University of Groningen

Unsupervised brain anomaly detection in MR images

Botter Martins, Samuel

DOI:
[10.33612/diss.144368886](https://doi.org/10.33612/diss.144368886)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2020

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):
Botter Martins, S. (2020). *Unsupervised brain anomaly detection in MR images*. [Thesis fully internal (DIV), University of Groningen]. University of Groningen. <https://doi.org/10.33612/diss.144368886>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

UNSUPERVISED BRAIN ANOMALY DETECTION IN MR IMAGES

SAMUEL BOTTER MARTINS

Cover: A brain picture which is divided into some parts whose edges form a graph.

Unsupervised Brain Anomaly Detection in MR Images

Samuel Botter Martins
PhD Thesis

This thesis is the result of a joint PhD between the University of Campinas and the University of Groningen.

/ university of
 groningen

Unsupervised Brain Anomaly Detection in MR Images

PhD thesis

to obtain the degree of PhD at the
University of Groningen
on the authority of the
Rector Magnificus Prof. C. Wijmenga
and in accordance with
the decision by the College of Deans,
and
to obtain the degree of PhD at the
University of Campinas
on the authority of the
Rector Magnificus Prof. M. Knobel.

Double PhD degree

This thesis will be defended in public on
Friday 27 November 2020 at 12.45 hours

by

Samuel Botter Martins

born on 15 October 1990
in Pompéia-SP, Brazil

Supervisors

Prof. A. C. Telea

Prof. A. X. Falcão

Assessment committee

Prof. N. Petkov

Prof. M. Biehl

Prof. R. Marcondes Cesar

Prof. R. da Silva Torres

I have a definition of success.
It is not about wealth, fame, and power.
It is about how many shining eyes I have around me.
— Benjamin Zander —

ABSTRACT

Brain disorders are characterized by morphological deformations in shape and size of (sub)cortical structures in one or both hemispheres. These deformations cause deviations from the *normal* pattern of brain asymmetries, resulting in asymmetric lesions that directly affect the patient’s condition. It is hence clinically crucial to define *normal* brain asymmetries for the identification and detection of these deformations (brain anomalies) early for proper diagnosis and treatment.

Most automatic computational methods in the literature rely on *supervised* machine learning to detect or segment anomalies in brain images. However, these methods require a large number of high-quality annotated training images, which is absent for most medical image analysis problems. Besides, they are only designed for the lesions found in the training set, and some methods still require weight fine-tuning (retraining) when used for a new set of images. In contrast, *unsupervised* methods aim to learn a model from unlabeled healthy images, so that an unseen image that breaks priors of this model, *i.e.*, an outlier, is considered an anomaly. As these methods do not use labeled images, they are less effective in detecting lesions from a specific disease when compared to supervised approaches trained from labeled images for the same disease. For the same reason, however, unsupervised methods are generic in detecting any lesions, *e.g.*, coming from multiple diseases, as long as these notably differ from healthy training images.

This thesis addresses the development of solutions to leverage unsupervised machine learning for the detection/analysis of abnormal brain asymmetries related to anomalies in magnetic resonance (MR) images. First, we propose an automatic probabilistic-atlas-based approach for anomalous brain image segmentation. Its goal is to define our target macro-regions of interest — *i.e.*, right and left hemispheres, cerebellum, and brainstem — to improve the preprocessing, restrict the analysis, and compute hemispheric asymmetries in some cases. Second, we explore an automatic method for the detection of abnormal hippocampi from abnormal asymmetries. Our solution uses deep generative networks and a one-class classifier to model normal hippocampal asymmetries inside pairs of 3D patches from healthy subjects and detect abnormal hippocampi. Third, we present a more generic framework to detect abnormal asymmetries in the entire brain hemispheres. Our approach extracts pairs of symmetric regions — called *supervoxels* — in both hemispheres of a test image under study. One-class classifiers then analyze the asymmetries present in each pair. This method is limited to detect asymmetric lesions only in the hemispheres. Finally, we generalize the previous solution for the detection of (a)symmetric lesions based on registration errors. Experimental results on 3D MR-T1 images from healthy subjects and patients with a variety of lesions show the effectiveness and robustness of the proposed unsupervised approaches for brain anomaly detection.

SAMENVATTING

Hersenaandoeningen worden gekenmerkt door morfologische vervormingen van vorm en grootte van (sub)corticale structuren in één of beide hemisferen. Deze vervormingen veroorzaken afwijkingen van het normale patroon van hersenasymmetrieën, resulterend in asymmetrische laesies die de conditie van de patiënt direct beïnvloeden. Het is daarom klinisch cruciaal om normale hersenasymmetrieën te definiëren voor het vroegtijdig identificeren en detecteren van deze vervormingen (hersenaafwijkingen) voor een juiste diagnose en behandeling.

De meeste automatische berekeningsmethoden in de literatuur zijn gebaseerd op supervised machine learning om afwijkingen in hersenscans te detecteren of te segmenteren. Deze methoden vereisen echter een groot aantal geannoteerde trainingsbeelden van hoge kwaliteit, die bij de meeste medische beeldanalyseproblemen ontbreken. Bovendien zijn ze alleen ontworpen voor de laesies die in de trainingsset voorkomen, en sommige methoden vereisen nog steeds fine-tuning van het gewicht (retraining) wanneer ze worden gebruikt voor een nieuwe set afbeeldingen. Daarentegen richten unsupervised methoden zich op het leren van een model van niet-gelabelde gezonde afbeeldingen, zodat een onbekende afbeelding dat de priors van dit model breekt, *i.e.*, een outlier, als een afwijking wordt beschouwd. Aangezien deze methoden geen gelabelde afbeeldingen gebruiken, zijn ze minder effectief in het detecteren van laesies van een specifieke ziekte in vergelijking met supervised methoden die zijn getraind op gelabelde afbeeldingen voor dezelfde ziekte. Om dezelfde reden zijn unsupervised methoden echter generiek voor het opsporen van laesies, *e.g.* afkomstig van meerdere ziekten, zolang deze verschillen van gezonde trainingsbeelden.

Dit proefschrift behandelt de ontwikkeling van oplossingen om unsupervised machine learning toe te passen voor de detectie / analyse van abnormale hersenasymmetrieën gerelateerd aan afwijkingen in magnetische resonantie (MR) -beelden. Ten eerste stellen we een automatische probabilistic-atlas-based methode voor voor afwijkende hersenbeeldsegmentatie. Het doel is om onze beoogde macroregio's te definiëren – *i.e.*, de rechter en linker hersenhelft, het cerebellum en de hersenstam - om de preprocessing te verbeteren, de analyse te beperken en in sommige gevallen hemisferische asymmetrie te berekenen. Ten tweede onderzoeken we een automatische methode voor de detectie van abnormale hippocampi vanuit abnormale asymmetrieën. Onze oplossing maakt gebruik van deep generative networks en een one-class classifier om normale hippocampale asymmetrieën in paren van 3D-patches van gezonde proefpersonen te modelleren en abnormale hippocampi te detecteren. Ten derde presenteren we een meer generiek raamwerk om abnormale asymmetrieën in de gehele hersenhelften te detecteren. Onze benadering extraheert paren van symmetrische regio's - supervoxels genaamd - in beide hemisferen van een bestudeerd testbeeld. One-class classifiers analyseren vervol-

gens de asymmetrieën die in elk paar aanwezig zijn. Deze methode is gelimiteerd tot het detecteren van asymmetrische laesies in de hemisferen. Ten slotte generaliseren we de vorige oplossing voor het detecteren van (a)symmetrische laesies op basis van registratiefouten. Experimentele resultaten op 3D MR-T1-afbeeldingen van gezonde proefpersonen en patiënten met een verscheidenheid aan laesies tonen de effectiviteit en robuustheid van de voorgestelde unsupervised methoden voor detectie van hersenafwijkingen.

RESUMO

Distúrbios cerebrais são caracterizados por deformações morfológicas na forma e tamanho de estruturas (sub)corticais em um ou ambos hemisférios. Estas deformações causam desvios do padrão de normal das assimetrias cerebrais, resultando em lesões assimétricas que diretamente afetam a condição do paciente. É clinicamente crucial, portanto, definir assimetrias cerebrais normais para a identificação e detecção precoce destas deformações (anomalias cerebrais) para um diagnóstico e tratamento adequados.

A maioria dos métodos computacionais presentes na literatura confiam em aprendizado de máquina supervisionado para detectar ou segmentar anomalias em imagens de cérebro. Entretanto, estes métodos requerem um grande conjunto de imagens de treinamento de alta qualidade anotadas, que é escasso para a maioria dos problemas de análise de imagens médicas. Além disso, eles são projetados para as lesões encontradas no conjunto de treinamento, sendo que alguns métodos ainda requerem refinamento dos pesos do modelo (retreinamento) quando usados por um novo conjunto de imagens. Em contraste, métodos não-supervisionados visam aprender um modelo a partir de imagens saudáveis não-rotuladas, de maneira que uma imagem inédita que quebre condições prévias deste modelo, *i.e.*, um *outlier*, é considerada uma anomalia. À medida que estes métodos não usam imagens rotuladas, eles são menos efetivos em detectar lesões de uma doença específica, quando comparados com abordagens supervisionadas treinadas a partir de imagens rotuladas para a mesma doença. Pela mesma razão, entretanto, métodos não-supervisionados são genéricos em detectar qualquer lesão, por exemplo lesões provenientes de múltiplas doenças, uma vez que elas notavelmente diferente de imagens de treinamento saudáveis.

Esta tese endereça o desenvolvimento de soluções para alavancar o aprendizado de máquina não-supervisionado para a detecção/análise de assimetrias cerebrais anormais relacionadas a anomalias em imagens de ressonância magnética (RM). Primeiramente, nós propomos uma abordagem automática baseada em atlas probabilístico para a segmentação de cérebros anormais. Seu objeto é definir nossas macrorregiões de interesse — *i.e.*, hemisfério esquerdo e direito, cerebelo e tronco cerebral — para, assim, melhorar o pré-processamento, restringir a análise e computar assimetrias cerebrais em alguns casos. Em segundo lugar, nós exploramos um método automático para a detecção de hipocampus anormais a partir de assimetrias anormais. Nossa solução usa redes neurais generativas e classificadores de classe única para modelar assimetrias hipocampais normais dentro de pares de janelas 3D de pessoas saudáveis, e então detectar hipocampus anormais. Em terceiro lugar, nós apresentamos um arcabouço mais genérico para detectar assimetrias anormais em todas as regiões dos hemisférios. Nossa abordagem extrai pares de regiões simétricas — chamadas *supervoxels* — em ambos os hemisférios

de uma imagem de teste sob análise. Classificadores de classe única então analisam as assimetrias presentes em cada par. A detecção deste método limita-se a lesões assimétricas encontradas nos hemisférios. Finalmente, nós generalizamos a solução anterior para a detecção de lesões (as)simétricas baseadas em erros de registro. Os resultados experimentais em imagens de RM 3D de pessoas saudáveis e pacientes com uma variedade de lesões mostram a efetividade e robustez das abordagens não-supervisionadas propostas nesta tese para a detecção de anomalias cerebrais.

PUBLICATIONS

This thesis is the result of the following publications:

- A. X. Falcão, T. V. Spina, S. B. Martins, and R. Phellan, “Medical image segmentation using object shape models: A critical review on recent trends, and alternative directions,” in *Eccomas Thematic Conference on Computational Vision and Medical Image Processing (VipIMAGE)*, pp. 9–15, 2015.
- S. B. Martins, T. V. Spina, C. L. Yasuda, and A. X. Falcão, “A multi-object statistical atlas adaptive for deformable registration errors in anomalous medical image segmentation,” in *SPIE Medical Imaging*, vol. 10133, pp. 691–698, 2017. **Honorable mention.**
- S. B. Martins, J. Bragantini, C. L. Yasuda, and A. X. Falcão, “An adaptive probabilistic atlas for anomalous brain segmentation in MR images,” *Medical Physics*, vol. 46, no. 11, pp. 4940–4950, 2019.
- S. B. Martins, B. C. Benato, B. F. Silva, C. L. Yasuda, and A. X. Falcão, “Modeling normal brain asymmetry in MR images applied to anomaly detection without segmentation and data annotation,” in *SPIE Medical Imaging*, vol. 10950, pp. 71–80, 2019.
- S. B. Martins, G. Ruppert, F. Reis, C. L. Yasuda, and A. X. Falcão, “A supervoxel-based approach for unsupervised abnormal asymmetry detection in MR images of the brain,” in *IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 882–885, 2019.
- S. B. Martins, A. C. Telea, and A. X. Falcão, “Extending supervoxel-based abnormal brain asymmetry detection to the native image space,” in *IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 450–453, 2019.
- S. B. Martins, A. C. Telea, and A. X. Falcão, “Investigating the impact of supervoxel segmentation for unsupervised abnormal brain asymmetry detection,” *Computerized Medical Imaging and Graphics*, vol. 85, pp. 101770, 2020.
- S. B. Martins, A. X. Falcão, and A. C. Telea, “BADRESC: Brain anomaly detection based on registration errors and supervoxel classification,” in *International Joint Conference on Biomedical Engineering Systems and Technologies: BIOIMAGING*, pp. 74–81, 2020. **Best student paper awards.**
- S. B. Martins, A. X. Falcão, and A. C. Telea, “Combining Registration Errors and Supervoxel Classification for Unsupervised Brain Anomaly Detection,” Accepted for publication in *Lecture Notes in Computer Science*.

Other publications during the development of this thesis include:

- A. Z. Peixinho, S. B. Martins, J. E. Vargas, A. X. Falcão, J. F. Gomes, and C. T. N. Suzuki, “Diagnosis of human intestinal parasites by deep learning,” in *Proc. of the Eccomas Thematic Conference on Computational Vision and Medical Image Processing (VipIMAGE)*, pp. 107–112, 2015.
- T. V. Spina, S. B. Martins, and A. X. Falcão, “Interactive medical image segmentation by statistical seed models,” in *Conference on Graphics, Patterns and Images (SIBGRAPI)*, pp. 273–280, 2016.
- S. B. Martins, G. Chiachia, and A. X. Falcão, “A fast and robust negative mining approach for enrollment in face recognition systems,” in *Conference on Graphics, Patterns and Images (SIBGRAPI)*, pp. 201–208, 2017.
- J. Bragantini, S. B. Martins, C. Castelo-Fernandez, and A. X. Falcão, “Graph-based image segmentation using dynamic trees,” in *Iberoamerican Congress on Pattern Recognition*, pp. 470–478, 2018.
- A. M. Sousa, S. B. Martins, A. X. Falcão, F. Reis, E. Bagatin, and K. Irion, “AL-TIS: A fast and automatic lung and trachea CT-image segmentation method,” *Medical Physics*, vol. 46, no. 11, pp. 4970–4982, 2019.

CONTENTS

1	INTRODUCTION	1
1.1	Medical Imaging	1
1.2	Brain Asymmetries	2
1.3	Analysis of Brain Disorders	4
1.3.1	Machine Learning	5
1.3.2	Automatic Brain Anomaly Detection	8
1.4	Thesis Problems and Approach	9
2	BACKGROUND	13
2.1	Basic Anatomical Concepts	13
2.1.1	Brain Anatomy	14
2.1.2	Anatomical Planes of Body	16
2.2	Basic Imaging Physics	17
2.2.1	Medical Image Resolution	17
2.2.2	Magnetic Field Strength	18
2.2.3	Medical Image Orientation	20
2.3	MRI Preprocessing	21
2.3.1	Noise Reduction	21
2.3.2	MSP Estimation	23
2.3.3	Bias Field Correction	24
2.3.4	Image Registration	25
2.3.5	Skull Stripping	26
2.3.6	Intensity Normalization	27
2.4	Image Foresting Transform	27
2.4.1	Preliminary Concepts	28
2.4.2	The General IFT Algorithm	30
2.5	Clustering by Optimum-Path Forest	32
2.6	Iterative Spanning Forest (ISF)	35
2.6.1	Theoretical Background	36
2.6.2	The ISF Algorithm	36
2.7	Conclusion	39
3	AUTOMATIC BRAIN IMAGE SEGMENTATION	41
3.1	Related Work	43
3.1.1	Probabilistic Atlas	44
3.1.2	Multi-Atlas Label Fusion (MALF)	46
3.2	Adaptive Probabilistic atlas (AdaPro)	48
3.2.1	Construction	48
3.2.2	Segmentation	52

3.3	Experimental Setup	53
3.3.1	Datasets	53
3.3.2	Evaluation Protocol	54
3.4	Results	56
3.5	Conclusion	60
4	DETECTION OF ABNORMAL HIPPOCAMPAL ASYMMETRIES	63
4.1	Autoencoders	64
4.2	Proposed Approach	66
4.2.1	3D Image Preprocessing	67
4.2.2	VOI Localization	67
4.2.3	Normal VOI Asymmetry Representation	67
4.2.4	VOI Classification	69
4.3	Experiments and Results	70
4.3.1	Datasets	70
4.3.2	Localization Model	71
4.3.3	Hippocampal Asymmetry Detection	72
4.4	Extension for Brain Asymmetry Detection	76
4.4.1	Proposed Extension	76
4.4.2	Preliminary Experiments	77
4.5	Conclusion	79
5	DETECTION OF ABNORMAL BRAIN ASYMMETRIES	81
5.1	Related Work	82
5.1.1	Atlas-based Methods	82
5.1.2	Supervised Learning with Hand-crafted Features	83
5.1.3	Discriminative Deep Learning	83
5.1.4	Unsupervised Approaches	84
5.1.5	Deep Generative Neural Networks	84
5.2	Description of SAAD	85
5.2.1	3D Image Preprocessing	85
5.2.2	Asymmetry Computation	87
5.2.3	Symmetric Supervoxel Segmentation	88
5.2.4	Feature Extraction and Classification	89
5.3	Experiments	91
5.3.1	Datasets	92
5.3.2	Evaluation Protocol	92
5.4	Results	94
5.4.1	Impact of Supervoxel Segmentation Quality on Abnormal Asymmetry Detection	94
5.4.2	Improving the end-to-end method	97
5.4.3	Per-supervoxel vs Global Classifier Design	102
5.5	Extending SAAD for the Native Image Space	103
5.5.1	Description of N-SAAD	104

5.5.1.1	Asymmetry Computation	104
5.5.1.2	Symmetric Supervoxel Segmentation in NIS	105
5.5.1.3	Feature Extraction and Classification	106
5.5.2	Preliminary Experiments	107
5.6	Conclusion	110
6	DETECTION OF GENERAL BRAIN ANOMALIES	113
6.1	Description of TUSCA	114
6.1.1	3D Image Preprocessing	114
6.1.2	Saliency Computation	114
6.1.3	Supervoxel Segmentation	117
6.1.4	Feature Extraction and Classification	118
6.2	Experiments and Results	119
6.2.1	Evaluation Protocol	119
6.2.2	Results and Discussion	120
6.3	Conclusion	123
7	CONCLUSION	125
7.1	Brain Image Segmentation	125
7.2	Abnormal Asymmetry Detection by Autoencoders and One-Class Classification	126
7.3	Unsupervised Supervoxel-based Abnormal Brain Asymmetry Detection	127
7.4	Towards Unsupervised Supervoxel Classification for Anomaly Detection	128
A	NOTATIONS AND DEFINITIONS	131
B	DATA	135
B.1	In-house Datasets	135
B.2	Public Datasets	135
C	METRICS	137
C.1	Image Similarity Measures	137
C.1.1	Mean Square Error (MSE)	137
C.1.2	Normalized Mutual Information (NMI)	137
C.2	Segmentation Metrics	139
C.2.1	Intersection over Union (IoU)	139
C.2.2	Dice	140
C.2.3	Average Symmetric Surface Distance (ASSD)	140
	BIBLIOGRAPHY	143
	ACKNOWLEDGMENTS	163

LIST OF FIGURES

Figure 1.1	Brain images from different modalities	3
Figure 1.2	Examples of normal and abnormal brain asymmetries.	5
Figure 1.3	The different appearance of brain anomalies.	6
Figure 1.4	Supervised machine learning.	7
Figure 1.5	Unsupervised machine learning and outlier detection.	8
Figure 1.6	General pipeline for unsupervised brain anomaly detection.	10
Figure 2.1	A simple diagram of the nervous system.	14
Figure 2.2	Brain regions and some of their corresponding responsibilities.	15
Figure 2.3	Anatomical planes of the brain.	17
Figure 2.4	Comparison between the same axial slice of an MR-T1 brain image with different spatial resolution.	19
Figure 2.5	Comparison between MR-T1 brain images with different field strengths.	19
Figure 2.6	Coordinate system with the <i>LPS+</i> orientation.	20
Figure 2.7	General preprocessing steps for MR brain images.	22
Figure 2.8	An axial slice of a noisy MR-T1 brain image and its filtered result by median filtering.	22
Figure 2.9	Steps for MSP definition by Ruppert <i>et al.</i>	23
Figure 2.10	Example of bias field correction.	24
Figure 2.11	Example of affine and non-rigid registration.	25
Figure 2.12	The proposed intensity normalization.	28
Figure 2.13	Examples of adjacency relation.	29
Figure 2.14	Multi-object image segmentation by IFT.	30
Figure 2.15	Example of the IFT seed competition with f_{max} .	32
Figure 2.16	The impact of the factors α and β for superpixel/super-voxel segmentation by ISF.	37
Figure 2.17	Example of the ISF execution on a 2D brain image.	38
Figure 3.1	Automatic image segmentation by the probabilistic-atlas-based method SOSM-S.	42
Figure 3.2	General steps for the construction and use of probabilistic atlases for automatic image segmentation.	44
Figure 3.3	General steps of Multi-Atlas Label Fusion for image segmentation.	46
Figure 3.4	Pipeline for the construction and use of AdaPro.	49
Figure 3.5	AdaPro segmentation.	50

Figure 3.6	Probabilistic atlases for the cerebellum, right hemisphere, and left hemisphere. 51
Figure 3.7	Selected voxels to design the texture binary classifier of AdaPro. 51
Figure 3.8	Examples of the considered datasets for brain segmentation. 55
Figure 3.9	Axial and coronal slices with the mean segmentation errors from the baselines. 59
Figure 4.1	Sagittal slice with an example of segmentation errors in the left hippocampus. 63
Figure 4.2	General structure of an autoencoder. 65
Figure 4.3	General pipeline of the proposed autoencoder-based approach to model normal brain asymmetries. 66
Figure 4.4	Scheme for the training and use of the proposed patch-based model (PBM) for VOI localization in 3D brain images. 68
Figure 4.5	Architecture of the convolution autoencoder (CAE) used for normal VOI asymmetry representation. 69
Figure 4.6	Datasets used for evaluation of abnormal hippocampal asymmetries. 71
Figure 4.7	The 2D t-SNE projection from the considered datasets for the CAE-based representations with PBM localization. 75
Figure 4.8	Uniform grid-sampling used to define the geometric centers of VOIs along the hemispheres. 76
Figure 4.9	Examples of false-positive abnormal asymmetries detected by the extended autoencoder-based method. 77
Figure 4.10	Pair of undetected VOIs intersecting an anomalous region of a postoperative image. 78
Figure 4.11	Pair of undetected VOIs covering a small lesion of an epilepsy patient image. 79
Figure 5.1	The pipeline of SAAD. 86
Figure 5.2	3D image preprocessing and registration steps. 86
Figure 5.3	Asymmetry computation on a standard image space. 87
Figure 5.4	The pipeline of SymmISF. 89
Figure 5.5	One-class classifier training for abnormal asymmetry detection. 90
Figure 5.6	Abnormal asymmetry detection by supervoxel classification. 91
Figure 5.7	Examples of false-positive supervoxels for two different brain slices. 94

Figure 5.8	Correlation between some characteristics of false-positive supervoxels. 101
Figure 5.9	t-SNE projection from texture feature vectors for the symmetric supervoxel extracted. 103
Figure 5.10	Results of the simplest extension of SAAD for the native image space (NIS). 104
Figure 5.11	Asymmetry computation of a 3D test stroke image in its own native image space. 105
Figure 5.12	SymmISF in NIS. 106
Figure 5.13	Results of N-SAAD on the ATLAS dataset. 109
Figure 6.1	The pipeline of TUSCA. 115
Figure 6.2	Registration error computation. 116
Figure 6.3	General pipeline for supervoxel segmentation. 118
Figure 6.4	Comparative results between SAAD and BADRES. 122
Figure 6.5	Results of BADRES for some images with stroke lesions in the cerebellum or brainstem. 123
Figure C.1	Example of two different grayscale images with the same Shannon Entropy. 138
Figure C.2	Illustration of IoU and Dice. 139
Figure C.3	Average Symmetric Surface Distance (ASSD). 139
Figure C.4	Two cases with approximately equal Dice, but different ASSD. 140

INTRODUCTION

The brain is the most complex organ in a vertebrate's body and serves the central nervous system (CNS) — a complex collection of billions of specialized nerves and cells known as neurons that transmit signals between different parts of the body [1, 2]. CNS represents a communication network of the organism that detects and responds to changes in its internal and external environment. Any dysfunctionality can severely impact a person's health and quality of life, resulting in problems as memory loss, motor skills, and mobility.

A *brain disorder* consists of any condition that affects one's brain. These conditions are mainly caused by genetic abnormalities, illness, and traumatic injuries [3]. Brain disorders are a major public health problem in the world [4]. According to reports presented in 2010 by the European Brain Council — an alliance of all major European organizations interested in brain diseases — about one-third of all European citizens had at least one brain disorder [4, 5]. Most cases consist of minor disorders such as migraine, whereas neuromuscular disorders and brain tumors are less prevalent. However, the diagnoses and treatments for the latter are more complex and very expensive. For example, the cost of the treatment of brain tumors per subject is 33,900 euros on average, whereas the one for migraine is about 662 euros [4].

Following the above, it is hence clinically crucial to detect brain lesions early for proper diagnosis and treatment. There is a variety of possible treatments, such as chemotherapy and surgical resection. The choice of treatment usually depends on the type of brain lesion, its anatomy, and location [6, 7]. This information is obtained from medical imaging.

1.1 MEDICAL IMAGING

Medical images are visual representations of physical features measured from the interior of a body for clinical analysis, medical diagnoses, and intervention [8]. They show attributes from such body structures in a noninvasive manner.¹

The first medical image dates the late 19th century from the discovery of X-rays by the German Wilhelm Röntgen. For the first time in history, an image — created by marked X-ray absorption — allowed noninvasive insights in the human body [9]. This imaging technique was called *radiography*. The more X-rays a tissue absorbs, the whiter it is in the X-ray image (Fig. 1.1a). Thus, dense tissues (e.g., bones) appear white, whereas fat and other soft tissues look gray or even black

¹ *Noninvasive* denotes a medical procedure that does not involve the introduction of instruments into the patient's body.

(e.g., the air inside the lungs). Soon after its introduction, radiography quickly became essential for medical diagnosis. Currently, digital X-ray images are widely used to examine bone fractures and detect certain diseases, e.g., pneumonia and pulmonary edema, in soft tissues [8].

New medical imaging techniques and technologies have emerged in the last 60 years, in particular, Computed Tomography (CT) and Magnetic Resonance Imaging (MRI). A *CT scanner* takes a series of X-rays emitted at different angles to generate a detailed volumetric image (3D image) of a particular section of the body. Elements of a 3D image are called *voxels*, by analogy to the *pixel* elements of a 2D image. Voxels are defined by their 3D coordinates and their corresponding values. CT images are more expensive to acquire than conventional X-ray images but yield a better way to separate between various types of tissues, atop the ability to reason about spatial structures in the body. Some common uses of CT images consist of diagnosing injuries from trauma, determining the location of a tumor, and detecting the location of blood clots.

MRI scanners do not use radiation during imaging. Instead, they produce a powerful fixed magnetic field around the patient so that radiofrequency waves excite protons within the body. As the excited protons relax back to their normal position, they emit signals that are captured and mapped into a 3D image [9, 10]. MRIs provide more detailed information about inner organs with superior soft-tissue contrast and anatomic detail compared to X-ray and CT images (Fig. 1.1). However, they are more expensive and take considerably more time to generate.² MRI is usually the commonly chosen image modality for structural brain analysis [11].

Different types of MR images can be obtained during the examination. The most common types are T1 and T2. Both types accentuate different characteristics of tissues resulting in images with distinct appearances. Water-rich structures — e.g., the cerebrospinal fluid (CSF) found in the brain and spinal cord — are dark in T1 and very bright in T2. Conversely, structures containing fat are considerably brighter in T1 than T2. For brain images, gray matter is darker than white matter in T1. The opposite is true for T2 — compare the pair of brain slices in Fig. 1.1c. Therefore, T1 images are more effective for analyzing anatomical structures, whereas T2 images are typically used when looking for areas of inflammation [12, 13]. This thesis focuses on the analysis of MR-T1 images of the brain for anomaly detection.

1.2 BRAIN ASYMMETRIES

The brain hemispheres can be distinguished visually by the longitudinal fissure (Fig. 2.3) — a membrane between both hemispheres filled with cerebrospinal fluid (CSF). Although they are, at a coarse scale, almost symmetrical in structure, subtle (finer-scale) anatomical differences between them exist [1, 14, 15]. These differ-

² A CT image takes 10 minutes on average depending on the body part being examined whereas an MR image takes between 45 minutes to 1 hour.

(a) X-ray image

(b) Computer Tomography (CT) image

(c) Magnetic Resonance Imaging (MRI)

Figure 1.1: Brain images from different modalities. (a) X-ray image. (b) Axial, sagittal, and coronal slices of a CT brain volumetric image. (c) Axial, sagittal, and coronal slices of MR T1 and T2 brain volumetric images of the same subject. CT and MR images provide superior soft-tissue contrast and anatomic detail compared to X-ray images. Water-rich structures are dark in T1 and very bright in T2, whereas structures containing fat are considerably brighter in T1 than T2.

ences are called *hemispheric asymmetries* or simply *brain asymmetries* and can be defined at functional and structural levels [16].

Functional differences between the hemispheres — so-called *hemispheric lateralization* — have been observed for several cognitive functions [17]. Both hemispheres are indeed specialized for separate tasks. The left hemisphere is more dominant for handedness and language than the right one. For instance, most humans are right-handed³, whose motor coordination is performed by the left hemisphere [17, 18]. Conversely, the right hemisphere is dominant, for example, for visuospatial processing, face recognition, music, and visual imagery [19, 20].

The realization of the functional differences between the brain hemispheres raises questions regarding the structural correlation of such lateralization [21]. **Structural differences** include changes in volume, shape, and size of (sub)cortical structures (e.g., sulci, cerebral lobes, and hippocampus) as well as a different amount of white and gray matter in the hemispheres [21, 22]. This thesis only focus on the analysis of *structural* differences.

Deviations from the *normal* pattern of brain asymmetries are useful insights about neurological pathologies [23]. Studies have shown that some neurological diseases — such as Alzheimer’s [24], schizophrenia [25, 26], epilepsy [27–29], and autism [30] — are indeed associated to abnormal brain asymmetries. Morphological changes in (sub)cortical in one or both hemispheres characterize these structural abnormalities, as illustrated in Fig. 1.2. Therefore, it becomes crucial to define *normal* brain asymmetries for the identification and detection of many abnormalities in the brain. We widely explore lesions associated with *abnormal asymmetries* throughout this thesis.

1.3 ANALYSIS OF BRAIN DISORDERS

Quantitative analysis of MR brain images has been used extensively for the characterization of brain disorders, such as stroke, tumors, and multiple sclerosis. Such methods rely on delineating objects of interest — (sub)cortical structures or lesions — trying to solve detection and segmentation simultaneously. Results are usually used for tasks such as quantitative lesion assessment (e.g., volume), surgical planning, and overall anatomic understanding [6, 31, 32]. Note that *segmentation* corresponds to the exact delineation of the object of interest, whereas *detection* consists of finding the rough location of such objects (e.g., by a bounding box around the object), in case they are present in the image.

The simplest strategy to detect brain anomalies consists of a visual slice-by-slice inspection by one or multiple specialists. This process is very time-consuming, error-prone, and even impracticable when a large amount of data needs to be processed.

The analysis of brain asymmetries commonly follows a similar strategy. First, the approach interactively segments structures of interest in the image, such as

³ Approximately 90% of the world population are right-handed [17, 18].

Figure 1.2: MR images and their corresponding asymmetry maps for (a) a healthy subject and (b) a stroke patient. Green borders indicate examples of pairs of regions with normal asymmetries, whereas red borders indicate abnormal asymmetries resulted from a stroke. The dashed yellow lines show mid-sagittal planes. Normal asymmetries are accentuated on the brain cortex (regions close to the borders). Both cases omit other regions with normal asymmetries.

hippocampi, amygdala, and putamen. Then, it computes morphometric measures from the segmented structures (*e.g.*, volume), and performs statistical analysis of these measures [33]. However, this strategy is also problematic since the interactive segmentation of brain structures may be very complicated, extremely susceptible to errors, and that demands much time from the expert. Thus, segmentation errors may severely impact the analysis.

Continuous efforts have been made for automatic anomaly detection that delineates anomalies with accuracy close to that of human experts. However, this goal is very challenging and complex due to the large variability in shape, size, and location present in different anomalies, even when the same disease causes these (see, *e.g.*, Fig. 1.3). All these difficulties have motivated the research and development of automatic brain anomaly detection methods based on *machine learning* algorithms, as discussed next.

1.3.1 Machine Learning

Machine learning (ML) can aid experts in detecting and classifying lesions from a brain image [35]. ML is based on algorithms that can learn from a dataset without being explicitly programmed to perform a task [36]. Each example from the dataset is called *sample*, and it is described by a set of features, called *feature vector*. For medical image analysis, a sample can be defined, for example, as a voxel, the image of a segmented object, or the shape attributes (descriptors) computed on this object. *Feature extraction* algorithms, in turn, are chosen according to the targeted

Figure 1.3: The different appearance of brain anomalies. **Top:** axial slices of three stroke patients with lesions (gold-standard borders in pink) that significantly differ in location, shape, and size. **Bottom:** slices of a 3D heatmap show the location frequency of stroke lesions across the brain. Although caused by the same disease, the lesions are sparsely distributed in the brain resulting in low-concentrated regions. The 3D heatmap was built from aligned manual lesion segmentation of stroke patients from the ATLAS dataset [34] after registration to a standard template.

problem and sample type. Texture [6, 37–40], shape features [41–43], and, more recently, deep-learning-based features [35, 44–46] are common feature examples adopted in medical image analysis problems.

Overall, machine learning can be either supervised or unsupervised. In *supervised learning*, the dataset is labeled, *i.e.*, each of its samples has an assigned class.⁴ For example, a dataset of MR brain images (samples) that is used in a classification task that aims to discriminate between normal and abnormal tissue will use two classes: normal and abnormal. A classification algorithm learns a *decision model* from labeled samples of a given training set by associating features to classes [47]. More generally, when the algorithm predicts a continuous value rather than a categorical class value, one says that it learns a *regression model*. In our work, we will mainly focus on decision models. New unseen samples are then classified according to the learned decision model. Fig. 1.4a shows a toy example of two easy separable classes with a linear classifier, *i.e.*, a classifier that assumes that the boundary between samples of the two existing classes is linear. Typically, linear classifiers are not sufficient to predict the correct classes of more complex sam-

⁴ Some classification problems might consider a sample with more than one label.

ple distributions in real-world data, as shown by the example in Fig. 1.4b. In such cases, *nonlinear* classifiers are used to properly split the feature space into areas corresponding to the two classes (Fig. 1.4c).

Figure 1.4: Toy example displaying the relation between feature 1 and feature 2 and two classes.⁵ (a) A linear classifier that can separate the given samples. (b) A linear classifier unable to separate other given samples. (c) A nonlinear classifier that separates the samples of (b).

Unsupervised machine learning algorithms aim at finding intrinsic structures in an *unlabeled/uncategorized* dataset [48]. The key added value of unsupervised methods as compared to supervised ones is that one does not need an expert to have created an annotated (labeled) training set. This is particularly essential in situations where labeling is expensive and requires specialist expertise, such as in the case of medical imaging datasets to be manually labeled by delineation by trained medical professionals. A potential drawback of unsupervised learning is that the structures extracted from an (image) dataset may not always be relevant to the expert [48]. *Clustering* is arguably the best known unsupervised strategy. It finds patterns in the feature space and uses these to divide the dataset into groups that exhibit high internal coherence and low similarity with other groups. Figs. 1.5a–b illustrate results produced by clustering for hypothetical data.

Outlier detection — also called anomaly detection — is another common problem in unsupervised machine learning.⁶ Techniques aim to detect *outliers* in an unlabeled dataset under the assumption that the majority of its samples are *normal* [49]. An *outlier* is a sample that differs significantly from the remainder of the dataset. Some authors also refer to outliers as anomalies, exceptions, noise, and novelties. Several applications use outlier detection, such as bank fraud detection, loan application processing, and medical condition monitoring [49]. Fig. 1.5c shows an example of outlier detection.

⁵ Figure inspired by the Ph.D. thesis of Jansen (2019) [36].

⁶ Some authors consider the term *supervised anomaly detection* when the training set has only two classes: normal and outlier [49]. A binary classifier is then trained for outlier detection.

Figure 1.5: (a) A hypothetical unlabeled dataset. (b) Resulting groups after performing a given clustering algorithm. Each color represents a different group. (c) Example of outlier detection. If an unseen test sample is far from the training set of normal samples (the yellow region with dashed borders), it is classified as an outlier.

Medical image analysis commonly uses outlier detection mainly for detecting anomalies (lesions). *One-class classification* (OCC) — also called unary classification — is a class of techniques commonly used for this purpose [40, 50–53]. Consider a training dataset with *only* medical images of *healthy* subjects — also known as *control* images. All training samples have the same single class: *healthy*. The OCC learns a classification boundary for the healthy class to classify new unseen images as *healthy* or *outlier*. Detected outliers are considered as anomalies, e.g., tumors, stroke, and cancer. OCC is different from and more challenging than the traditional classification problem, which tries to differentiate two or more classes from a labeled training set. In this thesis, we focus on unsupervised algorithms in particular one-class classification.

1.3.2 Automatic Brain Anomaly Detection

Most automatic methods in the literature rely on supervised machine learning to detect or segment brain anomalies. They train a classifier from training images — which must be previously labeled (e.g., lesion segmentation masks) by experts — to delineate anomalies by classifying voxels or regions of the target image. Traditional image features (e.g., edge detectors and texture features) and deep feature representations (e.g., convolutional features) are commonly used [6, 37–39, 46, 54–56].

However, these supervised methods commonly have three main limitations. First, they require a large number of high-quality annotated training images, which is absent for most medical image analysis problems [11, 35, 57]. Second, they are only designed for the lesions found in the training set. Third, some methods still require weight fine-tuning (retraining) when used for a new set of images

due to image variability across scanners and acquisition protocols, limiting its application into clinical routine.

All the above limitations of supervised methods motivate research on *unsupervised* anomaly detection approaches [13, 40, 53, 58, 59]. From a training set with images of *healthy* subjects *only*, these methods perform an outlier detection technique to identify anomalies in new images. Some of these methods can detect enormous lesions [58, 59], but show poor results with small lesions, which are the most challenging cases.

1.4 THESIS PROBLEMS AND APPROACH

As unsupervised brain anomaly detection methods do not use labeled samples, they are less effective in detecting lesions from a specific disease when compared to supervised approaches trained from labeled samples for the same disease. For the same reason, however, unsupervised methods are generic in detecting any lesions, *e.g.*, coming from multiple diseases, as long as these notably differ from healthy training samples.

Combining the pros and cons of unsupervised methods listed above, as well as the importance of identifying abnormal brain asymmetries associated to brain anomalies, we can now state the key research questions of this thesis:

RQ1: *Can we model normal brain asymmetries?*

RQ2: *Can we use the normal brain asymmetry model to detect brain anomalies?*

To illustrate how we approach answering these questions, let us consider the typical pipeline for brain image processing and analysis (Fig. 1.6). Given a 3D MR-T1 image, we first perform several preprocessing tasks (*e.g.*, noise filtering and intensity normalization) to overcome inherent acquisition issues, such as noise and inhomogeneity field. Next, we define the *volumes of interest* (VOI) to be analyzed: either the entire brain or some specific region. Features related to brain asymmetries are extracted from these VOIs and subsequently classified as *normal* or *abnormal* from the knowledge about normal asymmetries present in a training set of control images. We evaluate our approaches on MR-T1 images, mainly due to the greater availability of public datasets of healthy and abnormal brain volumetric images for this imaging modality. Public datasets of different imaging modalities exist. However, some only provide a subset of 2D slices for each image or interpolate slices to build a volume.

The structure of this thesis follows the considered steps of the pipeline in Fig. 1.6 in a bottom-up approach — starting with simpler, more specific problems, towards the more complex and general ones, as follows.

Chapter 2 presents background information on concepts explored in this work, such as brain anatomy concepts, imaging physics, and typical MRI preprocessing

Figure 1.6: General pipeline considered in this thesis to explore novel unsupervised brain anomaly detection approaches.

operations. Finally, the chapter also introduces the Image Forest Transform framework [60], as well as two algorithms derived from it, which serves as a basis for the design of some image operators used by the proposed solutions of this thesis.

Chapter 3 presents our solution for brain image segmentation. Its goal is to define our target macro-regions of interest — *i.e.*, right and left hemispheres, cerebellum, and brainstem — to improve the preprocessing, restrict the analysis, and compute hemispheric asymmetries in some cases. We start by exploring lesions associated with *abnormal hemispheric asymmetries* as detailed next in **Chapters 4** and **5**, as follows.

Chapter 4 proposes an automatic method for the detection of abnormal hippocampi from abnormal asymmetries. Our solution uses deep generative networks and a one-class classifier to model normal hippocampal asymmetries from healthy subjects and detect abnormal hippocampi. This is the first example of the usage of one-class classifiers for addressing the research questions of the thesis.

Chapter 5 presents a more generic solution that refines the proposal in **Chapter 4** to detect abnormal asymmetries in the entire brain hemispheres. Our approach extracts pairs of symmetric regions — called *supervoxels* — in both hemispheres of a test image under study. One-class classifiers then analyze the asym-

metries present in each pair. This method is limited to detect asymmetric lesions only in the hemispheres.

In [Chapter 6](#), we extend the previous solution from [Chapter 5](#) to detect lesions (symmetric or asymmetric) in the hemispheres, cerebellum, and brainstem. This new approach replaces asymmetries with any other *saliency map* that emphasizes brain anomalies. As proof of concept, we instantiated this solution with image registration errors to detect anomalies.

Finally, [Chapter 7](#) presents a compilation of our contributions and experimental findings, along with future research perspectives.

This chapter provides an overview of the basic concepts and techniques used in the next chapters. The chapter is targeted to a non-expert audience since it presents many basic and well-established topics on medical image analysis. Also, as the coming chapters detail the related work regarding their proposed methods, experienced readers are encouraged to skip this one and refer back whenever needed.

In [Section 2.1](#), we detail basic concepts about brain anatomy. [Section 2.2](#) provides an overview of medical imaging physics as well as which standards we adopted in this thesis. [Section 2.3](#) details the main preprocessing techniques used in MR image analysis.

[Section 2.4](#) introduces Image Foresting Transform (IFT) [60], a powerful methodology for the design of image operators based on optimum connectivity. IFT serves as the basis for the development of several algorithms used by the proposed solutions of this thesis, such as object delineation ([Section 3.2.2](#)), one-class classification ([Section 4.2.4](#)), and supervoxel segmentation ([Sections 5.2.3](#) and [6.1.3](#)). For better understanding the fundamentals of such algorithms, [Section 2.5](#) presents a clustering method derived from IFT, whereas [Section 2.6](#) details the Iterative Spanning Forest [61], a framework for superpixel segmentation also based on IFT. [Section 2.7](#) presents concluding remarks.

Appendices provide supplementary information to the main thesis as follows. [Appendix A](#) presents a quick reference about notations and definitions of terms used in this thesis.

To answer our research questions, we need datasets with isotropic 3D MR-T1 brain images from (i) healthy subjects, and (ii) with asymmetric anomalies of different sizes (especially small ones) and their gold-standard segmentation masks. As such, [Appendix B](#) presents a full description of all datasets used in the next chapters.

Finally, [Appendix C](#) describes all quantitative metrics adopted in this thesis to measure the accuracy and quality of our proposed solutions.

2.1 BASIC ANATOMICAL CONCEPTS

This section summarizes the main concepts related to brain anatomy. For a complete reference of the former, we recommend the books of Tortora and Derrickson [1], and Saladin [3]. More details about the latter can be found in the works of Hugdahl and Westerhausen [21], and Ocklenburg and Güntürkün [16].

2.1.1 Brain Anatomy

The *nervous system* is one of the most complex parts of the human body, yet its weight is equivalent to only 3% of the total body weight on average [1]. It is formed by a collection of specialized nerves and cells (neurons) that transmit signals to and from different parts of the body [1–3]. It acts as a communication network of the body that captures and interprets environmental stimuli, elaborating responses which may be converted, for example, in movements, sensations, and findings.

Structurally, the nervous system is organized in two main subdivisions: the *central nervous system* (CNS) and the *peripheral nervous system* (PNS), as shown in Fig. 2.1. The CNS consists of the brain and spinal cord. It processes different kinds of incoming sensory information, being responsible for all cognitive and affective capacities of humans. PNS, in turn, contains all the nerves that lie outside the CNS. Its leading role is to connect CNS to the organs, limbs, and skin so that CNS can receive and send information to these areas of the body [1, 2].

Figure 2.1: A simple diagram of the nervous system.

The *brain* is the interpreter of internal and external stimuli, containing about 85 billion neurons in an adult human [1]. Analogously, it is like the central processing unit (CPU) of a computer: it first receives and interprets different input information from our senses and internal organs and then provides appropriated responses. Thus, the brain provides control over body movement and regulates the operation of internal organs [1, 2, 14].

The *spinal cord*, in turn, is a long and fragile tube-like structure that is connected to the brain and extends down to the bottom of the spine. With about 100 million neurons, it is like a highway that carries incoming and outgoing messages between the brain and the rest of the body.

The brain consists of the cerebrum, cerebellum, and brainstem (Fig. 2.2). The *cerebrum* is the largest and uppermost portion of the brain. It contains two anatomically symmetrical hemispheres with several subcortical structures (e.g., hippocampus) [1]. The hemispheres are connected by a white matter structure called the *corpus callosum*. The cerebrum has an irregular appearance primarily due to gyri (elevations or ridges) and sulci (grooves or depressions).

Figure 2.2: (a) Brain regions and some of their corresponding responsibilities.¹ The four lobes from the hemispheres form the *cerebrum*. (b) Different axial slices of an MR-T1 image with the brain regions.²

The brain hemispheres consist of an inner core of nerve fibers called white matter and an outer cortex of gray matter. Each hemisphere can be divided into four lobes, as presented in Fig. 2.2. The *frontal lobe* is responsible for cognitive functions and the control of voluntary movements [14]. The *temporal lobe* is the location of

¹ Figure adapted from <http://picassowrites.blogspot.com/2019/03/any-exercise-great-for-aging-brain.html>.

² Figure adapted from https://commons.wikimedia.org/wiki/File:Brain_regions_on_T1-MRI.png.

the primary auditory cortex. It is the region where the sound is processed and where language and speech comprehension systems are located [62]. It is also involved with memory and emotion associations [63]. A deep lateral fissure called *Sylvian fissure* separates the temporal lobe from the parietal and frontal lobes (see the purple region in Fig. 2.2b). The *parietal lobe* is associated with linguistic and visuospatial functions. It helps to process the sense of touch and pain [14]. Finally, the *occipital lobe* is responsible for vision since the primary visual cortex is located within it.

The *cerebellum* is the second largest structure of the brain, located behind the temporal and occipital lobes [1]. It has an irregular and highly folded surface similar to the cerebrum. It plays a significant role in movement and acts in cognition and language processing [14]. Lastly, the *brainstem* connects the brain to the spinal cord and the rest of the body [14]. It receives and controls certain functions related to attention, temperature, heart rate, and breathing.

In this thesis, we focus on detecting lesions in structures inside the brain hemispheres, cerebellum, and brainstem. Chapter 3 details our approach for brain image segmentation.

2.1.2 Anatomical Planes of Body

To understand and describe the spatial organization of the body, we define positions and directions relative to standard anatomical planes and axes [64]. These planes are hypothetical geometric planes that divide the human body into sections. In human and animal anatomy, the body (or an organ) is sliced up using three planes: axial, coronal, and sagittal. In medical image analysis, a *slice* is a 2D image extracted from a 3D image along with one of these planes. Fig. 2.3 shows these planes for a brain.

For the sake of simplicity, suppose an upright subject. An *axial* plane divides the body into superior (upper) and inferior (lower) portions [1]. Such a plane is parallel to the floor and perpendicular to the long axis of the body. Slices are extracted from the feet to the head. When slicing the brain in this direction, we can see the left and right hemispheres (Fig. 2.3). This plane is also known as *transverse* or *horizontal* plane.

A *coronal* plane (also called *frontal* plane) divides the body into anterior (front) and posterior (back) portions [1]. Slices are extracted from the back to the front of the body. A coronal slice will show both brain hemispheres, like the axial slice.

Finally, a *sagittal* plane is a vertical plane that divides the body into right and left sides [1]. Indeed, slices are extracted from the right to the left side of the body. The *mid-sagittal* plane (MSP) is a plane that passes through the center of the body dividing it into approximately two symmetric parts [65] — see the coronal and axial MR slices in Fig. 2.3. Most structures on one side have a corresponding counterpart on the other side with similar shapes and relative locations [66]. Several applications, such as brain image registration [65, 67, 68] and, more importantly,

Figure 2.3: Anatomical planes of the brain.³ The dashed red line on the coronal and axial MR slice show their mid-sagittal planes.

brain asymmetry analysis [65, 66, 69] uses MSP. Likewise, some of our proposed methods will extensively use MSP as well. Section 2.3.2 provides a summary of automatic MSP extraction methods.

2.2 BASIC IMAGING PHYSICS

In this section, we present the main concepts of imaging physics and the standards used in this thesis. For a complete reference, we refer to the works of Runge *et al.* [10], Toennies [8], Larobina and Murino [70], and Brett *et al.* [71].

2.2.1 Medical Image Resolution

A medical image is a representation of some internal anatomical structures, or their functions, in the form of an array of picture elements called pixels for 2D and voxels for 3D. A *3D image* typically consists of a series of 2D images representing thin *slices* that form a volume (Section 1.1).⁴ It results from a sampling/reconstruction process that maps numerical values to voxels [8, 70]. For the sake of simplicity, let the term *image* be a 3D image and *slice* be a 2D image henceforth.

³ Figure adapted from https://www.wikiwand.com/en/Sagittal_plane.

⁴ It could also be a set of projections of an organ onto an image plane. Multiple acquisitions of the same volumetric image over time form a 4D medical image.

The smallest element of a slice is a *pixel*. It is defined by one or more *values* (also called *intensities*) and a *position* (2D coordinates; width and height) on the *image domain* [10]. It has dimensions along two axes in mm (e.g., a pixel size of $1 \times 1 \text{ mm}^2$). A *voxel*, in turn, is the volume element of an image. Its dimensions are given by the pixel and the *thickness* of the slice — i.e., the spacing/distance between two slices — which is measured along the third axis [10]. An image is *isotropic* when all its voxel dimensions are equal (e.g., a voxel size of $1 \times 1 \times 1 \text{ mm}^3$).

Voxel size is strongly related to spatial image resolution, which is an essential component of image quality. *Spatial image resolution* refers to the number of voxels in an image, or equivalently the number of pixels in a slice. The higher the number of voxels, the greater the resolution, and, consequently, the more detailed it is the image. Together with image contrast, spatial resolution determines the expert’s ability to distinguish one structure from others [72].

Altering voxel size impacts the spatial image resolution directly, as demonstrated in Fig. 2.4 that shows the same axial slice of an MR-T1 image from the same subject but acquired with different spatial resolution. For example, suppose an MRI scanner acquired a brain image by using a voxel size of $2 \times 2 \times 2 \text{ mm}^3$ and a given protocol. The resulting spatial resolution obtained was $128 \times 128 \times 128$ voxels. By appearance alone, the image is pixelated, grainy, and has jagged edges that make its analysis harder (Fig. 2.4a). In contrast, the same image was acquired with a smaller voxel size of $1 \times 1 \times 1 \text{ mm}^3$ in order to improve its quality. All other scanner parameters were the same. By halving the voxel size, the resulting image resolution doubled: $256 \times 256 \times 256$. Consequently, the image is sharper with improved anatomic details that considerably leverage its analysis (Fig. 2.4b). To achieve this higher quality, however, the imaging time approximately doubled.

A common practice in clinical routine to avoid long imaging times in MRI is to guarantee high-resolution for slices of a given direction (e.g., $1 \times 1 \text{ mm}^2$) but increase their thickness (e.g., 5 mm) [10]. The resulting number of slices can be considerably less depending on the chosen thickness. Such a practice results in two shortcomings: (i) small structures or lesions can be partially or even totally lost; and (ii) morphological measurements (e.g., volume) cannot be precisely computed. One might still interpolate slices to build a volume — as performed, for example, for the brain images from the popular BraTS dataset [73]. Nevertheless, this can create artifacts or textures that do not exist in the original image, impairing analysis. In this thesis, we only consider *isotropic* brain images for the development and evaluation of our methods. Appendix B details the considered brain image datasets.

2.2.2 Magnetic Field Strength

Field strength refers to the magnetic field strength used in the MRI scanner during image acquisition. Field strength is measured in teslas (T) and correlates image-quality factors [74], such as spatial-image resolution and artifacts. In general, a

Figure 2.4: Comparison between the same axial slice of an MR-T1 brain image with different spatial resolution. (a) Lower resolution: voxel size of $2 \times 2 \times 2 \text{ mm}^3$. (b) Higher resolution: voxel size of $1 \times 1 \times 1 \text{ mm}^3$. Highlighted regions indicate a lesion. The low-resolution slice is pixelated, grainy, and has jagged edges, whereas the high-resolution slice is sharper with improved anatomic details.

stronger field strength produces less noisy images with higher spatial resolution. Consequently, small and complex structures (e.g., hippocampus) are sharper, which makes their analysis more precise. However, some artifacts, like field inhomogeneity (Section 2.3.3), are more intense in high field strength.

Fig. 2.5 shows axial slices of MR-T1 brain images of 2T and 3T from different patients. Note that 2T images are noisier than 3T images, whereas field inhomogeneity is higher in 3T than 2T images. Brain structures are also sharper in 3T.

Figure 2.5: Comparison between axial slices from MR-T1 brain images of (a) 2T and (b) 3T.

In this thesis, we consider 3D MR-T1 brain images of 3T for the development and evaluation of most of the proposed methods. We only consider images of 2T during the evaluation of the automatic brain segmentation methods (Chapter 3).

2.2.3 Medical Image Orientation

MRI scanners can acquire thin slices at any angle or orientation within the body [10]. It is crucial to know the chosen orientation and coordinate system to interpret the voxels' positions in the image correctly. Although there is no single convention, some common concepts and terminologies are used to this end by popular medical image libraries [71, 75] and visualization tools [76, 77], as described below.

There are three conventional coordinate systems. The *world coordinate system* is the Cartesian coordinate system in which the subject is positioned. The *anatomical space* consists of the three planes that describe the standard anatomical position of a human (Section 2.1.2). The *image coordinate system* details how a medical image was acquired concerning the subject's anatomy and defines the voxels' coordinates. The conversion between the world and image coordinate systems commonly involves an affine transformation between both spaces.⁵

Suppose a subject is lying for a brain scan with his/her face up (Fig. 2.6). In this thesis, we consider that the world and image coordinate systems follow the *LPS+* orientation, that means:

- x-axis: from subject's right to Left;
- y-axis: from subject's anterior to Posterior; and
- z-axis: from subject's inferior to Superior.

Figure 2.6: Coordinate system with the *LPS+* orientation.

LPS+ is the usual convention for radiological visualization. The direction of the axes are given relative to the subject (e.g., “left” refers to the subject’s left). Each letter of the orientation reference is an abbreviation for the subject’s direction. The + symbol is a convention that defines which is the increasing direction along the

⁵ For more details, we refer to the manual of the NiBabel library [71] at https://nipy.org/nibabel/coordinate_systems.html.

corresponding axis. The considered origin for the image coordinate system — *i.e.*, the position of the voxel $(0,0,0)$ — is the upper-left corner toward the subject's feet (Fig. 2.6).

Regardless of how medical images are stored on disk, all images processed together must share the same coordinate system. Some medical image file formats,⁶ such as DICOM and Nifti, store the direction information that describes how the voxel data should be interpreted [10]. Consequently, one can reorient the images to be analyzed together to follow the same orientation. We reoriented all images used in this thesis to LPS+.

2.3 MRI PREPROCESSING

Automatic analysis of MR images is challenging due to typical acquisition artifacts — *e.g.*, noise, inhomogeneities, and variability of intensity and contrast — which negatively impact both medical diagnosis and automatic analysis. MRI preprocessing steps, in turn, aim to reduce these artifacts and, consequently, improve the image quality for subsequent analysis (Fig. 1.6).

In this section, we describe typical preprocessing steps applied to raw MR images [11, 40, 51, 78, 78–82] with a focus on the techniques used throughout this thesis. The combination of these steps is problem dependent and empirically estimated [82]. Fig. 2.7 presents the combination used in the next chapters. For a more detailed reference on MRI preprocessing, we refer to the book of Martí-Bonmatí and Alberich-Bayarri [81].

2.3.1 Noise Reduction

Even though significant improvements in imaging technology have been made in the past years, MR images are still prone to noise during acquisition [82–85]. Noise directly affects the accuracy of many automatic methods, such as segmentation, classification, and registration [83].

One strategy for *noise reduction*, also called *denoising*, consists of acquiring redundant images and averaging the outputs directly in the scanner. However, this option is uncommon in clinical routine since it increases the acquisition time significantly, which impacts the patient's comfort [83, 85]. Instead, filtering methods are the preferable alternatives in preprocessing pipelines [78, 82, 85].

Traditional denoising methods rely on *low-pass* filters to attenuate high-frequency signals in the image [83, 86]. One popular example is *median filtering*, which is effective at removing *salt-and-pepper* noise while preserving edges [87]. The filtered image is obtained by replacing each voxel with the median of all its neighboring voxels defined by an adjacency relation (*e.g.*, 26-neighborhood). Fig. 2.8 shows a noisy brain image and its filtered result by median filtering.

⁶ For a complete reference of medical image file formats, we refer to the work of Larobina and Murino [70].

Figure 2.7: General preprocessing steps for MR brain images. Native and Standard Image Space refer to, respectively, the coordinate space of the image being preprocessed and a given template.

Figure 2.8: (a) An axial slice of a noisy MR-T1 brain image and (b) its filtered result by median filtering.

Although median filtering is less effective for high levels of noise than modern and more complex denoising methods [88–90], it is suitable for the majority of the images used throughout this thesis. For simplicity, we do not consider the few extremely noisy images present in the datasets for the development and evaluation of our methods.

2.3.2 MSP Estimation

MRI scanners may produce tilted and misaligned brain images during acquisition due to factors as technicians’ inexperience, immobility of patients, and imprecise scanner calibration [65]. Tilt and misalignment may mislead visual inspection and affect the analysis of brain asymmetries since axial and coronal slices are no longer representing homologous structures [91].

Correcting the tilt of the head involves realigning the *mid-sagittal plane* (MSP) of the brain (Section 2.1.2) with the center of the image. A typical automatic approach for MSP estimation returns the plane that maximizes a given similarity measure between the two brain hemispheres [65, 66, 92, 93].

Ruppert *et al.* [66] propose a fast and accurate method that maximizes the symmetry between edges of the hemispheres in a multiscale-optimization scheme. The method starts by enhancing edges with the *Sobel operator* (Figs. 2.9a–b) followed by a thresholding step that selects the brightest edge voxels (Fig. 2.9c). The method then evaluates how symmetric the selected edges are with respect to each candidate plane. This measure is used to steer the position and orientation of the candidate MSP. This process is performed in a multiscale-optimization search to evaluate a high number of planes in interactive time. Each stage works on a different image scale, refining the solution of the previous stage. Fig. 2.9 presents some steps of this algorithm.

Figure 2.9: Steps for MSP definition by Ruppert *et al.* [66]. (a) An axial slice of a tilted MR-T1 brain image. (b) Enhanced edges by the Sobel operator. (c) Binary mask with the selected edge voxels for symmetry computation. (d) Resulting aligned image by its MSP (central white line).

We use the method of Ruppert *et al.* [66] to estimate the MSPs of the brain images considered in this thesis. We use MSPs to realign images before registration, and as a step to compute brain asymmetries for anomaly detection, as further detailed in Chapter 5.

2.3.3 Bias Field Correction

MR images are typically affected by *intensity inhomogeneities*, so-called *bias field*, which results from imperfections in the radio-frequency coils during image acquisition (Figs. 2.10a–b). This phenomenon makes the intensities of the same tissue vary in different locations within the image [94, 95] – *e.g.*, the intensity range of the white matter in one hemisphere is considerably different from the other hemisphere – or, more generally, there is a spatially-varying bias over the extent of the scanned brain. For instance, Fig. 2.10a shows such a case where the central brain area is overall brighter than the areas close to the cortex.

Figure 2.10: Example of bias field correction.

Most automatic analysis methods, such as segmentation and registration, assume that a given tissue presents similar voxel intensities throughout the image [95]. Thus, correcting inhomogeneities is crucial and must be performed before any quantitative MR analysis [94]. Fig. 2.10 presents an example of *bias field correction*.

Several methods have been proposed for bias field correction, typically for the analysis of MR brain images [96–99].⁷ Many approaches assume *a priori* knowledge about the image characteristics, such as the number of tissues or location, which make them limited to images from certain anatomical regions [100]. In contrast, the well-known and fully automatic N3 bias field correction [101] does not require any prior information about the MR data. Consequently, N3 is flexible and robust as it can accurately correct MR images of various anatomical structures from healthy subjects and patients. N3 interactively estimates the bias field by

⁷ We refer to the work of Vovk *et al.* [94] for a complete reference on bias field correction methods.

maximizing the high-frequency component of the image intensity distribution by using B-spline fitting [82, 94].

An optimized variant of N3, so-called N4 [102], proposes a faster B-spline approximation in a modified hierarchical optimization scheme. This variant is faster than N3 with similar accuracies. N4 is widely used in the literature and publicly available on the open-source ITK library [103]. We use N4 to correct intensity inhomogeneities in all images considered in this thesis.

2.3.4 Image Registration

Image registration is the task of establishing a spatial correspondence between images from the same context (e.g., MR images of the brain) by mapping them into the same coordinate system [104]. As such, the same anatomical structures present in these images will share an identical location in all images after registration. Typical MRI preprocessing pipelines use registration (i) to combine anatomical information from different imaging modalities — e.g., T1 and T2 images from the same patient — or (ii) to locate the images in a specific *standard space* to perform a population analysis [82].

The way image registration typically works is to deform an image, the *moving image* (Fig. 2.11a), to align with another one, the *fixed image* (Fig. 2.11b), also called *template* or *reference image*. A *cost function* — e.g., mutual information [105] (Appendix C.1) — assesses the quality of alignment, i.e., the similarity between the deformed moving image and the fixed image [36]. This cost function is iteratively optimized so that its best score results in the final registered image.⁸

Figure 2.11: Example of affine and non-rigid registrations. (a) Moving image. (b) Fixed image. (c) Registered images by affine registration. (d) Registered images by non-rigid registration.

⁸ We recommend the surveys of Litjens *et al.* [44], and Mani and Arivazhagan [106] for more details on medical image registration methods.

Image registration can rely on linear and non-linear transformations to deform a moving image. *Linear registration* applies the same transformation to every voxel in the moving image, whereas *non-linear registration* can apply a distinct transformation — so-called *deformation fields* — to different voxels [82]. The simplest example of the linear case is *rigid registration* that only uses rotations and translations [106]. *Affine registration* extends the degrees of freedom of the rigid approach by also considering scaling and shearing transformations. Finally, *non-rigid registration*, also known as *deformable registration*, relies on non-linear operations to align images where correspondences cannot be achieved without localized deformations. The choice of the registration approach is problem-dependent and usually considers constraints as alignment accuracy and processing time. Fig. 2.11 exemplifies the different types of image registration.

In this thesis, we consider affine and, most importantly, non-rigid registration for the development of the proposed image analysis methods. We use Elastix [107] — an open-source software widely used in the literature — to perform all registrations. Lastly, we consider the popular ICBM 2009c Non-linear Symmetric template [108] for our methods.

2.3.5 Skull Stripping

Most automatic MR brain image applications aim to analyze patterns in *only* brain tissues, *i.e.*, the ones in the hemispheres, cerebellum, and brainstem. In contrast, the inclusion of non-brain tissues (*e.g.*, skull, eyes, and neck) during analysis makes the processing time considerably slower, especially for 3D images. Besides, it can impair the resulting accuracies since these tissues have similar intensities to brain ones [109]. To circumvent these problems, most preprocessing pipelines rely on *skull-stripping* methods to extract the brain for subsequent analysis. The result can be either a new image with only brain voxels or a binary mask, which defines label 1 for brain voxels and label 0 for the remaining tissues [110].

A large number of methods proposed over the past years confirms the importance of skull-stripping [109–114]. However, these methods are not able to separate the right and left hemispheres, cerebellum, and brainstem — so-called *brain segmentation* — that allows a more specific analysis. For example, one can use segmented hemispheres to assess brain asymmetries [18], or as the first step for the segmentation of subcortical structures (*e.g.*, hippocampus) [80], or to compare morphological measures between the hemispheres [80]. Brain segmentation is more challenging than skull-stripping as the boundaries between the objects of interest are not well-defined on MR images, especially those between the hemispheres and cerebellum [41].

In this thesis, we propose a novel probabilistic-atlas-based method for the automatic segmentation of (ab)normal MR-T1 brain images (Chapter 3). Our method, named *AdaPro*, provides a *labeled image (segmentation mask)* that assigns a different label for each object of interest — *i.e.*, the right hemisphere, left hemisphere,

cerebellum, and brainstem. Different prior steps of the proposed anomaly detection methods use this segmentation mask: (i) for intensity normalization (Section 2.3.6); (ii) to warp both hemispheres for abnormal asymmetry detection on the native image space (Section 5.5); and (iii) to perform anomaly detection in each object of interest individually (Chapter 6). Chapter 3 provides a detailed review of brain image segmentation and presents AdaPro.

2.3.6 Intensity Normalization

In contrast with other medical imaging modalities, MR images even acquired with the same protocol and scanner typically do not share similar intensities [82]. This inter-image variability is problematic for automatic quantitative analysis because most methods expect that the intensity distribution of all considered images is the same. The previous methods for bias-field correction (Section 2.3.3) do not solve this problem, since they focus on correcting intra-image variability, *i.e.*, the unbalanced distribution of intensities, *e.g.*, from a given tissue across the image.

Intensity-normalization methods aim to correct the scanner-dependent variation by mapping intensities of all images into a standard reference [11]. Most traditional approaches rely on the use of the *histogram-matching* technique, which transforms the histogram from a source image to match a reference image's histogram [82, 115]. Recent methods incorporate *a priori* anatomical information by restricting the histogram matching to only some segmented brain tissues [11, 42, 109]. Consequently, non-brain tissues does not influence the quality of the final intensity normalization.

Inspired by the anatomical-based methods, we propose a novel intensity-normalization approach, as follows. Assume a source brain image I , a reference image R , and their corresponding brain segmentation masks after skull stripping (Section 2.3.5). First, we apply a linear intensity normalization into I by mapping all its intensities within $[0, 4095]$ (12-bits). We chose this intensity range since all priority (in-house) datasets were acquired within it. Finally, we apply a histogram matching between I and R by *only* considering the object brain voxels defined in their segmentation masks. We use our proposed automatic segmentation method, AdaPro, to obtain the brain segmentation masks from the source images (Chapter 3). Fig. 2.12 illustrates the proposed intensity-normalization approach.

2.4 IMAGE FORESTING TRANSFORM

The *Image Foresting Transform (IFT)* is a methodology for the design of image operators based on optimum connectivity [60]. For a given connectivity function and a graph derived from an image, the IFT algorithm minimizes (maximizes) a connectivity map to partition the graph into an *optimum-path forest* rooted at the minima (maxima) of the resulting connectivity map. The image operation resumes to a post-processing of the forest attributes, such as the root labels, optimum paths,

Figure 2.12: The proposed intensity normalization. (a) Segmented brain source image of a stroke patient. (b) Segmented reference image. (c) Resulting preprocessed image after intensity normalization.

and connectivity values. IFT has been successfully applied in different domains, such as image filtering [116], image descriptor [117, 118], segmentation [41, 119–123], superpixel segmentation [40, 61, 124], representation [125], (semi) supervised classification [126–128], and data clustering [129, 130].

In this thesis, we widely use IFT-based methods for brain image segmentation (Chapter 3), one-class classification (Chapter 4), and supervoxel segmentation (Chapters 5–6). This section presents preliminary concepts and introduces the IFT algorithm.

2.4.1 Preliminary Concepts

Image Graphs: A d -dimensional multi-band image is defined as the pair $\hat{I} = (D_I, \vec{I})$, where $D_I \subset Z^d$ is the *image domain* — i.e., a set of elements (pixels/voxels) in Z^d — and $\vec{I} : D_I \rightarrow \mathbb{R}^c$ is a mapping function that assigns a vector of c real-valued intensities $\vec{I}(p)$ — one value for each band (channel) of the image — to each element $p \in D_I$. For example, for 2D RGB-color images: $d = 2, c = 3$; for 3D grayscale images (e.g., MR images): $d = 3, c = 1$. We represent a *segmentation* of $\hat{I}$ by a label image $\hat{L} = (D_I, L)$, wherein the function $L : D_I \rightarrow \{0, 1, \dots, M\}$ maps every voxel of $\hat{I}$ to either the background (label 0) or one of the M objects of interest.

Most images, like the ones used in this thesis, typically represent their intensity values by *natural numbers* instead of *real numbers*. More specifically, $\vec{I} : D_I \rightarrow [0, 2^b - 1]$, where b is the *number of bits* (pixel/voxel depth) used to encode an intensity value.

An image can be interpreted as a graph $G_I = (D_I, \mathcal{A})$, whose *nodes* are the voxels and the *arcs* are defined by an *adjacency relation* $\mathcal{A} \subset D_I \times D_I$, with $\mathcal{A}(p)$ being the *adjacent set* of a voxel p . A *spherical adjacency relation* of radius $\gamma \geq 1$ is given by

$$\mathcal{A}_\gamma : \{(p, q) \in D_I \times D_I, \|q - p\| \leq \gamma\}. \quad (2.1)$$

The image operators considered in this thesis use two types of adjacency relations: $\mathcal{A}_1$ (6-neighborhood) and $\mathcal{A}_{\sqrt{3}}$ (26-neighborhood), as illustrated in Fig. 2.13.

Figure 2.13: Examples of adjacency relation for a given voxel p (red).

Paths: For a given image graph $G_I = (D_I, \mathcal{A})$, a *path* π_q with terminus q is a sequence of distinct nodes $\langle p_1, p_2, \dots, p_k \rangle$ with $\langle p_i, p_{i+1} \rangle \in \mathcal{A}$, $1 \leq i \leq k-1$, and $p_k = q$. The path $\pi_q = \langle q \rangle$ is called *trivial path*. The concatenation of a path π_p and an arc $\langle p, q \rangle$ is denoted by $\pi_p \cdot \langle p, q \rangle$.

Connectivity Function: A *connectivity function* (path-cost function) assigns a value $f(\pi_q)$ to any path π_q in the image graph $G_I = (D_I, \mathcal{A})$. A path π_q^* ending at q is optimum if $f(\pi_q^*) \leq f(\tau_q)$ for every other path τ_q . In other words, a path ending at q is optimum if no other path ending at q has lower cost.

Connectivity functions may be defined in different ways. In some cases, they do not guarantee the optimum cost mapping conditions [131], but, in turn, can produce effective object delineation [132]. In this thesis, we explore the *max-arc* path-cost function f_{max} , defined by

$$f_{max}(\langle q \rangle) = \begin{cases} 0 & \text{if } q \in \mathcal{S}, \\ +\infty & \text{otherwise.} \end{cases} \quad (2.2)$$

$$f_{max}(\pi_p \cdot \langle p, q \rangle) = \max\{f_{max}(\pi_p), w(p, q)\}$$

where $w(p, q)$ is the arc weight of $\langle p, q \rangle$, usually estimated from $\hat{I}$, and $\mathcal{S}$ is the labeled seed set.

2.4.2 The General IFT Algorithm

For multi-object image segmentation, IFT requires a labeled seed set $\mathcal{S} = \mathcal{S}_0 \cup \mathcal{S}_1 \cup \dots \cup \mathcal{S}_M$ with seeds for object i in each set $\mathcal{S}_i$ and background seeds in $\mathcal{S}_0$ (Fig. 2.14a). The algorithm then promotes an optimum seed competition so that each seed in $\mathcal{S}$ conquers its most closely connected voxels in the image domain. This competition considers a connectivity function f applied to any path π_q . In the case of a *seeded watershed transform* [133], as also adopted in this thesis, arc weights correspond to gradient image values of $\hat{I}$ (Fig. 2.14b).

Figure 2.14: Multi-object image segmentation by IFT. (a) Axial slice of a brain image with seeds $\mathcal{S}_0$ for the background (orange), $\mathcal{S}_1$ for the right ventricle (red), and $\mathcal{S}_2$ for the left ventricle (green). (b) Gradient image for (a) that defines the arc weights for seed competition. Arcs have high weights on object boundaries. (c) Resulting segmentation mask for the given seeds and arc weights. Red and green voxels represent object voxels, whereas the remaining ones are background.

Defining Π_q as the set of all possible paths with terminus q in the image graph, the IFT algorithm minimizes a path cost map

$$C(q) = \min_{\forall \pi_q \in \Pi_q} \{f(\pi_q)\}, \quad (2.3)$$

by partitioning the graph into an optimum-path forest P rooted at $\mathcal{S}$. That is, the algorithm assigns to q the path π_q^* of minimum cost, such that each object i is defined by the union between the seed voxels of $\mathcal{S}_i$ and the voxels of D_i that are rooted in $\mathcal{S}_i$, *i.e.*, conquered by such object seeds (Fig. 2.14c).

Algorithm 1 presents the general IFT approach. Lines 1–7 initialize maps, and insert seeds into the priority queue Q . The state map U indicates by $U(q) = \text{White}$ that the voxel q was never visited (never inserted into Q), by $U(q) = \text{Gray}$ that q has been visited and is still in Q , and by $U(q) = \text{Black}$ that q has been processed (removed from Q).

The main loop (Lines 8–20) performs the propagation process. First, it removes the voxel p that has *minimum path cost* in Q (Line 9). Ties are broken in Q using

Algorithm 1: The General IFT Algorithm

Input: Image $\hat{I} = (D_I, I)$, adjacency relation $\mathcal{A}$ connectivity function f , and seed set $\mathcal{S} \subset D_I$ labeled by λ .

Output: Optimum-path forest P , root map R , path-cost map C , and label map L .

Aux: Priority queue Q , state map U , and variable tmp .

```

1 foreach  $q \in D_I$  do
2    $P(q) \leftarrow \emptyset, R(q) \leftarrow q$ 
3    $C(q) \leftarrow f(\langle q \rangle), L(q) \leftarrow 0$ 
4    $U(q) \leftarrow White$ 
5   if  $q \in \mathcal{S}$  then
6     insert  $q$  into  $Q$ 
7      $L(q) \leftarrow \lambda(q), U(q) \leftarrow Gray$ 
8 while  $Q \neq \emptyset$  do
9   Remove  $p$  from  $Q$  such that  $C(p)$  is minimum
10   $U(p) \leftarrow Black$ 
11  foreach  $q \in \mathcal{A}(p)$  such that  $U(q) \neq Black$  do
12     $tmp \leftarrow f(\pi_p^* \cdot \langle p, q \rangle)$ 
13    if  $tmp < C(q)$  then
14       $P(q) \leftarrow p, R(q) \leftarrow R(p)$ 
15       $C(q) \leftarrow tmp, L(q) \leftarrow L(p)$ 
16      if  $U(q) = Gray$  then
17        update position of  $q$  in  $Q$ 
18      else
19        insert  $q$  into  $Q$ 
20         $U(q) \leftarrow Gray$ 
21 return  $(P, R, C, L)$ 

```

the first-in-first-out (FIFO) policy. The loop in Lines 11–20 then evaluates if a path with terminus p extended to its adjacent q is cheaper than the current path with terminus q and cost $C(q)$ (Line 13). If that is the case, p is assigned as the predecessor of q , and the root of p is assigned to the root of q (Line 14), whereas the path cost and the label of q are updated (Line 15). If q is in Q , its position is updated; otherwise, q is inserted into Q . The algorithm returns the optimum-path forest (predecessor map), root map, path-cost map, and the label map (object delineation mask). Fig. 2.15 illustrates the execution of the IFT algorithm with f_{max} .

Figure 2.15: Example of the IFT seed competition with f_{max} . **(a)** A 4-adjacent image graph with numbers indicating arc weights based on image gradient, and squares representing the seeds for two objects (red and yellow). **(b)–(h)** Iteration steps of IFT. Numbers inside circles and squares indicate path-cost values. The resulting optimum-path forest (OPF) is shown in **(h)**.

2.5 CLUSTERING BY OPTIMUM-PATH FOREST

This section presents OPF-clustering [129], a data-clustering algorithm based on optimum-path forests. OPF-clustering extends the IFT framework (Section 2.4) from the image domain to the feature space by interpreting a training set as a graph whose nodes are the samples, and their arcs are defined by an adjacency relation. The nodes are weighted by their probability density values (pdf), and a connectivity function is maximized, such that each local maximum of the pdf becomes the root of an optimum-path tree (cluster), composed by samples “more strongly connected” to that local maximum than to any other root. We use OPF-clustering for the development of our proposed one-class classifier, as presented in Section 4.2.4. The theoretical background and algorithm of OPF-clustering are presented next.

Let $\mathcal{Z}$ be a training set, and $s \in \mathcal{Z}$ a given training sample. A random choice of samples to compose $\mathcal{Z}$ makes the observations $\mathbf{x} = \mathbf{v}(s) \in \mathbb{R}^n$ a random field, whose probability density function (pdf) $\rho(\mathbf{x})$ can be estimated as

$$\rho(\mathbf{x}) = \frac{\sum_{\forall (s,t) \in \mathcal{A} | \mathbf{v}(s)=\mathbf{x}} \exp \left[\frac{-d^2(s,t)}{2\sigma^2} \right]}{\sum_{\forall s \in \mathcal{Z} | \mathbf{v}(s)=\mathbf{y}} \rho(\mathbf{y})}, \quad (2.4)$$

where $\sigma > 0$, $d(s, t) = \|\mathbf{v}(t) - \mathbf{v}(s)\|$, and $\mathcal{A}$ is an *adjacency relation* defined in the feature space by

$$\begin{aligned} \mathcal{A}: \{ (s, t) \in \mathcal{Z} \times \mathcal{Z} \mid s \neq t, \\ t \text{ is } k \geq 1 \text{ nearest neighbor of } s \text{ in } \mathbb{R}^n \}. \end{aligned} \quad (2.5)$$

A clustering in $\mathcal{Z}$ can be obtained by separating the domes of the pdf, such that the samples in each dome compose one cluster. The authors in [129] formulated this problem as an *optimum-path forest* problem in a graph. Their method selects one *root sample* per local maximum of the pdf such that each remaining sample is assigned to the cluster of the root that connects to it by an optimum path. A path is optimum in the sense that the minimum density value along it is maximum concerning the values of other paths to the same node. In order to guarantee that the roots will reach the remaining samples in the same dome, they first consider the extended adjacency relation $\mathcal{A}_e$, which is symmetric on the plateaus of the pdf.

$$\begin{aligned} \mathcal{A}_e: \{ (s, t) \in \mathcal{A} \cup \{ (t, s) \} \mid (t, s) \notin \mathcal{A} \\ \text{and } \rho(s) = \rho(t) \}, \end{aligned} \quad (2.6)$$

where $\rho(s) = \rho(\mathbf{v}(s)) = \rho(\mathbf{x})$. The parameter $\sigma = \max_{\forall (s, t) \in \mathcal{A}_e} \{d(s, t)\}$ is fixed and the graph $(\mathcal{Z}, \mathcal{A}_e)$ is weighted on the nodes $s \in \mathcal{Z}$ by $\rho(s)$ and on the arcs $(s, t) \in \mathcal{A}_e$ by $d(s, t)$. A path π_t with terminus t is a sequence $\langle s_1, s_2, \dots, s_m = t \rangle$ of nodes, such that $(s_i, s_{i+1}) \in \mathcal{A}_e$, $i \in [1, m]$, $\pi_t = \langle t \rangle$ is said a trivial path, and $\pi_t = \pi_s \cdot \langle s, t \rangle$ is the concatenation of π_s and $\langle s, t \rangle$ with the two joining instances of s merged into one. The path-value function f of minimum density is defined as

$$\begin{aligned} f(\langle t \rangle) &= \begin{cases} \rho(t) & \text{if } t \in \mathcal{R}, \\ \rho(t) - \delta & \text{otherwise,} \end{cases} \\ f(\pi_s \cdot \langle s, t \rangle) &= \min\{f(\pi_s), \rho(t)\}, \end{aligned} \quad (2.7)$$

where $\delta = \min_{\forall (s, t) \in \mathcal{A}_e} \{|\rho(t) - \rho(s)|\}$ and $\mathcal{R}$ is a root set with one sample per maximum of the pdf, as selected during the algorithm. The optimum-path forest algorithm has been first presented in [60] (Section 2.4), and the sufficient conditions for its correctness are established with proof in [131]. It can maximize a path-value map $V(t) = \max_{\forall \pi_t \in \Pi} \{f(\pi_t)\}$, where Π is the set of all paths in the graph, by partitioning the graph into an optimum-path forest P – an acyclic map that assigns to each node $t \in \mathcal{Z}$ its predecessor $P(t) = s \in \mathcal{Z}$ in the optimum path with terminus t or a marker $P(t) = \text{nil} \notin \mathcal{Z}$, when the node $t \in \mathcal{R}$ is a root of the map. Once the optimum-path forest is defined, new samples $t \notin \mathcal{Z}$ can be assigned to one of the obtained clusters by evaluating the values of the extended paths $\pi_s \cdot \langle s, t \rangle$, $\forall s \in \mathcal{Z}$. In [134], the authors simplify this process for the sake of efficiency by considering an adjacency radius $\Omega(s) = \max_{\forall t \in \mathcal{A}(s)} \{d(s, t)\}$, where $\mathcal{A}(s)$ is the set

of the k nearest neighbors of s , and giving a sample priority for $s \in \mathcal{Z}$, to conquer new samples t , proportional to its optimum path value $V(s)$. That is, t is assigned to the cluster of the root $R(s^*)$ of s^* whose

$$\rho(s^*) = \max_{\forall s \in \mathcal{L} | d(s, t) \leq \Omega(s)} \{\rho(s)\}, \quad (2.8)$$

where $\mathcal{L}$ is that priority list. When t does not satisfy the condition $d(s, t) \leq \Omega(s)$ for any $s \in \mathcal{L}$, t is assigned to the cluster of its closest sample in $\mathcal{L}$.

Algorithm 2: Clustering by Optimum Path Forest

Input: Graph $(\mathcal{Z}, \mathcal{A}_e)$ and the pdf ρ .

Output: Root map R and sorted list $\mathcal{L}$.

Aux: Path-value map V , a binary heap Q , and variable tmp .

```

1 foreach  $s \in \mathcal{Z}$  do
2    $R(s) \leftarrow s$ 
3    $V(s) \leftarrow \rho(s) - \delta$ 
4   insert  $s$  in  $Q$ 
5 while  $Q$  is not empty do
6   Remove  $s$  from  $Q$  such that  $s = \operatorname{argmax}_{t \in Q} \{V(t)\}$ 
7   Insert  $s$  in  $\mathcal{L}$ 
8   if  $R(s) = s$  then
9      $V(s) \leftarrow \rho(s)$ 
10  foreach  $t \in \mathcal{A}_e(s)$  and  $V(t) < V(s)$  do
11    Compute  $tmp \leftarrow \min\{V(s), \rho(t)\}$ 
12    if  $tmp > V(t)$  then
13       $R(t) \leftarrow R(s)$ 
14       $V(t) \leftarrow tmp$ 
    
```

Algorithm 2 presents the clustering by Optimum Path Forest. The process starts by defining all nodes as trivial paths $\langle s \rangle$ with values $f(\langle s \rangle) = \rho(s) - \delta$ (Lines 1–4, Eq. 2.7). In the main loop (Lines 5–14), the nodes are removed from Q in their non-increasing order of path value. When the first node of a pdf maximum is removed from Q , Line 9 updates its root value according to Equation 2.7. In the internal loop (Lines 10–14), the roots $R(s)$ conquer the remaining nodes t of the same plateau and dome of the pdf whenever the value $f(\pi_s \cdot \langle s, t \rangle)$ (Line 11, Eq. 2.7) is higher than the value $V(t)$ (Line 12) of the current path π_t , assigning t to the same cluster of s (Lines 13–14). In the end, only the roots of the forest have $R(s) = s$.

A last point for discussion is the choice of the parameter k . As suggested in [129], we find $k \in [1, k_{\max}]$ as the one that produces a minimum normalized cut $\text{cut}(k)$ in $(\mathcal{Z}, \mathcal{A}_e)$.

$$\begin{aligned} \text{cut}(k) &= \sum_{r \in \mathcal{R}} \frac{W'_r}{W_r + W'_r}, \\ W_r &= \sum_{(s,t) \in \mathcal{A}_e | R(s)=R(t)=r} \frac{1}{d(s,t)}, \\ W'_r &= \sum_{(s,t) \in \mathcal{A}_e | R(s)=r \neq R(t)} \frac{1}{d(s,t)}. \end{aligned} \tag{2.9}$$

The upper limit $k_{\max}$ is an application-dependent parameter — larger it is, fewer clusters are obtained. Therefore, the only parameters are the size of the removed trees and $k_{\max}$. In this thesis, we fixed $k_{\max}$ equal to 15% of the number of training samples and eliminated trees with less than five nodes from the training set.

Section 4.2.4 presents our proposed one-class classifier that extends OPF-clustering for anomaly detection.

2.6 ITERATIVE SPANNING FOREST (ISF)

A crucial step of the considered pipeline towards answering our research questions (Fig. 1.6) consists of selecting volumes of interest (VOIs) for the subsequent analysis. A strategy for VOI selection, which we explore throughout this thesis (Chapters 5 and 6), is *supervoxel segmentation*.

Supervoxels are groups of connected voxels with similar characteristics resulting from an oversegmentation of a 3D image or region of interest. Similarly, the term *superpixels* is used for 2D images. They preserve essential image information (e.g., the borders of tissues and lesions) and are used as an alternative to patches to define more meaningful VOIs for computer-vision problems [135, 136] and some medical image applications [6, 137]. For example, one can oversegment the hemispheres in multiple supervoxels for brain anomaly detection.

In this thesis, we consider the Iterative Spanning Forest (ISF) framework [61] for supervoxel segmentation. ISF is a recent approach for both superpixel and supervoxel segmentation that has shown higher effectiveness than several state-of-the-art counterparts, especially when used for 3D MR image segmentation of the brain [61]. ISF consists of three key steps: (i) seed sampling followed by multiple iterations of (ii) connected supervoxel delineation based on IFT [60] (Section 2.4), and (iii) seed recomputation to improve delineation. We next present the theoretical background for ISF as well as its algorithm. For the sake of clarity, we use the same nomenclature presented in Section 2.4.

2.6.1 Theoretical Background

Recall the pair $\hat{I} = (D_I, \vec{I})$ be a d -dimensional multi-band image, where $D_I \subset \mathbb{Z}^d$ is the *image domain*, and $\vec{I} : D_I \rightarrow \mathbb{R}^c$ is a mapping function that assigns a vector of c real-valued intensities $\vec{I}(p)$ — one value for each band (channel) of the image — to each element $p \in D_I$. For simplicity, assume that the term *voxels* represents the d -dimensional-image elements.

As outlined in [Section 2.4](#), an image can be interpreted as a graph $G_I = (D_I, \mathcal{A})$, whose *nodes* are the voxels, and the *arcs* are defined by an *adjacency relation* $\mathcal{A} \subset D_I \times D_I$, with $\mathcal{A}(p)$ being the *adjacent set* of a voxel p . In this work, we consider the 6-neighborhood adjacency for ISF ([Fig. 2.13a](#)).

For a given *initial seed set* $\mathcal{S}$, labeled with consecutive integer numbers $\{1, 2, \dots, |\mathcal{S}|\}$, and a *connectivity function* f , ISF computes each supervoxel as a *spanning tree* rooted at a *single seed*, such that the seeds compete among themselves by offering lower-cost paths to conquer their most strongly connected voxels. We use the following connectivity function f given by

$$f(\langle q \rangle) = \begin{cases} 0, & \text{if } q \in \mathcal{S}, \\ +\infty, & \text{otherwise,} \end{cases} \quad (2.10)$$

$$f(\pi_p \cdot \langle p, q \rangle) = f(\pi_p) + \left[\alpha \cdot \|\vec{I}(q) - \vec{I}(R(p))\| \right]^\beta + \|q - p\|,$$

where $\|\vec{I}(q) - \vec{I}(R(p))\|$ is the Euclidean distance between the intensity vectors at voxels $R(p)$ and q , $\|q - p\|$ the Euclidean distance between the voxels p and q , $\langle q \rangle$ is a trivial path, $\pi_p \cdot \langle p, q \rangle$ the extension of a path π_p with terminus q by an arc $\langle p, q \rangle$, and $R(p)$ the starting node (seed) of π_p . The factors α and β serve to control a compromise between supervoxel boundary adherence and shape regularity. Although the authors of ISF have fixed $\alpha = 0.5$ and $\beta = 12$ during the experiments [[61](#)], such factors are problem-dependent and should be optimized to yield more accurate supervoxels. [Fig. 2.16](#) shows the impact of α and β for the superpixel segmentation of a 2D brain image.

2.6.2 The ISF Algorithm

[Algorithm 3](#) presents a pseudo code for the Iterative Spanning Forest framework. At each iteration (Lines 2–4), ISF performs connected supervoxel delineation on the image I based on IFT (Line 3) — as described by [Algorithm 1](#) — from a given seed set $\mathcal{S}'$, adjacency relation $\mathcal{A}$, and the connectivity function f described by [Equation 2.10](#). The seed set at Iteration 1 is the initial seed set $\mathcal{S}$ (Line 1). Next, the seed set is recomputed by the function *SeedRecomputation* to improve delineation (Line 4). This process continues until reaching N iterations. The algorithm

Figure 2.16: The impact of the factors α and β for superpixel segmentation by ISF. Each superpixel is represented by a different color. For all cases, we performed ISF on the same 2D brain image with 10 iterations and identical 30 initial seeds selected by grid sampling.

Algorithm 3: Iterative Spanning Forest

Input: Image $\hat{I} = (D_I, \vec{I})$, adjacency relation $\mathcal{A}$, connectivity function f , initial seed set $\mathcal{S} \subset D_I$, and the maximum number of iterations $N \geq 1$.

Output: Optimum-path forest P , root map R , path-cost map C , and supervoxel label map L .

Aux: Seed set $\mathcal{S}'$, and the variable i .

```

1  $\mathcal{S}' \leftarrow \mathcal{S}$ 
2 for  $i \leftarrow 0$  to  $N - 1$  do
3    $(P, R, C, L) \leftarrow IFT(\hat{I}, \mathcal{A}, f, \mathcal{S}')$ 
4    $\mathcal{S}' \leftarrow \text{SeedRecomputation}(\hat{I}, \mathcal{S}', P, R, C)$ 
5 return  $(P, R, C, L)$ 

```

returns the optimum-path forest (predecessor map), root map, path-cost map, and the supervoxel label map. Fig. 2.17 illustrates the execution of ISF.

Figure 2.17: Example of the ISF execution (10 iterations with $\alpha = 0.5$ and $\beta = 12$) on a 2D brain image. (a)–(d) For iterations of ISF. For each iteration, we show its input seeds (red points) and the resulting obtained superpixels (each color represents a different superpixel). Iteration 0 shows the initial seed set obtained by grid sampling; the other seed sets are obtained by seed recomputation. As the insets show, most seeds do not change positions over iterations.

In this work, we adopted a seed-recomputation strategy proposed by Vargas-Muñoz *et al.* [61], as detailed next. At each iteration, we promote the *centroids* from the obtained supervoxels — *i.e.*, their geometric centers — to be the seeds of the next iteration. If a given centroid c_i is out of its supervoxel L_i — due to the non-convex shape of L_i — we select the voxel of L_i that is the closest to c_i . We refer to Vargas-Muñoz *et al.* [61] for more specific details.

A crucial step for the success of ISF consists of performing a robust *initial seed estimation*. This step, however, is problem-dependent, so that simple and general

strategies — *e.g.*, a grid sampling in the input image — can provide unsatisfactory results (*e.g.*, undersegmentating a lesion). [Chapters 5 and 6](#) introduce two ISF-based algorithms for supervoxel segmentation tailored to specific constraints. Both algorithms present different strategies to build a 2-band volumetric image from a 3D MR image, as well as robust initial seed estimations that lead to better final supervoxels (*e.g.*, a supervoxel that correctly fits a lesion).

2.7 CONCLUSION

In this chapter, we have summarized the main background information on concepts explored in this work, which involves basic brain anatomy, imaging physics, and typical MRI preprocessing operations. Besides, we introduced image foresting transform, a robust framework widely used for the development of several image operators throughout this thesis.

From the reviewed material, we can conclude that the pipeline of medical image analysis for brain data is complex, having many steps of various natures, and primarily demanding concepts on anatomy and imaging physics. Each of these steps is crucial, in its own way, to provide a good-quality final result, *i.e.*, a good anomaly detection, towards our research questions. In particular, we highlight the importance of the presented MRI preprocessing operations ([Section 2.3](#)) for the success of such a pipeline since MRI data typically presents acquisition artifacts of different characteristics — *e.g.*, noise, inhomogeneities, and variability of intensity and contrast — which negatively impact both medical diagnosis and automatic analysis.

The rest of this thesis is dedicated to improving the various steps of the considered medical image analysis pipeline.

The precise delineation of 3D objects (*e.g.*, organs and lesions) in 3D MR brain images has been an active research topic over the last decades [138, 139]. A common example is *brain segmentation* that aims to separate the so-called *macro brain structures*: the right and left hemispheres, cerebellum, and brainstem. Brain segmentation is typically used for a better understanding of neurological diseases, the study of brain asymmetries [18], morphological analysis of the hemispheres [80], surgical planning [29], and the development of computer-aided diagnosis systems [140]. Regarding our research questions, brain segmentation supports the proposed unsupervised anomaly detection approaches in different steps: (i) for intensity normalization (Section 2.3.6); (ii) to warp both hemispheres for anomaly detection on the native image space (Section 5.5); and (iii) to perform anomaly detection in each object of interest individually (Chapter 6).

The absence of well-defined boundaries between the macro brain structures in MR images makes brain segmentation challenging. Fully interactive segmentation methods require a high number of user intervention, becoming a tedious, time-consuming, and error-prone task, especially for studies involving large datasets [141]. Moreover, such methods require specialists with considerable experience in manual delineation. The use of prior anatomical information (*shape constraints*) can either mitigate the problem or eliminate the need for user interaction.

Atlases: *Probabilistic atlases* (PAs), also known as *statistical object shape models*, are popular and well-succeeded examples of shape constraints for *automatic* brain segmentation [41, 120, 141–144]. Methods based on PAs estimate the probability of a voxel to be part of a given object regardless of its intensity in the original image [41, 120, 143]. Some techniques, such as SOSM-S [120], combine these probabilities with a delineation algorithm to obtain better-refined object boundaries. However, the existing models do not adapt to possible object anomalies caused by the presence of a disease or a surgical procedure (Fig. 3.1). Such anomalies often alter the shape and texture of the brain structures, making them different

This chapter is based on the following publications:

- (i) A. X. Falcão, T. V. Spina, S. B. Martins, and R. Phellan, “Medical image segmentation using object shape models: A critical review on recent trends, and alternative directions,” *VipIMAGE*, pp. 9–15, 2015;
- (ii) S. B. Martins, T. V. Spina, C. L. Yasuda, and A. X. Falcão, “A multi-object statistical atlas adaptive for deformable registration errors in anomalous medical image segmentation,” in *SPIE Medical Imaging*, vol. 10133, pp. 691–698, 2017. **Honorable mention**;
- (iii) S. B. Martins, J. Bragantini, C. L. Yasuda, and A. X. Falcão, “An adaptive probabilistic atlas for anomalous brain segmentation in MR images,” *Medical Physics*, vol. 46, no. 11, pp. 4940–4950, 2019.

from the appearance of the model. One might post-process the resulting segmentation masks, for example, by removing cerebrospinal fluid (CSF) voxels obtained by tissue classification (Fig. 3.1f). However, post-processing does not fix existing segmentation errors of the model on gray matter (GM) and white matter (WM) voxels – *e.g.*, voxels between the hemispheres in Figs. 3.1e–f. Post-processing can still increase those segmentation errors (Fig. 3.1f).

Figure 3.1: Automatic brain segmentation by the probabilistic-atlas-based method SOSM-S [120]. (a) Coronal slice of a 3D brain image after left temporal lobe resection (arrow). (b) The corresponding slice of a prior probability map for the left hemisphere (blue and yellow voxels indicate the certainty and uncertainty regions, respectively). (c) Overlapping between (a) and (b). (d) Estimated seeds for object delineation. (e) Coronal slice of the resulting 3D object mask (blue) and the gold-standard border (red). Arrows indicate segmentation errors. (f) Post-processed object mask after removing voxels classified as CSF by expectation maximization algorithm [145].

Deep learning methods: Recent work proposes deep neural networks for segmenting GM, WM, and CSF [11, 32]; hippocampus [57]; brain lesions [35, 146]; and skull [114]. These approaches usually (i) take high processing times, (ii) depend on a large number of training images, which must be previously annotated by experts; and (iii) may require weight fine-tuning (retraining) when used in each new distinct set of images [11]. The latter is certainly a significant limitation, due to the difficulty of annotating medical image training sets with 3D objects [147], the image variability across scanners and acquisition protocols, and differences between healthy and pathological brain tissues. In this sense, methods that rely on

object-shape-based models and texture classifiers, which can be created from a few labeled voxels, are more attractive. Our proposed research falls into this category.

In this chapter, we present an effective and efficient adaptive probabilistic atlas, named *AdaPro*, to circumvent the above limitation. Our method adapts shape constraints on-the-fly according to the presence of detected anomalies in the target image. The adaptive model uses a binary texture classifier trained from a few background and object voxels on a *template* (reference 3D image). It relies only on voxel feature representation and texture classification to adapt its probabilistic atlas. Finally, *AdaPro* performs a new object-based delineation algorithm based on combinatorial optimization and diffusion filtering [148, 149] for shape smoothing.

We structure this chapter as follows. Section 3.1 presents related work on object-shape-based models for automatic brain image segmentation and details the base-lines considered during experiments. Section 3.2 introduces *AdaPro*. Section 3.3 describes the experimental setup, while Section 3.4 presents and discusses the results. Finally, Section 3.5 concludes this chapter.

3.1 RELATED WORK

Image segmentation involves two tightly coupled tasks: object *recognition* and object *delineation* [150]. *Recognition* indicates the whereabouts of the desired object, while *delineation* precisely defines its spatial extent in the image.¹ Some approaches can present different levels of automation for each task, varying from purely manual to fully automatic. This thesis only focuses on fully automatic approaches.

Atlas-based segmentation, also called object-shape-based segmentation, is one of the most widely-used and successful approaches for *automatic* brain image segmentation. These methods use the *a priori* knowledge about objects' shapes from a training set $X = \{A_1, \dots, A_n\}$ with n atlases. Each *atlas* $A_i = (I_i, M_i)$ consists of a source 3D image I_i (e.g., MR image) and its corresponding 3D label image M_i with the segmentation mask of each 3D object of interest. These segmentation masks, called *gold standards*, are obtained from manual or semi-automatic interactive segmentation by one or multiple experts. This process typically requires outlining the structures in a slice-by-slice fashion, resulting in a time-consuming, tedious, and error-prone task [11].

The simplest atlas-based methods [151–153] rely on a single atlas $A_r = (I_r, M_r)$ and segment a 3D test image I by propagating the labels from M_r to I after image registration between I and I_r (template). However, since registration does not perfectly align the borders of the registered image and the template (Section 2.3.4), a single atlas is insufficient to capture wide anatomical variations, especially when anomalies are present in I . This strategy evolved to *probabilistic atlases* (PAs) and then to *Multi-Atlas Label Fusion* (MALF), as discussed next.

¹ Some authors commonly use the term *segmentation* to refer to *delineation* so that recognition is implicitly assumed [150].

3.1.1 Probabilistic Atlas

Methods based on PAs estimate the probability of a voxel is part of a given object regardless of its intensity in the original image [41, 120, 143]. Fig. 3.2 shows the general pipeline of these methods. The construction of a PA initially requires the selection of a standard *template* I_r (Fig. 3.2, Step 1). One can select the most similar atlas to the others in the training set [120] or use a well-established coordinate space, such as the popular MNI template [108]. The training source images are then registered to I_r by non-rigid registration, and their segmentation masks are also mapped to I_r by applying the corresponding transformations (Fig. 3.2, Step 2). Once the training atlases are on the same coordinate space, one computes a *prior probabilistic map* P (Fig. 3.2, Step 3), where each voxel v has a *prior probability* of belonging to a given object of interest [144] — i.e., the frequency in which v assumes the label of such object in all training masks (see the resulting map for AdaPro posteriorly illustrated in Fig. 3.6). All voxels with probability within $(0, 1)$ form the *uncertainty region*, where the object’s boundaries are likely to fall. For segmentation, the unseen test image I_t is mapped on to the coordinate system of P (recognition), and delineation (Fig. 3.2, Steps 4–5) typically involves thresholding the prior probabilities associated to the voxels [104, 154], or estimating and thresholding a *posterior* probabilities [120, 155], or by using other image processing operators [120, 147, 156].

Figure 3.2: General steps for the construction and use of probabilistic atlases for automatic image segmentation.

PA-based methods have been actively investigated in the past decades [41, 120, 141, 144, 147, 154–158]. For example, the well-known FreeSurfer [144] automatically segments several brain anatomical structures by combining the use of a Markov random field and probabilistic atlases into a Bayesian framework [144,

157]. Despite being widely accepted as the *de facto* standard of brain segmentation in many anatomical studies [144, 159], FreeSurfer demands a considerably high processing-time to segment an image ($\sim 15\text{h}$), which makes it impracticable for clinical routine and studies involving large datasets.

Some strategies combine registration and delineation into a probabilistic framework [147]. Pohl *et al.* [155] proposed using the expectation-maximization (EM) algorithm [160] to find the optimum parameters for registration while labeling each voxel to a brain tissue. Ashburner and Friston [158] proposed a unified objective function to segment brain tissues (CSF, WM, and GM) while correcting the bias field and refining registration, whose parameters are derived from a mixture of Gaussians. Since this approach uses only healthy-shape priors, it fails to segment images with some anomalies.

The accuracy of PA-based approaches is very sensitive to registration errors, mainly caused by the inter-subject variability in anatomical patterns [147]. Some methods then use image processing operators to attenuate the impact of such errors while refining the final delineations of the brain structures. Grau *et al.* [156] use the watershed algorithm [133] from seeds corresponding to each brain-tissue class. Such seeds are estimated as the skeletons calculated from the probabilistic atlas.

Recently, Phellan *et al.* [120] have demonstrated that if we acknowledge registration errors, the accuracy of probabilistic atlases may be significantly improved as long as a local search for the object is performed with the model. For this purpose, their probabilistic atlas, named SOSM-S, uses a triple $\{P, D, F\}$ composed of the traditional *prior probability map* P (Fig. 3.6), a *delineation algorithm* D , and a *criterion function* F . The goal of P is to impose the object's shape learned from the training atlases, which aims to constrain the delineation with D to occur only for *uncertain voxels*. The delineation algorithm D , in turn, aims to adapt the shape constraints to the test image, rather than merely thresholding P after registration. SOSM-S uses the watershed transform by the Image Foresting Transform (IFT) algorithm [60] for object delineation (Section 2.4). Criterion F is used for the local object search to mitigate registration errors. The search requires the translation of P over the registered test image, delineating a candidate object with D at every position. Then, it evaluates the resulting mask with F , which is expected to be maximum when the candidate mask corresponds to the target object. Such synergistic operation between object localization and delineation constitutes the task of segmentation and stems from Fuzzy Object Shape Models [161].

Despite the success of SOSM-S for brain segmentation in control images, the method fails in the presence of anomalies since the model imposes the shape of healthy organs (Fig. 3.1). Moreover, only MR-T1 brain images of 2T were used for evaluation, ignoring inherent problems of higher quality 3T images (e.g., field inhomogeneity and noise). SOSM-S' local search, in turn, may require considerable processing time for performing IFT delineation for each candidate in the search region.

In this thesis, we revisited SOSM-S to present a faster adaptive model, named AdaPro, for anomalous brain segmentation with the following contributions:

1. We incorporate a *texture* classifier based on convex optimization that indicates *on-the-fly* the regions of the target 3D image where the shape constraints should be further adapted on the probabilistic atlases (*i.e.*, disregarded during object delineation). This strategy avoids the problems of applying the classifier as a post-processing operation (Fig. 3.1f);
2. We increase the uncertainty region of the statistical shape model (by erosion and dilation) to avoid local object search, as adopted by SOSM-S. This strategy makes the segmentation considerably faster;
3. We added the object relaxation procedure from Moya and Falcão [149] to improve segmentation accuracy and make the objects' appearances more pleasant (without jagged edges);
4. We extended the brain segmentation to the *native coordinate space* of the test images.

Section 3.2 introduces AdaPro and details the above contributions.

3.1.2 Multi-Atlas Label Fusion (MALF)

Methods based on *Multi-Atlas Label Fusion* (MALF) aim to suppress registration errors by considering that each training atlas $A_i = (I_i, M_i)$ produces one candidate segmentation per test image I_t . Each training atlas is registered on to I_t so that the fusion of all mapped segmentation masks generates the final segmentation [80, 104, 162–166]. MALF can be performed either in the *native* or *template space*. Fig. 3.3 presents the general steps of MALF.

Figure 3.3: General steps of Multi-Atlas Label Fusion for image segmentation.

Challenges: The computational bottleneck of MALF consists of registering the entire training atlas set to I_t , which demands the most significant processing time during segmentation [104]. Some MALF methods then select the subset of the k most similar training atlases to I_t (*atlas selection*) to make the segmentation faster, which is particularly important in scenarios where time is a significant

constraint [104]. Moreover, when removing training atlases that are anatomically unrepresentative to I_t , one might expect to improve the segmentation accuracy [167].

Pipeline: Most atlas-selection methods rank the relevance of the training atlas set X by employing a given metric based on image-similarity [167–169]. Aljabar *et al.* [167], for example, proposed the use of normalized mutual information (NMI) to rank X — see more details of NMI in Appendix C.1. Initially, the method chooses a reference image I_r from X and maps all training atlases to I_r . A test image I_t is then registered on to the space of I_r and the NMI between I_t and each image in X is computed. Finally, the method selects the k top-ranked training atlases for subsequent segmentation. This approach considerably reduces the number of registrations during atlas selection. Asman *et al.* [170] replace NMI with principal components analysis to define atlas similarity metrics. More recently, other works rely on clustering to select the most similar training atlases [171, 172] to I_t . For instance, Nouranian *et al.* [171] partition the set $I_t \cup X$ into k clusters by using the K-means algorithm [173] and select a subset of training atlases belonging to the same cluster of I_t .

Once the k most relevant training atlases are selected, their source images are registered on to the coordinate space of I_t , and their masks are also mapped to by applying their corresponding deformation fields. Although non-rigid registration is time-consuming, it is preferable in MALF applications for better capturing anatomical variation between different subjects [147], resulting in higher segmentation accuracies [104]. Typical non-rigid registration techniques are based on mathematic transformations, such as cosine-based functions [174], B-spline curves [107], and level set partial different equations [175].

Finally, the propagated labels from the registered segmentation masks are combined (*label fusion*) to generate the final segmentation. The simplest solution is *majority voting*, which selects the most frequent label at each voxel [176]. Other works assign a local or global weight for each registered training atlas during label fusion, which reflects the similarity between the test image and the atlas [112, 162, 163, 177]. For instance, Artachevarria *et al.* [177] compute global weights from the NMI between the training atlases and the test image. One of the most popular techniques is STAPLE [164], which weights each training atlas using the expectation-maximization algorithm. Alternative strategies estimate local weights by computing local cross-correlation [178], local mutual information [179], and local registration accuracy [180]. Recent label-fusion techniques involve the use of patches to compute weights at each voxel [181, 182].

Implementations: In the context of brain image segmentation, we highlight the recent software called volBrain [80], a solution that provides segmentation and structure asymmetry ratios at different scales for intracranial cavity (skull stripping), tissue volumes (GM, WM, and CSF), brain segmentation, lateral ventricles,

and subcortical GM structures. volBrain has reached superior segmentation results as compared to well-established and publically available solutions, such as FreeSurfer [144] and FIRST [183], serving as a starting point for works regarding brain image analysis.

volBrain uses NABS (Non-local Automatic Brain hemisphere Segmentation) [162] for brain segmentation. This method randomly selects 30 atlases from a training set of healthy subjects (with ~ 600 atlases), which are then preprocessed using the following operations: noise reduction, registration on to linear MNI space [108], skull stripping, inhomogeneity correction, and intensity normalization. After preprocessing the test image, NABS performs a tissue classification that only considers WM, GM, and CSF. Finally, the non-local label-fusion technique proposed by Coupe *et al.* [181] segments the hemispheres, cerebellum, and brainstem. This label fusion estimates the influence of each atlas for each voxel v by computing the Euclidean distance of a 3D patch around v in the test image and the source images from the chosen atlases. We considered volBrain as a baseline in our experiments (Section 3.3).

3.2 ADAPTIVE PROBABILISTIC ATLAS (ADAPRO)

We propose an Adaptive Probabilistic atlas (AdaPro) for anomalous brain image segmentation that incorporates a *texture classifier* during object delineation. This classifier indicates *on-the-fly* the regions of the target 3D image where the shape constraints should be further adapted (*i.e.*, regions disregarded during segmentation) due to the presence of abnormalities (Fig. 3.1a). Fig. 3.4 presents the steps for the construction and use of AdaPro, whereas Fig. 3.5 shows resulting images of these steps for the segmentation of both hemispheres and cerebellum. Although some figures presented in this section show examples with MR slices, AdaPro uses 3D MR images.

3.2.1 Construction

Recall a training set $X = \{A_1, \dots, A_n\}$ with n atlases of healthy subjects, where each *atlas* $A_i = (I_i, M_i)$ consists of a source 3D image I_i and its corresponding label image M_i with the mask of each 3D object of interest.

Template Selection

Since our method is based on registration, the first step is to select a standard reference coordinate space (*template*) where the probabilistic atlas will be constructed. This template might be a popular atlas, such as MNI [108], or the most similar training atlas to the entire training set. We apply the same procedure employed by Phellan *et al.* [120] that selects the training atlas $A_r = (I_r, M_r)$ whose

Figure 3.4: Pipeline for the construction and use of AdaPro.

union of the object masks in M_r has the least mean Average Symmetric Surface Distance (ASSD; [Appendix C.2](#)) to all others from X . The idea is to reduce the amount of deformation each image must undergo for model construction.

Preprocessing and Registration

As outlined in [Section 2.3](#), automated MR image segmentation is challenging due to inherent problems of image acquisition, such as noise, field inhomogeneity, and variability of the intensity ranges, mainly in high-resolution images from 3T scanners, for example. We then perform some image processing operations, as detailed in [Section 2.3](#), to improve the quality of the images.

We first apply noise reduction by median filtering, followed by MSP alignment, and an inhomogeneity correction with N4 [102]. The resulting images are then registered to the template I_r by non-rigid registration ([Section 2.3.4](#)), and their masks are also mapped to by applying their corresponding deformation fields ([Figs. 3.5a–b](#)).

Probabilistic Atlas Creation

For each object m , we build a *probabilistic atlas* P_m by counting the frequency of the label assignment from all training registered atlases $A_i \in X$ and keeping the mostly assigned label to each voxel. Label assignment only takes into account the object m and the background (label 0). This frequency corresponds to the *prior*

Figure 3.5: AdaPro segmentation. **(a)** Axial slice of a post-surgery 3D MR-T1 image (the arrow indicates an anomaly). **(b)** Object localization mask on the preprocessed and registered image. **(c)** Histogram matching of **(b)**. **(d)** Positive voxels (orange) after texture classification. **(e)–(g)** Estimated seeds from the adaptive models for the background and target objects. Magenta voxels indicate the forbidden regions imposed by **(d)**. **(h)** Gradient of **(c)**. **(i)–(j)** Delineated objects in 2D (axial slice) and 3D.

probability of the voxel to belong to object m . Fig. 3.6 depicts the probabilistic atlases for the cerebellum and brain hemispheres.

One might also build a single multi-object probabilistic atlas with all objects under study. However, a previous study showed that the use of a probabilistic

Figure 3.6: Axial, coronal, and sagittal slices of the combination (only for visualization) of probabilistic atlases for the cerebellum (red), right hemisphere (green), and left hemisphere (blue). The brighter the object's color, the greater its prior probability value.

atlas for each object results in more accurate segmentation as compared to the multi-object strategy [141].

Model Adaptation

We design a binary classifier C based on a linear Support Vector Machine [184] by interactively selecting training voxels on the 3D template I_r . SVM is fast and robust to classify high-dimensional data (like ours). Brain tissue voxels are considered positive samples and voxels with typically different intensities (CSF and image background) are considered negative samples. Each training voxel is represented by its intensity and the intensities of all neighbors inside a sphere of radius 5.0, resulting in a 515-dimensional feature vector. Fig. 3.7 shows an example of chosen voxels on a given I_r .

Figure 3.7: Slices of the chosen template (reference 3D image) with the selected voxels to design the texture binary classifier. The brain tissue voxels (orange scribbles) are the positive samples, whereas the voxels with different intensities (magenta scribbles) are the negative samples.

3.2.2 Segmentation

Let I_t represent a target 3D brain image after preprocessing and non-rigid registration on to I_r (Fig. 3.5b). Each object of interest m is independently segmented by following the steps, as detailed below. Then, all delineated objects are combined to result in the final 3D label image M_t .

Object Localization and Histogram Matching

As AdaPro relies on a texture classifier to detect anomalous regions, the wide differences in intensity and contrast among the 3D images, mainly when considering images provided by different scanners, must be attenuated to guarantee a similar range of intensities for the same tissue. One could then apply a histogram matching between I_t and I_r , but voxels from irrelevant tissues/organs for the addressed problem (e.g., neck and bones) can negatively impact this operation. AdaPro circumvents this problem by binarizing the probabilistic atlas P_m of each object m , followed by morphological closing to fill small holes (e.g., small gaps inside sulci; maximum volume of $8 \times 8 \times 8 \text{ mm}^3$), and merging them into a single 3D binary mask B (object localization mask). B can still be dilated if the user provides a dilation radius > 1 for seed estimation. Note that B contains all voxels from the certainty and uncertainty regions for all target objects (localization) that indeed define the regions for object delineation (Fig. 3.5b). Finally, AdaPro performs a histogram matching between I_t and I_r only inside the object voxels defined by B (Fig. 3.5c).

Texture Classification and Seed Estimation

We aim at estimating a seed set $\mathcal{S} = \mathcal{S}_0 \cup \mathcal{S}_m$, where $\mathcal{S}_m$, $m > 0$, contains seed voxels selected inside the object m , and $\mathcal{S}_0$ contains seed voxels selected in the background. The borders of the dilated and eroded certainty region of P_m form $\mathcal{S}_0$ and $\mathcal{S}_m$, respectively. The dilation and erosion slightly increase the uncertain region of P_m . Thus, seed estimation is simpler and faster as compared to SOSM-S (Section 3.1.1), for example, since it avoids performing several delineations as it is done during SOSM-S' local object search.

To identify regions on I_t where shape constraints should be adapted, AdaPro classifies I_t with the texture binary classifier C (Fig. 3.5d). Then, the residual image of B vs the classification forms a *forbidden region* F (magenta voxels in Figs. 3.5e–g) that eliminates its voxels from the competition between internal and external seeds during object delineation. This strategy will not work as a post-processing operation, because object delineation without the forbidden region might mislabel voxels from distinct objects.

Object Delineation

For the delineation of each object m , AdaPro uses a fast implementation [149] of the Relaxed Image Foresting Transform (RIFT) algorithm [148], an IFT-based algorithm (Section 2.4) that smooths the delineated object borders. This algorithm first computes a watershed transform from labeled seeds defined by $\mathcal{S} = \mathcal{S}_0 \cup \mathcal{S}_m$, exactly as in SOSM-S. The seed competition takes into account the gradient image of I_t (Fig. 3.5h), whose voxels are expected to be brighter along the objects' boundaries than elsewhere. Therefore, the voxels are conquered by the seed, which offers the path whose maximum intensity along it is minimum as compared to any other path from the remaining seeds. A fast diffusion filter [149] then smooths the resulting mask. Since the segmentation of each object m is performed independently, when a voxel is assigned to multiple labels, its final label is the one with higher probability value in the probabilistic atlas (Figs. 3.5i–j).

Segmentation on Native Image Space

AdaPro can segment I_t on its own native image space (NIS). The only change for this is to map the probabilistic atlases to I_t by applying the transformations (deformation fields) resulting from the non-rigid registration between the template I_r on to I_t . AdaPro then performs the remaining steps as previously presented.

One might use AdaPro to segment I_t on the template coordinate space and then inversely mapping it to NIS. However, this approach cannot tackle segmentation errors resulting from the inverse mapping. This strategy proposed by AdaPro is different, as it relies on object delineation to circumvent those registration errors.

3.3 EXPERIMENTAL SETUP

To assess the performance of AdaPro, we conducted a set of experiments. This section describes the MR-T1 image datasets, compared methods, and the evaluation protocol considered for the experiments. All computations were performed on the same Intel Core i7-7700 CPU 3.60GHz with 64GB of RAM.

3.3.1 Datasets

We used six in-house datasets of 3D MR-T1 brain images of healthy subjects and patients before and after temporal lobe resections, as presented in Appendix B.1. All images were provided by the Neuroimaging Laboratory (LNI) at the University of Campinas (UNICAMP), Brazil. The datasets were divided into two groups according to their field strengths.

Group A has three datasets of 2T images acquired with a 2T Elscint scanner and voxel size $0.98 \times 0.98 \times 0.98 \text{ mm}^3$: (HEALTHY-2T) 19 images from healthy subjects, (PRE-2T) 20 pre-operative images from epilepsy patients, and (POST-2T) 20 post-

operative images from epilepsy patients. This group has 20 epilepsy patients with a pair of pre- and post-operative images for each one.

Group B has three datasets of 3T images acquired with a 3T Siemens scanner and voxel size $1 \times 1 \times 1 \text{ mm}^3$: (HEALTHY-3T) 20 images from healthy subjects, (PRE-3T) 30 pre-operative images from epilepsy patients, and (POST-3T) 60 post-operative images from epilepsy patients. This group has 30 epilepsy patients to which there are one pre- and two post-operative images.

A template obtained from HEALTHY-2T, with $165 \times 255 \times 255$ voxels, was considered for group A, whereas a template obtained from HEALTHY-3T, with $180 \times 240 \times 240$, was used for group B. Therefore, our study involved a total of 169 images.

A neurologist from LNI has carefully delineated the cerebellum (C), right hemisphere (RH), and left hemisphere (LH) to build the gold-standard segmentation of all datasets. Consequently, brainstem is not considered in the experiments. Resulting segmentation masks consider that background voxels have label 0 and each object has label $m = 1, 2, \dots, L$, where L is the total number of objects. Fig. 3.8 presents examples of the considered datasets.

3.3.2 Evaluation Protocol

We compared AdaPro² with three atlas-based methods: SOSM-S [120], volBrain [80], and a MALF technique instantiated with the popular atlas selection of Aljabar *et al.* [167] and STAPLE label fusion [164]. All these methods are detailed in Section 3.1.

The quality, resolution, and brain tissues' intensities are quite different in 2T and 3T MR-T1 images (Section 2.2.2), even after applying the same preprocessing operations detailed in Section 3.2.1. For example, 2T images are noisier than 3T images, whereas field inhomogeneity is higher in 3T than 2T images — compare the examples in Fig. 3.8. Thus, a single texture classifier — designed from voxels selected in a *template* acquired with a given field strength — may not be accurate when classifying anomalies in images acquired with different field strength. Therefore, we evaluated Groups A and B independently (Section 3.3.1). For each one, we considered its dataset of healthy subjects as the training set, which is used for SOSM-S, MALF, and AdaPro. volBrain uses its own training atlas set since we do not have access to it.

We considered the entire training set to construct the probabilistic map of SOSM-S and AdaPro. Due to processing-time constraints, MALF used all images of HEALTHY-2T for evaluation in group A, and it selected the 20 top-ranked training atlases from HEALTHY-3T for each test image during validation in group B. All methods used the same 3D template (reference image) which consists of the

² All binaries of AdaPro can be found on https://github.com/lidsunicamp/MedicalPhysics19_AdaPro

Figure 3.8: Examples of the considered datasets for brain segmentation. Each example shows an axial slice of the 3D MR-T1 image after preprocessing, and its gold-standard segmentation. Arrows indicate removed tissues after lobe resection.

most similar atlas to the others in the training set. We performed all registrations by the Elastix³ software [107].

The linear SVM texture classifier from AdaPro was trained from positive and negative voxels chosen on the template of each group (Fig. 3.7). We set the SVM penalty parameter C to 10^{-5} . For each object m , AdaPro requires the choice of the radii d and e for dilation and erosion of the probabilistic atlas, respectively. We used the training set of each group to determine the best choice for these parameters by grid search. The values $(d_C, d_{RH}, d_{LH}), (e_C, e_{RH}, e_{LH})$ found respectively for the cerebellum (C), right hemisphere (RH), and left hemisphere (RH) are: group A (0, 0, 0), (1, 2, 2) and group B (0, 0, 0), (1, 2, 2). These values were then fixed for all experiments.

Since not all baselines work in both *template* and *native* image spaces, we made some adaptations. The software volBrain outputs brain segmentation masks in native and MNI linear space (after mapping the input image by affine registration). To obtain the segmentation masks in the coordinate space of the chosen templates for the groups, we used the registered test images as input for volBrain. Conversely, SOSM-S only performs the segmentation in the template space. Then, for each test image, we applied the inverse transformation on its segmentation masks in order to have them in the native space of each test image.

3.4 RESULTS

This section presents and discusses the results of the quantitative evaluation of the methods on the two datasets of epilepsy patients from each group. In our evaluation, we will rely more on the Average Symmetric Surface Distance (ASSD; Appendix C.2) score (in millimeters) as segmentation accuracy measure than on global measures such as Dice (Appendix C.2). ASSD better captures segmentation errors along the segmented boundaries than Dice, especially in the case of post-operative images with removed portions of the brain (Fig. 3.8). However, we also include Dice in the results due to its wide popularity in the literature. Finally, we performed an analysis of variance (ANOVA) to compare the resulting mean scores between the considered methods.

Tables 1 and 2 present the ASSD score (lower is better) for the epilepsy datasets of 2T images and 3T images, respectively. Tables 3 and 4, in turn, present the Dice score (higher is better) for the same datasets. The numbers correspond to the mean and standard deviation values of all instances of each object of interest. Resulting p-values from the ANOVA test are shown in parenthesis for each evaluation scenario. Fig. 3.9 shows the mean segmentation errors in the template space for the considered baselines.

SOSM-S and MALF perform worse on post-operative images because they cannot capture abnormalities, as evidenced in Fig. 3.9. Note that errors occur mainly

³ We use the *par0000* files available on http://elastix.bigr.nl/wiki/index.php/Parameter_file_database

Table 1: Comparison of ASSD (mm) for the Cerebellum (C), Right Hemisphere (RH), and Left Hemisphere (LH) of the pre- and post-operative images of **2T**. Lower ASSD means better accuracy (bold scores are the best with statistical significance). Resulting p-values from ANOVA test are in parenthesis.

Epilepsy Datasets (Group A) - ASSD							
		PRE-2T			POST-2T		
		C	RH	LH	C	RH	LH
template space	SOSM-S	0.90 ± 0.08	0.99 ± 0.12	1.07 ± 0.14	0.79 ± 0.15	1.41 ± 0.20	1.32 ± 0.39
	MALF	0.88 ± 0.18	0.93 ± 0.08	0.91 ± 0.11	1.15 ± 0.37	1.36 ± 0.22	1.38 ± 0.39
	volBrain	0.97 ± 0.18	1.52 ± 0.32	1.47 ± 0.30	0.98 ± 0.19	1.32 ± 0.17	1.27 ± 0.21
	AdaPro	0.86 ± 0.11 <i>(p = 0.098)</i>	0.92 ± 0.11 <i>(p < 0.001)</i>	0.90 ± 0.12 <i>(p < 0.001)</i>	0.96 ± 0.19 <i>(p < 0.001)</i>	1.17 ± 0.19 <i>(p < 0.001)</i>	1.13 ± 0.30 <i>(p = 0.013)</i>
native space	SOSM-S	0.97 ± 0.10	1.07 ± 0.11	1.14 ± 0.12	0.96 ± 0.13	1.24 ± 0.20	1.22 ± 0.22
	MALF	0.79 ± 0.25	0.87 ± 0.09	0.86 ± 0.11	0.86 ± 0.30	1.01 ± 0.11	1.04 ± 0.19
	volBrain	1.00 ± 0.21	1.46 ± 0.33	1.40 ± 0.30	1.01 ± 0.23	1.43 ± 0.24	1.41 ± 0.28
	AdaPro	0.83 ± 0.15 <i>(p = 0.001)</i>	0.86 ± 0.12 <i>(p < 0.001)</i>	0.85 ± 0.11 <i>(p < 0.001)</i>	0.86 ± 0.16 <i>(p = 0.070)</i>	0.90 ± 0.13 <i>(p < 0.001)</i>	0.90 ± 0.20 <i>(p < 0.001)</i>

Table 2: Comparison of ASSD (mm) for the Cerebellum (C), Right Hemisphere (RH), and Left Hemisphere (LH) of the pre- and post-operative images of **3T**. Lower ASSD means better accuracy (bold scores are the best with statistical significance). Resulting p-values from ANOVA test are in parenthesis.

Epilepsy Datasets (Group B) - ASSD							
		PRE-3T			POST-3T		
		C	RH	LH	C	RH	LH
template space	SOSM-S	0.91 ± 0.22	1.04 ± 0.14	1.02 ± 0.13	0.90 ± 0.19	1.28 ± 0.20	1.25 ± 0.18
	MALF	1.02 ± 0.31	1.07 ± 0.18	1.06 ± 0.17	1.00 ± 0.28	1.30 ± 0.16	1.25 ± 0.20
	volBrain	1.02 ± 0.08	1.42 ± 0.11	1.36 ± 0.10	1.01 ± 0.09	1.51 ± 0.14	1.47 ± 0.13
	AdaPro	0.75 ± 0.07	1.03 ± 0.18	1.01 ± 0.14	0.75 ± 0.09	1.08 ± 0.18	1.10 ± 0.18
		(<i>p</i> < 0.001)	(<i>p</i> < 0.001)	(<i>p</i> < 0.001)	(<i>p</i> < 0.001)	(<i>p</i> < 0.001)	(<i>p</i> < 0.001)
native space	SOSM-S	0.97 ± 0.26	1.16 ± 0.17	1.15 ± 0.16	0.94 ± 0.19	1.29 ± 0.29	1.26 ± 0.24
	MALF	1.05 ± 0.40	1.14 ± 0.23	1.12 ± 0.23	1.01 ± 0.38	1.24 ± 0.26	1.18 ± 0.26
	volBrain	1.06 ± 0.09	1.44 ± 0.16	1.42 ± 0.16	1.05 ± 0.11	1.56 ± 0.16	1.55 ± 0.17
	AdaPro	0.68 ± 0.09	0.94 ± 0.16	0.94 ± 0.13	0.67 ± 0.10	0.99 ± 0.19	1.02 ± 0.17
		(<i>p</i> < 0.001)	(<i>p</i> < 0.001)	(<i>p</i> < 0.001)	(<i>p</i> < 0.001)	(<i>p</i> < 0.001)	(<i>p</i> < 0.001)

in the temporal lobes, where a portion has been removed. Their ASSD scores in the cerebellum are similar for pre- and post-operative images because it is the only object that remains intact after resection surgery.

Table 3: Comparison of Dice for the Cerebellum (C), Right Hemisphere (RH), and Left Hemisphere (LH) of the pre- and post-operative images of **2T**. Higher Dice means better accuracy (bold scores are the best with statistical significance). Resulting p-values from ANOVA test are in parenthesis.

		Epilepsy Datasets (Group A) - Dice					
		PRE-2T			POST-2T		
		C	RH	LH	C	RH	LH
template space	SOSM-S	0.944 ± 0.006	0.961 ± 0.005	0.957 ± 0.006	0.949 ± 0.013	0.952 ± 0.008	0.951 ± 0.014
	MALF	0.953 ± 0.011	0.969 ± 0.002	0.966 ± 0.004	0.930 ± 0.025	0.951 ± 0.008	0.948 ± 0.015
	volBrain	0.943 ± 0.014	0.955 ± 0.009	0.954 ± 0.009	0.935 ± 0.026	0.949 ± 0.010	0.951 ± 0.015
	AdaPro	0.949 ± 0.008 (<i>p</i> = 0.012)	0.967 ± 0.004 (<i>p</i> < 0.001)	0.966 ± 0.004 (<i>p</i> < 0.001)	0.940 ± 0.015 (<i>p</i> = 0.039)	0.954 ± 0.008 (<i>p</i> = 0.452)	0.953 ± 0.013 (<i>p</i> = 0.726)
native space	SOSM-S	0.933 ± 0.014	0.955 ± 0.004	0.952 ± 0.006	0.935 ± 0.017	0.949 ± 0.009	0.946 ± 0.012
	MALF	0.951 ± 0.021	0.971 ± 0.003	0.969 ± 0.004	0.950 ± 0.027	0.965 ± 0.008	0.961 ± 0.012
	volBrain	0.941 ± 0.020	0.960 ± 0.009	0.959 ± 0.009	0.940 ± 0.022	0.957 ± 0.008	0.956 ± 0.010
	AdaPro	0.948 ± 0.015 (<i>p</i> = 0.015)	0.969 ± 0.005 (<i>p</i> < 0.001)	0.969 ± 0.004 (<i>p</i> < 0.001)	0.947 ± 0.016 (<i>p</i> = 0.124)	0.966 ± 0.006 (<i>p</i> < 0.001)	0.963 ± 0.010 (<i>p</i> < 0.001)

Table 4: Comparison of Dice for the Cerebellum (C), Right Hemisphere (RH), and Left Hemisphere (LH) of the pre- and post-operative images of **3T**. Higher Dice means better accuracy (bold scores are the best with statistical significance). Resulting p-values from ANOVA test are in parenthesis.

		Epilepsy Datasets (Group B) - Dice					
		PRE-3T			POST-3T		
		C	RH	LH	C	RH	LH
template space	SOSM-S	0.941 ± 0.016	0.953 ± 0.007	0.952 ± 0.007	0.942 ± 0.015	0.947 ± 0.013	0.947 ± 0.011
	MALF	0.939 ± 0.021	0.956 ± 0.007	0.954 ± 0.007	0.940 ± 0.020	0.949 ± 0.011	0.950 ± 0.010
	volBrain	0.930 ± 0.013	0.937 ± 0.009	0.939 ± 0.008	0.932 ± 0.012	0.939 ± 0.009	0.938 ± 0.008
	AdaPro	0.952 ± 0.008 (<i>p</i> < 0.001)	0.955 ± 0.007 (<i>p</i> < 0.001)	0.955 ± 0.005 (<i>p</i> < 0.001)	0.952 ± 0.008 (<i>p</i> < 0.001)	0.956 ± 0.006 (<i>p</i> < 0.001)	0.954 ± 0.006 (<i>p</i> < 0.001)
native space	SOSM-S	0.931 ± 0.027	0.941 ± 0.010	0.941 ± 0.009	0.934 ± 0.022	0.936 ± 0.015	0.937 ± 0.013
	MALF	0.935 ± 0.034	0.952 ± 0.010	0.951 ± 0.011	0.938 ± 0.032	0.946 ± 0.013	0.947 ± 0.014
	volBrain	0.928 ± 0.016	0.936 ± 0.008	0.936 ± 0.008	0.931 ± 0.017	0.936 ± 0.009	0.936 ± 0.009
	AdaPro	0.956 ± 0.013 (<i>p</i> < 0.001)	0.956 ± 0.007 (<i>p</i> < 0.001)	0.954 ± 0.007 (<i>p</i> < 0.001)	0.958 ± 0.011 (<i>p</i> < 0.001)	0.956 ± 0.007 (<i>p</i> < 0.001)	0.954 ± 0.008 (<i>p</i> < 0.001)

Since CSF fills the removed portions in the temporal lobes after surgery, volBrain can correctly identify them during segmentation due to its tissue classification that separates voxels of WM, GM, and CSF. However, its ASSD is higher than the other baselines because it also misclassifies several GM voxels from the hemi-

Figure 3.9: Axial and coronal slices with the mean segmentation errors from the baselines (0 means no error, and 1 means errors in all images) on template coordinate space for the entire datasets.

spheres as CSF, resulting in more segmentation errors (Fig. 3.9). Its preprocessing probably was not able to guarantee the same range of values for the tissues of the testing images, resulting in misclassification.

In contrast, AdaPro’s classification can detect the anomalies without missing object voxels in the brain hemispheres and cerebellum. Its preprocessing operations are fundamental for that since it is responsible for ensuring that all voxels of the same tissue will have intensities within the same range. Moreover, the relaxation process on its delineated objects results in smoother and well-defined object boundaries, thus avoiding an effect of serrated borders. The texture classification and delineation perform better in 3T images due to their higher quality when compared to 2T images. Consequently, the segmentation results are better in images of 3T than 2T.

Regarding the ASSD scores, AdaPro is more accurate than the baselines in the post-operative images for the hemispheres (Fig. 3.9) – objects affected by surgical procedures – and for the pre-operative images of 3T. AdaPro is equivalent to MALF for PRE-2T and the cerebellum in POST-2T. We have evidence, based on the p-values provided by the ANOVA test, that the mean scores of AdaPro differ from the baselines (see the bold scores in the Tables 1–4).

Conversely, all evaluated methods present similar Dice scores in most scenarios, even for post-operative images. AdaPro presents equivalent results compared to the baselines for 2T images and the PRE-3T dataset, being superior for 3T images in the *native* image space. Dice is not sensitive to capture local segmentation errors, such as the regions surgically removed in the temporal lobes (Fig. 3.8). It is also sensitive to the size of the object – differences in Dice values for small objects are less significant than the same values for large objects. Therefore, we prefer to draw conclusions based on ASSD than on Dice.

volBrain’s segmentation is performed on its own online platform and takes around 12 minutes. It is also limited to 10 free executions per day. SOSM-S takes around 110 seconds (50 seconds for registration and 60 seconds for object delineation). MALF is the slowest approach with about 16 minutes for segmentation in native space and 3 minutes in the template space. AdaPro is the fastest approach with around 90 seconds to complete its entire pipeline in any coordinate space, which includes 20 seconds for preprocessing, 50 seconds for registration and histogram matching, 15 seconds for texture classification, and 5 seconds for object delineation, respectively.

3.5 CONCLUSION

In this chapter, we presented a fast and effective solution, named AdaPro, for the automated segmentation of brain structures in anomalous 3D MR images. AdaPro was used to delineate the brain hemispheres and cerebellum in 3D MR-T1 images of 2T and 3T from epilepsy patients before and after temporal lobe resections, being statistically more accurate and considerably more efficient than three atlas-

based methods, SOSM-S, MALF, and volBrain. Moreover, it can segment new images in the template and native image spaces.

Future work may extend AdaPro to other organs and imaging modalities, as well as evaluate other tissue classification for other different anomalies. Another worthwhile goal is investigating other preprocessing techniques to use a single version of AdaPro on 3D MR-T1 images regardless of the field strength.

We initially designed AdaPro to support collaborating neurologists to estimate morphological metrics for the cerebellum and brain hemispheres. The primary neurologists' interest was to study how these structures change in volume and shape in epilepsy patients after temporal lobe resection. Consequently, the brainstem was ignored during the creation of gold-standard segmentation of the considered in-house datasets, which made its evaluation impossible. However, we posteriorly obtained a new subset of MR-T1 of healthy images with gold-standard segmentation that also included the brainstem, which made it possible to extend AdaPro to segment all macro brain structures, as described in the next chapters.

Although many brain regions present *normal* brain asymmetries, studies have shown that some neurological diseases — such as Alzheimer’s [24], schizophrenia [25, 26], and epilepsy [27–29] — are associated to *abnormal* brain asymmetries (Fig. 1.2). Morphological changes in (sub)cortical structures in one or both hemispheres characterize such asymmetries. Detecting these abnormalities in MR brain images is useful, for example, to support neurologists during medical diagnosis, as well as to investigate structural changes in brain tissues after surgical procedures [29].

Most image analysis methods rely on the segmentation of target structures of interest (e.g., hippocampus) to quantify variations in shape, size, and texture [33, 185]. One can interactively segment these structures — which is error-prone and time-consuming — or rely on an automated segmentation tool, such as FreeSurfer [144] and volBrain [80] (Section 3.1). However, segmentation errors may severely affect the reliability of the computed morphometric measures and, consequently, the asymmetry analysis (Fig. 4.1).

Figure 4.1: Sagittal slice with an example of segmentation errors in the left hippocampus. The border (blue) of the object segmented by volBrain [80] is larger than the border (red) of the ground truth.

Other methods train a discriminative model (supervised learning) to classify the entire image or a volume of interest (VOI), for example, as *healthy* or *abnormal* [186–190]. The VOI analysis is more attractive for anomaly detection due to its flexibility in defining VOIs in any region of the brain, especially in those where there are no available segmentation tools. However, these methods typically re-

This chapter is based on the publication:

S. B. Martins, B. C. Benato, B. F. Silva, C. L. Yasuda, and A. X. Falcão, “Modeling normal brain asymmetry in MR images applied to anomaly detection without segmentation and data annotation,” in *SPIE Medical Imaging*, vol. 10950, pp. 71–80, 2019.

quire a considerable effort from experts to annotate a reasonable number of examples (images or VOIs), which are then used to train discriminative models. Moreover, the models are specific for the anomalies related to the disease(s) present in the training set.

In this chapter, we propose (i) an unsupervised framework to model *normal* brain asymmetries from healthy subjects — which explores the MR image analysis of corresponding VOIs in the left and right hemispheres — and (ii) the use of the model to detect abnormal asymmetries. As proof of concept, we instantiate the framework for the detection of abnormal hippocampal asymmetries from epilepsy patients. Then, we extend this framework for the entire hemispheres.

Our approach starts localizing VOIs around structures of interest. These VOIs may be defined as 3D bounding boxes from segmentation, whenever they are available. However, we propose a fast and accurate *3D patch-based model* for VOI localization in new images. Next, a *generative deep neural network* — a *convolutional autoencoder* (CAE) [191, 192] — is used to learn the image transformation from the left VOI to the flipped right VOI and vice-versa. We concatenate the outputs of the intermediate layers from CAE to form each observation (feature vector) of a *normal* structural asymmetry. Finally, we train a *one-class classifier* (OCC) to detect outliers as abnormal asymmetries. Our method localizes the corresponding VOIs in both hemispheres of a given a test brain image, extracts image features by CAE, and uses OCC to determine if the VOI pair represents normal or abnormal asymmetries.

We organize the chapter as follows. [Section 4.1](#) briefly provides an overview of autoencoders. [Section 4.2](#) details our proposed unsupervised approach. [Section 4.3](#) presents the experiments and results for abnormal hippocampal asymmetry detection. [Section 4.4](#) extends the proposed framework for abnormal asymmetry detection in the entire hemispheres, as well as it presents preliminary qualitative experiments. [Section 4.5](#) summarizes our findings and suggests future work.

4.1 AUTOENCODERS

Autoencoders (AE)s are *unsupervised* neural networks (also called generative neural networks) designed to reconstruct input data while compressing it in a low-dimensionality representation (*latent space*) [193]. AEs are based on the *encoder-decoder* paradigm [194]. The **encoder** stage transforms the high-dimensional input data into a latent low-dimensional representation (code). The **decoder** stage reconstructs the input data to the original space from the code. Thereby, the data is represented faithfully, but with removal of redundancies. [Fig. 4.2](#) presents the general structure of an autoencoder.

To formulate an AE model, let $\mathbf{x} = \{x_1, x_2, \dots, x_m\}$ be an unlabeled high-dimensional training set, where m is the number of samples, and $x_i \in \mathbb{R}^N$ is the i -th N -dimensional sample. The simplest autoencoder architecture consists of a

Figure 4.2: General structure of an autoencoder.

fully-connected feedforward neural network with a single hidden layer. AEs first map $\mathbf{x}$ to the latent space $\mathbf{z}$ (*encoder*) by using the encoding function, given by

$$\mathbf{z} = f_\theta(\mathbf{x}) = \mathbf{w}\mathbf{x} + b, \quad (4.1)$$

where $\theta = \{\mathbf{w}, b\}$, $\mathbf{w}$ are the encoding-layer weights, b is the bias, and $\mathbf{z} \in \mathbb{R}^d$ is the low-dimensional *latent representation (code)* of $\mathbf{x}$, $d \ll N$. This code is then used to *reconstruct* the input data (*decoder*) by a decoding function, given by

$$\mathbf{x}' = g_\phi(\mathbf{z}) = \mathbf{w}'\mathbf{z} + b', \quad (4.2)$$

where $\phi = \{\mathbf{w}', b'\}$, $\mathbf{w}'$ are the *decoding-layer weights*, b' is the *bias*, and $\mathbf{x}' \in \mathbb{R}^N$ is the resulting *reconstruction* of $\mathbf{x}$ from decoding $\mathbf{z}$. Putting together both equations:

$$\mathbf{x}' = g_\phi(f_\theta(\mathbf{x})). \quad (4.3)$$

The training of AEs consists of optimizing the parameters (θ, ϕ) by minimizing the *reconstruction errors* of the input data, given by

$$\theta^*, \phi^* = \arg \min_{\theta, \phi} \mathcal{L}(\mathbf{x}, \mathbf{x}'). \quad (4.4)$$

where $\mathcal{L}$ is a given *loss function*. Cross entropy and mean squared error (Appendix C.1) are examples of loss functions widely used to train AEs [193].

Medical image analysis applications typically use AEs for feature extraction [52, 58, 59, 195], image denoising [196], image reconstruction [197, 198]. Most of these methods rely on Convolutional Autoencoders (CAEs) [192], which extends the conventional AEs by incorporating convolutional layers [193]. CAEs share weights among all locations in the input images, preserving spatial locality. Consequently, its reconstruction results from the linear combination of basic image patches based on the latent code [192].

4.2 PROPOSED APPROACH

This section presents the steps involved in the creation of a model of normal structural brain asymmetry and its use for anomaly detection. Fig. 4.3 shows the general pipeline of our solution. The numbered steps in Fig. 4.3 are henceforth referred to as Step 1, Step 2, and so on, in the remainder of this chapter.

Figure 4.3: General pipeline of the proposed autoencoder-based approach to model normal brain asymmetries (left blue block) and to use that model (right pink block).

4.2.1 3D Image Preprocessing

The pipeline starts preprocessing the training set with only control images (Step 1) by following the operations described in [Section 2.3](#). We apply noise reduction by median filtering, followed by MSP alignment, and an inhomogeneity correction with N4 [\[102\]](#). We perform skull-stripping by AdaPro ([Chapter 3](#)), followed by intensity normalization. We consider the same reference image (template) used by AdaPro ([Section 3.2.1](#)), which has $180 \times 240 \times 240$ voxels and a voxel size of $1 \times 1 \times 1 \text{ mm}^3$. We then register all images on this template by non-rigid registration using Elastix [\[107\]](#). Finally, we perform intensity normalization to standardize the brain tissue intensities from the registered training images with the template. We apply the same preprocessing tasks to the testing images (Step 6).

4.2.2 VOI Localization

A *volume of interest* (VOI) is a sub-image with reduced background-size around a given structure under study. The simplest and most used case is a 3D patch, *i.e.*, an axis-aligned 3D box (parallelepiped). One could first segment the structure by using any automated segmentation tool, such as FreeSurfer [\[144\]](#), to define the VOI. However, it might be essential to analyze VOIs that do not include any specific object segmented by these available tools [\[40, 51\]](#). We then propose a patch-based model (PBM) to localize VOIs without any segmentation, as illustrated in [Fig. 4.4](#).

Firstly, one expert interactively specifies the left and right VOIs for each training control image in the reference space by selecting their first and last points ([Fig. 4.3](#), Step 2). Left and right localization models are VOIs of equal sizes, which is defined as follows. For each hemisphere, we compute the minimum bounding box that covers all its specified VOIs. The largest minimum bounding box defines the size of the localization models ([Fig. 4.3](#), Step 3). The initial location of each model is the geometric center of its respective VOIs in the training set. The differences in location among the training VOIs define possible translations within a small *search region* around the center of each model ([Fig. 4.3](#), Step 3). For a new test 3D image in the reference space, the method independently searches the translation from the center of each model, which maximizes the normalized mutual information (NMI) between the template and test image ([Appendix C.1](#)) inside the model ([Fig. 4.3](#), Step 7).

4.2.3 Normal VOI Asymmetry Representation

We aim at learning one transformation (reconstruction) from the left VOI to the right VOI, one for each hemisphere, and vice-versa. We first flip the right VOIs on the mid-sagittal plane to keep the same orientation between the left and right VOIs (Step 4). The reconstruction takes into account *normal asymmetries* that may affect any of the sides. The mean squared error between reconstructed and reals

Figure 4.4: Scheme for the training and use of the proposed patch-based model (PBM) for VOI localization in 3D brain images. For simplicity, the figure only considers VOIs in the right hemisphere.

VOIs expects to be minimum, given that the training VOI pairs represent normal asymmetries. We propose one *convolutional autoencoder* (CAE) for each transformation [192]. From a set of *source* 3D images, CAE learns to reconstruct a set of *target* 3D images by applying a set of linear and non-linear transformations (e.g., convolutions, activation, and pooling) to the *source* 3D images.

Given a training set with N images, we extract the desired flipped right and left VOIs (R_i, L_i) for each training image i . The considered *source* set consists of the union between the N flipped right VOIs and the N left VOIs. The *target* set, in turn, consists of the union between the N left VOIs and the N flipped right VOIs, ensuring that each flipped right VOI will be reconstructed as its corresponding left VOI and vice-versa (output from Step 5). Since CAE can reconstruct the desired images from the output of the last layer of their encoders, the concatenation of those intermediate outputs (*latent representation*) for the flipped right and left VOIs (output from Step 8) forms a suitable *feature vector* (normal VOI asymmetry representation).

Fig. 4.5 shows the considered architecture used for normal *hippocampal asymmetry* representation. Each VOI is a multivariate volume with $32 \times 64 \times 32 \times 1$ values. CAE contains three 3D convolutional layers with 16, 8, and 8 filters of $3 \times 3 \times 3$ weights each, respectively, followed by *ReLU* activation [199] and 3D max-pooling in the encoder, and the corresponding reconstruction operations in the decoder. Since the input images are normalized within $[0, 1]$, a *sigmoid* function is used as activation in the last layer of the decoder, and the output values are multiplied by 4095 to obtain the reconstructed image. The mean squared error between reconstructed and expected VOI is minimized by using the *nadam* gradient optimizer [200], which is popular in the field of deep learning mainly because it achieves good results fast. The reconstruction reaches errors around 0.0049 in the training set after 500 epochs.

Figure 4.5: Architecture of the convolution autoencoder (CAE) used for normal VOI asymmetry representation.

4.2.4 VOI Classification

The observations (vectors with 2048 features) from the training set lead to a cloud of points in $\mathbb{R}^n$, $n=2048$, each one representing a healthy subject. We must then train a one-class classifier to detect *outliers* as abnormal asymmetries. We have evaluated one-class classifiers based on Support Vector Machine (OC-SVM) [201] — due to its efficiency and robustness in classifying multidimensional data — and a simple variant of the Optimum-Path Forest clustering (OPF-clustering) [129], previously presented in Section 2.5. OC-OPF is detailed next.

One-Class OPF Classifier (OC-OPF)

A typical strategy to design outlier detection methods consists of first clustering a given training dataset so that if an unseen test data sample is far from the

clusters, from a given defined threshold, it is considered an outlier. In this work, we followed this strategy to design the one-class OPF (OC-OPF), an outlier detection approach based on OPF-clustering (Section 2.5). We chose this clustering algorithm primarily due to its good results in problems related to brain image analysis [134].

Recall $\mathcal{Z}$ be a training set interpreted as a graph, where each sample $s \in \mathcal{Z}$ is a node and $\mathcal{A}$ a given adjacency that defines the arcs between the nodes. Recall a path π_t with terminus t be a sequence $\langle s_1, s_2, \dots, s_m = t \rangle$ of nodes, such that $(s_i, s_{i+1}) \in \mathcal{A}$, $i \in [1, m]$, and $\pi_t = \pi_s \cdot \langle s, t \rangle$ be the concatenation of π_s and the arc $\langle s, t \rangle$ with the two joining instances of s merged into one.

OC-OPF starts clustering $\mathcal{Z}$ by defining its optimum-path forest, such that each obtained cluster is an optimum-path tree (as performed by OPF-clustering). We then remove clusters with a few samples (e.g., 5) to avoid noise in the training set, which makes the unsupervised classification more conservative — OPF-clustering does not originally perform this step. Given a new test sample $t \notin \mathcal{Z}$, the algorithm tries to assign one cluster to t by evaluating the values of the extended paths $\pi_s \cdot \langle s, t \rangle$, $\forall s \in \mathcal{Z}$. OC-OPF then extends OPF-clustering for *outlier detection* by considering that:

- If the cost of the found optimum path $\pi_t^* = \pi_s^* \cdot \langle s, t \rangle$ is greater than a given threshold, t is an *outlier*; otherwise, it is classified as *normal*.

Following the formal definition of OPF-clustering, as presented in Section 2.5, this threshold is defined by $\Omega(s)$. In this work, we consider that $\Omega(s)$ is the median distance to the k nearest neighbors of s . Section 4.3 compares OC-OPF and OC-SVM for detecting abnormal hippocampal asymmetries.

4.3 EXPERIMENTS AND RESULTS

This section describes the MR-T1 image datasets, experiments, and results obtained for the evaluation of the automatic abnormal hippocampal asymmetry detection. All computations were performed on the same Intel Core i7-7700 CPU 3.60GHz with 64GB of RAM, and an NVIDIA Titan Xp 12GB.

4.3.1 Datasets

The experiments considered five in-house datasets from healthy subjects and epilepsy patients with abnormal hippocampal asymmetries, as detailed in Appendix B.1. All images were provided by the Neuroimaging Laboratory (LNI) at the University of Campinas (UNICAMP), Brazil. Fig. 4.6 shows some examples of these datasets.

We separated 575 3D MR-T1 brain images of 3T from healthy subjects with ages between 25 and 65 years old (CONTROLS). All images have a voxel size of $1 \times 1 \times 1 \text{ mm}^3$ and do not present any structural lesion in any part of the brain.

Figure 4.6: Coronal slices that show: (a) normal hippocampal asymmetries, (b) left hippocampal atrophy, (c) right hippocampal atrophy, and (d) postoperative hippocampus.

We also considered 3D MR-T1 brain images of 3T, voxel size of $1 \times 1 \times 1 \text{ mm}^3$, from epilepsy patients which are divided into four in-house datasets: (PRE) preoperative patients with unilateral hippocampal atrophy (47 images); (POST) postoperative patients (88 images); (RHA) patients with right hippocampal atrophy (34 images); and (LHA) patients with left hippocampal atrophy (37 images). Therefore, the experiments involved a total of 781 images.

4.3.2 Localization Model

To validate our patch-based model (PBM) for VOI localization, we separated a subset of 60 control images from CONTROLS and the entire PRE and POST datasets. A neurologist from LNI has then manually delineated both hippocampi of all these images, generating the gold-standard segmentation for quantitative validation. This selection was made to avoid the high user-effort and required time to accurately segment both hippocampi in all available images (total of 781 images).

We used 50% of the selected control images for training and the remaining images from all the three datasets for testing, along 5 random splits of training and testing sets. For each training control image, the same neurologist has interactively specified its left and right VOIs (3D patches) around the hippocampi (Fig. 4.4). These VOIs are then used to train PBM.

We used two baselines for comparison. The first one (TEMP) assumes that the VOIs of the template, as interactively specified, can represent the VOIs of the test images, since they are all in the standard reference space. The second approach (SSEG) uses the minimum bounding box around each hippocampus as VOI, after segmentation by volBrain [80] (Section 3.1).

Table 5 presents the *percentage of missed foreground* — i.e., the percentage of voxels from the hippocampi that fall outside the detected VOIs — for each baseline. The analysis of variance (ANOVA), using the post-hoc Tukey honestly significant

difference (HSD) test, indicates that PBM and SSEG are equivalent for most cases, being both superior to TEMP (the p-value is 0.05). This result makes sense because the segmentation by volBrain tends to add background voxels around each hippocampus. It also proves that the background-size in PBM is reduced, with the advantage of not depending on segmentation. PBM also presents the best scores for control images.

Table 5: Percentage of missed foreground (%) of the considered localization approaches for the Right (RH) and Left Hippocampus (LH) in the considered datasets in 5 random splits. Bold values show better scores with statistical significance (**Note, for example, that 0.04 means 0.04% below**).

		CONTROLS	PRE	POST
RH	TEMP	1.78 ± 0.42	1.25 ± 1.87	1.42 ± 1.63
	SSEG	0.11 ± 0.38	0.15 ± 0.32	0.13 ± 0.36
	PBM	0.04 ± 0.13	0.11 ± 0.39	0.10 ± 0.45
LH	TEMP	2.01 ± 1.82	1.12 ± 1.35	1.94 ± 0.74
	SSEG	0.18 ± 0.32	0.07 ± 0.31	0.28 ± 0.38
	PBM	0.12 ± 0.25	0.09 ± 0.19	0.14 ± 0.28

4.3.3 Hippocampal Asymmetry Detection

In this section, we evaluate how accurate is the classification of hippocampal asymmetries of healthy subjects and epilepsy patients. We combined the convolutional-autoencoder (CAE) representation (Fig. 4.5) with each one-class classifier (OC-SVM and OC-OPF), which were trained in the *original feature space* and the *two-dimensional spaces* after non-linear projection by t-SNE [202], a well-popular projection algorithm widely used in several machine-learning problems.

The reason for considering a projection space in the experiments is that, by analogy with AEs, t-SNE (or, actually, any other projection method) is a similar kind of tool that reduces dimensionality. The only differences with regards to AEs are that (i) the latent space is 2D, and (ii) the cost/error functions are quite different – projections aim to preserve *relations* between samples, and not the *information* encoded in *each* sample (in the sense of being able to decode the n-D sample from the low-D one). Besides, according to Rauber *et al.* [203], the separation among groups (classes) in the two-dimensional projection space, as created by the t-SNE, is a strong indication of their separation in the original feature space.

An important aspect of t-SNE is that every execution of the algorithm creates a distinct projection. Therefore, it can only make sense when there is a set of testing images to be classified. As an unsupervised technique, the t-SNE algorithm can

project the training and testing sets with no knowledge of true labels, and then the one-class classifier can be trained in the projection space and used to classify the testing samples.

Experimental Protocol

Initially, we used PBM (Section 4.2.2) to automatically localize both hippocampi (VOIs) in all datasets (Section 4.3.1). We then used the resulting VOIs for feature learning, extraction, and the design of the classifiers. We considered 5 evaluation splits on CONTROLS by randomly selecting 240 controls images for *training*, 60 for *validation*, and the remaining 275 images for *testing*. We also considered all other images from epilepsy patients for testing. Finally, we compared our CAE-based representation with the simple *absolute difference* (ABS-DIFF) between left and flipped right VOIs.

Since our framework is *unsupervised*, we can only count on the training set of *controls images* for parameter optimization. Consequently, we optimized the parameters of each classifier by grid search aiming at maximizing the *detection scores* for these images — *i.e.*, healthy hippocampi classified as normal hippocampal asymmetries. The best parameters found for each case were:

- ABS-DIFF/OC-SVM: kernel = *rbf*, nu = 0.01, and gamma = 0.007;
- CAE/OC-SVM: kernel = *linear*, and nu = 0.1;
- ABS-DIFF/OC-OPF: the best k was found within [20, 45];
- CAE/OC-OPF: the best k was found within [15, 45].

Results in the Original Feature Space

Table 6 presents the average detection scores in the *original feature space* for hippocampal asymmetries. These scores show the percentage of the classification hits of controls images as normal hippocampal asymmetries and patient images as abnormal hippocampal asymmetries.

Although the difference between the best scores of ABS-DIFF and CAE is small for the CONTROLS dataset, CAE-based representation is consistently superior to ABS-DIFF in all cases. These results confirm the efficiency of CAE in providing representative features for the addressed problem. The combination CAE/OC-SVM provides slightly better accuracies than CAE/OC-OPF in most datasets, with no errors for PRE and POST.

Results in the Two-dimensional Projection Space

To further improve the accuracy of the CAE-based representation, we repeated the same experiments in the *two-dimensional projection* spaces created by the t-SNE algorithm. Each sample is now represented by two features resulting from

Table 6: Anomaly detection scores (%) in the original feature space for the instances of the proposed framework. These scores show the percentage of the classification hits of controls images as normal hippocampal asymmetries and patient images as abnormal hippocampal asymmetries.

	Original Feature Space			
	CAE		ABS-DIFF	
	OC-SVM	OC-OPF	OC-SVM	OC-OPF
CONTROLS	86.61 $\pm$ 2.30	89.04 $\pm$ 2.55	87.45 $\pm$ 1.34	86.35 $\pm$ 1.13
PRE	100.0 $\pm$ 0.00	98.40 $\pm$ 2.04	93.62 $\pm$ 0.00	61.70 $\pm$ 9.67
POST	100.0 $\pm$ 0.00	100.0 $\pm$ 0.00	22.44 $\pm$ 0.57	100.0 $\pm$ 0.00
RHA	99.26 $\pm$ 1.47	97.06 $\pm$ 4.16	94.85 $\pm$ 1.47	42.65 $\pm$ 5.09
LHA	99.32 $\pm$ 1.35	97.30 $\pm$ 5.41	94.59 $\pm$ 0.00	27.70 $\pm$ 8.37

this projection. As a non-linear and unsupervised projection approach, we must project training and testing sets together to design the classifiers from the training samples and use them to label the testing set. This approach is indicated whenever there is a set of testing images to be tagged.

Table 7 shows the resulting detection scores for the projection space. This strategy has a surprisingly positive impact on the results: the detection scores for control images have considerably increased (99.72% against 89.04% for the original space), and there are no classification errors for patient images in both classifiers.

Table 7: Anomaly detection scores (%) in the projection space for the instances of the proposed framework. These scores show the percentage of the classification hits of controls images as normal hippocampal asymmetries and patient images as abnormal hippocampal asymmetries.

	Projection Space by t-SNE	
	CAE / OC-SVM	CAE / OC-OPF
CONTROLS	96.22 $\pm$ 2.72	99.72 $\pm$ 0.18
PRE	100.0 $\pm$ 0.00	100.0 $\pm$ 0.00
POST	100.0 $\pm$ 0.00	100.0 $\pm$ 0.00
RHA	100.0 $\pm$ 0.00	100.0 $\pm$ 0.00
LHA	100.0 $\pm$ 0.00	100.0 $\pm$ 0.00

Since the datasets have different sizes, Table 8 presents the mean values of Cohen Kappa for each feature space taking into account all datasets. Results confirm that the design of both classifiers for hippocampal asymmetry detection, based on the proposed unsupervised framework, is more accurate on the projection space. Furthermore, the proposed OC-OPF classifier outperforms OC-SVM in both spaces.

Finally, to better understand these results and their impact on developing intelligent and interactive virtual environments in neuroscience, Fig. 4.7 presents

Table 8: Cohen Kappa for the considered framework instances with PbM localization.

	CAE / OC-SVM	CAE / OC-OPF
Original Space	0.846 ± 0.03	0.863 ± 0.02
Projection Space	0.956 ± 0.03	0.995 ± 0.00

a t-SNE projection space from one of the splits. Here, each point consists of hippocampal asymmetries of a given pair of hippocampi (VOI) as extracted by CAE. By clicking on any sample (point on the plot), the user can see the corresponding slice across the centers of the VOIs. It is also possible to navigate in the image around that location for inspection and annotation of the anomaly type. From that, one can train discriminative neural networks to allow the detection of abnormal brain asymmetries, followed by the identification of their anomaly type.

Figure 4.7: The 2D t-SNE projection from the considered datasets for the CAE-based representations with PbM localization. The expert can select any observation from the projection to inspect its image slices (see cropped slices of selected observations).

Fig. 4.7 shows a clear visual separation between the projected features from normal hippocampal asymmetries (green circles) and the abnormal cases (orange marks). Rauber *et al.* [203] showed that the visual separability of classes in a t-SNE projection is highly correlated with the ability of a classifier to separate classes in the original feature space. Hence, we can conclude that both classes are also well-separated in the original space, which also confirms the high detection scores in this space, as presented in Table 6.

Note that Fig. 4.7 also highlights a false-positive case — the bottom-right green circle associated with the bottom cropped slice — *i.e.*, a sample of normal hippocampal asymmetries classified as abnormal. Although this test sample did not have much impact on the final accuracy, its possible use as a training one is problematic for the design of the one-class classifiers, leading to worse detection scores. This impact may be severe for OC-OPF since it does not have any treatment for noise in the training set. Therefore, OC-OPF deserves further investigation to prevent this scenario.

The computational time to analyze a new image takes in the worst case (with the slowest methods) around 92 s, which is a proper time for the clinical routine. This time includes 80 s for preprocessing, 4 s for VOI localization by PBM, 7 s for t-SNE projection, 0.5 s for feature extraction, and 0.5 s for classification.

4.4 EXTENSION FOR BRAIN ASYMMETRY DETECTION

This section presents our first attempt to detect abnormal brain asymmetries automatically. For that, we have extended the proposed unsupervised abnormal hippocampal asymmetry detection for the entire hemispheres, as detailed next.

4.4.1 Proposed Extension

Consider a *template* T with a pre-defined *segmentation mask* M for its right hemisphere. We start automatically selecting pairs of VOIs along the entire hemispheres, as illustrated in Fig. 4.8. We apply a uniform grid-sampling within M and select the corresponding symmetric voxels in the left hemisphere according to the mid-sagittal plane of T . This forms the set $C = \{(CR_1, CL_1), (CR_2, CL_2), \dots, (CR_m, CL_m)\}$ with m pairs of symmetric voxels, where (CR_i, CL_i) is the i -th selected voxel in the right and left hemisphere, respectively.

Figure 4.8: Uniform grid-sampling used to define the geometric centers of VOIs along the hemispheres. The dashed red lines in the slices correspond to mid-sagittal planes (MSPs).

We consider the voxels of each pair $(CR_i, CL_i) \in C$ as the initial geometric centers of the corresponding VOIs whose asymmetries will be analyzed. We then perform our previous method (Section 4.2) for each of these pairs, independently, with a simple modification: all VOIs along the hemispheres have the same size so that the *localization model* now only consists of a local search within a fixed *search region* of $5 \times 5 \times 5$. Note that, consequently, this extended methodology analyzes asymmetries along the entire hemispheres by performing multiple executions of our previous method, according to the number of pairs selected in C .

4.4.2 Preliminary Experiments

We carried out preliminary experiments to qualitatively evaluate the extended methodology for brain anomaly detection. We adopted the same network architecture as illustrated in Fig. 4.5 so that we fixed the size of $32 \times 64 \times 32$ for all VOIs extracted along the hemispheres. We considered the training set of control images and the template used in Section 4.3, and the one-class classifier with the optimized parameters found for the hippocampi.

Fig. 4.9 shows axial slices with resulting detected abnormal asymmetries (VOIs) for (a) a control image from CONTROLS, (b) a postoperative image from POST, and (c) a patient image from LHA with a small lesion in the left hemisphere. All these detected VOIs are, in fact, false positives, *i.e.*, normal asymmetries classified as abnormal. Note that all of them are close to the hemispheres' borders, which are locations typically asymmetric. The control image is supposed not to present any abnormal asymmetries, whereas the method should detect the anomalies (orange circles) in the remaining images. By inspecting other slices from these 3D images, we also found several other false-positive detected VOIs.

Figure 4.9: Examples of false-positive abnormal asymmetries (red patches) detected by the proposed method. Axial slice of (a) a control image, (b) a postoperative image after temporal lobe resection (orange circle), and (c) an epilepsy patient with a small lesion in the left hemisphere (orange circle). Each patch in one hemisphere has a corresponding patch in the other hemisphere.

To better understand the reasons for these false positives, we first inspected the VOIs that cover the anomaly (removed portion) of the postoperative image of Fig. 4.9b. Fig. 4.10 shows an axial slice with the pair of undetected VOIs (blue patches) with the highest intersection with the anomaly (orange circle). Since the anomaly is not entirely inside the VOIs, its asymmetry representation is poor, and, consequently, the method cannot detect it. Thus, we have a problem: **how to estimate VOI positions adequately?** This task is indeed hard — since anomalies can be found in different locations in the brain — and widely investigated in the computer-vision literature [204, 205].

Figure 4.10: Pair of undetected VOIs (blue patches) intersecting an anomalous region (orange circle) of a postoperative image. Note that the anomaly is not entirely inside the VOIs, which undermines the effectiveness of your asymmetry representation.

We now inspected the VOIs that cover the small lesion of the epilepsy patient image of Fig. 4.9c. Although a VOI entirely covers this lesion (Fig. 4.11), the huge amount of background information within this VOI impairs the asymmetry representation of the lesion. In the example of Fig. 4.11, the volume of the anomaly is only about 5% of its covering VOI volume. Consequently, the asymmetry representation is biased for the background, which leads to another problem: **how to estimate VOI sizes effectively?** Each brain structure and anomaly have different shapes that demand different VOI sizes, which makes their automatic estimation challenging.

Given the high complexity of effectively defining 3D patches, we decided to use *supervoxel segmentation* to define volumes of interest for asymmetry analysis. Supervoxels offer significant freedom and control on defining such asymmetric volumes. At the same time, their application comes with additional complications such as algorithmic choices and parameter setting choices. As such, we explore this approach separately in two dedicated chapters, [Chapters 5](#) and [6](#).

Figure 4.11: Pair of undetected VOIs (blue patches) covering a small lesion (orange circle) of an epilepsy patient image. Although the lesion is completely covered by the VOI, the huge amount of background information inside the VOI impairs the lesion asymmetry representation.

4.5 CONCLUSION

This chapter presented our first solution for addressing the research questions of the thesis. From a simple interactive specification of a 3D patch around a structure of interest in training control images, we presented an unsupervised framework that defines (i) the corresponding VOI localization model; (ii) image features based on a generative deep neural network (convolutional autoencoder) for asymmetry representation; and (iii) trains a one-class classifier that corresponds to our *normal* brain asymmetry model, thus addressing our first research question **RQ1**. We use this one-class classifier to detect outliers as abnormal brain asymmetries, which addresses our second research question **RQ2**. Our solution does not need the segmentation of the target structures and data annotation. We also proposed a novel one-class classifier based on the Optimum Path Forest algorithm [134].

We first instantiated and validated the proposed unsupervised framework to detect abnormal *hippocampal* asymmetries in MR-T1 3D images of healthy subjects and epilepsy patients with unilateral hippocampal atrophies. Experimental results by using the original feature space and a two-dimensional space showed high detection scores and kappa values, especially considering some difficult cases that only a trained expert can visually identify.

Next, we extended the framework for the detection of brain asymmetries along the entire hemispheres. From a set of pairs of VOIs automatically extracted in the hemispheres by uniform grid-sampling, we performed our initial solution for each pair. We selected a few control and epilepsy patient images for preliminary qualitative validation. Experiments show poor results with several false-positive asymmetries detected for all considered images (Fig. 4.9). The extended approach did not also detect true abnormal asymmetries caused by lesions. We found two main reasons for these: (i) 3D patches did not entirely cover anomalies due to poor localization; (ii) there was a high amount of background information inside

the 3D patches. The effective estimation of VOI localization and size is indeed challenging and demands further investigation. Therefore, we indicate our proposed framework for the analysis of well-defined structures of interest inside the brain, where one carefully specifies training 3D patches around these structures. All these limitations motivated us to use *supervoxel segmentation* to estimate VOIs for brain asymmetry detection, as detailed in the next chapter.

For future work, we first suggest further investigating the impact of the t-SNE projection in the design of classifiers and the extension of the framework to other well-defined brain structures. Second, one should explore faster (and yet robust) projection algorithms than t-SNE in order to make the asymmetry analysis in the two-dimensional space feasible for large-scale studies. Third, the proposed approach could also be evaluated for other organs and/or by using different medical imaging modalities. Finally, one may develop interactive visual tools to support the detection, inspection, annotation, and identification of brain anomalies based on abnormal asymmetries.

In [Chapter 4](#), we have presented our first unsupervised solution for addressing the research questions of the thesis. This solution models healthy asymmetries between a pair of 3D patches — volumes of interest (VOIs) — by using a convolutional autoencoder [192] and a one-class classifier, detecting *outliers* as abnormal asymmetries associated to anomalies. Although this first approach was well-succeed in detecting abnormal *hippocampal* asymmetries in epilepsy patients, its result for the remaining regions of the hemispheres was unsatisfactory (Section 4.4.2). Two main reasons explain these results: (i) 3D patches did not entirely cover anomalies due to poor localization; (ii) there was a high amount of background information inside the 3D patches. The effective estimation of sizes and localization of 3D patches are indeed challenging and widely investigated in the computer-vision literature [204, 205]. An alternative strategy relies on using *supervoxel segmentation*.

Supervoxels are groups of voxels with similar characteristics resulting from an oversegmentation of a 3D image or region of interest. They preserve intrinsic image information (e.g., the borders of tissues and lesions) and are used as an alternative to patches to define more meaningful VOIs for computer-vision problems [135, 136] and some medical image applications [6, 137]. For example, one can oversegment the hemispheres in multiple supervoxels for brain anomaly detection. Supervoxels overcome the two main problems of 3D patches for our target problem, as (i) they better fit lesions and tissues to VOIs, and (ii) their voxels contain minimum heterogeneous information. The irregular shapes of supervoxels, however, prevent the use of recent deep-learning-based techniques that only work with regular 2D/3D patches.

This chapter proposes an automatic unsupervised supervoxel-based framework for detecting abnormal asymmetries associated with anomalies in 3D MR brain images. By having all images registered to the same symmetric template, the proposed framework, called *Supervoxel-based Abnormal Asymmetry Detection* (SAAD), computes asymmetries between hemispheres by using their mid-sagittal

This chapter is based on the following publications:

- (i) S. B. Martins, G. Ruppert, F. Reis, C. L. Yasuda, and A. X. Falcão, “A supervoxel-based approach for unsupervised abnormal asymmetry detection in MR images of the brain,” in *IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 882–885, 2019;
- (ii) S. B. Martins, A. C. Telea, and A. X. Falcão, “Extending supervoxel-based abnormal brain asymmetry detection to the native image space,” in *IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 450–453, 2019;
- (iii) S. B. Martins, A. C. Telea, and A. X. Falcão, “Investigating the impact of supervoxel segmentation for unsupervised abnormal brain asymmetry detection,” *Computerized Medical Imaging and Graphics*, vol.85, 101770, 2020.

plane as reference. It then extracts pairs of *symmetric supervoxels* from the left and right hemispheres for each test image, so that each pair generates a one-class classifier. This classifier is next trained on control images of healthy subjects to find supervoxels with abnormal asymmetries. We propose a supervoxel segmentation method, named SymmISF, based on the *Iterative Spanning Forest* (ISF) framework [61], to extract symmetrical supervoxels from left and right hemispheres. We also extend SAAD to perform in the *native* image space of testing images. Both versions of SAAD are validated to detect abnormal brain asymmetries in 3D MR-T1 images of stroke patients.

The remaining of this chapter is organized as follows. Section 5.1 shows the related work of brain anomaly detection/segmentation. Section 5.2 presents SymmISF and SAAD for abnormal brain asymmetry detection on *standard* image space, whereas Section 5.3 details experimental protocols, and Section 5.4 discusses the results. Section 5.5 presents the extension of SAAD for the *native* image space, as well as reports experimental results. Finally, Section 5.6 summarizes our contributions and suggests future work.

5.1 RELATED WORK

From a certain point of view, automatic brain lesion detection/segmentation methods can be grouped into five classes. From the least to the most versatile, these are as follows.

5.1.1 Atlas-based Methods

These methods use the *a priori* knowledge about the object’s shapes in a training atlas set registered on a standard template, where each atlas consists of a source 3D image and its corresponding 3D label image with the mask of each 3D object of interest [41, 42, 80]. Directly encode anomaly shape-constraints from an atlas set and use these models to segment anomalies is not effective as they vary greatly in size, shape, and location [147] (Fig. 1.3). Thus, some atlas-based methods aim to combine these prior shape-constraints and texture segmentation to precisely delineate anomalies [147, 206–210].

Some solutions incorporate the prior healthy shape-constraints in a supervised framework [207–209]. For example, Zijdenbos *et al.* [207] proposed an automatic technique for multiple sclerosis segmentation that trained an artificial neural network with images from healthy subjects and patients. The input features included three MRI modalities and three spatial-tissue priors from probabilistic atlases. In contrast, Prastawa *et al.* [210] use atlases to detect outliers as anomalies by selecting abnormal tissue samples while estimating healthy ones. They proposed using the Minimum Covariance Determinant to estimate probability density functions for CSF, WM, and GM using healthy images. Any outlier to this estimation

is considered an abnormal tissue. The next sections detail more supervised and unsupervised approaches for anomaly detection/segmentation.

Although atlas methods show impressive segmentation results for healthy tissues [41, 80, 144, 183], their results for anomaly segmentation are still unsatisfactory, especially in the presence of anomalies with arbitrary shapes and locations.

5.1.2 Supervised Learning with Hand-crafted Features

These methods use different classifiers trained from various *hand-crafted* image features (e.g., edge detectors and texture features) to delineate anomalies by classifying voxels or regions of the target image [6, 37–39]. Such supervised methods usually do not generalize well for anomalies from different disorders since their considered features have limited representation capability considering the significant variation of the anomalies’ appearances. Moreover, these methods work well only for detecting anomalies related to diseases present in the training set.

For example, Goetz *et al.* [37] rely on ExtraTrees classifiers [211] to classify voxels as healthy or tumors. From four different MRI modalities, the method computes 54 hand-crafted features (e.g., local histograms, and first-order statistics) for each voxel. Soltaninejad *et al.* [6], in turn, proposed to segment anomalies in FLAIR images by extracting and classifying superpixels across slices of FLAIR images. The method extracts several hand-crafted features for each supervoxel, such as intensities, fractal, and curvature features. This solution is accurate to segment medium and large anomalies in FLAIR images *only*, where such lesions are considerably highlighted. This thesis, however, aims to automatically detect lesions (especially the small ones) in 3D MR-T1 images.

5.1.3 Discriminative Deep Learning

These techniques have emerged as a powerful alternative to the previous class of methods, given their ability to learn highly discriminative features for a particular task. In particular, Convolutional Neural Networks (CNNs) [212] have become a mainstay of the computer vision community due to breakthrough performance in several applications [213] as compared to approaches using hand-crafted features. Deep learning has gained popularity in medical image analysis as well [35, 45, 46, 54, 214]. Such methods learn deep feature representations (e.g., convolutional features) in a data-driven way without any feature engineering being required. Nevertheless, deep-learning-based methods have some limitations:

- (a) they require a large number of training images that must be previously annotated by specialists (e.g., lesion segmentation masks);
- (b) they typically require weight fine-tuning (retraining) when used for a new set of images due to image variability across scanners and acquisition protocols;

- (c) they are only designed for the anomalies found in the training set, just as the supervised methods outlined before;
- (d) the success of such methods on new images is limited by the absence of large, high-quality, annotated training sets for most medical image analysis problems [11].

5.1.4 *Unsupervised Approaches*

All the above limitations of supervised methods motivate research on *unsupervised* anomaly detection approaches based on *outlier detection* [13, 40, 53, 58, 215]. These methods aim to learn a model from control images of healthy subjects *only* by encoding general knowledge or assumptions (*priors*) from healthy tissues. This model is next used to guide brain segmentation so that *outliers* who break such general priors are considered anomalies [13].

For example, Shen *et al.* [215] proved that the voxel-intensity-based segmentation and the spatial-location-based tissue distribution (based on a probabilistic atlas) in the lesions are inconsistent with those in healthy tissues. They use the conventional Fuzzy C-Mean algorithm [216] and probabilistic maps from a template to quantify such inconsistencies, and then apply a threshold to obtain a binary lesion segmentation. Juan-Albarracin *et al.* [78], in turn, propose a more complex method that uses four different MRI modalities of a given patient to segment the brain tumors of the BraTS dataset [73]. The method groups all voxels in a few clusters that are classified as normal or outlier, based on probabilistic maps from a template.

As unsupervised brain anomaly detection methods do not use labeled samples, they are less effective in detecting lesions from a specific disease when compared to supervised approaches trained from labeled samples for the same disease. For the same reason, however, unsupervised methods are generic in detecting any lesions, *e.g.*, coming from multiple diseases, as long as these notably differ from healthy training samples. Our proposed methods, presented in the next sections, aim to combine the pros and cons of unsupervised learning for the detection of anomalies associated with abnormal brain asymmetries.

5.1.5 *Deep Generative Neural Networks*

Also known as Encoder-Decoder Neural Networks or autoencoders (AEs; [Section 4.1](#)), these methods have been used for unsupervised anomaly detection by modeling the distribution of healthy brain tissues and next detecting anomalies as outliers. The underlying hypothesis is that this model can reconstruct normal brain anatomies while failing to reconstruct anomalies in images with some disorder. AEs learn to reconstruct training images from healthy individuals *only* by first compressing (encoding) them into a low-dimensional representation (latent

features) and then decompressing that representation to minimize the reconstruction error between the input data and its reconstruction.

Some methods [53, 58, 59, 217, 218] delineate anomalies by thresholding the resulting reconstruction errors, *i.e.*, the residual image between the input image and its reconstruction. Baur *et al.* [53] determine as threshold a given percentile of the reconstruction errors on the training control images. Chen *et al.* [59] evaluate different AEs for the reconstruction of MR-T1 and T2 images, but they do not detail how they choose the thresholds. Other AE-based methods [52] train a one-class classifier from latent features to classify if an image (or region of interest) has some anomaly [52]. Although all these methods can detect large lesions in MR-T2, FLAIR, and CT [53, 58, 59], they show inferior results in MR-T1 images and completely fail with small lesions, which are the most challenging cases. In Section 5.4, we compare our proposed solution in the standard image space with an AE-based approach derived from Baur *et al.* [53] and Chen *et al.* [59].

5.2 DESCRIPTION OF SAAD

We next describe the SAAD method (Fig. 5.1) for abnormal brain asymmetry analysis in the *standard* image space (SIS), *i.e.*, the coordinate space of a template. SIS is chosen when considering a group of reference images from healthy subjects and/or patients for comparison during the analysis. All images are registered in the same coordinate, ensuring spatial correlation (position and shape) among same structures of interest across all images.

SAAD consists of four steps: 3D image preprocessing, asymmetry computation, supervoxel segmentation, and classification, described next.

5.2.1 3D Image Preprocessing

The pipeline first preprocesses the training control image set and the test image (Steps 1 and 4) to improve image quality for subsequent analysis. Fig. 5.2 illustrates the preprocessing results for a given stroke 3D image. We follow the same preprocessing steps presented in Section 2.3, summarized as follows.

For each 3D image (Fig. 5.2a), we start performing noise reduction by median filtering, followed by MSP alignment and bias field correction by N4 [102]. As voxels from irrelevant tissues/organs for the addressed problem (*e.g.*, neck and bones) can negatively impact the image registration and intensity normalization, we use AdaPro (Section 3.2) to segment our macro-regions of interest: right and left hemispheres, cerebellum, and brainstem (Fig. 5.2b). To correctly estimate asymmetries from anomalies, we ignore the tissue classification step of AdaPro, so that anomalies and typically dark tissues (*e.g.*, CSF) are now also considered during analysis.

To attenuate differences in brightness and contrast among images, we perform intensity normalization by first linearly mapping all intensities within [0, 4095] followed by a histogram matching between the segmented images and the template

Figure 5.1: The pipeline of SAAD. Steps 1 to 3 (blue part) are performed offline. Steps 4 to 8 (pink part) are computed for each 3D test image (detection stage). The output images from Steps 3, 5, 6, and 8 are visualized as a symmetrical image. However, the method can consider just one hemisphere.

Figure 5.2: 3D image preprocessing and registration steps. (a) Axial slice of a raw test 3D image. The dashed line shows its mid-sagittal plane (MSP) and the arrow indicates a stroke lesion. (b) Test image after noise filtering, MSP alignment, bias field correction, and brain segmentation. (c) Axial slice of the symmetric brain template (reference image). (d) Histogram matching between (b) and the template (intensity normalization). (e) Final preprocessed image after non-rigid registration and histogram matching with the template.

inside its predefined brain segmentation mask (Figs. 5.2c–d). This operation only considers voxels inside the brain.

We perform non-rigid registration to place all images in the coordinate space of the ICBM 2009c Nonlinear Symmetric template [108], popularly known as MNI, whose hemisphere masks and MSP are well defined. We rely on Elastix [107] to

perform all registrations.¹ Finally, we apply another histogram matching between the registered images and the template, and we use the brain segmentation mask from the template and its MSP to separate the left and right brain hemispheres in each image for further asymmetry computation (Fig. 5.2e).

5.2.2 Asymmetry Computation

Let X be the set of registered training 3D images (output of Step 1) and I the test 3D image after preprocessing (output of Step 4). We obtain the set of asymmetry maps A_X for all X by computing the voxel-wise absolute differences between left and right hemispheres with respect to the template's MSP (Step 2). One might argue on estimating the MSP of each registered image, separately. However, we confirmed that, when mapping all images to the same *symmetric* template by *non-rigid* registration, each registered image's MSP was either exactly or slightly tilted from the template's MSP, which does not impact the asymmetry computation.

Next, we create a *normal asymmetry map* $\overline{A_X}$ (Step 3) by averaging the absolute difference values of A_X (Fig. 5.3a). We use this map to reduce the detection of false-positive asymmetries in I in commonly asymmetric brain regions (e.g., cortex), as detailed next in Section 5.3. Finally, we compute voxel-wise absolute differences between the hemispheres for I (Figs. 5.3b–c) and then subtract $\overline{A_X}$ from them. Resulting positive values form a final asymmetry map A_I (Fig. 5.3d) for the test image I (Step 5).

Figure 5.3: Asymmetry computation on a standard image space. (a) Axial slice of the normal asymmetry map for healthy subjects. (b) Axial slice of a 3D test stroke image after preprocessing and non-rigid registration on a symmetric template. (c) Asymmetries of (b) by computing voxel-wise absolute differences between the hemispheres with respect to its MSP. (d) Final attenuated asymmetries: positive values of the subtraction between (c) and (a).

¹ We used the *par0000* files available at http://elastix.bigr.nl/wiki/index.php/Parameter_file_database

5.2.3 Symmetric Supervoxel Segmentation

Directly comparing the flipped, segmented, and registered hemispheres is not helpful as it will not tell us where *small-scale* asymmetries occur. At the other extreme, comparing every voxel pair in these hemispheres is risky, since individual voxels contain too little information to capture asymmetries. These difficulties motivate the use of *supervoxels* as the unit of comparison (Step 6).

An ideal supervoxel segmentation should create precisely one supervoxel per anomaly. This case is, of course, highly unlikely to succeed, given the high variability of size, shape, and position of anomalies (Fig. 1.3). At any rate, too small supervoxels should be avoided as they *oversegment* larger anomalies and thus cannot capture their essence, and also will confuse the end-users when visually exploring the results. Too large supervoxels, in contrast, should be avoided as they cannot precisely delineate small-scale anomalies from the background (*undersegmentation*).

We propose a new method, named *SymmISF*, that extracts symmetrical supervoxels from left and right brain hemispheres simultaneously. *SymmISF* is based on the recent *Iterative Spanning Forest* (ISF) framework [61] (Section 2.6) for superpixel segmentation and has three steps: (i) seed sampling followed by multiple iterations of (ii) connected supervoxel delineation, and (iii) seed recomputation to improve delineation (Fig. 5.4), as follows.

As outlined in Section 2.6, *initial seed estimation* is a crucial step for the success of ISF. The adopted strategy for that, however, depends on the addressed problem, which, in turn, may have specific constraints. For the problem of this chapter, the simplest approach to find initial seeds is to select N seeds uniformly distributed in the right hemisphere defined by a segmentation mask for the template. We call this strategy *Uniform SymmISF next*. However, there are no guarantees that this strategy will place *at least* one seed within each asymmetric anomaly, so this can easily lead to undersegmentation. We then propose a new strategy, called next *Asymmetry-guided SymmISF*, that is guided by the hemispheric asymmetries of the image when selecting one seed per local maximum in A_I (see the asymmetry-guided seeds in Fig. 5.4). It computes the local maxima of the foreground of a binarized A_I at $\gamma \times \tau$, where τ is Otsu's threshold [219]. The higher the factor γ is, the lower is the number of asymmetric components in the binarized A_I . This seed-set is next extended with a fixed number of seeds (e.g., 100) by uniform grid sampling the low-asymmetry regions of the binarized image. A detailed comparison of Uniform SymmISF with Asymmetry-guided SymmISF is presented next in Section 5.4.

As the cortex is typically very asymmetric, we can still remove seeds placed very close to the hemisphere borders to reduce the number of false positives in such regions and also to weigh the normal asymmetry map to attenuate other asymmetries further. Both strategies are evaluated and discussed in detail in Section 5.4.

Figure 5.4: The pipeline of SymmISF with two possible initial seed estimation strategies (red points represent seeds). The method flips the input test 3D image (volume) by using its MSP and builds a 2-band volume by stacking both volumes. Then, the ISF framework [61] estimates supervoxels inside the hemisphere mask from the initial seeds. The resulting label map is flipped to form the final label map with pairs of symmetrical supervoxels.

By stacking the right hemisphere with the left hemisphere — flipped using the MSP — as the input 2-band volume (Fig. 5.4), SymmISF applies ISF only inside the right hemisphere from the initial seeds. ISF relies on a cost function controlled by two parameters: α and β (Section 2.6.1). This process yields a label map in which each supervoxel is given a distinct label. Finally, SymmISF flips these supervoxels to obtain the symmetrical supervoxels in the left hemisphere, which yields the final label map L (output of Step 6). Note that one can proceed conversely, *i.e.*, apply SymmISF on the left hemisphere, and map the result to the right hemisphere.

5.2.4 Feature Extraction and Classification

SAAD presents a novel approach for outlier detection (Section 1.3.1) — here instantiated for abnormal asymmetry detection — that designs a set of *specialized* one-class classifiers (OCCs) specific for each test 3D image, as shown in Fig. 5.5. For each 3D test image, each pair of symmetrical supervoxels is used to create a *specialized* one-class classifier (OCC) using as feature vector the *normalized histogram* of the asymmetry values inside the pair (Step 7). Classifiers are trained

from *control images* only, thus locally modeling normal asymmetries for the entire hemispheres. Finally, SAAD uses the trained OCCs to find supervoxels with abnormal asymmetries in the corresponding testing image (Step 8). Fig. 5.6 illustrates the supervoxel classification.

By default, SAAD yields *pairs* of symmetric supervoxels corresponding to the detected abnormal asymmetries. This output is useful for subsequent visual analysis as an expert can compare such regions in both hemispheres as well as their computed asymmetries. To output only the supervoxel that covers the detected asymmetric anomaly, one may simply compute the similarity from the test image with the template inside each supervoxel of the pair. The less similar supervoxel contains the anomaly.

Figure 5.5: One-class classifier (OCC) training for abnormal asymmetry detection. For each pair of symmetric supervoxels from a given test 3D image, SAAD trains an OCC from the training normal asymmetry maps previously computed.)

When *dynamically* designing specialized one-class per-supervoxel classifiers for each test image, SAAD implicitly considers the *position* of the supervoxels in the hemispheres when deciding upon their asymmetry. The central premise for this is that a *single* global classifier cannot separate normal and anomalous tissues by only using texture features. Experimental results concerning this hypothesis are presented in Section 5.4.3.

Even though the proposed classification scheme demands a higher processing time compared to using a single global classifier trained offline, this time is not too high (≈ 2 min) and still feasible for clinical purposes as SAAD relies on a simple and fast feature extraction (histogram) and the one-class linear Support Vector Machine [201]. More details are presented in Section 5.4.

Figure 5.6: Abnormal asymmetry detection of a test 3D image by supervoxel classification. For each pair of symmetrical supervoxels, SAAD uses the corresponding one-class classifier to classify the asymmetries inside the pair.

5.3 EXPERIMENTS

SAAD claims that the combination of supervoxels as volumes of interest and specialized one-class per-supervoxel classifiers is a better option for outlier detection. This assumption leads to two key questions for abnormal asymmetry detection:

- (i) What is the impact of supervoxel segmentation in SAAD for the quality of the abnormal asymmetry detection?
- (ii) Why use a specialized one-class classifier for each supervoxel instead of a global classifier?

To evaluate SAAD and answer these questions, we performed a detailed evaluation of different scenarios for supervoxel segmentation and classification for the automatic detection of abnormal asymmetries in 3D MR-T1 brain images. We chose a public and challenging dataset with 3D brain images of stroke patients with a variety of strokes in terms of shape, size, location, and texture. Each image has a segmentation mask with its stroke. Although different automatic methods for lesion detection/segmentation exist (Section 5.1), most of them typically capture only specific lesions. They also require multiple imaging modalities or a considered number of segmented images for training (especially for MR-T1 brain images), which is particularly rare for most applications on medical image analysis. Moreover, details, available tools, and trained models for most of these methods are absent, which makes their comparison difficult. In contrast, our method aims to detect any kind of (asymmetric) lesion. We then considered as baselines

for SAAD the state-of-the-art *unsupervised* method (like ours) for the considered stroke image dataset.

This section describes the setup of the entire experiments. We detail the MR-T1 image datasets (Section 5.3.1), followed by the evaluation protocol (Section 5.3.2), including the compared methods and quantitative metrics.

5.3.1 Datasets

To answer our key questions, we need datasets with volumetric MR-T1 brain images (i) from *healthy* subjects (for training), and (ii) with hemispheric *asymmetric* lesions of different sizes (especially small ones) and their segmentation masks. For this, we first considered the CamCan dataset [220], which has 653 3D MR-T1 images of 3T from healthy men and women between 18 and 88 years. For each 3D MR-T1 image, CamCan also has a corresponding 3D MR-T2 image, which we do not use in the experiments. To our knowledge, CamCan is the largest public dataset with 3D images of healthy subjects acquired from different scanners. We visually inspected all MR-T1 images and removed images with bad acquisition or artifacts, yielding 524 images.

Public datasets with different brain lesions exist. However, some only provide a subset of 2D slices for each image or interpolate slices to build a volume (e.g., BraTS [73]); others provide 3D images with only very symmetric lesions (e.g., MSSEG [221]). Given these limitations, we settled on using the Anatomical Tracings of Lesions After Stroke (ATLAS) public dataset release 1.2 [34] in our experiments.

ATLAS is a rather challenging dataset with a large variety of manually annotated lesions and images acquired from different scanners. It contains lesions ranging from very small to large ones, located in several parts of the brain (see Fig. 1.3 for examples). All images have a mask with the primary stroke lesion. Some images also have additional masks with other stroke lesions. Current state-of-the-art *segmentation* results for ATLAS are inaccurate yet [59, 214]. We are not affected by this problem since we aim to *detect*, and not *segment*, the lesions.

Since SAAD is designed to detect abnormal *hemispheric* asymmetries and the considered training images have a 3T field strength, we selected all 3T images from ATLAS, which contain only lesions in the hemispheres (total of 229 images). All images were registered into the coordinate space of ICBM 2009c Nonlinear Symmetric template [108] and preprocessed as outlined in Section 5.2.1 (see an example in Fig. 5.2).

5.3.2 Evaluation Protocol

Baselines: In the absence of details, available tools, and trained models for automatic anomaly detection, we compared SAAD against the convolutional-autoencoder-based approach (CAE) from Chen *et al.* [59] (Section 5.1.5), which is,

as far as we know, the current state-of-the-art *unsupervised* method for the ATLAS dataset. We refer to [Section 4.1](#) for details about autoencoders.

We considered the 2D axial slices of all preprocessed training images to train CAE, which has the following architecture: three 2D convolutional layers with 16, 8, and 8 filters of patch size 3×3 , respectively, followed by *ReLU* activation and 2D max-pooling in the *encoder*, and the corresponding operations in the *decoder*. The nadam gradient optimizer minimized the mean squared error between reconstructed and expected 2D axial slices during training. The method detects anomalies by thresholding the resulting residual image of between the input image vs its reconstruction to obtain a binary segmentation, similarly to Baur *et al.* [53] and Chen *et al.* [59]. We followed Baur *et al.* [53] and selected three thresholds as the 85th, 90th, and 95th percentile from the histogram of reconstruction errors on the considered training set, resulting in the brightness of 143, 194, and 282, respectively. For simplicity, we call these three versions of the method as CAE-85, CAE-90, and CAE-95, respectively, based on the chosen percentiles.

Metrics: Although SAAD detects abnormal asymmetries regardless of the type of anomalies, we can compute quantitative scores only over those lesions that are labeled in ATLAS, which are a subset of what SAAD can detect. For these lesions, we first computed the *detection rate* based on at least 15% overlap between lesions labeled in ATLAS with detected volumes of interest (VOIs) with abnormal asymmetries ([Tables 9–11](#), row 1), as detected by SAAD (supervoxels) and CAE (segmented regions). We then computed the *true positive rate* (recall) that measures the percentage of lesion voxels correctly classified as abnormal ([Tables 9–11](#), row 2). Although our focus is on *detecting* abnormal asymmetries, we also measured the *Dice score* ([Appendix C.2](#)) between lesions and the detected VOIs to check SAAD’s potential as a *segmentation method* ([Tables 9–11](#), row 3). However, observe that truly abnormal asymmetries detected by our method that are not annotated as lesions in the ground-truth masks will be incorrectly considered as false-positive and, thus, underestimating the Dice score. We could then consider only supervoxels overlapped with the annotated lesions to compute Dice scores, but this would be unfair to the considered baselines.

Highly accurate detection methods are useful only if their false positive count is quite low. Otherwise, one needs to manually inspect the many positives to validate them, which is very costly. To gauge this, we provided false-positive (FP) scores in terms of both voxels and supervoxels concerning the ground-truth stroke lesions of ATLAS. Hence, some regions with true abnormal asymmetries but with no labeled masks in ATLAS are considered FP (e.g., see the deformed ventricles of [Fig. 5.7a](#)). This is the best we can do in the absence of labeled masks for all kinds of abnormalities in this dataset.

To evaluate the detection quality, we proposed a set of fine-to-coarse metrics, as follows. At the finest level, we first computed the *mean number of FP voxels*, i.e.,

incorrectly classified as abnormal (Tables 9–11, row 4). We normalized this count with respect to all classified voxels (Tables 9–11, row 5), *i.e.*, the total number of voxels inside the right hemisphere for SAAD, and both hemispheres for CAE.

At the next level, we estimated FP supervoxels as those whose voxels overlap less than 15% with ground-truth lesion voxels. We computed the mean number of FP supervoxels and their proportions to the total number of supervoxels (Tables 9–11, rows 6 and 7). The first metric gives us an estimation of the visual-inspection user effort. The second metric checks how imprecise is the detection regarding the total number of regions that the user has to visually analyze.

When visually analyzing FP supervoxels, it is harder to check many disconnected supervoxels spread across the brain (Fig. 5.7b) than a few connected ones (Fig. 5.7a). Hence, at the coarsest level, we gauge visual analysis user-effort by evaluating the two metrics outlined above on the level of connected FP supervoxel components (Tables 9–11, rows 8 and 9).

Finally, we also computed the *mean processing times* of each method (Tables 9–11, row 10) for preprocessed images, thus excluding the mean time of the preprocessing step (Fig. 5.1, Step 4), which is 90 seconds on average. All experiments were executed on an Intel i7 3.60GHz PC with 64GB RAM and an NVIDIA Titan XP 12GB GPU.

Figure 5.7: Examples of false-positive supervoxels for two different brain slices.

5.4 RESULTS

We next discuss our results from the perspective of our key questions, stated in Section 5.3.

5.4.1 *Impact of Supervoxel Segmentation Quality on Abnormal Asymmetry Detection*

To check the effectiveness of SAAD for abnormal brain asymmetry detection and if the supervoxel segmentation influences this task, we used two variants of the SymmISF method to extract pairs of symmetric supervoxels (one for each

hemisphere), as follows.

Uniform SymmISF: This method chooses N initial seeds *uniformly distributed* in a hemisphere, with N defined by the user (Section 5.2.3). While simple, this strategy does not guarantee to place *at least* one seed within each asymmetric anomaly, especially when N is small. In turn, this leads to undersegmentation — the missed lesions will be assimilated to background. Conversely, when N is too large, this easily leads to oversegmentation of larger lesions into many supervoxels, which have too little individual information to capture asymmetries.

Asymmetry-guided SymmISF: To better fit supervoxels with asymmetric anomalies of various morphologies, this strategy first seeds the highest-asymmetry-value brain regions (where anomalies are more likely to occur) and then seeds the remaining, more symmetric, areas with a fixed number of extra seeds (Section 5.2.3). Since asymmetries vary for each image, the final number of supervoxels is dynamically obtained. For the experiments, we fixed a number of 100 extra seeds uniformly distributed on low asymmetric regions of the images.

We quantitatively compared the above two seeding strategies by using 5-fold cross-validation on ATLAS, considering one subset for *validation* (46 images) and the remaining four subsets for *testing* (183 images) in each fold. For this initial experiment, we first used the following parameters for SAAD, empirically obtained by observing a few training images: $\alpha = 0.08$, $\beta = 3.0$, $\gamma = 2.0$, asymmetry histograms of 128 bins. We considered the one-class Support Vector Machine classifier [201] with the best parameters found for hippocampal asymmetry detection, as stated in Section 4.3.3: kernel = *linear*, and $nu = 0.1$. For Uniform SymmISF, we considered five different numbers of seeds N : {100, 250, 400, 550, 700}.

Table 9 shows the mean results of SAAD with Uniform SymmISF for the primary stroke lesions of ATLAS by considering all five folds, as well as selected visual results. As expected, Uniform SymmISF presents poor detection results for low N values since anomalies are covered by large supervoxels that mix lesion and background voxels (see images in Table 9). As N increases, the chance of placing at least one seed inside each lesion is higher, even for smaller lesions, which leads to better results: We see how the detection rates and mean recall monotonically increase with N in Table 9. Likewise, the number of FP voxels also increases with N , which explains the similar Dice scores from all methods. Yet, there is no guarantee that increasing N yields increasingly-better fitting supervoxels to lesions. This is visible in the results for image 2 (insets) in Table 9, where we see that a small lesion was missed for $N = 700$ but found for $N = 550$. Moreover, the number of FP voxels and supervoxels also increase as N increases — compare rows 4, 6, and 8 of Table 9. This results in considerably high FP rates for large N values. Hence, visual inspection becomes difficult even when the detection rate is high (compare

Table 9: Experimental results for SAAD with Uniform SymmISF for different numbers of seeds. **Top part:** higher values mean better accuracies. **Bottom part:** lower values mean better accuracies. Each result contains a box (inset) surrounding the lesion whose border color indicates if the lesion was detected (green) or missed (red). The abbreviation k denotes thousands.

		Uniform SymmISF				
		N=100	N=250	N=400	N=550	N=700
1	Detection rate	0.389 ± 0.058	0.62 ± 0.052	0.738 ± 0.041	0.808 ± 0.033	0.86 ± 0.028
2	True positive rate (mean recall)	0.159 ± 0.024	0.293 ± 0.025	0.364 ± 0.022	0.409 ± 0.017	0.447 ± 0.018
3	Dice	0.12 ± 0.022	0.13 ± 0.02	0.127 ± 0.02	0.126 ± 0.019	0.123 ± 0.018
4	# FP voxels	$14k \pm 1.36k$	$26k \pm 1.69k$	$32k \pm 1.71k$	$36k \pm 1.75k$	$40k \pm 1.77k$
5	FP voxel rate	0.018 ± 0.002	0.033 ± 0.002	0.041 ± 0.002	0.045 ± 0.002	0.05 ± 0.002
6	# FP supervoxels	8.83 ± 0.27	30.54 ± 0.84	49.97 ± 1.20	67.12 ± 1.86	90.91 ± 2.19
7	FP supervoxel rate	0.103 ± 0.003	0.122 ± 0.003	0.128 ± 0.003	0.128 ± 0.004	0.13 ± 0.003
8	# FP connected supervoxels	6.39 ± 0.17	18.40 ± 0.37	26.11 ± 0.47	32.67 ± 0.68	41.03 ± 0.76
9	FP connected supervoxel rate	0.077 ± 0.002	0.078 ± 0.002	0.071 ± 0.001	0.067 ± 0.001	0.063 ± 0.001
10	Mean processing time (in secs)	39.35 ± 0.98	59.08 ± 0.90	83.68 ± 1.00	111.77 ± 1.49	162.24 ± 12.59
image 1	Ground-truth segmentation
	supervoxels
image 2	Ground-truth segmentation
	supervoxels
image 3	Ground-truth segmentation
	supervoxels

row 1 and rows 4–9 for $N = \{550, 700\}$). Also, the more supervoxels we extract, the longer is the processing time, as shown in Table 9, row 10.

We next compare the Uniform and Asymmetry-guided SymmISF versions for SAAD against three versions of CAE, as presented in Table 10. For simplicity, we considered only the version of Uniform-SymmISF with the best detection accuracy. Also, note that only SAAD reports false-positive supervoxel-based metrics due to its VOI representation.

CAE-85 and CAE-90 present considerably higher detection scores, 0.995 and 0.943, respectively, than the two versions of SAAD: Uniform SymmISF (0.86) and Asymmetry-guided SymmISF (0.851). However, these impressive results are misleading as CAE reports drastically (about 20x) more false-positive voxels than SAAD — compare rows 4 and 5 in Table 10. For instance, although CAE-85 almost detects all lesions, it misclassifies 40% of the hemispheres as abnormal, which is far from being reasonable and hinders the visual analysis (we expect just a small portion of the brain, e.g., 1%). These high FP rates explain the poor Dice scores for CAE in Table 10, which in turn are compatible with the ones reported in [59]. Additionally, CAE is speedy (running time about 2s per image) and yields very noisy detected regions, especially in regions with transitions between white and gray matter (e.g., the cortex), that hinder the subsequent visual inspection (see the results for the considered images in Table 10). Even though the FP voxels decrease as higher thresholds are considered, the detection score can be hugely impacted; for example, the threshold at the 95th percentile approximately halves both the detection score and FP voxels rates compared with the results for the 90th percentile in Table 10. CAE might present better results by using a considerable large training set and/or some additional post-processing, but this is not considered in [53, 59]. CAE presents better results for other medical imaging modalities, such as CT and T2 [53, 59].

Asymmetry-guided SymmISF has a slightly worse detection rate (0.851) compared to Uniform SymmISF (0.86). It is also able to find small abnormal asymmetries (Table 10, images 1-2). However, it fails to detect very subtle and/or tiny asymmetries (Table 10, image 3). Also, this seeding strategy has lowest number of FP (connected) supervoxels and FP voxel scores. However, the expert still has to unnecessarily analyze about 29 FP connected supervoxels per image, which may take a considerable time. The next section details our strategy to improve SAAD with Asymmetry-guided SymmISF to yield higher detection rates and still attenuate FP scores.

5.4.2 Improving the end-to-end method

SAAD with Asymmetry-guided SymmISF is more suitable for our task since the hemispheric asymmetries of each image guide its supervoxel estimation. Moreover, as Table 10 shows, the asymmetry-guided seeding scales computationally better, being roughly twice as fast as uniform seeding for a comparable quality. Hence, we decided to improve Asymmetry-guided SymmISF by (i) optimizing its parameters by grid-search aiming to increase detection accuracy; and (ii) propos-

Table 10: Quantitative and qualitative comparison between the two versions of SAAD and CAE (with different thresholds) for the ATLAS dataset. **Top part:** higher values mean better accuracies. **Bottom part:** lower values mean better accuracies. Each result contains a box (inset) surrounding the lesion whose border color indicates if the lesion was detected (green) or missed (red). The abbreviation k denotes thousands.

	Uniform SymmISF N=700	Asymmetry-guided SymmISF	CAE-85	CAE-90	CAE-95
1 Detection rate	0.86 ± 0.028	0.851 ± 0.016	0.995 ± 0.002	0.943 ± 0.018	0.55 ± 0.03
2 True positive rate (mean recall)	0.447 ± 0.018	0.436 ± 0.009	0.439 ± 0.01	0.333 ± 0.011	0.199 ± 0.011
3 Dice	0.123 ± 0.018	0.132 ± 0.02	0.018 ± 0.003	0.017 ± 0.003	0.016 ± 0.002
4 # FP voxels	$40k \pm 1.77k$	$28k \pm 0.55k$	$644k \pm 2.55k$	$428k \pm 2.1k$	$206k \pm 3.68k$
5 FP voxel rate	0.05 ± 0.002	0.035 ± 0.001	0.4 ± 0.002	0.267 ± 0.001	0.129 ± 0.002
6 # FP supervoxels	90.91 ± 2.19	58.21 ± 1.83	method does not use supervoxels		
7 FP supervoxel rate	0.13 ± 0.003	0.194 ± 0.004			
8 # FP connected supervoxels	41.03 ± 0.76	29.81 ± 0.88			
9 FP connected supervoxel rate	0.063 ± 0.001	0.111 ± 0.002			
10 Mean processing time (in secs)	162.24 ± 12.59	63.03 ± 6.73	2.13 ± 0.08	2.09 ± 0.08	2.04 ± 0.16
image 1 Ground-truth segmentation result	supervoxels		no supervoxels

ing a false-positive-attenuation (FPA) strategy. We describe these optimizations next.

For parameter optimization, we considered the *validation* set of each fold (Section 5.3) and the following search space: $\alpha \in \{0.04, 0.06, 0.08, 0.1, 0.12\}$, $\beta \in \{1.0, 3.0, 5.0, 7.0, 9.0\}$, and $\gamma \in \{0.5, 1.0, 1.5, 2.0, 2.5\}$. As cost function, we considered the *Intersection over Union* (IoU; Appendix C.2) metric that computes the overlap of a supervoxel with each lesion. Indeed, when this overlap is maximal, each lesion is accurately covered by precisely *one* supervoxel. The best parameters found by the grid search were $\alpha = 0.12$, $\beta = 5$, and $\gamma = 0.5$. Note that we used IoU as cost function, and not the metrics listed in Tables 9–10, since it generically looks at how supervoxels fit lesions, whereas those metrics gauge higher-level, more task-specific, concerns.

We repeated the same experiment by considering the optimized parameters. Table 11 presents the results for this seeding strategy, called next Optimized-SymmISF. We replicated the results of CAE-85 and CAE-90, and Asymmetry-guided SymmISF in Table 11 to make the comparison easier. We see that Optimized-SymmISF has a higher detection rate (0.939) than Asymmetry-guided SymmISF (0.851), being slightly worse than CAE-90 (0.943). Optimized-SymmISF also presents the highest true positive rate (0.4889) among *all* compared methods. However, it still has high FP rates and has a considerable increase for the mean number of FP supervoxels and connected supervoxels (Table 11, rows 6 and 8) than Asymmetry-guided SymmISF.

To attenuate FPs, we first performed an analysis of their characteristics. Fig. 5.8 presents two parallel coordinate plots (PCPs) correlating the following three metrics on each FP supervoxel s : (i) distance d of the centroid of s to the right hemisphere border; (ii) volume of s ; and (iii) the mean asymmetry value inside s . Both PCPs are identical except by their highlighted examples (in red).

Fig. 5.8a highlights FP supervoxels close to the hemisphere’s border ($d < 5$ mm), *i.e.*, in the cortex. These supervoxels are relatively small with high variability of mean asymmetries inside them and usually cover gyri and sulci (see the brain slices beside the PCP), which are naturally asymmetric. Conversely, larger FP supervoxels are farther from the hemisphere border (Fig. 5.8b), although their mean asymmetries have high variability. By visually inspecting them, we can also find true abnormal regions deformed by the stroke lesions in the dataset (see the ventricles in the brain slices beside the second PCP). Hence, it seems reasonable to reduce false-positive supervoxels in the cortex.

To do this, we propose a *false-positive-attenuation* (FPA) strategy that accentuates the normal asymmetry map (Section 5.2.2) by adding the standard-deviation asymmetries from the training set to it. As a result, the asymmetry map of the test image (output of Step 5 in Fig. 5.1) is more attenuated so that only highly asymmetric supervoxels will be detected as abnormal. Next, we remove the initial seeds found by Asymmetry-guided SymmISF, whose distance to the hemisphere border is less or equal to 5 mm, as suggested in Fig. 5.8a. We repeated the same

Table 11: Quantitative and qualitative comparison between SAAD after parameter optimization, and the baselines for the ATLAS dataset. **Top part:** higher values mean better accuracies. **Bottom part:** lower values mean better accuracies. Each result contains a box (inset) surrounding the lesion whose border color indicates if the lesion was detected (green) or missed (red). The abbreviation k denotes thousands.

	CAE-85	CAE-90	Asymmetry-guided SymmISF	Optimized- SymmISF	Optimized- SymmISF with FPA
1 Detection rate	0.995 $\pm$ 0.002	0.943 $\pm$ 0.018	0.851 $\pm$ 0.016	0.939 $\pm$ 0.008	0.862 $\pm$ 0.013
2 True positive rate (mean recall)	0.439 $\pm$ 0.01	0.333 $\pm$ 0.011	0.436 $\pm$ 0.009	0.489 $\pm$ 0.006	0.451 $\pm$ 0.006
3 Dice	0.018 $\pm$ 0.003	0.017 $\pm$ 0.003	0.132 $\pm$ 0.02	0.123 $\pm$ 0.018	0.19 $\pm$ 0.018
4 # FP voxels	644k $\pm$ 2.55k	428k $\pm$ 2.1k	28k $\pm$ 0.55k	39k $\pm$ 1.87k	11k $\pm$ 1.17k
5 FP voxel rate	0.4 $\pm$ 0.002	0.267 $\pm$ 0.001	0.035 $\pm$ 0.001	0.049 $\pm$ 0.002	0.014 $\pm$ 0.002
6 # FP supervoxels	method does not use supervoxels		58.21 $\pm$ 1.83	98.15 $\pm$ 1.53	21.19 $\pm$ 0.87
7 FP supervoxel rate			0.194 $\pm$ 0.004	0.176 $\pm$ 0.003	0.065 $\pm$ 0.005
8 # FP connected supervoxels			29.81 $\pm$ 0.88	54.59 $\pm$ 0.73	15.98 $\pm$ 0.56
9 FP connected supervoxel rate			0.111 $\pm$ 0.002	0.107 $\pm$ 0.001	0.049 $\pm$ 0.004
10 Mean processing time (in secs)	2.13 $\pm$ 0.08	2.09 $\pm$ 0.08	63.03 $\pm$ 6.73	111.97 $\pm$ 13.65	72.36 $\pm$ 9.19
image 1 Ground-truth segmentation	supervoxels	no supervoxels
	result
	supervoxels	no supervoxels
image 2 Ground-truth segmentation	supervoxels	no supervoxels
	result
	supervoxels	no supervoxels
image 3 Ground-truth segmentation	supervoxels	no supervoxels
	result
	supervoxels	no supervoxels

parameter optimization for SAAD using FPA, finding the optimal values $\alpha = 0.06$,

(a) Highlighted lines for distances within $[0, 5]$ mm. (b) Highlighted lines for distances within $(5, 20]$ mm.

Figure 5.8: Correlation between some characteristics of false-positive supervoxels. Each false-positive supervoxel is a line in each plot that correlates the distance from its centroid to the right hemisphere’s borders, its volume, and the mean asymmetry inside it. Red lines indicate supervoxels with distances (a) in $[0, 5]$ mm and (b) in $(5, 20]$ mm. An example of these corresponding supervoxels is shown beside each plot.

$\beta = 5$, and $\gamma = 0.5$. Then, we repeated the full detection experiment for the optimal parameter method (called Optimized-SymmISF with FPA).

Table 11 (rightmost column) shows the results. Optimized-SymmISF with FPA has slightly better detection rate (0.862) and TP rate (0.451) to Asymmetry-guided SymmISF. Also, it can detect subtle and tiny asymmetric lesions (Table 11, image 3), which indeed are well-defined by its supervoxels. Although its detection rate is lower than Optimized-SymmISF, it attains the *lowest* FP rates from all considered methods (compare rows 4–9 in Table 11). This method yields, on average, only 1.40% of all voxels as FPs, and these cover only 4.9% of all connected supervoxels. Moreover, Optimized-SymmISF with FPA yields about from twice to three times less FP connected supervoxels for visual analysis than the other versions of SAAD,

which decreases the user effort. Hence, Optimized-SymmISF with FPA has the best balance between high detection rates and low FP rates from all studied methods.

Next, we compared Optimized-SymmISF with FPA, our best method so far, with other related methods in the literature. Such methods are usually designed for the segmentation of, e.g., organs or lesions. As we do not have access to implementations of these methods running on the same dataset as ours, except for CAE, we cannot compute all metrics shown in Table 11. The best we can do is to compare our method with these alternatives as a *segmentation* tool, using segmentation scores. However, note, again, that our method is designed primarily for detection, not segmentation.

Optimized-SymmISF with FPA yields the best Dice score (0.19) among all compared methods in the experiments. As outlined in Section 5.3.2, however, this score is underestimated as truly abnormal asymmetries detected by our method, which are not labeled as lesions in the ground-truth masks, are considered false-positive. If we considered only symmetric supervoxels overlapped with the annotated lesions, such a Dice score leverages to 0.411. While still low, this score is not far from state-of-the-art results (Dice score 0.4867) on the ATLAS dataset from a *supervised* method based on U-Net [214]. Interestingly, our method is noticeably superior to CAE, which is an *unsupervised* method (like ours), reporting Dice scores of 0.018 and 0.017 for thresholds at the 85th and 90th percentile, respectively. Our method reports drastically fewer FP voxels than CAE. Also, note that these compared methods yield their above-reported Dice scores by segmenting quite *large* lesions; in contrast, we focus on the more challenging problem of finding many *small* lesions (see, e.g., image 3 in Table 11).

5.4.3 Per-supervoxel vs Global Classifier Design

We now investigate our second key question, *i.e.*, whether a per-supervoxel classifier design is indeed preferable to a global classifier. Suppose (hypothetically) that supervoxel segmentation is *entirely irrelevant* for the final detection accuracy. Thus, the features we use (normalized histogram of absolute asymmetries for each symmetrical supervoxel) should be able to yield robust texture features for detection regardless of supervoxel quality. Hence, *only* a single classifier — not a (specialized) classifier per supervoxel for each test image — trained from texture features of training images should be enough to obtain similar results to those in Tables 9–10.

To test this hypothesis, we first chose a brain image from ATLAS (Fig. 5.9) with a large asymmetric stroke lesion, which is not as challenging to detect as a small one. If our hypothesis were correct, this lesion should be classified *easily* by global classifier. If global classification failed, then the situation would be even worse for smaller, harder to detect, lesions. To investigate this further, we projected the texture feature vectors of all symmetric supervoxels extracted by Optimized-SymmISF with FPA using t-SNE [202] (Fig. 5.9). Here, each point rep-

resents a symmetric supervoxel, colored by its overlapping percentage with the ground-truth lesion. We see that there is no clear separation between the high-overlap supervoxels (warm-colored points) and healthy-tissue supervoxels (cool colors), even though the considered lesion is very well-defined by a single supervoxel (compare the brain slices in Fig. 5.9). Rauber *et al.* [203] showed that the visual separability of classes in a t-SNE projection is highly correlated with the ability of a classifier to separate classes in the original feature space. Hence, since we do not find good visual separation, we conclude that a single classifier only based on texture features is not sufficient to detect lesions, even large ones.

Figure 5.9: t-SNE projection (perplexity of 50) from texture feature vectors (normalized histogram of absolute asymmetries) for the symmetric supervoxel extracted by Optimized-SymmISF with FPA for a given stroke image. The overlapping percentage between the ground-truth lesion segmentation and each supervoxel is color-encoded in the plot.

5.5 EXTENDING SAAD FOR THE NATIVE IMAGE SPACE

This section extends SAAD for brain anomaly detection in the own *native* image space (NIS) of each test image. NIS is commonly used in clinical routine to provide diagnosis, quantification of disease severity, and treatment planning.

The simplest approach consists of performing SAAD in the *standard* image space (SIS), as proposed in Section 5.2, and then mapping the results to NIS by applying the inverse transformation, obtained by the image registration. Even though this strategy presents reasonable results for most regions in the brain (Fig. 5.10), others are “normalized” due to image registration, *i.e.*, they seem to be healthier than they are — compare the deformed left ventricle in NIS in Fig. 5.10a with the one in Fig. 5.10b, after non-rigid registration to a given template. Consequently, this normalization might attenuate abnormal asymmetries that, in turn,

results in misdetection (Figs. 5.10c–d). Therefore, we need a strategy that preserves real brain asymmetries in NIS for the subsequent analysis.

Figure 5.10: Results of the simplest extension of SAAD for the native image space (NIS). (a) Axial slice of a 3D MR-T1 stroke image (gold-standard borders in red) in NIS. (b) Resulting axial slice after non-rigid registration from (a) to a standard image space (SIS) from a given template. (c) Detected abnormal asymmetries for (b) by SAAD. (d) Resulting detected abnormal asymmetries for (a) after mapping (c) to NIS by the inverse transformation. The arrows point to the deformed left ventricle, damaged by the stroke. The severe deformation in NIS is mitigated in SIS, which impairs its detection.

The key challenge is finding corresponding regions between hemispheres, as these differ in shape, size, and positioning in NIS. Thus, we cannot use MSPs to symmetrically separate such structures, mainly when severe morphological deformations exist (Fig. 5.10). The next sections present the proposed approach, called *Native Supervoxel-based Abnormal Asymmetry Detection* (N-SAAD), and report the experimental results.

5.5.1 Description of N-SAAD

We extend SAAD to perform asymmetry detection in the native image space (NIS) by basically changing only how we compute brain asymmetries for the *test image*, its symmetrical supervoxel segmentation, and its feature extraction. The remaining pipeline steps are the same as presented in Section 5.2 and illustrated in Fig. 5.1. We next detail these proposed changes.

5.5.1.1 Asymmetry Computation

Let X be the set of preprocessed and registered 3D training images of healthy subjects, I the 3D test image after preprocessing except registration (Section 5.2.1), and T the chosen template. We obtain the set of asymmetry maps A_X for all X exactly as proposed in Section 5.2.2, *i.e.*, by computing the voxel-wise absolute differences between left and right hemispheres with respect to their MSPs. We also create the *normal-asymmetry map* $\overline{A_X}$ by averaging the absolute difference

values in A_X . We map $\overline{A_X}$ to I — by applying the deformation fields resulting from the registration between T and I — to attenuate the test asymmetries, as described next.

Since hemispheres in the NIS can be very different in shape, size, and positioning — compare both hemispheres in Fig. 5.11a — we cannot solely rely on the MSP to compute asymmetries for I , as performed by SAAD (Section 5.2.2). Thus, we first flip the left segmented hemisphere (Fig. 5.11c) to the right one, by using the MSP, to keep both hemispheres in the same direction. Next, we register the flipped left hemisphere (Fig. 5.11c) to the right one (Fig. 5.11b) and histogram-match them to guarantee spatial correlation between them (Fig. 5.11d). Section 5.5.2 evaluates N-SAAD with two different registration types for this step. Finally, we compute voxel-wise absolute differences between the correlated hemispheres and subtract the mapped $\overline{A_X}$ from them. Resulting positive values form a final asymmetry map A_I for the test image I (Fig. 5.11e).

Figure 5.11: Asymmetry computation of a 3D test image in its own native image space. (a) Axial slice of a 3D test stroke image after preprocessing and segmentation of the right (red borders) and left hemisphere (green borders). (b) Segmented right hemisphere. (c) Segmented left hemisphere flipped to the right hemisphere using the mid-sagittal plane. (d) Resulting left hemisphere after registering (c) to (b), by non-rigid registration, and histogram-matching them. (e) Asymmetry map resulting from the subtraction (only positive values) between (d) and (b) and attenuation with a precomputed mapped normal-asymmetry map.

5.5.1.2 Symmetric Supervoxel Segmentation in NIS

Since SymmISF considers that the MSP separates brain hemispheres equally after registration on a given symmetric template, we proposed a simple change to adapt it to work in NIS. Given a 3D test image (Fig. 5.12a) and its computed asymmetry map (Fig. 5.12b), both in NIS, we first perform SymmISF (Section 5.5.1.1) *only* inside the right hemisphere by using its corresponding segmentation mask. The result is a label map in which each supervoxel is assigned to a distinct label (Fig. 5.12c). Finally, we map these supervoxels by using the inverse transformation from the corresponding hemisphere registration — previously performed during

asymmetry computation (Section 5.5.1.1) — to obtain the symmetrical supervoxels in the left hemisphere (Fig. 5.12d), which yield the *final label map* L (Fig. 5.12e). We could proceed conversely, *i.e.*, apply SymmISF on the left hemisphere, and map the result to the right hemisphere. Finally, note that the “symmetrical” supervoxels of each pair do not have the precise shape and volume — compare the supervoxels in Fig. 5.12c and their corresponding symmetrical ones in Fig. 5.12d — as obtained by SymmISF in SIS (Section 5.2.3).

Figure 5.12: SymmISF in NIS. (a) Axial slice of a 3D test stroke image after preprocessing and segmentation of the right (red borders) and left hemisphere (green borders). (b) Asymmetry map of (a). (c) Resulting supervoxels for the right hemisphere. (d) Corresponding symmetrical supervoxels for the left hemisphere. (e) Final label map with the symmetrical supervoxels: combination of (c) and (d).

5.5.1.3 Feature Extraction and Classification

In Section 5.5.1.1, we computed the set of training asymmetries A_X — *i.e.*, our adopted knowledge about healthy brains — on the standard image space of a given template T , whereas we computed the asymmetries A_I of the 3D test image I in its own native image space (NIS). As such, we cannot directly use the symmetrical supervoxels L , which is segmented in NIS, to extract asymmetry features for A_X . We need to map L on T to guarantee spatial correlation between supervoxels in both coordinate spaces. This is the main difference compared to the feature-extraction step of SAAD (Section 5.2.4), where, in turn, there is a single coordinate space — and, consequently, a single symmetrical supervoxel map — to extract features from the training and testing asymmetry maps.

Initially, we register I on to T and map L by using the resulting deformation fields, resulting in the map L_T . Note that L and L_T have the same symmetrical supervoxels, only mapped in different coordinate spaces. As in Section 5.2.4, we train a one-class classifier (OCC) for each pair of symmetrical supervoxels in L_T , by using as feature vector the *normalized histogram* of the training asymmetries of A_X inside the pair. Next, we extract the same features for the testing asymmetries in A_I inside each pair of symmetrical supervoxel in L . Finally, we use the trained OCCs to classify their corresponding pairs in the test image as normal or abnormal.

One might claim to compute the asymmetries of each training image in its specific NIS, as performed for the test image. However, this would make the feature extraction and classifiers' training considerably slow, since we would have to map L to the coordinate space of each training image (multiple registrations for each new test image). Instead, our strategy simplifies this process when considering the coordinate space of T to compute all training asymmetries, which is drastically faster.

5.5.2 Preliminary Experiments

We conducted preliminary experiments to assess the accuracy of N-SAAD, as follows.

Experimental Protocol

We adopted the same evaluation protocol presented in [Section 5.3](#): 524 MR-T1 control images from the CamCan dataset [220] for training, and 229 MR-T1 images of stroke patients from the ATLAS dataset [34] for testing. Although we do not perform any parameter optimization in these preliminary experiments — as carried out for SAAD in [Section 5.4.2](#) — we considered the same 5-fold cross-validation on ATLAS, defined in [Section 5.4.1](#), to simplify comparisons between SAAD and N-SAAD. We also used the exact evaluation metrics considered in [Section 5.3](#).

We evaluated N-SAAD using both *affine* and *non-rigid* registration for the entire pipeline ([Section 5.5.1](#)) — *i.e.*, registration of the training images and test image to the template, and registration between the test image's hemispheres. We performed all image registrations by Elastix [107]. Finally, we instantiated N-SAAD with the same initial parameters used for SAAD ([Section 5.4.1](#)): $\alpha = 0.08$, $\beta = 3.0$, $\gamma = 2.0$, asymmetry histograms of 128 bins, and one-class support vector machine classifiers [201] with *linear* kernel and $nu = 0.1$.

All experiments were executed on an Intel i7 3.60GHz PC with 64GB RAM and an NVIDIA Titan XP 12GB GPU.

Results

[Table 12](#) compares experimental results for N-SAAD against two versions of SAAD, *Asymmetry-guided SymmISF* and *Optimized SymmISF with FPA*, presented in [Section 5.4](#). Observe that the former is the initial version of SAAD, which uses the same parameters of N-SAAD, whereas the latter uses optimized parameters and a false-positive-attenuation scheme. For the sake of simplicity, we call the former as *initial-SAAD* and the latter as *optimized-SAAD*. Lastly, [Fig. 5.13](#) shows some visual results for N-SAAD.

N-SAAD with *affine* registration reports the best detection scores, detecting 89.6% of the lesions with a considerably higher mean recall (0.579) than the compared methods. This superiority, however, contrasts with its false-positive (FP)

Table 12: Quantitative comparison of N-SAAD, instantiated with *affine* and *non-rigid* registration, and two versions of SAAD. All scores of SAAD were copied from [Table 11](#): *initial-SAAD* refers to Asymmetry-guided SymmISF, whereas *optimized-SAAD* refers to Optimized-SymmISF with FPA. **Top part**: higher values mean better accuracies. **Bottom part**: lower values mean better accuracies. The abbreviation *k* denotes thousands.

	SAAD (initial)	SAAD (optimized)	N-SAAD (affine registration)	N-SAAD (non-rigid registration)
1 Detection rate	0.851 $\pm$ 0.016	0.862 $\pm$ 0.013	0.896 $\pm$ 0.021	0.856 $\pm$ 0.016
2 True positive rate (mean recall)	0.436 $\pm$ 0.009	0.451 $\pm$ 0.006	0.579 $\pm$ 0.019	0.421 $\pm$ 0.018
3 Dice	0.132 $\pm$ 0.02	0.19 $\pm$ 0.018	0.18 $\pm$ 0.021	0.284 $\pm$ 0.018
4 # FP voxels	28k $\pm$ 0.55k	11k $\pm$ 1.17k	35k $\pm$ 2.45k	6427 $\pm$ 639
5 FP voxel rate	0.035 $\pm$ 0.001	0.014 $\pm$ 0.002	0.03 $\pm$ 0.002	0.001 $\pm$ 0.001
6 # FP supervoxels	58.21 $\pm$ 1.83	21.19 $\pm$ 0.87	29.709 $\pm$ 1.271	10.809 $\pm$ 0.78
7 FP supervoxel rate	0.194 $\pm$ 0.004	0.065 $\pm$ 0.005	0.121 $\pm$ 0.005	0.047 $\pm$ 0.003
8 # FP connected supervoxels	29.81 $\pm$ 0.88	15.98 $\pm$ 0.56	21.804 $\pm$ 0.566	9.246 $\pm$ 0.56
9 FP connected supervoxel rate	0.111 $\pm$ 0.002	0.049 $\pm$ 0.004	0.093 $\pm$ 0.002	0.041 $\pm$ 0.002
10 Mean processing time (in secs)	63.03 $\pm$ 6.73	72.36 $\pm$ 9.19	225.5 $\pm$ 10.5	380 $\pm$ 7.24

scores: for example, it presents the highest number of FP voxels (35,000) regarding all baselines, and it has worse FP scores compared to optimized-SAAD and N-SAAD with non-rigid registration. Conversely, when compared with initial-SAAD, the other FP scores are consistently better (compare rows 5–9 for both methods in [Table 12](#)).

In contrast, N-SAAD with *non-rigid* registration presents an equivalent detection rate (0.856) and mean recall (0.421) with both versions of SAAD, while it yields drastically lower FP scores ([Table 12](#), rows 4–9). For instance, it reports roughly half of FP voxels (only 0.1% of hemispheric voxels) compared to optimized-SAAD, which, in turn, has a false-positive-attenuation scheme. These scores explain the higher Dice (0.284) compared to the baselines. Moreover, the expert will analyze only about 9 FP connected supervoxels per image, which leads to the least user effort among all baselines.

Both instances of N-SAAD detect well-defined abnormal asymmetries for visual inspection, which may be related to anomalies ([Fig. 5.13](#), image 1). They can also find small abnormal asymmetries ([Fig. 5.13](#), image 2). However, N-SAAD cannot detect tiny anomalies ([Fig. 5.13](#), image 3) and pairs of similar and symmetric anomalies in the same region in both hemispheres ([Fig. 5.13](#), image 4), due to the lack of asymmetries (see the asymmetries for [Fig. 5.13](#), image 4 pointed by the arrows).

Figure 5.13: Results of N-SAAD on the ATLAS dataset. Each row is an axial slice of a 3D test image. First column: Ground-truth lesion segmentations. Columns 2 and 4: Results of N-SAAD methods. Columns 3 and 5: N-SAAD asymmetry maps with affine and non-rigid registration, respectively. Arrows indicate undetected lesions.

The good results of N-SAAD indicate that the registration between the hemispheres to spatially correlate them is promising. Non-rigid registration yields less accentuated asymmetry maps than affine registration (Fig. 5.13, Columns 3 and 5), especially in commonly asymmetric regions in the brain (e.g., the cortex), since it uses localized deformations to align hemispheres better. Despite yielding fewer FP anomalies, non-rigid registration is slower than affine registration and may also suppress real abnormal asymmetries (e.g., ventricles of Fig. 5.13, image 1). Fig. 5.11 better illustrates this effect: compare the ventricles of (c) and (d).

Additional registrations make N-SAAD noticeably slower than SAAD (from 3.5x to 6x), especially for non-rigid registration — Table 12, row 10 — which may hin-

der, for example, large-scale studies. Thus, some optimization or faster registration methods are desirable. Besides, as optimized-SAAD has improved both detection and false-positive scores from initial-SAAD — by performing a parameter optimization and a false-positive-attenuation scheme — N-SAAD may follow a similar strategy to further improve its results.

5.6 CONCLUSION

In this chapter, we addressed our research questions by presenting a fully unsupervised approach for abnormal asymmetry detection based on supervoxel segmentation and one-class classification. We also introduced the SymmISF method that extracts symmetrical supervoxels in the brain. Our proposed approach, named SAAD, detects abnormal asymmetric lesions of a given target image — previously registered to the standard image space of a template — by classifying pairs of symmetric supervoxels by using a model (one-class classifier) trained for each pair, independently, from normal brain anatomies only.

We used SAAD to detect stroke lesions on 3D MR-T1 brain images from a wide range of different symmetric supervoxels extracted by two different instances of SymmISF. SAAD achieved higher detection scores and considerably lower false-positive rates compared to an autoencoder-based approach (also unsupervised like ours). SAAD can accurately detect from large to small asymmetric anomalies, which indeed are the most challenging ones. Experimental results confirmed that the quality of supervoxel segmentation *truly* impacts anomaly detection, especially for small anomalies. They also showed that a single global classifier only based on texture features is not sufficient to detect even large anomalies, since their textures are similar to some healthy brain tissues. Putting together our experimental insights, we conclude that (1) a good fit of symmetrical supervoxels to lesions and (2) using a per-supervoxel classifier *are* beneficial design decisions for a good detection of abnormal asymmetries.

Finally, we extended SAAD to perform asymmetry detection in the native image space (NIS). Instead of working on a single coordinate space defined by a template, the extended approach, called N-SAAD, estimates asymmetries and symmetric supervoxels for each test image in its own NIS. We compared N-SAAD instantiated with affine and non-rigid registration — image registration is mainly crucial during asymmetry computation — on the stroke images previously considered. No parameter optimization was performed for N-SAAD. Results show that both instances of N-SAAD have similar or better detection scores compared to the best version of SAAD, with drastically lower false-positive scores by using non-rigid registration. These results indicate that the registration between the hemispheres to correlate them spatially is promising. However, N-SAAD is considerably slower than SAAD (from 3.5x to 6x) for performing multiple image registration along its pipeline.

Our both proposed methods, however, have two main shortcomings: (i) for relying on brain asymmetries for anomaly detection, the analysis is limited only to asymmetric anomalies in the hemispheres; and (ii) they cannot detect similar and “symmetric” anomalies located roughly in the same region in both hemispheres (lack of accentuated asymmetries). These limitations motivated us to further explore the proposed supervoxel classification methodology by replacing asymmetries with registration errors. We address precisely this solution in the next chapter.

Possibilities for future work may include to:

1. Improve asymmetry computation;
2. Refine the symmetric supervoxel segmentation to define supervoxels in subtle lesions;
3. Investigate other feature-extraction techniques;
4. Investigate other one-class classifiers;
5. Explore other false-positive-attenuation strategies (especially for N-SAAD);
6. Evaluate N-SAAD after optimizing its parameters;
7. Evaluate both approaches in other medical imaging modalities.

Another worthwhile goal is using additional visual-analytics techniques to find challenging cases where both methods fail to detect complex small-scale lesions, and then support improving all the above possibilities for future work. One may also investigate fast projection to perform our approaches on the 2D feature space, similarly to our solution presented in [Chapter 4](#).

In the previous chapter, we proposed a fully unsupervised supervoxel-based framework, so-called SAAD, designed to automatically detect anomalies associated with abnormal brain asymmetries, both in the *standard* and *native* image spaces. SAAD detects outliers as abnormal asymmetries by using a supervoxel segmentation, combined with a one-class per-supervoxel classifier.

SAAD obtains higher detection scores and considerably lower false-positive rates compared to a state-of-the-art unsupervised method. However, its analysis is limited to asymmetric anomalies in the brain hemispheres (for only relying on brain asymmetries), ignoring lesions in the cerebellum and brainstem. Also, it cannot detect similar and “symmetric” anomalies located roughly in the same region in both hemispheres, because of the lack of accentuated asymmetries.

Although SAAD uses brain asymmetries in two steps – to guide the symmetrical supervoxel segmentation and to extract features for the symmetrical supervoxels – its pipeline is not strongly related to asymmetries. In fact, after establishing a spatial correspondence between training control and test images (e.g., by image registration), the novelty of this pipeline consists of:

- (1) extracting supervoxels as meaningful volumes of interest (VOIs) specialized for each test image; and
- (2) for each VOI, generating a local one-class classifier, trained on control images, to classify such a VOI as healthy or abnormal on the test image – outliers are considered anomalies.

Thus, we may generalize such a pipeline by replacing asymmetry maps with any other distinct *saliency map*, which must indicate what is more important to analyze in the images, according to a specific problem; in our target problem, brain anomalies. This map can then improve (1) and (2), and also extend the detection of (a)symmetric anomalies for the cerebellum and brainstem, thus overcoming SAAD’s limitations. As such, we may refine our two research questions to:

Can we detect brain anomalies by modeling healthy brain patterns based on image similarity?

This chapter is based on the publications:

- (i) S. B. Martins, A. X. Falcão, and A. C. Telea, “BADRESC: Brain anomaly detection based on registration errors and supervoxel classification,” in *International Joint Conference on Biomedical Engineering Systems and Technologies: BIOIMAGING*, pp. 74–81, 2020. **Best student paper awards.**
- (ii) S. B. Martins, A. X. Falcão, and A. C. Telea, “Combining Registration Errors and Supervoxel Classification for Unsupervised Brain Anomaly Detection,” Accepted for publication in *Lecture Notes in Computer Science*. **Selected paper.**

Following the above observations, this chapter introduces TUSCA — *Towards Unsupervised Supervoxel Classification for Anomaly detection* — a general unsupervised anomaly detection framework based on supervoxel classification. As proof of concept, we instantiate TUSCA by considering *registration errors* as its saliency map. We assume that registration errors for *anomalies* are considerably different from the registration errors for *healthy* tissues. We originally named this instantiated method in [51] of BADRESC — *Brain Anomaly Detection based on Registration Errors and Supervoxel Classification*.

We structure this chapter as follows. [Section 6.1](#) introduces the general pipeline of TUSCA and details its instance based on registration errors. [Section 6.2](#) presents experimental results for anomaly detection in the standard image space, while [Section 6.3](#) concludes the chapter.

6.1 DESCRIPTION OF TUSCA

This section describes the TUSCA framework, instantiated to use *registration errors* as its saliency map (so-called BADRESC method), for the detection of anomalies in the right and left hemispheres, cerebellum, and brainstem. Note that, for our target problem, one could use any other saliency map that emphasizes brain anomalies. TUSCA performs anomaly detection in the *standard* image space. [Fig. 6.1](#) presents the pipeline of TUSCA.

Since TUSCA is a generalization of SAAD, this section presents a similar structure, examples, and diagrams as those presented in [Section 5.2](#) that describe SAAD. Consequently, making this section self-contained implies some repetition, but repetition also serves as a review for people reading the thesis linearly. Therefore, we encourage the reader to compare both sections whenever needed.

6.1.1 3D Image Preprocessing

TUSCA performs precisely the same preprocessing operations described in [Section 5.2.1](#), which, in turn, follows the pipeline presented in [Section 2.3](#). Therefore, we refer to [Section 5.2.1](#) for a complete description of such operations and [Fig. 5.2](#) for a visual example of the preprocessing results for a given stroke 3D image.

6.1.2 Saliency Computation

TUSCA generalizes SAAD by mainly replacing brain asymmetries with any other kind of saliency map tailored to brain anomaly detection. In short, a *saliency map* is an image whose voxels indicate what is more important (*i.e.*, *saliencies*) to analyze in a given test image. For our target problem, an *ideal* saliency map should *only* emphasize brain anomalies, —*e.g.*, by representing them with high grey levels — so that simple thresholding on this map would perfectly segment such anomalies. *How these maps are estimated, or which algorithm is used for that, does not matter*

Figure 6.1: The pipeline of TUSCA. Steps 1 to 3 (blue part) are performed offline. Steps 4 to 8 (pink part) are computed for each 3D test image (detection stage). Step 3 is optional and depends on the target problem. Similarly, the template can be used in any step if needed. The presented saliency map (output of Step 5) is a registration-error map for the test image. One may use a different saliency map according to the target problem.

to perform TUSCA. Note that, for example, SAAD uses a kind of saliency map that indicates the asymmetries of brain images.

To make the pipeline of TUSCA presented in Fig. 6.1 more general, we include an optional step to estimate some information about the training saliency maps (Step 3). For example, one may compute the average saliency map for the training control images, which may be used to attenuate false positives, similarly to SAAD. In practice, this strategy serves to improve the quality of saliency maps by removing undesired saliencies.

The estimation of robust asymmetry maps is a challenging task that has been intensively investigated, especially in computer vision problems [222]. This section proposes to compute *registration errors* as the saliency maps for TUSCA. We hypothesize that registration errors for anomalies are considerably different from the registration errors for *healthy* tissues, which favors outlier detection. We next detail such a strategy.

Registration Error Computation

When registering images to a standard template with *only* healthy tissues, we expect that registration errors (REs) — *i.e.*, voxel-wise absolute differences between the registered image and the template — are lower and present a different pattern compared to anomalies (Fig. 6.2e). However, some healthy structures in the cortex,

such as gyri and sulci, present high REs due to their complex shapes and very large variability between subjects — observe the cortex of the template and the registered image in Figs. 6.2a and 6.2d; note its resulting REs in Fig. 6.2e. As such, we need to apply some attenuation process to avoid detecting false positives in this region.

Figure 6.2: Registration error computation. (a) Axial slice of the brain template. (b) Euclidean Distance Transform (EDT) normalized within $[0, 1]$ computed for the brain segmentation mask defined for the template. Brain borders are shown only for illustration purposes. (c) Common registration errors for control images. (d) Axial slice of a test stroke image after preprocessing and registration in (a). The arrow indicates the stroke lesion. (e) Registration errors. (f) Attenuation of (e) for the cortex based on the EDT. (g) Final registration errors for the test image: positive values of the subtraction between (f) and (c).

Let T be the template (Fig. 6.2a) and M_T its predefined brain segmentation mask for the right hemisphere, left hemisphere, cerebellum, and brainstem (background voxels have label 0, and each object has a different label). Let $X = \{X_1, \dots, X_k\}$ be the set of k registered training images (output of Step 1 in Fig. 6.1) and I the test image after preprocessing and registration (output of Step 4 in Fig. 6.1; see also Fig. 6.2d).

Firstly, we compute the Euclidean Distance Transform (EDT) for each object of M_T and normalize the distances within $[0, 1]$ to build the map E (Fig. 6.2b). Next, we obtain the set of registration errors R_X for all X by computing the voxel-wise absolute differences between X and T (Fig. 6.1, Step 2; see also Fig. 6.2e). For each training image $X_i \in X$, we attenuate REs in its cortex such that for each voxel $v \in X_i$,

$$\begin{aligned} f(v) &= 1 - (E(v) - 1)^\lambda \\ A_{X_i}(v) &= R_{X_i}(v) \cdot f(v), \end{aligned} \tag{6.1}$$

where $E(v)$ is the Euclidean distance for the voxel v , $f(v)$ is its attenuation factor within $[0, 1]$, λ is the exponential factor of the function, and A_{X_i} is the map with the attenuated REs for X_i . In this work, we considered $\lambda = 4$. Thus, REs of voxels close to the brain borders are extremely attenuated, whereas those from voxels far from the borders are slightly impacted (Fig. 6.2f). A downside of this approach is that subtle lesions in the cortex tend to be missed due to the lack of REs.

In order to even ignore REs caused by noises or small intensity differences in regions/tissues far from the cortex, we create a *common registration error map* $\overline{A_X}$ by averaging the attenuated REs from A_X (output of Step 3 in Fig. 6.1; see also Fig. 6.2c). Finally, we repeat the same steps to compute the attenuated REs for the test image I and then subtract $\overline{A_X}$ from them. Resulting positive values form a final attenuated registration error map A_I for I (output of Step 5 in Fig. 6.1; see also Fig. 6.2g).

When comparing the final registration error map of Fig. 6.2g with the corresponding asymmetry map for the same test image (Fig. 5.3d), we can observe that both emphasize the anomaly while attenuating saliencies on the cortex. However, the former computes saliencies for each object of interest independently, which makes it possible to detect “symmetric” anomalies.

6.1.3 Supervoxel Segmentation

Inspired by the SymmISF method (Section 5.2.3) used in SAAD for symmetrical supervoxel segmentation, we propose a more general approach that extracts supervoxels in the entire brain guided by generic saliency maps, as shown in Fig. 6.3. Our supervoxel segmentation is also based on the recent Iterative Spanning Forest (ISF) framework, presented in Section 2.6. We next detail the proposed approach and its instance to use registration errors.

Recall a template T , its predefined brain segmentation mask M_T (macro-objects of interest), a preprocessed 3D test image I registered on T , and A_I be the saliency map for I (e.g., the final attenuated registration errors). Equivalently to SymmISF, we find initial seeds by selecting one seed per local maximum in A_I (see the seeds in Fig. 6.3). We compute the local maxima of the foreground of a binarized A_I at $\gamma \times \tau$, where τ is Otsu’s threshold [219]. The higher the factor γ is, the lower is the number of components in the binarized A_I . We extend the seed set with a fixed number of seeds (e.g., 100) by uniform grid sampling the regions with low REs of the binarized image, resulting in the final seed set $\mathcal{S}$.

Next, we could then either perform ISF directly on I or a different input image improved according to specific constraints for the target problem. For instance, when relying on registration errors as saliency maps, we can even force a matching

Figure 6.3: General pipeline for supervoxel segmentation. The considered examples in this figure come from the registration-error-based instance of TUSCA: the saliency map A_I consists of the attenuated registration errors for a given 3D test image I (segmented objects are colored), and the input image I' is a 2-band volume formed by stacking I and its corresponding template. The preparation of the input image is optional so that different instances of TUSCA could follow a different strategy of simply perform ISF directly on I .

between regions in I and T during supervoxel segmentation by stacking them as the input 2-band volume I' of ISF (Fig. 6.3). This strategy, however, is optional and problem-dependent.

Finally, we perform ISF inside each macro-object of interest in M_T , *separately*, from the initial seeds. The results are label maps wherein each supervoxel is assigned to a distinct number/color. We then combine and relabel the resulting supervoxels to build the final supervoxel map L (output of Step 6 in Fig. 6.1).

6.1.4 Feature Extraction and Classification

TUSCA follows SAAD when relying on designing a set of *specialized* one-class per-supervoxel classifiers (OCCs) specific for each test 3D image to detect anomalies. Each supervoxel estimated for a given 3D test image is then used to train an OCC. This strategy implicitly considers the *position* of the supervoxels in the brain during classification. TUSCA is flexible for feature extraction so that one can choose any algorithm to compute the feature vector of each supervoxel. The classifiers are trained from control images *only* and used to identify outlier supervoxels as anomalies (Fig. 6.1, Step 8).

For the considered instance of TUSCA, the feature vector of each supervoxel L_i consists of the normalized histogram of the attenuated registration errors inside

L_i (Fig. 6.1, Step 7). We consider the one-class linear Support Vector Machine [201] for classification. Yet, note that different feature extractors and one-class classifiers could be used instead the here considered.

6.2 EXPERIMENTS AND RESULTS

This section reports the experiments conducted to evaluate our general pipeline for brain anomaly detection. We adopted the same evaluation protocol presented in Section 5.3, which is briefly summarized next.

6.2.1 Evaluation Protocol

Considered Methods: As a proof of concept, we instantiated TUSCA to detect abnormal registration errors as anomalies, henceforth called BADRESC — *Brain Anomaly Detection based on Registration Errors and Supervoxel Classification*. We compared BADRESC against the primary proposed instance of TUSCA: the SAAD method (Section 5.2).

We considered two versions of SAAD, *Asymmetry-guided SymmISF* and *Optimized SymmISF with FPA*, as presented in Section 5.4. The former is the initial version of SAAD that uses parameters empirically obtained, whereas the latter uses optimized parameters and a false-positive-attenuation scheme. For the sake of simplicity, we call the former as *initial-SAAD* and the latter as *optimized-SAAD*.

For these preliminary experiments, BADRESC uses the identical initial parameters previously considered for SAAD-initial in Section 5.4.1: $\alpha = 0.08$, $\beta = 3.0$, $\gamma = 2.0$, histograms of 128 bins, and one-class support vector machine classifiers [201] with *linear* kernel and $nu = 0.1$.

Datasets: To train the considered methods, we considered the same subset of 524 MR-T1 control images from the CamCan dataset [220] used in Section 5.3.1. As far as we know, CamCan is the largest public dataset with 3D images of healthy subjects acquired from different scanners. After visually inspecting the entire dataset (653 images), we have removed some images with artifacts or bad acquisitions, yielding this subset.

For testing, we also chose the Anatomical Tracings of Lesions After Stroke (ATLAS) public dataset release 1.2 [34] in our experiments. Although we do not perform any parameter optimization for BADRESC in these preliminary experiments — as carried out for SAAD in Section 5.4.2 — we considered the same 5-fold cross-validation on ATLAS (total of 229 images), defined in Section 5.4.1, to simplify comparisons between SAAD and BADRESC. Such folds consist of 3T images that only contain lesions in the hemispheres. Additionally, we also evaluated BADRESC for all the 3T images from ATLAS with stroke lesions in the cerebellum and brainstem (total of 41 images). Therefore, our study involved a total of 794 images.

We performed all non-rigid image registrations by Elastix [107]. All images were registered into the coordinate space of ICBM 2009c Nonlinear Symmetric template [108] and preprocessed as outlined in Section 5.2.1 (see an example in Fig. 5.2).

Metrics: We consider the same evaluation metrics defined in Section 5.3.2, which, in short, measures detection and false-positive (FP) scores. Therefore, we refer to Section 5.3.2 for a full description of all adopted evaluation metrics.

6.2.2 Results and Discussion

Table 13 summarizes all quantitative results for the considered instances of TUSCA, whereas Figs. 6.4 and 6.5 present some visual results. Both versions of SAAD present better detection rates and mean recall for hemispheric lesions compared to BADRES (Table 13, rows 1 and 2), although the difference between the scores is not high — e.g., SAAD-initial has detection score of 0.851 while BADRES has 0.824. BADRES reports a better Dice score (0.169) than SAAD-initial (0.132), being slightly worse than SAAD-optimized (0.19). As outlined in Section 5.3.2, however, this score is underestimated since real unlabeled anomalies detected by the methods are considered false-positive. Regardless, there is a large room for improvements to make these methods accurate for segmentation.

Both BADRES and SAAD can accurately detect small asymmetric lesions in the hemispheres (Fig. 6.4, Image 1). Their considered saliency maps — asymmetries for SAAD and registration errors for BADRES — can successfully emphasize this kind of lesion (see the saliency maps for Image 1 in Fig. 6.4). SAAD cannot detect lesions with low asymmetries, while BADRES does not have this limitation (compare the results and saliency maps for Image 2 in Fig. 6.4). However, both methods are ineffective in detecting tiny anomalies (Fig. 6.4, Image 3) since their saliency maps are not able to highlight such anomalies. Poor saliencies result in the undersegmentation of supervoxels covering anomalies, which compromises their detection.

BADRES is a bit faster and reports fewer false-positive (FP) voxels than SAAD (Table 13, rows 4, 5, and 10), with a considerable difference to SAAD-initial: an average of 8,980 FP voxels against $\approx 28,000$. Concerning FP supervoxel scores, BADRES is consistently better than SAAD-initial (scores roughly twice higher) — compare rows 6–9 in Table 13. For instance, SAAD-initial incorrectly classifies 58.21 supervoxels on average — which consists of 19.4% of all analyzed supervoxels and 3.5% of the analyzed voxels in the hemisphere, respectively. BADRES, in turn, reports an average of 21.43 FP supervoxels, which corresponds to 10.5% of analyzed supervoxels and only 1% of voxels in the entire brain. When compared with SAAD-optimized, BADRES yields to a similar number of FP (connected) supervoxels (Table 13, rows 6 and 8), but it proportionally detects more FP supervoxels to the total number of analyzed supervoxels (Table 13, rows 7 and 9).

Table 13: Quantitative results for images from the ATLAS dataset with stroke lesions in the hemispheres, cerebellum, and brainstem. All scores of SAAD were copied from Table 11: *initial-SAAD* refers to Asymmetry-guided SymmISF, whereas *optimized-SAAD* refers to Optimized-SymmISF with FPA. Since we performed BADRESF on all images with lesions in the cerebellum and brainstem, its detection rate is the exact rate (not a mean). **Top part:** higher values mean better accuracies. **Bottom part:** lower values mean better accuracies. The abbreviation *k* denotes thousands.

	Hemispheres			Cerebellum and Brainstem
	SAAD (initial)	SAAD (optimized)	BADRESF	BADRESF
1 Detection rate	0.851 $\pm$ 0.016	0.862 $\pm$ 0.013	0.824 $\pm$ 0.017	0.683
2 True positive rate (mean recall)	0.436 $\pm$ 0.009	0.451 $\pm$ 0.006	0.4 $\pm$ 0.012	0.26 $\pm$ 0.26
3 Dice	0.132 $\pm$ 0.02	0.19 $\pm$ 0.018	0.169 $\pm$ 0.014	0.11 $\pm$ 0.151
4 # FP voxels	28k $\pm$ 0.55k	11k $\pm$ 1.17k	8.98k $\pm$ 748	8.78k $\pm$ 7.94k
5 FP voxel rate	0.035 $\pm$ 0.001	0.014 $\pm$ 0.002	0.01 $\pm$ 0.001	0.01 $\pm$ 0.01
6 # FP supervoxels	58.21 $\pm$ 1.83	21.19 $\pm$ 0.87	21.43 $\pm$ 1.114	25.63 $\pm$ 15.64
7 FP supervoxel rate	0.194 $\pm$ 0.004	0.065 $\pm$ 0.005	0.105 $\pm$ 0.005	0.0971 $\pm$ 0.049
8 # FP connected supervoxels	29.81 $\pm$ 0.88	15.98 $\pm$ 0.56	16.58 $\pm$ 0.654	18.414 $\pm$ 10.2
9 FP connected supervoxel rate	0.111 $\pm$ 0.002	0.049 $\pm$ 0.004	0.083 $\pm$ 0.003	0.079 $\pm$ 0.034
10 Mean processing time (in secs)	63.03 $\pm$ 6.73	72.36 $\pm$ 9.19	54.17 $\pm$ 1.3	52.62 $\pm$ 2.5

BADRESF is less accurate when detecting lesions in the cerebellum and brainstem (detection rate of 0.6829). Some of these lesions are indeed challenging, especially in the cerebellum, whose appearances are similar to their surrounding tissues (Fig. 6.5, Image 6). Although its FP scores are similar to those of hemispheric lesions — compare rows 4-9 for BADRESF in Table 6.5 — the considered registration-error attenuation (Eq. 6.1) seems to be very strong for the cerebellum and brainstem, which impairs the representation of the lesions.

The two evaluated instances of TUSCA show its flexibility and potential for unsupervised brain anomaly detection. Although the preliminary experimental results indicate that SAAD-optimize is the best option for the detection of asymmetric anomalies in the hemispheres, its parameters were previously optimized for such a problem. In contrast, BADRESF used parameters empirically obtained, which, even so, resulted in competitive results to SAAD-optimized with the advantage of (i) performing in the cerebellum and brainstem, and (ii) detecting symmetric lesions.

Figure 6.4: Comparative results between SAAD and BADRES. All images have a stroke lesion in the hemisphere. We present an axial slice of the result and corresponding saliency map for both methods. Each image contains an inset surrounding the lesion whose border color indicates if the lesion was detected (green) or missed (red).

Figure 6.5: Results of BADRESC for some images with stroke lesions in the cerebellum or brainstem. Each image contains an inset surrounding the lesion whose border color indicates if the lesion was detected (green) or missed (red).

6.3 CONCLUSION

This chapter presented TUSCA, a general fully unsupervised framework for anomaly detection. This framework generalizes the previous SAAD method (Section 5.2) to use any other distinct *saliency maps* that emphasizes brain anomalies, instead of only *brain asymmetries*. Consequently, we can extend the anomaly detection to the cerebellum and brainstem, as well as detecting lesions with low asymmetries, in contrast to SAAD. This generalization refines our research questions, which are focused on abnormal brain asymmetry detection.

As proof of concept, we instantiated TUSCA to use *registration errors* (REs) as saliency maps, so that abnormal REs are considered as anomalies. We compared this instance, named BADRESC, with two different versions of SAAD on 3T MR-T1 images of stroke patients. BADRESC reports a bit lower detection scores than SAAD for hemispheric lesions, but it attains similar false-positive scores to the most accurate version of SAAD, being superior to the other version. BADRESC also detects lesions in the cerebellum and brainstem with promising results. Consequently, we have shown it is possible to model healthy brain patterns, based on image similarity, for the detection of brain anomalies, thus

addressing the question introduced in this chapter.

Future work: The two evaluated instances of TUSCA show its flexibility and potential for unsupervised anomaly detection. As such, future work may address several extensions to this work. Initially, one could investigate other different maps for our target problem, since the quality of the saliency maps directly impacts the detection accuracy. Likewise, one could focus on designing saliency maps for a given specific problem — *e.g.*, a map that only highlights tumors — so that detection and segmentation scores will be maximized for such a problem.

Even though we performed TUSCA on brain images, its applicability is not limited to brains. Indeed, its only major requirement consists of *having all images aligned on the same coordinate space*. A robust image registration guarantees such spatial correspondences among the images. A few other adjustments, such as different preprocessing operators, can also be necessary according to the application. Therefore, future work include extending and evaluating TUSCA to detect anomalies in other organs (*e.g.*, lungs).

Some possibilities for future work outlined in [Section 5.6](#) can also be investigated for TUSCA. In particular, one could explore other feature-extraction techniques, since we only considered a simple normalized histogram for that.

Concerning BADRESC, one could first optimize its parameters and evaluate the resulting quantitative and qualitative impact. Since brain hemispheres, cerebellum, and brainstem have different constraints, the use of a different registration-error attenuation for each of these objects should be further investigated.

BADRESC depends on the quality of image registration in order to generate high-quality saliency maps. In this work, we register the entire brain of images (after preprocessing and skull-stripping) to the template. One could then investigate if the registration of each object, independently, provides more accurate results, *i.e.*, low registration errors, which yields better detection results.

CONCLUSION

We conclude here our work by revisiting our two research questions stated in [Section 1.4](#):

RQ1: *Can we model normal brain asymmetries?*

RQ2: *Can we use the normal brain asymmetry model to detect brain anomalies?*

To answer these questions, this thesis successfully explored different fully unsupervised approaches, as based on the general pipeline presented by [Fig. 1.6](#), to detect abnormal brain asymmetries associated to anomalies in 3D MR-T1 images. We successfully presented solutions that showed how unsupervised machine learning can leverage brain anomaly detection. Our key contributions include:

- (1) A novel automatic brain image segmentation method used in different steps of the proposed approaches (*e.g.*, image preprocessing);
- (2) An autoencoder-based method to model normal brain asymmetries of a given brain structure and detect outliers as anomalies;
- (3) An unsupervised supervoxel-based framework to detect *abnormal brain asymmetries* as anomalies.

Although this thesis focused on abnormal brain asymmetry detection, as the last contribution, we still extended the solution (3) – which is only designed for asymmetric lesions in the hemispheres – for the detection of (a)symmetric brain anomalies in the entire brain (hemispheres, cerebellum, and brainstem) regardless their asymmetries.

We next summarize such contributions towards our research questions, their results, limitations, as well as we present plenty of opportunities for future work.

7.1 BRAIN IMAGE SEGMENTATION

Brain image segmentation is an essential task in many applications, such as the study of brain asymmetries, morphological analysis of the hemispheres, or for a better understanding of neurological diseases. It consists of the precise segmentation of the right and left hemispheres, cerebellum, and brainstem. Different prior steps of the proposed anomaly detection approaches rely on such segmentation so that a robust segmentation method is required.

[Chapter 3](#) introduced a fast and effective solution, named *AdaPro*, for the automated segmentation of brain structures in control and anomalous MR 3D images.

The novelty of AdaPro is it incorporates a texture classifier based on convex optimization that *dynamically* indicates the regions of the target 3D image where the probabilistic atlases (shape constraints) — built from healthy structures/organs — should be further adapted, due to the presence of anomalies. AdaPro then delineates the objects of interest by a new algorithm based on combinatorial optimization and diffusion filtering.

Experimental results on 3D MR-T1 images of 2T and 3T from epilepsy patients before and after temporal lobe resections showed the superiority of AdaPro than three other atlas-based methods, in terms of accuracy and efficiency. Moreover, AdaPro can segment new images in the native and reference image spaces.

Future work: One opportunity for future work is employing AdaPro to other organs and imaging modalities. More robust tissue classification for other different anomalies could be also evaluated. To further increase the efficiency of AdaPro, one may investigate faster (and yet accurate) image registration methods.

7.2 ABNORMAL ASYMMETRY DETECTION BY AUTOENCODERS AND ONE-CLASS CLASSIFICATION

While the human brain presents natural structural asymmetries between both hemispheres, some neurological diseases, such as epilepsy, are associated with *abnormal asymmetries*. The simplest strategy to detect such anomalies consists of a visual slice-by-slice inspection in a 3D brain image by one or multiple specialists. However, this manual analysis is very time-consuming, error-prone, and even impracticable when a large amount of data needs to be processed.

[Chapter 4](#) presented our first unsupervised solution regarding the research questions. The proposed automatic framework exploits convolutional autoencoders (CAEs) and a one-class classifier to model normal asymmetries from healthy subjects, thus addressing **RQ1**. We then addressed **RQ2** by using such a model to detect outliers as abnormal asymmetries. We use the intermediate layers from CAEs (latent features) to represent asymmetries (feature vector). As proof of concept, we instantiated the framework to analyze hippocampal asymmetries from 3D patches around the hippocampi in both hemispheres. We also proposed an automatic method to localize these 3D patches, as well as a novel one-class classifier based on optimum-path forests.

We evaluated the framework using MR-T1 images from healthy subjects and epilepsy patients with unilateral hippocampal atrophy. We considered two feature spaces to train the classifiers: the original one and a 2D space created by non-linear projection. The latter facilitates the understanding of the data distribution, sample inspection, and annotation of the detected anomaly type. Results reported high detection scores, especially considering some difficult cases that only a trained expert can visually identify. Regarding our research questions, we therefore showed our

first well-succeeded unsupervised solution to model normal brain asymmetries to detect brain anomalies.

When extending the framework for the detection of brain asymmetries along the entire hemispheres, however, the reported results were unsatisfactory. Experiments showed poor results with several false-positive asymmetries detected for all considered images. Also, such an extension did not detect true abnormal asymmetries caused by lesions. All these limitations are directly related to the poor estimation of 3D patches (size and location).

Putting together our experimental insights, the proposed autoencoder-based framework is a robust solution for the asymmetry analysis of a specific structure of interest, since one may (automatically or manually) design high-quality 3D patches for it. Multiple analysis of general structures, with distinct shapes and locations in the brain, should rely on one more precise volumes of interest, such as supervoxels, which we next explored in this thesis.

Future work: The simplest possibility for future work involves performing the proposed framework to other brain structures, organs, or different medical imaging modalities than MR-T1. The impressive results in the 2D projection space suggest further investigating the impact of the projection in the design of classifiers. One could also explore fast projection algorithms to make the asymmetry analysis in the two-dimensional space feasible for large-scale studies. Finally, one may develop interactive visual tools to support the detection, inspection, annotation, and identification of brain anomalies based on abnormal asymmetries.

7.3 UNSUPERVISED SUPERVOXEL-BASED ABNORMAL BRAIN ASYMMETRY DETECTION

The limitations of analyzing asymmetries from 3D patches motivated us to explore *supervoxel segmentation* to estimate more precise volumes of interests (VOIs). *Supervoxels* are groups of voxels with similar characteristics resulting from an over-segmentation of a 3D image or region of interest, which preserves intrinsic image information (e.g., the borders of tissues and lesions).

Chapter 5 presented the key contribution of this thesis: a general fully unsupervised framework for abnormal asymmetry detection based on supervoxel segmentation, and specialized one-class per-supervoxel classifiers for outlier detection in the *standard* image space. Such a framework, called SAAD (*Supervoxel-based Abnormal Asymmetry Detection*), uses brain asymmetries to guide the supervoxel segmentation and extract a feature vector for each supervoxel. SAAD relies on a novel method, so-called SymmISF, that extracts pairs of symmetric supervoxels in both hemispheres for the subsequent analysis.

Chapter 5 provided an in-depth evaluation of SAAD to detect stroke lesions on 3D MR-T1 brain images by considering different scenarios for supervoxel segmentation, parameter optimization, and presenting a false-positive-attenuation strat-

egy for SAAD. Experimental results showed that SAAD attained higher detection scores and considerably lower false-positive rates than a state-of-the-art unsupervised method (like ours). They also showed the effectivity of using a set of specialized per-supervoxel classifier instead of a single global one. However, SAAD is limited to detecting asymmetric anomalies in the hemispheres.

Additionally, [Chapter 5](#) also introduced an extension of SAAD for abnormal asymmetry detection in the *native* image space. We evaluated the new approach, named N-SAAD (‘N’ stands for *native*), from a set of parameters empirically obtained. N-SAAD presented competitive detection scores against an optimized version of SAAD, with less false-positive rates.

Future work: All experimental insights in this thesis open several potential directions for future work, as follows. One may further investigate other feature-extraction techniques and one-class classifiers to model better normal asymmetries for (N-)SAAD. Since false positives are problematic in both methods, one should consider investigating different attenuation strategies. Besides, an optimization procedure should be performed for N-SAAD in order to improve its results further. Another worthwhile goal is using additional visual-analytics techniques to find challenging cases where both methods fail to detect complex small-scale lesions, and then support improving all the above possibilities for future work.

Although our unsupervised approach can attain high detection accuracies, *supervised* techniques are usually more accurate for specific tasks. However, these methods require a large number of manually annotated training images, which is absent for most medical image analysis problems. Thus, one interesting direction is to use SAAD to facilitate data annotation for supervised problems, as follows.

One can initially perform SAAD on training images from patients with a given anomaly (e.g., tumor). Next, an expert visually removes the false-positive detected supervoxels and refines the segmentation of the remaining detected supervoxels if needed. Resulting supervoxels form the gold-standard segmentation for the training images, which may then be used to train a supervised model for the target problem. This novel data annotation process rely on considerably fewer user interactions to yield a high-quality labeled dataset. One may still use these interactions (expert’s inputs) to improve further the unsupervised anomaly detection of SAAD via some visual-analytics technique.

7.4 TOWARDS UNSUPERVISED SUPERVOXEL CLASSIFICATION FOR ANOMALY DETECTION

SAAD is only designed to detect asymmetric lesions in the brain hemispheres. Although SAAD uses brain asymmetries to guide the symmetrical supervoxel segmentation and to extract features for the symmetrical supervoxels, its pipeline is not strongly related to asymmetries. Thus, [Chapter 6](#) presents a fully unsupervised

framework that extends SAAD to detect lesions (symmetric or asymmetric) in the hemispheres, cerebellum, and brainstem. This solution relies on any other *saliency map* that emphasizes brain anomalies, instead of only *brain asymmetries*.

A *saliency map* is an image (or a heatmap) that indicates what is more important to analyze in a given target image, according to a specific problem; in the context of this thesis, *brain anomalies*. The proposed framework, called TUSCA (*Towards Un-supervised Supervoxel Classification for Anomaly detection*), uses saliencies to guide the supervoxel segmentation and extract a feature vector for each supervoxel.

Chapter 6 presented and evaluated an instance of TUSCA, called BADRESC, that considers *registration errors* as saliencies. Our premise was that the registration errors for *anomalies* are considerably different from those registration errors for *healthy* tissues. BADRESC reports slightly lower detection scores than SAAD for hemispheric lesions, with similar false-positive scores. Nevertheless, BADRESC can detect lesions in the cerebellum and brainstem with promising results.

Future work: Since the only major requirement from TUSCA consists of *having all images aligned on the same coordinate space*, one opportunity for future work is to extend TUSCA to detect anomalies in other organs (e.g., lungs). Initially, one could simply evaluate BADRESC in the other target organ or designing a specific saliency map that emphasizes its lesions, taking into account specific characteristics of these anomalies.

Like SAAD, one may investigate other feature-extraction techniques and one-class classifiers to model better registration errors or other saliencies. As estimating high-quality supervoxels that accurately cover the lesions is crucial for the accuracy of TUSCA, one may also focus on improving this segmentation, especially for complex small-scale lesions. Regarding BADRESC, an optimization procedure should be performed to improve its results further. Finally, one may use some instance of TUSCA to assist data annotation, as previously suggested for SAAD.

NOTATIONS AND DEFINITIONS

ADAPRO	Adaptive Probabilistic Atlas proposed for anomalous brain image segmentation (Chapter 3).
AE	Abbreviation of Autoencoders (Section 4.1).
ASYMMETRY MAP	A 3D (volumetric) image with brain asymmetries (Chapter 5).
ATLAS	Pair of a source image and its segmentation mask for objects of interest (Chapter 3).
AUTOENCODER	A generative neural network for image reconstruction (Section 4.1).
BRAIN ASYMMETRY	Anatomical differences between the two brain hemispheres.
BADRESC	An instance of the TUSCA framework for <u>B</u> rain <u>A</u> nomaly <u>D</u> etection based on <u>R</u> egistration <u>E</u> rrors and <u>S</u> upervoxel <u>C</u> lassification (Chapter 6).
CAE	Convolutional Autoencoder (Section 4.1).
CEREBRUM	The largest and uppermost of the brain containing the cerebral cortex of the two hemispheres as well as several subcortical structures such as the hippocampus (Section 2.1.1).
CONTROL IMAGE	An image with no pathologies.
CSF	Cerebrospinal fluid: it is a clear, colorless body fluid found in the brain and spinal cord (Section 2.1.2).
DETECTION	A visual indication about the location of an object of interest (<i>e.g.</i> , an organ or lesion). It is commonly defined as a simple 2D or 3D bounding box around the object or a rough object segmentation mask of it.
ELASTIX	A popular software for (non)rigid image registration (Section 2.3.4).

GM	Gray Matter tissue present in the brain (Section 2.1.1).
GOLD-STANDARD	Term used to describe segmentation masks resulting from fully manual or semi-automatic image segmentation by one or multiple experts (Section 3.1).
IFT	Image Forest Transform framework commonly used to design image processing operators (Section 2.4).
ISF	An Iterative Spanning Forest framework for super-pixel and supervoxel segmentation (Section 2.6).
ISOTROPIC	An image is isotropic if its pixel/voxel size is the same in every dimension (Section 2.2.1).
MR-T1	T1 Magnetic-resonance image modality (Section 1.1).
MRI	Magnetic Resonance Imaging (Section 1.1).
MSP	Mid-Sagittal Plane — the median vertical longitudinal plane that approximately divides a bilaterally symmetrical brain into right and left hemispheres (Section 2.1.2).
N-SAAD	Version of the SAAD method for analysis on the native coordinate space of a given image (Section 5.5).
NIS	Native Image Space. It corresponds to the coordinate space of a given image (Section 5.5).
OPF	Optimum-Path Forest — A graph-based framework designed for clustering and classification methods (Section 2.5).
PBM	The proposed Patch-Based Model for VOI location (Section 4.2.2).
REFERENCE IMAGE	An image — commonly created by averaging several control images — that is used as standard coordinate space for image visualization and analysis. Also known as template (Section 2.3.4).
ROI	Region of Interest. Commonly a 2D bounding box or a segmentation mask of a given object/region.

SAAD	Supervoxel-based Abnormal Asymmetry Detection method (Section 5.2).
SALIENCY MAP	An image (or heatmap) that indicate what is more important to analyze in a given target image (Chapter 6) — in the context of this thesis, brain anomalies.
SEGMENTATION	Process of partitioning an image into multiple objects (segments) by precisely delineating their borders.
SIS	Standard Image Space. It corresponds to the standard coordinate space commonly defined by a template (Section 5.2).
SLICE	A 2D image extracted from a 3D image along one of its three orthogonal planes (Section 2.1.2).
SUPERPIXEL	A group of pixels of a 2D image that shares common characteristics.
SUPERVOXEL	Equivalent to superpixel but grouping voxels in a 3D image instead.
SYMMISF	Symmetric ISF algorithm that extracts pairs of symmetric supervoxels in objects <i>e.g.</i> , brain hemispheres (Section 5.2.3).
TEMPLATE	Equivalent to Reference Image.
TUSCA	Towards Unsupervised Supervoxel Classification for Anomaly detection — a general anomaly detection framework based on supervoxel classification (Chapter 6).
VOI	Volume of Interest. Equivalent to ROI but for volumes in 3D images.
WM	White Matter tissue present in the brain (Section 2.1.1).

DATA

A key challenge in medical image analysis is the absence of publicly large annotated datasets for different applications [11, 44, 147]. *Annotations* are problem-dependent and typically consist of labels that categorize entire images (e.g., the image is healthy or abnormal), bounding boxes around objects of interest (e.g., organs and anomalies), or the precise segmentation of these objects. Providing annotations for 3D images is time-consuming since they require manually slice-by-slice inspection or labeling by one or multiple experts. To answer our research questions, we need datasets with isotropic 3D MR-T1 brain images from (i) healthy subjects, and (ii) with anomalies of different sizes (especially small ones) and their *gold-standard* segmentation masks, as detailed next.

B.1 IN-HOUSE DATASETS

This thesis has the collaboration of neurologists from the Neuroimaging Laboratory (LNI) at the University of Campinas (UNICAMP), Brazil, which are particularly interested in (i) investigating abnormal asymmetries in epilepsy patients, and (ii) estimating morphometric measures for the hemispheres and cerebellum of those patients. LNI has provided brain images of healthy subjects and epilepsy patients from the Clinical Hospital at UNICAMP. Table 14 presents such datasets, grouped according to the task they are used in this thesis: Brain segmentation (Chapter 3), and Analysis of Hippocampal Asymmetries (Chapter 4).

B.2 PUBLIC DATASETS

We next present the public datasets used to train and evaluate our supervoxel-based approaches for anomaly detection, as detailed in Chapters 5 and 6.

CAMCAN

CamCan dataset [220] has 653 3D MR-T1 images of 3T from healthy subjects between 18 and 88 years. For each 3D MR-T1 image, it also has a corresponding 3D MR-T2 image, which we do not use in this thesis. To our knowledge, CamCan is the largest public dataset with 3D control images acquired from different scanners.

ATLAS

Anatomical Tracings of Lesions After Stroke (ATLAS) public dataset release 1.2 [34] is a rather challenging dataset with a large variety of manually annotated lesions and images acquired from different scanners. It contains lesions ranging

Table 14: In-house datasets used in this thesis for Brain Segmentation (top) and Analysis of Hippocampal Asymmetries (bottom). A neurologist from LNI has carefully delineated the cerebellum, right and left hemispheres of all brain-segmentation datasets to build the gold-standard segmentation. There is no intersection among the datasets presented in this table.

	Name	Description	No. of Images	Voxel Size (mm ³)	Field Strength
Brain Segmentation	HEALTHY-2T	Control images acquired with a 2T Elscint scanner.	19	0.98 × 0.98 × 0.98	2T
	PRE-2T	<i>Pre-operative</i> images from epilepsy patients took with a 2T Elscint scanner. Each patient has one image.	20	0.98 × 0.98 × 0.98	2T
	POST-2T	<i>Post-operative</i> images from the same patients of PRE-2T, after temporal lobe resection. Each patient has one image.	20	0.98 × 0.98 × 0.98	2T
	HEALTHY-3T	Control images acquired with a 3T Siemens scanner.	20	1 × 1 × 1	3T
	PRE-3T	<i>Pre-operative</i> images from epilepsy patients took with a 3T Siemens scanner. Each patient has one image.	30	1 × 1 × 1	3T
	POST-3T	<i>Post-operative</i> images from the same patients of PRE-3T, after temporal lobe resection. Each patient has two image.	60	1 × 1 × 1	3T
Analysis of Hippocampal Asymmetries	CONTROLS	<i>Control images</i> (subjects are between 25 and 65 years old) took with a 3T Siemens scanner. A neurologist from LNI has delineated the hippocampi of a subset of 60 images.	575	1 × 1 × 1	3T
	PRE	<i>Pre-operative</i> images from epilepsy patients took with a 3T Siemens scanner. A neurologist from LNI has carefully delineated the hippocampi of all images.	47	1 × 1 × 1	3T
	POST	<i>Post-operative</i> images from epilepsy patients after temporal lobe resection took with a 3T Siemens scanner. A neurologist from LNI has carefully delineated the hippocampi of all images.	88	1 × 1 × 1	3T
	RHA	Images from epilepsy patients with <i>right</i> hippocampal atrophy, acquired with a 3T Siemens scanner.	34	1 × 1 × 1	3T
	LHA	Images from epilepsy patients with <i>left</i> hippocampal atrophy, acquired with a 3T Siemens scanner.	37	1 × 1 × 1	3T

from very small to large ones, located in several parts of the brain. All images have a mask with the primary stroke lesion. Some images also have additional masks with other stroke lesions.

METRICS

Throughout this thesis, we present several methods for different tasks related to image analysis. All of these rely on computing some quantitative metric in some step, *e.g.*, for quantitative evaluation or fine-tuning purposes. For instance, the proposed autoencoder-based approach in [Section 4.2](#) uses the mean squared error as the loss function in the training. Moreover, we rely on different metrics to quantitatively compare our proposed methods against baselines in the experiments. All these metrics are next detailed in this appendix.

Basic notation: Assume that a *3D grayscale image* is defined as $\hat{I} = (D_I, I)$, where $D_I \subset Z^3$ is the *image domain* — *i.e.*, a set of elements in Z^3 denominated *voxels* — and $I(p) \in Z$ is the *gray value* assigned to every voxel $p \in D_I$. Similarly, a *binary mask* $\hat{B} = (D_B, B)$, which describes a given object of interest, considers that background voxels have label 0 and object voxels have label 1.

C.1 IMAGE SIMILARITY MEASURES

For the metrics presented in this section, consider that we want to measure the similarity between two *grayscale images*, $\hat{I} = (D_I, I)$ and $\hat{J} = (D_J, J)$, where $D_I = D_J$.

C.1.1 Mean Square Error (MSE)

Mean Squared Error (MSE) measures the mean voxel-wise intensity difference (error) between two images. The higher is MSE, the more different is such images. MSE can be mathematically defined by

$$MSE(\hat{I}, \hat{J}) = \frac{1}{|D_I|} \sum_{p \in D_I} [I(p) - J(p)]^2, \quad (C.1)$$

where $|D_I|$ is the number of voxels of $\hat{I}$, which is identical for $\hat{J}$.

We use MSE as the loss function in the training of our proposed autoencoder-based solution for abnormal hippocampal asymmetry detection ([Section 4.2](#)). In this case, we use MSE to compute the mean reconstruction errors for the training images.

C.1.2 Normalized Mutual Information (NMI)

Normalized Mutual Information (NMI) scales the *Mutual Information* (MI) score between 0 (no mutual information) and 1 (perfect correlation). The definition of

NMI involves some other concepts as detailed next.

Shannon Entropy (SE): It measures the dispersion of the gray values of an image $\hat{I}$. For instance, a gray image $\hat{I}$ that is entirely filled by a single gray value has $SE(\hat{I}) = 0$. Mathematically, SE is given by

$$SE(\hat{I}) = - \sum_{\forall x \in I} (p_x \cdot \log p_x), \quad (C.2)$$

where I represents the set of gray values of $\hat{I}$, and p_x is the probability associated with the gray value x , i.e., p_x is the frequency that x appears in $\hat{I}$ divided by $|D_I|$ (the total number of voxels of $\hat{I}$). The log function in Eq. C.2 can be a $\log_2$ or a $\ln$, having the entropy units in *bits* (binary units) or *nats* (natural units), respectively.

Note that *spatial information* is not taken into account to compute SE. Consequently, two different images can have the same SE without conveying the same visual information, as illustrated in Fig. C.1.

Figure C.1: Example of two different grayscale images with the same Shannon Entropy.

Joint Intensity Histogram (JIH): It is a two-dimensional graphic where the value assigned to each point (x, y) corresponds to the number of voxels with gray value x in $\hat{I}$, whose corresponding voxels in $\hat{J}$ have gray value y [223].

The frequencies from the joint intensity histogram can be used to calculate the Shannon Entropy of that *joint distribution* of gray values. In this case, the mathematical expression is defined as

$$SE(\hat{I}, \hat{J}) = - \sum_{\forall (x,y) \in I \times J} (p_{x,y} \cdot \log p_{x,y}), \quad (C.3)$$

where $p_{x,y}$ represents the relative frequency of each point from the joint intensity histogram between $\hat{I}$ and $\hat{J}$.

Finally, **Normalized Mutual Information** can be defined by

$$NMI(\hat{I}, \hat{J}) = \frac{SE(\hat{I}) + SE(\hat{J})}{SE(\hat{I}, \hat{J})}. \quad (C.4)$$

We use NMI for image registration (Section 2.3.4), and VOI localization (Section 4.2). Some of the related work also uses NMI in their methods.

C.2 SEGMENTATION METRICS

For the metrics presented in this section, consider that we want to measure the similarity between a binary mask $\hat{M} = (D_M, M)$, resulting from a given segmentation, and its corresponding ground-truth binary mask $\hat{G} = (D_G, G)$, where $D_M = D_G$. For simplicity, consider that M and G are the sets of object voxels for $\hat{M}$ and $\hat{G}$, respectively. As such, $|M|$ corresponds to the volume of the object from $\hat{M}$, for example. Figs. C.2–C.3 illustrate the considered segmentation metrics.

Figure C.2: Illustration of IoU and Dice.

Figure C.3: Average Symmetric Surface Distance (ASSD). ASSD is calculated using the object surfaces, $\beta(\hat{M})$ and $\beta(\hat{G})$, from two binary images, $\hat{M}$ and $\hat{G}$, respectively. For each surface voxel from the object surface of $\hat{M}$, the Euclidean distance to the closest surface voxel of $\hat{G}$ is calculated. The ASSD is the average of all distances calculated from $\hat{M}$ to $\hat{G}$ and $\hat{G}$ to $\hat{M}$.

c.2.1 Intersection over Union (IoU)

IoU, also known as Jaccard index, is a similarity metric between two regions, *e.g.*, the objects of two binary images (Fig. C.2b). It is widely used in object detection problems and ranges from 0 to 1, where a higher value indicates a better matching

between the regions. We use IoU as the cost function of the parameter optimization in [Section 5.4.2](#). Formally, IoU is given by

$$IoU(\hat{M}, \hat{G}) = \frac{|M \cap G|}{|M \cup G|}. \quad (C.5)$$

c.2.2 Dice

Sørensen–Dice coefficient, or simply Dice score, is one of the evaluation metrics most frequently used in medical image segmentation. It measures the amount of overlap between the objects of two binary images ([Fig. C.2c](#)), with values within $[0, 1]$ with 1 indicating perfect segmentation (matching). We use Dice in this thesis in the experiments from [Chapters 3, 5, and 6](#). We can define Dice as

$$Dice(\hat{M}, \hat{G}) = \frac{2|M \cap G|}{|M| + |G|}. \quad (C.6)$$

Dice only gives a global similarity impression between the binary images, and can over/underestimate the real level of matching between them.

c.2.3 Average Symmetric Surface Distance (ASSD)

ASSD is a measure based on the Euclidean distance between the object surfaces of two binary images. ASSD can better capture local differences (e.g., segmentation errors) along the segmented boundaries than Dice ([Fig. C.4](#)). We rely on ASSD for the evaluation of the proposed brain image segmentation approach in [Chapter 3](#).

Figure C.4: Two cases with approximately equal Dice, but different ASSD.

For each surface voxel from the object boundaries of $\hat{M}$, the Euclidean distance to the closest surface voxel of $\hat{G}$ is calculated. The ASSD is the average of all distances calculated from $\hat{M}$ to $\hat{G}$ and $\hat{G}$ to $\hat{M}$ ([Fig. C.3](#)). Formally, we define ASSD by

$$ASSD(\hat{M}, \hat{G}) = \frac{ASD(\hat{M}, \hat{G}) + ASD(\hat{G}, \hat{M})}{2}, \quad (C.7)$$

where $ASD(\hat{M}, \hat{G})$ is the Average Surface Distance from $\hat{M}$ to $\hat{G}$. ASD is defined by

$$ASD(\hat{M}, \hat{G}) = \frac{1}{|\beta(\hat{G})|} \cdot \sum_{p \in \beta(\hat{G})} E(p, \beta(\hat{M})), \quad (C.8)$$

where $\beta(\hat{G})$ is the object surface of $\hat{G}$, $|\beta(\hat{G})|$ denotes its number of voxels, and $E(p, \beta(\hat{M}))$ is the Euclidean distance from the voxel p to $\beta(\hat{M})$, measured in the direction of the local surface normal in $\beta(\hat{M})$.

BIBLIOGRAPHY

- [1] G. J. Tortora and B. H. Derrickson, *Principles of anatomy and physiology*. John Wiley & Sons, 2018.
- [2] R. Lent, *Cem bilhões de neurônios: conceitos fundamentais de neurociência*. Atheneu, 2004.
- [3] K. S. Saladin, *Human anatomy*, vol. 3. McGraw-Hill Education, 5 ed., 2016.
- [4] M. DiLuca and J. Olesen, “The cost of brain diseases: a burden or a challenge?,” *Neuron*, vol. 82, no. 6, pp. 1205–1208, 2014.
- [5] J. Olesen, A. Gustavsson, M. Svensson, H.-U. Wittchen, B. Jönsson, C. S. Group, and E. B. Council, “The economic cost of brain disorders in europe,” *European journal of neurology*, vol. 19, no. 1, pp. 155–162, 2012.
- [6] M. Soltaninejad, G. Yang, T. Lambrou, N. Allinson, T. L. Jones, T. R. Barrick, F. A. Howe, and X. Ye, “Automated brain tumour detection and segmentation using superpixel-based extremely randomized trees in FLAIR MRI,” *International Journal of Computer Assisted Radiology and Surgery*, vol. 12, no. 2, pp. 183–203, 2017.
- [7] M. L. Seghier, A. Ramackhansingh, J. Crinion, A. P. Leff, and C. J. Price, “Lesion identification using unified segmentation-normalisation models and fuzzy clustering,” *Neuroimage*, vol. 41, no. 4, pp. 1253–1266, 2008.
- [8] K. D. Toennies, *Guide to medical image analysis*. Springer, 2017.
- [9] E. Bercovich and M. C. Javitt, “Medical imaging: From roentgen to the digital revolution, and beyond,” *Rambam Maimonides Medical Journal*, vol. 9, no. 4, 2018.
- [10] V. M. Runge, W. R. Nitz, and J. Heverhagen, *The Physics of Clinical MR Taught Through Images*. Thieme, 2018.
- [11] Z. Akkus and et al., “Deep learning for brain MRI segmentation: state of the art and future directions,” *Journal of Digital Imaging*, vol. 30, no. 4, pp. 449–459, 2017.
- [12] J. Liu, M. Li, J. Wang, F. Wu, T. Liu, and Y. Pan, “A survey of MRI-based brain tumor segmentation methods,” *Tsinghua Science and Technology*, vol. 19, no. 6, pp. 578–595, 2014.

- [13] D. Guo, J. Fridriksson, P. Fillmore, C. Rorden, H. Yu, K. Zheng, and S. Wang, "Automated lesion detection on MRI scans using combined unsupervised and supervised methods," *BMC Medical Imaging*, vol. 15, no. 1, p. 50, 2015.
- [14] B. Baars and N. M. Gage, *Fundamentals of cognitive neuroscience: a beginner's guide*. Academic Press, 2013.
- [15] A. R. Palmer, "Symmetry breaking and the evolution of development," *Science*, vol. 306, no. 5697, pp. 828–833, 2004.
- [16] S. Ocklenburg and O. Gunturkun, "Hemispheric asymmetries: the comparative view," *Frontiers in Psychology*, vol. 3, p. 5, 2012.
- [17] M. C. Corballis, "The evolution and genetics of cerebral asymmetry," *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, vol. 364, no. 1519, pp. 867–879, 2009.
- [18] A. W. Toga, K. L. Narr, P. M. Thompson, and E. Luders, "Brain asymmetry: Evolution," in *Encyclopedia of Neuroscience*, pp. 303–311, Academic Press, 2009.
- [19] J. J. Vogel, C. A. Bowers, and D. S. Vogel, "Cerebral lateralization of spatial abilities: a meta-analysis," *Brain and Cognition*, vol. 52, no. 2, pp. 197–204, 2003.
- [20] P. Kalavathi, M. Senthamilselvi, and V. B. Prasath, "Review of computational methods on brain symmetric and asymmetric analysis from neuroimaging techniques," *Technologies*, vol. 5, no. 2, p. 16, 2017.
- [21] K. Hugdahl and R. Westerhausen, *The two halves of the brain: information processing in the cerebral hemispheres*. MIT press, 2010.
- [22] K. Amunts, "Structural indices of asymmetry," *The two halves of the brain*, pp. 145–176, 2010.
- [23] A. A. Woolard and S. Heckers, "Anatomical and functional correlates of human hippocampal volume asymmetry," *Psychiatry Research: Neuroimaging*, vol. 201, no. 1, pp. 48–53, 2012.
- [24] A. Convit, M. J. De Leon, C. Tarshish, S. De Santi, W. Tsui, H. Rusinek, and A. George, "Specific hippocampal volume reductions in individuals at risk for Alzheimer's disease," *Neurobiology of Aging*, vol. 18, no. 2, pp. 131–138, 1997.
- [25] J. G. Csernansky, M. K. Schindler, N. R. Splinter, L. Wang, M. Gado, L. D. Selemon, D. Rastogi-Cruz, J. A. Posener, P. A. Thompson, and M. I. Miller, "Abnormalities of thalamic volume and shape in schizophrenia," *American Journal of Psychiatry*, vol. 161, no. 5, pp. 896–902, 2004.

- [26] L. Wang, S. C. Joshi, M. I. Miller, and J. G. Csernansky, "Statistical analysis of hippocampal asymmetry in schizophrenia," *Neuroimage*, vol. 14, no. 3, pp. 531–545, 2001.
- [27] N. Farid, H. M. Girard, N. Kemmotsu, M. E. Smith, S. W. Magda, W. Y. Lim, R. R. Lee, and C. R. McDonald, "Temporal lobe epilepsy: quantitative MR volumetry in detection of hippocampal atrophy," *Radiology*, vol. 264, no. 2, pp. 542–550, 2012.
- [28] N. Bernasconi, A. Bernasconi, Z. Caramanos, S. B. Antel, F. Andermann, and D. L. Arnold, "Mesial temporal damage in temporal lobe epilepsy: a volumetric MRI study of the hippocampus, amygdala and parahippocampal region," *Brain*, vol. 126, no. 2, pp. 462–469, 2003.
- [29] C. L. Yasuda, C. Valise, A. V. Saúde, A. R. Pereira, F. R. Pereira, A. L. F. Costa, M. E. Morita, L. E. Betting, G. Castellano, C. A. M. Guerreiro, H. Tedeschi, E. Oliveira, and F. Cendes, "Dynamic changes in white and gray matter volume are associated with outcome of surgical treatment in temporal lobe epilepsy," *Neuroimage*, vol. 49, no. 1, pp. 71–79, 2010.
- [30] L. J. Lotspeich and R. D. Ciaranello, "The neurobiology and genetics of infantile autism," in *International Review of Neurobiology*, vol. 35, pp. 87–129, Elsevier, 1993.
- [31] K. Kamnitsas, C. Ledig, V. F. J. Newcombe, J. P. Simpson, A. D. Kane, D. K. Menon, D. Rueckert, and B. Glocker, "Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation," *Medical Image Analysis*, vol. 36, pp. 61–78, 2017.
- [32] H. Chen, Q. Dou, L. Yu, J. Qin, and P.-A. Heng, "VoxResNet: Deep voxelwise residual networks for brain segmentation from 3D MR images," *Neuroimage*, vol. 170, pp. 446–455, 2018.
- [33] M. R. Herbert, D. A. Ziegler, C. K. Deutsch, L. M. O'Brien, D. N. Kennedy, P. A. Filipek, A. I. Bakardjiev, J. Hodgson, M. Takeoka, N. Makris, *et al.*, "Brain asymmetries in autism and developmental language disorder: a nested whole-brain analysis," *Brain*, vol. 128, no. 1, pp. 213–226, 2004.
- [34] S.-L. Liew and *et al.*, "A large, open source dataset of stroke anatomical brain images and manual lesion segmentations," *Scientific Data*, vol. 5, p. 180011, 2018.
- [35] M. Havaei, A. Davy, D. Warde-Farley, A. Biard, A. Courville, Y. Bengio, C. Pal, P.-M. Jodoin, and H. Larochelle, "Brain tumor segmentation with deep neural networks," *Medical Image Analysis*, vol. 35, pp. 18–31, 2017.
- [36] M. J. A. Jansen, *Analysis of liver lesions in dynamic contrast enhanced MR images*. PhD thesis, Utrecht University, 2019.

- [37] M. Goetz, C. Weber, J. Bloecher, B. Stieltjes, H.-P. Meinzer, and K. Maier-Hein, "Extremely randomized trees based brain tumor segmentation," *Proc. of BRATS challenge-MICCAI*, pp. 006–011, 2014.
- [38] E. Geremia, O. Clatz, B. H. Menze, E. Konukoglu, A. Criminisi, and N. Ayache, "Spatial decision forests for MS lesion segmentation in multi-channel magnetic resonance images," *Neuroimage*, vol. 57, no. 2, pp. 378–390, 2011.
- [39] A. Pinto, S. Pereira, H. Correia, J. Oliveira, D. M. Rasteiro, and C. A. Silva, "Brain tumour segmentation based on extremely randomized forest with high-level features," in *IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 3037–3040, 2015.
- [40] S. B. Martins, G. Ruppert, F. Reis, C. L. Yasuda, and A. X. Falcão, "A supervoxel-based approach for unsupervised abnormal asymmetry detection in MR images of the brain," in *IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 882–885, 2019.
- [41] S. B. Martins, J. Bragantini, C. L. Yasuda, and A. X. Falcão, "An adaptive probabilistic atlas for anomalous brain segmentation in MR images," *Medical Physics*, vol. 46, no. 11, pp. 4940–4950, 2019.
- [42] J. M. P. Lötjönen, R. Wolz, J. R. Koikkalainen, L. Thurfjell, G. Walde-mar, H. Soininen, D. Rueckert, A. D. N. Initiative, *et al.*, "Fast and robust multi-atlas segmentation of brain magnetic resonance images," *Neuroimage*, vol. 49, no. 3, pp. 2352–2365, 2010.
- [43] E. I. Zacharaki, S. Wang, S. Chawla, D. Soo Yoo, R. Wolf, E. R. Melhem, and C. Davatzikos, "Classification of brain tumor type and grade using mri texture and shape in a machine learning scheme," *Magnetic Resonance in Medicine*, vol. 62, no. 6, pp. 1609–1618, 2009.
- [44] G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, J. A. Van Der Laak, B. Van Ginneken, and C. I. Sánchez, "A survey on deep learning in medical image analysis," *Medical Image Analysis*, vol. 42, pp. 60–88, 2017.
- [45] A. V. Vasilakos, Y. Tang, and Y. Yao, "Neural networks for computer-aided diagnosis in medicine: A review," *Neurocomputing*, vol. 216, pp. 700–708, 2016.
- [46] S. Aslani, M. Dayan, V. Murino, and D. Sona, "Deep 2D encoder-decoder convolutional neural network for multiple sclerosis lesion segmentation in brain MRI," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pp. 132–141, 2018.
- [47] R. F. Mello and M. A. Ponti, *Machine learning: a practical approach on the statistical learning theory*. Springer, 2018.

- [48] M. E. Celebi and K. Aydin, *Unsupervised learning algorithms*. Springer, 2016.
- [49] V. Hodge and J. Austin, “A survey of outlier detection methodologies,” *Artificial Intelligence Review*, vol. 22, no. 2, pp. 85–126, 2004.
- [50] J. Mourão-Miranda, D. R. Hardoon, T. Hahn, A. F. Marquand, S. C. Williams, J. Shawe-Taylor, and M. Brammer, “Patient classification as an outlier detection problem: an application of the one-class support vector machine,” *Neuroimage*, vol. 58, no. 3, pp. 793–804, 2011.
- [51] S. B. Martins, A. X. Falcão, and A. C. Telea, “BADRESC: Brain anomaly detection based on registration errors and supervoxel classification,” in *International Joint Conference on Biomedical Engineering Systems and Technologies: BIOIMAGING*, pp. 74–81, 2020. Best student paper awards.
- [52] Y. Tang, Y. Tang, M. Han, J. Xiao, and R. M. Summers, “Abnormal chest X-ray identification with generative adversarial one-class classifier,” in *IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 1358–1361, 2019.
- [53] C. Baur, B. Wiestler, S. Albarqouni, and N. Navab, “Deep autoencoding models for unsupervised anomaly segmentation in brain MR images,” in *International MICCAI Brainlesion Workshop*, pp. 161–169, 2018.
- [54] T. Kooi, G. Litjens, B. Van Ginneken, A. Gubern-Mérida, C. I. Sánchez, R. Mann, A. den Heeten, and N. Karssemeijer, “Large scale deep learning for computer aided detection of mammographic lesions,” *Medical Image Analysis*, vol. 35, pp. 303–312, 2017.
- [55] Y. Gao, T. Riklin-Raviv, and S. Bouix, “Shape analysis, a field in need of careful validation,” *Human Brain Mapping*, vol. 35, no. 10, pp. 4965–4978, 2014.
- [56] M. Shakeri and et al., “Statistical shape analysis of subcortical structures using spectral matching,” *Computerized Medical Imaging and Graphics*, vol. 52, pp. 58–71, 2016.
- [57] B. Thyreau, K. Sato, H. Fukuda, and Y. Taki, “Segmentation of the hippocampus by transferring algorithmic knowledge for large cohort processing,” *Medical Image Analysis*, vol. 43, pp. 214–228, 2018.
- [58] D. Sato, S. Hanaoka, Y. Nomura, T. Takenaga, S. Miki, T. Yoshikawa, N. Hayashi, and O. Abe, “A primitive study on unsupervised anomaly detection with an autoencoder in emergency head ct volumes,” in *SPIE Medical Imaging*, p. 105751P, 2018.
- [59] X. Chen, N. Pawlowski, M. Rajchl, B. Glocker, and E. Konukoglu, “Deep generative models in the real-world: An open challenge from medical imaging,” *arXiv preprint arXiv:1806.05452*, 2018.

- [60] A. X. Falcão, J. Stolfi, and R. de Alencar Lotufo, “The image foresting transform: Theory, algorithms, and applications,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 26, no. 1, pp. 19–29, 2004.
- [61] J. E. Vargas-Muñoz, A. S. Chowdhury, E. B. Alexandre, F. L. Galvão, P. A. V. Miranda, and A. X. Falcão, “An iterative spanning forest framework for superpixel segmentation,” *IEEE Transaction on Image Processing*, vol. 28, no. 7, pp. 3477–3489, 2019.
- [62] B. J. Baars and N. M. Gage, *Cognition, brain, and consciousness: Introduction to cognitive neuroscience*. Academic Press, 2010.
- [63] P. A. Abhang, B. W. Gawali, and S. C. Mehrotra, *Introduction to EEG-and speech-based emotion recognition*. Academic Press, 2016.
- [64] D. Purves, *Neuroscience*. NCBI bookshelf, Palgrave Macmillan, 2001.
- [65] H. Rehman and S. Lee, “An efficient automatic midsagittal plane extraction in brain MRI,” *Applied Sciences*, vol. 8, no. 11, p. 2203, 2018.
- [66] G. C. S. Ruppert, L. Teverovskiy, C.-P. Yu, A. X. Falcão, and Y. Liu, “A new symmetry-based method for mid-sagittal plane extraction in neuroimages,” in *IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 285–288, 2011.
- [67] S. Minoshima, R. A. Koeppe, K. A. Frey, and D. E. Kuhl, “Anatomic standardization: linear scaling and nonlinear warping of functional brain images,” *Journal of Nuclear Medicine*, vol. 35, no. 9, pp. 1528–1537, 1994.
- [68] R. S. Alves and J. M. R. S. Tavares, “Computer image registration techniques applied to nuclear medicine images,” in *Computational and experimental biomedical sciences: methods and applications*, pp. 173–191, Springer, 2015.
- [69] M. B. Stegmann, K. Skoglund, and C. Ryberg, “Mid-sagittal plane and mid-sagittal surface optimization in brain MRI using a local symmetry measure,” in *SPIE Medical Imaging*, pp. 568–579, 2005.
- [70] M. Larobina and L. Murino, “Medical image file formats,” *Journal of Digital Imaging*, vol. 27, no. 2, pp. 200–206, 2014.
- [71] M. Brett and et al., “nipy/nibabel: 3.0.0,” December 2008.
- [72] S. H. Joshi, A. Marquina, S. J. Osher, I. Dinov, and J. Darrell, “Image resolution enhancement and its applications to medical image processing,” *Laboratory of Neuroimaging University of California, Los Angeles, CA*, vol. 90095, 2008.

- [73] B. H. Menze, A. Jakab, S. Bauer, J. Kalpathy-Cramer, K. Farahani, J. Kirby, Y. Burren, N. Porz, J. Slotboom, R. Wiest, *et al.*, “The multimodal brain tumor image segmentation benchmark (brats),” *IEEE Transaction on Medical Imaging*, vol. 34, no. 10, pp. 1993–2024, 2014.
- [74] B. K. Rutt and D. H. Lee, “The impact of field strength on image quality in MRI,” *Journal of Magnetic Resonance Imaging*, vol. 6, no. 1, pp. 57–62, 1996.
- [75] M. Jenkinson, C. F. Beckmann, T. E. J. Behrens, M. W. Woolrich, and S. M. Smith, “FSL,” *Neuroimage*, vol. 62, no. 2, pp. 782–790, 2012.
- [76] P. A. Yushkevich, J. Piven, H. C. Hazlett, R. G. Smith, S. Ho, J. C. Gee, and G. Gerig, “User-guided 3D active contour segmentation of anatomical structures: significantly improved efficiency and reliability,” *Neuroimage*, vol. 31, no. 3, pp. 1116–1128, 2006.
- [77] R. Kikinis, S. D. Pieper, and K. G. Vosburgh, “3D Slicer: a platform for subject-specific image analysis, visualization, and clinical support,” in *Intraoperative imaging and image-guided therapy*, pp. 277–289, Springer, 2014.
- [78] J. Juan-Albarracín, E. Fuster-Garcia, J. V. Manjón, M. Robles, F. Aparici, L. Martí-Bonmatí, and J. M. García-Gómez, “Automated glioblastoma segmentation based on a multiparametric structured unsupervised classification,” *PLoS One*, vol. 10, no. 5, p. e0125143, 2015.
- [79] S. B. Martins, A. C. Telea, and A. X. Falcão, “Extending supervoxel-based abnormal brain asymmetry detection to the native image space,” in *IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 450–453, 2019.
- [80] J. V. Manjón and P. Coupé, “volBrain: An online MRI brain volumetry system,” *Frontiers in Neuroinformatics*, vol. 10, 2016.
- [81] L. Martí-Bonmatí and A. Alberich-Bayarri, *Imaging biomarkers: development and clinical integration*. Springer, 2017.
- [82] J. V. Manjón, “MRI preprocessing,” in *Imaging Biomarkers*, pp. 53–63, Springer, 2017.
- [83] J. Mohan, V. Krishnaveni, and Y. Guo, “A survey on the magnetic resonance image denoising methods,” *Biomedical signal processing and control*, vol. 9, pp. 56–69, 2014.
- [84] H. M. Ali, “A new method to remove salt & pepper noise in magnetic resonance images,” in *International Conference on Computer Engineering & Systems (ICCES)*, pp. 155–160, 2016.
- [85] I. Diaz, P. Boulanger, R. Greiner, and A. Murtha, “A critical review of the effects of de-noising algorithms on MRI brain tumor segmentation,” in *IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 3934–3937, 2011.

- [86] C.-Y. Ning, S.-f. Liu, and M. Qu, "Research on removing noise in medical image based on median filter method," in *IEEE International Symposium on IT in Medicine & Education*, vol. 1, pp. 384–388, 2009.
- [87] L. Tan and J. Jiang, "Image processing basics," in *Digital Signal Processing*, pp. 683–765, Academic Press, second edition ed., 2013.
- [88] J. V. Manjón, P. Coupé, A. Buades, D. L. Collins, and M. Robles, "New methods for MRI denoising based on sparseness and self-similarity," *Medical Image Analysis*, vol. 16, no. 1, pp. 18–27, 2012.
- [89] P. Coupé, J. V. Manjón, M. Robles, and D. L. Collins, "Adaptive multiresolution non-local means filter for three-dimensional magnetic resonance image denoising," *IET Image Processing*, vol. 6, no. 5, pp. 558–568, 2012.
- [90] H. M. Golshan, R. P. R. Hasanzadeh, and S. C. Yousefzadeh, "An MRI denoising method using image data redundancy and local SNR estimation," *Magnetic Resonance Imaging*, vol. 31, no. 7, pp. 1206–1217, 2013.
- [91] S. X. Liu, "Symmetry and asymmetry analysis and its implications to computer-aided diagnosis: A review of the literature," *Journal of biomedical informatics*, vol. 42, no. 6, pp. 1056–1064, 2009.
- [92] H. Wu, D. Wang, L. Shi, Z. Wen, and Z. Ming, "Midsagittal plane extraction from brain images based on 3D SIFT," *Physics in Medicine & Biology*, vol. 59, no. 6, p. 1367, 2014.
- [93] R. J. Ferrari, C. H. V. Pinto, and C. A. F. Moreira, "Detection of the midsagittal plane in mr images using a sheetness measure from eigenanalysis of local 3D phase congruency responses," in *International Conference on Image Processing (ICIP)*, pp. 2335–2339, 2016.
- [94] U. Vovk, F. Pernus, and B. Likar, "A review of methods for correction of intensity inhomogeneity in MRI," *IEEE Transaction on Medical Imaging*, vol. 26, no. 3, pp. 405–421, 2007.
- [95] B. Belaroussi, J. Milles, S. Carme, Y. M. Zhu, and H. Benoit-Cattin, "Intensity non-uniformity correction in MRI: existing methods and their validation," *Medical Image Analysis*, vol. 10, no. 2, pp. 234–246, 2006.
- [96] J. D. Gispert, S. Reig, J. Pascau, J. J. Vaquero, P. García-Barreno, and M. Desco, "Method for bias field correction of brain T1-weighted magnetic resonance images minimizing segmentation error," *Human Brain Mapping*, vol. 22, no. 2, pp. 133–144, 2004.
- [97] B. Likar, M. A. Viergever, and F. Pernus, "Retrospective correction of MR intensity inhomogeneity by information minimization," *IEEE Transaction on Medical Imaging*, vol. 20, no. 12, pp. 1398–1410, 2001.

- [98] U. Vovk, F. Pernuš, and B. Likar, "MRI intensity inhomogeneity correction by combining intensity and spatial information," *Physics in Medicine & Biology*, vol. 49, no. 17, p. 4119, 2004.
- [99] J. V. Manjón, J. J. Lull, J. Carbonell-Caballero, G. García-Martí, L. Martí-Bonmatí, and M. Robles, "A nonparametric MRI inhomogeneity correction method," *Medical Image Analysis*, vol. 11, no. 4, pp. 336–345, 2007.
- [100] C. T. Larsen, J. E. Iglesias, and K. Van Leemput, "N3 bias field correction explained as a Bayesian modeling method," in *Bayesian and graphical models for biomedical imaging*, pp. 1–12, 2014.
- [101] J. G. Sled, A. P. Zijdenbos, and A. C. Evans, "A nonparametric method for automatic correction of intensity nonuniformity in MRI data," *IEEE Transaction on Medical Imaging*, vol. 17, no. 1, pp. 87–97, 1998.
- [102] N. J. Tustison, B. B. Avants, P. A. Cook, Y. Zheng, A. Egan, P. A. Yushkevich, and J. C. Gee, "N4ITK: improved N3 bias correction," *IEEE Transaction on Medical Imaging*, vol. 29, no. 6, pp. 1310–1320, 2010.
- [103] H. J. Johnson, M. M. McCormick, and L. Ibanez, *The ITK Software Guide Book 2: Design and Functionality-Volume 2*. Kitware, Inc., 2015.
- [104] J. E. Iglesias and M. R. Sabuncu, "Multi-atlas segmentation of biomedical images: a survey," *Medical Image Analysis*, vol. 24, no. 1, pp. 205–219, 2015.
- [105] W. M. Wells, P. Viola, H. Atsumi, S. Nakajima, and R. Kikinis, "Multi-modal volume registration by maximization of mutual information," *Medical Image Analysis*, vol. 1, no. 1, pp. 35–51, 1996.
- [106] V. R. S. Mani and S. Arivazhagan, "Survey of medical image registration," *Journal of Biomedical Engineering and Technology*, vol. 1, no. 2, pp. 8–25, 2013.
- [107] S. Klein, M. Staring, K. Murphy, M. A. Viergever, and J. P. W. Pluim, "elastix: A toolbox for intensity-based medical image registration," *IEEE Transaction on Medical Imaging*, vol. 29, no. 1, pp. 196–205, 2010.
- [108] V. S. Fonov, A. C. Evans, R. C. McKinstry, C. R. Almli, and D. L. Collins, "Unbiased nonlinear average age-appropriate brain templates from birth to adulthood," *Neuroimage*, vol. 47, p. S102, 2009.
- [109] J. V. Manjón, S. F. Eskildsen, P. Coupé, J. E. Romero, D. L. Collins, and M. Robles, "Nonlocal intracranial cavity extraction," *International journal of biomedical imaging*, vol. 2014, 2014.
- [110] I. Despotović, B. Goossens, and W. Philips, "MRI segmentation of the human brain: challenges, methods, and applications," *Computational and mathematical methods in medicine*, vol. 2015, 2015.

- [111] S. M. Smith, “Fast robust automated brain extraction,” *Human Brain Mapping*, vol. 17, no. 3, pp. 143–155, 2002.
- [112] J. Doshi, G. Erus, Y. Ou, B. Gaonkar, and C. Davatzikos, “Multi-atlas skull-stripping,” *Academic radiology*, vol. 20, no. 12, pp. 1566–1576, 2013.
- [113] S. Roy, J. A. Butman, D. L. Pham, A. D. N. Initiative, *et al.*, “Robust skull stripping using multiple MR image contrasts insensitive to pathology,” *Neuroimage*, vol. 146, pp. 132–147, 2017.
- [114] J. Kleesiek, G. Urban, A. Hubert, D. Schwarz, K. Maier-Hein, M. Bendszus, and A. Biller, “Deep MRI brain extraction: A 3D convolutional neural network for skull stripping,” *Neuroimage*, vol. 129, pp. 460–469, 2016.
- [115] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*. Pearson Education, 3 ed., 2007.
- [116] A. X. Falcão, B. S. Cunha, and R. A. Lotufo, “Design of connected operators using the image foresting transform,” in *SPIE Medical Imaging*, vol. 4322, pp. 468–479, 2001.
- [117] F. A. Andaló, P. A. V. Miranda, R. S. Torres, and A. X. Falcão, “Shape feature extraction and description based on tensor scale,” *Pattern Recognition*, vol. 43, no. 1, pp. 26–36, 2010.
- [118] A. M. Freitas, R. S. Torres, and P. A. V. Miranda, “TSS & TSB: Tensor scale descriptors within circular sectors for fast shape retrieval,” *Pattern Recognition Letters*, vol. 83, pp. 303–311, 2016.
- [119] P. A. V. Miranda, A. X. Falcão, and T. V. Spina, “Riverbed: A novel user-steered image segmentation method based on optimum boundary tracking,” *IEEE Transaction on Image Processing*, vol. 21, no. 6, pp. 3042–3052, 2012.
- [120] R. Phellan, A. X. Falcão, and J. K. Udupa, “Medical image segmentation via atlases and fuzzy object models: Improving efficacy through optimum object search and fewer models,” *Medical Physics*, vol. 43, no. 1, pp. 401–410, 2016.
- [121] J. Bragantini, S. B. Martins, C. Castelo-Fernandez, and A. X. Falcão, “Graph-based image segmentation using dynamic trees,” in *Iberoamerican Congress on Pattern Recognition (CIARP)*, pp. 470–478, 2018.
- [122] A. M. Sousa, S. B. Martins, A. X. Falcão, F. Reis, E. Bagatin, and K. Irion, “AL-TIS: A fast and automatic lung and trachea CT-image segmentation method,” *Medical Physics*, vol. 46, no. 11, pp. 4970–4982, 2019.
- [123] M. A. T. Condori, F. A. M. Cappabianco, A. X. Falcão, and P. A. V. Miranda, “An extension of the differential image foresting transform and its application to superpixel generation,” *Journal of Visual Communication and Image Representation*, p. 102748, 2020.

- [124] F. Belém, L. Melo, S. J. F. Guimarães, and A. X. Falcão, “The importance of object-based seed sampling for superpixel segmentation,” in *Conference on Graphics, Patterns and Images (SIBGRAPI)*, pp. 108–115, 2019.
- [125] A. X. Falcão, L. da Fontoura-Costa, and B. S. Da Cunha, “Multiscale skeletons by image foresting transform and its application to neuromorphometry,” *Pattern Recognition*, vol. 35, no. 7, pp. 1571–1582, 2002.
- [126] J. P. Papa, A. X. Falcão, and C. T. N. Suzuki, “Supervised pattern classification based on optimum-path forest,” *International Journal of Imaging Systems and Technology*, vol. 19, no. 2, pp. 120–131, 2009.
- [127] W. P. Amorim, A. X. Falcão, J. P. Papa, and M. H. Carvalho, “Improving semi-supervised learning through optimum connectivity,” *Pattern Recognition*, vol. 60, pp. 72–85, 2016.
- [128] W. P. Amorim, A. X. Falcão, and J. P. Papa, “Multi-label semi-supervised classification through optimum-path forest,” *Information Sciences*, vol. 465, pp. 86–104, 2018.
- [129] L. M. Rocha, F. A. M. Cappabianco, and A. X. Falcão, “Data clustering as an optimum-path forest problem with applications in image analysis,” *International Journal of Imaging Systems and Technology*, vol. 19, no. 2, pp. 50–68, 2009.
- [130] A. E. Montero and A. X. Falcão, “A divide-and-conquer clustering approach based on optimum-path forest,” in *Conference on Graphics, Patterns and Images (SIBGRAPI)*, pp. 416–423, 2018.
- [131] K. C. Ciesielski, A. X. Falcão, and P. A. V. Miranda, “Path-value functions for which dijkstra’s algorithm returns optimal mapping,” *Journal of Mathematical Imaging and Vision*, vol. 60, no. 7, pp. 1025–1036, 2018.
- [132] P. A. V. Miranda and L. A. C. Mansilla, “Oriented image foresting transform segmentation by seed competition,” *IEEE Transaction on Image Processing*, vol. 23, no. 1, pp. 389–398, 2014.
- [133] J. Cousty, G. Bertrand, L. Najman, and M. Couprie, “Watershed cuts: Thinnings, shortest path forests, and topological watersheds,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 32, no. 5, pp. 925–939, 2009.
- [134] F. A. M. Cappabianco, A. X. Falcão, C. L. Yasuda, and J. K. Udupa, “Brain tissue MR-image segmentation via optimum-path forest clustering,” *Computer Vision and Image Understanding*, vol. 116, no. 10, pp. 1047–1059, 2012.
- [135] J. Yan, Y. Yu, X. Zhu, Z. Lei, and S. Z. Li, “Object detection by labeling superpixels,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5107–5116, 2015.

- [136] D. Stutz, A. Hermans, and B. Leibe, “Superpixels: An evaluation of the state-of-the-art,” *Computer Vision and Image Understanding*, vol. 166, pp. 1–27, 2018.
- [137] W. Wu, A. Y. C. Chen, L. Zhao, and J. J. Corso, “Brain tumor detection and segmentation in a CRF (conditional random fields) framework with pixel-pairwise affinity and superpixel-level features,” *International Journal of Computer Assisted Radiology and Surgery*, vol. 9, no. 2, pp. 241–253, 2014.
- [138] M. S. Atkins and B. T. Mackiewicz, “Fully automatic segmentation of the brain in MRI,” *IEEE Transaction on Medical Imaging*, vol. 17, no. 1, pp. 98–107, 1998.
- [139] N. Gordillo, E. Montseny, and P. Sobrevilla, “State of the art survey on MRI brain tumor segmentation,” *Magnetic Resonance Imaging*, vol. 31, no. 8, pp. 1426–1438, 2013.
- [140] E.-S. A. El-Dahshan, H. M. Mohsen, K. Revett, and A.-B. M. Salem, “Computer-aided diagnosis of human brain tumor through MRI: A survey and a new algorithm,” *Expert Systems with Applications*, vol. 41, no. 11, pp. 5526–5545, 2014.
- [141] S. B. Martins, T. V. Spina, C. L. Yasuda, and A. X. Falcão, “A multi-object statistical atlas adaptive for deformable registration errors in anomalous medical image segmentation,” in *SPIE Medical Imaging*, vol. 10133, pp. 691–698, 2017.
- [142] J. C. Mazziotta, A. W. Toga, A. Evans, P. Fox, J. Lancaster, *et al.*, “A probabilistic atlas of the human brain: theory and rationale for its development,” *Neuroimage*, vol. 2, no. 2, pp. 89–101, 1995.
- [143] A. Guimond, J. Meunier, and J.-P. Thirion, “Average brain models: A convergence study,” *Computer Vision and Image Understanding*, vol. 77, no. 2, pp. 192–210, 2000.
- [144] B. Fischl, “Freesurfer,” *Neuroimage*, vol. 62, no. 2, pp. 774–781, 2012.
- [145] B. B. Avants, N. J. Tustison, J. Wu, P. A. Cook, and J. C. Gee, “An open source multivariate framework for n-tissue segmentation with evaluation on public data,” *Neuroinformatics*, vol. 9, no. 4, pp. 381–400, 2011.
- [146] M. Havaei, N. Guizard, H. Larochelle, and P.-M. Jodoin, “Deep learning trends for focal brain pathology segmentation in MRI,” in *Machine Learning for Health Informatics: State-of-the-Art and Future Challenges*, pp. 125–148, 2016.

- [147] M. Cabezas, A. Oliver, X. Lladó, J. Freixenet, and M. B. Cuadra, "A review of atlas-based segmentation for magnetic resonance brain images," *Computer methods and programs in biomedicine*, vol. 104, no. 3, pp. e158–e177, 2011.
- [148] F. Malmberg, I. Nyström, A. Mehnert, C. Engstrom, and E. Bengtsson, "Relaxed image foresting transforms for interactive volume image segmentation," in *SPIE Medical Imaging*, vol. 7623, p. 762340, 2010.
- [149] N. Moya, "Interactive segmentation of multiple 3D objects in medical images by optimum graph cuts," Master's thesis, University of Campinas, Brazil, 2015.
- [150] Y. Tong, J. K. Udupa, D. Odhner, C. Wu, S. J. Schuster, and D. A. Torigian, "Disease quantification on pet/ct images without explicit object delineation," *Medical Image Analysis*, vol. 51, pp. 169–183, 2019.
- [151] G. E. Christensen, S. C. Joshi, and M. Miller, "Volumetric transformation of brain anatomy," *IEEE Transaction on Medical Imaging*, vol. 16, pp. 864–877, 1997.
- [152] D. L. Collins, C. J. Holmes, T. M. Peters, and A. C. Evans, "Automatic 3D model-based neuroanatomical segmentation," *Human Brain Mapping*, vol. 3, no. 3, pp. 190–208, 1995.
- [153] D. L. Pham, C. Xu, and J. L. Prince, "Current methods in medical image segmentation," *Annual Review of Biomedical Engineering*, vol. 2, pp. 315–337, 2000.
- [154] H. Park, A. Hero, P. Bland, M. Kessler, J. Seo, and C. Meyer, "Construction of abdominal probabilistic atlases and their value in segmentation of normal organs in abdominal CT scans," *IEICE Transactions on Information and Systems*, vol. 93, no. 8, pp. 2291–2301, 2010.
- [155] K. M. Pohl, J. Fisher, W. E. L. Grimson, R. Kikinis, and W. M. Wells, "A bayesian model for joint segmentation and registration," *Neuroimage*, vol. 31, no. 1, pp. 228–239, 2006.
- [156] V. Grau, A. U. J. Mewes, M. Alcaniz, R. Kikinis, and S. K. Warfield, "Improved watershed transform for medical image segmentation using prior information," *IEEE Transaction on Medical Imaging*, vol. 23, no. 4, pp. 447–458, 2004.
- [157] B. Fischl, D. H. Salat, E. Busa, M. Albert, M. Dieterich, C. Haselgrove, A. van Der Kouwe, R. Killiany, D. Kennedy, S. Klaveness, *et al.*, "Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain," *Neuron*, vol. 33, no. 3, pp. 341–355, 2002.
- [158] J. Ashburner and K. J. Friston, "Unified segmentation," *Neuroimage*, vol. 26, no. 3, pp. 839–851, 2005.

- [159] Y. Huo, A. J. Plassard, A. Carass, S. M. Resnick, D. L. Pham, J. L. Prince, and B. A. Landman, "Consistent cortical reconstruction and multi-atlas brain segmentation," *Neuroimage*, vol. 138, pp. 197–210, 2016.
- [160] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 39, no. 1, pp. 1–22, 1977.
- [161] P. A. V. Miranda, A. X. Falcão, and J. K. Udupa, "Cloud bank: A multiple clouds model and its use in mr brain image segmentation," in *IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 506–509, 2009.
- [162] J. E. Romero, J. V. Manjón, J. Tohka, P. Coupé, and M. Robles, "NABS: Non-local automatic brain hemisphere segmentation," *Magnetic Resonance Imaging*, vol. 33, no. 4, pp. 474–484, 2015.
- [163] H. Wang, J. W. Suh, S. R. Das, J. B. Pluta, C. Craige, and P. Yushkevich, "Multi-atlas segmentation with joint label fusion," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 35, no. 3, pp. 611–623, 2013.
- [164] S. K. Warfield, K. H. Zou, and W. M. Wells, "Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation," *IEEE Transaction on Medical Imaging*, vol. 23, no. 7, pp. 903–921, 2004.
- [165] G. Sanroma, G. Wu, Y. Gao, and D. Shen, "Learning to rank atlases for multiple-atlas segmentation," *IEEE Transaction on Medical Imaging*, vol. 33, no. 10, pp. 1939–1953, 2014.
- [166] E. Konukoglu, B. Glocker, D. Zikic, and A. Criminisi, "Neighbourhood approximation using randomized forests," *Medical Image Analysis*, vol. 17, no. 7, pp. 790–804, 2013.
- [167] P. Aljabar, R. A. Heckemann, A. Hammers, J. V. Hajnal, and D. Rueckert, "Multi-atlas based segmentation of brain images: atlas selection and its effect on accuracy," *Neuroimage*, vol. 46, no. 3, pp. 726–738, 2009.
- [168] P. Aljabar, R. Heckemann, A. Hammers, J. V. Hajnal, and D. Rueckert, "Classifier selection strategies for label fusion using large atlas databases," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pp. 523–531, 2007.
- [169] Q. Xie and D. Ruan, "Low-complexity atlas-based prostate segmentation by combining global, regional, and local metrics," *Medical Physics*, vol. 41, no. 4, 2014.

- [170] A. J. Asman and B. A. Landman, "Hierarchical performance estimation in the statistical label fusion framework," *Medical Image Analysis*, vol. 18, no. 7, pp. 1070–1081, 2014.
- [171] S. Nouranian, S. S. Mahdavi, I. Spadinger, W. J. Morris, S. E. Salcudean, and P. Abolmaesumi, "A multi-atlas-based segmentation framework for prostate brachytherapy," *IEEE Transaction on Medical Imaging*, vol. 34, no. 4, pp. 950–961, 2014.
- [172] T. R. Langerak, F. F. Berendsen, U. A. van der Heide, A. N. T. J. Kotte, and J. P. W. Pluim, "Multiatlas-based segmentation with preregistration atlas selection," *Medical Physics*, vol. 40, no. 9, p. 091701, 2013.
- [173] A. K. Jain, "Data clustering: 50 years beyond K-means," *Pattern Recognition Letters*, vol. 31, no. 8, pp. 651–666, 2010.
- [174] J. Ashburner and K. J. Friston, "Nonlinear spatial normalization using basis functions," *Human Brain Mapping*, vol. 7, no. 4, pp. 254–266, 1999.
- [175] B. C. Vemuri, J. Ye, Y. Chen, and C. M. Leonard, "A level-set based approach to image registration," in *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, pp. 86–93, 2000.
- [176] R. A. Heckemann, J. V. Hajnal, P. Aljabar, D. Rueckert, and A. Hammers, "Automatic anatomical brain MRI segmentation combining label propagation and decision fusion," *Neuroimage*, vol. 33, no. 1, pp. 115–126, 2006.
- [177] X. Artaechevarria, A. Muñoz-Barrutia, and C. Ortiz-de Solórzano, "Efficient classifier generation and weighted voting for atlas-based segmentation: two small steps faster and closer to the combination oracle," in *SPIE Medical Imaging*, vol. 6914, pp. 6914 – 6914 – 9, 2008.
- [178] X. Artaechevarria, A. Munoz-Barrutia, and C. Ortiz-de Solórzano, "Combination strategies in multi-atlas image segmentation: application to brain MR data," *IEEE Transaction on Medical Imaging*, vol. 28, no. 8, pp. 1266–1277, 2009.
- [179] J. Nie and D. Shen, "Automated segmentation of mouse brain images using multi-atlas multi-roI deformation and label fusion," *Neuroinformatics*, vol. 11, no. 1, pp. 35–45, 2013.
- [180] R. Datteri, A. J. Asman, B. A. Landman, and B. M. Dawant, "Applying the algorithm" assessing quality using image registration circuits" (aquire) to multi-atlas segmentation," in *SPIE Medical Imaging*, vol. 9034, p. 90341F, 2014.

- [181] P. Coupé, J. V. Manjón, V. Fonov, J. Pruessner, M. Robles, and D. L. Collins, "Patch-based segmentation using expert priors: Application to hippocampus and ventricle segmentation," *Neuroimage*, vol. 54, no. 2, pp. 940–954, 2011.
- [182] G. Sanroma, G. Wu, K. Thung, Y. Guo, and D. Shen, "Novel multi-atlas segmentation by matrix completion," in *International Workshop on Machine Learning in Medical Imaging*, pp. 207–214, 2014.
- [183] B. Patenaude, S. M. Smith, D. N. Kennedy, and M. Jenkinson, "A bayesian model of shape and appearance for subcortical brain segmentation," *Neuroimage*, vol. 56, no. 3, pp. 907–922, 2011.
- [184] C. Cortes and V. Vapnik, "Support-vector networks," *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.
- [185] J. Ashburner, C. Hutton, R. Frackowiak, I. Johnsrude, C. Price, and K. Friston, "Identifying global anatomical differences: Deformation-based morphometry," *Human Brain Mapping*, vol. 6, no. 5-6, pp. 348–357, 1998.
- [186] P. Rajpurkar and et al., "Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning," *arXiv preprint arXiv:1711.05225*, 2017.
- [187] E. Yates, L. Yates, and H. Harvey, "Machine learning "red dot": open-source, cloud, deep convolutional neural networks in chest radiograph binary normality classification," *Clinical radiology*, vol. 73, no. 9, pp. 827–831, 2018.
- [188] I. Beheshti, H. Demirel, H. Matsuda, A. D. N. Initiative, et al., "Classification of Alzheimer's disease and prediction of mild cognitive impairment-to-Alzheimer's conversion from structural magnetic resource imaging using feature ranking and a genetic algorithm," *Computers in biology and medicine*, vol. 83, pp. 109–119, 2017.
- [189] S. Basaia, F. Agosta, L. Wagner, E. Canu, G. Magnani, R. Santangelo, M. Filippi, A. D. N. Initiative, et al., "Automated classification of Alzheimer's disease and mild cognitive impairment using a single MRI and deep neural networks," *Neuroimage: Clinical*, vol. 21, p. 101645, 2019.
- [190] Y. R. Fung, Z. Guan, R. Kumar, J. Y. Wu, and M. Fiterau, "Alzheimer's disease brain MRI classification: Challenges and insights," *arXiv preprint arXiv:1906.04231*, 2019.
- [191] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.-A. Manzagol, "Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion," *Journal of Machine Learning Research*, vol. 11, no. Dec, pp. 3371–3408, 2010.

- [192] J. Masci, U. Meier, D. Cireşan, and J. Schmidhuber, “Stacked convolutional auto-encoders for hierarchical feature extraction,” in *International Conference on Artificial Neural Networks*, pp. 52–59, 2011.
- [193] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [194] M. Ranzato, F. J. Huang, Y. L. Boureau, and Y. LeCun, “Unsupervised learning of invariant feature hierarchies with applications to object recognition,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1–8, 2007.
- [195] M. Chen, X. Shi, Y. Zhang, D. Wu, and M. Guizani, “Deep features learning for medical image analysis with convolutional autoencoder neural network,” *IEEE Transactions on Big Data*, 2017.
- [196] L. Gondara, “Medical image denoising using convolutional denoising autoencoders,” in *IEEE Intern. Conf. on Data Mining Workshops*, pp. 241–246, 2016.
- [197] J. Mehta and A. Majumdar, “RODEO: robust DE-aliasing autoencOder for real-time medical image reconstruction,” *Pattern Recognition*, vol. 63, pp. 499–510, 2017.
- [198] V. Ghodrati, J. Shao, M. Bydder, Z. Zhou, W. Yin, K.-L. Nguyen, Y. Yang, and P. Hu, “MR image reconstruction using deep learning: evaluation of network structure and loss functions,” *Quantitative imaging in medicine and surgery*, vol. 9, no. 9, p. 1516, 2019.
- [199] V. Nair and G. E. Hinton, “Rectified linear units improve restricted boltzmann machines,” in *International Conference on Machine Learning (ICML)*, pp. 807–814, 2010.
- [200] I. Sutskever, J. Martens, G. Dahl, and G. Hinton, “On the importance of initialization and momentum in deep learning,” in *International Conference on Machine Learning (ICML)*, pp. 1139–1147, 2013.
- [201] L. M. Manevitz and M. Yousef, “One-class SVMs for document classification,” *Journal of Machine Learning Research*, vol. 2, no. Dec, pp. 139–154, 2001.
- [202] L. v. d. Maaten and G. Hinton, “Visualizing data using t-SNE,” *Journal of Machine Learning Research*, vol. 9, pp. 2579–2605, Nov 2008.
- [203] P. E. Rauber, A. X. Falcão, and A. C. Telea, “Projections as visual aids for classification system design,” *Information Visualization*, vol. 17, no. 4, pp. 282–305, 2018.

- [204] J. Dai, Y. Li, K. He, and J. Sun, "R-FCN: Object detection via region-based fully convolutional networks," in *Advances in neural information processing systems*, pp. 379–387, 2016.
- [205] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, real-time object detection," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [206] S. Shen, A. J. Szameitat, and A. Sterr, "An improved lesion detection approach based on similarity measurement between fuzzy intensity segmentation and spatial probability maps," *Magnetic Resonance Imaging*, vol. 28, no. 2, pp. 245–254, 2010.
- [207] A. P. Zijdenbos, R. Forghani, and A. C. Evans, "Automatic" pipeline" analysis of 3-D MRI data for clinical trials: application to multiple sclerosis," *IEEE Transaction on Medical Imaging*, vol. 21, no. 10, pp. 1280–1291, 2002.
- [208] X. Tomas and S. K. Warfield, "Fully-automatic generation of training points for automatic multiple sclerosis segmentation," *Work. Med. Image Anal. Mult. Scler*, pp. 5–14, 2009.
- [209] D.-J. Kroon, E. Van Oort, and K. Slump, "Multiple sclerosis detection in multispectral magnetic resonance images with principal components analysis," *Grand Challenge Work.: Multiple Sclerosis. Lesion Segmentation Challenge*, pp. 1–14, 2008.
- [210] M. Prastawa, E. Bullitt, S. Ho, and G. Gerig, "A brain tumor segmentation framework based on outlier detection," *Medical Image Analysis*, vol. 8, no. 3, pp. 275–283, 2004.
- [211] P. Geurts, D. Ernst, and L. Wehenkel, "Extremely randomized trees," *Machine Learning*, vol. 63, no. 1, pp. 3–42, 2006.
- [212] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, p. 436, 2015.
- [213] S. Pouyanfar, S. Sadiq, Y. Yan, H. Tian, Y. Tao, M. P. Reyes, M.-L. Shyu, S.-C. Chen, and S. S. Iyengar, "A survey on deep learning: Algorithms, techniques, and applications," *ACM Computing Surveys*, vol. 51, no. 5, p. 92, 2018.
- [214] K. Qi, H. Yang, C. Li, Z. Liu, M. Wang, Q. Liu, and S. Wang, "X-net: Brain stroke lesion segmentation based on depthwise separable convolution and long-range dependencies," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pp. 247–255, 2019.
- [215] S. Shen, A. J. Szameitat, and A. Sterr, "Detection of infarct lesions from single MRI modality using inconsistency between voxel intensity and spatial location—a 3D automatic approach," *IEEE Transaction on Information Technology in Biomedicine*, vol. 12, no. 4, pp. 532–540, 2008.

- [216] J. C. Bezdek, *Pattern recognition with fuzzy objective function algorithms*. Springer Science & Business Media, 2013.
- [217] X. Chen and E. Konukoglu, “Unsupervised detection of lesions in brain mri using constrained adversarial auto-encoders,” *arXiv preprint arXiv:1806.04972*, 2018.
- [218] A. Atlason, H. E. Love, S. Sigurdsson, V. Gudnason, and L. M. Ellingsen, “Unsupervised brain lesion segmentation from MRI using a convolutional autoencoder,” in *SPIE Medical Imaging*, vol. 10949, pp. 372 – 378, 2019.
- [219] N. Otsu, “A threshold selection method from gray-level histograms,” *IEEE Trans. on systems, man, and cybernetics*, vol. 9, no. 1, pp. 62–66, 1979.
- [220] J. R. Taylor, N. Williams, R. Cusack, T. Auer, M. A. Shafto, M. Dixon, L. K. Tyler, R. N. Henson, *et al.*, “The cambridge centre for ageing and neuroscience (Cam-CAN) data repository: structural and functional mri, meg, and cognitive data from a cross-sectional adult lifespan sample,” *Neuroimage*, vol. 144, pp. 262–269, 2017.
- [221] O. Commowick and et al., “Objective evaluation of multiple sclerosis lesion segmentation using a data management and processing infrastructure,” *Scientific reports*, vol. 8, no. 1, p. 13650, 2018.
- [222] K. Simonyan, A. Vedaldi, and A. Zisserman, “Deep inside convolutional networks: Visualising image classification models and saliency maps,” *arXiv preprint arXiv:1312.6034*, 2013.
- [223] J. P. W. Pluim, J. B. A. Maintz, and M. A. Viergever, “Mutual-information-based registration of medical images: a survey,” *IEEE Transaction on Medical Imaging*, vol. 22, no. 8, pp. 986–1004, 2003.

ACKNOWLEDGMENTS

To my family: For supporting me in my dreams, and for their endless love. I am really grateful to my parents, Beth and Cesar, for your love and hard work in providing me with many opportunities in my life. I also thank my sweet grandmother, Rosalina, who I feel is looking after me somewhere.

To my beloved Fabiana: For your endless love and support over these long years. There were so many moments of concern in which a simple hug from you calmed me and lifted me up. I thank God for having you. Thank you, my love, for giving the best gift of my life, our little girl, Margot, who comes to bless our family.

To Alex: For being a great example of a professional, leader, and supervisor. I was privileged to have worked with you. Thank you for your support and countless pieces of advice from which I have learned a lot.

To Falcão: For being an exceptional supervisor. You have introduced me to the research journey and always believed in me. I will always be grateful for your support, guidance, and patience.

To IFSP: For the opportunity to flourish my passion for teaching and for supporting me to complete my doctorate. In particular, my friend, Tiago Carvalho, an example of a researcher for me. Thank you for your support as a co-worker and friend. To the IFSP Human Resources team, Joana and Guilherme, for their exceptional work and support throughout this time.

To the UNICAMP staff: For the exceptional support during my PhD. I especially thank Cris, Wilson, and Denise for always being so attentive and kind to me.

To my LIDS friends: For being part of my life during this journey. I especially thank Azael, Barbara, and my little brother, Jordão, for all great moments, discussions, and invaluable help with my PhD.

To my UU friends: For sharing all great moments together. In special, my other little brother, Eduardo, for the company and memorable table tennis matches; Federica, my proactive and enthusiastic Italian friend who started her path at UU with me; and Anouk, my dear Dutch friend for her support in my thesis and sharing great cultural moments with me.

To Nikolas and Marilia: For being our family in the Netherlands. I will never forget so many great moments and memories together. I hope to be with you soon.

I would also like to thank the research funding provided by FAPESP (São Paulo Research Foundation, grant numbers 2014/12236-1) and CAPES (Coordination for the Improvement of Higher Education Personnel, financial code 001).

COLOPHON

This document was typeset using the typographical look-and-feel `classicthesis` developed by André Miede. The style was inspired by Robert Bringhurst’s seminal book on typography “*The Elements of Typographic Style*”. `classicthesis` is available for both $\text{\LaTeX}$ and $\text{\LyX}$:

<http://code.google.com/p/classicthesis/>

Final Version as of October 6, 2020 (`classicthesis`).