number of labels in \mathcal{L} is large, as it does not preserve discontinuities. The penalty that a convex function $V_{i,j}$ imposes on discontinuities is large, and in an optimal labelling f these discontinuities are smoothed out with a ‘ramp’. Even if the cost of few steps are equal to the cost of a single large one, using the absolute difference measure, in practice over-smoothed disparities are seen. Below, it is shown how to minimise discontinuity preserving functions using graph-cuts.

A possible way to construct a discontinuity-preserving function $V_{i,j}$ is to bound its maximum value. An example is given by the Potts model, defined by $V_{i,j} = K \cdot T(f_i \neq f_j)$, where f_i, f_j represents the labels of the nodes i and j and the function T returns 1 if the labels are different, 0 otherwise [110]. Unfortunately, performing energy minimisation with a Potts model is NP-hard [16].

Boykov et al. [16] proposed two different procedures, or ‘moves’, that can be independently used to minimise, locally, this kind of energy functions: the ‘swap move’ and the ‘expansion move’ which will be described in rest of the section. There are also other approximation algorithms that have been developed in order to minimise non-submodular functions, for example the quadratic pseudo-boolean optimisation (QPBO) [72, 113], QPBOP, where the last P stands for ‘probing’ [13], and the QPBOI where the I stands for improving [113].

Both these moves are generally applied to multi labelling problems, and they allow many pixels to change their label simultaneously. The swap involves exactly two labels, α and β , say simultaneously. During an α/β -swap move, pixels with the label α can get the new label β , and vice versa, while all the other pixels retain their current label, see Figure 4.4(b). This type of move can be applied when the so-called inequality $V_{i,j}(\alpha, \alpha) + V_{i,j}(\beta, \beta) \leq V_{i,j}(\alpha, \beta) + V_{i,j}(\beta, \alpha)$ holds.

The expansion move considers one label at the time, as for example α . The α -expansion move allows any image pixel to change its label to α , Figure 4.4(c). The expansion move can be used when $V_{i,j}(\alpha, \alpha) + V_{i,j}(\beta, \gamma) \leq V_{i,j}(\alpha, \gamma) + V_{i,j}(\beta, \alpha)$ holds for all $\alpha, \beta, \gamma \in \mathcal{L}$. In this case the function $V_{i,j}$ is said to satisfy the expansion inequality.

Both of these algorithms are based on an initial segmentation which provides a labelling of the image, e.g. the Ishikawa construction. In addition both of these moves find a local minimum with respect to their move, and the two algorithms are quite similar in their structure. First it is important to define the meaning of ‘local minimum’. For each labelling f , the set of moves is defined as M_f . Then f is a local minimum if there exists a labelling f' such that $E(f) \leq E(f')$. The structure of the algorithms can be seen in Algorithm 3. The sets of moves M_f involves, in general, both α/β -swap or α -expansion move. The Algorithm 3 is generally repeated for all α/β pairs in the label space or for all individual α labels. Because the number of expansion/swap moves for each label, or pair of labels, is exponential in the number of vertices/pixels, a direct search for the optimal expansion/swap move is unfeasible. However, it is possible to compute the optimal α -expansion or the optimal α/β -swap using the max-flow/min-cut algorithm

Algorithm 3

- 1: Start with an arbitrary label f
- 2: **repeat**
- 3: $f'' = \arg \min E(f')$ among $f' \in M_f$
- 4: if $E(f'') < E(f)$, then set $f = f''$
- 5: **until** $E(f'') \geq E(f)$

presented above, Section 4.1. This is because the computation of any of these moves is a binary minimisation problem which is regular/submodular when the swap/expansion inequality holds. In the remainder of this section it is described how to find the optimal move for both algorithms.

4.5.1 Swap Move

In order to compute the optimal move for a α/β -swap, let consider only that part of the graph which contains the nodes with labels α and β , denoted by $G_{\alpha,\beta} = (\mathcal{V}_{\alpha,\beta}, \mathcal{E}_{\alpha,\beta})$. Since it is a binary segmentation, one constructs the graph as for the max-flow algorithm, with the only difference that the terminal *source* vertex is renamed α and the terminal *sink* vertex is renamed β (Figure 4.5(a)).

For simplicity let consider two neighbouring pixels, corresponding to two vertices in the graph, the i and j vertices. The weights of the t-links are defined as:

$$t_i^\alpha = D_i(\alpha) + \sum_{\substack{j \in N_i \\ j \notin \alpha\beta}} V_{i,j}(\alpha, f_j) \quad t_i^\beta = D_i(\beta) + \sum_{\substack{j \in N_i \\ j \notin \alpha\beta}} V_{i,j}(\beta, f_j) \quad (4.5)$$

where N_i represents the neighbourhood of i . From Equation (4.5) one sees that the terminal weights depend on the data term $D_i(\cdot)$ as well as on the interaction with nodes




Figure 4.5: Two types of move as introduced in [16]: (a) α/β -swap graph, (b) α -expansion graph constructions.




Figure 4.6: Example of α -expansion graphs. In (a) the graph for $i = \alpha$ and $j = \beta$, while in (b) the graph for $i = \beta$ and $j = \gamma$.

which are neither α or β . After the graph is constructed the max-flow/min-cut algorithm can be used to perform an optimal cut. An example of the α/β -swap move can be seen in Figure 4.4(a).

4.5.2 Expansion Move

The α -expansion move is a binary segmentation where, for simplicity, the *source* and the *sink* vertices in the graph are respectively denoted as α and $\bar{\alpha}$. Note that $\bar{\alpha}$ represents any possible label in the label space different than α . Since all vertices can, potentially, change label, all the vertices are included into it $G = \{\mathcal{V}, \mathcal{E}\}$. In order to understand how to construct the graph, two neighbouring vertices i and j , are considered. The energy to be minimised based on this pair of vertices is:

$$V_{i,j} = V_{0,0}\bar{i}\bar{j} + V_{0,1}\bar{i}j + V_{1,0}i\bar{j} + V_{1,1}ij \quad (4.6)$$

$$\begin{aligned} &= V_{0,0}(1-i)\bar{j} + V_{0,1}(1-i)(1-\bar{j}) + V_{1,0}i\bar{j} + V_{1,1}(1-i)\bar{j} \\ &= i\bar{j} \underbrace{(V_{0,1} + V_{1,0} - V_{0,0} - V_{1,1})}_{\text{submodular term}} + i \underbrace{(V_{1,1} - V_{0,1})}_{\text{linear term}} + \bar{j} \underbrace{(V_{0,0} - V_{0,1})}_{\text{linear term}} + \underbrace{V_{0,1}}_{\text{const.}} \end{aligned} \quad (4.7)$$

where the terms $V_{i,j} = V_{i,j}(\cdot, \cdot)$ of Equation (4.3). The notation \bar{i} , or \bar{j} , indicates a label change, e.g. if i has label α , then \bar{i} indicates label $\bar{\alpha}$. Equation (4.7) contains a submodular/regular term, two linear terms, one each for the i and j vertices, and a constant term which represent a cost that needs to be added to the graph. The subscript indices, 0 and 1, represent two possible behaviours: 0 if the label does not change, and 1 if the label was $\bar{\alpha}$ and becomes α . The graph construction can be seen in Figure 4.5(b).

There are three possible label configurations for any pair of vertices i, j . The simplest configuration occurs for $i = j = \alpha$. Since both nodes are already labelled as α , in Equation 4.6 the terms proportional to $V_{0,1}$ and $V_{1,0}$ vanish because they are respectively multiplied by \bar{i} and \bar{j} , both of which are zero. The two remaining terms




Figure 4.7: Object interactions proposed in [37]. In (a) A contains B, in (b) A attracts B, and in (c) A excludes B.

$V_{0,0}, V_{1,1}$ do not contribute to the cost since the two nodes have the same label and the cut will not separate them. Therefore, in this simple case, no additional costs must be added to the graph.

The second case occurs when one of the two vertices is labelled as α and the other is labelled differently, for example $i = \alpha$ and $j = \bar{\alpha}$. The costs are now given by $V_{0,0} = V_{\alpha,\bar{\alpha}}$, $V_{0,1} = V_{\alpha,\alpha} = 0$, $V_{1,0} = V_{\alpha,\bar{\alpha}}$ and $V_{1,1} = V_{\alpha,\alpha} = 0$. Both the submodular term and the linear term for i vanish, while the linear term for j becomes $V_{0,0} - V_{0,1} = V_{\alpha,\bar{\alpha}}$ and needs to be added to the graph, Figure 4.6(a).

The third case occurs when neither vertex is labelled as α , for example $i = \beta$ and $j = \gamma$. In this case $V_{0,0} = V_{\beta,\gamma}$, $V_{0,1} = V_{\beta,\alpha}$, $V_{1,0} = V_{\alpha,\gamma}$ and $V_{1,1} = V_{\alpha,\alpha} = 0$. The resultant linear term for i is $-V_{\beta,\alpha}$, whereas for j this is $V_{\beta,\gamma} - V_{\beta,\alpha}$ and the submodular term becomes $V_{\beta,\alpha} + V_{\alpha,\gamma} - V_{\beta,\gamma}$, Figure 4.6(b). After the graph is constructed the max-flow algorithm can be used to perform an optimal cut. An example of the α -expansion move is shown in Figure 4.4(c).

4.6 Further Developments

Further developments in graph-cuts applications have addressed the study of object interactions and introduced, for example, object constraints such as connectivity.

4.6.1 Objects Interaction

Delong and Boykov analysed different object interactions in a multi-labelling framework [37]. Given two objects A and B , the authors studied the following three interactions: ‘ A contains B ’, ‘ A attracts B ’, and ‘ A excludes B ’ represented in Figure 4.7.

The first two interactions are submodular, whereas the exclusion constraint is not. These interactions are very common in medical applications, since they can represent the relative position of different organs or tissues. Therefore, all three interactions are analysed next.

To impose the inclusion constraint, the labels need to be nested. To see this, consider the label space $\mathcal{L} = \{1, 2, 3\}$ with a two layered graph, as prescribed by the Ishikawa construction, Equation 4.4. The Ishikawa construction already eliminated the 1 0 configuration by introducing an ∞ -edge between the two layers, but this does not suffice




Figure 4.8: Inclusion and attraction constraints proposed on [37]. The inclusion is imposed by the ∞ -edges, where the attraction is imposed via an edge weight α .

to enforce an inclusion constraint. To remedy this, Delong and Boykov included additional ∞ -edges between the layers that connect a given vertex in layer 0 with the neighbourhood of the corresponding vertex in layer 1, Figure 4.8.

As seen in Figure 4.8, pixel i_0 fan out several ∞ -edges reaching the neighbours of i_1 . The effect of these edges can be seen by considering two nodes i and j with different configurations: i is labelled as 0 1, and j as 1 1. The ∞ -edge from i_0 to j_1 imposes that around the nodes labelled as 1 1, there can only be vertices labelled 0 1 in addition to vertices labelled 1 1. Therefore the inclusion constraint is obtained.

The second interaction is ‘ A attracts B ’. This constraint is similar to the inclusion constraint, with the difference that instead of the ∞ -edges, through-layer edges are defined with a positive weight α , as shown in Figure 4.8.

The third and more complicated type of interaction is ‘ A excludes B ’. Consider the pair of vertices i, j , each belonging to a different object. Table 4.1 lists the cost of the four possible configurations, *i.e.* when $i_0 = 0$ or $i_0 = 1$, and $j_1 = 0$ or $j_0 = 1$:

i_0	j_1	$V_{i,j}$
0	0	0
0	1	0
1	0	0
1	1	∞

Table 4.1: Exclusion constraint costs for two neighbouring vertices i and j .

As it is possible to see $V_{i,j}(1, 1) = \infty$, therefore this interaction is not graph representable, as the regularity condition in Equation (4.3) does not hold. The energy minimisation is NP-hard. In this case, approximate solutions such as QPBO or α/β swap, described above, can be used.

Based on these simple interactions, it is possible to create more complex ones such as a hierarchy of nested regions, mutually exclusive regions, or combinations of these, Figure 4.9. Note that since the exclusion constraint is not graph-representable the objects B and C in Figure 4.9(b) have a non-empty intersection.

Multi label methods that take into account object interactions are very useful in medical image segmentation. Different organs and tissue types can be modelled as separate




Figure 4.9: Advanced inclusion constraints. In (a) inclusion and exclusion interactions. In (b) hierarchical configuration: where A contains both B and C and B excludes C .

objects, and their position or mutual interactions can be represented in graph form. For example, the heart can be represented as an object which contains two ventricles that exclude one another. A similar example is that of kidney segmentation, where the medullar regions are independent, mutually exclusive objects but are all contained inside the kidney [37]. In Chapter 6 will be shown how such constraints can be used to model the various brain tissues simultaneously.

4.6.2 Connectivity Constraint

Yet another useful extension of graph-cuts consists in the introduction of a so-called connectivity prior. The idea is to enforce that the object of interest consist solely of connected parts, that is, for this object no separate, independent parts are allowed in the segmentation. This kind of prior is very relevant to many medical applications: it can model, for example, bone or soft tissue connectivity. Connectivity priors were first introduced by Vicente et al. in [146]. The idea proposed by these authors is to perform an initial, constraint-free segmentation of the image, followed by manual selection of the points which must be connected to the main object of interest. The connection between the points and the main region is performed iteratively, where, at each iteration, the vertex which gives the best energy value is added to the object. This algorithm is very time consuming, in particular for 3D computations. In this thesis, a new type of connectivity prior has been developed that provides improved performance in terms of time and segmentation quality. This is described in Chapter 6.

Graph-cut methods, with their various extensions, have found widespread use in segmentation applications. In particular, they have been used in brain segmentation, including: segmentation of the neonate brain [122], segmentation of internal brain structures such as the hippocampus [142] and the subcortical nuclei [151], and tumour segmentation. In this thesis, a new method, based on graph-cuts, has been developed to segment brain white matter lesions. This is described in Chapter 6.

Chapter 5

Localisation of Focal Cortical Dysplasia

Focal Cortical Dysplasia (FCD) is a malformation of cortical development, and is one of the major causes of pharmacologically intractable epilepsy. This disease causes an abnormal development of the cortical grey matter characterised by an atypical stratification [101]. Patients affected by FCD suffer from recurrent seizures and only after a successful surgical treatment they can achieve seizure freedom. Therefore the precise pre-surgical localisation and complete surgical removal of the dysplastic tissue is pivotal for a successful outcome and for the avoidance of any additional interventions [41]. The four subtypes of FCD lesions are 1a, 1b, 2a, 2b [101]: types 2a and 2b are more visible in brain images compared to FCD of type 1a-1b. The development of high-quality brain imaging methods, in particular MRI, has contributed to the routine treatment of epilepsy. An example of an axial MR image of a patient with a typical FCD lesion (highlighted in red) is shown in Figure 5.1.

In this chapter we present our method to localise FCD lesions in brain MR images which has been published in [127] and is here reported in an extended format. The method is based on the extraction of multiple features from the patient data, and on the comparison of these with the same features extracted from healthy volunteers. We tested the algorithm on several confirmed FCD patients, and for all a correct localisation of the lesion has been obtained.




Figure 5.1: An axial slice of a typical example of an FCD lesion. In (a) the FLAIR image, in (b) the T_1 -weighted (T_1 -w) and in (c) the segmented lesion superposed on the T_1 -w image.

5.1 Related Work

There has been a considerable amount of work aimed at the localisation of FCD lesions, mainly by extracting various image features. FCD lesions show a variation in the intensities, changes in texture, cortical thickening, as well as blurring of the grey-matter/white-matter interface.

Most of the proposed methods rely solely on the information extracted from T_1 -weighted (T_1 -w) images, and to a lesser extent also on FLAIR images [46]. A large number of these methods follow some general steps: initial intensity inhomogeneity correction, segmentation into Grey Matter (GM), White Matter (WM) and CerebroSpinal Fluid (CSF), feature extraction, mainly based on GM tissue, and a final classification step to identify the lesion.

The method proposed by Louek et al. [88] extracts several features from T_1 -w images. As an initial step the method performs a segmentation into GM, WM and CSF using an atlas based approach [3]. All the relevant features are then computed in the GM mask, as FCD lesions appear only in cortex. The first-order features employed include mean, median, and variance of the GM intensities, as well as measures computed on the histogram such as skewness and kurtosis. Second-order texture properties are based on co-occurrence matrices [57], such as contrast and homogeneity. The last set of features is based on run-length-based techniques which calculate textural measures from grey-level run lengths in different image directions. All these features are subsequently used to train an SVM classifier, which separates the FCD lesion voxels from the healthy cortex voxels. A limitation of this method is its dependence on an atlas based approach to segment the GM, which may give imprecise results along boundaries.

A similar approach has been proposed by Antel et al. [2] where they also use a T_1 -w image to detect the lesions. The extracted features are the cortical thickness, blurring between WM/GM interface, hyper-intense signal compared to the GM-WM peaks, and texture information extracted with the co-occurrence matrices. The identification of the FCD lesions is performed using two Bayesian classifiers consecutively. The first classifier considers cortical thickness, blurring and hyper-intense signal, while the second classifier, based on the result of the first one, discriminates based on textural features.

The method proposed by Colliot et al. [32] combined a probabilistic estimation of FCD with a level-set approach. The method extracts features as in [2] and combines them to create a probability map. Subsequently a manual selection of a FCD region is needed to initialise the deformable model which then delineates the lesion boundary.

The method proposed by Bergo et al. [10] computes, among other features, the mid-sagittal plane which is then used to compare patches situated to the right and left of this plane. For each patch, the method extracts features such as mean and standard deviation of the intensities, contrast, entropy and homogeneity based on the co-occurrence matrix. The final classification step is based on a Reduced Coulomb Energy classifier [40].

Huppertz et al. [61] proposed to use a set of control images as a reference for normal appearing tissue, to determine an FCD probability map, or so-called 'junction image' I_J . The authors initially perform a tissue segmentation based on T_1 -w images, and then use

the GM and WM maps to identify two different thresholds: $T_{\text{low}} = \mu(\text{GM}) + 0.5 \cdot \sigma(\text{GM})$, and $T_{\text{high}} = \mu(\text{WM}) + 0.5 \cdot \sigma(\text{WM})$. All voxels in the T_1 -w image with intensities between T_{low} and T_{high} are initially selected to create a binary image, and subsequently a cubic smoothing kernel is applied to highlight clustered voxels. In their study, the authors processed each instance of healthy volunteer data to obtain a binary image, as described above. Based on this, mean and standard deviation images were subsequently calculated for the entire dataset of healthy volunteers. The junction image, or I_J , is obtained by voxel-wise subtraction of the mean image of the healthy volunteer database from the binary image of the patient, and dividing by the standard deviation image.

We propose a method that is based on an extended set of image features, extracted from T_1 -w and FLAIR images, which are then combined in a probabilistic framework. Similarly to the method proposed by Huppertz et al. [61], we also employ statistics of normal tissue from a database of healthy volunteers. An intensity feature is extracted from both FLAIR and T_1 -w images, whereas texture and form features are only extracted from T_1 -w image. The features are used for the characterisation of both controls and patients. The reference for normal statistics is obtained voxel-wise from the set of co-registered control images and is used to infer the probability of FCD. The FCD lesions are identified with high precision and increased sensitivity compared to the commonly used clinical method of Huppertz et al. [61]. The reference for validation consists in a visual evaluation by an expert physician as well as the resected region that results in freedom from seizures.

5.2 Pre-processing

In order to identify and localise FCD lesions two types of MR image contrast, *i.e.*, T_1 -w and FLAIR, are commonly used in clinical practice, as they provide complementary information. In what follows, the T_1 -w image is denoted as $I'_{T_1} : \Omega \rightarrow \mathbb{R}$, and the FLAIR image as I'_F . In addition to patient data, identified as P , a dataset of healthy volunteers data C are used which provide a control reference for the FCD lesions.

The following pre-processing steps are applied both to patient and to volunteer data. The first pre-processing steps are performed with the atlas-based tool SPM8 [124]: intra-subject registration to remove possible misalignment between the FLAIR and the T_1 -w images, patient normalisation to MNI space [42], and tissue segmentation into GM, WM and CSF regions (performed with the method proposed by Ashburner and Friston [5] and described in Section 3.2.6). To improve the robustness of the registration, the Anterior Commissure (AC) point is provided as input to SPM8 for inter-subject registration, Figure 5.2(a). Additionally, a brain localisation has been performed with the BET tool [121] using the T_1 -w image. This tool also provides a binary brain mask, which is applied after registration to the FLAIR image to remove the skull and the non-brain tissues. The intra-subject rigid registration \mathcal{T}_R has been estimated for all patients P and controls C . The reference, or fixed, image is I'_{T_1} while the moving image is I'_F , and




Figure 5.2: Pre-processing steps. The input images in (a) and (b), normalised I_F in (c) and I_{T1} (d), and tissue segmentation into WM, GM, and CSF from (e) to (h), respectively. The red circle in (a) surrounds the anterior commissure.

the image distance is the normalised mutual information for multi-contrast registration. Finally all subjects are normalised to T_1 -w MNI space with an affine registration \mathcal{T}_A . The equations below show that the registration is applied to both the FLAIR and the T_1 -w images:

$$I_{F,\mathcal{V}}(x) = \mathcal{T}_{A,\mathcal{V}}^{-1} \cdot \mathcal{T}_{R,\mathcal{V}}^{-1} I'_{F,\mathcal{V}}(x) \quad I_{T1,\mathcal{V}}(x) = \mathcal{T}_{A,\mathcal{V}}^{-1} I'_{T1,\mathcal{V}}(x),$$

where $\mathcal{V} = 0 \dots P - 1, P, \dots, P + C - 1$.

The non-uniform radio frequency field of the head coil results in intensity inhomogeneities in the images which have been jointly removed in the I_{T1} and I_F images using the co-occurrence method [56]. The final tissue segmentation into GM, WM and CSF is estimated on the basis of the bias-corrected I_{T1} image. In addition, intensity standardisation has been performed for all the patient and control data by matching their intensity ranges using three continuous piecewise linear segments for the dynamic ranges. The three linear segments are defined by the zero intensity point, the mode intensity of GM, the mode intensity of WM, and the highest intensity value. These are computed from the average image of the control set. The complete dynamic range of image intensities is then normalised to $[0, \dots, 1]$.

5.3 Extraction of Features

The proposed method uses a variety of intensity-based, texture-based, and form-based features which are extracted from I_{T1} and I_F . The various features are computed for each spatial location $x \in \Omega$, for both the images of the control dataset and for the patient data. In this framework a voxel x can be classified as belonging into one of two classes $\{H, FCD\}$, where H represents the class of healthy tissue, and the pathology can only be FCD. The features computed on the control data provide a spatial voxel-based reference for the normal appearing statistics H . Every feature f_i , where i is an index over the features, is modelled as a Gaussian distribution. Therefore, for each feature and each spatial location x , the mean $\mu(f_i(x))$ and the standard deviation $\sigma(f_i(x))$ are estimated, from which the probability that $f_i(x)$ follows a normal distribution, is denoted as $G(f_i(x); \mu(f_i(x)), \sigma(f_i(x)))$.




Figure 5.3: Axial slices of a FCD lesion in the yellow circle. In (a) the I_F , in (b) the I_{T1} , in (c) the intensity feature, in (d) the gradient magnitude, and in (e) $f_{\|\nabla I\|_2}$.

In the following subsections a description of the specific features and the probabilistic framework used to estimate the FCD probability is provided.

5.3.1 Intensity-based Features

In I_F , the FCD lesions appear hyper-intense compared to the normal-appearing GM tissue surrounding the lesions. Therefore, the intra-patient mean value $\mu_{GM}(I_F)$ and the standard deviation $\sigma_{GM}(I_F)$ over the GM region are estimated. Assuming that the probability of the normal appearing tissue follows a Gaussian distribution, the probability of voxel x being classified as FCD is equal to the complement of G evaluated at the mean value $\mu_{GM}(I_F)$. Thus, for I_F , the intensity-based feature f_F , is expressed as:

$$f_F(x) = \begin{cases} 1 - G(I_F(x); \mu_{GM}(I_F), \sigma_{GM}(I_F)) & \text{if } I_F(x) > \mu_{GM}(I_F) \\ 0 & \text{otherwise.} \end{cases} \quad (5.1)$$

In I_{T1} images, FCD lesions may show intensities similar to those in GM and healthy WM. Therefore, in order to model these intensities a larger range from $\mu_{GM}(I_{T1})$ to $\mu_{WM}(I_{T1})$ is considered. In this case the Gaussian distribution is modelled with $\mu = 0.5 \cdot (\mu_{WM}(I_{T1}) + \mu_{GM}(I_{T1}))$ and with a standard deviation of $\sigma = 0.5 \cdot (\sigma_{WM}(I_{T1}) + \sigma_{GM}(I_{T1}))$:

$$f_{T1}(x) = \begin{cases} 1 - G(I_{T1}(x); \mu(I_{T1}(x)), \sigma(I_{T1}(x))) & \text{if } \mu_{GM}(I_{T1}) < I_{T1}(x) < \mu_{WM}(I_{T1}) \\ 0 & \text{otherwise.} \end{cases} \quad (5.2)$$

Both f_F and f_{T1} are combined into a 2D distribution to model the hypo-intensity in I_{T1} and the hyper-intensity in I_F , Section 5.3.4. An example of the joint intensity feature for an FCD patient is shown in Figure 5.3(c).

5.3.2 Texture-based Features

FCD lesions are associated with a blurring of the boundary between GM and WM. In order to extract features to represent such characteristics, first a computation of the spa-




Figure 5.4: Axial slices of an FCD lesion in the yellow circle. In (a) the I_{T1} and in (b) the orientation feature $f_{\angle \nabla I}$.

tial derivatives of I_{T1} along the three axis X, Y, Z is computed: $\nabla I_{T1} = [I_X, I_Y, I_Z]$. The first feature intended to represent this texture in FCD regions is the *gradient magnitude* of the I_{T1} image (Figure 5.3(d)). The probability for a voxel to be in a lesioned area can be taken to be proportional to the negative of the log of the gradient magnitude [153]. The log scale is introduced to enhance the low gradient values typical of FCD lesions:

$$f_{\|\nabla I\|_2}(x) = -\log \|\nabla I_{T1}(x)\|_2. \quad (5.3)$$

Regions with high gradient magnitude will have a low $f_{\|\nabla I\|_2}$ and vice-versa, Figure 5.3(e).

The second texture feature highlights the diffuse boundary separation between GM and WM. It is expressed as the variance of the spatial orientation of the gradient, computed over I_{T1} [153]. Assuming that variations in gradient orientation, when projected onto the three co-ordinate planes, are mutually uncorrelated, this feature is computed as:

$$f_{\angle \nabla I} = \sigma_{X,Y,Z} \approx \sigma_{XY}(\angle(I_Y, I_X)) + \sigma_{XZ}(\angle(I_Z, I_X)) + \sigma_{YZ}(\angle(I_Z, I_Y)) \quad (5.4)$$

where $\angle(I_X, I_Y)$ represents, for example, the angle of orientation of the gradient on the X, Y plane, and so on, Figure 5.4.

5.3.3 Form-based Features

Focal cortical dysplasia lesions appear as regions of amorphous brain tissue or, equivalently, as a merging of WM and GM regions with no clear separating boundary between them. This is in contrast to a well-formed cortex that has a sharp boundary. The first form-based feature to represent this boundary is the *fractional anisotropy* [8], defined




Figure 5.5: Axial slices of a FCD lesion in the yellow circle. In (a) the I_{T1} and in (b) the orientation feature f_{th} .

considering the eigenvalues of the Hessian matrix, *i.e.* e_0, e_1, e_2 , in a window size of $(2\text{mm})^3$:

$$f_{FA}(x) = \sqrt{\frac{2}{3}} \cdot \frac{\sqrt{(e_0(x) - \mu(x))^2 + (e_1(x) - \mu(x))^2 + (e_2(x) - \mu(x))^2}}{\sqrt{e_0(x)^2 + e_1(x)^2 + e_2(x)^2}} \quad (5.5)$$

where $\mu(x) = (e_0(x) + e_1(x) + e_2(x))/3$.

The second feature computes the skewness of the histogram of local cortical thickness [153]. The thickness $t(x)$ of the GM is approximated as the shortest straight line segment passing from each voxel and ending on the two closest opposite cortical boundaries:

$$t(x) = \min_{l(x)} \int_{l(x')} I_{GM}(x') dl(x') \quad (5.6)$$

where l is the line segment of the two points connecting the WM and the CSF, and I_{GM} is an image representing the GM tissue provided from the GM segmentation.

The skewness of a voxel x is defined by considering a surrounding neighbourhood of size N , and estimating the mean, $\mu_t(x)$, and the standard deviation, $\sigma_t(x)$, of the thickness inside the neighbourhood. The sample skewness is defined by the expression:

$$f_{th}(x) = \frac{\sum_{n=1}^N (t_n(x) - \mu_t(x))^3}{(N-1)\sigma_t(x)^3}. \quad (5.7)$$

Note that Equation (5.7) can assume positive and negative values. An example of f_{th} computed on an FCD patient is shown in Figure 5.5.

5.3.4 Spatial Probability Map for FCD Lesions

The various features described in the previous paragraphs are computed for all the registered control and patient data sets. The features calculated for the patient data are compared, voxel-wise, to the corresponding normal statistics $\mu(f(x))$ and $\sigma(f(x))$.

A Naive Bayesian classifier is here used, and a further assumption of independency among the probability distributions of the various features is made [153]. Therefore the probability of H (healthy) can be expressed as:

$$P(H|f_0, f_1, \dots, f_{m-1}) \propto P(f_0, f_1, \dots, f_{m-1}|H)P(H) = P(H) \prod_{i=0}^{m-1} P(f_i|H) \quad (5.8)$$

where m is the number of features, and $f_i \in \{f_{th}, f_{FA}, f_{\angle \nabla I}, f_{\|\nabla I\|_2}\}$.

The voxel-wise probability of FCD is derived as the complement of the probability of healthy tissue given in Equation (5.8):

$$P(FCD|f_0, f_1, \dots, f_{m-1}) = 1 - P(H|f_0, f_1, \dots, f_{m-1}) = 1 - P(H) \prod_{i=0}^{m-1} P(f_i|H). \quad (5.9)$$

The probability of lesion in the 2D intensity joint distribution of f_F and f_{T1} is expressed as:

$$P(FCD|f_F, f_{T1})(x) = 1 - \prod_{k \equiv T1}^{FLAIR} G(f_k(x); \mu(f_k(x)), \sigma(f_k(x))). \quad (5.10)$$

Combining all features, gives the spatial probability map, P_{FCD} for FCD lesions:

$$P_{FCD} = P(FCD|f_F, f_{T1}) \cdot \prod_i P(FCD|f_i) \quad (5.11)$$

where $i \in \{th, FA, \angle \nabla I, \|\nabla I\|_2\}$. In order to constrain the FCD lesion only to the GM region an additional template is used, *i.e.* the Automated Anatomical Labelling (AAL) template [140], as a spatial prior. This template provided a parcellation of the GM tissue, and as it is already in MNI space, no additional registration is needed to align it to the patient data. The use of this atlas as a binary mask on the P_{FCD} is straightforward and considerably reduces the number of false positives. To further remove remaining noisy voxels, a Gaussian smoothing is added to P_{FCD} . The final FCD probability map is thereby:

$$P_{FCD} = P(FCD|f_F, f_{T1}) \cdot \prod_i P(FCD|f_i) P(AAL \equiv 1) P_G(FCD|FCD_N) \quad (5.12)$$

where $P(AAL \equiv 1)$ is the probability of being GM given from the AAL template, and the last term $P_G(FCD|FCD_N)$ is the spatial Gaussian smoothing computed over a neighbourhood N around each voxel.




Figure 5.6: From top to bottom: axial, sagittal, and coronal slices showing a typical example of FCD2b lesion. From the left to right: I_{T1} , I_F , P_{FCD} , I_J and the RV. The blue arrows point to the FCD lesion.

5.4 Results

The proposed method has been tested on an in-house database, composed by healthy controls and patients affected by FCD. First the used measures to quantify the performances are described, and subsequently the detailed results on the patient data. The method proposed here has been also compared with the clinically validated FCD feature called ‘junction image’, I_J [61].

In order to quantify the quality of the proposed P_{FCD} a comparative ground truth needs to be defined. Here, the post-operative MR image has been used to localise the resection volume (RV), or rather the brain area that has been surgically removed. In the pre-operative surgical planning, especially for refractory cryptogenic epilepsy, *i.e.* no visible MRI lesion, the signals from an electroencephalogram (EEG) are analysed to localise the lesions [136]. During a seizure the EEG usually reveals a spike in correspondence of the location where the seizure started. With this additional information, a more precise delineation of the FCD lesion can be performed. Note that this area, for many patients, is much larger than the actual MRI visible FCD lesion. This is mainly because, during surgery, all suspected diseased brain tissue is removed to avoid further interventions.




Figure 5.7: From top to bottom: axial, sagittal and coronal slices showing a typical FCD2a lesion. From the left to right: I_{T1} , I_F , P_{FCD} , I_J and the RV images. The blue arrows point to the FCD lesion.

In order to perform quantitative measures, a binary mask of the resection volume has been created for all patients using the ITK-SNAP tool [157]. In order to compare this mask with our results, the P_{FCD} has been thresholded at a value of 0.5. Since in most patients, the resection volume is larger than the actual lesion, the P_{FCD} may, in some cases, not cover the entire region but only part of it. The used quantitative measures are the Precision P and the Specificity S . The precision is defined as:

$$P = \frac{TP}{TP + FP}$$

where TP denotes the number of true positives, and FP is the number of false positives. The precision is a measure of the extent to which lesion voxels, *i.e.*, true positives, are correctly identified by the segmentation algorithm. The highest possible precision of 100% is diminished by the presence of false positives. The specificity is defined as:

$$S = \frac{TN}{TN + FN}$$

where TN is the number of true negatives, and FN is the number of false negatives. The best sensitivity value is 100%; this decreases with the the number of lesion voxels




Figure 5.8: From top to bottom: axial, sagittal and coronal slices showing a typical FCD1a lesion. From the left to right: I_{T1} , I_F , P_{FCD} , I_J and the resection volume (RV) images. The blue arrows point to the FCD lesion.

incorrectly missed out (false negatives) in the segmentation. Note that the specificity increases with the number of true negatives, which can itself be increased easily by doubling the imaging resolution. In our framework the resolution has been kept constant. The validity of this measure, however, remains considerably high owing to the difficulty of detecting FCD lesion.

Additionally, a validation was performed by an expert physician who visually compared the estimated P_{FCD} with the post-operative MRI which presented in the figures as the resection volume (RV), and verified the high specificity of the method.

It is very common for imaging data of epilepsy patients, in particular for children, to present motion artefacts. Where motion artefacts have been identified, the data has been excluded from the study. Here, $P = 11$ patients (4 women and 7 men; mean age 14.75; age range 5–38 years) who suffered from epilepsy caused by FCD of types 1a, 1b, 2a, and 2b have been considered. All of them underwent a surgical treatment that led to seizure freedom for at least two years. The histopathology that has been performed in the resected tissue confirmed in all patients the presence of FCD lesion(s) with their type. The study also included images from $C = 20$ controls (11 women and 9 men; mean age 27.2 years; age range 23–32 years). The study has been approved by the local review board and both patients and volunteers provided informed consent for the

Table 5.1: Quantitative evaluation of the proposed method (P_{FCD}) and the method presented in [61] (I_J) with respect to precision P and specificity S . The values are expressed in percentage.

Patient	P		S	
	P_{FCD}	I_J	P_{FCD}	I_J
01 Figure 5.6	46	18	68	76
02 Figure 5.7	57	16	80	74
03 Figure 5.8	50	17	50	65
04 FCD2b	48	20	87	70
05 FCD2a	49	16	70	70
06 FCD2b	50	12	79	52
07 FCD2a	53	18	67	70
08 FCD1a	45	19	60	61
09 FCD2a	54	15	70	72
10 FCD1b	53	17	61	65
11 FCD1b	54	15	60	60
Mean	51	17	68	67

analysis of their imaging data. Brain images have been acquired in a 3T Siemens Trio MRI Scanner equipped with a head coil. The field of view of the MR images covered of the head and the neck region. The acquisition protocol consists of a 3D MPRAGE T_1 -w sequence with TR/TE/FA = 1390ms / 2.15ms / 15°, and with a matrix size = 156 × 512 × 512; a 3D FLAIR image with TR/TE/FA = 5000ms / 338ms / 120°, and with a matrix size of 156 × 512 × 512.

The presented results include four different FCD subtypes, namely FCD1a, FCD1b, FCD2a, and FCD2b. For each patient the I_{T_1} , I_F , P_{FCD} , I_J and the RV are shown in all three projections, *i.e.* axial, sagittal and coronal. Figure 5.6 shows an FCD of type 2b, Figure 5.7 shows a lesion of type 2a, while Figure 5.8 shows a FCD of type 1a. In the probability maps P_{FCD} shown in Figure 5.6(c)(h)(o), and in Figure 5.7(c)(h)(o), the lesions have a high probability and are inside the RV. The high probability voxel are also clustered spatially in contrast to other irregularly placed voxels in the image.

Figure 5.8 illustrates an example of imaging data for an FCD1a type lesion. Both FCD1a and FCD2a type lesions are comparatively less conspicuous in MRI images [94]. As is shown in the first two columns of Figure 5.8 the lesion is not clearly distinguishable by visual inspection of the anatomical images, I_{T_1} and I_F , whereas the probability maps P_{FCD} (third column) do highlight a small fraction of the lesion that lies within the resection volume. Additionally, I_J shows the presence of stripe artefacts due to the cubic smoothing kernel of uniform weights applied after the generation of the binary map.

The quantitative evaluation for all eleven patients is shown in Table 5.1. For all patient data, both the precision and the specificity of the method proposed here are higher, on average, than those obtained from the I_J method.

5.5 Summary and Discussion

A new method for the 3D localisation of FCD lesions using MR images has been presented. The proposed method uses different image features which include information related to spatial variations in image intensity and to the shape of the lesions, as well as some particular characteristics such as the blurring of the GM-WM border and the increase in cortical thickness.

An intensity feature has been evaluated for both FLAIR and T_1 -w images, Section 5.3.1. For FLAIR images, this highlights regions with an intensity higher than the mean intensity of the GM. For T_1 -w images, this feature highlights regions where image intensities lie between those for GM and WM.

Focal cortical dysplasia lesions present blurred boundaries with the WM. This anatomical characteristic can be studied with the use of texture features, Section 5.3.2. The first feature involves the gradient of image intensity, $f_{\|\nabla I\|_2}$, which is comparatively low in lesioned areas (therefore the log gradient is used to enhance the contrast in these regions). A second feature which enhances the blurring of FCD lesions is the local incoherency of the gradient orientation $f_{\angle \nabla I}$. Lastly, the shape of the lesions is described by two form-based features, namely the fractional anisotropy f_{FA} , and the skewness, f_{th} , of the distribution of cortical thickness, Section 5.3.3.

The combination of all features has been able to represent the location of the lesions in all investigated cases. Therefore, all the features extracted from the patient data have been combined into a probabilistic map P_{FCD} which represents, for each voxel, the probability of being an FCD lesion.

All the above features have been computed also for a control database, to obtain the reference probability of normal cortical tissue. Comparison with the control database is necessary and fundamental to avoid potential false positives due to normal tissue variability. Additionally a spatial prior has been introduced, the *i.e.* ALL template, to constraint the P_{FCD} to cortical regions.

Due to the presence of motion artefacts several patient data have been excluded from the study. However the method has been tested on 11 FCD patients of types 1a, 1b, 2a and 2b, all of whom underwent surgical resection which confirmed the presence of FCD lesions. FCD2a and FCD1a lesion types are less conspicuous in MR images [94]. As is shown in the first and second columns of Figures 5.7 and 5.8, the lesions are not clearly visible in the FLAIR and T_1 -w images. On the other hand, the probability maps P_{FCD} in the third column of Figures 5.7, and 5.8 highlight regions corresponding not only to the resected volume but to other regions as well.

A visual comparison of our results with the method proposed by Huppertz et al. [61] which provides the probability map I_J has been performed. As shown in Figures 5.7, 5.6 and 5.8 the P_{FCD} shows high probability values in locations contained in the resected volume and has a reduced number of false positives compared to I_J , therefore it is possible to observe that the lesions are identified with a higher precision. Additionally a quantitative evaluation based on precision P and specificity S has been computed using the resection volume as ground truth. The latter, in general, covers a brain area which

is much bigger than the actual lesion visible in MR images. However, the results in Table 5.1 showed that our approach performs best, on average, in both measures. Note that even if the I_J produces a better segmentation in Figure 5.6 the precision value is low due to the very high number of false positive detected. For FCD2b and for FCD1a, Figure 5.6 and 5.8, the proposed method has a lower specificity compared to the I_J . This may be due to the extended resection volume, which is much larger than the actual lesion.

The method developed in this work significantly improves the localisation of FCD lesions, as compared to a clinically evaluated standard [61]. These quantification results are dependent on the resection volume, which can be, as already discussed, much larger than the actual lesion. Therefore, a manual annotation, provided by an expert physician, can help to have a more precise quantitative evaluation. Potential improvements include the quantification of the contribution of the single features to the total probability P_{FCD} and the enlargement of the control database. This, on the one hand, is expected to improve the estimation of local probabilities of healthy tissue; and it would allow a better matching of control data and patient data for any given case, thereby enhancing the probability of correct lesion detection.

Chapter 6

Segmentation of Multiple Sclerosis Lesions

Multiple Sclerosis (MS) is a neurodegenerative disease of the central nervous system, and its main hallmarks are lesions which are appearing in the white and grey brain matter. The pathogenesis of lesions is not completely clear, and understanding the nature and mechanism of tissue injury remains one of the most challenging aspects of the disease. Accordingly, accurate lesion identification and volumetry are crucial for quantifying the burden of the disease. In the clinical routine the lesions are manually annotated. This rather time-consuming process is prone to subjective variability and therefore offers poor reproducibility. Lesions present in grey matter usually cannot be well differentiated either in FLAIR or T_1 -weighted (T_1 -w) images, since they have very similar signal intensities to those of normal grey matter tissue. This chapter is concerned solely with lesions appearing in the white matter. An example of a typical MS patient is shown in Figure 6.1.

In this chapter our method for performing lesion segmentation in the brain white matter is described. Our algorithm is based on a geometric model of the brain which simultaneously segments the normal-appearing tissues, such as grey matter and white matter, and the MS lesions. Novel features of the proposed algorithm include the use of prior information such as the connectivity of the cortical grey matter and a minimum lesion size. The algorithm has been tested on two databases with 15 and 22 patient data respectively, and competitive results have been achieved compared to state of the art methods. The work presented in this chapter has been published in [131] and is here reported in an extended format.




Figure 6.1: Axial slice of a patient with a typical pattern of MS lesions. In (a) the T_1 -w image, in (b) the FLAIR image and in (c) the manually segmented lesions superimposed on the FLAIR.

6.1 Related Work

Lesion segmentation of the brain white matter is an extensively studied field of medical imaging post-processing. Most segmentation methods strongly rely on available atlas information, although a number of atlas-free algorithms also have been presented [49, 51, 149]. Other methods use machine learning techniques to discriminate between normal tissue and lesions. In the following, some of the most interesting methods published over the last years are discussed, pointing out the main differences with the segmentation method proposed in this chapter. For a detailed literature review the reader is referred to [50, 86].

As described in Section 3.2.6, atlas-based methods usually start with a registration of the patient data to an atlas in order to determine the locations of the major brain tissue, *i.e.* GM, WM, and CSF. The accuracy of the registration is important to initially discriminate healthy tissue from the lesions, and due to the high inter-brain variability this step may result in an erroneous tissue classification.

A publicly available atlas-based segmentation tool is Lesion-TOADS (Topology-preserving Anatomical Segmentation) [118], which is part of the MIPAV software [93]. The underlying method uses an intensity atlas and a topological atlas, both of which are computed from healthy volunteer data. In Lesion-TOADS the brain segmentation is formulated as an energy optimisation, and the minimisation iterates between the optimisation of the two atlases with respect to the patient data, to obtain both an intensity and a topologically consistent segmentation. Therefore, in this first step only the main tissue types are segmented. Since the lesions topologically belong to the WM, a two-class separation between WM and lesion voxels is performed subsequently. The lesion segmentation results are highly dependent on the atlas information and on the quality of the first segmentation step. In contrast to this approach, our method models the location of WM lesions simultaneously with the segmentation of the other brain tissues and, in addition, topological information is incorporated without the use of an atlas.

The method proposed by Souplet et al. [123] is less sensitive to atlas-driven misclassification. As a first step it performs an atlas-driven segmentation to obtain the GM, WM, and CSF masks of the patient data. Next, it creates a contrast enhanced FLAIR image which highlights the lesions. Based on this new image, outliers with respect to the WM mask are initially identified as potential lesions. Lastly, a morphological post-processing step renders the final lesion segmentation.

The Lesion Segmentation Tool (LST) proposed by Schmidt et al. [116] embodies an extension of the above method. The software is freely available as a MATLAB plugin [78]. This method computes outliers with respect to each of the tissue types, resulting in a so-called ‘belief map’ which contains an upper bound of the possible lesions. Hyper-intense voxels on this map are identified as seeds for a region-growing approach which produces the final lesion segmentation.

The Model of Population and Subjects (MOPS) [137] extends the classical atlas-based approach by using local intensity models instead of global ones. This method performs a tissue segmentation into GM, WM, CSF by coupling global and local Gaussian Mixture

Models (GMM). The global GMMs are computed on the patient data, one per each tissue type. The local models are extracted from a dataset of healthy volunteers, estimating a Gaussian Mixture Model (GMM) per pixel. These GMMs are then used to calculate in the patient data the likelihood of each tissue type. Lesions are initially identified as voxels having a low likelihood with respect to the local intensity models, and subsequently, a graph-cuts segmentation is performed to refine the segmentation.

Machine learning techniques, most notably classifiers, have been widely applied to segmentation problems. Usually these methods start with a segmentation of the brain into GM, WM, and CSF tissues. This information is then input to a pixel-wise classification step. In the presence of small lesions these classifiers may mix lesion and WM voxels. Zijdenbos et al. [162] created a pipeline based on an Artificial Neural Network (ANN) classifier to discriminate lesion and non-lesion voxels, using multiple image contrasts and three tissue segmentation (GM, WM, CSF) calculated from an atlas of healthy volunteers.

Geremia et al. [53] proposed a method in which a lesion classifier is trained using a random forest approach [19]. The method extracts many different features from the multiple image contrasts incorporating neighbourhood information for every pixel. Even if the neighbourhood is considered, the final decision is made independently for every pixel. An extension of this method was proposed by Cabezas et al. [26] which combines more classifiers into a Gentleboost framework [48].

Another interesting machine learning technique is the patch-based approach (Section 3.2.2). The size of the patch plays an important role in the representation of both lesion and healthy tissue. Weiss et al. [149] sparsely encode the patient image using a dictionary consisting of $3 \times 3 \times 3$ -patches. An ‘artificial’ image is then created from this dictionary, assuming that the amount of lesion-patches, in the dictionary, is small compared to the patches extracted from other tissues. The artificial image should therefore contain no lesion. The pixel-wise difference between the original and the artificial images highlights the lesions and some inherent reconstruction error. A final threshold applied to the difference image provides the lesion segmentation. Guizard et al. [55] propose a non-local mean approach introducing a rotational invariant distance measure (RMNMS) to account for the diversity of MS lesions.

Recently, non atlas-based methods have also been proposed for lesion segmentation. Gao et al. [49] presented a regularised segmentation method which encodes each tissue type with a constant intensity value while taking a global bias correction field into account. A fuzzy segmentation is computed by formulating a global energy functional that is convex in each variable. A local minimum is computed by an iterative minimisation approach. A graph-cuts binary segmentation method was proposed by Garcia-Lorenzo et al. [51] that discriminates MS lesions from healthy tissue. Post-processing steps were also included in order to remove lesions appearing in external CSF and in areas close the brain boundaries.

While non atlas-based methods clearly do not suffer from atlas-related misclassification, an important drawback is that they do not encode geometric information. Unlike the above mentioned methods, the non atlas-based method developed in this thesis does

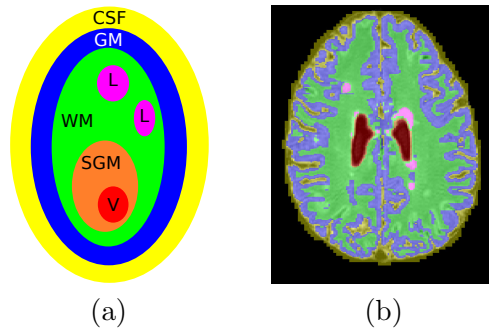


Figure 6.2: In (a) the brain geometric model used in the proposed method, and in (b) an example of a segmentation obtained using the geometric model.

incorporate geometric/topologic information, which, together with a graph-cuts minimisation approach, will be shown to produce more robust segmentation results.

6.2 Notation

The proposed method is based on the brain topologic assumption shown in Figure 6.2(a). The various tissue types are represented as independent circles with a relative specific position. Therefore, the cortical CSF is assumed to surround the cortical grey matter, which itself is adjacent to the white matter. Inside the white matter there are the lesions and the subcortical nuclei which contain the ventricles. The following notation is used to describe the proposed lesion segmentation method.

As already presented in Chapter 3, a medical image is defined as a mapping $I: \Omega \rightarrow \mathbb{R}^c$. In this work, $c = 2$ refers to the available FLAIR and T_1 -w image contrasts.

The goal of the brain segmentation method presented in this thesis is to determine a labelling $f: \Omega \rightarrow \mathcal{L}$ that assigns to each voxel x a label $f(x) \in \mathcal{L}$ of the label space $\mathcal{L} := \{\text{CSF}, \text{GM}, \text{WM}, \text{SGM}, \text{V}, \text{L}\}$. These labels refer, respectively, to sulcal CSF, grey matter, white matter, subcortical grey matter nuclei, ventricles and lesions.

In addition, the region identified by the label ℓ is denoted as R_ℓ , Section 3.1. Lastly, the hard constraints voxels used to guide the minimisation process are denoted as $S_\ell = \{x \in \Omega | f(x) = \ell\}$ for $\ell \in \{\mathcal{L} - \text{L}\}$.

In this chapter the following colour code is used for labelling: CSF appears in yellow, GM in blue, WM in green, SGM in orange, V in red and L in magenta. In the case that only the binary mask of the lesions is presented, then green is used for the manual annotation and orange for the proposed segmentation.

6.3 Pre-processing

The analysis of MR brain images requires a number of pre-processing steps in order to correct for intensity inhomogeneities (which are typical artefacts for this kind of images) and to eliminate potential sources (like the skull and the non-brain tissues) of registration errors. In this section all the necessary steps are described. First the image registration is performed to align the input images, then inhomogeneity correction is computed. In a final step, a partial segmentation of the brain is performed and it is used as hard constraint for the proposed method.

6.3.1 Image Registration

As already discussed in Section 3.2.6, the quality of the registration is fundamental for a correct segmentation. Since the c different patient image contrasts, used as the input to our method, may have a relative misalignment of a few voxels, a rigid registration is used to align them. This transformation is estimated using the default parameters of the ITK libraries [62].

Subsequently skull stripping is performed using the BET toolkit [121]; this step removes the signal from the skull and from the non-brain tissues, avoiding any potential miss-classification in this region. This tool also provides a brain mask which can be used to speed up the computations by considering only brain tissue voxels, avoiding any computation over the background voxels.

6.3.2 Inhomogeneity Correction

Due to the inhomogeneities inherent to the MR acquisition process, a correction for the bias field in the images is needed. In this thesis, this is performed by extending the method proposed by Li et al. [80] to the case of $c > 1$ image contrasts. This method has been originally proposed for $c = 1$ and in the following this assumption is used; whenever a modification is needed to allow for $c > 1$, this will be clearly indicated. The method [80] models the bias field in the image and generates a segmentation into the three main brain tissue types, *i.e.* GM, WM, and CSF. It starts using the following formulation for an image:

$$I(x) = b(x) \cdot J(x) + n(x)$$

where I denotes the observed image, b is the bias field, J is the noise- and bias- free image and n represents additive Gaussian noise. In the ideal image J the i th tissue type can be represented by a constant value v_i , that is, J is piecewise constant:

$$J(x) = \sum_{i=1}^N v_i u_i(x)$$

where N is the number of tissues, and u_i is a membership function such that:

$$u_i(x) = \begin{cases} 1, & \text{if } x \in \Omega_i \\ 0, & \text{otherwise} \end{cases} \quad \sum_{i=1}^N u_i(x) = 1$$

where Ω_i represents the region for the tissue i . These membership functions u_i are actually the tissue segmentation masks. The bias field b is assumed to vary slowly and is represented as a linear combination of a set of basis functions:

$$b(x) = \sum_{m=1}^M w_m g_m(x)$$

where w_m are the weights for cosine basis functions g_m . The energy that the method in [80] seeks to minimise is given by:

$$F(U, \mathbf{v}, \mathbf{w}) = \sum_{i=1}^N \int_{\Omega} |I(x) - (\mathbf{w}^T G(x))v_i|^2 u_i(x) dx, \quad \text{with } b(x) = \mathbf{w}^T G(x)$$

where $\mathbf{w} = (w_1, \dots, w_M)^T$, $G(x) = (g_1(x), \dots, g_M(x))^T$, $U(x) = (u_1(x), \dots, u_N(x))^T$ and $\mathbf{v} = (v_1, \dots, v_N)^T$. This energy can be minimised iteratively by performing interleaved minimisation with respect to each variable. The minimisation with respect to the vector of tissue intensities, \mathbf{v} , regards U and \mathbf{w} as fixed, and amounts simply to the computation of the average intensity value with respect to every tissue type:

$$\hat{v}_i = \frac{\int_{x \in \Omega_i} I(x)b(x)u_i(x)dx}{\int_{x \in \Omega_i} b^2(x)u_i(x)dx}. \quad (6.1)$$

Since in our configuration there are multiple images, a mean value per tissue per image contrast is computed.

The minimisation of F with respect to the basis function weights, \mathbf{w} , is computed by taking the derivatives of F with respect to \mathbf{w} , with U and \mathbf{v} held constant. This results in:

$$\frac{\partial F}{\partial \mathbf{w}} = -2\mathbf{s} + 2A\mathbf{w}, \quad \mathbf{s} = \sum_{x \in \Omega} I(x)G(x)J(x) \quad \text{and} \quad A = \sum_{x \in \Omega} G(x)G(x)^T J^2(x)$$

where $J = \mathbf{v}^T U$. The solution for \mathbf{w} is:

$$\hat{\mathbf{w}} = A^{-1}\mathbf{s}. \quad (6.2)$$

Finally the minimisation of the membership functions u_i is computed according to the formula below, with \mathbf{v} and \mathbf{w} kept fixed:

$$i_{\min} = \arg \min_i |I(x) - (\mathbf{w}^T G(x))v_i|^2.$$

Note that for multiple input images, *i.e.* $c > 1$, the best i value for the voxel x is computed considering all images:

$$i_{\min} = \arg \min_i \sum_{im=1}^c |I_{im}(x) - (\mathbf{w}^T G(x)v_i)|^2. \quad (6.3)$$

The steps for performing inhomogeneity correction and obtaining an initial segmentation are summarised in Algorithm 4.

The outputs of this algorithm are the bias free image/images, and the segmentation masks of the main tissue types. The value for N has been set to $N = 3$ and therefore an initial segmentation of the brain tissues with just three labels is obtained, Figure 6.3(a)-(e). These labels are combinations of the six labels in \mathcal{L} , *i.e.*, $csf = \{\text{CSF}, \text{V}\}$, $gm = \{\text{GM}, \text{SGM}\}$, $wm = \text{WM}$ as shown in Figure 6.3(e).

6.3.3 Partial Segmentation

Given the wm , gm , csf masks, a partial segmentation with respect to a subset of \mathcal{L} is subsequently computed, that is used as hard constraint for the proposed method.

In order to separate V from CSF, a binary graph-cuts segmentation, Section 4.2, is performed on the csf mask. The seed points necessary to guide the graph-cuts segmentation are extracted according to a distance criterion with respect to the brain boundaries. In order to identify the brain boundaries, the brain mask provided by the BET tool is used. The seeds for the CSF are those voxels in the csf mask which lie closest to the brain boundaries. The seeds for the V are computed taking into account the central location of the V, therefore these voxels need to have a distance to the brain boundaries which reaches the brain centre. Once the seeds voxels are identified, a binary segmentation which separates the lateral ventricles from the cortical CSF is computed. In our setting, the label V includes, anatomically, the two lateral ventricles. Note that at this step of the segmentation process, a precise separation is not needed since these are only initial estimations that will be corrected later.

The separation of SGM from GM on the gm mask is similarly performed via a graph-cuts segmentation. In this case one can use the segmentations CSF and V obtained in the previous step in order compute the seeds for SGM and GM. More precisely the seeds

Algorithm 4 Inhomogeneity correction and initial segmentation

- 1: Initialise U , \mathbf{v} and \mathbf{w}
 - 2: **repeat**
 - 3: Update v according to \hat{v} , Equation (6.1)
 - 4: Update \mathbf{w} according to $\hat{\mathbf{w}}$, Equation (6.2)
 - 5: Update U according to Equation (6.3)
 - 6: **until** Memberships u_i do not change
-




Figure 6.3: Based on the original FLAIR (a) and T_1 -w (b), the bias field (c) is computed [80] in order to correct the input images, for example the FLAIR (d). The bias field computation also provides coarse segmentation mask of gm , wm and csf (e). These masks are subsequently refined in order to obtain masks for SGM and V (f). Strongly eroding these masks $S_\ell^{(0)}$, provides the masks S_ℓ that serve as hard constraints (g) for our segmentation framework (Section 6.4.1).

for the SGM are voxels in the gm mask close to V, and the seeds for GM are the voxels in gm close to CSF.

These steps produce an initial labelling $f_0: \Omega \rightarrow \mathcal{L} - \{L\}$ and an initial segmentation $S_\ell^{(0)} := \{x \in \Omega | f_0(x) = \ell\}$, Figure 6.3(f).

The method of Li et al. [80] provides three masks: gm , wm , and csf . Therefore, lesions are inside $S_{WM}^{(0)}$, $S_{GM}^{(0)}$, and $S_{SGM}^{(0)}$. To have a correct initial segmentation, the lesion voxels need to be removed from these segmentation masks. Assuming that lesions are hyperintense in FLAIR, one can take, as a conservative threshold, the median value of the FLAIR intensities. Hence, all voxels with an intensity value higher than the threshold are eliminated from $S_{WM}^{(0)}$, $S_{GM}^{(0)}$, $S_{SGM}^{(0)}$. To be completely sure that these segmentations do not contain any lesion voxels, an erosion of each segmentation with a sphere of radius 2 mm is performed, which results in the segmentation masks S_ℓ , Figure 6.3(g). Note that due to the strong erosion the S_ℓ have a coverage of 15% of the image domain. The S_ℓ are hard constraints for our method, so they will not change their label, and furthermore they will guide the segmentation process.

The following section describes the energy that is minimised to obtain a segmentation of the brain which represents the lesions. In the optimisation process it is described how to rely on the few voxels of the hard constraints S_ℓ to guide the minimisation, Section 6.5.

6.4 Geometric Brain Model

In this section the proposed geometric model to represent the brain is described. Moreover, it is explained how it is included in the energy formulation.

Given an input image I , the objective is to determine the most likely probable labelling f . The Bayesian formulation presented in Section 3.1 is briefly recalled below for ease of reference. The logarithm of the posterior probability of the labelling f is given by:

$$-\log P(f|I) = -\log(P(I|f)) - \log(P(f)) + \log(P(I)) \quad (6.4)$$

where $P(I)$ denotes the probability of the observed image I , which is treated as constant, and $P(I|f)$ describes the likelihood that the image I ‘originates’ from the labelling f . A standard way to estimate this likelihood is to compute the respective probabilistic intensity models for each tissue type. For this reason, in the sequel, the term $E_{\text{App}}(f) := -\log(P(I|f))$ is referred to as the *appearance energy*. The term $P(f)$ models the prior information of the geometric model. Two independent geometric properties are modelled, *i.e.*, $P(f) = P_{\text{Loc}}(f) \cdot P_{\text{Geo}}(f)$. The $P_{\text{Loc}}(f)$ models the fact that only certain transitions between different tissue types are allowed, while the term $P_{\text{Geo}}(f)$ enforces the global geometric properties which the brain model must satisfy. Defining $E_{\text{Loc}}(f) := -\log(P_{\text{Loc}}(f))$ as the *local geometric energy* and $E_{\text{Geo}}(f) := -\log(P_{\text{Geo}}(f))$ as the *global geometric energy*, our model attempts to minimise the following energy:

$$E(f) := E_{\text{App}}(f) + E_{\text{Loc}}(f) + E_{\text{Geo}}(f). \quad (6.5)$$

In the remainder of this section, every energy is described in detail. The minimisation of Equation (6.5) will be addressed in Section 6.5.

6.4.1 Appearance Energy

In brain lesion segmentation a standard way to represent tissue intensities is to model each tissue type with a Gaussian Mixture Model (GMM) [137, 143], Section 3.2.4. Here, the GMM probability of a tissue $\ell \in \mathcal{L}$ is denoted by p_ℓ . The probability to assign the label ℓ to a given voxel x with intensity $y = I(x)$ is given by:

$$p_\ell(y) = \sum_{i=1}^{K^\ell} \alpha_i^\ell G(\mu_i^\ell, \Sigma_i^\ell; y) \quad \sum_{i=1}^{K^\ell} \alpha_i^\ell = 1. \quad (6.6)$$

Here K^ℓ is the number of Gaussian models that are mixed in order to model the tissue type ℓ . Each Gaussian model $G(\mu_i^\ell, \Sigma_i^\ell; \cdot)$ is defined by its mean intensity μ_i^ℓ and its covariance matrix Σ_i^ℓ . The mixture parameters α_i^ℓ determine the importance of each Gaussian model in the distribution p_ℓ . Note that the GMM are multi-dimensional, as in this work two image contrasts have been considered.

Here the first usage of the initial segmentations S_ℓ from Section 6.3 is encountered. Since the voxels in these masks are hard constraints, they will not change label. Therefore the probability that a voxel x is labelled ℓ is:

$$P_\ell(x) = \begin{cases} 1 & \text{if } x \in S_\ell \\ 0 & \text{if } \exists \ell' \neq \ell : x \in S_{\ell'} \\ p_\ell(I(x)) & \text{otherwise.} \end{cases} \quad (6.7)$$

As it is possible to see from Equation (6.7), the probability of the voxels belonging to S_ℓ is unity, meaning that there is no need to compute the intensity probability value for those voxels.

In the appearance model P_ℓ , the assumption is made that, for each voxel x that belongs to the tissue type ℓ , the voxel intensity is an independent, identically distributed random variable. Thus the *appearance energy* can be written as:

$$E_{\text{App}}(f) = -\log \left(\prod_{x \in \Omega} P_{f(x)}(I(x)) \right) = \sum_{x \in \Omega} -\log (P_{f(x)}(I(x))) \quad (6.8)$$

where $f(x)$ represents the label for voxel x . From this equation one can see how various S_ℓ contribute to the overall energy. In particular, assigning to a given voxel $x \in S_\ell$ a label $\ell' \neq \ell$ is penalised with an infinite cost ($-\log(0) = +\infty$). This implies that $S_\ell \subset R_\ell$ for all $\ell \in \mathcal{L} - \{\text{L}\}$, therefore the region R_ℓ must always contain S_ℓ . Thus, the inclusion constraint $S_\ell \subset R_\ell$ can be considered as a global constraint that is at all times enforced during the optimisation process.

6.4.2 Local Geometric Model

The minimisation of Equation 6.5 with respect to the *appearance energy* alone would result in a noisy segmentation, since E_{App} encodes no neighbourhood or geometric information. Therefore, to avoid overly noisy labelling it is common to penalise neighbouring voxels if their labels disagree. This penalty is represented as a pair-wise term between adjacent voxels. In the method developed in this thesis, these pair-wise terms are also used to additionally impose local topological constraints. These topological constraints represent our geometric model, see Figure 6.4 (r.h.s.): each tissue schematically represented in the figure by a colour region, has a specific relative position. As an example the CSF lies on the outermost part of the brain, and can only be adjacent to the cortical grey matter (GM).

In order to impose these constraints and to achieve a smooth segmentation, a 6-neighbourhood $N(x)$ is used. The local geometric model is formulated as:

$$E_{\text{Loc}}(f) = \sum_{x \in \Omega} \sum_{y \in N(x)} \lambda_{f(x), f(y)} \cdot g(x, y) \quad (6.9)$$

where $g : \Omega \times \Omega \rightarrow \mathbb{R}^+$ is a gradient based function which provides a high penalty between pixels of similar intensities [15, 17, 70]. This function is similarly defined as the one in Equation 4.2:

$$g(x, y) = \exp \left(-0.5 \cdot \left(\frac{|I(x) - I(y)|}{\sigma} \right)^{1.5} \right) \cdot \frac{1}{\text{dist}(x, y)} + \epsilon.$$

The positive scalar value $\epsilon = 0.01$ guarantees that the $g(x, y)$ always exceeds the threshold ϵ , σ guarantees that the method is not sensitive to the dynamic intensity range of image I , and the distance $\text{dist}(\cdot, \cdot)$ between the two voxels takes into account the anisotropic resolution of I . For any two voxels x and y with different labels ($f(x) \neq f(y)$), if $|I(x) - I(y)| < \sigma$ then $g(x, y)$ represents a relatively high penalty compared to

	CSF	GM	WM	SGM	L	V
CSF	0	$\lambda_{\text{CSF,GM}}$	∞	∞	∞	∞
GM	$\lambda_{\text{CSF,GM}}$	0	$\lambda_{\text{GM,WM}}$	∞	∞	∞
WM	∞	$\lambda_{\text{GM,WM}}$	0	$\lambda_{\text{WM,SGM}}$	$\lambda_{\text{WM,L}}$	∞
SGM	∞	∞	$\lambda_{\text{WM,SGM}}$	0	∞	$\lambda_{\text{SGM,V}}$
L	∞	∞	$\lambda_{\text{WM,L}}$	∞	0	∞
V	∞	∞	∞	$\lambda_{\text{SGM,V}}$	∞	0




Figure 6.4: The relationship between two different tissue layers ℓ_1 and ℓ_2 is modelled by weighting their interface with a factor $\lambda_{\ell_1,\ell_2} \geq 0$. An infinite cost prevents two regions from having a common boundary. The λ -values are based on the geometric model shown in Figure 6.2(a) and here reported in the right hand side.

$|I(x) - I(y)| > \sigma$. Moreover, for a given intensity difference $|I(x) - I(y)|$, the assigned penalty is higher the smaller the distance $dist(x, y)$. The symmetric weighting factor $\lambda_{f(x),f(y)}$ of Equation (6.9) handles different label transitions. To favour smooth segmentations $\lambda_{\ell_1,\ell_2} = 0$ if $\ell_1 = \ell_2$. Therefore there is no additional costs in our energy in case two neighbouring voxels have the same label, *i.e.* $f(x) = f(y)$. In addition, only those segmentations that satisfy the topological restrictions described below are allowed (see also Figure 6.4):

- **CSF** is adjacent only to GM, then $\lambda_{\text{CSF},\ell} = \infty$ for $\ell \in \{\text{WM, SGM, L, V}\}$.
- **GM** is adjacent to the WM and it is surrounded by CSF, then $\lambda_{\text{GM},\ell} = \infty$ for $\ell \in \{\text{SGM, L, V}\}$.
- **WM** surrounds the lesions and is adjacent to the GM as well as the SGM. Hence, $\lambda_{\text{WM},\ell} = \infty$ only for CSF and V.
- **L** are only found inside the WM, then $\lambda_{\text{L},\ell} = \infty$ for $\ell \in \{\text{CSF, GM, SGM, V}\}$.
- **SGM** is adjacent to the ventricles and the WM, then $\lambda_{\text{SGM},\ell} = \infty$ for $\ell \in \{\text{CSF, GM, L}\}$.
- **V** are modelled as being surrounded by the subcortical GM. Therefore, $\lambda_{\text{V},\ell} = \infty$ for $\ell \in \{\text{CSF, GM, WM, L}\}$.

The different transitions are illustrated in the sketch of Figure 6.4, while the values for λ_{ℓ_1,ℓ_2} are summarised in the table of Figure 6.4. Note that the ∞ values represent impossible transitions. In our geometric model there are six labels which occupy specific relative positions leading to five allowed transitions: CSF-GM, GM-WM, WM-L, WM-SGM, and SGM-V.




Figure 6.5: FLAIR image in (a). Resultant segmentation without (b) and with enforcement of the minimal size constraint in Equation (6.13)(c). Upon enforcing the constraint, a small lesion can either disappear or grow to satisfy the constraint (arrow in (b,c)).

6.4.3 Global Geometric Model

The intensities of the cortical and subcortical grey matter, and of lesions are all very similar both in FLAIR and in T_1 -w images, therefore the probabilities p_ℓ of these three tissues are not sufficiently discriminative to generate a correct segmentation. If one would compute a segmentation based only on E_{App} and E_{Loc} , most of the lesions would be segmented as either GM or SGM since the resultant energy costs would be smaller than the cost of introducing the label L, Figure 6.11(b).

To overcome this limitation, additional global constraints to the regions R_ℓ of these labels, *i.e.*, R_{GM} , R_{SGM} , are introduced. Anatomically, the GM is continuously connected; this feature will also be assumed to hold for SGM (Figure 6.4). Therefore, a connectivity prior on R_{GM} and R_{SGM} can be enforced, resulting in the prior $P_{\text{CC}}(f)$.

In addition, noise in MR images can appear as local hyper-intensities in the WM, and these small regions can be wrongly segmented as lesions. To limit these false positives a minimum size constraint for the lesions has been introduced, enforcing that each lesion contains at least a ball of 1 pixel radius. This results in the prior $P_{\text{min}}(f)$. These priors are modelled as mutually independent probabilities, obtaining:

$$P_{\text{Geo}}(f) = P_{\text{CC}}(f) \cdot P_{\text{min}}(R_L). \quad (6.10)$$

Taking the negative logarithm, the global energy is obtained:

$$E_{\text{Geo}}(f) = -\log(P_{\text{Geo}}(f)) = -\log(P_{\text{CC}}) - \log(P_{\text{min}}(R_L)). \quad (6.11)$$

Let $E_{\text{CC}}(f) = -\log(P_{\text{CC}})$ and $E_{\text{min}}(f) = -\log(P_{\text{min}}(R_L))$, then the global energy can be re-written as:

$$E_{\text{Geo}}(f) = E_{\text{CC}}(f) + E_{\text{min}}(f)$$

with:

$$E_{\text{CC}}(f) = \begin{cases} \infty & \text{if } R_{\text{GM}} \text{ or } R_{\text{SGM}} \text{ are not connected} \\ 0 & \text{otherwise} \end{cases} \quad (6.12)$$




Figure 6.6: FLAIR image in (a). The minimisation of $E_0(f)$ results in the labelling shown in (b), which contains the *false labels* F_1 (grey) and F_2 (dark green). In (c) the removal of the *false labels*. Final labelling at low(d), middle (e), and original (f) resolution.

and

$$E_{\min}(f) = \begin{cases} 0 & \forall A \subset R_L \text{ is connected component, } \exists x \in A : \forall y \notin A : \|x - y\| > 1 \\ \infty & \text{otherwise.} \end{cases} \quad (6.13)$$

An example of the minimal size prior is shown in Figure 6.5. Lesions with a radius smaller than 1 voxel can either disappear or increase in size, Figure 6.5(c).

6.5 Energy Optimisation

The energy in Equation (6.5) is not submodular [73], because the exclusion constraint between L and SGM, Figure 6.4, is not graph-representable, and the regularity condition (Section 4.3) in Equation (4.3) is not satisfied. Therefore in order to minimise Equation (6.5) it is necessary to search for a local optimal solution. In this thesis, a coarse-to-fine approach is followed whereby at the lowest resolution level an optimal solution is found by excluding certain constraints. As explained in detail below, based on this solution it is feasible to recover all the constraints that our model imposes, and due to local optimisation a final segmentation in the image original resolution is achieved.

The coarse-to-fine approach subsamples the image $I : \Omega \rightarrow \mathbb{R}^c$ and its image domain Ω by a factor of $s_{x_1}^{(1)} \times s_{x_2}^{(1)} \times s_{x_3}^{(1)}$ to obtain $I_1 : \Omega_1 \rightarrow \mathbb{R}^c$. By subsampling I_1 by another factor of $s_{x_1}^{(2)} \times s_{x_2}^{(2)} \times s_{x_3}^{(2)}$, a second image $I_2 : \Omega_2 \rightarrow \mathbb{R}^c$ is obtained. All the subsampling factors are greater than or equal to 1, and their actual value depends on the resolution of the data. As a result, a pair of lower resolution images, I_1 and I_2 , is obtained from the original image I . Note that in our method with two image contrasts, both input images are subsampled.

In Section 6.5.1 the input images for the proposed method are described, Figure 6.7, while in Section 6.5.2 an initialisation for the labelling $f|\Omega_2$ is computed, Figure 6.6(b). In Section 6.5.3 an improvement with respect to $f|\Omega_2$ is obtained, Figure 6.6(c). After finding a local minimum for $f|\Omega_2$, Figure 6.6(d), the resolution is iteratively increased in




Figure 6.7: Original FLAIR and T_1 -w images in (a) and (c). In (b) and (d) the result of the denoising algorithm proposed by Chambolle [29]. Red arrows point to two representative lesions.

order to find a local minima for $f|\Omega_1$, and finally for f itself, Figure 6.6(e)-(f). The λ_{ℓ_1, ℓ_2} values of Section 6.4.2 are fixed parameters with the exception of $\lambda_{WM, L}$. In Section 6.5.4 two globally optimal hyper-parameters are learned that allow to choose the correct value for $\lambda_{WM, L}$.

6.5.1 Input Images

Additionally to the T_1 -w and FLAIR images, the method considers a denoised version of both images for the computation of the appearance energy in Equation (6.8). This additional information can help to better distinguish the intensity model of each label compared to the others.

In order to smooth an image while simultaneously preserving the edges, the method proposed by Chambolle [29] and presented in Section 3.3.1, has been used. The parameter σ , which controls the tradeoff between smoothness and fidelity to the original data, has been kept constant for all the image resolutions and empirically estimated, $\sigma = 0.25$. The value for λ depends on the actual image resolution: a large value would preserve all structures in the image, while a small value would result in an image with an average intensity. Therefore, λ has been empirically set to $\lambda = 20$ for I_2 , $\lambda = 40$ for I_1 $\lambda = 60$ for I . The additional parameter τ is determined according to the inequality in Section 3.3.1. These values have all been chosen according to the smoothness and the intensity difference between the tissue types present in the denoised image. An example of this denoising method applied to both FLAIR and T_1 -w images is shown in Figure 6.7. The WM is seen to be more uniform in the denoised image Figure 6.7(b)-(d), while small objects as the lesions are still visible, arrows in Figure 6.7(b).

With these two additional images the appearance energy E_{App} is, at each resolution, computed over four images, *i.e.*, the original and the denoised FLAIR and T_1 -w images.

6.5.2 Computation of the Initial Labelling

The minimisation of Equation 6.5 starts with an initial estimation, for each label, of its associated GMM, which consists of a mixture of two Gaussians. The GMMs are initially estimated from the results of the initial segmentations S_ℓ (Section 6.3.3) and are subsequently used in E_{App} . For each S_ℓ , its GMM is determined by first using a K -means clustering, Section 3.2.4 to obtain an initial partition of the voxels. Subsequently an Expectation-Maximisation algorithm (Section 3.2.4), is initialised with the mean values of the clusters obtained with the K -means. After the EM converges the final GMM model parameters are obtained. This procedure is implemented in the `mlpack` library [35], and is here used with the default parameters.

Since, at this stage, no lesion segmentation S_L is available, it is not possible to use the above procedure to estimate the GMM model for the lesions. The lesions appear hyper-intense in the FLAIR image, and hypo-intense in the T_1 -w image, and for both images one can use the mean intensity of S_{WM} and S_{GM} to estimate the GMM for the lesions, Figure 6.7. The mean lesion values for the FLAIR image, $\mu_L(\text{FLAIR})$, and for the T_1 -w image, $\mu_L(T_1)$, are estimated as:

$$\begin{aligned}\mu_L(\text{FLAIR}) &= \mu(S_{\text{GM}}(\text{FLAIR})) + (\mu(S_{\text{GM}}(\text{FLAIR})) - \mu(S_{\text{WM}}(\text{FLAIR}))) \\ \mu_L(T_1) &= \mu(S_{\text{WM}}(T_1)) - (\mu(S_{\text{WM}}(T_1)) - \mu(S_{\text{GM}}(T_1)))\end{aligned}$$

The covariance matrix is estimated based on the variance of the S_{WM} in both images:

$$\Sigma_L = \begin{bmatrix} \sigma^2(S_{\text{GM}}(\text{FLAIR})) & 0 \\ 0 & \sigma^2(S_{\text{GM}}(T_1)) \end{bmatrix} \quad (6.14)$$

Delong and Boykov proved in [37] that a simplified version of the problem stated in Equation (6.5) does fit into the Ishikawa formulation, Section 4.4. To achieve this, the label space must be modified (see below) and the following energy $E_0(f)$ must be considered instead, which does not incorporate geometric information:

$$E_0(f) := E_{\text{App}}(f) + E_{\text{Loc}}(f). \quad (6.15)$$

For the altered label space, it is actually possible to either choose a smaller label space or an extended label space. Consider both the maximal lower bound $\check{\mathcal{L}} \subset \mathcal{L}$, a smaller label space, and the minimal upper bound $\hat{\mathcal{L}} \supset \mathcal{L}$, an extended label space, that satisfy the Ishikawa construction. The maximal lower bound $\check{\mathcal{L}} = \mathcal{L} - \{L\}$ ignores the lesion label, and the energy $E(f)$, with this label space, can be directly minimised with graph-cuts. However, since our method requires that the lesions are segmented, the label space $\hat{\mathcal{L}}$, which includes additional labels, is chosen:

$$\hat{\mathcal{L}} := \{\text{CSF}, \text{GM}, \text{WM}, \text{L}, \{\text{L}, \text{SGM}\}, \text{SGM}, \text{V}, \text{F}_2\}. \quad (6.16)$$




Figure 6.8: Ishikawa construction for our six labels. On the left, the inter-layer ∞ -edges required to generate the labels on the right. In addition to the six labels two *false labels*, F_1 and F_2 appear.

As already shown in Section 4.4, the Ishikawa construction requires multiple copies of the input image to encode the different labels. Since our method models six labels, five copies of each image are needed. The number of possible label configurations is $2^5 = 32$, which can be reduced to 8 if ∞ -edges, or rather an ordering between the labels, are introduced. The left-hand side of Figure 6.8 shows the configuration used in the model presented here; note that x_0, \dots, x_4 denote a copy of voxel x in the zeroth to fourth layer, respectively. The right-hand side of Figure 6.8 lists all the possible labels that can be obtained imposing the constraints depicted on the left.

This scheme, Equation 6.16, produces two labels that were not intended to be modelled. To understand this, let consider the voxel on layer 0, *i.e.* x_0 . If x_0 has label 1, then also x_1 must have label 1 due to the ∞ -edge between the two layers. This constraint is applied to all the pair of voxels which are connected by an ∞ -edge. All the possible configurations, given the 5 layers and the ∞ -inter layer edges, are the eight ones presented in Figure 6.8 (r.h.s.).

As a result, the lesions and the SGM are not separated, leading to the *false label* $F_1 := \{L, SGM\}$. The *false label* F_2 is an artefact from the ventricle model and is embedded inside V, Figure 6.6(b).

In the following it will be explained how to construct the graph. Both E_{App} and E_{Loc} in $E_0(f)$ have to be considered, Equation (6.15). The appearance energy adds weight on the terminal links, *source* and *sink*, while the local energy adds weights on the edges between neighbouring voxels. Let D_0, D_1, \dots, D_5 denote the appearance of each tissue type, assuming that D_0 represent CSF probability, D_1 represents the GM probability, D_2 the WM, D_3 the SGM, D_4 the L and D_5 the V. Then the weights on the terminal links of the graph follow the configuration in Figure 6.9(a), where the layer order is consistent with the one in Figure 6.8 (l.h.s.). The hard constraint voxels included in S_ℓ , will be also in the graph, like every other voxel in the image, but they will have specific weights. These weights are either 0 or ∞ , thus guaranteeing that in the final segmentation their labels remain unchanged. An example of a voxel belonging to S_{GM} is shown in Figure 6.9(b).

The local energy, E_{Loc} , considers neighbourhood interactions, which are represented by edges connecting two neighbouring voxels. The smoothness of the segmentation is




Figure 6.9: Appearance energy weights for our model. In (a) the general case, while in (b) the weights for a voxel $x \in S_{GM}$. The ∞ -edges force the voxel to be segmented as GM, and determine the shown cut (black dotted line). The edges without weight take the value zero.

encoded by the weights computed with Equation (6.9). A representation of these edges can be seen in Figure 6.10(a). Additionally with E_{Loc} , a minimal distance of 1 voxel between non-neighbouring tissues is encoded: for example, between CSF and WM a layer of GM must be present. This is valid also for all other pairs of non adjacent tissues: between GM and L, as well as between L and SGM, there must be a WM region. Note that in the transition from L to V there must be a layer of WM and a layer of SGM. This minimal distance can be encoded in the graph by adding ∞ -edges as shown in Figure 6.10(b). For every voxel, the edges associated with both the appearance energy and the local energy must be created.

The minimisation of the graph so constructed converges into the global optimum. Therefore one can use this optimum $\hat{f}: \Omega_2 \rightarrow \hat{\mathcal{L}}$ of $E_0(\cdot)$ to find a feasible labelling $f: \Omega_2 \rightarrow \mathcal{L}$ with respect to $E(\cdot)$, hence a labelling that satisfies $E(f) < \infty$.

As a first step the two *false labels* are removed. Since F_2 uses the same appearance model as V, F_2 is simply replaced with V. Removing F_1 is more complicated since it is a combination of L and SGM. Therefore these two labels need to be separated. In order to disallow the transition between SGM and L, some of the WM has to appear between SGM and L. This can be efficiently performed with the α/β -swaps algorithm, Section 4.5.1, in the region $R_L \cup R_{SGM} \cup R_{F_1}$ with respect to the labels $\{WM, SGM, L\}$. For convenience, initially all the voxels belonging to F_1 are assigned to L. Next, WM/SGM-swaps and WM/L-swaps are performed until convergence is reached. This occurs when the energy between two consecutive iterations does not decrease.

This process leads to a mapping $f: \Omega_2 \rightarrow \mathcal{L}$ with $E_0(f) < \infty$, which might still violate some global geometrical constraints, resulting in $E(f) = \infty$. In this case, the minimal lesion size constraint cannot be violated, for the subsampling that results in lesion diameters larger than 1 voxel at the original resolution. This means that either




Figure 6.10: Local energy weights between two neighbouring voxels x and y , according to our model. In (a) the edges encoding the smoothness of the segmentation, while in (b) the edges encoding the minimal distance constraint between non adjacent tissues.

R_{GM} or R_{SGM} are not connected. Without loss of generality it is assumed that R_{GM} is not connected. In the next paragraph the proposed connectivity prior is explained, with a brief summary of the already proposed priors.

Connectivity priors are in general NP-hard. Vicente et al. [146] proposed a connectivity constraint for binary image segmentation and for user selected points in the image, where the selected points are iteratively connected to the segmented object. Since this method is limited to binary segmentation, is it hardly applicable to our multi-labelling framework. In addition, the manually selected points are in contrast with an automated process. In a recent approach, Stühmer et al. [132] proposed a connectivity prior which introduces a distance measure based on image intensities. This is a feasible approach provided that the distribution of intensities for the connected object is clearly distinguishable from those of labels/objects. Since the GM, SGM, and the L have similar intensity distributions this method is hardly applicable to our framework. Therefore, a distance function based on the gradient values is proposed, Figure 6.11(f). CSF voxels are assigned a distance of zero. For voxels in other brain regions, their distance increases as a function of their Euclidean distance to the CSF and of the image gradient. In order to incorporate this constraint into the graph, ∞ -edges are added to the layer which ‘controls’ the GM. Such an edge connects a given voxel x to the closest neighbour y , Figure 6.12(a).

The proposed connectivity constraint is able to correctly connect the GM tissue and recover some L labels, Figure 6.11(g). The connectivity constraint is similarly applied to the SGM, Figure 6.12(b). Note that in this case the distance function is computed from the ventricles. After encoding these two additional constraints, an initial segmentation $f: \Omega_2 \rightarrow \mathcal{L}$ is obtained that satisfies $E(f) < \infty$, Figure 6.6 (c).




Figure 6.11: FLAIR image in (a). Segmentation obtained without the connectivity constraint may lead to a disconnected GM (b) (arrow in (b)). Enforcing the connectivity prior as proposed in [146] creates a thin artefact connection path (c). The prior proposed in [132] uses an appearance based distance map (d) to enforce connectivity, but fails in our multi-labelling setting (arrow in (e)). Using the proposed gradient based distance map (f) instead results in a connected GM (g) and an improved lesion segmentation (circle in (e) and (g)).

6.5.3 Local Energy Minimisation

In Section 6.5.2 it has been shown how to obtain a feasible initialisation $f: \Omega_2 \rightarrow \mathcal{L}$, *i.e.*, a labelling that satisfies $E(f) < \infty$ at the lowest image resolution. After reaching a global optimum with respect to $\hat{\mathcal{L}}$, Equation (6.16), the *false labels* had to be removed via local optimisation, which was limited to three labels, *i.e.* WM, L, and SGM. Because changes in these labels can also potentially affect the other labels, it is necessary to improve the initial labelling in order to determine a (local) minimum for the entire label space \mathcal{L} at all image resolutions. This procedure is described in this section. To find such a minimum, all possible α/β -swaps are performed, on the whole label space \mathcal{L} . Note that even though there are six labels, only five out of 15 different α/β -swaps are necessary, due to topological restrictions, Figure 6.4: CSF / GM, GM / WM, WM / L, WM / SGM, and SGM / V.

As explained in Section 4.5.1, the α/β -swaps algorithm for binary segmentation converges to a local optimum solution. In our configuration, the five possible α/β -swaps are applied in cascade fashion. Starting from the outermost tissue, *i.e.*, the CSF, CSF / GM-swaps is applied, and as a result both R_{CSF} and R_{GM} are updated. Next, GM / WM-swaps is applied and R_{GM} and R_{WM} are thereby updated. The remaining α/β -swaps are applied following this order: WM / SGM-swaps, WM / L-swaps and SGM / V-swaps. Each α/β -swap lowers the energy, and convergence is reached when the energy value between consecutive iterations no longer decreases.

Upon convergence, the GMMs probabilities p_ℓ of Section 6.4.1 are updated in an EM-like fashion, Section 3.2.4 to a new distribution q_ℓ . The GMMs are initialised with respect to the hard constraint sets S_ℓ , Section 6.5.2 and, after convergence of the α/β -swaps, the new regions R_ℓ are used to compute the new GMM models, *i.e.* q_ℓ . In order to avoid large jumps in the intensity models, only ‘half of the gradient update’ is here used. To understand how this is computed, initially a pure Gaussian model is considered, and subsequently the extension to GMMs is explained. If pure Gaussian models are used,




Figure 6.12: Connectivity constraint for GM in (a) and SGM in (b). The ∞ -edges are between a voxel x and its closest neighbour y ; the position of the ∞ -edge corresponds to the layer in the Ishikawa construction which ‘controls’ the GM and SGM labels respectively. In (c) the ∞ -edges enforcing the minimal size constraint for the lesions in the high resolution image are shown. Note that the central voxel represents the centre of mass of a lesion.

then $p_\ell = G(\mu_\ell, \Sigma_\ell)$ and $q_\ell = G(\mu_\ell, \Sigma_\ell)$ are combined. In this case, the mean value can be computed with respect to the Bhattacharyya distance, to obtain $G(\mu, \Sigma)$:

$$\mu = \frac{\mu_p + \mu_q}{2}, \quad \text{and} \quad \Sigma = \Sigma_p \cdot U \cdot \Lambda^{\frac{1}{2}} \cdot U^T, \quad (6.17)$$

where $U \cdot \Lambda \cdot U^T$ is the diagonalisation of $\Sigma_p^{-1} \cdot \Sigma_q$. Note that, in this case, the $G(\mu, \Sigma)$ represents the model that would be used.

In the case of GMMs the models $p_\ell = G(\mu_p, \Sigma_p, \alpha_p)$ and $q_\ell = G(\mu_q, \Sigma_q, \alpha_q)$ have multiple Gaussians that need to be matched. The matching is computed by a linear assignment approach [96] using the Bhattacharyya distance. After the matching is performed the mean and covariance matrix are updated for each Gaussian with Equation 6.17. The mixing coefficients are updated with $\alpha = \frac{\alpha_p + \alpha_q}{2}$. Therefore, instead of using the q_ℓ model directly, our method uses the model resulting from the described update.

The local optimisation with respect to α/β -swaps and GMM re-estimation is iterated until convergence. Upon convergence, the resolution from Ω_2 to Ω_1 is extended and re-iterated, α/β -swaps and GMM re-estimation, until a locally optimal labelling $f|\Omega_1$ is found. For this and the highest resolutions, the minimal lesion size constraint, Equation (6.13), may be violated. In this case, for the centre of mass c_A of each connected component $A \subset R_L$ it is enforced that $f(c_A) = L$, which implies that $f(x) = L$ for all $\|x - c_A\| \leq 1$ voxel, which can be easily integrated into the graph cut optimisation [37], Figure 6.12(c). This means that for every lesion the centre of mass is calculated, and from it ∞ -edges are added, that link the centre of mass to its neighbourhood. After introducing these constraints, any small lesion will either be automatically removed or extended to include a ball of diameter 3 voxel, Figure 6.5.

Algorithm 5

-
- 1: Find initial labelling according to Section 6.5.2
 - 2: **repeat**
 - 3: Local optimisation w.r.t. α/β -swaps and GMM re-estimation
 - 4: **until** convergence
 - 5: Increase resolution if possible and **goto Line 2**
-

The final step requires that the resolution increases from Ω_1 to Ω . After convergence of α/β -swaps and GMM re-estimation, the optimal labelling $f: \Omega \rightarrow \mathcal{L}$ is found. The complete method is summarised in Algorithm 5, and the process is shown in Figure 6.6.

6.5.4 Hyper-parameter Learning

The various GMM parameters *i.e.* the mean, the covariance and the mixing coefficients, are optimised for every tissue using an Expectation Maximisation algorithm [35]. The remaining parameters are the five different λ terms that specify the local geometric model (Section 6.4.2). Since the main goal is to achieve lesion segmentation, all λ values, except $\lambda_{\text{WM,L}}$ are fixed empirically. In order to make this latter adaptive to the contrast in the image, it is modelled as being affinely dependent on the maximum magnitude:

$$\text{mg}(I) = \max_{x \in \Omega} \|\nabla I\|, \quad \text{where} \quad \nabla I(x) = \left(\frac{\partial I}{\partial x_1}, \frac{\partial I}{\partial x_2}, \frac{\partial I}{\partial x_3} \right).$$

Thus, $\lambda_{\text{WM,L}} = m \cdot \text{mg} + b$ where m and b are so-called hyper-parameters that are to be optimally chosen. The optimal values for m and b are learnt from a given training set using the parametric max-flow framework [71].

For notation simplicity, in the following let assume $\lambda = \lambda_{\text{WM,L}}$. The max-flow algorithm shows that, for a binary segmentation, given two submodular energies E_1 and E_2 , the set of optimal segmentations (s) with respect to the energy

$$E(\lambda) = \min_s E_1(s) + \lambda E_2(s) \tag{6.18}$$

is finite and can be computed efficiently. Note that the max-flow algorithm has been proposed for binary segmentations, therefore in our multi-labelling approach the R_L is considered as foreground, and the remaining tissues are background. Moreover, the energy E_1 represents our E_{App} plus E_{Loc} for all non-lesion voxels, while the second energy E_2 is our E_{Loc} of the lesion voxels. The E_{Geo} is not considered since its possible values are only 0 or ∞ , hence given a feasible segmentation the energy does not increase due to the geometric energy term.

Each of the optimal segmentations covers a certain interval $[\lambda_i, \lambda_{i+1}]$ of feasible values for λ . Since our method iteratively solves submodular problems, it is only necessary to track the subdivision of the intervals in order to obtain any possible final segmentation for any value of λ . In order to find the λ values corresponding to the subdivision points,




Figure 6.13: Graphic representation of the energy functions obtained with two different values of λ , *i.e.* λ_0 and λ_1 . The intersection point identifies a subdivision point, λ^* .

first the range of possible λ values is defined. The smaller value is usually set to zero, $\lambda_0 = 0$, whereas the highest value must correspond to a segmentation where no lesions are segmented, *e.g.* $\lambda_1 = 100$. Therefore the best $\lambda_{\text{best}} \in [\lambda_0, \lambda_1]$. Given the energy functions $E(\lambda) = E_1(s_{\lambda_0}) + \lambda E_2(s_{\lambda_0})$ for λ_0 , and $E(\lambda) = E_1(s_{\lambda_1}) + \lambda E_2(s_{\lambda_1})$ for λ_1 , a new subdivision $\lambda^* \in [\lambda_0, \lambda_1]$ can be defined as the interaction lines between the two functions, Figure 6.13. Note that the blue line, which corresponds to λ_1 , is parallel to the λ -axis since $E_2(s_{\lambda_1}) = 0$. The value of λ^* can be analytically calculated as:

$$E_1(s_{\lambda_0}) + \lambda^* E_2(s_{\lambda_0}) = E_1(s_{\lambda_1}) + \lambda^* E_2(s_{\lambda_1}) \quad (6.19)$$

re-arranging the above equation results in:

$$\lambda^* = \frac{E_1(s_{\lambda_1}) - E_1(s_{\lambda_0})}{E_2(s_{\lambda_0}) - E_2(s_{\lambda_1})} \quad (6.20)$$

Once the first λ^* is found, then the intervals $[\lambda_0, \lambda^*]$ and $[\lambda^*, \lambda_1]$ are recursively analysed, until no new λ^* is identified. The final intervals $[\lambda_i, \lambda_{i+1}]$ correspond to a convex domain A_i in the 2D domain of feasible values for m and b . For each domain the energy value of the correspondent segmentation is stored. For each patient in the training set a different separation of the domain is computed. The intersection of these different separations of the domain (one for each patient) leads to a partition that contain a finite number of regions, where in each region the final energy value is the sum of the corresponding energies in the training set segmentations, Figure 6.14.

Let A^* be the region with the lowest cumulative energy value. By choosing (m, b) as the centre of mass of A^* , the globally optimal hyper-parameters for our algorithm are obtained, Figure 6.14.




Figure 6.14: Example of m - b domain parcellation in (a), and a close-up in (b). The magenta dots represent the optimal values.

6.6 Results

Our MS lesion segmentation method has been evaluated on the MS *grand challenge clinical dataset* [133], and on an in-house database. Additionally the proposed method has been compared against several state of the art methods.

In Section 6.6.1 a description of commonly used measures for evaluating the quality of a segmentation method is provided. In addition, in this thesis a new measure called *Distance Dice* D_θ has been introduced. The goal of this new measure is to be more robust with respect to perturbations that occur within a small distance θ along the lesion segmentation boundaries.

For all the experiments the T_1 -w and the FLAIR images of each patient have been used. The coefficients of the local geometric model have been empirically set as: $\lambda_{\text{CSF,GM}} = 1$, $\lambda_{\text{GM,WM}} = 2$, $\lambda_{\text{WM,SGM}} = 10$ and $\lambda_{\text{SGM,V}} = 1$. The subsampling factors for the MICCAI database have been set to $s_{x_1}^{(1)} = s_{x_2}^{(1)} = s_{x_3}^{(1)} = s_{x_1}^{(2)} = s_{x_2}^{(2)} = s_{x_3}^{(2)} = 2$, while for the in-house database have been set to $s_{x_3}^{(1)} = 1$ and $s_{x_1}^{(1)} = s_{x_2}^{(1)} = s_{x_1}^{(2)} = s_{x_2}^{(2)} = s_{x_3}^{(2)} = 2$. Note that for I_1 no downsampling was performed in the through-plane direction x_3 since the in-house data have a lower resolution in this direction (Section 6.6.3).

6.6.1 Measures

In order to evaluate the quality of a segmentation it is common to compare the segmentation against a manual annotation of the lesions. Usually the manual annotation consists in a binary mask defined by an expert physician, in which lesion voxels have a positive value (congenitally set to unity). It is also desirable to have multiple manual annotations which can be beneficial to identify the lesions on which physicians agree and also to weigh the likelihood of the lesions according to the level of agreement.

There are two classes of measures to compare a given manual annotation, also called Ground Truth (GT), with a lesion segmentation. The first class regards the whole lesion as unity, therefore even if the manual annotation and the segmentation differ by few

voxels the score is still 100%. This class of measures give more importance to the actual detection of the lesion rather than to the precision of the segmentation. On the other hand, the second class of measures considers the voxel as the unity for comparison, so the focus is on the precision of the lesion segmentation.

In the following the lesion segmentation is identified as S , and the manual annotation as GT , Figure 6.15(a). The measures defined on the lesion level are the True Positive Rate (TPR), and the False Positive Rate (FPR) defined as:

$$\text{TPR} = \frac{|S \cap GT|}{|GT|} \quad \text{FPR} = \frac{|S| - |S \cap GT|}{|S|} = 1 - \text{TPR} \frac{|GT|}{|S|}$$

where $|S|$ and $|GT|$ are the number of lesions in the segmentation and in the ground truth, respectively; $|S \cap GT|$ denotes the number of lesions in GT which overlap with at least one lesion in S . The best value of the TPR is 100%, so when all the lesions in GT have an overlap with at least one lesion in S . The best value of the FPR is 0, hence when there are no lesions in S without an overlap to GT .

At the voxel level various measures have been defined. The Volume Difference VD considers the actual volume of the lesions is S , *i.e.* V_S , and compares it with the volume of the lesions in GT , V_{GT} :

$$\text{VD} = \frac{|V_S - V_{GT}|}{V_{GT}}$$

The volume can be computed by counting the voxels in the segmentation. The best value is 0, thus when the total volume of S is equal to the total volume in GT .

The symmetric Surface Distance (SD) considers the distance between the surface voxels of every lesion in both S and GT [133]:

$$\text{SD} = \frac{\sum_{x \in \partial GT} d(x, \partial S) + \sum_{y \in \partial S} d(\partial GT, y)}{|\partial S| + |\partial GT|}$$

where ∂S and ∂GT represent the surface voxels of S and GT respectively, and d represents the Euclidean distance. Surface voxels are defined as those which have at least one neighbour which is not a lesion. The best value is 0.

The additional measures defined on the voxel level require a precise identification of missed, oversegmented and properly segmented voxels. Let define the (number of) True Positives (TP) as the number of correctly identified voxels, *i.e.*, voxels in S which coincide with the corresponding voxel in GT . Define the False Positives (FP) as the number of oversegmented voxels in S , hence those voxels which do not coincide with their corresponding voxel in GT . Finally, define the False Negatives (FN) as the number of lesion voxels missed out in the segmentation, *i.e.*, the number of voxels in GT which do




Figure 6.15: In (a) a representation of the ground truth GT and segmentation S , and in (b) the classical FN , FP and TP . In (c) the proposed distance dice D_θ , where the introduced tolerance zone is shown, (dashed lines) in which the TP can increase in number.

not overlap with a voxel in S , Figure 6.15(b). Based on these terms it is possible to define the Recall factor R , also called $TPRv$ on the voxel level:

$$R = \frac{TP}{TP + FN} = \frac{TP}{|GT|}.$$

The best value for R is 100%. This measure is sensitive to the number of FN, hence to the percentage of missed voxels. The Precision P , already presented in Section 5.4 is defined as:

$$P = \frac{TP}{TP + FP} = \frac{TP}{|S|}$$

the best value being 100%. Note that $P = 1 - FPRv$, where $FPRv$ is the FPR at the voxel level. A high precision P means that the number of false positives is small in comparison to the number of true positives. This, however, does not necessarily imply that the region S will closely match GT (consider the case where $S \subset GT$). In general, the goodness of the segmentation, as compared with the ground truth, requires that both R and P be high. This is captured by the Dice coefficient D , which combines P and R and is defined as:

$$D = 2 \frac{P \cdot R}{P + R} = \frac{2 \cdot TP}{FP + FN + 2 \cdot TP}$$

and the best value is 100%.

Since GT is often inaccurate at the lesion boundary, in this thesis a new distance-based Dice coefficient D_θ has been introduced, where θ defines a tolerance zone both inwards and outwards of the GT border. The voxels included in the tolerance zone are considered neither as FN nor as FP , but instead as TP , resulting in $D_\theta \geq D$, Figure 6.15. D_θ is similar to the *slack borders* measure that was introduced in [145] to evaluate lung segmentations. This measure ignores the information around the boundary for TP as well as for FN and FP. Therefore, the *slack borders* measure may be smaller than D . On the other hand, the distance-based Dice coefficient D_θ always provides a higher or equal measure than the classical Dice coefficient.




Figure 6.16: Example of results produced by our method for T_1 -w and FLAIR images from the MICCAI test dataset. The results are shown as orange overlay on the FLAIR image. The first three columns show data from patient CHB15, the other columns from patient UNC03.

6.6.2 MICCAI Dataset

The MS lesion challenge proposed in 2008 [133] consists of 20 patient data in the training set and 22 in the test set. The patients were scanned at the University of North Carolina (UNC) and at the Children’s Hospital, Boston (CHB). While the acquisition protocol consisted of T_1 -w, T_2 -w and FLAIR images, in the method proposed in this thesis only T_1 -w and FLAIR images have been used. All images have an isotropic resolution of $0.5 \times 0.5 \times 0.5 \text{ mm}^3$. The training set used here consists of 10 CHB and 10 UNC patients. One and two manual annotations were respectively available for CHB and UNC patient data. On the other hand, the test set was evaluated online using two manual annotations. The training set has been used to estimate the hyper-parameters, m and b , for the $\lambda_{L,WM}$ and the test set to evaluate the performance of our method.

These data have been evaluated with respect to VD, SD, TPR and FPR. The website of the challenge also provides a score based on average values of all four measures. Our method has been compared with Lesion-TOADS [118], LST [116], with the method proposed by Geremia et al. [53], the MOPS method [137], and the RMNMS method [55]. The results are shown in Table 6.1, where boldface values represent the best value for each metric and each dataset. All the methods produce similar results, however our method gives the best performance with respect to the FPR, *i.e.*, provides the most conservative segmentation with respect to all other state of the art methods. In Figure 6.16 a visual result of the proposed method for two patients from the test set is provided: CHB15 is presented in the first three columns and UNC03 in the last three.

6.6.3 In-House Database

The in-house database contains of data from 20 MS patients (mean age 37.2; age range 28 – 55 years). It is divided up into two datasets: the first contains data from 5 patient data, the second from 15 patient data.

Both datasets were acquired with a 3T Siemens Trio MR System. The acquisition protocol for both consists of a 3D MPRAGE T_1 -w image and a 3D FLAIR image. The first dataset has a matrix size of $480 \times 512 \times 160$ with a resolution of $0.5 \times 0.5 \times 1.2 \text{ mm}^3$. The MPRAGE parameters are: $TR/TE/FA = 1390 \text{ ms} / 2.15\text{ms} / 9^\circ$; and FLAIR

Table 6.1: Comparative results for the MICCAI MS grand challenge test dataset, and for the two manual annotations provided (rat.) [133]. The SD is expressed in millimetres while all the other measures are expressed in percentage.

Method	Hospital	VD		SD		TPR		FPR		Score
		1 rat.	2 rat.	1 rat.	2 rat.	1 rat.	2 rat.	1 rat.	2 rat.	
Lesion-TOADS [118]	UNC	45.6	128.2	8.3	12.7	44.1	61.5	68.1	72.8	79.96
	CHB	75.3	57.6	6.5	5.3	52.2	52.2	79.2	69.1	
LST [116]	UNC	63.3	110.5	9.5	13.5	48.5	59.7	73.6	73.4	80.00
	CHB	50.4	69.3	5.9	4.4	50.0	55.9	78.8	65.8	
Geremia et al [53]	UNC	42.1	56.7	8.0	10.7	43.5	66.8	74.7	82.8	82.07
	CHB	47.6	49.7	4.2	2.0	56.1	54.0	78.0	64.3	
MOPS [137]	UNC	29.4	54.1	9.9	11.3	36.1	62.3	46.2	54.7	84.46
	CHB	43.3	53.0	5.2	6.4	45.8	45.0	42.7	38.9	
RMNMS [55]	UNC	51.2	56.6	5.7	7.3	44.9	63.1	39.1	53.2	86.11
	CHB	43.2	47.9	5.4	4.4	48.4	46.0	46.3	34.7	
Proposed method	UNC	77.7	215.1	11.5	13.1	30.8	49.7	37.8	46.7	83.93
	CHB	45.0	55.7	4.5	4.8	41.6	39.1	32.8	21.2	

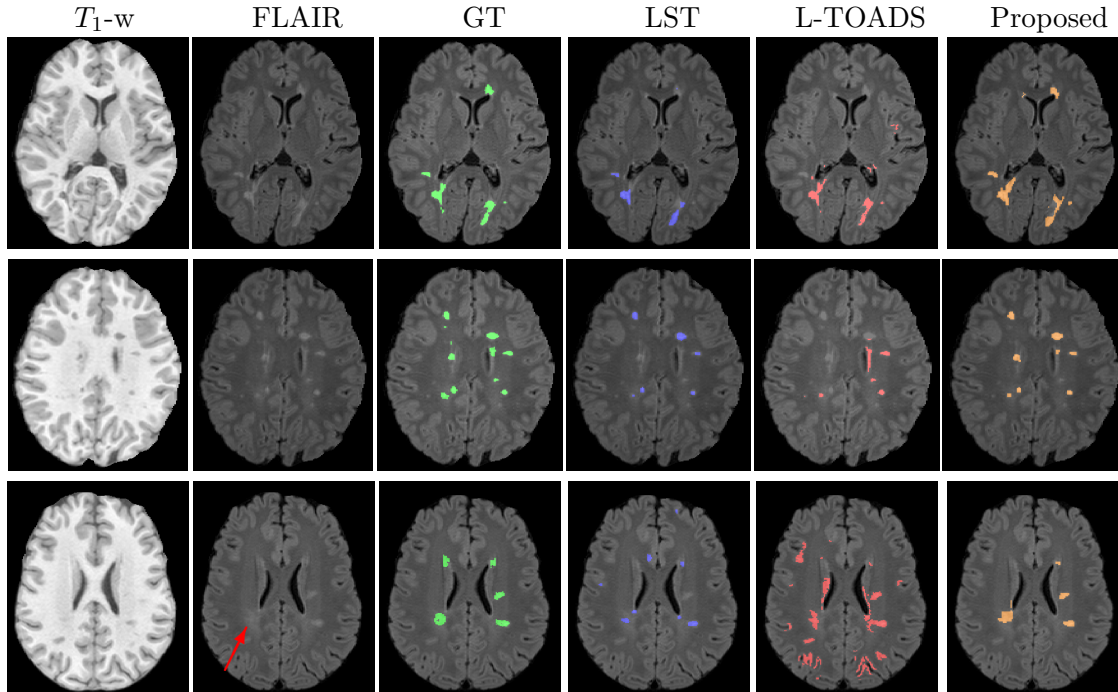


Figure 6.17: Results for the in-house dataset, consisting of T_1 -w and FLAIR images. The results are overlaid on the FLAIR image. The ground truth GT (green) is shown together with the results obtained with LST (blue), Lesion-TOADS (red) and the proposed method (orange). Each row presents a different patient and the red arrow in the third row points to a ‘dirty’-appearing-WM region.

parameters are $TR/TE/FA = 5000\text{ms} / 388\text{ms} / 120^\circ$. The second dataset has a matrix size of $512 \times 512 \times 160$ with a resolution of $0.49 \times 0.49 \times 1 \text{ mm}^3$. The MPRAGE parameters are $TR/TE/FA = 2300\text{ms} / 2.98\text{ms} / 12^\circ$; the FLAIR parameters are $TR/TE/FA = 5000\text{ms} / 394\text{ms} / 120^\circ$.

Two expert neuroradiologists, one for each dataset, provided a manual GT segmentation on the FLAIR images. The first dataset has been used for the hyper-parameter learning phase and the second one for the testing phase. For each image, the lesion load has been computed as the sum of the total volume lesion volume in GT : for the first dataset, total lesion volume is $1.01 - 9.8 \text{ cm}^3$, for the second $0.9 - 24 \text{ cm}^3$.

The method presented here has been compared against Lesion-TOADS [118] and LST [116]. For both external methods the parameters have been fine-tuned in order to obtain the best results, Figure 6.17. In Table 6.2 the results for each of the presented measures are shown, with boldface values representing the best value for each metric. It observed that the method proposed in this thesis performs best in terms of VD, TPR, R , D and $D_{\theta=0.5\text{mm}}$. The LST method performs best for the SD and P measures, whereas Lesion-TOADS gives the best performance for the FPR. In general, the D_θ scores improve with respect to D for all methods.

Table 6.2: Comparative results of our method with two state of the art methods on the in-house dataset of 15 patients. The SD is expressed in millimetres while all the other measures are expressed in percentage.

Patient	Method	VD	SD	TPR	FNR	P	R	D	$D_{0.5}$
Pat 01	L-TOADS	48	6	65	10	30	45	36	44
	LST	53	4	75	70	60	29	39	68
	Proposed	37	5	100	80	59	37	45	69
Pat 02	L-TOADS	50	7	60	42	30	43	36	42
	LST	52	4	78	59	72	34	46	71
	Proposed	50	4	85	70	72	35	47	74
Pat 03	L-TOADS	63	5	76	42	77	39	51	70
	LST	89	4	90	72	88	10	18	41
	Proposed	9	3	56	72	48	35	40	44
Pat 04	L-TOADS	40	6	60	36	10	43	16	20
	LST	48	2	54	62	74	38	51	73
	Proposed	7	2	48	53	49	44	46	72
Pat 05	L-TOADS	170	9	10	50	10	44	16	35
	LST	67	3	44	77	84	28	42	68
	Proposed	44	2	67	77	81	45	58	83
Pat 06	L-TOADS	106	6	13	74	20	15	17	20
	LST	86	5	56	86	76	17	28	63
	Proposed	78	8	85	83	86	19	31	67
Pat 07	L-TOADS	86	6	20	50	5	47	9	30
	LST	67	3	64	60	91	41	57	84
	Proposed	35	4	80	75	78	51	61	83
Pat 08	L-TOADS	92	10	50	33	10	43	17	30
	LST	64	2	32	11	50	29	37	55
	Proposed	24	3	66	18	48	59	53	66
Pat 09	L-TOADS	31	7	50	40	10	31	15	30
	LST	25	1	75	21	64	71	68	84
	Proposed	26	1	66	20	77	57	65	85
Pat 10	L-TOADS	53	6	65	55	79	37	50	70
	LST	63	2	69	69	87	43	58	79
	Proposed	60	1	61	60	47	75	58	78
Pat 11	L-TOADS	73	3	60	74	66	41	51	71
	LST	74	2	85	60	87	31	46	76
	Proposed	24	1	72	61	75	57	64	85
Pat 12	L-TOADS	94	5	10	46	10	41	16	40
	LST	45	1	73	30	71	56	63	81
	Proposed	36	2	76	30	52	60	60	73
Pat 13	L-TOADS	49	2	65	60	67	35	46	74
	LST	52	1	55	34	77	53	63	80
	Proposed	23	2	54	30	58	71	64	78
Pat 14	L-TOADS	25	5	25	58	30	30	30	43
	LST	66	1	43	63	85	42	56	80
	Proposed	39	2	65	50	80	51	62	87
Pat 15	L-TOADS	77	5	51	54	91	17	29	57
	LST	57	3	33	70	32	23	27	50
	Proposed	55	2	75	64	27	50	35	55
Mean	L-TOADS	70	6	45	48	36	37	29	45
	LST	61	2	62	56	73	36	47	70
	Proposed	36	3	70	56	62	50	52	73




Figure 6.18: Axial slices with FLAIR at the level of the cerebral peduncle. The intensity difference between the corticospinal tract (red arrow) and two brainstem lesions is seen on the left. The proposed method segments only the lesions.

It is noted that certain brain regions, like the so-called ‘dirty’-appearing white matter and the corticospinal tract, exhibit a diffuse hyper-intense appearance in MRI images. This complicates the correct delineation of these areas, parts of which can wrongly be classified as lesions. The ‘dirty’-appearing white matter is usually considered as a lesion, even if the precise boundaries are hard to detect precisely. An example is shown in Figure 6.17 third row. In Figure 6.18 the lesions close to the corticospinal tract have been correctly segmented, while the tract itself has not been detected.

6.6.4 Robustness against Noise

The robustness of the proposed method against increasing levels of noise has been tested. Rician noise has been artificially added to both FLAIR and the T_1 -w images on three representative patients with low, moderate and high lesion loads, respectively. The percentage of the added noise is based on the brightest tissue, *i.e.*, GM for the FLAIR and WM for the T_1 -w. The noise levels tested are 1%, 3%, 5%, 7%, and 9% over the mean value of the brightest tissue, and additionally a comparison against Lesion-TOADS and LST has been performed. For all three methods, the optimal parameters estimated for the selected patient data (without artificially added noise) have been used. In Figure 6.19 the Dice coefficient D and $D_{0.5}$ versus the noise level are plotted for LST, Lesion-TOADS and for the proposed method.

For noise levels between 1%–5%, the performance of the proposed method is minimally affected, whereas for a noise level of 9% the coefficient D decreases by 20% in the patient data with low and moderate lesion loads. Note that both the D and $D_{0.5}$ values for 0% of noise are the same as for 1% for all the methods. For a high lesion load, all three methods were barely sensitive to the noise level. Figure 6.20 presents comparative image segmentation results, at all noise levels, for the patient data with a moderate lesion load discussed in this section.




Figure 6.19: D (solid lines) and $D_{0.5}$ (dashed lines) values for increasing levels of noise. Lesion-TOADS (red), LST (blue), and the proposed method (orange) for three representative patients with high (a), moderate (b), and low (c) lesion loads.

6.7 Summary and Discussion

A new method for 3D lesion segmentation has been presented which does not require the use of an atlas. The proposed method formulates the segmentation as an energy minimisation problem consisting of three different energies, *i.e.* appearance energy, local geometric model, and global geometric model. The appearance energy encodes the image intensity for each tissue type, as modelled by tissue-specific GMM models. The local energy encodes a geometric model of the brain which fixes the relative position of each tissue. Compared to an atlas-based approach, this model is more flexible and easier to adapt to varying brain shapes. Therefore, the artefacts that may appear during an atlas-based registration are avoided. The global energy, on the other hand, represents additional constraints, specifically, the connectivity of the GM and SGM and a constraint on minimum lesion size. Both these constraints aid in the process of lesion segmentation by reducing, through the connectivity constraint E_{CC} that some lesion voxels get erroneously assigned to GM or to SGM, and that noisy hyper-intense voxels get segmented as lesion.

One of the main contributions of this thesis is the proposed connectivity constraint, which uses a distance function based on the gradient of image intensities. Similar to other methods the connectivity prior is imposed as a hard constraint. However, whereas the formulation of Stühmer et al. [132] is applicable to binary segmentation, the more general formulation proposed here is designed for multi-labelling segmentation. In Section 6.4.3 it has been shown that the method of Stühmer [132] fails to enforce a connectivity prior in our multi-labelling setting, whereas our method properly connects the GM, Figure 6.11.

The second global constraint introduced in this work is the minimum lesion size. For the high resolution step, this is fixed to a diameter of 3 voxels. As shown in Figure 6.5, owing to the effect of this constraint, the resulting lesion segmentation becomes less noisy, however small lesions may be grown to the predetermined size dictated by the constraint or they disappear.

The overall energy in Equation (6.5) is minimised following a coarse-to-fine approach. The input images have been down-sampled twice, resulting in I_2 and I_1 . Following the Ishikawa construction [64] it has been possible to globally optimise the energy $E_0(f)$,

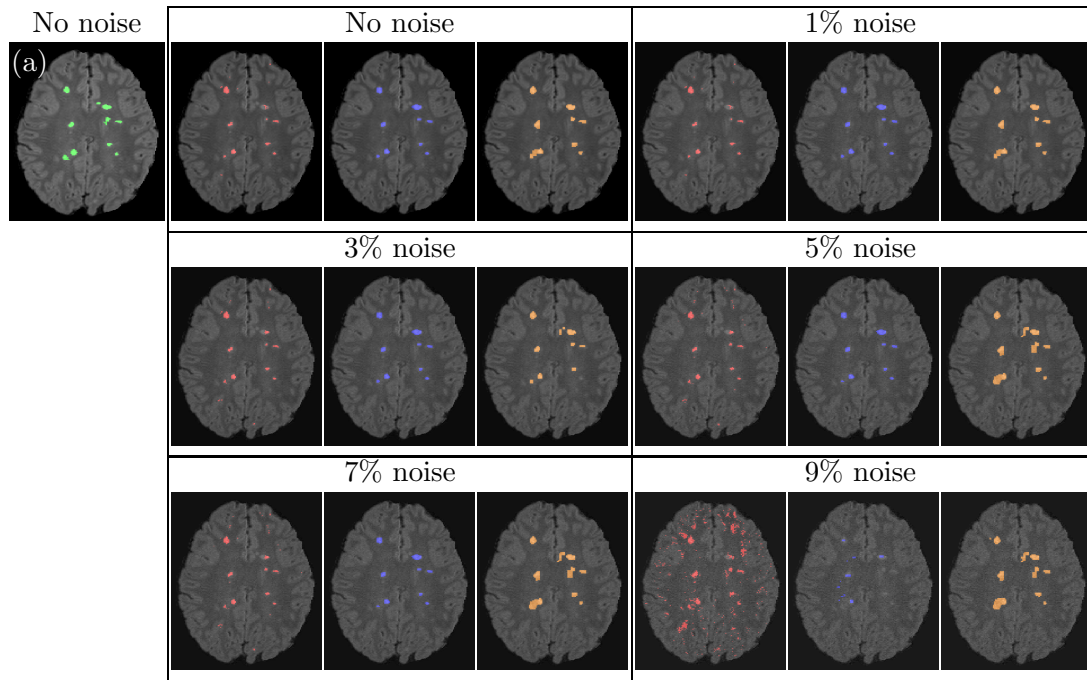


Figure 6.20: Results of the segmentation in the presence of noise. In (a) the original FLAIR image with GT. Then the results with increasing levels of noise. The comparison is made with respect to Lesion-TOADS (red), LST (blue), and the proposed method (orange).

consisting of the appearance and local geometry energies, in the image I_2 with lowest resolution, while ignoring the global energy E_{Geo} . Delong and Boykov [37] showed that combining an inclusion constraint with an exclusion constraint, for example in our model WM includes both L and SGM which are excluding themselves, results in an expression for the energy that is no longer submodular. Hence, the exclusion constraint between L and SGM could not be modelled in our framework. As a result of the failure to model that exclusion constraint, the minimisation of $E_0(f)$ generated two *false labels*, F_1 and F_2 , Section 6.5.2. Using the α/β -swap algorithm [16] it has been possible to repair all the erroneously assigned labels and to find a feasible solution for the overall energy $E(f)$. Then, the resolution is increased to I_1 and a local optimum is found alternating between all possible α/β -swap and GMM re-estimations. A similar procedure has been applied at the original resolution and a final labelling for the original images was found.

Our method is based on few parameters, namely the λ_{l_1, l_2} values for each neighbouring tissue. These values are empirically estimated and kept constant for all our experiments with the exception of $\lambda_{WM, L}$. This term has been modelled to be affinely dependent on the maximum amplitude of the image gradient, in order to make it adaptive to varying image contrasts, Section 6.5.4. In the hyper-parameter learning step, for which global optimality is guaranteed, the affine parameters m and b have been estimated that give

the best result with respect to the training set. These values of m and b have been used to find the best $\lambda_{\text{WM,L}}$ for each individual item of the patient data in the test set.

The proposed method uses FLAIR and T_1 -w images: T_1 -w images provide comparatively more information about the contrast between different tissue types, while FLAIR images are more sensitive to lesion intensities. Additional image contrasts can be easily incorporated into our framework. This only affects the computation of the GMMs, which can nevertheless be efficiently computed for multiple image contrasts. On the other hand, the size of the graphs used is virtually independent of the number of different contrasts available. The run time of the proposed method is about 3 hours for the in-house database and about 7 hours for the MICCAI data.

A limitation of the proposed method is that the geometric model is not entirely anatomically correct, since the ventricles may also be adjacent to the WM and the lesions. The reason of this modelling choice is that there are signals from the ventricle that are not correctly suppressed in MR images. These voxels have a high intensity value in the FLAIR image and can therefore be represented by the SGM-model. In practice, our model may create a thin artefactual layer of subcortical GM around the ventricles. While the resulting labelling might not be anatomically correct, it only slightly affects the lesion segmentation. From Figure 6.17 (first row) and Figure 6.16 it is seen that the lesions around the ventricles have been properly segmented.

The proposed method is biased towards detection and segmentation of lesions appearing in the WM, however lesions which are appearing in the subcortical GM can still be segmented. In our database no lesion in the main deep GM structures (basal ganglia, thalami) has been found. However, in order to segment a lesion in this area, the method must create an artefactual rim of WM around such lesions to satisfy the geometric constraints. The proposed method may not entirely segment the dirty-appearing WM completely, which in FLAIR images typically appears as a smooth hyper-intense region in the periventricular WM, Figure 6.17 (second row). Because the boundaries of these regions are difficult to detect, they may be partially segmented as healthy WM. This limitation might be overcome by lowering the value of $\lambda_{\text{L,WM}}$. Since the training set did not contain extended regions of dirty-appearing-WM, the learning process could not take this effect into account.

The lesion load of the in-house database ($0.9 \text{ cm}^3 - 24 \text{ cm}^3$) showed that the proposed method performs robustly in the presence of mild, moderate and severe lesion loads. The robustness of the method against noise has been tested for three representative patients with low, moderate and high lesion loads, respectively. In moderate and low lesion loads the performances of the proposed method are slightly affected at low noise levels, while at a noise level of 9% the Dice coefficient D has a drop of 20%. In this case most of the lesions are still detected, but the number of FP increases, indicating an over-segmentation, Figure 6.20. A possible explanation of this behaviour is that the λ values are estimated on data without artificially added noise, therefore they are not optimal for noisy images.

To increase the robustness of the proposed method with respect to small changes along the borders of the segmentations, a new metric D_θ has been introduced (Section 6.6).

The method has been tested on the MS grand challenge [133] and on an in-house dataset. The in-house datasets have a different image quality and contrast as compared to the challenge data, which may be due to more modern hardware and optimised acquisition protocols. Consequently, a better performance of the proposed method, of LST [123] and of Lesion-TOADS [118] for the in-house database has been observed. The results on the MS grand challenge, Table 6.1, show that our method is competitive with state of the art methods and outperforms all methods tested with respect to the FPR. On the other hand, the number of TPR is lower compared to the other methods, which is in compliance with the low values of the FPR as usually a low score in one of these measures corresponds to a low score in the other. Note that for the proposed method only two image modalities have been used, *i.e.* T_1 -w and FLAIR, while the other methods compute the lesions segmentation using also the T_2 -w image. An additional image contrast may help to improve performances. The results for the in-house database, Table 6.2, show that the propose method produces the best results for most metrics and that in general the D_θ scores improved with respect to D for all methods.

Chapter 7

Quantification of Tumour Vessel Abnormality

Glioblastomas are one of the most frequent, and aggressive, kinds of brain tumour. In order to be able to maintain its aggressive proliferation pattern, glioblastomas initiate the formation of new blood vessels (neo-angiogenesis). Neo-angiogenesis is characterised by the expression of molecular growth factors such as vascular endothelial growth factor (VEGF), which cause the rapid formation of arterial blood vessels in and around the tumour. These newly formed vessels are often characterised by an increased number and size, an irregular shape, immature vessel organisation, and, in particular, by a high vascular permeability. Many anti-angiogenic therapeutic approaches attempt to inhibit the formation and maturation of these blood vessels, and thus to disconnect the growing tumour from its vascular supply of nutrients. Unfortunately, the therapeutic success of anti-angiogenic therapies depends on many factors, and thus, an early detection and quantification of changes in the vascular architecture before and under therapy would be highly desirable to improve and stratify anti-angiogenic therapies. An example of an axial slice of a glioblastoma tumour is shown in Figure 7.1.

In this chapter we present our method to quantify vessel abnormality in glioblastoma tumours, which has been published in [128, 129, 130] and is here reported in an extended format. We define two different measures that, based on the local vessel orientations, are able to quantify the extent of vascular abnormality.




Figure 7.1: An axial slice of a typical glioblastoma tumour. In (a) the T_2 -weighted (T_2 -w) image, in (b) the TOF and in (c) the maximum intensity projection calculated for TOF image. The red contour identifies the tumour.

7.1 Related Work

One of the first problems encountered when dealing with these aggressive tumours is to create an image in which the vessels are visible. Typically MR images are used, contrast-enhanced MR angiography (CE-MRA) being currently the most widely used imaging method to assess vascular anatomy. In CE-MRA, a contrast agent is used to increase the signal of the blood flow, yielding an excellent vessel-to-background contrast. Unfortunately, in neo-angiogenic blood vessels CE-MRA often cannot be used, as the high vessel permeability leads to a rapid extravasation of the administered contrast agent so that the required concentration gradient between vessel and surrounding tissue cannot be achieved. Time-of-flight (TOF) MR angiography is an alternative to CE-MRA, which uses the inherent signal difference between the inflowing unsaturated arterial blood and the reduced signal intensity of the surrounding tissue during steady state imaging. In the past, several studies have attempted to demonstrate that it is possible to use TOF-MRA for the detection of tumour vessels in glioblastomas; however, this technique is still not routinely used in clinical glioblastoma MR protocols. Recently, Bock et al. [12] have shown that TOF-MRA can delineate tumour blood vessels at a field strength of 7 Tesla, where an isotropic spatial resolution of up to $300\mu m$ could be achieved. Therefore in our method TOF-MRA data have been used to localise the tumour vessels.

A major limitation of all vascular imaging studies of malignant tumours is the lack of adequate quantification methods to assess the amount of neo-angiogenic vasculature, and its possible changes under therapy. The methods proposed in the literature for the quantification of vessel abnormality follow some common steps. Initially a vessel segmentation is performed, subsequently certain discriminative features are extracted both from the tumour vessels and from selected normal appearing vessels. The comparison with healthy or normal appearing vessels can be performed based on the same patients or by comparing with healthy volunteers. In the first case the healthy vessels can be taken from the contralateral hemisphere where the tumour is not present, and a normal vasculature is expected. In case the vessel comparison is made with healthy volunteers then usually the vessels which are in the same location as the tumour are taken. This procedure, however, is generally undesirable since it requires a registration step between the healthy volunteers' images and the patient data, which can be complicated due to the high brain deformation typical of glioblastoma patients.

Bullitt et al. [24] proposed to discriminate between healthy and tumour vessels using several features such as the distance metric, inflection count metric, and the tortuosity Sum Of Angular Moments (SOAM). The distance metric gives the ratio between the path length of a curve or vessel skeleton and the Euclidian distance between the start and end points of the skeleton. The second measure is defined as the number of inflection points of a curve times the distance metric. Therefore it is more sensitive to sinuous curves. The tortuosity SOAM is calculated from in- and through-plane, or torsion, angles between three points along the vascular skeleton; it was reported to be the most discriminative measure for vessel abnormality quantification [23].

The approach proposed by Radbruch et al. [111], first enhances the brain vessels in TOF images with the Frangi filter [47], and subsequently a manual threshold is selected to remove noisy voxels and to determine the final vessel segmentation mask. As second step, the authors extract several features such as vessels length, surface and volume of the vessels, diameter, and number of branches. The latter was identified as the most discriminative abnormality measure.

In a recent review [39] several morphometric abnormality measures are compared. The most promising are, for example, vessels density, vasculature area, number of ramifications, and fractal dimension. The latter is presented as the most discriminative, even though tumour vasculature may not be successfully represented a fractal due to its complex morphological structure [6].

In the following section our novel measures, based on local vessel direction, are presented. These measures aim at differentiating normal mature blood vessels in the brain from neo-angiogenic arterial vessels in gliomas. Based on TOF-MRA data sets acquired in glioma patients, the method is tested, Section 7.4, by comparing vessel structures at the perimeter of the tumour against normal appearing vessels.

7.2 Method

The quantification of vascular abnormality in TOF-MRA data involves as a first step the vessels segmentation, followed by the identification of the vessel directions. The abnormality measures proposed in this thesis are based on the local (in)coherency of the vessels directions, namely the normalised dot product D and the two spherical angles Φ and Θ .

In order to quantify vessel abnormality for each patient both a TOF-MRA and a T_2 -weighted (T_2 -w) image are used. The TOF-MRA image is used to visualise the vasculature and the T_2 -w image provides a very good contrast between the healthy tissues and the tumour region, Figure 7.2(a)-(b). In the following the TOF-MRA is denoted as $I_{TOF} : \Omega \rightarrow \mathbb{R}$, where Ω is the image domain and $x \in \Omega$ where $x = (x_1, x_2, x_3)$ is a voxel in the image, and the T_2 -w image is denoted as I_{T_2} .

7.2.1 Pre-processing and Vessel Segmentation

As a first step a brain extraction on the I_{T_2} is performed with the BET tool [121] to eliminate vessels in the non-brain regions. This tool also provides a binary brain mask which has positive values in the regions corresponding to brain tissues. The I_{T_2} is subsequently used to manually delineate the tumour region, which results in a binary mask, *i.e.* *tumour* (Figure 7.2(a)). For local comparison, a normal-appearing region, *i.e.* *nac*, is also manually defined on the contralateral side of the tumour Figure 7.2(a). In addition a region is defined over the total brain volume without the *tumour* region, *i.e.* *natotal*. This last region is computed from the brain mask provided with the BET tool




Figure 7.2: Axial slices of a glioblastoma patient. In (a) the I_{T2} , in (b) I_{TOF} , in (c) the maximum intensity projection calculated for I_{TOF} , and in (d) the three regions *tumour* (red), *nac* (blue), and *natotal* (blue + white).

where the *tumour* region is subtracted. The I_{T2} is then rigidly registered to I_{TOF} with the default parameters of the ITK library [62], and the same registration is also applied to the *tumour*, *nac*, and *natotal* regions, Figure 7.2(b)-(d).

To quantify the abnormality of the tumour vessels, it is essential first to identify the location of the vessels. Therefore a segmentation of the vasculature in the I_{TOF} is performed with the Efficient Monte-Carlo Image-analysis for the Location Of Vascular Entity (EMILOVE) method [120]. EMILOVE is a probabilistic method based on simulated annealing. The vessels are modelled locally with particles which are connected to create tubular structures. The segmentation is formulated as a minimisation problem that consist of an external and an internal energy:

$$E\{I, \mathcal{M}\} = E_{\text{ext}}\{I, \mathcal{M}\} + E_{\text{int}}\{\mathcal{M}\} \quad (7.1)$$

where I is the input image, and $\mathcal{M} = (\mathcal{P}, \mathcal{E})$ is the used model to represent the segmentation composed by a set of particles \mathcal{P} and a set of edges \mathcal{E} . The external energy term depends on the location of the particles, their scale and their orientation:

$$E_{\text{ext}}\{I, \mathcal{M}\} = \sum_{v \in \mathcal{P}} C_{\mathcal{P}}\{I\}(v) \quad (7.2)$$

with

$$C_{\mathcal{P}}\{I\}(v) = \underbrace{-\alpha_1 V\{I\}(v)}_{\text{vesselness}} + \underbrace{\alpha_2 s_v^2}_{\text{particle costs}} + \underbrace{\alpha_3 G_V\{I\}(v)}_{\text{gradient length in particle direction}} + \underbrace{\alpha_4 G_{V^\perp}\{I\}(v)}_{\text{gradient length perpendicular to particle direction}} + c. \quad (7.3)$$

The model parameter α_1 controls the location of the particles in background regions or in vessels with irregular shape, and is set to $\alpha_1 = 2$. The particle scale parameter is set to $\alpha_2 = 4$. The α_3 and α_4 parameters penalise particles with an orientation which is inconsistent with the vessel direction or diverging from vessel centrelines; they are set to




Figure 7.3: Reference I_{TOF} slice of a glioblastoma patient in (a). Vessel segmentation at an intermediate temperature in (b) and at the final temperature in (c).

100 and 20, respectively. The particle likeliness c controls the total number of particles and is set to $c = 40$.

The internal energy enforces topological constraints on the particles and controls the particles interactions:

$$E_{\text{int}}\{\mathcal{M}\} = C_{\text{model}}(\mathcal{M}) + \sum_{u \in \mathcal{E}} C_{\text{edge}}^L(u) \quad (7.4)$$

with

$$C_{\text{edge}}^L(u) = \beta_1 C_{\text{con}}(u) + \beta_2 C_{\text{scale}}(u) + L \quad (7.5)$$

where L controls the reward for edges. The C_{model} ensures absence of intersections, loops, and direct connection between parallel vessels. The particle interaction C_{edge}^L controls the length of the edges between particles (β_1) and the difference in orientation and scale of neighbouring particles (β_2); these parameters have been set to $\beta_1 = 2$ and $\beta_2 = 1$. For more details the reader is referred to [120].

EMILOVE is an iterative method where the segmentation is updated by adding a new vessel or modifying the current ones until an equilibrium state is reached. Then, the temperature T (a surrogate parameter for the energy) is decreased and a new equilibrium state is reached. The method converges when T cannot be lowered further.

In this thesis, the EMILOVE algorithm has been applied on I_{TOF} resulting in a vascular skeleton and a binary vessels mask. The surface reconstruction of the binary mask is shown for an intermediate temperature level in Figure 7.3(b) and the final temperature in Figure 7.3(c). To find the optimal segmentation, the temperature level is automatically decreased until no change in segmentation is detectable; this resulted in $T < 0.01$ for all the analysed patient data.

Following the vessel segmentation the Frangi filter [47] is applied to I_{TOF} data at the locations defined by the vessels binary mask. The filter computes voxel-wise a Hessian matrix, using different scale values s to cover all the possible vessel sizes. This parameter s is the standard deviation of the Gaussian used to approximate the second order derivatives. In our method the parameter s has been set to $s = \{1, 2, \dots, 5\}$.




Figure 7.4: Reference I_{TOF} slice of a glioblastoma patient. Vessels orientation in *tumour* (red box) and in *nac* (blue box).

The filter then computes the eigenvectors and eigenvalues, *i.e.* $\lambda_1, \lambda_2, \lambda_3$, of the Hessian matrix. Assuming that $\lambda_1 < \lambda_2 < \lambda_3$, then a voxel that belong to a vessel would have a small value for λ_1 (ideally zero), and a large value for the other two eigenvalues. The eigenvector that corresponds to λ_1 points in the direction of minimum intensity variation, therefore along the vessel direction, while the other two eigenvectors point into orthogonal directions. The method proposed by Frangi et al. provides a ‘vesselness’ measure which depends on the eigenvalues values:

$$V(s) = \begin{cases} 0 & \text{if } \lambda_2 > 0 \text{ or } \lambda_3 > 0 \\ \left(1 - \exp\left(-\frac{R_A^2}{2\alpha_F}\right)\right) \exp\left(-\frac{R_B^2}{2\beta_F}\right) \left(1 - \exp\left(-\frac{S^2}{2c_F}\right)\right) & \end{cases}$$

where R_A and R_B are two geometric ratios. R_B accounts for the deviation of the vessels from a blob-like structure, while R_A is able to distinguish between plate-like and line-like structures. The vesselness measure also depends on two filter parameters which are empirically estimated in [47], and are kept constant for all data: $\alpha_F = 0.5$, $\beta_F = 0.5$. In addition, the intensity parameter c_F , which depends on the grey scale range of the input image, has been set to $c_F = 500$. For more details the reader is referred to [47].

After computing the filter over all the vessels, in each voxel, the Hessian matrix with the highest vesselness response $V(s)$ is selected, and the eigenvector corresponding to the lowest eigenvalue, *i.e.* λ_1 , is stored as the local vessel direction $\vec{d}(x)$. The vector field $\vec{d}(x)$ at each location is then normalised to unit length, as the subsequent calculations only require this directional information. An example of vessel directions in the *tumour* and a contralateral normal-appearing region *nac* is shown in Figure 7.4, where is seen that the vessels in the *tumour* region have incoherent orientations compared to those in the *nac* region.

7.2.2 Measures of Vessel Abnormality

This section describes the abnormality measures introduced in this thesis. To quantify the abnormality of the vessels in the *tumour*, regional vascular parameters are calculated from the field of vessel orientations \vec{d} . The abnormality measures that are here proposed

are based on the local (in)coherency of the vessels direction, namely the normalised dot product D between a pair of nearby vessel directions and the spherical angles Φ and Θ . The polar angle Θ and the azimuthal angle Φ are defined as:

$$\Theta(x) = \arccos \left(\frac{d_{x_3}(x_1, x_2, x_3)}{\|\vec{d}(x_1, x_2, x_3)\|} \right) \quad \Phi(x) = \arctan \left(\frac{d_{x_2}(x_1, x_2, x_3)}{d_{x_1}(x_1, x_2, x_3)} \right). \quad (7.6)$$

Note that both the polar angle Θ and the azimuthal angle Φ can assume values over the range from 0° to 180° .

To define the first abnormality measure, the polar angle variation $\Delta\Theta(x)$ is computed in a $3 \times 3 \times 3$ neighbourhood:

$$\Delta\Theta(x) = \frac{1}{3^3 - 1} \cdot \sqrt{\sum (\Theta(x_1 \pm \Delta x_1, x_2 \pm \Delta x_2, x_3 \pm \Delta x_3) - \Theta(x_1, x_2, x_3))^2}. \quad (7.7)$$

This measure is computed for each location x . Similarly, the variation $\Delta\Phi(x)$ is defined as:

$$\Delta\Phi(x) = \frac{1}{3^3 - 1} \cdot \sqrt{\sum (\Phi(x_1 \pm \Delta x_1, x_2 \pm \Delta x_2, x_3 \pm \Delta x_3) - \Phi(x_1, x_2, x_3))^2}. \quad (7.8)$$

From these variations, the mean values over the defined ROI s, where $ROI \in \{tumour, nac, natotal\}$, are calculated as the first abnormality measures:

$$\overline{\Delta\Theta}_{ROI} = \frac{1}{N_{ROI}} \cdot \sum_{x \in ROI} \Delta\Theta(x) \quad \overline{\Delta\Phi}_{ROI} = \frac{1}{N_{ROI}} \cdot \sum_{x \in ROI} \Delta\Phi(x) \quad (7.9)$$

where N_{ROI} is the number of voxels in each ROI . In a ‘coherence location’, that is, where the orientation vectors are pointing in similar directions, the variation of the spherical angles is expected to be lower than in a ‘incoherence location’ such as the *tumour* region. In addition to the mean values, the standard deviations are also computed for each ROI :

$$\begin{aligned} \sigma(\Delta\Theta_{ROI}) &= \sqrt{\frac{1}{N_{ROI}} \cdot \sum_{x \in ROI} (\Delta\Theta(x) - \overline{\Delta\Theta}_{ROI})^2} \\ \sigma(\Delta\Phi_{ROI}) &= \sqrt{\frac{1}{N_{ROI}} \cdot \sum_{x \in ROI} (\Delta\Phi(x) - \overline{\Delta\Phi}_{ROI})^2}. \end{aligned} \quad (7.10)$$

To define the second abnormality measure, the average of the dot product between the orientation vector $\vec{d}(x)$ within each of the $3 \times 3 \times 3$ neighbouring vectors $\vec{d}(x_1 \pm \Delta x_1, x_2 \pm \Delta x_2, x_3 \pm \Delta x_3)$ is calculated as:

$$D(x) = \frac{1}{3^3 - 1} \sum \vec{d}(x_1, x_2, x_3) \cdot \vec{d}(x_1 \pm \Delta x_1, x_2 \pm \Delta x_2, x_3 \pm \Delta x_3). \quad (7.11)$$

The dot product of two unit vectors equals the cosine of the relative angle between the vectors; thus, the maximum value is one (for parallel vectors), and lower values are




Figure 7.5: Plots of $\overline{\Delta\Phi}$ in (a) and $\overline{\Delta\Theta}$ in (b), for the three regions, *i.e.* *tumour* (blue), *nac* (red), and *natotal* (green) and for each patient. The error bars represent $\sigma(\Delta\Phi)$ and $\sigma(\Delta\Theta)$, respectively.

expected for incoherent directions. Again, the mean value over the defined *ROIs* is calculated as the second abnormality measure:

$$\overline{D}_{ROI} = \frac{1}{N_{ROI}} \sum_{x \in ROI} D(x). \quad (7.12)$$

Similarly also the standard deviation of the dot product is computed as:

$$\sigma(D_{ROI}) = \sqrt{\frac{1}{N_{ROI}} \sum_{x \in ROI} (D(x) - \overline{D}_{ROI})^2}. \quad (7.13)$$

The two abnormality measures $\{\overline{\Delta\Theta}, \overline{\Delta\Phi}\}$ and \overline{D} are calculated in the three defined *ROIs*. For comparison, the tortuosity measure SOAM, initially proposed by [24], is calculated from the vessel centrelines that are available as an output of the EMILOVE segmentation. A self-implementation of the SOAM algorithm has been performed. Each vessel is analysed independently, looping through triplets of points along the vessels. These points can be regarded as two consecutive vectors, and for which both the corresponding in- and through-plane angles are computed. The SOAM is given by the average, over the entire vessels, of a combination of the two angles. The standard deviation $\sigma(\text{SOAM})$ has also been computed.

For statistical analysis, a Student t-test is applied to all tumour-related parameters versus the parameters from both the normal-appearing contralateral side and the total brain volume excluding the tumour.

7.3 Comparison to Histology

Patients diagnosed with this type of malignant tumour undergo brain surgery within 24 hours of diagnosis. During the removal of the tumour several samples are analysed to




Figure 7.6: Plot, for each patient, of \bar{D} for the three regions, *i.e.* *tumour* (blue), *nac* (red), and *natotal* (green). The error bars represent $\sigma(D)$.

determine the type of the tumour and additionally several histological and gene markers are extracted.

The following histological parameters are determined (Table 7.1): the proliferation marker Ki-67 MIB, and the endothelial proliferation marker *ep*. The MIB is the percentage of all tumour cells in a mitotic state [87], while *ep* gives the percentage of vessels in the specimen showing endothelial proliferates [109]. In a healthy subject the MIB should be close to zero and the number of $ep = 0$.

These histologic markers are plotted against the \bar{D} measure. The \bar{D} as a function of MIB and *ep* for both the *tumour*, and the *nac* regions. Note that MIB and *ep* values are not calculated in any *nac* since the tissue sample is extracted only from the tumour, therefore these values are plotted with the \bar{D}_{nac} as a reference.

7.4 Results

The proposed method has been tested on 12 patients with an age range between 40 and 77 years, Table 7.1. A retrospective evaluation of the histology has been done on a database collected in the Dept. of Neurosurgery, University Medical Centre Freiburg, between 2012 and 2013. All except one patient underwent tumour resection at the Dep. of Neurosurgery. On the intraoperative tumour samples histopathological analysis was performed with the following results: 10 glioblastoma multiforme, and 1 anaplastic oligodendroglioma.

The study has been approved by the local review board, and the glioma patients gave informed consent for the acquisition and analysis of the MRA data. Brain images have been acquired as part of a pre-operative MR imaging assessment using a 3T Siemens Trio MRI scanner. The acquisition protocol consists of a 3D T_2 -w sequence with TR/TE/FA = 4390 ms / 99ms / 140°, with a matrix size = 448 × 512 × 23, and with a resolution of 0.43 × 0.43 × 6mm³; and a 3D TOF MRA sequence with TR/TE/FA = 23 ms / 4.12ms / 18°, with a matrix size = 438 × 512 × 84, and with a resolution of 0.39 × 0.39 × 0.55mm³.




Figure 7.7: Plot of SOAM measure for the three regions, *i.e.* *tumour* (blue), *nac* (red), and *natotal* (green), for each patient. The error bars represent $\sigma(SOAM)$.

Abnormality quantification could be successfully performed on all 12 patient datasets. The results for the abnormality measures $\overline{\Delta\Theta}$ and $\overline{\Delta\Phi}$ for each patient are shown in Figure 7.5. In all patients, the mean values in the *tumour*, *i.e.* $\overline{\Delta\Theta}_{tumour}$ and $\overline{\Delta\Phi}_{tumour}$, are larger than those in the *nac* and in the *natotal* regions. For $\overline{\Delta\Theta}$, its average value in *tumour* regions over all patients is about 34% higher than in the *nac*, and about 20% higher than in the *natotal* region. For $\overline{\Delta\Phi}$, even higher increases of 49% of the *tumour* region over the *nac* region and of 32% over the *natotal* region are seen. Both results indicate that in the tumour the orientation vectors are less coherent in comparison to normal appearing regions, as shown in Figure 7.4. It is also observed that the mean value of both $\overline{\Delta\Theta}$ and $\overline{\Delta\Phi}$ in *natotal* regions is in general higher than in *nac* regions, since in the brain there are also normal appearing small and tortuous vessels which are segmented.

The Student t-test shows a statistically significant difference, with p -value < 0.005 , between $\overline{\Delta\Theta}$ and $\overline{\Delta\Phi}$ in the *tumour* against corresponding values in the *nac* and *natotal* regions.

The results of the second abnormality measure, *i.e.* the dot product, are shown for every patient in Figure 7.6. When averaged over all patients, the value of \overline{D} in the *tumour* is between 24% and 50% lower than in the *nac* region, and between 22% and 48% lower than in the *natotal* region. Again, Student's t-test shows a statistically significant difference, with p -value < 0.001 , between \overline{D}_{tumour} and \overline{D}_{nac} , and between \overline{D}_{tumour} and $\overline{D}_{natotal}$.

The results of the tortuosity measure SOAM are shown in Figure 7.7. For all patients, the SOAM values in the *tumour* are higher than or equal to those found in the *nac*. However, for certain patients the SOAM values in the *natotal* region are higher than those in the *tumour* region, whereas for other patients the situation is opposite. This unclear trend has been confirmed by the statistical analysis. The Student t-test found a significant difference between *tumour* and *nac* values, with p -value < 0.03 , whereas no statistically significant difference has been found between *tumour* and *natotal* values.




Figure 7.8: \bar{D} as a function of MIB (a), and of endothelial proliferation (ep) (b). A negative trend is seen between \bar{D}_{tumour} and MIB, and a positive correlation is seen between \bar{D}_{tumour} and ep . The nac values are reported only for comparison.

The comparison of the dot product with histology shows a slight negative linear, but not statistically significant, correlation between MIB and \bar{D}_{tumour} , *i.e.* $\bar{D}_{tumour} = (0.43 \pm 0.03) + (-0.003 \pm 0.001) \cdot \text{MIB}$, Figure 7.8(a). Note that for the nac region, only the mean value of \bar{D}_{nac} over the patients is plotted as a constant value. The comparison between \bar{D}_{tumour} and the endothelial proliferation ep shows a weak positive correlation which is significant at $p\text{-value} < 0.05$, *i.e.* $\bar{D}_{tumour} = (0.23 \pm 0.03) + (3.1 \pm 1.1)10^{-3} \cdot ep$, Figure 7.8(b).

7.5 Summary and Discussion

In this thesis, novel abnormality quantification measures have been presented to discriminate between tumour and normal-appearing vessels in TOF-MRA data. The proposed method has been tested on 12 glioma patients, for all of which a statistically significant discrimination between tumour and normal-appearing vessels has been found. One patient did not undergo surgery, hence a glioma diagnosis could not be histologically confirmed. Nevertheless, the patient showed all imaging signs of a glioma and image data has been thus included in the abnormality calculations.

The first step of the proposed method is to perform a vessel segmentation with the EMILOVE algorithm of Skibbe et al. [120]. EMILOVE is based on simulated annealing and provides highly consistent vessel structures. A major advantage of this algorithm is that, once an initial estimate of the parameters has been obtained, it is able to compute the segmentation for the whole patient dataset without additional interaction on the part of the user. This is in contrast with the pipeline proposed in [111] where a manual threshold has to be set for every patient data in order to get a vessel segmentation. Additionally, in their original work Skibbe and co-workers, already applied the EMILOVE algorithm to TOF-MRA data of the brain. Thus, their parameter values could be directly used in this study to segment vessels in normal brain regions. Even though the algorithm has not been trained on the abnormal vasculature in the periphery of gliomas,

Table 7.1: Overview of relevant patients information: age, tumour size and location, histological result (Histo), tumour grade, MIB and endothelial proliferation (ep) values. GBM stands for glioblastoma, and aOG for anaplastic oligodendroglioma.

Patient	Age	Tumour size [cm ³]	Tumour location	Histo	Tumour grade	MIB[%]	ep
Pat. 01	64	11.52	Precuneus left	GBM	4	15	24.7
Pat. 02	72	57.06	Frontal right	GBM	4	25	14.8
Pat. 03	65	32.62	Frontal left	aOG	3	30	14.1
Pat. 04	48	31.73	Fronto-medial right	GBM	4	30	35.9
Pat. 05	73	7.31	Precentral right	GBM	4	35	28
Pat. 06	77	49.44	Temporal left	GBM	4	30	33.3
Pat. 07	77	14.35	Occipital right	GBM	4	20	9.9
Pat. 08	63	24.11	Temporal right	GBM	4	60	13
Pat. 09	40	59.97	Frontal right	GBM	4	30	29.1
Pat. 10	63	4.56	Parahippocampal gyrus	GBM	4	40	18.8
Pat. 11	59	7.85	Temporal left	GBM	4	30	18.8
Pat. 12	67	20.17	Temporal occipital right	n.a.	n.a.	n.a.	n.a.

the results of the segmentation in all patients show that the same algorithm parameters can also be used to extract tumour vasculature.

After the step of vessel segmentation, a vector field has been created, based on the main vessels direction extracted from the output of the Frangi filter. Based on this vector field, local angular variations have been estimated, *i.e.* $\overline{\Delta\Theta}$, $\overline{\Delta\Phi}$, and \overline{D} . Both the polar angular variation $\overline{\Delta\Theta}$ and the azimuthal variation $\overline{\Delta\Phi}$ show significantly higher values in the *tumour* region compared to both *nac* and *natotal* regions. The second abnormality measure proposed here is the dot product, *i.e.* \overline{D} . This measure is also clearly able to discriminate between tumour and normal-appearing vessels, and this result is in line with the angular measures. In general, the dot product is advantageous over the individual angular measures, as it is rotationally invariant, and it is represented by a single value.

All the results are consistent with the observation that tumour vessels are less organised and less coherent in their direction than normal brain vessels, and thus should show a higher directional variation. Interestingly, both the angular variational measures and the dot product not only provide a good relative discrimination between tumour vessels and normal-appearing vasculature, but the values of the normal brain vessels are also very coherent over all patients. This constancy of the values for the normal vessels indicates that the proposed method is robust against variations in the acquisition parameters, and image background from different local receive coil sensitivities. Thus, the variational measures in the normal vasculature might serve as reference values for a later vessel classification. Unfortunately no data from healthy controls was available to quantify a baseline for normal-appearing vessels in volunteers. However, it is expected that

the values computed in the *nac* and *natotal* regions of a patient data are not statistically different than the one computed in volunteers.

The proposed abnormality measures have been compared to the tortuosity measure SOAM [24], where a self-implementation has been used. The SOAM measure has been calculated in the *tumour*, *nac* and *natotal* regions of the same patient data. The statistical analysis revealed that there is a statistically significant difference between *tumour* and the *nac* regions, whereas no difference has been found between *tumour* and the *natotal* regions. The latter result may be related to the fact that tortuous vessels are, on average, also present in normal-appearing vasculature. A difference between SOAM and the dot product is the dependence on the choice of the coordinate system: SOAM values are not rotationally invariant, and thus depend on the orientation of the local slice coordinate system, whereas the dot product is defined by the relative angle between two vectors in space, which does not change under rotation.

A connection between microvascular malformation and high MIB values has already been published [134]. The variation in \overline{D} with MIB found in the present analysis, although not statistically significant, agrees qualitatively with the results of that study.

The positive, although weak, correlation with the endothelial proliferation marker ep and \overline{D} , is, however, unexpected. With decreasing ep values, \overline{D}_{tumour} should increase towards \overline{D}_{nac} which would result in a negative correlation. However, ep does not account for the direction of the vessels. Therefore, this correlation should be further investigated.

Chapter 8

Conclusion

This chapter summarises the research conducted in this thesis, its main contributions as well as known limitations and suggestions for potential future work. The major contribution of this thesis is the development of segmentation techniques for two specific brain pathologies, namely focal cortical dysplasia and multiple sclerosis. Additionally, novel measures for the quantification of abnormality in the brain vasculature have been developed for patients affected by glioblastoma tumours.

8.1 Summary of Contributions

After a brief introduction to the topic of medical imaging post-processing, the most relevant segmentation techniques have been reviewed (Chapters 3, 4). These include a simple thresholding approach as well as more advanced algorithms, such as the level-set method and atlas-based techniques.

An initial contribution based on the combination of level-set and atlas-based methods, has been presented (Chapter 3) that improves the segmentation of the brain grey matter compared to pure atlas-based approaches. The MR brain images used for this approach have been acquired at a 7T MR scanner. These images have a higher spatial resolution compared to 3T images, but on the other hand they present a stronger intensity inhomogeneity. Therefore, a multi-resolution approach has been developed to cope with the high computational costs of the bias-field estimation and with the large number of basis functions. The output of the atlas-based initial segmentation is used to initialise a level-set approach to find the grey matter boundaries.

Chapter 5 presented our contribution to the detection of focal cortical dysplasia. After performing an atlas-based pre-processing of the input T_1 -w and FLAIR images, several novel features that represent shape, intensity and texture of focal cortical dysplasia lesions are extracted. Due to the complexity in the detection of this kind of lesions, a baseline for normal-appearing tissue is created based on a database of healthy individuals. This is encoded in a probabilistic formulation, the probability of lesion being the complement of the probability of normal tissue. The proposed method has been tested on 11 patients and in all cases the performance of the detection was found to have a high specificity.

The detection of multiple sclerosis lesions is presented in Chapter 6 with our novel method based on graph-cuts. The main contribution is the use of a multi-labelling

approach to model the main brain tissue types and to constrain the lesion segmentation to the white matter. A novel connectivity constraint has been incorporated that forces the grey matter to remain a connected region. This resulted in lesions in white matter being correctly segmented that might otherwise be wrongly classified as grey matter due to their similar intensities. A second prior, namely a minimum lesion size, has been integrated to avoid potential noisy lesion segmentations. The proposed method has been tested on two different databases, the MICCAI and an in-house database with 15 patients. For both databases competitive results have been achieved with respect to state-of-the-art methods.

The more clinical oriented study concerning quantification has been presented in Chapter 7. Specifically, the vasculature of glioblastoma tumours was analysed. After performing an initial segmentation of the brain vasculature, two novel measures have been defined to represent the (in)coherency of vessel orientations, *i.e.* the dot product and the spherical angular variations. In order to quantify the abnormality/incoherency of tumour vessels, normal-appearing vessels were considered for comparison. In all tested cases, all the proposed measures were able to discriminate normal-appearing from tumour vessels. Additionally the correlation found between the dot product metric and histologic markers may hint to its potential use as marker for histologically determined vessel creation and constitution.

8.2 Limitations and Future Work

A main feature of all existing segmentation methods is their strong dependence on the quality of MR images. Therefore, the correct MR sequence parameters must be chosen so as to achieve highest contrast-to-noise and intensity-to-noise ratios. However, if during a measurement a patient moves, this results in strong artefacts in the acquired images. These artefacts corrupt image intensities, texture and structure, so that it may no longer be possible to correctly identify the boundary of the lesions or of the tumours. Another main limitation is that of data availability, which is usually limited to few patient data. This poses a strong limitation to method validation. To help remedy this situation, a publicly available database containing a large number of patient data instances will be created incorporating a number of brain diseases.

The probabilistic model for the detection of focal cortical dysplasia, presented in Chapter 5, strictly relies on the accuracy of the probability of normal-appearing tissue. This probability was computed via an initial registration of several healthy volunteers' data to a common space, followed by the combination of the extracted features. The registration can be inaccurate at the tissue boundaries, due to the high inter-brain variability. Moreover, the accuracy of the normal tissue probability depends on the amount of healthy data used. Therefore, in the future, the quantification of potential registration misalignment should be evaluated, as well as its impact on the quality of the results. In addition, larger databases should be used that embody more comprehensively the brain inter-subject variability.

In this thesis, the choice of extracted features has been based on the visual inspection of patient data. Thus, a quantification of the discriminant effectiveness of each single feature should be performed, thereby potentially leading to a higher performance and shorter computation time. The probability map for focal cortical dysplasia has been compared to the available resection volume which, for most patients, is much larger than the actual lesion. Therefore, the specificity of the proposed method, for some patients, has a low value. Consequently, comparison with a manual segmentation only of the lesion can help to have a more precise quantitative evaluation.

The geometric brain model developed for the segmentation of multiple sclerosis lesions, Chapter 6, is biased towards the lesions appearing in the white matter. This may result in lesions located close to the grey matter being erroneously segmented. Additionally, this geometric model is anatomically incorrect around the ventricles, and the segmentation of the lesions adjacent to them may be under-estimated, thus having some false negatives. A more anatomically accurate model may result in a performance increase.

The performance of the proposed method depends on the accuracy of the initial brain segmentation, which provides hard constraints: if at this step some relevant voxels are misclassified, the final lesion segmentation may be inaccurate. As future work, the robustness of the proposed method against incorrect initial segmentations may be performed. Moreover, an investigation of alternative algorithms, or procedures, is required in order to define initial hard constraints that are able to handle more flexibly variations in data quality and acquisition parameters.

The quantification of abnormality in the vasculature of brain tumour patients, presented in Chapter 7, starts with a segmentation of the vasculature. This is a fundamental step, as a poor segmentation can lead to incorrect results. Moreover, different methods for vessel segmentation may give different quantitative values. Future work may investigate the impact of the initial vessel segmentation on the quantification of vascular abnormality. Another point to note is that the proposed method relies on the manual identification of both the tumour regions and their contralateral counterparts. This time consuming process demands a lengthy training of the physicians. In future, this manual task could be replaced with either semiautomatic or fully automatic algorithms for region identification.

Bibliography

- [1] Abragam, A. (1961). *The principles of nuclear magnetism*. Oxford university press.
- [2] Antel, S., D. Collins, N. Bernasconi, F. Andermann, R. Shinghal, R. Kearney, D. Arnold, and A. Bernasconi (2003). Automated detection of focal cortical dysplasia lesions using computational models of their MRI characteristics and texture analysis. *Neuroimage* 19(4), 1748–1759.
- [3] Ashburner, J. and K. Friston (1997). Multimodal image coregistration and partitioning a unified framework. *Neuroimage* 6(3), 209–217.
- [4] Ashburner, J. and K. Friston (2000). Voxel-based morphometry-the methods. *Neuroimage* 11(6), 805–821.
- [5] Ashburner, J. and K. J. Friston (2005). Unified segmentation. *Neuroimage* 26(3), 839–851.
- [6] Avnir, D., O. Biham, D. Lidar, and O. Malcai (1998). Is the geometry of nature fractal? *Science* 279(5347), 39–40.
- [7] Barcelos, C. A. Z., M. Boaventura, and E. C. Silva Jr (2003). A well-balanced flow equation for noise removal and edge detection. *IEEE Transactions on Image Processing (TIP)* 12(7), 751–763.
- [8] Basser, P., C. Pierpaoli, et al. (1996). Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *JMR-Series B* 111(3), 209–219.
- [9] Bauer, S., C. Seiler, T. Bardyn, P. Buechler, and M. Reyes (2010). Atlas-based segmentation of brain tumor images using a markov random field-based tumor growth model and non-rigid registration. *IEEE International Conference of the Engineering in Medicine and Biology Society (EMBC)*, 4080–4083.
- [10] Bergo, F., A. Falcao, C. Yasuda, and F. Cendes (2008). FCD segmentation using texture asymmetry of MR-T1 images of the brain. *IEEE International Symposium on Biomedical Imaging (ISBI)*, 424–427.
- [11] Blake, A. and A. Zisserman (1987). *Visual reconstruction*, Volume 2. MIT press Cambridge.
- [12] Bock, M., A. Nagel, S. Schmitter, S. Johst, L. Gerigk, A. Biller, A. Milker-Zabel, L. Schuster, M. Essig, H.-P. Schlemmer, et al. Time-of-flight MRA at 7 Tesla.

- [13] Boros, E., P. Hammer, and G. Tavares (2006). Preprocessing of unconstrained quadratic binary optimization. *Technical Report RRR 10-2006, RUTCOR*.
- [14] Boykov, Y. and V. Kolmogorov (2004). An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* 26(9), 1124–1137.
- [15] Boykov, Y. and O. Veksler (2006). Graph cuts in vision and graphics: Theories and applications. *Handbook of mathematical models in computer vision*, 79–96.
- [16] Boykov, Y., O. Veksler, and R. Zabih (2001). Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* 23(11), 1222–1239.
- [17] Boykov, Y. Y. and M.-P. Jolly (2001). Interactive graph cuts for optimal boundary & region segmentation of objects in ND images. *IEEE Int. Conf. on Computer Vision (ICCV)*, 105–112.
- [18] Bredies, K., K. Kunisch, and T. Pock (2010). Total generalized variation. *SIAM Journal on Imaging Sciences* 3(3), 492–526.
- [19] Breiman, L. (2001). Random forests. *Machine learning* 45(1), 5–32.
- [20] Bresson, X., S. Esedoglu, P. Vandergheynst, J. Thiran, and S. Osher (2005). Global minimizers of the active contour/snake model. Technical report.
- [21] Brox, T., O. Kleinschmidt, and D. Cremers (2008). Efficient nonlocal means for denoising of textural patterns. *IEEE Transactions on Image Processing (TIP)* 17(7), 1083–1092.
- [22] Buades, A., B. Coll, and J.-M. Morel (2005). A non-local algorithm for image denoising. *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 60–65.
- [23] Bullitt, E., M. Ewend, J. Vredenburgh, A. Friedman, W. Lin, K. Wilber, D. Zeng, S. R. Aylward, and D. Reardon (2009). Computerized assessment of vessel morphological changes during treatment of glioblastoma multiforme: report of a case imaged serially by MRA over four years. *Neuroimage* 47, T143–T151.
- [24] Bullitt, E., G. Gerig, S. M. Pizer, W. Lin, and S. R. Aylward (2003). Measuring tortuosity of the intracerebral vasculature from MRA images. *IEEE Transactions on Medical Imaging (TMI)* 22(9), 1163–1171.
- [25] Cabezas, M., A. Oliver, X. Lladó, J. Freixenet, and M. B. Cuadra (2011). A review of atlas-based segmentation for magnetic resonance brain images. *Computer methods and programs in biomedicine* 104(3), e158–e177.

-
- [26] Cabezas, M., A. Oliver, S. Valverde, B. Beltran, J. Freixenet, J. C. Vilanova, L. Ramió-Torrentà, À. Rovira, and X. Lladó (2014). BOOST: A supervised approach for multiple sclerosis lesion segmentation. *Journal of neuroscience methods* 237, 108–117.
- [27] Caselles, V., F. Catté, T. Coll, and F. Dibos (1993). A geometric model for active contours in image processing. *Numerische mathematik* 66(1), 1–31.
- [28] Caselles, V., R. Kimmel, and G. Sapiro (1997). Geodesic active contours. *International journal of computer vision* 22(1), 61–79.
- [29] Chambolle, A. (2004). An algorithm for total variation minimization and applications. *Journal of Mathematical imaging and vision* 20(1-2), 89–97.
- [30] Chambolle, A. and P.-L. Lions (1997). Image recovery via total variation minimization and related problems. *Numerische Mathematik* 76(2), 167–188.
- [31] Chan, T. F., L. Vese, et al. (2001). Active contours without edges. *IEEE Transactions on Image Processing (TIP)* 10(2), 266–277.
- [32] Colliot, O., T. Mansi, P. Besson, N. Bernasconi, and A. Bernasconi (2006). Improved segmentation of focal cortical dysplasia lesions on MRI using expansion towards cortical boundaries. *IEEE International Symposium on Biomedical Imaging (ISBI)*, 323–326.
- [33] Coupé, P., J. V. Manjón, V. Fonov, J. Pruessner, M. Robles, and D. L. Collins (2011). Patch-based segmentation using expert priors: Application to hippocampus and ventricle segmentation. *NeuroImage* 54(2), 940–954.
- [34] Cox, I. J., S. B. Rao, and Y. Zhong (1996). "Ratio Regions": A technique for image segmentation. *Int. Conf. on Pattern Recognition*, 557–564.
- [35] Curtin, R. R., J. R. Cline, N. P. Slagle, W. B. March, P. Ram, N. A. Mehta, and A. G. Gray (2013). MLPACK: A scalable C++ machine learning library. *Journal of Machine Learning Research* 14, 801–805.
- [36] De Coene, B., J. V. Hajnal, P. Gatehouse, D. B. Longmore, S. J. White, A. Oatridge, J. Pennock, I. Young, and G. Bydder (1992). MR of the brain using fluid-attenuated inversion recovery (FLAIR) pulse sequences. *American journal of neuroradiology* 13(6), 1555–1564.
- [37] Delong, A. and Y. Boykov (2009). Globally optimal segmentation of multi-region objects. *IEEE Int. Conf. on Computer Vision (ICCV)*, 285–292.
- [38] Deshmukh, M. and U. Bhosle. A survey of image registration. *International Journal of Image Processing (IJIP)* 5(3), 245.
- [39] Di Ieva, A. (2010). Angioarchitectural morphometrics of brain tumors: are there any potential histopathological biomarkers? *Microvascular research* 80(3), 522–533.

- [40] Duda, R. O., P. E. Hart, and D. G. Stork (2012). *Pattern classification*. John Wiley & Sons.
- [41] Engel Jr, J. (1996). Surgery for seizures. *New England Journal of Medicine* 334(10), 647–653.
- [42] Evans, A., D. Collins, S. Mills, E. Brown, R. Kelly, and T. Peters (1993). 3D statistical neuroanatomical models from 305 MRI volumes. *Proc. of NSS and MIC.*, 1813–1817.
- [43] Evans, A., D. Collins, P. Neelin, D. MacDonald, M. Kamber, and T. Marrett (1994). Three-dimensional correlative imaging: applications in human brain mapping. *Functional neuroimaging: technical foundations*, 145–162.
- [44] F. Catte, P. Lions, J. Morel, T. C. (1992). Image selective smoothing and edge detection by nonlinear diffusion. *SIAM Journal of Numerical Analysis* 5(1), 530–558.
- [45] Fiot, J.-B., L. D. Cohen, P. Raniga, and J. Fripp (2013). Efficient brain lesion segmentation using multi-modality tissue-based feature selection and support vector machines. *International journal for numerical methods in biomedical engineering* 29(9), 905–915.
- [46] Focke, N., M. Symms, J. Burdett, and J. Duncan (2008). Voxel-based analysis of whole brain FLAIR at 3T detects focal cortical dysplasia. *Epilepsia* 49(5), 786–793.
- [47] Frangi, A. F., W. J. Niessen, K. L. Vincken, and M. A. Viergever (1998). Multi-scale vessel enhancement filtering. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 130–137.
- [48] Friedman, J., T. Hastie, R. Tibshirani, et al. (2000). Additive logistic regression: a statistical view of boosting. *The annals of statistics* 28(2), 337–407.
- [49] Gao, J., C. Li, C. Feng, M. Xie, Y. Yin, and C. Davatzikos (2014). Non-locally regularized segmentation of multiple sclerosis lesion from multi-channel MRI data. *Magnetic Resonance Imaging* 32(8), 1058–1066.
- [50] García-Lorenzo, D., S. Francis, S. Narayanan, D. L. Arnold, and D. L. Collins (2013). Review of automatic segmentation methods of multiple sclerosis white matter lesions on conventional magnetic resonance imaging. *Med. Image Analysis* 17(1), 1–18.
- [51] García-Lorenzo, D., J. Lecoer, D. L. Arnold, D. L. Collins, and C. Barillot (2009). Multiple sclerosis lesion segmentation using an automatic multimodal graph cuts. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 584–591.
- [52] Geman, S. and D. Geman (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* (6), 721–741.

-
- [53] Geremia, E., O. Clatz, B. H. Menze, E. Konukoglu, A. Criminisi, and N. Ayache (2011). Spatial decision forests for MS lesion segmentation in multi-channel magnetic resonance images. *Neuroimage* 57(2), 378–390.
- [54] Guidotti, P. (2009). A new nonlocal nonlinear diffusion of image processing. *Journal of Differential Equations* 246(12), 4731–4742.
- [55] Guizard, N., P. Coupé, V. S. Fonov, J. V. Manjón, D. L. Arnold, and D. L. Collins (2015). Rotation-invariant multi-contrast non-local means for MS lesion segmentation. *NeuroImage: Clinical* 8, 376–389.
- [56] Hadjidemetriou, S., M. Buechert, U. Ludwig, and J. Hennig (2011). Joint restoration of bi-contrast MRI data for spatial intensity non-uniformities. *Information Processing in Medical Imaging (IPMI)*, 346–358.
- [57] Haralick, R. M., K. Shanmugam, and I. H. Dinstein (1973). Textural features for image classification. *IEEE Transactions on Systems, Man and Cybernetics* (6), 610–621.
- [58] Held, K., E. R. Kops, B. J. Krause, W. M. Wells III, R. Kikinis, and H.-W. Muller-Gartner (1997). Markov random field segmentation of brain MR images. *IEEE Transactions on Medical Imaging (TMI)* 16(6), 878–886.
- [59] Herlidou-Meme, S., J. Constans, B. Carsin, D. Olivie, P. Eliat, L. Nadal-Desbarats, C. Gondry, E. Le Rumeur, I. Idy-Peretti, and J. de Certaines (2003). MRI texture analysis on texture test objects, normal brain and intracranial tumors. *Magnetic resonance imaging* 21(9), 989–993.
- [60] Höllig, K. and J. A. Nohel (1983). *A diffusion equation with a nonmonotone constitutive function*. Springer.
- [61] Huppertz, H.-J., J. Wellmer, A. M. Staack, D.-M. Altenmüller, H. Urbach, and J. Kröll (2008). Voxel-based 3D MRI analysis helps to detect subtle forms of subcortical band heterotopia. *Epilepsia* 49(5), 772–785.
- [62] Ibanez, L., W. Schroeder, L. Ng, J. Cates, et al. (2003). *The ITK software guide*. Kitware.
- [63] Ilea, D. E. and P. F. Whelan (2011). Image segmentation based on the integration of colour–texture descriptors—A review. *Pattern Recognition* 44(10), 2479–2501.
- [64] Ishikawa, H. (2003). Exact optimization for Markov random fields with convex priors. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* 25(10), 1333–1336.
- [65] Islam, A., S. Reza, and K. M. Iftekharuddin (2013). Multifractal texture estimation for detection and segmentation of brain tumors. *IEEE Transactions on Biomedical Engineering* 60(11), 3204–3215.

- [66] Jin, J. S., Y. Wang, and J. Hiller (2000). An adaptive nonlinear diffusion algorithm for filtering medical images. *IEEE Transactions on Information Technology in Biomedicine* 4(4), 298–305.
- [67] Kass, M., A. Witkin, and D. Terzopoulos (1988). Snakes: Active contour models. *International journal of computer vision* 1(4), 321–331.
- [68] Kaus, M., S. K. Warfield, A. Nabavi, E. Chatzidakis, P. M. Black, F. A. Jolesz, and R. Kikinis (1999). Segmentation of meningiomas and low grade gliomas in MRI. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 1–10.
- [69] Kichenassamy, S., A. Kumar, P. Olver, A. Tannenbaum, and A. Yezzi (1995). Gradient flows and geometric active contour models. *IEEE Int. Conf. on Computer Vision (ICCV)*, 810–815.
- [70] Kohli, P. and P. H. Torr (2005). Efficiently solving dynamic markov random fields using graph cuts. *IEEE Int. Conf. on Computer Vision (ICCV)*, 922–929.
- [71] Kolmogorov, V., Y. Boykov, and C. Rother (2007). Applications of parametric maxflow in computer vision. *IEEE Int. Conf. on Computer Vision (ICCV)*, 1–8.
- [72] Kolmogorov, V. and C. Rother (2007). Minimizing nonsubmodular functions with graph cuts—a review. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* 29(7), 1274–1279.
- [73] Kolmogorov, V. and R. Zabini (2004). What energy functions can be minimized via graph cuts? *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* 26(2), 147–159.
- [74] Laboratory of Neuro Imaging (UCLA), International Consortium for Brain Mapping. <http://lony.ucla.edu/Atlases/>.
- [75] Lao, Z., D. Shen, D. Liu, A. F. Jawad, E. R. Melhem, L. J. Launer, R. N. Bryan, and C. Davatzikos (2008). Computer-assisted segmentation of white matter lesions in 3D MR images using support vector machine. *Academic radiology* 15(3), 300–313.
- [76] Lee, C., S. Huh, T. A. Ketter, and M. Unser (1998). Unsupervised connectivity-based thresholding segmentation of midsagittal brain MR images. *Computers in biology and medicine* 28(3), 309–338.
- [77] Lee, C.-H., M. Schmidt, A. Murtha, A. Bistriz, J. Sander, and R. Greiner (2005). Segmenting brain tumors with conditional random fields and support vector machines. *Computer vision for biomedical image applications*, 469–478.
- [78] Lesion Segmentation Tool. <http://statisticalmodelling.de/lst.html>.
- [79] Levitt, M. H. (2001). *Spin dynamics: basics of nuclear magnetic resonance*. John Wiley & Sons.

-
- [80] Li, C., C. Gatenby, L. Wang, and J. C. Gore (2009). A robust parametric method for bias field estimation and segmentation of MR images. *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 218–223.
- [81] Li, C., C.-Y. Kao, J. C. Gore, and Z. Ding (2008). Minimization of region-scalable fitting energy for image segmentation. *IEEE Transactions on Image Processing (TIP)* 17(10), 1940–1949.
- [82] Li, C., C. Xu, C. Gui, and M. D. Fox (2005). Level set evolution without re-initialization: a new variational formulation. *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 430–436.
- [83] Li, C., C. Xu, C. Gui, and M. D. Fox (2010). Distance regularized level set evolution and its application to image segmentation. *IEEE Transactions on Image Processing (TIP)* 19(12), 3243–3254.
- [84] Li, H., M. Kallergi, L. Clarke, V. Jain, and R. Clark (1995). Markov random field for tumor detection in digital mammography. *IEEE Transactions on Medical Imaging (TMI)* 14(3), 565–576.
- [85] Li, S. Z. (2012). *Markov random field modeling in computer vision*. Springer Science & Business Media.
- [86] Lladó, X., A. Oliver, M. Cabezas, J. Freixenet, J. C. Vilanova, A. Quiles, L. Valls, L. Ramió-Torrentà, and À. Rovira (2012). Segmentation of multiple sclerosis lesions in brain MRI: a review of automated approaches. *Information Sciences* 186(1), 164–185.
- [87] Louis, D., S. Edgerton, A. Thor, and E. Hedley-Whyte (1991). Proliferating cell nuclear antigen and Ki-67 immunohistochemistry in brain tumors: a comparative study. *Acta neuropathologica* 81(6), 675–679.
- [88] Loyek, C., F. Woermann, and T. Nattkemper (2008). Detection of focal cortical dysplasia lesions in MRI using textural features. *Bildverarbeitung für die Medizin*, 432–436.
- [89] Magnin, B., L. Mesrob, S. Kinkingnéhun, M. Péligrini-Issac, O. Colliot, M. Sarazin, B. Dubois, S. Lehericy, and H. Benali (2009). Support vector machine-based classification of alzheimers disease from whole-brain anatomical MRI. *Neuroradiology* 51(2), 73–83.
- [90] Mahmoudi, M. and G. Sapiro (2005). Fast image and video denoising via nonlocal means of similar neighborhoods. *IEEE Signal Processing Letters* 12(12), 839–842.
- [91] Malladi, R., J. A. Sethian, and B. C. Vemuri (1994). Evolutionary fronts for topology-independent shape modeling and recovery. *European Conf. on Computer Vision (ECCV)*, 1–13.

- [92] Marroquín, J. L., B. C. Vemuri, S. Botello, F. Calderón, and A. Fernandez-Bouzas (2002). An accurate and efficient bayesian method for automatic segmentation of brain MRI. *IEEE Transactions on Medical Imaging (TMI)* 21(8), 934–945.
- [93] McAuliffe, M. J., F. M. Lalonde, D. McGarry, W. Gandler, K. Csaky, and B. L. Trus (2001). Medical image processing, analysis and visualization in clinical research. *IEEE International Symposium on Biomedical Imaging (ISBI)*, 381–386.
- [94] Mühlebner, A., R. Coras, K. Kobow, M. Feucht, T. Czech, H. Stefan, D. Weigel, M. Buchfelder, H. Holthausen, T. Pieper, et al. (2012). Neuropathologic measurements in focal cortical dysplasias: validation of the ILAE 2011 classification system and diagnostic implications for MRI. *Acta neuropathologica*, 1–14.
- [95] Mumford, D. and J. Shah (1989). Optimal approximations by piecewise smooth functions and associated variational problems. *Communications on pure and applied mathematics* 42(5), 577–685.
- [96] Munkres, J. (1957). Algorithms for the assignment and transportation problems. *SIAM Journal on Imaging Sciences* 5(1), 32–38.
- [97] Nitzberg, M. and T. Shiota (1992). Nonlinear image filtering with edge and corner enhancement. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* (8), 826–833.
- [98] Osher, S. and R. Fedkiw (2006). *Level set methods and dynamic implicit surfaces*, Volume 153. Springer Science and Business Media.
- [99] Osher, S. and J. A. Sethian (1988). Fronts propagating with curvature-dependent speed: algorithms based on hamilton-jacobi formulations. *Journal of computational physics* 79(1), 12–49.
- [100] Pal, N. R. and S. K. Pal (1993). A review on image segmentation techniques. *Pattern recognition* 26(9), 1277–1294.
- [101] Palmieri, A., I. Najm, G. Avanzini, T. Babb, R. Guerrini, N. Foldvary-Schaefer, G. Jackson, H. Lüders, R. Prayson, R. Spreafico, et al. (2004). Terminology and classification of the cortical dysplasias. *Neurology* 62(6 suppl 3), S2–S8.
- [102] Paragios, N. and R. Deriche (2002). Geodesic active regions and level set methods for supervised texture segmentation. *International Journal of Computer Vision (IJCV)* 46(3), 223–247.
- [103] Park, J. G. and C. Lee (2009). Skull stripping based on region growing for magnetic resonance brain images. *NeuroImage* 47(4), 1394–1407.
- [104] Peng, B., L. Zhang, and D. Zhang (2013). A survey of graph theoretical approaches to image segmentation. *Pattern Recognition* 46(3), 1020–1038.

-
- [105] Penny, W. D., K. J. Friston, J. T. Ashburner, S. J. Kiebel, and T. E. Nichols (2011). *Statistical parametric mapping: the analysis of functional brain images*. Academic press.
- [106] Perona, P. and J. Malik (1990). Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* 12(7), 629–639.
- [107] Pham, D. L., C. Xu, and J. L. Prince (2000). Current methods in medical image segmentation. *Annual review of biomedical engineering* 2(1), 315–337.
- [108] Pilutti, D., M. Strumia, and S. Hadjidemetriou (2014). Bi-modal non-rigid registration of brain MRI data with deconvolution of joint statistics. *IEEE Transactions on Image Processing (TIP)* 23(9), 3999–4009.
- [109] Plate, K. and H. Mennel (1995). Vascular morphology and angiogenesis in glial tumors. *Experimental and Toxicologic Pathology* 47(2), 89–94.
- [110] Potts, R. B. (1952). Some generalized order-disorder transformations. *Mathematical proceedings of the cambridge philosophical society* 48(01), 106–109.
- [111] Radbruch, A., O. Eidel, B. Wiestler, D. Paech, S. Burth, P. Kickingereeder, M. Nowosielski, P. Bäumer, W. Wick, H.-P. Schlemmer, et al. (2014). Quantification of tumor vessels in glioblastoma patients using time-of-flight angiography at 7 Tesla: a feasibility study. *PloS one* 9(11), e110727.
- [112] Ronfard, R. (1994). Region-based strategies for active contour models. *International Journal of Computer Vision (IJCV)* 13(2), 229–251.
- [113] Rother, C., V. Kolmogorov, V. Lempitsky, and M. Szummer (2007). Optimizing binary MRFs via extended roof duality. *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 1–8.
- [114] Roy, S., S. Nag, I. K. Maitra, and S. K. Bandyopadhyay (2013). A review on automated brain tumor detection and segmentation from MRI of brain. *arXiv preprint*.
- [115] Rudin, L. I., S. Osher, and E. Fatemi (1992). Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena* 60(1), 259–268.
- [116] Schmidt, P., C. Gaser, M. Arsic, D. Buck, A. Förchler, A. Berthele, M. Hoshi, R. Ilg, V. J. Schmid, C. Zimmer, B. Hemmer, and M. Mühlau (2012). An automated tool for detection of FLAIR-hyperintense white-matter lesions in MS. *Neuroimage* 59(4), 3774–3783.
- [117] Shi, J. and J. Malik (2000). Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* 22(8), 888–905.

- [118] Shiee, N., P.-L. Bazin, A. Ozturk, D. S. Reich, P. A. Calabresi, and D. L. Pham (2010). A topology-preserving approach to the segmentation of brain images with multiple sclerosis lesions. *Neuroimage* 49(2), 1524–1535.
- [119] Singleton, H. R. and G. M. Pohost (1997). Automatic cardiac MR image segmentation using edge detection by tissue classification in pixel neighborhoods. *Magnetic Resonance in Medicine* 37(3), 418–424.
- [120] Skibbe, H., M. Reiser, S.-i. Maeda, M. Koyama, S. Oba, K. Ito, and S. Ishii (2015). Efficient monte carlo image analysis for the location of vascular entity. *IEEE Transactions on Medical Imaging (TMI)* 34(2), 628–643.
- [121] Smith, S. (2002). Fast robust automated brain extraction. *Human Brain Mapping* 17(3), 143–155.
- [122] Song, Z., N. Tustison, B. Avants, and J. C. Gee (2006). Integrated graph cuts for brain MRI segmentation. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 831–838.
- [123] Souplet, J.-C., C. Lebrun, N. Ayache, and G. Malandain (2008). An automatic segmentation of T2-FLAIR multiple sclerosis lesions. *MIDAS Journal*.
- [124] Statistical Parametric Mapping. www.fil.ion.ucl.ac.uk/spm/.
- [125] Steinwart, I. and A. Christmann (2008). *Support vector machines*. Springer Science & Business Media.
- [126] Strumia, M., D. Feltell, N. Evangelou, P. Gowland, C. Tench, and L. Bai (2011). Grey matter segmentation of 7T MR images. *IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, 3710–3714.
- [127] Strumia, M., G. Ramantani, I. Mader, J. Henning, L. Bai, and S. Hadjidemetriou (2013). Analysis of structural MRI data for the localisation of focal cortical dysplasia in epilepsy. *MICCAI-Workshop Clinical Image-Based Procedures. From Planning to Intervention*, 25–32.
- [128] Strumia, M., W. Reichardt, I. Mader, J. Hennig, and M. Bock (2014). Glioblastoma vessels abnormality quantification in TOF-MRA. *International Society for Magnetic Resonance in Medicine (ISMRM)*.
- [129] Strumia, M., W. Reichardt, O. Staszewski, H. Heiland, A. Weyerbrock, I. Mader, and M. Bock (2016a). Correlation of vessel abnormality quantification from TOF MRA with histologic and gene markers. *International Society for Magnetic Resonance in Medicine (ISMRM)*.
- [130] Strumia, M., W. Reichardt, O. Staszewski, H. Heiland, A. Weyerbrock, I. Mader, and M. Bock (2016b). Glioblastoma vessels abnormality quantification using time-of-flight mr angiography. *Magnetic Resonance Materials in Physics, Biology and Medicine (MAGMA)*.

-
- [131] Strumia, M., F. R. Schmidt, C. Anastasopoulos, C. Granziera, G. Krueger, and T. Brox (2016). White matter MS-lesion segmentation using a geometric brain model. *IEEE Transactions on Medical Imaging (TMI)*.
- [132] Stühmer, J., P. Schroder, and D. Cremers (2013). Tree shape priors with connectivity constraints using convex relaxation on general graphs. *IEEE Int. Conf. on Computer Vision (ICCV)*, 2336–2343.
- [133] Styner, M., J. Lee, B. Chin, M. Chin, O. Commowick, H. Tran, S. Markovic-Plese, V. Jewells, and S. Warfield (2008). 3D segmentation in the clinic: A grand challenge II: MS lesion segmentation. *In MIDAS Journal 2008*, 1–6.
- [134] Sun, H., D. Guo, Y. Su, D. Yu, Q. Wang, T. Wang, Q. Zhou, X. Ran, and Z. Zou (2014). Hyperplasia of pericytes is one of the main characteristics of microvascular architecture in malignant glioma. *PloS one* 9(12), e114246.
- [135] Talairach, J. and P. Tournoux (1998). Co-planar stereotaxic atlas of the Human brain. *Transaction in Thieme*.
- [136] Tassi, L., N. Colombo, R. Garbelli, S. Francione, G. L. Russo, R. Mai, F. Cardinale, M. Cossu, A. Ferrario, C. Galli, et al. (2002). Focal cortical dysplasia: neuropathological subtypes, EEG, neuroimaging and surgical outcome. *Brain* 125(8), 1719–1732.
- [137] Tomas-Fernandez, X. and S. K. Warfield (2015). A model of population and subject (MOPS) intensities with application to multiple sclerosis lesion segmentation. *IEEE Transactions on Medical Imaging (TMI)* 34(6), 1349–1361.
- [138] Tošić, I. and P. Frossard (2011). Dictionary learning. *IEEE Signal Processing Magazine* 28(2), 27–38.
- [139] Tsitsios, C. and M. Petrou (2013). On the choice of the parameters for anisotropic diffusion in image processing. *Pattern recognition* 46(5), 1369–1381.
- [140] Tzourio-Mazoyer, N., B. Landeau, D. Papathanassiou, F. Crivello, O. Etard, N. Delcroix, B. Mazoyer, and M. Joliot (2002). Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *Neuroimage* 15(1), 273–289.
- [141] Unger, M., T. Pock, W. Trobin, D. Cremers, and H. Bischof (2008). Tvseg-interactive total variation based image segmentation. *BMVC* 31, 44–46.
- [142] Van der Lijn, F., T. den Heijer, M. M. Breteler, and W. J. Niessen (2008). Hippocampus segmentation in MR images using atlas registration, voxel classification, and graph cuts. *Neuroimage* 43(4), 708–720.
- [143] Van Leemput, K., F. Maes, D. Vandermeulen, A. Colchester, and P. Suetens (2001). Automated segmentation of multiple sclerosis lesions by model outlier detection. *IEEE Transactions on Medical Imaging (TMI)* 20(8), 677–688.

- [144] Van Leemput, K., F. Maes, D. Vandermeulen, and P. Suetens (1999). Automated model-based tissue classification of MR images of the brain. *IEEE Transactions on Medical Imaging (TMI)* 18(10), 897–908.
- [145] Van Rikxoort, E. and B. van Ginneken (2011). Automatic segmentation of the lungs and lobes from thoracic CT scans. *Fourth international workshop on pulmonary image analysis*, 261–268.
- [146] Vicente, S., V. Kolmogorov, and C. Rother (2008). Graph cut based image segmentation with connectivity priors. *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 1–8.
- [147] Wang, S. and J. M. Siskind (2001). Image segmentation with minimum mean cut. *IEEE Int. Conf. on Computer Vision (ICCV)*, 517–524.
- [148] Weickert, J. (1998). *Anisotropic diffusion in image processing*, Volume 1. Teubner Stuttgart.
- [149] Weiss, N., D. Rueckert, and A. Rao (2013). Multiple sclerosis lesion segmentation using dictionary learning and sparse coding. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 735–742.
- [150] Weiss, P., L. Blanc-Féraud, and G. Aubert (2009). Efficient schemes for total variation minimization under constraints in image processing. *SIAM journal on Scientific Computing* 31(3), 2047–2080.
- [151] Wolz, R., P. Aljabar, D. Rueckert, R. Heckemann, A. Hammers, et al. (2009). Segmentation of subcortical structures and the hippocampus in brain MRI using graph-cuts and subject-specific a-priori information. *IEEE International Symposium on Biomedical Imaging (ISBI)*, 470–473.
- [152] Wu, Z. and R. Leahy (1990). Tissue classification in MR images using hierarchical segmentation. *IEEE Nuclear Science Symposium*, 1410–1414.
- [153] Yang, C., M. Kaveh, and B. Erickson (2011). Automated detection of focal cortical dysplasia lesions on T1-weighted MRI using volume-based distributional features. *IEEE International Symposium on Biomedical Imaging (ISBI)*, 865–870.
- [154] Yoshizawa, S., A. Belyaev, and H.-P. Seidel (2006). Smoothing by example: Mesh denoising by averaging with similarity-based weights. *IEEE International Conference on Shape Modeling and Applications*, 9–9.
- [155] Yu, J., Y. Wang, and Y. Shen (2008). Noise reduction and edge detection via kernel anisotropic diffusion. *Pattern Recognition Letters* 29(10), 1496–1503.
- [156] Yu, Y. and S. T. Acton (2004). Edge detection in ultrasound imagery using the instantaneous coefficient of variation. *IEEE Transactions on Image Processing (TIP)* 13(12), 1640–1655.

- [157] Yushkevich, P. A., J. Piven, H. Cody Hazlett, R. Gimpel Smith, S. Ho, J. C. Gee, and G. Gerig (2006). User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. *Neuroimage* 31(3), 1116–1128.
- [158] Zhang, J., K.-K. Ma, M.-H. Er, and V. Chong (2004). Tumor segmentation from magnetic resonance imaging by learning via one-class support vector machine. *International Workshop on Advanced Image Technology*, 207–211.
- [159] Zhang, Y., M. Brady, and S. Smith (2001). Segmentation of brain MR images through a hidden markov random field model and the expectation-maximization algorithm. *IEEE Transactions on Medical Imaging (TMI)* 20(1), 45–57.
- [160] Zhang, Y. and L. Wu (2012). An MR brain images classifier via principal component analysis and kernel support vector machine. *Progress In Electromagnetics Research* 130, 369–388.
- [161] Zhao, H.-K., T. Chan, B. Merriman, and S. Osher (1996). A variational level set approach to multiphase motion. *Journal of computational physics* 127(1), 179–195.
- [162] Zijdenbos, A. P., R. Forghani, and A. C. Evans (2002). Automatic “pipeline” analysis of 3-D MRI data for clinical trials: application to multiple sclerosis. *IEEE Transactions on Medical Imaging (TMI)* 21(10), 1280–1291.
- [163] Zitova, B. and J. Flusser (2003). Image registration methods: a survey. *Image and vision computing* 21(11), 977–1000.