

Master thesis : Matching ex vivo MRI scan and histological slices from multiple sclerosis patients

Auteur : Guittet, Thomas

Promoteur(s) : Phillips, Christophe; LOMMERS, Emilie

Faculté : Faculté des Sciences appliquées

Diplôme : Master en ingénieur civil biomédical, à finalité spécialisée

Année académique : 2025-2026

URI/URL : <http://hdl.handle.net/2268.2/26038>

Avertissement à l'attention des usagers :

Tous les documents placés en accès ouvert sur le site le site MatheO sont protégés par le droit d'auteur. Conformément aux principes énoncés par la "Budapest Open Access Initiative"(BOAI, 2002), l'utilisateur du site peut lire, télécharger, copier, transmettre, imprimer, chercher ou faire un lien vers le texte intégral de ces documents, les disséquer pour les indexer, s'en servir de données pour un logiciel, ou s'en servir à toute autre fin légale (ou prévue par la réglementation relative au droit d'auteur). Toute utilisation du document à des fins commerciales est strictement interdite.

Par ailleurs, l'utilisateur s'engage à respecter les droits moraux de l'auteur, principalement le droit à l'intégrité de l'oeuvre et le droit de paternité et ce dans toute utilisation que l'utilisateur entreprend. Ainsi, à titre d'exemple, lorsqu'il reproduira un document par extrait ou dans son intégralité, l'utilisateur citera de manière complète les sources telles que mentionnées ci-dessus. Toute utilisation non explicitement autorisée ci-avant (telle que par exemple, la modification du document ou son résumé) nécessite l'autorisation préalable et expresse des auteurs ou de leurs ayants droit.

Matching ex vivo MRI scan and histological slices from multiple sclerosis patients

Guittet Thomas

Thesis presented to obtain the degree of :
Master of Science in Biomedical Engineering

Thesis supervisor :
Phillips Christophe
Lommers Emilie

Academic year: **2025 - 2026**

Acknowledgments

This Master's thesis represents the end of my studies in Biomedical Engineering and my first real immersion into the world of scientific research. Throughout the development of this project, I faced multiple technical and methodological challenges, but I always felt supported and guided by the right people. I certainly could not have completed this work without the contributions of several important individuals whom I would like to sincerely thank.

First and foremost, I would like to express my sincere gratitude to my supervisor, Prof. Christophe Phillips. Thank you for giving me the opportunity to work on this project. You always offered excellent scientific advices but what I appreciated the most was the patience, kindness, and understanding with which you delivered it. Thank you for always leaving your door open, for taking the time to listen to my ideas, and for pointing me in the right direction. Your guidance and your feedback made all the difference in bringing this work to life.

I am equally grateful to my co-supervisor, Dr. Emilie Lommers, for introducing me to this project. Given my engineering background, your clinical expertise provided a new point of view that we rarely have the opportunity to encounter as students, allowing me to deepen my understanding of a complex disease like multiple sclerosis. Furthermore, this opportunity enabled me to discover and learn about essential medical research techniques, particularly histology. Thank you for bridging the gap between engineering algorithms and clinical realities, and for providing the medical context that gave true meaning to this entire pipeline.

I would also like to extend my sincere thanks to the entire team at the Cyclotron Research Centre (CRC) for their warm welcome. Working in such a scientific environment alongside their researchers was a true privilege. Thank you to everyone who provided technical advice, data access, or simply a friendly and motivating atmosphere during my time in the lab.

Lastly, I cannot express enough my gratitude to my family and friends for their support throughout this project, and more generally throughout all stages of my studies. I feel incredibly lucky to be surrounded by all those people. I am particularly grateful to my parents, and especially my mother, for their trust, for always pushing me to do my best, and for keeping me on track during the most intense phases of my degree.

Summary

In the context of Multiple sclerosis (MS) research, in order to detect the underlying cortical pathologies, we can rely on ultra-high-field (7T) quantitative magnetic resonance imaging (qMRI) as an alternative to the histological modality. Establishing a precise correlation between those 2 modalities can pave the way for improved in vivo detection and monitoring of the disease. Establishing this correspondence is very challenging due to the dimensionality mismatch between 3D MRI volumes and 2D histological sections, further compounded by severe non-linear tissue deformations such as shrinkage, stretching, and tearing introduced during chemical fixation and sections slicing.

This master's thesis presents the development of an automated, computational co-registration pipeline designed to accurately align 7T *ex vivo* MRI scans with corresponding histological slices. The pipeline introduces a robust, multi-modal pre-processing stage featuring adaptive 3D brain masking and 2D colorimetric segmentation to strictly isolate cerebral tissue from background noise. The core 2D-to-3D affine registration is driven by a hybrid function that balances intensity-based Mutual Information (MI) with structural gradient matching. Optimized via a multi-resolution Z-axis grid search and Powell's method, this approach successfully retrieves the arbitrary cutting plane within the 3D volume. Finally, a 2D-to-2D non-linear elastic registration, using Thin Plate Splines (TPS) and multi-resolution B-splines, is implemented to correct non-linear morphological distortions.

The pipeline was evaluated on two post-mortem MS brain datasets, demonstrating a strong ability to recover the correct 3D cutting plane and to accurately align the outer cortical contours. Furthermore, the study highlights methodological key insights regarding multimodal optimization. We can mention the standard signal homogenization techniques such as N4 bias field correction were found to degrade the gradient-driven MI metric by flattening essential structural contrast. Additionally, while B-spline elastic deformations effectively matched outer cortical silhouettes, the lack of internal regularization anchors led to artificial warping of the internal cellular architecture.

Finally, this work serves as a promising proof of concept, providing a standardized and reproducible mathematical framework for multimodal neuroimaging alignment. While the pipeline demonstrates excellent matching capabilities on high-quality datasets, it also highlights the need for enhanced algorithmic robustness to handle lower-quality or artifact-heavy scans. By bridging the gap between engineering algorithms and neuropathological realities, it establishes a solid baseline for researchers and paves the way for future improvements in histological segmentation.

Contents

Objective	1
1 Introduction	2
2 Theoretical Background	4
2.1 Multiple sclerosis (MS)	4
2.2 Physics of Ultra-High Field MRI (7 Tesla)	5
2.2.1 Gain of the Signal-to-Noise Ratio (SNR)	5
2.2.2 Increased Sensitivity to Magnetic Susceptibility and Iron Detection	6
2.2.3 Quantitative MRI and Multi-Parameter Mapping (MPM)	6
2.3 Histology	8
2.3.1 Tissue Processing and Induced Deformations	8
2.4 Blockface Imaging as a Registration Intermediary	10
3 State of the art	11
3.1 Hardware-based models	11
3.2 Computational Approaches: The Limitations of Feature-Based Methods	13
3.3 Justification for a Custom Registration Pipeline	15
3.4 Intensity-Based Registration and MI	15
3.4.1 The EMMA Formulation (Viola and Wells)	16
3.4.2 Adaptation to Medical Imaging	17
3.4.3 Application to Histology-MRI Co-registration	18
4 Materials	20
4.1 MRI dataset description	20
4.1.1 Tissue Preparation and Fixation Protocol	20
4.1.2 Scanning protocol	21
4.1.3 Data Quality Challenges and Pre-processing Requirements	22
4.2 Histology dataset description	24
4.2.1 Tissue Processing and Blockface Imaging	24
4.2.2 Microscopic slicing and Staining Protocol	24
4.2.3 Slide Digitization and Ground Truth Annotation	25
5 Method	26
5.1 MRI Pre-processing: 3D Brain Masking	26
5.2 MRI Pre-processing: N4 Bias Field Correction	30
5.2.1 Phase 1: Permissive Pre-Masking and Bias Field Estimation	31

5.2.2	Phase 2: Signal Homogenization and Precise Skull-Stripping . . .	31
5.3	MRI pre-processing summary	32
5.4	Histological Pre-processing: 2D Tissue Masking and Contrast Alignment for stained sections	32
5.5	Blockface Pre-processing: HSV Color Segmentation and Ventricle Preser- vation	35
5.6	Histological pre-processing summary	38
5.7	Azimuthal Translation	38
5.7.1	Physical Coordinate Frames and RAS Orientation	39
5.7.2	Algorithm	40
5.8	Automated Artifact Detection	41
5.9	Volumetric Registration Pipeline	42
5.9.1	Data upload, Masking, and Intensity Normalization	42
5.9.2	Spatial Sampling Strategy for Mutual Information Evaluation . .	44
5.9.3	Affine Transformation and 3D Spatial Mapping	45
5.9.4	Objective Function: Combined Metrics and Artifact Penalization	46
5.9.5	The Multi-Step Pipeline	49
5.9.6	Summary of the Decoupled Thresholding Strategy	54
5.10	Elastic corrections	54
5.10.1	Geometric Initialization via Thin Plate Splines (TPS)	55
5.10.2	Multi-Resolution Silhouette B-Spline Transformation	56
5.11	Summary of the global Co-registration pipeline	57
5.12	Code Availability	58
6	Results	59
6.1	Pre-processing Robustness	59
6.1.1	3D Brain Masking on Multi-Contrast MRI (Patient 784)	59
6.1.2	3D Brain Masking for Patient 854	62
6.1.3	2D Blockface and Histological Masking (Yellow Masks)	65
6.2	Matching assessment	68
6.2.1	Patient 784	68
6.2.2	Patient 854 with no bias correction	70
6.2.3	Patient 854 with bias correction	73
6.2.4	Elastic corrections	75
6.3	Unbounded matching	77
6.4	Stains matching	79
6.5	Summary of Coregistration Results	81
6.5.1	Pre-processing and Tissue Masking	81
6.5.2	Rigid and Affine Initialization	81
6.5.3	Bias Field Correction Impact	82
6.5.4	Geometric Initialization via TPS	82
6.5.5	Elastic Non-Rigid Deformation	82
6.5.6	Direct Histological Stain Matching	82

7	Discussion and Perspectives	83
7.1	Achievements and pipeline robustness	83
7.2	Impact of the Bias Correction on MI	85
7.3	Hyperparameters Sensitivity	86
7.3.1	Metric Weighting (α)	86
7.3.2	Sampling points	88
7.4	Limitations of the Pipeline	89
7.4.1	Elastic corrections	89
7.4.2	MRI data masking	90
7.4.3	Matching limitations	91
7.4.4	Histological data quality	91
7.5	Future work and perspectives	92
7.5.1	Machine Learning-Driven Segmentation	92
7.5.2	Feature-Based Elastic Registration	92
7.5.3	Stain adaptation of the pipeline	92
8	Conclusion	94
	AI Usage Declaration	95
	Bibliography	100

List of Figures

2.1	Multiple sclerosis illustration	4
2.2	Pulse sequence diagram of a Multi-Echo FLASH acquisition.	8
2.3	Example of a qualitative assessment of a histological slice. Green: de-myelinated WM lesions. Red: Demyelinated cortical lesions. Blue: sign of repair within cortex (sign of cortical remyelination).	10
3.1	Flowchart of the hardware-based co-registration procedure developed by Jung & <i>al.</i> [20]	12
3.2	Blockface imaging setup proposed by Wang & <i>al.</i> [18]. The front and top views illustrate the specific optical configuration utilizing Brewster’s angle (57° to the surface normal) between the light source, the tissue block, and the camera. A humidifier is applied prior to sectioning to maintain tissue integrity.	13
4.1	Quantitative MRI maps (MPM) at 300 μm resolution of Patient 784. Visual representation of the four physical parameters extracted from the 7T multi-echo FLASH sequences	23
4.2	View of raw scan of Patient 854, first echo of Proton Density (PD) contrast	24
4.3	Histological processing	25
5.1	Application of mask in R^2 contrast scan of patient 784, coronal view . .	30
5.2	Application of mask in MT contrast scan on Patient 784, coronal view .	30
5.3	Detailed block diagram of the 3D MRI pre-processing and skull-stripping pipeline.	32
5.4	Brain histological slices	34
5.5	Computed 2D tissue masks	34
5.6	Final pre-processed histological section ready for registration of Patient 784	35
5.7	Raw blockface surface photographs	37
5.8	HSV-based binary masks	38
5.9	Final processed blockface representations	38
5.10	Detailed block diagram illustrating the divergent 2D pre-processing pipelines tailored for chemically stained histology and raw blockface photographs. .	39

5.11	Automated artifact detection for rotational initialization. On the left side is depicted a slice chosen (arbitrary) and on the right side is shown the results of bar detector. The blue line highlights the brain tissue contours, the red bounding box isolates the detected reference bar, and the green line represents the computed orientation axis used to extract the rotation angle.	43
5.12	Global block diagram of the multimodal coregistration pipeline. The distinct pre-processing steps are grouped to highlight the macroscopic data flow leading to the affine and elastic registrations.	58
6.1	Preprocessing robustness comparison. The top row displays the original, uncorrected slices where the MT_{sat} contrast exhibits significantly higher background noise compared to the R_2^* . Despite these extreme initial differences, the adaptive 3D masking pipeline successfully isolates the brain tissue in both modalities (middle and bottom rows).	61
6.2	Frontal view of the bias correction pipeline results, proton-density contrast at first echo-time	63
6.3	Sagittal view of the bias correction pipeline results for Patient 854, proton-density contrast at first echo-time	63
6.4	Effect of brain mask applied on Patient 854, coronal view with no bias correction applied, slice 288, scan of first echo time	64
6.5	Effect of brain mask applied on Patient 854, coronal view (after bias correction), slice 288, scan of first echo time	65
6.6	masks for 784 histo blocks	67
6.7	Binary masks generated for the blocks of Patient 854.	67
6.8	Resulting masked MRI volumes for Patient 854 (blocks 18, 20 and 23).	68
6.9	Binary masks for the blocks 24 and 27 of Patient 854.	68
6.10	Resulting masked MRI volumes for Patient 854 (blocks 24 and 27)	68
6.11	Rigid matching results for Patient 784. For each block, the panels display from left to right: the extracted MRI slice, the corresponding histological image, and the color overlay showing the spatial alignment.	70
6.12	Patient 854: best matchings obtained before bias correction.	71
6.13	Lower quality results for Patient 854 before bias correction.	72
6.14	Matchings of Patient 854 after bias correction keeping the same parameters for the bad matchings	73
6.15	Matchings of Patient 854 after bias correction keeping the same parameters for the best matchings	74
6.16	Overview of elastic corrections applied to Patient 784.	76
6.17	Results of the 2D non-linear elastic registration pipeline applied to histological sections of Patient 784.	77
6.18	Visual assessment of the 2D elastic deformation process for Patient 854.	77
6.19	Rigid matching for Patient 784, unbounded conditions	78

6.20	Rigid matching results for Patient 784 (Block 11) using different macroscopic stains under unbounded conditions.	80
6.21	Patient 784: Rigid matching failures on HE-stained histological sections for Blocks 18 and 19, demonstrating systematic algorithmic breakdown. .	81
7.1	Visual assessment of the alpha effect on the matching pipeline of Patient 784	87
7.2	Visual assessment of the alpha effect on the matching pipeline of Patient 854	88

List of Tables

4.1	Technical specifications of the 7T ex vivo MPM protocol.	22
5.1	Chronological summary of the decoupled intensity thresholds used across the 2D-to-3D co-registration pipeline.	55
6.1	Hyperparameters of the 3D automated brain masking pipeline across MRI contrasts (Patient 784).	60
6.2	Parameters utilized for the N4ITK bias field correction algorithm.	62
6.3	Hyperparameter adaptations for 2D blockface and histological masking.	66
6.4	Optimization hyperparameters used for the affine volumetric registration (Patient 784).	69
6.5	Optimization hyperparameters used for the affine volumetric registration (Patient 854).	71
6.6	Hyperparameters of the non-linear B-Spline elastic registration pipeline.	75
7.1	Sensitivity analysis of the α hyperparameter. The table compares the final converged Z-slice and the corresponding metric score. While various configurations were evaluated, setting $\alpha = 1.0$ generally yielded a higher MI score.	87

Objective

The primary objective of this master's thesis is to design and develop an automated computational co-registration pipeline capable of aligning ultra-high-field (7T) *ex vivo* quantitative magnetic resonance imaging volumes with corresponding 2D histological slices of human brain tissue. By establishing a robust spatial correlation between macroscopic imaging and microscopic ground truth, this project aims to provide a reliable mathematical framework to validate MRI signatures of cortical pathologies in multiple sclerosis research.

Due to the dimensionality mismatch and non-linear morphological distortions introduced during tissue preparation, achieving this goal requires addressing several specific technical sub-objectives:

- **Multimodal Pre-processing:** To develop a masking technique capable of isolating cerebral tissue across different modalities.
- **Macroscopic Spatial Orientation (3D Affine Registration):** To implement a hybrid objective function balancing intensity-based Mutual Information with structural gradient matching to successfully retrieve the arbitrary 2D cutting plane within the 3D MRI volume.
- **Microscopic Morphological Alignment (2D Non-Linear Registration):** To apply, evaluate and correct elastic deformation models

Ultimately, this thesis aims to deliver a functional proof of concept that is able to translate complex image processing algorithms into practical tools for medicine research, introducing a method for fully automated registration tools in the future.

Chapter 1

Introduction

Currently, *in vivo* detection and monitoring of cortical pathology in patients with multiple sclerosis (MS) remain a challenge for conventional clinical magnetic resonance imaging (1.5T or 3T for example). Indeed, as said by Nistri & al [1], due to their small size and low contrast, cortical lesions are hard to see using conventional clinical MRI sequences (e.g., T1w, T2W or FLAIR). Some other techniques are said to have improved the visualization but those techniques remain limited in routine diagnostic. Thanks to a higher signal-to-noise ratio and increased spatial resolution, quantitative 7T MRI also allows for the highly accurate identification of small structural changes in cortical myelin, thereby providing a better understanding of disease progression.

The importance of combining MRI scans and histological data is well-known, in particular to better assess and validate what can be seen from MRI scans. In 2012, Annese & al. [2] highlighted the need to establish a correlation between MRI and microscopic techniques to overcome the limited spatial resolution of the scanners. Providing them histological slices brings data with a higher resolution, but those data cannot provide a 3D image of the brain. The 3D spatial information provided by MRI combined with the cellular and molecular features given by histology could be crucial for further research. To summarize, it has been a lot of years since this issue is known with some tools that appeared, but none of them can be directly applied in the context of Dr Lommers' study.

In her current research on multiple sclerosis (MS) and its impact on the brain, Dr Emilie Lommers aims to establish precise correspondences between microscopic histological slices and macroscopic MRI scans.

To do so, she performed quantitative MRI on the 7T MRI, available at the GIGA *in vivo* imaging platform (at the Cyclotron Research Centre) of the University of Liege, of *ex vivo* brains from MS patients.

Then, the histological slices were processed and scanned at the (*Division of Neuropathology and Neurochemistry* at the Medical University of Vienna). This multimodal approach is designed to investigate the link between the histological quantification of cortical demyelination and remyelination. To achieve this, a multiparametric MRI protocol is used, allowing for the extraction of quantitative maps, such as the effective transverse relaxation rate (R_2^*), Proton Density (PD), and Magnetization Transfer saturation (MT_{sat}). These

specific parameters provide important insights into the microstructural integrity and myelin content of the tissue.

A major issue arises because the brain undergoes important spatial misalignment between the 3D high-resolution MRI volumes and the 2D slices obtained through histological sectioning, making direct comparisons between these modalities nearly impossible.

This project aims to develop an automated co-registration pipeline that bridges the gap between *ex vivo* MRI and histological sections. By transposing microscopic ground truth onto macroscopic MRI volumes, this tool will enable Dr. Lommers to identify previously undetectable MS lesions, deepening the understanding of cortical pathology. Accomplishing this requires addressing several technical challenges inherent to multimodal alignment. We can already mention:

- **Difference in angle:** Interpolation to account for the difference in angle between the MRI slices and histological cuts.
- **Physical deformations:** Deformation that could have occurred during the slicing itself or all other manipulations made on the brain (during transport, for example).
- **Resolution and dimension:** The difference of resolution and dimension is also an issue while trying to match the data. Matching two different modalities that are also at different dimensions (2D for the histology and 3D for MRI) is not straightforward, especially if we want to keep the higher resolution on histological slices.
- **Thickness and layers:** Differences of thickness or even the number of layers will have to be taken into account for the histological blocks.
- **Artifacts and distortions:** The brain could have suffered from dehydration, paraffin embedding, and other distortions that brought more artifacts to be accounted for when the histological slices were coated by the contrast stains.

Chapter 2

Theoretical Background

2.1 Multiple sclerosis (MS)

The cortex has the ability to regenerate the myelin sheaths through a process called remyelination (RM), where oligodendrocyte progenitor cells and probably also mature oligodendrocytes produce new myelin sheaths, a repair mechanism that is more widespread in the cortical gray matter compared to the white matter [3]. However, this process tends to decrease and even disappear for the MS patients. This is illustrated on Figure 2.1

In the healthy central nervous system, oligodendrocytes are the most iron-rich cells, with iron being predominantly concentrated within their associated myelin sheaths [4]. However, in Multiple Sclerosis, this physiological iron homeostasis is disrupted. The disease progression can lead to a high loss of iron due to the destruction of myelin and the degeneration of oligodendrocytes. Concurrently, iron pathologically accumulates in specific focal structures. For instance, this accumulation occurs at the edges of demyelinated plaques, forming what is called: chronic active lesions or "iron rim" lesions. This iron can be used as biomarker during MRI scans because active/inactive MS lesions have a propensity to accumulate iron at their edge in microglia and macrophages and imaging studies suggests a lower frequency of remyelination in iron rim or active/inactive lesions, when compared with lesions without iron rims [5].




Figure 2.1: Multiple sclerosis illustration

Despite the advancements in MRI technology, magnetic resonance remains an indirect macroscopic measurement. To confirm that the variations in MRI signals (like $R2^*$ or $R1$)

truly correspond to demyelination or iron-loaded microglia, researchers must correlate the scans with microscopic cellular data. Histology remains the gold standard for this purpose. An other issue comes from the histological process involving intense chemical and physical manipulations. The brain tissue is fixed, dehydrated, embedded in paraffin, and sliced using a microtome. These steps introduce severe non-linear deformations, including tissue shrinkage, stretching, and tearing. To mitigate the loss of the 3D spatial context during slicing, modern protocols employ *Blockface Imaging*. By taking high-resolution photographs of the paraffin block's surface right before each cut often using specific lighting setups like the Brewster angle to eliminate glare and accurately capture the surface shape researchers obtain an undistorted 2D visual reference of the slice (will be detailed in section 3.1).

Conventional clinical MRI (1.5T or 3T) remains the standard imaging technique for MS diagnosis, as it effectively visualizes macroscopic lesions within the white matter. Specific sequences, such as T_2 -weighted fluid-attenuated inversion recovery (FLAIR), can partially improve the detection of cortical and leukocortical lesions [6] but conventional imaging faces fundamental limitations. Firstly, it lacks the sensitivity required to detect more discrete disease features, including purely cortical lesions, diffuse neuroinflammation, and progressive neurodegeneration. Secondly, and most important, the signal intensity changes observed on conventional scans are not pathologically specific. Because they lack sensitivity to microstructural changes, standard contrasts are insufficient to distinguish active demyelination from the remyelination process.

To overcome these diagnostic barriers, ultra-high-field 7T MRI provides an efficient alternative. By using quantitative mapping, such as T_1 and T_2^* analyses, 7T MRI has demonstrated excellent reproducibility and an increased sensitivity to cortical MS pathology compared to conventional scanners [7].

2.2 Physics of Ultra-High Field MRI (7 Tesla)

As mentioned, typical MRI scans are at 1.5T or 3T. Here, a much higher magnetic field of 7T is used, this field is used for *ex vivo* MRI, allowing for much longer scan times, and thus providing much higher sensitivity along with a much higher resolution of the obtained data.

2.2.1 Gain of the Signal-to-Noise Ratio (SNR)

The primary benefit of 7 Tesla MRI lies in the significant increase in the Signal-to-Noise Ratio (SNR). In this context, the equilibrium net magnetization (M_0) is directly proportional to the strength of the static magnetic field (B_0). Consequently, as B_0 increases, the net magnetization vector grows, leading to a proportionally larger signal induced in the

receiver coils.

This SNR gain offers several advantages compared with lower magnetic field scans, it enables to obtain an ultra-high spatial resolution of $300\ \mu\text{m}$. However, it must be mentioned that reaching this level of detail is not only the result of the ultra-high magnetic field. It also requires an ultra long data acquisition time (exceeding 12 hours) to collect the large amount of data necessary to obtain a good SNR at such microscopic scales. This type of scan is only possible for *ex vivo* imaging. However, this high resolution makes it possible to identify cortical layers and micro-lesions that remain invisible at 3T, enabling a direct correlation between ultra-high field MRI and neuropathological microscopy [8].

2.2.2 Increased Sensitivity to Magnetic Susceptibility and Iron Detection

Beyond the resolution benefits, 7T MRI also offers superior sensitivity to magnetic susceptibility. This effect is very important for Multiple Sclerosis research, as it allows for the non-invasive visualization of iron-rich structures.

Magnetic susceptibility causes local distortions of the magnetic field in the presence of paramagnetic substances such as iron. The resulting spin dephasing, characterized by the effective transverse relaxation rate (R_2^*), is directly proportional to the field strength B_0 . Consequently, iron-related contrast is highly magnified at 7T. Iron accumulations within macrophages and microglia at the borders of active MS lesions will then create an enhanced contrast. This allows for a far more accurate characterization of the lesions compared to conventional clinical systems.

2.2.3 Quantitative MRI and Multi-Parameter Mapping (MPM)

Conventional clinical MRI mostly relies on T_1 -weighted or T_2 -weighted sequences, which provide excellent qualitative contrast for routine diagnostics but are also affected by multiple tissue properties and scanner-dependent parameters. Consequently, the pixel intensities in these weighted images are not consistent between patients and this decreases the physical specificity and consistency bringing a lack of reproducibility and interpretability between patients, sites and so it makes any study more complex to continue across researchers.

In order to achieve a true method to extract quantitative data, following [9], it is necessary to use quantitative MRI (qMRI). The Multi-Parameter Mapping (MPM) protocol is a highly specialized qMRI framework designed to acquire the real physical properties of the brain's microstructure. The MPM protocol typically provides four distinct quantitative maps:

- **Longitudinal relaxation rate ($R_1 = 1/T_1$):** Sensitive to myelin content [5], but also to iron, lipid, and free water.
- **Effective transverse relaxation rate ($R_2^* = 1/T_2^*$):** Highly sensitive to local magnetic susceptibility gradients, R_2^* serves as a quantitative biomarker for mapping both paramagnetic iron and diamagnetic myelin, the presence of both components increases R_2^* values [4]. This parameter is also influenced by the strength of the main magnetic field (B_0) as well as the spatial orientation of the myelin fibers relative to this field.
- **Magnetization Transfer saturation (MT_{sat}):** Provides a measure of the macro-molecular tissue fraction, offering high specificity for myelin density [5].
- **Effective proton density (PD^*):** Reflects the tissue's free water content.

In [10], Weiskopf & *al* demonstrated the reproducibility of the results within the same site and across different sites weighted MRI scans are highly susceptible to instrumental biases: with inter-site variations that can reach up to 20% for example. In the paper, they use the Multi-Parametric Mapping (MPM) to map the four majors tissue properties mentioned here above. The results show a drop in the inter-site bias (remaining below 5%). Furthermore, the inter-site coefficient of variation (CoV) did not exceed 8% (except for R_2^* where it goes up to 20%). Also non-quantitative techniques, such as conventional T_1 -weighted imaging, yielded significantly higher variability but it's expected since it as a non-quantitative technique. It also justifies its application in our study's dataset.

However, extracting these pure physical parameters from raw multi-echo FLASH (Fast Low Angle SHot) acquisitions requires complex biophysical modelling. As demonstrated in [11], images acquired with 3D FLASH sequences are highly susceptible to signal and contrast-to-noise ratio (CNR) inhomogeneities, primarily driven by spatial variations in the radio-frequency coil configuration. At ultra-high fields (7 Tesla), these inhomogeneities in the transmit field (B_1^+) can heavily bias the actual flip angles. If left uncorrected, these hardware-induced variations shift tissue intensity distributions, blurring the boundaries between structures like gray and white matter and ultimately corrupting the quantitative maps.

To understand how quantitative maps are generated, it is important to look at the underlying acquisition sequence. As mentioned, the MPM protocol relies on 3D multi-echo FLASH sequences. This sequence is performed as follows:

After a single Radio Frequency (RF) excitation pulse (with flip angle α), the readout gradient ($G_{readout}$) rapidly alternates to acquire multiple echoes at successive Echo Times (TE) within a single Repetition Time (TR). This multi-echo sampling is fundamental for MPM, as the exponential decay of the signal across these successive echoes allows for the precise quantification of the R_2^* relaxation rate over time. By tracking the signal decay

across these multiple readouts, we can first extend the decay model to estimate the original signal at $TE=0s$ and the biophysical models can separate the true tissue relaxation (R_2^*) from hardware biases, leading to the highly reproducible quantitative maps used in this study as shown on Figure 2.2.

We can also mention that R_2^* is a key biomarker for iron accumulation in multiple sclerosis.




Figure 2.2: Pulse sequence diagram of a Multi-Echo FLASH acquisition.

To overcome any hardware-induced biases, the raw data must undergo rigorous processing. As highlighted by Wang & *al.* [11], correcting these FLASH acquisitions requires spatial registration alongside the accurate estimation of the actual B_1^+ transmit field and the receive coil sensitivity. Furthermore, to extract the pure physical parameters from these corrected signals, parameter optimization is required. In modern neuroimaging, this mathematical reconstruction is intensive but standardized using specialized software frameworks, such as the widely adopted hMRI toolbox [12], ensuring that the resulting R_1 , R_2^* , PD^* , and MT_{sat} maps genuinely reflect the underlying cellular architecture rather than scanner artifacts.

2.3 Histology

2.3.1 Tissue Processing and Induced Deformations

Although histology provides essential microscopic validation, the extensive protocol required to prepare brain tissue for microscopic examination introduces severe structural alterations. All the manipulations made from the 3D post-mortem brain to a 2D glass slide are highly invasive and alters the spatial integrity of the tissue.

Based on the specific neuropathological protocol used for this dataset, the brain tissue undergoes several chemical and physical steps:

1. **Fixation:** The whole brain undergoes immersion fixation in a buffered 4% paraformaldehyde (PFA) solution for a duration ranging from two weeks to three months. This

chemical cross-linking process stops degradation but also induces global tissue shrinkage.

2. **MRI acquisition:** Once the brain is fixed, it undergoes the *ex vivo* scan allowing to capture the whole brain structure before the cutting procedure.
3. **Macroscopic Slicing and Embedding:** The brain is first cut into 1 cm thick coronal sections. These thick sections are then progressively dehydrated using alcohols, cleared with solvents, and embedded in paraffin wax. This embedding step hardens the tissue but can also cause further tissue shrinkage.
4. **Microscopic sectioning** The hardened paraffin blocks are cut into ultra-thin, 10 μm thick whole double-hemispheric slices using a heavy-duty tetrandr microtome.
5. **Mounting and Staining:** The 10 μm slices are floated on a water bath, collected on large glass plates, and put in contact to specific stains.

The slicing and mounting stages can cause several deformations. The mechanical force of the microtome blade slicing a histological block 10 μm thick, combined with the floating process, induces non-linear deformations in the 2D slice. These include stretching, compression, tearing, and folding. Furthermore, the angle at which the manipulators made the cuts of 1 cm coronal slabs is manual and rarely aligns perfectly with reference planes used in the 3D MRI acquisition.

Several specific stainings are applied on the histological slices to highlight distinct features of multiple sclerosis:

- **Hematoxylin & Eosin (H&E):** used for the general evaluation of pathology and cellular cytoarchitecture [13].
- **Kluver-PAS:** employed for the depiction of myelin to precisely assess areas of demyelination, also referred as Luxol fast blue (LFB) [13].
- **DAB-enhanced Turnbull blue:** a specific staining optimized for cortical iron visualization.
- **GFAP immunohistochemistry (IHC):** used to highlight cortical astrogliosis [14].
- **MAP2 IHC:** applied to evaluate neuronal density [15].
- **CD68 IHC:** targets microglia and macrophages, which is crucial for visualizing chronic cortical inflammation [16].
- **PLP IHC:** used to assess intracortical myelin density, including the detailed cortical myeloarchitecture [13].
- **TPPP/p25 IHC:** applied to measure mature cortical oligodendrocyte density [16].

- **BCAS1 IHC:** a specific marker utilized to identify remyelinating oligodendrocytes and assess the remyelination process [17].

An example of an histological analysis made by Dr Lommers is shown on Figure 2.3.




Figure 2.3: Example of a qualitative assessment of a histological slice. Green: demyelinated WM lesions. Red: Demyelinated cortical lesions. Blue: sign of repair within cortex (sign of cortical remyelination).

2.4 Blockface Imaging as a Registration Intermediary

To mitigate the severe non-linear deformations introduced during histological processing, modern neuropathological pipelines introduce an intermediate photographic reference known as *blockface imaging*. This technique consists of capturing a high-resolution photograph of the paraffin-embedded tissue block surface immediately before each cut [18].

Because these images are acquired prior to any physical slicing, they inherently preserve the macroscopic geometry of the tissue [18]. Consequently, the blockface photographs serve as a reliable geometric bridge between the 3D MRI volume and the deformed 2D histological slices. This allows for the decoupling of the complex co-registration issue into two distinct, more manageable steps: a 2D-to-2D registration (to map the distorted histology back to its undistorted blockface shape) preceded by a 2D-to-3D rigid or affine registration (to locate the blockface plane within the 3D MRI volume) [19]. The practical computational implementations of this intermediary strategy, as well as their technical limitations, will be further discussed in section 3.1.

Chapter 3

State of the art

Bridging the gap between 3D MRI volumes and 2D histological slices remains a challenge due to tissue deformations. This chapter reviews the state of the art in multimodal co-registration, categorizing existing strategies into physical hardware-based setups and computational software pipelines.

3.1 Hardware-based models

Different studies previously attempted to co-register MRI scans and histology. We can mention one from Jung & al[20], where they provide a hardware-based method for an alignment of the images from both method in several steps: They acquired first an *in vivo* scan of the brain, followed by an *ex vivo* one in order to create a rotation matrix after registration between the scans. Using this rotation matrix the brain was placed in an agarose block and trimmed to physically match its original *in vivo* orientation. Subsequently, a second *ex vivo* scan was performed using a localizer to plan 4 mm-thick slices. Finally, the histological slices were cut at 4mm intervals to ensure that their orientations remained aligned with the initial *in vivo* scan as illustrated in Figure 3.1.

However, this physical alignment method presents some practical limitations. First, it requires a preliminary *in vivo* scan, and also to place the brain in the agarose block. All those steps are often very complex or even impossible to implement .

An other issue that emerges from this method is that it does not accounts for all the tissue deformations that occurred during the earlier steps of the process. Indeed, following the standard pipeline used in our research, brain tissues could already have undergone a lot of deformations proving their fragility. Using the physical method presented by Jung & al[20]: these subsequent elastic deformations are left completely unmeasured and uncorrected.

A more recent and computationally advanced approach was proposed by Wang & al. [18] that was published last year. They enhanced the standard blockface photography to develop a technique named **Brewster Blockface Quantification (BBQ)**. Their method begins with an *ex vivo* MRI scan (with a 7T MRI), after which the brain is embedded in paraffin and sliced. The major upgrade it provides is that they use what is called *Brewster angle* (setup from Wang & al. shown on Figure 3.2). It is used to capture the surface




Figure 3.1: Flowchart of the hardware-based co-registration procedure developed by Jung & al.[20]

shape before cutting. This specific optical configuration helps noticing deformations that could occur during the slicing.

Following the slicing, a first manual co-registration is performed with re-orientation using ITK-SNAP to provide a coarse alignment. The precise, automated registration is then executed using the Tensor Image Registration Library (TIRL) [19]. TIRL applies sophisticated non-linear transformations to correct for the more complex tissue deformations that have occurred during the manipulations, mapping 2D histological pixel coordinates into the 3D physical voxel space of the MRI. demonstrated excellent accuracy, achieving spatial overlaps with Dice similarity coefficients frequently exceeding 90%.

This technique once again is very efficient but needs a prior alignment of the data using the optical hardware not used here. Since the Brewsters blockface quantification was not performed on the present data, this matching methods is neither not applicable.

An other trial to solve this alignment issue was to address it since the MRI scan itself. For instance, Absinta & al [21] developed a method using individualized, 3D-printed cutting boxes based on post-mortem MRI surface renderings to physically guide the brain slicing process along specific MRI planes. This approach is however hardware-based and must be executed before cutting the histological slices to improves the initial macroscopic alignment. It presents several limitations: it requires custom 3D printing for every single brain




Figure 3.2: Blockface imaging setup proposed by Wang & *al.* [18]. The front and top views illustrate the specific optical configuration utilizing Brewster’s angle (57° to the surface normal) between the light source, the tissue block, and the camera. A humidifier is applied prior to sectioning to maintain tissue integrity.

specimen, and most importantly, it does not account for the subsequent microscopic tissue deformations such as shrinkage, stretching, and tearing that inevitably occur later during chemical fixation and paraffin embedding.

Consequently, even if this approach could resolve the difference of angles issue, it will not take all the deformation and dimensional issues that still must be taken into account. This still need computational co-registration, relying on multi-modal registration algorithms and intermediate blockface imaging as will be developed in this thesis, in order to achieve precise, microscopic MRI-histology alignment.

3.2 Computational Approaches: The Limitations of Feature-Based Methods

Since physical and prospective alignment methods are often impractical for standard tissue processing, the registration often relies primarily on computational algorithms. Usually, most of these solutions have depended on feature-based registration techniques.

An alternative family of registration methods relies on the detection and matching of local image features rather than on global intensity statistics. The most prominent example is the **Scale-Invariant Feature Transform** (SIFT), initially introduced by Lowe [22] and later refined by him in [23]. SIFT extracts distinctive keypoints from an image such as corners, edges and particular locations, describes them with a local gradient histogram that is invariant to image scale, rotation, and partially invariant to illumination changes and affine distortions [23]. Correspondences between two images are then established by matching these descriptors using a nearest-neighbour strategy, providing a set of point pairs from which a geometric transformation can be estimated.

Significant efforts have been made to adapt these techniques for medical neuroimaging. For instance, Rister & *al.* [24] successfully extended the SIFT algorithm to three dimensions for volumetric medical image registration. Their 3D-SIFT implementation proved highly effective for aligning 3D intra-modal volumes (e.g., matching brain MRI scans to existing anatomical atlases using affine transformations), and was even validated on simulated MR images featuring multiple sclerosis lesions [24].

The issue is that their algorithms fundamentally struggle with extreme multi-modal challenges, such as matching a 7T MRI slice to a blockface photograph or a stained histological slice. The contrast mechanisms generating these images are entirely distinct, as the intensities and key features that are usually essential are now no more, leading to several critical failures:

- **Intensity inversion and absence of shared gradients:** a prominent edge in a quantitative MRI map may not produce a corresponding geometric edge in a standard blockface photograph depending on the contrast of the MRI (R_1 or R_2^* for MRI provide opposite contrast for example). Furthermore, the gradients might appear completely different, inverted or even absent depending on the specific histological stain used (e.g., PLP for myelin vs DAB for iron).
- **Background inconsistencies and artefactual keypoints:** feature extractors are easily confused by non-anatomical background noise. Even if they may look the same, backgrounds around *ex vivo* MRI, blockface images or stainings slices are very different (in colours, intensity, composition,...). Algorithms like SIFT often detect thousands of "keypoints" in those backgrounds that are totally irrelevant for matching, leading to several misalignments.
- **Microscopic deformations:** the tissue tearing and stretching induced by the manipulations physically alter the local geometry. Consequently, a keypoint's surrounding gradient neighbourhood in the deformed histology slice no longer matches its spatial neighbourhood in the undistorted MRI.

Because homologous anatomical landmarks cannot be reliably extracted across such distinct modalities, feature-based methods are inadequate to be used directly for precise MRI-histology co-registration. This necessitates the use of alternatives that may have similarities to SIFT but that will be much more consistent across samples and are also less specific to only one modality (MRI in the case of the SIFT) but also to match images coming from different modalities.

3.3 Justification for a Custom Registration Pipeline

All those tools already seem to be working well, however they are primarily optimized for standard 3D-to-3D registration. Applying these generalized tools directly to our specific dataset would give poor results and co-registration failures, which necessitates to develop an original and new pipeline that would take all the parameters we have that are more conventional but also less well prepared to fit into their tools.

The requirement for a custom algorithm arises from some computational challenges:

1. **The 2D-to-3D Registration issue:** As mentioned, standard algorithms are designed for data matching of corresponding dimensions. Registering a single 2D histological slice into a whole 3D MRI volume requires a more intensive data processing capable of iteratively extracting 2D planes defined by arbitrary angles.
2. **Avoidance of outliers:** optimization algorithms used in standard software are highly susceptible to local minima if the initial spatial alignment is poor. Because the cutting angle of the histological block is arbitrary and unknown, a standard optimizer will not be able to match slices with such a large relative angle between them. Our custom pipeline will be able to take this angle into account and then find the global best matching across the angular space before initiating the fine-tuning optimization algorithm.
3. **Background Artifacts and Strict Masking:** Algorithms provided by the different authors attempt to register all image intensities, including the background. In our multi-modal dataset, the background comprises completely unrelated materials that bring some noise to the image and tends to falsify the true intensity distribution across the corresponding slices. Developing a custom script allowed for the rigorous implementation of brain masks, forcing the Mutual Information metric to be evaluated exclusively on the brain tissue (different masks must be created depending on the nature of the data available).

Consequently, developing a programme framework using Python is essential to achieve the precision required to match ultra-high field MRI signals with histopathology.

3.4 Intensity-Based Registration and Mutual Information

An alternative to the feature-based matching exploits a totally different approach: instead of matching geometric structures based on their intensities and shapes, these methods measure the statistical dependency between the intensity distributions of the two images. In other words, the algorithm aims to predict the intensities around pixels/voxels of an image A based on the intensities around the corresponding pixel/voxel in the image B. To achieve this, the registration process uses **Mutual Information** (MI), a highly robust

metric that will be the basis of this multimodal image alignment.

This approach was first introduced by Viola & al that used it for a broad range of applications (visages, car, all sort of objects) [25]. It demonstrated that this matching technique has a wide range of possibilities but, once again, all the images compared had the same modality as well as the same general characteristics. However comparing two brains must differentiate finer and more precise fluctuations.

Shortly afterwards, Maes & al. [26] specifically adapted and validated the metric even for multi-modality. This comes from the fact that no assumptions are made about the nature of the intensity correspondence, making MI robust to the modalities that makes the feature-based technique like SIFT fail.

3.4.1 The EMMA Formulation (Viola and Wells)

The first concrete and efficient formulation of MI-based alignment for computer vision was proposed by Viola and Wells [25] under the name **EMMA** (EMpirical entropy Manipulation and Analysis). Here, MI treats the image intensities themselves as continuous random variables. To estimate the underlying probability density $p(z)$ of a specific intensity value z , EMMA uses the *Parzen Window* method: superimposing Gaussian densities G_ψ over the intensity values z_j of a randomly drawn sample set of pixels A :

$$p(z) \approx \frac{1}{N_A} \sum_{z_j \in A} G_\psi(z - z_j)$$

The differential entropy $h(z) = - \int p(z) \ln p(z) dz$ is then approximated using the sample average over a second independent sample B :

$$h(z) \approx -\frac{1}{N_B} \sum_{z_i \in B} \ln \left(\frac{1}{N_A} \sum_{z_j \in A} G_\psi(z_i - z_j) \right)$$

This continuous and differentiable formulation allows the analytical derivative of the entropy with respect to the spatial transformation to be computed, driving alignment via stochastic gradient. A key practical contribution of Viola and Wells [25] is the use of stochastic sampling: in fact rather than evaluating MI over all image points, a small random subset is drawn at each step (arbitrary number of samples taken). This reduces computational cost while preserving convergence. Moreover, Viola & al have found that issues that may arise from taking a whole dataset because it may cause local minima to appear. But the stochastic sampling overcome this issue, the noise introduced by sampling can penetrate small local minima [25].

3.4.2 Adaptation to Medical Imaging

EMMA was proven to be working well on general problem but Maes & *al.* [26] specifically adapted and validated the MI paradigm for the complex intensity distribution of neuroimaging. Rather than using continuous Parzen windows, they proposed a robust discrete formulation based on the joint image intensity histogram.

To compute the Mutual Information, we first bin the intensity pairs (f, r) of the floating image F and the reference image R into a 2D joint histogram, denoted as $h_\alpha(f, r)$, which depends on the transformation α . From this histogram, we estimate the joint and marginal probability distributions by normalizing $h_\alpha(f, r)$ as follows:

$$\begin{aligned} p_{FR}(f, r) &= \frac{h_\alpha(f, r)}{\sum_{f,r} h_\alpha(f, r)} \\ p_F(f) &= \sum_r p_{FR}(f, r) \\ p_R(r) &= \sum_f p_{FR}(f, r) \end{aligned} \quad (3.1)$$

The intensities of both images are first linearly rescaled to a predefined range ($[0, n_F - 1]$ and $[0, n_R - 1]$), resulting in an $n_F \times n_R$ joint histogram. The resolution used by [26] is 256×256 bins.

Using this discrete framework, Maes & *al.* expressed the MI criterion through the Kullback-Leibler divergence form:

$$I(\alpha) = \sum_{f,r} p_{FR}(f, r) \log_2 \frac{p_{FR}(f, r)}{p_F(f) p_R(r)}$$

Kullback-Leibler divergence measures the statistical distance between the joint probability distribution of the two images (p_{FR}) and the distribution that would be expected if the two images were statistically independent ($p_F \times p_R$). By maximizing this divergence, the statistical independence between the images is minimized, forcing the registration algorithm to maximize the information shared between the MRI and the histological reference. A registration is considered optimal when statistical consistency between the aligned images is maximized. Mathematically, the optimal transformation α^* is found by determining the parameters that maximize this Mutual Information score:

$$\alpha^* = \arg \max_{\alpha} I(\alpha) \quad (3.2)$$

An issue that arise comes when a transformation α results in sample points that do not coincide with the discrete grid points of the reference image. In other words, standard interpolation methods are insensitive to translations smaller than one voxel and create discrete jumps in the joint histogram that introduce noise into the optimization process.

It can also introduce new intensity values because it calculates the intensity at a non-grid position as a weighted average of the intensities of the eight nearest neighbour voxels (i.e., smoothing). If the transformation makes a pixel travel to a position surrounded by various intensity values, its intensity will then be changed.

To ensure that the histogram varies smoothly with the transformation α , they introduce **trilinear partial volume (PV) interpolation**. Instead of assigning a sample's contribution to a single histogram bin, the algorithm now distributes the weight of a single intensity pair across the entries of the eight nearest neighbors in the reference grid, using the same weights as trilinear interpolation. This approach ensures that after the transformation α , the pixel's intensity is now spread on their neighbours (the closest have the largest weight), considerably improving the robustness of the optimizers. Validated on rigid-body CT/MR and PET/MR registration, this approach achieved sub-voxel accuracy, completely automatically and without any prior segmentation.

Finally, in both the continuous and discrete formulations, no assumptions are made about the nature of the intensity correspondence between modalities. This makes MI much more robust to the contrast inversions and background artifacts than feature-based methods like SIFT.

3.4.3 Application to Histology-MRI Co-registration

The application of MI to the specific problem of histology-MRI registration was demonstrated by Meyer & *al.* [27], who registered a single H&E-stained histological slide to an *in vivo* 7T MRI volume of a rat brain. Their pipeline used blockface photography as an intermediate registration step and relied on MI maximization for the alignment, a strategy that is directly linked to the approach developed in the present work.

More recently, Aquino & *al.* [28] applied a similar MI-based slice-to-volume framework to *ex vivo* 7T MRI and histological slices from epilepsy patients. Their work is particularly relevant here, as it shares the same data modalities and 2D-to-3D registration and in the previous papers, the slice-to-volume registrations (2D to 3D matching) were not discussed. However, their method employs a 2D correlation coefficient as a similarity metric, which requires high contrast consistency. As they say, "being MRI planes and histological slicing already in geometrical correspondence by means of the previously mentioned cutting-box." meaning that their workflow is working on data already pre-matched. However designing and manufacturing a custom 3D-printed box for each brain is time-consuming and expensive, making this approach highly impractical for large number of data.

Despite its robustness, MI-based optimization has also some limitations. As a global metric evaluated over sampled image points, it can be highly sensitive to the initial estimation and may converge to local optima when the starting point of the analysis is far from the global maximum. In this context, we can mention that if the initial estimation is made in the peripheral background regions, those regions will be compared first and quickly reach a local maximum if they are similar. This limitation has a huge impact on the matching

results by making the two images correspond based on irrelevant information from each one.

Chapter 4

Materials

Before detailing the specific imaging and histological protocols, it is important to add some context about the origin of the anatomical specimens used in this methodological project. The data used for the development and validation of the co-registration algorithms originate from researches investigating the multi-scale pathology of Multiple Sclerosis (MS). The post-mortem brain specimens were carefully extracted and provided by the Division of Neuropathology and Neurochemistry at the Medical University of Vienna, Austria. For the scope of this thesis, a subset consisting of two post-mortem MS brains (designated as Patient 784 and Patient 854) was given. The complete experimental pipeline applied to these specimens involves two physical phases conducted across different places: an initial ultra-high field *ex vivo* MRI acquisition performed at the Cyclotron Research Centre (CRC) in Liège, followed by an the histological sectioning and digitization carried out subsequently in Vienna.

The following sections detail the technical protocols, data structures, and specific image quality challenges associated with each of these two modalities.

4.1 MRI dataset description

As already mentioned, instead of relying on conventional weighted imaging (T_1 -w or T_2 -w), the acquisition protocol utilized a Multi-Parameter Mapping (MPM) framework [20].

4.1.1 Tissue Preparation and Fixation Protocol

The preparation of the post-mortem brain followed a protocol designed to stabilize the tissue while optimizing the signal properties during the whole scanning steps. Firstly, both brains underwent a fixation in a buffered 4% paraformaldehyde (PFA) solution in phosphate-buffered saline (PBS) at a pH of 7.4. The fixation duration was controlled, ranging from two weeks to three months.

To ensure the highest possible image quality, a washing step was implemented: 72 hours prior to the MRI scan, the brains were taken from the fixative to be placed into a pure PBS solution. This procedure serves to remove free formaldehyde monomers from the tissue.

By clearing these monomers, it allows to obtain improved image contrast and scans quality, ensuring that the quantitative maps reflect the real biological structure without having any chemicals residuals that would impact the local environments around them, decreasing the overall quantitative information.

Finally, to prepare the specimens for the long *ex vivo* MRI acquisitions, the brains were submerged in an immersion fluid within a custom 3D-printed plastic container. To prevent any microscopic tissue movement during the 12-hour scanning protocol and to serve as an internal geometric reference, a plastic stabilization rod (referred to as the "mounting bar") was inserted alongside the tissue in the first dataset (Patient 784).

4.1.2 Scanning protocol

The magnetic resonance imaging data used in this project was acquired *ex vivo* using the 7 Tesla MRI scanner (Magnetom Terra, Siemens Healthineers) at the Cyclotron Research Centre (CRC), University of Liege. The combination of an *ex vivo* scanning protocol which completely eliminates patients motion artifacts and the use of a highly sensitive 32-channel radio-frequency head coil enables the achievement of the ultra-high isotropic spatial resolution of $300\ \mu\text{m}$. Each post-mortem brain being scanned for a total of at least 12 hours to achieve an optimal Signal-to-Noise Ratio (SNR). This level of detail is very important for identifying MS lesions and ensure a precise spatial match with the histological slices.

The Multi-Parameter Mapping (MPM) framework, implemented using a 3D multi-echo FLASH sequence, involves three distinct acquisitions with varying flip angles (α) to generate different tissue weightings:

- **PD-weighted and MT-weighted:** Acquired with a flip angle of 11° .
- **T1-weighted:** Acquired with a flip angle of 59° .

As illustrated in Figure 2.2, each repetition cycle (TR) of 50 ms to 60 ms included the acquisition of 6 to 8 equidistant echoes. For the 7T post-mortem protocol, the echo times (TE) ranged from 3.66 ms to 22.56 ms with a constant increment of 2.7 ms. This multi-echo sampling allows for the precise estimation of the R_2^* decay and the extrapolation of the signal amplitude to $TE = 0$, which is necessary for removing hardware-related biases from the final quantitative maps. The specific parameters for the 3D multi-echo FLASH sequences are detailed in Table 4.1.

Following the acquisition, the raw 3D MRI data were exported and converted into the NIfTI (Neuroimaging Informatics Technology Initiative) format (.nii). Unlike standard image formats, this specific volumetric format is designed to retain both the 3D discrete grid of intensity values and the spatial metadata (including resolution, orientation among all), most notably the affine transformation matrix. This structural property ensures that the physical dimensions (the $300\ \mu\text{m}$ isotropic resolution) and the spatial orientation of the

Table 4.1: Technical specifications of the 7T ex vivo MPM protocol.

Parameter	Value
Isotropic Resolution	0.3 mm ($300\mu m$)
Repetition Time (TR)	50–60 ms
Echo Times (TE)	8 echoes (3.66 to 22.56 ms)
Flip Angles (α)	11° (PDw/MTw), 59° (T1w)
B1+ Mapping	Bloch-Siebert Shift method

brain scans are mathematically preserved. In other words, this 3D format provides a robust, spatially-aware basis for all subsequent computational pre-processing and co-registration steps.

4.1.3 Data Quality Challenges and Pre-processing Requirements

Despite the theoretical advantages of 7T ultra-high field imaging, the acquired datasets exhibited significant practical limitations. The raw scans from both Patient 784 and Patient 854 were characterized by substantial noise levels and hardware-induced artifacts, primarily stemming from B_1^+ transmit inhomogeneities and B_1^- receive coil sensitivity biases. These high-level noise and contrast gradients posed a major challenge for the registration pipeline: standard intensity-based algorithms were frequently biased by irrelevant background noise and artefactual high-intensity values, precluding convergence. This necessitated the implementation of a rigorous, custom-designed pre-processing workflow, including brain masking and outlier removal, to isolate the anatomical signal.

A visual inspection of the multi-parameter maps for Patient 784 (Figure 4.1) highlights these challenges. Despite the expected SNR gain, the raw volumes exhibit a highly granular texture and high noise levels, particularly pronounced in the MT_{sat} map. Furthermore, aliasing artifacts affect all quantitative maps, appearing as geometric distortions in the bottom part of the scans (indicated by red arrows in Figure 4.1a, Figure 4.1b, and Figure 4.1d). When looking at the superior aspect of the brain, the tissue extends outside the Field of View (FOV), causing this aliasing effect likely a consequence of the acceleration factors used during acquisition.

Patient 854 presents a different challenge: massive spatial intensity inhomogeneities (bias field). As seen in Figure 4.2, the outer brain edges exhibit hyper-intensities, while the deeper structures appear significantly darker. From a physical perspective, these contrast gradients are explained by the complex radio-frequency (RF) behaviour at 7 Tesla. As demonstrated by Van de Moortele *et al.* [29], these patterns are not solely due to "dielectric resonances" but are primarily driven by pronounced spatial phase patterns and destructive interferences within the tissue. This phenomenon creates non-uniform distributions for both the transmit (B_1^+) and receive (B_1^-) fields. The superposition of these

fields leads to signal cancellation in some regions (darkening the deeper structures) and relative enhancement at the edges, as observed in Patient 854.

The combination of background noise, intensity bias, and geometric aliasing significantly complicates the optimization of the registration pipeline. The matching metric is frequently biased, as these artifacts create erroneous local maxima that trap the optimizer. Consequently, these data quality issues rigorously justify the implementation of the custom-designed pre-processing workflow including brain masking and field bias correction to isolate the relevant anatomical signal before attempting any registration with the histological slices.




Figure 4.1: Quantitative MRI maps (MPM) at $300\ \mu\text{m}$ resolution of Patient 784. Visual representation of the four physical parameters extracted from the 7T multi-echo FLASH sequences




Figure 4.2: View of raw scan of Patient 854, first echo of Proton Density (PD) contrast

4.2 Histology dataset description

Following the *ex vivo* MRI acquisitions, the brain specimens were transferred to the Division of Neuropathology and Neurochemistry at the Medical University of Vienna, Austria, for the following histological processing.

4.2.1 Tissue Processing and Blockface Imaging

The fixed whole brains were initially cut into 1 cm thick coronal slabs and subsequently embedded in paraffin. The paraffin allows for preservation of the undistorted macroscopic geometry of the tissue before any microscopic slicing occurred, intermediate blockface images were systematically acquired for the specific embedded tissue blocks (three different blocks in the dataset of Patient 784 and five for Patient 854). These geometric references included high-resolution scans digitized at 1200 dpi, complemented by macroscopic photographs. Those high-resolution blocks scans will be used during the matching part.

4.2.2 Microscopic slicing and Staining Protocol

The paraffin-embedded blocks were sliced using a large Reichert tetrandr microtome, a technique specifically chosen to allow for the extraction of whole double-hemispheric slices. These large tissue slices were then mounted on 15 cm × 12 cm glass plates. To accommodate the specific requirements of various brain markers (see section 2.3), the slice thicknesses were precisely adapted:

- **10 μ m slices:** Prepared for standard structural pathological evaluations. These whole double-hemispheric slices were stained with Haematoxylin and Eosin (HE) for general cellular cytoarchitecture, Klüver-Barrera/Periodic Acid-Schiff (KPAS) to assess demyelination, and Toluidine Blue (TBB). Other markers typically evaluated at this thickness include DAB-enhanced Turnbull blue for cortical iron visualization,

GFAP for astrogliosis, MAP2 for neuronal density, and CD68 for visualizing chronic cortical inflammation.

- **6 μm slices:** Prepared for specialized immunohistochemical (IHC) labelling requiring thinner slices. These included Proteolipid Protein (PLP) staining to assess intracortical myelin density and cortical myeloarchitecture, alongside BCAS-MBP markers to identify remyelinating oligodendrocytes and evaluate the remyelination process.


4.2.3 Slide Digitization and Ground Truth Annotation

Following the staining and mounting procedures, the histological sections were fully digitized to obtain high-resolution images. The visualization and processing of these images were performed using QuPath, a software specifically designed for digital pathology analysis. This tool allows to easily highlight the important regions of interest (ROIs) and export their spatial coordinates, providing the essential microscopic "ground truth" required to validate our subsequent MRI-to-histology co-registration pipeline.

Figure 4.3 shows an example of a block image and a corresponding staining. Figure 4.3a shows the blockface photograph of brain tissue before slicing it into the different stainings glasses, preserving a true spatial geometry. While Figure 4.3b shows the corresponding 10 μm histological slice stained with Haematoxylin and Eosin (HE). It is visible that the tissue macroscopic morphology and contrast are not the same. Some deformations including tissue shrinkage, macroscopic tearing, and separation of the hemispheres, can also be seen (for the bottom part and the ventricles especially) in the mounted slice when compared to the undistorted blockface reference. This illustrates the importance to make a blockface scan as an intermediate step to achieve accurate MRI-to-histology co-registration.



(a) Blockface photograph



(b) Histological slide (HE stain)

Figure 4.3: Histological processing

Chapter 5

Method

The co-registration of *ex vivo* 3D MRI volumes with 2D histological sections presents significant geometric, photometric, and computational challenges. To address these issues and ensure a robust multimodal spatial alignment, the methodology developed in this thesis is structured around a sequential pipeline divided into three primary pillars:

1. **Data Pre-processing and Standardization:** Before any matching can occur, both modalities must be cleaned to isolate the relevant anatomical signal from severe background noise, mounting artifacts, and contrast variations. This foundational step includes 3D skull-stripping and bias field correction for the MRI volumes, alongside colorimetric segmentation and photometric inversion for the 2D histological and blockface sections.
2. **Volumetric Affine Registration (2D-to-3D):** This phase aims to mathematically bridge the geometric gap between the modalities by finding the exact spatial location and orientation of the histological cut within the intact MRI volume. It relies on a custom multi-step optimization pipeline guided by a hybrid Mutual Information metric—to extract the optimal virtual oblique MRI slice.
3. **Non-Rigid Elastic Correction (2D-to-2D):** Once the global cutting plane is identified and extracted, this final phase compensates for the localized physical distortions (e.g., tissue shrinkage, mounting deformations) induced by the histological sectioning process. It uses Thin Plate Splines (TPS) and multi-resolution B-Splines to achieve a highly precise, pixel-to-voxel topological alignment.

By decoupling the global volumetric search from the local morphological corrections, this architecture ensures both computational efficiency and high spatial accuracy. The following sections detail the mathematical and algorithmic implementation of each phase, beginning with the critical preparation of the raw MRI datasets.

5.1 MRI Pre-processing: 3D Brain Masking

As previously illustrated in Figure 4.1, the raw 7T *ex vivo* MRI volumes contain a non-negligible quantity of unwanted signals. These include the brain container, the immersion fluid (PBS or Fomblin), and significant background noise that completely skew the

quantitative maps. If left uncorrected, these intensity artifacts would severely bias the co-registration metric during the co-registration process, as the algorithm would be tempted to use those falsified information as a baseline of matching.

To resolve this, a custom 3D automated brain masking (skull-stripping) algorithm was developed and implemented in Python using the `SciPy` and `scikit-image` libraries. The pipeline operates directly on the 3D NIFTI volumes through the following sequential steps:

1. **Volumetric Noise Reduction (3D Median Filtering):** The first step of the skull-stripping algorithm involves the application of a 3D median filter to mitigate the high-frequency spatial noise. Unlike standard 2D image processing which operates on a flat 3×3 pixel grid (evaluating a square of 9 pixels), our approach is made to work on a 3D analysis.

For a given target voxel, the algorithm evaluates a cube of $3 \times 3 \times 3$ voxels. The filter extracts the intensity values of all 27 voxels within this local neighborhood and sorts them in increasing order. The original voxel (the center of the cube) takes the median value of this ranking (the value at the center: the 14th position).

This operation is particularly advantageous for brain masking: it is highly effective for the "salt-and-pepper" noise (extreme outlier intensities) that we can find in the backgrounds of our images while the brain structures, since they are wider than 3 voxels, will not undergo a strict filter modification. Furthermore, the kernel size was chosen in order to have the less impact possible on each quantitative map. A cube of $3 \times 3 \times 3$ voxels kernel was sufficient for the relatively clean R_2^* and PD maps, whereas a larger $5 \times 5 \times 5$ voxels cubic kernel (evaluating a neighbourhood of 125 voxels) was required to treat the background of the MT_{sat} acquisitions since this was a very noisy one.

2. **Adaptive Bounding Box Identification:** Following the initial noise reduction, the algorithm performs a coarse foreground-to-background separation aiming to isolate the brain itself compared to the background. An initial global mean intensity is calculated across the entire filtered volume. A binary mask is then generated by thresholding the data (the mean is the threshold). Voxels with intensities exceeding the global mean are labelled as potential tissue.

The algorithm identifies the spatial coordinates of all non-zero voxels in this initial mask to determine the edges of the brain in the 3D coordinates system. These coordinates define a 3D bounding box that will encapsulate the brain as tightly as possible (will then also include its immediate surroundings as security) while discarding the vast majority of the empty background. This step is useful because the subsequent region of interest defined will ensure that further thresholding calculations are not biased by the millions of zero-intensity voxels in the surrounding empty space.

3. **Adaptive Threshold Calculation:** Once the scan is restricted to the brain area, the algorithm computes an updated mean intensity only within this region of interest.

This local average is significantly more representative of the specimen's brain intensity as it is no longer skewed by the surrounding empty space.

To determine the optimal threshold for brain extraction, the algorithm identifies all voxels within the bounding box that possess intensities lower than the new mean (mean of the region of interest). These voxels typically represent non-cerebral elements within the brain tissue such as the immersion fluid, the sulci or the plastic reference rod used to stabilize the specimen (explained in subsection 4.1.1). A supplementary mean of the intensities of this subset of darker points within the region of interest is then used as the final adaptive threshold that will serve as the final brain mask [30]. This mask is subsequently generated by retaining only the voxels from the filtered volume that exceed this last threshold, isolating the brain tissue from its immediate non-anatomical surroundings.

4. **Morphological Cleaning and Structural Isolation:** To decouple the brain tissue from adjacent non-anatomical structures or artifacts, a supplementary 3D binary opening operation is applied. Using a 3D structuring element with a connectivity of 1, the algorithm performs 6 iterations of erosion followed by 6 iterations of dilation. This process aims to break the small pieces retained that should not be while preserving the large brain structures [30].

Subsequently, the volumetric size of each independent 3D object are calculated, the algorithm identifies and retains only the two largest contiguous structures. In the dataset of Patient 784, these primary components correspond to the two sides of the brain around the large plastic bar, while all smaller structures (including pieces of the plastic bar that have hyper-intensities) are left out. This ensures that the final mask is strictly limited to the major anatomical regions of interest. For patient 854, the brain itself can be retained as 1 large structures but there may be some tissue that have been detached from the cortex during the slicing so it ensures that those structures are also kept.

5. **Hole Filling and Final Tissue Interchapter:** Following the isolation of the cerebral hemispheres, a final 3D binary closing operation (10 iterations) is applied to the mask. This consists of a dilatation followed by an erosion. This step ensures a filling of internal dark cavities within the mask, such as the ventricles or macroscopic tissue tears induced by the fixation process, ensuring a continuous and solid anatomical envelope.

Once the optimal mask is generated, it is retained as a boolean format and applied directly to the raw, uncorrected MRI volume.

This operation sets all non-anatomical regions (including the surrounding fluid, the background noise, and the 3D-printed container) to an absolute intensity of zero, while conserving the original signal of the cerebral tissue. This skull-stripping step is the basis for the co-registration as all those non-zero voxels outside the brain would have participate in

the matching and influence the comparison voxel to voxel, this would have as consequence to false the results. Consequently, the subsequent co-registration optimization algorithm will evaluate the alignment strictly based on the actual brain tissue rather than matching artefactual background structures.

Application of the 3D brain masking pipeline on the R2 contrast of the same slice as Figure 4.1b. On Figure 5.1a, it is represented the mask itself generated after adaptive intensity thresholding and 3D morphological operations. On Figure 5.1b, the final skull-stripped R_2^* volume resulting from the element-wise multiplication of the raw data (previously shown in Figure 4.1) and the binary mask. The cerebral tissue is still not perfectly isolated, letting custom 3D-printed specimen container still visible, this is explained because this structure is a very large one and mostly because it is attached directly to the brain and so counted as part of it. The erosion will not be large enough to remove it afterwards. However, the background is now perfectly removed as the large holes within the brain, allowing the data to be good enough to go through the matching.

Figure 5.2 shows the application of the 3D brain masking pipeline on an MT_{sat} map. Figure 5.2a shows the binary mask generated after slight adaptations on the intensity thresholding and 3D morphological operations. As the MT_{sat} had a much larger background (see Figure 4.1c, a larger $5 \times 5 \times 5$ voxels median filter was used in the initial step and the final mean of the voxels to retain was also a bit raised (by a 1.1 factor)). Figure 5.2b represents the final skull-stripped MT_{sat} volume resulting from the element-wise multiplication of the raw data and the binary mask. The custom 3D-printed container is here nearly totally rejected but a bit of high-frequency noise is still present in the bottom of the image. This is due once again to the noise voxels directly attached to the brain structure itself.




Figure 5.1: Application of mask in R_2^* contrast scan of patient 784, coronal view




Figure 5.2: Application of mask in MT contrast scan on Patient 784, coronal view

5.2 MRI Pre-processing: N4 Bias Field Correction

As previously identified in subsection 4.1.3, ultra-high field 7T MRI acquisitions particularly those from the second dataset (Patient 854) suffer from severe spatial intensity inhomogeneities, commonly referred to as B_1^+ bias fields. These low-frequency gradients cause extreme hyper-intensities on the superficial structures and artificial darkening in the deeper ones, which would inevitably disrupt the co-registration metric during matching. To

homogenize the tissue intensities, an automated N4 Bias Field Correction algorithm was implemented using the *SimpleITK* framework[31].

Because accurate bias field estimation requires isolating the tissue from the background, and precise masking relies on homogeneous intensities, a custom pipeline was implemented (mainly inspired by the N4ITK algorithm of [31] with the code explained in [32]).

5.2.1 Phase 1: Permissive Pre-Masking and Bias Field Estimation

Initially, a highly permissive mask is generated using a minimal intensity threshold (10% of the global mean in this case). This first step aims to broadly discard the empty background space without attempting to extract a definitive brain envelope. However since the scans have a very high resolution, applying the 3D median filter along with the iterative morphological operations (such as 3D openings and closings) across millions of voxels is computationally heavy.

To optimize execution time, the strategy of sub-sampling the data is used. Both the raw MRI volume and the permissive pre-mask are downsampled using a shrink factor of 4 across all spatial dimensions. Because the bias field is inherently a smooth, low-frequency phenomenon, this subsampling preserves most of the spatial profile while drastically accelerating the computation efficiency. The B-spline routine of the N4 algorithm is then executed on the subsampled volume across four hierarchical resolution levels. This multi-resolution pipeline progressively refines the B-spline control grid, going from a coarse mesh intended to capture the global macroscopic B_1^+ gradient, to a thinner mesh able to localize field variations without confounding them with the real tissue voxels. The optimizer is constrained to a maximum of 50 iterations per level and a strict convergence threshold of 0.001 (if minus than 0.1% of change between two iterations before the 50 maximum iterations are reached, the algorithm stops).

5.2.2 Phase 2: Signal Homogenization and Precise Skull-Stripping

Once the optimal low-frequency bias field is estimated, it is upsampled back to the original discrete voxel grid. The full-resolution raw MRI volume is then divided by this bias field calculated, effectively flattening the contrast gradient. This operation successfully recovers the anatomical signal.

Finally, this uniformly contrasted volume is given to the automated 3D masking pipeline detailed previously. Because the signal is now (relatively more) homogeneous, the adaptive tissue mean is much more stable, allowing for precise structural tissue preservation while reducing the risk of tissue exclusion from darker zones.

5.3 MRI pre-processing summary

To visually synthesize the different steps involved in the preparation of the MRI data, the complete 3D MRI pre-processing workflow is illustrated in Figure 5.3. The diagram overall illustrates how median filtering, adaptive thresholding, and morphological operations are sequentially applied to eliminate background artifacts and isolate the cerebral tissue.




Figure 5.3: Detailed block diagram of the 3D MRI pre-processing and skull-stripping pipeline.

5.4 Histological Pre-processing: 2D Tissue Masking and Contrast Alignment for stained sections

Just as the 3D MRI volumes required skull-stripping to eliminate background noise, the 2D histological sections and stainings must be rigorously pre-processed too. The bright background of the glass slide, mounting media, and microscopic debris could also lead to

a big drop of the Mutual Information metric during 2D-to-3D registration.

To address this, the previously established volumetric masking algorithm (in section 5.1) was adapted into a flexible 2D pipeline (`mask2D.py`). While the core statistical logic relying on a 3×3 pixel median filter and an adaptive bounding box remains identical to the MRI processing pipeline, some modifications had to be implemented to ensure a good mask of the different stainings:

- **Background removing:** Unlike MRI where the background is dark, histological slices have a much more luminous background, giving a background with intensity above the stain intensity. To account for this, the adaptive thresholding logic isolates the bright background and remove the brain tissue. A local mean is computed within the bounding box and the initial mask captures the background. The final boolean tissue mask is therefore obtained by logically inverting the result, ensuring the tissue is assigned a value of 1 and the glass slide a value of 0.
- **Smaller structures retention:** The histological processes (slicing, apply the staining in warm water,...) can easily induce structure's alterations, such as micro-tears, separation of tissue, or fragmented slice borders. To prevent the accidental deletion of valid anatomical structures by the masking algorithm, the morphological operations were significantly scaled down compared to the 3D pipeline (where the largest structures were kept with a high threshold for their sizes). The binary opening was reduced to a single iteration, and the closing to 3 iterations. Furthermore, instead of strictly retaining only the two largest structures, the algorithm preserves any tissue fragment whose volume exceeds 10% of the largest detected object. This ensures that detached but valid cortical fragments are preserved for registration while smaller structures that would not bring useful information anymore are removed.
- **Intensity Contrast Inversion:** To ensure a successful matching computation, the histological slice must mimic the contrast profile of an MRI (bright tissue on a dark background). While the initial masking step successfully forces the bright glass slide to become an absolute black background (intensity = 0), the chemically stained tissue inherently retains its native low-intensity profile. To correct this, an intensity inversion is applied ($I_{\text{inverted}} = 255 - I_{\text{original}}$) so the dark tissue becomes bright. However, applying this mathematical operation globally would inadvertently turn the newly black background back to bright white ($255 - 0 = 255$). To prevent this side effect, the mathematically inverted image is multiplied one final time by the binary tissue mask. This guarantees that only the brain tissue is inverted to a bright contrast, while the surrounding background is forced to remain strictly black.

This pre-processing step outputs a cleaned histological image where the tissue appears bright against an absolute black background. By successfully mimicking the macroscopic contrast gradient of the MRI slices, this geometric and photometric standardization provides an optimal starting point for the non-linear 2D-to-3D matching algorithms.

Figure 5.4 shows Myelin-stained (PLP) and cell-stained (HE) slices exhibiting the characteristic bright background and some tissue differences with the MRI scans as explained. On Figure 5.5 we can see that the algorithm successfully isolates the anatomical tissue from the background. Unfortunately for Figure 5.5a, it also shows the algorithm limitations: the brighter zones of the stain are also masked out while for Figure 5.5b, the ventricles are retained, they can be a good anchor for the matching process to compare both modalities so this lack of masking is not ideal. Following the application of the binary masks, the backgrounds are forced to absolute zero. The final slices are shown on Figure 5.6.




Figure 5.4: Brain histological slices




Figure 5.5: Computed 2D tissue masks




Figure 5.6: Final pre-processed histological section ready for registration of Patient 784

5.5 Blockface Pre-processing: HSV Color Segmentation and Ventricle Preservation

To bridge the geometric gap between the intact 3D MRI and the 2D histological slices, which encounter physical degradation as explained, Blockface photographs were acquired alongside digital scans of the corresponding tissue blocks. These high-resolution images, captured just after the slicing, retain the true macroscopic geometry of the brain before it undergoes the severe mechanical stresses of mounting on glass slides, making it a potential good ally to match the 2 modalities. However, the raw blockface images contain significant background noise, primarily consisting of the surrounding embedding medium and also some entrapped air pockets (due to the paraffin for example).

Unlike the gray-scale intensity thresholding used for MRI or digitized slides, the blockface images present a totally different colorimetric challenge. The embedding medium and background do not share exactly the same colour with the cerebral tissue but in the gray-scale they look very identical making a classical mask as presented in section 5.1 and section 5.4 inefficient because the yellow nuances of the tissue and the white/gray colours cannot easily be differentiated. To accurately isolate the tissue structures, a custom segmentation pipeline based on the Hue, Saturation, Value (HSV) colour space was developed trying to effectively isolate the yellow/brown colours of the brain and separate it from the white/gray background.

The HSV colour space decouples pure colour information from local lighting conditions. This decoupling offers to our method a much more consistent and precise algorithm of the brain colour identification, which exhibits varying illumination but a consistent yellow-based color. The segmentation relies on the independent thresholding of three distinct

channels:

- **Hue (H):** This channel isolates the pure chromatic tone. As the paraffin or resin embedding medium typically presents a yellowish-brown hue, the boundaries were explicitly constrained to $H \in [10, 80]$. This ensures the algorithm exclusively targets these colour tones while disregarding other tones.
- **Saturation (S):** Representing colour purity or intensity, this channel acts as the primary discriminant between the anatomical tissue and the rest that must be masked out. While the cerebral tissue exhibit strong and distinctive yellow tones, the rest of the image (background, artifact) are much less colourful and closer to gray for most of them. By enforcing a minimum saturation threshold ($S \geq 130$), the algorithm effectively ignores the pale background elements and exclusively isolates the vivid and saturated brain tissues.
- **Value (V):** This channel encodes absolute luminosity. To account for shadows that could appear on the pictures, a broad luminosity range was permitted ($V \in [30, 255]$). The lower bound ($V_{\min} = 30$) prevents the misclassification of the low-intensity background, while the maximum upper bound accommodates highly reflective entrapped air pockets that appear hyper-intense in the pictures.

The pipeline follows some distinct steps that are for some very similar to those of the previous masks explained:


- **HSV Color Thresholding:** The original BGR images were converted into the HSV color space. The Empirical thresholds were defined ($H \in [10, 80]$, $S \in [130, 255]$, $V \in [30, 255]$) to create an initial mask targeting the non-tissue background.
- **Morphological Smoothing:** Due to the extremely high spatial resolution of the blockface photographs (exceeding dozen of Mb), noise artifacts are numerous (given the number of pixels, some artifacts can be imperceptible to the naked eye). Elliptical structuring elements with size of 15×15 pixels were initially used for sequential morphological opening and closing operations to smooth the mask boundaries. The elliptical shape prevent the boundaries from developing artificial, square-like edges while the size (3 times larger that for the MRI mask) is needed to catch sufficient noise (the 5 pixel-wide element compared to the total amount of pixels would be mathematically insignificant and would go through relevant spatial noise). Subsequently, an algorithm was implemented to retain exclusively the single largest contiguous object, effectively discarding external artifacts and disconnected debris.
- **Ventricular Detection (Hole Drilling):** The most difficult part to implement for this pipeline was the handling of internal anatomical cavities, such as the ventricles and deep sulci. Standard morphological operations could fill all internal voids, resulting in an inaccurate, solid silhouette. To preserve the true internal topology, a hierarchical contour detection algorithm was implemented. The algorithm establishes a two-level hierarchy that distinguishes between external

boundaries and internal cavities. Rather than a global filling approach, a selective drilling logic was implemented: the script iterates through each detected internal contour and calculates its area relative to the total brain surface. A dynamic threshold of 10^{-6} of the total area was empirically established to differentiate between noise and legitimate anatomical structures. Only those internal voids exceeding this threshold are masked out. This ensures that the final mask preserves essential landmarks like the lateral ventricles and the inter-hemispheric fissure, which are very important landmarks used for the matching.


A final 21×21 pixels morphological opening ensures smooth and continuous anatomical boundaries. By applying this refined mask, the output is a true-colour blockface representation of the tissue over an absolute black background, providing an indispensable ground-truth reference for the 3D volumetric reconstruction.

The images on Figure 5.7 capture the brain immediately after slicing. The background exhibits non-uniform tone with sometimes colours similar to the brain (edges of Figure 5.7a). The mask is presented on Figure 5.8, the successful preservation of the internal topology can be mentioned: the algorithm identifies and excludes the ventricular cavities while maintaining smooth and continuous anatomical boundaries through elliptical morphological operations.

The final masking results, illustrated in Figure 5.9, reveal a mixed performance regarding the internal topology. While the large ventricular cavities are successfully excluded from the tissue mask, the narrower cortical sulci are not entirely removed. Furthermore, the algorithm inadvertently excludes the brown tissue lesion visible in Figure 5.7a. Those defects remain small and the results are satisfying enough to go through the matching algorithm and obtain good results.




(a) Raw blockface photograph (Block 11)




(b) Raw blockface photograph (Block 18)

Figure 5.7: Raw blockface surface photographs




(a) Block 11 Binary Mask



(b) Block 18 Binary Mask

Figure 5.8: HSV-based binary masks



(a) Processed Block 11



(b) Processed Block 18

Figure 5.9: Final processed blockface representations

5.6 Histological pre-processing summary

In order to summarize the specific preparation steps required for the histological data, the corresponding algorithmic architecture is presented in Figure 5.10. This block diagram shows how the pre-processing pipeline splits into two distinct paths depending on the type of input image. It distinguishes the approach applied to chemically stained histological slices relying on gray-scale thresholding and contrast inversion from the adaptive HSV colorimetric segmentation and selective ventricular drilling tailored for raw blockface sections.

5.7 Azimuthal Translation

To extract the corresponding 2D planes from the 3D MRI volume, a framework is required. This section details the mathematical and geometric steps involved in this azimuthal translation, starting with the definition of the reference coordinate system.




Figure 5.10: Detailed block diagram illustrating the divergent 2D pre-processing pipelines tailored for chemically stained histology and raw blockface photographs.

5.7.1 Physical Coordinate Frames and RAS Orientation

An important aspect of this plane extraction is the strict separation between the image voxel space (the 3D data array) and the physical coordinate frame. Inspired by standard neuroimaging co-registration frameworks such as MNE-Python, all spatial transformations are strictly computed in the physical MRI coordinate space (expressed in millimeters) rather than in voxel indices. This allows a more intuitive way to represent the data.

By utilizing the NIfTI header's affine transformation matrix (*Affine*), the mathematical origin is projected into the standard **RAS** (Right, Anterior, Superior) continuous coordinate system, this is done before the matching process itself. In this physical space:

- The X-axis corresponds to the Left-to-Right anatomical axis.
- The Y-axis corresponds to the Posterior-to-Anterior anatomical axis.
- The Z-axis corresponds to the Inferior-to-Superior anatomical axis.

This mapping ensures that the azimuthal angles (α and ϕ) carry direct consistent and practical anatomical significance. Furthermore, computing the oblique plane's coordinates in the continuous physical space before mapping them back to the discrete voxel space via

the inverse affine matrix guarantees that the extraction process is entirely independent of the respective resolutions that vary within the modality and between the modalities. The corresponding physical plane will be extracted no matter the initial voxel resolution of the input MRI volume.

5.7.2 Algorithm

To perform the 2D-to-3D registration, the algorithm must iteratively extract virtual 2D slices from the 3D *ex vivo* MRI volume. During the first step of the optimization loop, the position of the cutting plane is governed by an 8-parameter affine transformation matrix encompassing Cartesian translations (t_x, t_y, t_z), Euler rotations (R_x, R_y, R_z), and 2D scaling (S_x, S_y) that can zoom in or zoom out on the slices.

Being mathematically efficient, all those 8 parameters are not instinctive neither easy to manipulate for all small transformations someone would want to make. To establish a robust, reproducible, and human-readable definition of the final oblique plane, an Azimuthal Coordinate System was implemented, all the affine transformations can be reduced to a 3 parameter space that is much more instinctive to handle.

Once the co-registration optimization converges to a match, the resulting slice is not aligned with the initial planes making it very difficult to see using the initial raw scans. The exact position and inclinations are transmitted from the final Cartesian matrix into three independent spherical parameters relative to the volumetric centre:

1. **Radius (r):** The orthogonal distance (in millimetres) from the from the volumetric centre to the plane.
2. **Azimuth (α):** The horizontal rotation angle (in degrees) of the plane's normal vector.
3. **Elevation (ϕ):** The vertical rotation angle (in degrees) of the plane's normal vector.

To orient the cutting plane, the input elevation angle (ϕ) is first converted to a standard spherical colatitude ($\theta = \frac{\pi}{2} - \phi$). The final spatial orientation is then achieved by computing a composite rotation matrix ($R = R_z(\alpha) \times R_y(\theta)$), which sequentially rotates the plane around the Z-axis (azimuth) and the Y-axis (elevation). This matrix is then multiplied by the unit vector $(0, 0, 1)$ to transmit the orientation to the spatial coordinates.

By extracting the normal vector (n_x, n_y, n_z) from the final Euler rotation matrix, the azimuthal angles are computed using trigonometric projections ($\alpha = \arctan(n_y/n_x)$ and $\phi = \arcsin(n_z)$), those operations are carried out during the final steps of the matching algorithm. This algorithm provides a standardized mathematical framework to extract the precise location and orientation of histological cuts and, by noting the parameter's value, allows for the direct, arbitrary re-extraction of identical virtual planes without needing to

go through all the previous and following steps which is time consuming.

5.8 Automated Artifact Detection

Iterative optimization algorithms, such as those used in this multimodal co-registration, have shown a very high sensitivity to their initial parameters during the trials. A poor initialization can lead the algorithm to converge into a local optimum, resulting in large spatial misalignment and thus mismatches. To mitigate this, a fast pre-alignment step was introduced to provide a robust initial orientation estimation.

Our *ex vivo* MRI scan from Patient 784 contains a massive, linear signal voids (a plastic bar) resulting from the physical mounting apparatus used during the preparation of the scanning process. Because these artifacts maintain a consistent geometric relationship with the brain tissue, their orientation serves as an excellent ground for exclusion since this device is not present in the whole histological slices. A detection pipeline was implemented to isolate this artifact to serve as additional parameter during the matching process:

- **Restricted Area:** To prevent the misidentification of external background noise, the search is strictly confined within the brain tissue. The previously computed binary MRI mask (see section 4.1) is using some computer vision functions of the OpenCV library, surrounded by a bounding polygon that goes all around the brain. This imaginary line is made around the outermost extremities of the brain ensuring that all the points are located within this zone. The zone is eroded using a 15×15 pixels elliptical kernel in order to have a strict confinement of the brain in the zone to create a safe, internal Region of Interest (ROI) slightly offset from the cortical edges.
- **Intensity and Morphological Filtering:** The physical bars manifest as absolute signal voids within a higher intensity zone. An adaptive intensity threshold is applied, isolating pixels within the lowest 8th percentile of the tissue's grayscale distribution (chosen empirically). Small anatomical voids (like micro-fissures) are then filtered out using sequential morphological closing and opening operations (as in section 5.4).
- **Angle Extraction:** A *Connected Component Analysis* extracts the contours of the remaining dark regions. A geometric filter is applied to retain only highly elongated objects with a length-to-width ratio > 3.0 and a minimum length of 24 mm (corresponding to 80 pixels at 0.3 mm isotropic resolution), effectively filtering out residual sulci.

Finally, a robust linear regression is applied to the spatial coordinates of the valid contour points (`cv2.fitLine`). By using the L2-norm (Ordinary Least Squares method), the algorithm minimizes the sum of squared Euclidean distances between the scatter of artifact pixels (pixels detected as belonging to the barre) and the fitted line. This mathematical optimization effectively extracts the true principal ge-

ometric axis of the artifact, inherently smoothing out any morphological irregularities along its edges that would not be relevant. The directional vector (V_x, V_y) of this optimized line is then extracted to compute the artifact's absolute orientation angle ($\arctan(V_y/V_x)$). The resulting directional vector yields the principal angle of the artifact ($\arctan(V_y/V_x)$).

By automatically extracting this angle prior to the co-registration pipeline, the algorithm is provided with an adaptative rotational initialization that provide information about the rotation that can not be applied for the matching. This drastically shrinks the required angular search space, preventing convergence failure and it also effectively accelerates the overall matching pipeline by exclusion of parameters to analyse.

On Figure 5.11, the *left figure* represents the raw *ex vivo* MRI slice ($Z=388$, chosen arbitrary) exhibiting a massive linear signal void caused by the plastic bar.

On the *right side*, a visual output of the algorithm. The blue polygon delineates the eroded convex hull, establishing the restricted internal Region of Interest (ROI). The red contour highlights the successfully segmented artifact following intensity thresholding and morphological filtering. The green line represents the true principal axis computed via Ordinary Least Squares regression (L2 distance), yielding a highly accurate initial rotation angle (74.6°) for the subsequent matching pipeline. It is important to note that the artifact detection illustrated at slice $Z = 388$ serves strictly as a proof of concept. In practice, the exact spatial orientation of the mounting apparatus may vary along the longitudinal axis. Therefore, applying the angular initialisation from one extreme of the volume (e.g., $Z = 388$) to another (e.g., $Z = 200$) would be geometrically inaccurate. To address the slice to look at, the detection algorithm is implemented dynamically within the matching pipeline: it is executed only after a preliminary coarse spatial alignment has identified the approximate target slice, providing an angular initialization for that specific Z -coordinate prior to the fine optimization.

5.9 Volumetric Registration Pipeline

This chapter is the core of the 2D-to-3D registration algorithms and is handled by an iterative optimization pipeline designed to find the exact spatial configuration of a virtual MRI slice that maximizes its structural similarity with the 2D histological reference.

5.9.1 Data upload, Masking, and Intensity Normalization

Before initiating the optimization loop, both the 3D MRI volume and the 2D target image must be loaded into a standardized computational space. As detailed in subsection 5.7.1, the 3D NIfTI volumes (both the raw data and the binary mask) are first reoriented to the RAS coordinate system. To reduce the registration metrics to the biological tissue and eliminate background noise, the matching algorithm directly uses the masked MRI volume,




Figure 5.11: Automated artifact detection for rotational initialization. On the left side is depicted a slice chosen (arbitrary) and on the right side is shown the results of bar detector. The blue line highlights the brain tissue contours, the red bounding box isolates the detected reference bar, and the green line represents the computed orientation axis used to extract the rotation angle.

where all non-brain voxels have already been masked out.

Simultaneously, the 2D histological or blockface image is imported and put on a gray-scale matrix. A systematic horizontal reflection can be applied if necessary (since the histological sections can be reversed compared to the MRI slices). In this study, this reflection was systematically applied to the 2D image for both patients to correct for the mirroring effect induced by the scanning protocols. Manually aligning the chiral space of both modalities is computationally required, as standard gradient-based optimizers can theoretically inverse itself the image but it can easily bring issues and make it during the affine transformations is too late to have a good matching before.

While the Mutual Information (MI) metric is theoretically robust to non-linear intensity mappings between different modalities. In practice, quantitative MRI volumes often exhibit large floating-point ranges, whereas histological images are constrained to a discrete 8-bit integer space ($[0, 255]$). Large differences of the intensities value can bring more incertitude and be more computationally intensive since the matching should take more values into account (each of the 256 values of histology can correspond to many values of MRI).

To address this, an intensity normalization (specific to each modality) is applied to both modalities prior to the registration loop. The algorithm first isolates the relevant biological tissue by applying an intensity threshold to both modalities, excluding the previously masked

background ($I > 0$ for the MRI volume, and $I > 15$ for the histological slice). The statistical mean (μ) and standard deviation (σ) are then computed for the tissue pixels. Subsequently, a Z-score standardization is applied, followed by a linear rescaling to shift the distributions into a gray-scale space, centred at 128 with a standard deviation of 40:

$$I_{\text{normalized}} = \left(\frac{I_{\text{raw}} - \mu_{\text{tissue}}}{\sigma_{\text{tissue}}} \right) \times 40 + 128 \quad (5.1)$$

Finally, the resulting matrices are restricted to the $[0, 255]$ integer range. This procedure mathematically aligns the contrast profiles of both the MRI and the 2D slice, ensuring that the background remains perfectly black while the tissue intensities variations are optimally distributed across the histogram bins for the co-registration.

5.9.2 Spatial Sampling Strategy for Mutual Information Evaluation

As initially proposed by Viola and Wells in their EMMA framework [25], computing the Mutual Information (MI) metric over every single voxel of high-resolution datasets is computationally too intensive (since it makes millions of voxels to analyse). To ensure practical execution times, the function will process a reduced number of points.

The limitations of stochastic sampling

The standard implementation relied on a stochastic sampling strategy, where a specific number of points (a few thousand or tens of thousands in our case) are randomly drawn without replacement from the entire 2D image grid. While this random sampling effectively penetrates local minima in general computer vision tasks [25], it presents two major drawbacks for our specific multi-modal dataset:

1. **Background evaluation waste:** A random selection across the entire rectangular grid blindly picks points within the empty background. Since our strict masking algorithms set the background to zero, evaluating MI on these empty regions wastes samples that will not bring useful information. This can also bring bias to the results since the backgrounds will match with each other and increases the matching score.
2. **Stochastic noise:** Because a new random subset is drawn at every step of the optimization, the MI score exhibits slight statistical fluctuations even if the transformation matrix remains identical. This inherent noise can cause early convergence or erratic gradient estimations.

Deterministic and tissue-restricted sampling

To overcome these limitations and stabilize the co-registration pipeline, a custom deterministic and targeted sampling algorithm was implemented. Instead of picking a large number

of random points, the algorithm extracts a smaller set of structural points exclusively from the relevant tissue.

The procedure is executed as follows:

1. **Tissue Isolation:** The algorithm first identifies all the coordinates (X, Y) within the 2D reference slice that possess an intensity strictly greater than the background threshold ($I > 15$). This creates an array containing exclusively valid tissue pixels.
2. **Sampling step calculation:** To ensure that the sampled points are uniformly distributed across the entire brain structure rather than clustered in a specific region, a constant sampling step is calculated. This step is defined as the integer division of the total number of available tissue pixels by the target number of points ($\text{Step} = N_{\text{total}}/N_{\text{target}}$).
3. **Sequential Extraction:** The final subset is generated by iterating through the valid tissue array using the computed step, yielding exactly the desired number of relevant points. Finally, these coordinates are mapped into matrices to allow for rapid linear algebra multiplications during the latter affine transformations.

This technique allows a much quicker computing of the algorithm and, furthermore, the absence of random variations ensures that the matching score remains perfectly stable for a given transformation, effectively eliminating stochastic noise. This algorithmic refinement improved the convergence speed and the spatial accuracy of the iterative optimization algorithms.

5.9.3 Affine Transformation and 3D Spatial Mapping

Before the Mutual Information can be evaluated, the 2D histological coordinates must be projected into the 3D space of the MRI volume. This spatial mapping is governed by an iterative optimizer that, at each iteration, updates a set of transformation parameters to find the best alignment possible.

At each iteration, the algorithm constructs a 4×4 pixel homogeneous affine matrix. To ensure a geometrically logical transformation, the operation follows this mathematical sequence:

1. **Centering (T_1):** The 2D slice coordinates are first translated so that the geometric origin $(0, 0)$ aligns with the centre of the image. This ensures that potential rotations occur around the anatomical center rather than the corner of the image.
2. **Scaling and Shearing (S and Sh):** A 3D scaling matrix adjusts the dimensions of the slice to account for tissue shrinkage or stretching, while a shear matrix compensates for non-rigid slicing deformations.
3. **Rotation (R):** Euler rotations (R_x, R_y, R_z) are applied to orient the plane in the

3D space. It will be useful especially to find the cut angle of the histological slices relative to the MRI planes.

4. **Translation (T_2):** Finally, the plane is translated (t_x, t_y, t_z) to its target location within the MRI volume to assure the best overlay of the two slices possible.

The final composite matrix ($M = T_2 \cdot R \cdot Sh \cdot S \cdot T_1$) is multiplied by the homogeneous coordinates of the sampled points. These projected 3D coordinates are then passed to an interpolation function, which extracts the corresponding continuous MRI intensities at these precise voxels locations to evaluate the objective function and the new matching score.

5.9.4 Objective Function: Combined Metrics and Artifact Penalization

Once the spatial coordinates of the 2D slice are transformed and mapped into the 3D MRI volume, the algorithm must evaluate the quality of the alignment.

Standard Intensity-Based Mutual Information

Unlike simpler metrics such as Sum of Squared Differences (SSD), which assume a linear relationship between image intensities, **Mutual Information** (MI) assumes no prior relationship. Instead, it measures the statistical dependence between the intensities of the two images, making it the best method for multimodal registration where tissue contrasts drastically differ (e.g., dark MRI tissue versus bright stained histological tissue).

As described in Section 3.4.2, the algorithm computes the Mutual Information (MI) by extracting the joint and marginal probability distributions from the overlapping tissue pixels. This quantifies the spatial alignment between the interpolated MRI slice and the histological section. The algorithm seeks to find the spatial transformation that maximizes this MI score, meaning it searches for the alignment where the intensities of one image best predict the intensities of the other, minimizing their joint entropy.

While standard Mutual Information (MI) is highly effective, relying solely on raw intensities can lead to misalignments due to the specific artifacts present in the *ex vivo* datasets. For instance, the histological slices can suffer from mechanical tears or missing tissue for example. Conversely, the MRI volume have the plastic bar or can also suffer from artifacts. To prevent the Mutual Information metric from maximizing over these corresponding signal issues, a custom and multi-feature objective function was developed.

Anti-Artifact Filtering and Overlap Penalization

An important issue arises when the optimization algorithm maps valid histological tissue into an MRI signal void (such as the background or the dark plastic mounting bar). Because these voids present a perfectly uniform intensity (near zero), standard MI can artificially spike, guiding the optimizer into a local maximum where it believes the alignment is highly consistent (i.e., mapping all tissue to black).

To prevent this, a filter was implemented within the MI evaluation loop:

1. **Histological Tissue Baseline:** The algorithm isolates the valid tissue coordinates on the 2D histology slice using a threshold ($I_{\text{histo}} > 5$) lower than the one used for the mask application. While a strict threshold was previously required to ensure that only pure tissue pixels influenced the statistical mean, this lower threshold is applied during the registration loop to retain all non-zero pixels and maximize the alignment accuracy at the structural boundaries.
2. **MRI Signal Validation:** For these exact coordinates, the algorithm checks the corresponding MRI intensities. A voxel is only considered as a potential valid overlap if its intensity has a sufficient signal intensity ($I_{\text{MRI}} > 25$) which represents the 1% lower intensities of a normal distribution centered on 128.

If a transformation pushes a significant portion of the slice into a void, the algorithm immediately penalizes the objective function. The raw MI score is multiplied by an *Overlap Ratio*, defined as the number of valid overlapping pixels divided by the total number of histological tissue pixels. This severe penalty forces the optimizer to actively seek regions where both modalities present actual physical tissue. Furthermore, the severity of this penalty is dynamically adapted based on the inherent quality of the dataset. For high-quality acquisitions featuring distinct and crisp anatomical boundaries (such as Patient 784, which includes the mounting bar), a squared penalty (Ratio^2) is enforced. This aggressive constraint heavily punishes partial overlaps, forcing a strict adherence to the mask boundaries. In the other hand, for datasets with lower signal-to-noise ratios and blurrier mask boundaries (such as Patient 854), a linear penalty (Ratio^1) is applied. This more permissive approach prevents the optimizer from failing due to minor, inevitable masking uncertainties, allowing it to still prioritize intensity and gradient matching.

Gradient-Based Mutual Information

Standard intensity-based MI relies exclusively on the voxels intensity to build its joint probability distribution. By matching absolute voxel intensities, it effectively correlates regions of uniform signal. However, this approach lacks spatial awareness; the optimizer can successfully align large homogeneous tissue masses while completely failing to align their actual borders due to the sharp changes of intensity that could be a few pixels/voxels away.

To inject structural context into the information-theoretic framework, a Gradient-Based Mutual Information metric was integrated. Rather than computing the joint probability distribution over raw intensities which blindly correlates regional intensity mappings, this approach forces the algorithm to evaluate the topology of the tissues. In other words, it shifts the focus from matching uniform regions to explicitly aligning structural contours, ensuring that a sharp boundary in the MRI volume directly corresponds to a sharp morphological edge in the histological reference.

Mathematically, the algorithm achieves this by computing the spatial gradient magnitude for both the interpolated MRI slice and the histological reference:

$$G_F = |\nabla F| \quad \text{and} \quad G_R = |\nabla R| \quad (5.2)$$

Because numerical differentiation amplifies high-frequency noise (a difference of one value on the gray-scale is not relevant for the algorithm but derivating them can give a much larger value) and expands the range of values possible (calculating a gradient can give negative values), the resulting gradient arrays are independently rescaled to a discrete space $[0, 255]$ via a min-max normalization (G_{norm}). This step is necessary to ensure the numerical stability and density of joint histogram estimation.

Finally, the Mutual Information is evaluated over these normalized gradients. By shifting the mutual information calculation to the derivative domain, the objective function provides a very precise tool to compare the boundaries of each modality. While translating one slice for a few pixels, standard MI score remains relatively stable but the gradient score drops precipitously since the edges will not be aligned anymore.

The Hybrid Objective Function

To take advantage of both the global contrast mapping and the local edge alignment, the final objective function minimized by the optimizer is a weighted linear combination of both metrics:

$$\text{Metric} = \alpha \cdot MI_{\text{intensity}} + (1 - \alpha) \cdot MI_{\text{gradient}} \quad (5.3)$$

The hyperparameter α dictates the balance between the two features. By combining regional statistics with structural derivatives, with this parameter offers a dynamic way to calculate the mutual information. This hybrid objective function creates a highly stable and discriminant optimization landscape, effectively guiding the algorithm toward the true physical alignment.

Optimization Vector and Cost Function

To autonomously navigate the optimization landscape and find the optimal spatial alignment, the iterative solver (e.g., Powell's method) requires a scalar objective function to

minimize. For this registration phase, an 8-parameter optimization vector was defined. The 8 parameters that will be dynamically updated by the function are:

- **Translations:** t_x, t_y, t_z
- **Rotations:** R_x, R_y, R_z
- **In-plane Scaling:** S_x, S_y

Although a full 3D affine matrix would require 12 parameters, scaling along the Z-axis (S_z) has no physical meaning when iteratively analysing 2D slice and at this point the shearing deformations (Sh_x, Sh_y, Sh_z) are not taken into account and set to 0. Therefore, the function completes the parameter array by locking the Z-scale to 1.0 and setting the three shear parameters to 0.0.

The 8 parameters (translations, rotations, scaling) are sent with their value to the function that will create the matrix of transformation. This matrix will multiply the temporary chosen slice to update it. Finally, the function returns the combined metric score (using mutual information and gradient mutual information) to update the new slice score.

At each step of the iterative process, the cost function performs three operations:

1. **Matrix Assembly:** The 8 active parameters and 4 fixed parameters are aggregated to construct the 4×4 pixel homogeneous affine transformation matrix (M).
2. **Coordinate Projection:** By multiplying the MRI slice coordinates by the affine matrix ($P_F = M \cdot PR_{\text{homo}}$), the function put a new cut orientation to the 3D MRI space. This defines the slice whose spatial locations intensities must be interpolated to be compared with the static histological reference.
3. **Score Inversion:** Standard numerical optimizers (calling the function with *scipy.optimize.minimize*) are designed to find the global minimum of a function. Because we want to maximize a score, the objective function computes the hybrid metric score (using the weighting factor α) and returns its exact negative value ($-\text{Metric}$).

5.9.5 The Multi-Step Pipeline

Spatial Initialization and Raw Data Retrieval

To ensure the iterative optimization algorithms converge toward the true anatomical alignment rather than falling into local minima, a spatial initialization has been set. Starting the search from the geometric origin (0,0,0) would first place the virtual plane in the empty background, causing the Mutual Information gradient to be near 0 which could lead to local maxima of the mutual information function.

To prevent this, the algorithm computes automatically the volumetric center of the brain. This is achieved by isolating the valid tissue within the masked MRI volume using a strict intensity threshold (retaining only voxels exceeding 30% of the mean brain intensity, the

rest being the background). The spatial coordinates (X, Y, Z) of these tissue voxels are averaged to extract the true centre of the brain. These coordinates serve as the initial translation parameters (t_x, t_y, t_z) for the affine matrix, guiding the matching process immediately at the core of the brain.

The pipeline also computes the initial rotational parameters using the automated artifact detection detailed in section 5.8. Because the working MRI volume has already been masked (with the background and the plastic mounting bar set to an absolute intensity of zero), to optimize memory allocation and computational efficiency, the algorithm exploits the geometry of the masked data itself instead of reloading the raw scan. By generating a convex hull around the outer boundaries of the 2D brain slice, the masked plastic bar manifests as a massive, linear signal void strictly contained within the anatomical envelope. The artifact detection function directly analyzes this pre-masked volume to extract the orientation of this geometrical void. By aligning the initial rotational parameters with the bar's orientation, the algorithm ensures that the matching process operates in a search space that actively avoids this artifact. Even if the histological slice was not sectioned exactly along this plane, initializing the rotation parallel to the mounting apparatus guarantees that the virtual plane does not intersect the signal void, providing a robust starting point.

It is important to note that the presence of the mounting apparatus is not present across all acquired datasets. For *ex vivo* scans that does not have this artifact, the automated detection function is bypassed. In such cases, the algorithm uses a predefined set of empirical initial angles (chosen by the users) to perform a broader grid search during the first scanning phase. This logic ensures that the code also begins the matching with a good first approximation with or without bar.

Dimensional Initialization and Scale Factor

Before applying any iterative spatial transformations, the fundamental dimensional differences between the two modalities must be addressed. The 3D MRI volume is defined in a macroscopic voxel grid, whereas the digitized histological sections consist of high-resolution pixel matrices. To establish a consistent initial spatial correspondence, a global baseline scale factor is computed. This factor is geometrically defined as the ratio between the transversal MRI width (in voxels) and the corresponding histological slice width (in pixels). Integrating this baseline scaling parameter normalizes the two modalities into a shared dimensional magnitude. This initialization is a computational prerequisite, ensuring that the subsequent affine optimizer starts from a geometrically plausible configuration rather than attempting to match spaces with orders-of-magnitude differences in scale.

Z-Axis Scanning and Grid Search

Searching the perfect slice of a high-resolution 3D MRI volume with the perfect parameters in a single iteration is computationally very demanding but can also stop at a local maximum. To narrow the search space and be more computationally efficient, the pipeline first executes a discrete, coarse grid search to identify the most promising volumetric regions before committing to heavy computations.

The algorithm iterates along the longitudinal Z-axis (depth) within a predefined anatomical range (a few dozen of slice identified as potential best match by the user). At each Z-coordinate, the algorithm calls the automated artifact detection function. As previously established, if the mounting apparatus is identified on that specific slice, the rotational search space (R_x) is immediately restricted to the computed artifact angle found. This dynamic parameter pruning significantly reduces the computational weight. If the artifact is absent or undetectable on that slice, the pipeline safely falls back to a predefined array of empirical angles as mentioned in subsection 5.9.5.

For every valid Z-slice, the grid search systematically explores combinations of spatial offsets and orientations:

- **Rotations:** Combinations of sagittal (R_x) and frontal (R_y) angles. R_x being mostly set by the presence of the bar and R_y by slight angles user-defined (across -15 and $+15$ in the code).
- **Translations:** Slight offsets (Y-axis) around the previously computed centre of mass to account for tissue alignment variations.

For each discrete combination, the algorithm constructs the corresponding affine transformation matrix, projects the histological point cloud into the MRI space, and evaluates the hybrid Mutual Information score.

Once the global scan is complete, all tested configurations are compiled and sorted in descending order based on their objective scores. Instead of retaining only the absolute best match, the algorithm isolates the Top N highest-scoring parameter combinations (the top 5 slices with the highest matching score in this case). Retaining multiple distinct potential slices is a safeguard: it prevents the pipeline from committing too early to a false positive because once the whole transformations performed, a slice a priori less good could have become better than the original best peak depending the transformations. These top peaks serve as highly robust, pre-aligned initial seeds for the final continuous optimization phase.

Fine 3D Alignment and Continuous Optimization

While the initial grid search is meant to identify the most probable anatomical regions and avoids global (angle) misalignments, its discrete nature only provides a basic slice from the initial scan. To achieve an optimal, spatial matching, the pipeline will now perform a fine-tuning optimization phase.

This phase utilizes SciPy’s implementation of Powell’s conjugate direction method (`scipy.optimize.minimize`). Powell’s method is a highly robust, derivative-free optimization algorithm, which is necessary for this specific multi-modal problem. When the Mutual Information (MI) metric is evaluated across a complex voxel grid as the one of a brain scan, the necessary spatial interpolations will eventually create high-frequency noise tiny artifacts in the objective function’s landscape. Standard optimization algorithms based on gradient descent calculate the local derivative (the slope) to find the minimum. Consequently, these gradient-based solvers are easily biased by this interpolation noise: they detect a false local slope and converge to a local minimum.

Powell method is independent of this high-frequency noise. Rather than computing local mathematical derivatives, it relies on sequential 1-dimensional line searches along bi-directional vectors. By evaluating the metric further along a line rather than just checking the immediate slope, the algorithm effectively steps over the high-frequency grid noise and smoothly navigates toward the true macroscopic optimum.

To maximize the chances of finding the true global maximum, the Powell optimizer is executed independently for each of the top N slices primary found during the first grid search. For each iteration, the solver is initialized with the slice original specific parameters $(t_x, t_y, t_z, R_x, R_y, R_z)$ and a default scale of 1.0.

The major feature of this step is the application of the boundary constraints (bounds) on the 8-parameter optimization vector:

- **Translations:** Confined to a strict local neighbourhood around the initial geometric guess. In-plane translations (t_x, t_y) are restricted to ± 20 voxels to ensure the slice remains centered on the target tissue mass. The depth search (t_z) is constrained to a margin of ± 10 to ± 15 slices, with tighter bounds automatically applied to datasets where the initial artifact detection provided a higher confidence in the starting Z-coordinate.
- **Rotations:** Restricted to small angular corrections to tune the orientation without drastically altering the oblique cutting plane. The Z-axis rotation (R_z) is strictly confined to $\pm 10^\circ$. For the remaining axes, the constraints dynamically adapt to the dataset’s initialization quality: when the mounting apparatus is detected, the sagittal

tilt (R_x) is tightly constrained ($\pm 5^\circ$) to respect the robust initial angle provided by the artifact, while granting more flexibility to perpendicular rotation angle ($R_y \in [-10^\circ, 10^\circ]$). In datasets lacking this physical anchor, these rotational freedoms are logically inverted to allow broader sagittal exploration ($\pm 10^\circ$ for R_x).

- **Scaling:** Constrained within a plausible biological and physical deformation range (0.8 to 1.2) to accommodate tissue shrinkage or stretching without inducing large and unrealistic distortions. This constraint is important to accurately align the macroscopic boundaries of both modalities, as unconstrained scaling could lead to truncated voxel-pixel comparisons and artificially leading to biased alignment scores.

While a complete affine transformation would also include three shearing parameters, these were deliberately excluded from this affine optimization. Allowing them would lead the algorithm to maximize the Mutual Information score by inducing parallelogram-like deformations of the brain structure. Such morphological distortions like local tissue tearing are mostly due to non-elastic deformations and cannot be accurately corrected using only linear transformations. Instead, they are left to be resolved in the subsequent non-rigid elastic registration phase (see section 5.10).

By confining the solver to these boundaries, the algorithm is forced to refine the specific oblique plane rather than wandering across the entire brain volume (more efficient). Once the optimization has converged for all N original slices, the pipeline compares their final maximized objective scores. The specific configuration that yields the absolute highest hybrid Mutual Information score is retained as the global optimal spatial alignment.

Slice extraction for visual validation

Once the optimal affine transformation matrix and the best initial slice are identified, the final task of the pipeline is to generate and extract the transformed MRI slice that corresponds exactly to the histological block. Because the datasets are computationally demanding, the algorithm creates a discrete spatial matrix characterized by the exact dimensions and pixel spacing of the blockface scan to serve as the final output. Given the high spatial dimensionality of this target matrix, multiplying the entire coordinate grid by the affine matrix simultaneously would cause severe memory overflow and lead to a system failure.

To overcome this issue, a sequential row-by-row extraction algorithm was implemented. The pipeline iterates along the height (Y-axis) of the target 2D space optimized by Powell's method. For each iteration, an array representing a single horizontal row of pixels for the final matrix is generated. These homogeneous coordinates are multiplied by the optimized affine matrix to project this specific line into the 3D MRI volume space. This matrix acts as the spatial mapping function that defines the exact position of each histological pixel within the 3D coordinate system of the MRI.

However, these projected 3D coordinates are continuous (floating-point values) and rarely

align perfectly with the discrete integer centers of the original MRI voxels. To assign a precise intensity value to each new pixel, the projected coordinates are passed to a trilinear interpolation function (`scipy.ndimage.map_coordinates`). The role of the interpolation will be here to have a smoothed and continuous final extracted slice because the projection will nearly never match a voxel of the 3D scan and could contain pixelation artifacts. To counter this effect, the algorithm calculates a weighted average of the 8 nearest surrounding voxels.

Finally, for the qualitative visual assessment, the algorithm will plot 3 different figures creating a comparative panel.

1. The fixed, pre-processed 2D histological block (as reference).
2. The newly extracted, optimally aligned MRI slice.
3. A direct spatial superposition of both modalities.

The superposition is mostly to compare to ventricles position or the edges of the brain. Even if the slices are still not exactly matched (because the elastic deformations are still not corrected), the coarse features of both modalities should be aligned.

5.9.6 Summary of the Decoupled Thresholding Strategy

Throughout the multimodal registration pipeline, distinct intensity thresholds are applied. Rather than relying on a single global value, this decoupled approach ensures that each mathematical operation ranging from intensity normalization to spatial initialization and fine topological alignment is governed by a specifically tailored limit. Table 5.1 provides a comprehensive overview of these thresholds in their chronological order of execution, detailing their respective stages within the computational pipeline and their mathematical justifications.

5.10 Elastic corrections

While the 3D-to-2D affine registration successfully aligns the global orientation and extracts the corresponding optimal MRI slice, the matching algorithm is limited to 8 linear transformations (3 translation, 3 rotation, and 2 global scaling). It cannot account for the non-linear morphological distortions such as tissue shrinkage, ventricles deformations or other physical deformations that appeared during the physical histological sectioning and mounting processes. To optimize the spatial correspondence and resolve these complex distortions, a final 2D-to-2D non-rigid (elastic) registration step is required.

This final phase operates in two main steps: a geometric initialization followed by a multi-resolution B-Spline deformation model that will improve the edges alignment. The pipeline

Parameter	Threshold	Pipeline Stage	Mathematical & Physical Justification
Statistical Baseline	$I_{\text{histo}} > 15$	Intensity Normalization	Ensures statistical tissue purity. Excludes the background and microscopic noise to accurately compute the Z-score parameters (μ, σ) without artificially lowering the true tissue mean.
Centroid Initialization	$> 30\% \text{ of } \mu$	Spatial Initialization	Acts as a coarse geometric filter. Isolates the bulk tissue mass to compute a the centre of gravity, preventing disconnected edge artifacts from skewing the spatial starting coordinates.
Topological Anchor	$I_{\text{histo}} > 5$	Optimization Loop (Gradient MI)	Prevents morphological erosion. Permissive enough to capture the true cortical boundaries for gradient matching, yet strict enough to filter out low-intensity artifacts generated by the bilinear spatial interpolations.
MRI Signal Validation	$I_{\text{MRI}} > 25$	Optimization Loop (Overlap Penalty)	Acts as a statistical safeguard. Excludes the darkest 1% of the normalized MRI tissue, effectively rejecting partial volume effects and non-anatomical signal voids during the overlap validation.

Table 5.1: Chronological summary of the decoupled intensity thresholds used across the 2D-to-3D co-registration pipeline.

is executed as follows.

5.10.1 Geometric Initialization via Thin Plate Splines (TPS)

The first phase corrects major macroscopic tissue displacements. An automated contour-detection algorithm is once again applied to the binary masks of both the MRI slice and the histological image. To find robust key-points, the algorithm geometrically divides the images into left and right hemispheres (just by cutting at the geometrical centre of the image). For each isolated hemisphere, the algorithm calculates its geometric centre (centre of mass) using spatial image moments (`cv2.moments`).

These two computed centroids serve as primary internal anatomical anchors. However, relying solely on internal points could cause the outer cortical boundaries to undergo severe and unnecessary deformations. To prevent this, eight additional static anchor points are strategically placed along the image boundaries, ensuring that the deformations maintain

the global image shape. A Thin Plate Spline (TPS) shape transformer evaluates these corresponding point sets to compute an initial non-linear warp.

Theoretically, the Thin Plate Spline (TPS) algorithm is a spline-based spatial interpolation method. As established by Bookstein [33], it constructs a mapping function by combining a global affine transformation with a linear combination of local radial basis functions centred at defined anchor points. The algorithm computes a deformation field that precisely aligns these landmarks while minimizing a penalty defined as the integral of the squared second-order partial derivatives of the spatial mapping thereby ensuring a maximally smooth deformation. The use of such landmark-based optimization strategies for non-linear brain image registration has already been established to handle structural variations [34]. The specific application to histological-to-MRI co-registration is also well documented in the literature [27].

In the practical implementation of our pipeline, a TPS shape transformer evaluates our specifically defined corresponding point sets (the extracted centroids and boundary anchors) to compute this initial non-linear warp. This applied transformation step forces the macroscopic internal tissue masses of the hemispheres to align accurately while effectively maintaining the global geometric integrity of the cortical perimeter. Finally, this computed initial deformation field is applied to resample both the high-resolution histological image and its corresponding binary mask.

An alternative strategy involved placing anchor points within the ventricles. However, this approach relies on accurate ventricular segmentation, thus necessitating a highly robust masking pipeline regardless of the initial MRI scan quality. With Patient 854, where the ventricles could not be well masked out, this technique is non applicable. Furthermore, positioning anchors at the ventricles further from the brain contours allows for much larger geometric variations, leading to non physiologically possible deformations.

5.10.2 Multi-Resolution Silhouette B-Spline Transformation

Following the geometric initialization, the next step of the pipeline uses the *SimpleITK* framework to compute a highly precise non-rigid spatial warp. Unlike earlier methods relying on pixel intensities or discrete landmarks, this final optimization is driven entirely by the geometrical alignment of the tissue shapes (binary masks).

A *B-Spline* transformation matrix is initialized over the MRI mask using a highly flexible 15×15 pixel control point grid (using `BSplineTransformInitializer`). The optimizer aims to minimize the Mean Squares metric between the fixed MRI mask and the histological mask to correct. The matching step of section 5.9 uses the histological section as reference to go through the whole scan et match the best slice possible but now. The extracted MRI slice now serves as the fixed reference to account for the elastic defor-

mations that occurred in the histological tissue, ensuring a highly precise, pixel-to-voxel spatial alignment between the two modalities.

To ensure a robust convergence and prevent too large foldings, a multi-resolution process is implemented. The registration is performed across three hierarchical levels of Gaussian smoothing and subsampling (shrink factors of 4, 2, and 1). Beginning with the highest subsampling gives to the algorithm only the coarse details to be aligned. Then by reducing it, the coarse shape is aligned and now the algorithm has access to more details that it can align easier since the coarse shapes are aligned. Finally the smallest details are also corrected allowing the global image to be perfectly aligned without large and unnecessary deformations

Driven by a Gradient Descent Line Search optimizer, the B-Spline model generates a deformable continuous grid to apply the deformations to blocks of pixels instead of one pixel at a time which would be too intensive. Finally, once the optimal non-linear transformation field is computed from the masks, this exact B-Spline transform is applied to resample the high-resolution, corrected histological tissue. This produces the final co-registered image, where both the global cutting angle and the local microscopic deformations are perfectly aligned with the MRI reference, making it ready for precise clinical and neuropathological analysis.

5.11 Summary of the global Co-registration pipeline

To provide a global a overview of the methodologies described throughout this chapter, the multimodal registration workflow is summarized in Figure 5.12.

The block diagram illustrates the pipeline's three phases.

First, the parallel pre-processing of both modalities: the 3D MRI volumes undergo bias correction and structural skull-stripping, detailed in Figure 5.3, while the 2D histological sections are subjected to background removal and photometric inversion as detailed in Figure 5.10.

Second, the iterative 2D-to-3D affine matching, which relies on automated artifact detection and a hybrid Mutual Information metric to extract the optimal corresponding MRI slice.

Finally, the 2D-to-2D non-linear elastic correction phase, which uses *Thin Plate Splines* and multi-resolution *B-Splines* to resolve the localized tissue distortions and generate the final aligned images.




Figure 5.12: Global block diagram of the multimodal coregistration pipeline. The distinct pre-processing steps are grouped to highlight the macroscopic data flow leading to the affine and elastic registrations.

5.12 Code Availability

To facilitate the reproducibility of the proposed methodology, the complete pipeline described in this chapter including the 3D skull-stripping, the hybrid MI optimization, and the non-rigid elastic warping algorithms has been implemented as a Python framework. The full source code is available on the project's GitLab repository¹ for further investigation and adaptation.

¹https://gitlab.uliege.be/CyclotronResearchCentre/Studies/p2025_histoms/match_histo_evqmri

Chapter 6

Results

The efficacy and accuracy of the volumetric co-registration pipeline inherently depend on the quality of the input data. As detailed in chapter 4, the raw datasets presented severe structural and photometric artifacts. This chapter first evaluates the robustness of the automated pre-processing algorithms, before presenting the final 2D-to-3D matching performances and their clinical applicability.

6.1 Pre-processing Robustness

This section first evaluates the ability of the developed pipelines to preserve anatomical integrity while effectively suppressing non-cerebral noise in both the 3D MRI volumes and the 2D histological sections.

6.1.1 3D Brain Masking on Multi-Contrast MRI (Patient 784)

The automated 3D skull-stripping pipeline was evaluated on the MRI acquisitions of Patient 784. The primary computational challenge was to achieve consistent tissue isolation across quantitative maps (PD , R_1 , R_2^* , and MT_{sat}) that exhibit drastically different Signal-to-Noise Ratios (SNR) and background characteristics.

As shown in Figure 6.1, the algorithm demonstrated high consistency across modalities without requiring manual tuning the contrast parameters. For structural maps with relatively low background noise, such as the R_2^* contrast Figure 6.1a, the pipeline successfully excluded the surrounding immersion fluid, the ventricles, the largest sulci and the physical mounting container while preserving the cortical edges Figure 6.1e.

Furthermore, the algorithm also shows a high efficacy on the highly degraded MT_{sat} map Figure 6.1b. Despite the presence of the high-frequency background noise and signal inhomogeneities that severely degraded the brain boundaries in the raw volume, the adaptive thresholding which calculates the cut-off dynamically based on the local tissue mean effectively isolated the cerebral hemispheres Figure 6.1f.

Visual assessment shows that the resulting binary masks really well surrounds the anatomical tissue across both contrasts, setting the background arrays to absolute zero. This skull-stripping provides a clean and standardized anatomical basis for the subsequent spatial registration steps.

Table 6.1: Hyperparameters of the 3D automated brain masking pipeline across MRI contrasts (Patient 784).

Pipeline Parameter	PD^*	R_1	R_2^*	MT_{sat}
3D Median Filter Kernel Size	$3 \times 3 \times 3$	$3 \times 3 \times 3$	$3 \times 3 \times 3$	$5 \times 5 \times 5$
Adaptive Threshold Multiplier	$0.85 \times \mu$	$0.85 \times \mu$	$1.1 \times \mu$	$1.1 \times \mu$
Structuring Element Connectivity	1	1	1	1
3D Binary Opening Iterations	6	6	6	6
3D Binary Closing Iterations	10	10	10	10
Target Connected Components	2	2	2	2

Note: μ represents the dynamically updated local mean calculated within the adaptive 3D bounding box to isolate the tissue from zero-intensity background.







(a) Raw R_2^* / PD slice(b) Raw MT_{sat} slice(c) Computed R_2^* mask(d) Computed MT_{sat} mask(e) Final masked R_2^* (f) Final masked MT_{sat}

Figure 6.1: Preprocessing robustness comparison. The top row displays the original, uncorrected slices where the MT_{sat} contrast exhibits significantly higher background noise compared to the R_2^* . Despite these extreme initial differences, the adaptive 3D masking pipeline successfully isolates the brain tissue in both modalities (middle and bottom rows).

6.1.2 3D Brain Masking for Patient 854

The results of the pipeline for the field bias correction (section 5.2) are presented on Figure 6.2 and Figure 6.3. The bias field is effectively corrected allowing for a much more homogeneous intensity distribution. We can observe aliasing effects on the lateral sides of both view. Figure 6.2b and Figure 6.3b also have the mask applied, however the background is not perfectly rejected and the cage still has high voxels intensities.

To clearly illustrate the limitations of our initial skull-stripping algorithm on this specific dataset, Figure 6.4 presents the masking results applied directly to the raw volumes (prior to any field bias correction). Due to the low-frequency intensity fluctuations present across the uncorrected 7T volume, the algorithm fails to accurately isolate the brain tissue. The algorithm relying on global intensity thresholds, the signal drop in the deeper parts of the brain causes it to mask true anatomical structures, treating them as background. The hyper-intensities at the outer cortical surfaces prevent the morphological operations from accurately masking the non-anatomical space between the brain cage and the brain tissue, leaving unwanted background artifacts attached to the final masked volume.


In an attempt to resolve these intensity gradients, the masking pipeline was applied to the volumes after performing the N4ITK bias correction, as shown in Figure 6.5. The N4ITK bias field correction algorithm, using the parameters summarized in Table 6.2.

Table 6.2: Parameters utilized for the N4ITK bias field correction algorithm.


Parameter	Value	Description
Shrink Factor	4	Subsampling factor
Number of Fitting Levels	4	Number of hierarchical multi-resolution levels used by the B-Spline optimization grid.
Maximum Iterations	50	Defines the stopping criterion if convergence is not reached.
Convergence Threshold	0.1%	Stopping criterion based on the variation between subsequent bias field estimates.
Spline Order	3 (Cubic)	Degree of the B-Spline polynomials generating the continuous deformable grid (default algorithmic value).

However, this approach introduced new, contrast-specific errors. For the Proton Density (PD) contrast, flattening the bias field resulted in a mask that became significantly less restrictive (too permissive). Consequently, the background located within the deeper

anatomical structures and internal cavities is no longer masked. For the Magnetization Transfer (MT) contrast, the intensity distribution after the correction resulted in an too aggressive masking, leading to an even larger proportion of the deeper anatomical structures erased compared to the uncorrected raw scan. These results demonstrate that a standard intensity-based masking approach, is not suitable for the complex contrast issues of Patient 854.




(a) Raw scan




(b) Corrected scan

Figure 6.2: Frontal view of the bias correction pipeline results, proton-density contrast at first echo-time




(a) Sagittal view of the raw scan




(b) Sagittal view of the corrected scan


Figure 6.3: Sagittal view of the bias correction pipeline results for Patient 854, proton-density contrast at first echo-time




(a) Mask applied on Patient 854 (PD contrast, first echo time)



(b) PD contrast scan of Patient 854, brain masked




(c) Mask applied on Patient 854 (MT contrast, first echo time)




(d) MT contrast scan of Patient 854, brain masked


Figure 6.4: Effect of brain mask applied on Patient 854, coronal view with no bias correction applied, slice 288, scan of first echo time




(a) Mask applied on Patient 854 (PD contrast, first echo time)



(b) PD contrast scan of Patient 854, brain masked



(c) Mask applied on Patient 854 (MT contrast, first echo time)



(d) MT contrast scan of Patient 854, brain masked

Figure 6.5: Effect of brain mask applied on Patient 854, coronal view (after bias correction), slice 288, scan of first echo time

6.1.3 2D Blockface and Histological Masking (Yellow Masks)

Following the standardized 3D MRI skull-stripping, the segmentation of the corresponding high-resolution 2D blockface histological images is performed using the algorithm called "yellow masks". The primary computational objective was to isolate the cerebral tissue from the photographic background, which consists of the laboratory surface along with potential shadows and other reflections.

As shown in Figure 6.6, the pipeline successfully isolated the brain tissue but was intentionally set to a permissive configuration regarding background objects. Visual comparison reveals that while the intensity thresholding efficiently separated the biological signal from the darker background, the large morphological opening operator (25×25 pixels) was sometimes insufficient to fully mask out brighter artifacts (mounting residues reflexion in the sulci) from some parts of the brain. Consequently, the algorithm remains permissive to prevent the loss of real cortical signal, particularly in the tiny sulci, the parameters were also configured to preserve the two largest connected components. As shown in

Figure 6.6, the primary cerebral mass is entirely preserved and correctly isolated, ensuring that no brain tissue is lost.

However, this strategy results in residual background artifact (in the sulci on the edges of the brain tissue) being included in the final binary mask. While non-ideal for true segmentation, this result is accepted as a necessary trade-off: a slightly larger search area for the subsequent registration is preferable to a loss of anatomical reference.


Table 6.3: Hyperparameter adaptations for 2D blockface and histological masking.

Pipeline Parameter	Standard Configuration	Specific Slices (e.g., 11_3, 15_2)
Median Blur Kernel Size	35×35	15×15
Binary Intensity Threshold	45	65
Morphological Opening Kernel	15×15	25×25
Preserved Connected Components	1	2


Note: Specific adjustments were required for a subset of images exhibiting altered lighting conditions, reflections or other mounting residues. The morphological opening was intensified and the threshold made stricter to counter these artifacts, while two connected components were preserved to ensure no genuine cortical tissue was eroded.

This strategy was then applied to the second dataset (Patient 854) to assess the robustness of the "yellow masks" pipeline under different experimental conditions. As illustrated in Figure 6.7 and Figure 6.8, the algorithm maintained its performance despite variations in tissue contrast and background lighting even if some tissue zones a bit darker have been masked out.

However, some histological blocks (such as blocks 24 and 27, shown in Figure 6.9 and Figure 6.10) presented more complex artifacts and still masked out darker zones. Even in these cases, the permissive configuration of the pipeline successfully preserved the cerebral morphology. While these additional background residuals remained present, the consistency of the masking results across both patients validates the scalability of the proposed method for multi-site histological studies.




(a) histo block 11 mask




(b) histo block 11 masked




(c) histo block 18 mask



(d) histo block 18 masked




(e) histo block 19 masked




(f) histo block 19 masked


Figure 6.6: masks for 784 histo blocks



(a) Block 18 mask



(b) Block 20 mask



(c) Block 23 mask

Figure 6.7: Binary masks generated for the blocks of Patient 854.




Figure 6.8: Resulting masked MRI volumes for Patient 854 (blocks 18, 20 and 23).




Figure 6.9: Binary masks for the blocks 24 and 27 of Patient 854.



Figure 6.10: Resulting masked MRI volumes for Patient 854 (blocks 24 and 27)

6.2 Matching assessment

6.2.1 Patient 784

The primary objective of the 3D-to-2D affine registration pipeline was to autonomously navigate the continuous spatial parameters (translations, rotations, and scaling) to extract the virtual MRI slice that perfectly corresponds to the arbitrary cutting angle of the histological blocks. The matching performance was initially assessed on the dataset of Patient

784 using the $R2^*$ contrast.

As illustrated in Figure 6.11, the automated optimization pipeline successfully recovered the global spatial orientation of the different histological sections. The qualitative visual assessment of the matching demonstrates satisfactory macroscopic alignment. The principal anatomical landmarks, including parts such as the ventricles, global cortex shape or large sulci are well matched. This confirms that the hybrid matching function combining intensity-based Mutual Information with structural gradient matching reached a good matching efficiency.


However, slight adaptations of the parameters were needed to achieve the best match possible in each case. As detailed in Table 6.4, the weighting factor (α) of the hybrid metric, the initial Z-axis search range, and number of sampling points were adjusted to optimize the accuracy of each match.

Table 6.4: Optimization hyperparameters used for the affine volumetric registration (Patient 784).


Histological Block	α (MI / Gradient)	Slice Init (Z)	Slice End (Z)	Sampled Points (N)
Block 11	0.7	380	410	20 000
Block 18	0.8	260	300	10 000
Block 19	0.2	230	260	20 000

Note: The Z-slice search range represents the initial and final slices where the algorithm searched for the optimal slice (coarse range determined with a quick analysis of the 3D scan using **ITK-Snap**), while the number of sampled points dictates the tissue sampling.


The global matching is satisfactory but also shows the limitation of only making the rigid affine optimization. While the macroscopic orientation is correct, localized geometric features remain not perfectly superposed. These mismatches are not algorithmic failures, but rather the direct visual manifestation of the non-linear deformations. This results are a very good basis but would need to use the pipeline in section 5.10 to be more accurately superposed.



(a) Block 11



(b) Block 18



(c) Block 19

Figure 6.11: Rigid matching results for Patient 784. For each block, the panels display from left to right: the extracted MRI slice, the corresponding histological image, and the color overlay showing the spatial alignment.


6.2.2 Patient 854 with no bias correction

The algorithm was then evaluated on the raw dataset of Patient 854, using the first echo time of the Proton Density (PD) contrast. Firstly without any initial bias correction, the objective was to assess if the hybrid metric is efficient with data of lower quality mask-


ing. The pipeline using the data from Patient 854 showed variable results. As shown in Figure 6.12, the algorithm successfully achieved satisfactory spatial alignments for one section (section 20) and promising results for an other one (section 24) which has its left part with good matches while the right one is less well matched. The best results were the one when using a higher alpha.

Table 6.5: Optimization hyperparameters used for the affine volumetric registration (Patient 854).

Histological Block	α (MI / Gradient)	Slice Init (Z)	Slice End (Z)	Sampled Points (N)
Block 18	0.6	400	430	20 000
Block 20	0.6	350	400	20 000
Block 23	0.9	280	320	30 000
Block 24	0.9	240	290	10 000
Block 27	0.3	190	230	20 000



(a) Block 20 matching



(b) Block 24 matching


Figure 6.12: Patient 854: best matchings obtained before bias correction.

However, those data also highlighted the algorithm limits, resulting in lower quality alignments for other sections, as illustrated in Figure 6.13 (blocks 18, 23, 27). Those blocks show nearly nothing in common with the corresponding slices found by the algorithm. Since


the registration algorithm matches the high-intensity regions, the hyperintense signal of the brain cage improves the matching score for wrong slices, causing the algorithm to false matchings.

These mixed outcomes highlight the limits of the **Mutual information** metric when confronted with poorer quality data. It can be noted that the affine correction well aligned higher-intensity zones in all cases showing a good baseline of the algorithm affine correction phase.


This justifies the exploration of the N4ITK bias field correction (to homogenize the search space) aiming to reduce those intensity variations artefacts.



(a) Block 18



(b) Block 23




(c) Block 27


Figure 6.13: Lower quality results for Patient 854 before bias correction.

6.2.3 Patient 854 with bias correction


Following the mixed outcomes observed on the raw dataset for Patient 854, the N4-ITK Bias Field Correction pipeline (detailed in section 5.2) was integrated before the affine registration for this dataset, still using the first-echo time of the Proton Density (PD) contrast. The theoretical results would be that correcting the low-frequency bias field would homogenize the image intensities, preventing the Mutual Information optimizer from converging on false local maxima caused by the hyperintense structures around the brain tissue.



(a) Block 18



(b) Block 23




(c) Block 27


Figure 6.14: Matchings of Patient 854 after bias correction keeping the same parameters for the bad matchings

However, contrary to these expectations, visual assessment reveals that the signal homogenization did not improve the spatial alignment; on the contrary it lead to a loss in matching accuracy. As illustrated in Figure 6.14, the bias field correction did not resolve the initial registration failures. On the contrary, the misalignments are even more clear, resulting in even larger spatial deviations for these problematic blocks. Furthermore, when applied to sections that previously showed good matching on the raw data, the optimizer also exhibits a worse matching for the block 20 while the optimizer converges to a suboptimal slice matched for Block 24 (on Figure 6.15), making the co-registration poorer.

This counter-intuitive outcome provides a significant methodological insight: for a gradient-driven Mutual Information framework, preserving the native, high-contrast tissue boundaries is more important for optimizer convergence than achieving a perfectly flat global intensity profile. Consequently, the N4 bias field correction was deliberately excluded from the volumetric registration pipeline.



(a) Block 20 matching with bias correction



(b) Block 24 matching with bias correction

Figure 6.15: Matchings of Patient 854 after bias correction keeping the same parameters for the best matchings

6.2.4 Elastic corrections

An overview of the elastic registration outcomes for Patient 784 is presented in Figure 6.16. As observed in Figure 6.16a, Figure 6.16b and Figure 6.16c for Patient 784, the elastic registration algorithm successfully fulfills its primary geometric constraint: driving the boundaries of the histological tissue to match the corresponding MRI mask to achieve a perfect macroscopic superposition. However, a critical analysis of the deformation fields reveals significant limitations depending on the initial data quality.

The parameters used are presented in Table 6.6.

Table 6.6: Hyperparameters of the non-linear B-Spline elastic registration pipeline.


SimpleITK Parameter	Value	Functional Description
B-Spline Grid Size	15×15	Spatial density of the deformable control points.
Shrink Factors	[4, 2, 1]	Subsampling factors for the 3-level multi-resolution pyramid.
Smoothing Sigmas	[2, 1, 0] mm	Gaussian blur applied per level to avoid high-frequency local minima.
Learning Rate	1.0	Initial step size for the gradient descent spatial exploration.
Max Iterations	50	Maximum allowed optimization steps per resolution level.
Convergence Threshold	10^{-6}	Minimum required metric improvement to continue optimization.
Convergence Window	10	Number of consecutive iterations evaluated for the stopping criterion.

On this example, the contour matching is generally excellent. Specifically, as observed in Block 11 (Figure 6.16a), a minor geometric artifact (leakage) occurs on the superior edge. This is likely caused by an initial anchor point erroneously attaching to the external embedding cage rather than the true cortical surface. Consequently, the local tissue including the visibly distinct brown structural region was artificially deformed by the algorithm as it tend to compensate to close the spatial gap with the opposite hemisphere. Block 18 (Figure 6.16b) displays an accurate superposition driven by the elastic corrections algorithm. However, a defect remains: the histological section overlaps too much with the cerebellum. It is important to note that this is not a failure of the elastic non-rigid algorithm, which successfully superpose both modalities. Instead, this artifact originates from the initial MRI skull-stripping phase, where the cerebellum was not fully masked out. Similarly, Block 19 (Figure 6.16c) suffers from a comparable overlap with the remaining cerebellar tissue, while the ventricular cavities at this level tend to collapse under the boundary alignment.


More importantly, while the macroscopic silhouette is perfectly superimposed, this forced boundary matching comes at a high cost for the internal structural integrity. Because the optimization relies exclusively on the external mask contours without any internal regularization points (such as the ventricles or deep gray matter boundaries), the algorithm distributes the tissue directly into the whole brain tissue mask no matter the morphology. The B-spline deformations severely warp the internal architecture, compressing or unnaturally stretching the tissue, ultimately making the final histological images completely unreadable for reliable structural or clinical analysis, as shown in Figure 6.17. These lim-

itations are even more pronounced for Patient 854 (Figure 6.18). In this dataset, the algorithm aggressively stretches the tissue to fill the central regions. Because the MRI masks are compromised by the previously discussed intensity bias artifacts, the fixation cage and other non-anatomical regions are artificially bright. The algorithm blindly follows this target, forcing the tissue to over-expand into these artefactual spaces.


Consequently, the elastic deformation step was intentionally omitted for the remaining sections of Patient 854 that had already failed during the initial rigid alignment. Applying complex elastic transformations to fundamentally misaligned and artifact-heavy data would only amplify the errors without providing any scientifically relevant results.



(a) Block 11: Elastic corrections



(b) Block 18: Elastic corrections



(c) Block 19: Elastic corrections

Figure 6.16: Overview of elastic corrections applied to Patient 784.




Figure 6.17: Results of the 2D non-linear elastic registration pipeline applied to histological sections of Patient 784.






Figure 6.18: Visual assessment of the 2D elastic deformation process for Patient 854.

6.3 Unbounded matching


To empirically validate the necessity of the axial search space restriction discussed in the methodology, a supplementary test was conducted on the dataset of Patient 784. The registration pipeline was executed for histological sections 11, 18, and 19 of Patient 784

without any slice boundaries, allowing the optimizer to explore the entire 3D MRI volume blindly.


As anticipated, while the algorithm successfully converged toward a mathematical optimum (meaning a high Mutual Information score), the resulting spatial alignments were anatomically incorrect, as illustrated in Figure 6.19. Driven by the macroscopic self-similarity of the cortical structures along the longitudinal axis, the optimizer was systematically trapped in false local maxima.



(a) Block 11



(b) Block 18



(c) Block 19

Figure 6.19: Rigid matching for Patient 784, unbounded conditions

For instance, the unconstrained search for Block 19 converged at an MRI slice of $Z = 407.2$ instead of its true anatomical location near $Z = 256.6$. Similarly, Block 11 matched with a slice at $Z = 266.1$ rather than $Z = 379.8$. In these failure cases, the algorithm matched the blockface to entirely different sets of sulci and ventricular cavities that presented a statistically similar macroscopic intensity distribution.

It therefore confirms that relying solely on unsupervised Mutual Information optimization without prior macro-localization inevitably leads to severe misalignments. This strictly justifies the implementation of a constrained Z-axis search window in the final pipeline to safeguard the geometrical convergence.

6.4 Stains matching

An additional test was performed to evaluate the pipeline's robustness when applied directly to the chemically stained histological slices rather than the intermediate blockface photographs. The objective was to determine if the hybrid Mutual Information metric could overcome the severe contrast alterations introduced by chemical processing. The parameters used were the same as those of the optimal results obtained.


The evaluation was conducted on slices from Block 11 using three distinct macroscopic stains: Haematoxylin and Eosin (HE), Kluver-PAS (KPAS), and Turnbull blue (TBBu). It can be mentioned that other stains available in the dataset were deliberately excluded from this specific evaluation because they exhibited macroscopic damages such as large tears and missing tissue fragments introduced during the mounting process. Including them would not have provided any additional valuable data.

The matching results, illustrated in Figure 6.20, highlight a severe breakdown of the intensity-based MI metric. Because specific stains selectively emphasize distinct cellular structures while completely suppressing the signal of others, the resulting 2D contrast profile becomes uncorrelated to the quantitative MRI maps. Consequently, the algorithm failed to establish consistent statistical co-occurrences within the joint histogram.


Therefore, the optimizer was unable to match the corresponding slices accurately. As seen in Figure 6.20, the algorithm converged on completely different and incorrect slices depending on the stain used (e.g., $Z = 403.5$ for HE and $Z = 408.7$ for KPAS, instead of the true anatomical target near $Z = 379.8$). Even when the spatial plane was closer to the target (as with TBBu at $Z = 377.9$), the resulting superposition demonstrated significant in-plane rotational and translational misalignments, trapping the optimizer in false local maxima.

To confirm that this failure is only driven by the chemical contrast, the HE-stained sections


of Blocks 18 and 19 were also evaluated, as illustrated in Figure 6.21. The results demonstrate that the algorithmic breakdown is systematic. For Block 18, the optimizer found a spatially proximate plane (converging at $Z = 287.3$ instead of the true $Z = 284.4$) but suffered from severe in-plane rotational and translational errors, failing to properly superimpose the hemispheric boundaries. For Block 19, the algorithm experienced a complete volumetric mismatch, converging at $Z = 240.1$ instead of the true anatomical target near $Z = 256.6$, resulting in a completely erroneous alignment of the lateral ventricles. This confirms that the photometric distortion introduced by histological staining systematically compromises the Mutual Information metric, regardless of the anatomical region.



(a) Haematoxylin and Eosin (HE)




(b) Kluver-PAS (KPAS)



(c) Turnbull blue (TBBu)

Figure 6.20: Rigid matching results for Patient 784 (Block 11) using different macroscopic stains under unbounded conditions.



(a) Block 18 (HE stained)



(b) Block 19 (HE stained)

Figure 6.21: Patient 784: Rigid matching failures on HE-stained histological sections for Blocks 18 and 19, demonstrating systematic algorithmic breakdown.

6.5 Summary of Coregistration Results

This section provides a factual and quantitative overview of the pipeline's performance across the different datasets and experimental conditions.

6.5.1 Pre-processing and Tissue Masking

The automated contour-detection algorithm successfully isolated the cerebral tissue from the background. Standard parameters (e.g., 35×35 pixels median blur, threshold at 45) performed well for optimal scans but histological sections showing prominent reflections or mounting residues required stricter, case-specific morphological adjustments (thresholds up to 65 and 25×25 pixels kernels) to preserve only the cortical tissue.

6.5.2 Rigid and Affine Initialization

- **High-quality data (Patient 784):** The Mutual Information-based rigid registration successfully established the macroscopic spatial correspondence between the intermediate blockface photographs and the MRI volumes.

- **Lower-quality data (Patient 854):** The current robustness of the pipeline is insufficient to produce reliable outcomes. While 2 out of 5 blocks (Blocks 20 and 24) achieved acceptable initial alignments, the remaining 3 blocks (Blocks 18, 23, and 27) completely failed. The optimizer systematically matched the hyperintense MRI signal of the external brain cage rather than the cortical tissue.

6.5.3 Bias Field Correction Impact

Applying the N4ITK bias field correction to Patient 854 upgraded the intensity homogeneity but it did not help recovering fine alignments. Instead, it makes it worse for the 3 failed blocks and even degraded the previously acceptable results for Blocks 20 and 24, forcing the optimizer to reach suboptimal configurations.

6.5.4 Geometric Initialization via TPS

For the successfully matched slices (Patient 784), the Thin Plate Spline (TPS) transformation, constrained by 8 outer boundary anchors and 2 internal center-of-mass anchors, successfully provided an initial non-linear macroscopic alignment. It effectively forced the general hemispheres to overlap while maintaining the global geometric integrity.

6.5.5 Elastic Non-Rigid Deformation

- **External Boundary Matching:** The 15×15 multi-resolution B-Spline registration successfully fulfilled its geometric constraint on Patient 784, achieving a visually good superposition of the external cortical contours. Minor artifacts were strictly limited to external initializations (e.g., cage leakage on Block 11) or unmasked MRI structures (e.g., cerebellum overlap on Blocks 18 and 19).
- **Internal Structural Integrity:** Driven only by external binary masks without internal regularization anchors, the algorithm compressed and stretched the internal tissue architecture, leading the final histological images far from analysable and consistent with real brain tissue. Consequently, the elastic step not applied for Patient 854.

6.5.6 Direct Histological Stain Matching

Applying the rigid pipeline directly to chemically stained slices (HE, KPAS, TBBu) resulted in systematic algorithmic failures due to the stainings and the deformations that occurred. The Mutual Information metric consistently failed to identify the correct spatial planes. For example, on Block 11, the optimizer converged at false volumetric depths ($Z = 403.5$ for HE and $Z = 408.7$ for KPAS) instead of the true anatomical target ($Z = 379.8$). This wrong slice finding was replicated on Block 19 (converging at $Z = 240.1$ instead of $Z = 256.6$), accompanied by large plane rotational and translational errors in all tested cases.

Chapter 7

Discussion and Perspectives

7.1 Methodological Achievements and Pipeline Robustness

The primary ambition of this master's thesis was to engineer an automated, reproducible, and computationally efficient 3D-to-2D multimodal registration pipeline capable of aligning high-resolution *ex vivo* 7T MRI volumes with thin, distorted neuropathological sections. Achieving this spatial mapping is important to provide researchers with a reliable bridge between macroscopic imaging signatures and microscopic cellular features in multiple sclerosis research. To demonstrate the feasibility of such a pipeline, I had the rare opportunity to access a unique, high-resolution *ex vivo* dataset which, even if restricted to only two patients, allowed for the development of this multimodal co-registration pipeline.

First, by masking the data, the developed pipeline successfully met its primary requirement. By designing an adaptive skull-stripping algorithm, the framework demonstrated a high capacity to isolate complex cerebral geometry across different datasets, effectively accommodating prominent anatomical variations. Especially for Patient 784 where the background noise has effectively been cancelled, the large sulci and ventricles also have been successfully masked out. Moreover, the masking pipeline inadvertently created holes within the cerebral tissue of Patient 854. Although the N4ITK pipeline effectively corrects the global transmit field bias, hyper-intensities persist around the cortical surface and the embedding cage, causing the threshold-based algorithm to misidentify deep brain tissue as background. While the overall signal is successfully homogenized, the automated masking of the brain and its ventricles remains non optimal. This failure occurs because the algorithm classifies the bright peripheral artifacts as brain tissue due to their elevated pixel values. Furthermore, the external cage exacerbates this issue particularly for Patient 784 by introducing smooth signal variations around the brain.

Furthermore, during the developmental phase of the pipeline, an alternative optimization approach utilizing the Dice Similarity Coefficient (DSC) as the primary registration metric was investigated. The premise was to drive the spatial alignment exclusively by maximizing the geometric overlap between the binary tissue masks of both modalities. However, this

approach yielded inconclusive and highly inaccurate matching results.

Because the anatomy of the cortex shows similarities along the longitudinal axis, multiple distinct MRI slices possess nearly identical morphological brain contours. By evaluating only the binary silhouettes and entirely ignoring the internal intensity distributions, the DICE-based optimizer frequently converged on incorrect anatomical planes. This empirical failure definitively confirmed that relying exclusively on binary mask overlap is insufficient for complex multimodal slice-to-volume registration. This matching also highly relies on the initial brain masking, since the sulci are more masked in the MRI scan data and not as much as are masked for the histological sections, this brings a lack of correlation between modalities.

The histological masking shows excellent results, providing a very accurate background cancelling but also masking the large sulci. The pathological zones as in Figure 6.6b in the upper part can be masked which brings a lot of important information for matching, this zone is a key feature that brings an important differentiation with the other slices.

This robust geometric initialization directly helps to maximize the outcomes of the hybrid cost function which simultaneously maximizes intensity-based Mutual Information and aligns gradient tensors that maximise the superposition of slices. Consequently, even when the histological sections which are not cut consistently, the multi-resolution affine architecture successfully recovered the true tissue morphology to isolate it. This achievement underscores the efficiency of this pipeline to isolate both the tissue in both modalities but also to cancel the noise of the background.

A notable added value of the developed pipeline lies in the explicit extraction of the azimuthal spatial orientation parameters that mathematically define the matched 2D cutting plane within the 3D MRI volume. In the current workflow, these numerical coordinates are not manipulated, as the optimal matching slice is automatically reconstructed and visually rendered upon convergence. However, systematically recording these angular coordinates presents a great advantage for long-term data reproducibility. By setting the exact spatial coordinates of the transformation matrix, researchers have access to a standardized spatial reference for every histological section. This feature allows them to instantly retrieve, reconstruct, or re-evaluate the exact same anatomical plane in future studies without the need to run the computationally expensive algorithm again, this can help creating a reliable spatial database of the processed pathological tissues.

Ultimately, these outcomes demonstrate that while the morphological silhouettes are well and consistently aligned, the pipeline's core strength lies in its structural standardization. Because the framework directly binds the different imaging modalities through their shape and macroscopical features, attempting to inject different kinds of contrasts or raw tissue contrasts directly into the primary intensity matching process could destabilize the optimization pipeline. Without all the pre-processing implemented in this workflow, introducing divergent contrast mechanisms would inevitably corrupt the statistical distributions within

the joint histogram, invariably inducing severe false matchings.

7.2 Impact of the Bias Correction on Mutual Information

A major counter-intuitive result observed during the pre-processing phase lies in the impact of the N4ITK bias field correction on the registration accuracy. Restoring intensity homogeneity by filtering B_1 transmit field variations seemed here a fundamental prerequisite given the inhomogeneities in the scans. While the N4 pipeline successfully fulfilled this visual objective effectively flattening global intensity artifacts and making the *ex vivo* 7T volumes more uniform for human inspection. It systematically degraded or completely compromised the performance of the automated affine optimizer, particularly for artifact-heavy datasets like Patient 854.

This phenomenon can be mathematically explained by the operational mechanics of the hybrid cost function used in this pipeline, which heavily relies on Mutual Information (MI) and structural intensity gradients. Mutual Information does not evaluate absolute pixel intensities: instead, it optimizes the statistical co-occurrence of intensities between the two modalities by maximizing their joint histogram. When the N4 algorithm homogenizes the 7T MRI volume, it also flattens the local contrast variations between adjacent tissue structures, in particular at the boundary between the cortical gray matter (GM) and the corresponding white matter (WM). Indeed, by smoothing out these intensity variations to retrieve a global uniformity, the N4 filter also altered the underlying probability distribution of the MRI voxels. In consequence, the distinct statistical features that the MI metric uses to establish cross-modal correspondences were also altered. This intensity flattening drastically reduced the magnitude of internal structural gradient tensors. Reducing the number of structural landmarks and distinct statistical boundaries to guide the spatial exploration, the gradient descent search space became unstable, often trapping the transformation matrices in false local minima.

This finding highlights an important trade-off in multimodal image: pre-processing steps optimized for visual enhancement can alter the registration metrics. This issue justifies the methodological decision implemented in the final pipeline architecture, where the raw, masked 7T MRI data preserving tissue gradients and local contrast was preferred over the homogenized outputs to ensure a robust and physically accurate global affine registration.

7.3 Hyperparameters Sensitivity

Since the registration relies on raw data rather than visually smoothed images, the performance is heavily dependent on the chosen cost function. To ensure this function correctly interprets the multimodal contrast, a sensitivity analysis was performed on the hyperparameter weighting.

7.3.1 Metric Weighting (α)


To empirically validate the necessity of the hybrid cost function introduced in the methodology, a sensitivity analysis was performed on the weighting parameter α . Table 7.1 summarizes the registration outcomes, in particular the final objective score and the converged Z-slice for four different histological blocks across three configurations: a purely gradient-driven approach ($\alpha = 0.0$), the optimized hybrid approach, and a purely intensity-driven approach ($\alpha = 1.0$). It is visually confirmed by Figure 7.1 and Figure 7.2, optimizing at these mathematical extremes brings different types of spatial misalignments. When driven exclusively by raw intensity ($\alpha = 1.0$), the optimizer captures the global macroscopic mass of the brain but critically lacks boundary precision drifting the algorithm toward registration errors (particularly visible in Patient 854). Conversely, relying only on structural derivatives ($\alpha = 0.0$) forces the optimizer to anchor into high-frequency gradients. The lack of the global statistical stability provided by the raw intensities brings the gradient-only approach to become attracted by the masking noise. These visual assessments definitively confirm that a balanced hybrid approach is necessary to achieve the best match between different modalities.

These results highlight an other mathematical and qualitative trade-off governing the optimization landscape. From a strictly numerical perspective, the absolute Mutual Information score systematically decreases as the metric relies more heavily on structural gradients (i.e., when α is closer 0.0). Because raw intensity distributions ($\alpha = 1.0$) consist of broad, homogeneous statistical regions, their joint histogram is dense and mathematically easier to maximize, leading to a higher MI scores. On the other side, spatial gradients represent high-frequency edge data; their statistical co-occurrence is much stricter, which naturally reduces the MI score returned by the algorithm.


However, this statistical divergence does not always means that it is correlate with the qualitative registration accuracy. It means that even if the gradient-driven metric produces statistically lower scores, it is qualitatively indispensable. By rigidly anchoring the optimization to the sharp structural boundaries of the cortex and ventricles, the hybrid configuration successfully recovers the true anatomical slice. The optimal α is therefore a dataset-dependent compromise. It also means that the pipeline's parameters cannot be set automatically but will still need a human visual confirmation. Even if the MI score is a good indication of the matching quality, the slices matched with comparable alpha are well ranked by it but depending on the choice of alpha, it should not be the final assessment method, the visual assessment must keep this place.

Dataset	$\alpha = 0.0$ (Gradient Only)	Optimal Hybrid α	$\alpha = 1.0$ (Intensity Only)
784 – Block 11	Score: [0.1217] Z = 397.5	Score: [0.2512] Z = 379.8 ($\alpha = 0.7$)	Score: [0.2799] Z = 382.8
784 – Block 18	Score: [0.2058] Z = 286.6	Score: [0.2890] Z = 284.4 ($\alpha = 0.8$)	Score: [0.3044] Z = 291.0
784 – Block 19	Score: [0.1215] Z = 248.5	Score: [0.1411] Z = 256.6 ($\alpha = 0.2$)	Score: [0.2649] Z = 252.4
854 – Block 20	Score: [0.1303] Z = 359.9	Score: [0.2722] Z = 378.1 ($\alpha = 0.6$)	Score: [0.2077] Z = 385.3
854 – Block 24	Score: [0.0999] Z = [272.0]	Score: [0.1836] Z = 378.1 ($\alpha = 0.9$)	Score: [0.1983] Z = [281.3]

Table 7.1: Sensitivity analysis of the α hyperparameter. The table compares the final converged Z-slice and the corresponding metric score. While various configurations were evaluated, setting $\alpha = 1.0$ generally yielded a higher MI score.



(a) Block 11, $\alpha = 0.0$



(b) Block 11, $\alpha = 1.0$

Figure 7.1: Visual assessment of the α effect on the matching pipeline of Patient 784




Figure 7.2: Visual assessment of the alpha effect on the matching pipeline of Patient 854

7.3.2 Sampling points

An additional important hyperparameter for the pipeline's efficiency is the total number of sampled points extracted from the 2D histological slice. A computational trade-off must be managed during this selection. If the number of sampling points is too low, the algorithm fails to capture sufficient structural information, leading to an inaccurate representation of the tissue boundaries and a convergence toward a local maximum. Conversely, increasing the number of points beyond a certain threshold does not yield any significant improvement in the final registration accuracy, as the macroscopic structural information is already fully represented. However, this over-sampling also increases the computational time during the iterative spatial mapping and trilinear interpolation phases. Therefore, an empirically optimized target typically ranging between 10,000 and 30,000 points were selected to guarantee robust anatomical representation while maintaining highly efficient execution times. Using this configuration, the iterative optimization converges in approximately 5 to 15 minutes per block (evaluated on an AMD Ryzen 7 3700U processor with 16 GB of RAM), yielding highly satisfactory spatial alignments.

7.4 Limitations of the Pipeline

The proposed pipeline successfully automates the macroscopic cross-modal alignment process and it also demonstrated specific technical and methodological boundaries. These limitations, in particular, concern the internal tissue integrity and the reliance on specific contrast properties, are discussed in the following subsections.

7.4.1 Elastic corrections

The most significant limitation emerges during the final non-linear elastic registration phase. As shown in the results, driving the continuous B-spline optimization exclusively through external morphological silhouettes creates a fundamental biomechanical trade-off: perfect macroscopic contour matching is achieved at the severe expense of internal structural integrity.

Because the current framework relies entirely on binary tissue masks to compute the deformation fields, the elastic optimizer lacks internal structure information about the tissue, deep gray matter boundaries, or specific pathological lesions. Consequently, the algorithm extends the peripheral stretching required to match the MRI mask deep in the whole brain. This mathematically unconstrained internal warping compresses and distorts the raw cellular architecture of the histological sections, ultimately limiting their exploitability for fine neuropathological analysis and cell counting. If the algorithm could recognize different tissue elastic properties, they could be used in order to weight the elastic deformation.

An other limitation within the matching optimization lies in the simplistic penalty scoring system utilized by the cost function. Currently, the algorithm applies an arbitrary squared penalty inversely proportional to the overlapping of the two modalities, i.e., a larger overlap of brain tissue results in a lower penalty. While this constraint successfully prevents the slice from drifting entirely outside the brain mask during the spatial exploration, it lacks a rigorous mathematical foundation. In certain spatial configurations, this arbitrary weighting could overpower the scores of the Mutual Information metrics, inadvertently pushing the optimizer away from the true anatomical alignment to satisfy the background constraint. To guarantee a highly accurate registration, this penalty term should ideally be replaced with a mathematically robust penalization method. Such methods would not compromising the integrity of the primary intensity-based metric.

Furthermore, the pipeline exhibits a strong sensitivity to the initial conditions established during the affine registration. The 2D-to-3D cross-modal search space is non homogeneous, interpreting the matching scores to achieve perfect registration is not straightforward. If the global affine alignment is trapped in a local minimum due to histological tissue tearing, embedding artifacts or due to cut angles not found, the subsequent elas-

tic step cannot recover the true anatomical plane. Instead, the B-spline grid will warp the misaligned tissue to fit an incorrect target space, generating transformations that are mathematically optimal but biologically meaningless. This limitation emphasizes that the elastic deformations should only be applied to slices that have already achieved a sufficiently accurate rigid superposition.

7.4.2 MRI data masking

Another fundamental limitation of the proposed framework stems from the inherent quality of the raw 7T *ex vivo* MRI data and the impact it will have all along the 3D automated masking process. Quantitative 7T MRI provides a ultra-high spatial resolution, which is highly beneficial for resolving fine neuropathological structures. However, it is also susceptible to B_1 transmit field inhomogeneities and magnetic susceptibility artifacts. In the context of this research, these physical constraints were a large issue for the data pre-processing. Moreover, the brain samples were placed within a physical embedding cage. This configuration introduced artificial intensity gradients and persistent hyper-intensities at the immediate periphery of the cortical tissue preventing this cage to be masked by during the scan cleaning.

These raw data artifacts have deteriorated the performance of the automated 3D skull-stripping algorithm. Since the masking pipeline relies on intensity thresholds and morphological operations to isolate the true brain tissue, it is highly vulnerable when background noise or cage artifacts share overlapping intensity profiles with the cortical matter.

The current pipeline performs reliably on standard datasets but lacks the contextual awareness required to distinguish between cerebral tissue and complex non-anatomical artifacts in 7T scans. As observed in the case of Patient 854, the reliance on global intensity thresholds caused the algorithm to erroneously classify peripheral hyper-intensities as legitimate brain tissue, resulting in a flawed mask. A more robust implementation perhaps incorporating local spatial priors or a deep learning-based segmentation model could help to mitigate this susceptibility to local intensity gradients.

Nevertheless, within the constraints of the current approach, any initial masking failure propagates downstream: because the downstream matching relies on these specific intensity gradients, the cage is effectively integrated as an artificial boundary, misleading the multi-resolution affine optimization and ultimately compromising the accuracy of the spatial mapping.

7.4.3 Matching limitations

A final operational limitation concerns the restriction of the axial search space during the slice-to-volume matching process. To mitigate the severe risk of converging toward false local minima, the algorithm does not evaluate the entire 3D MRI volume for a given histological section. Instead, the spatial exploration is intentionally constrained to a restricted sub-volume, typically spanning a few dozen slices around the estimated target region. Those slices must first be coarsely evaluated using an other imagery software. The macroscopic similarity of the human cortex across adjacent planes, combined with the mentioned intensity artifacts, leads to an unconstrained global search susceptible to initial misalignment. Consequently, spatial bounding is necessary to ensure both convergence robustness and computational efficiency. Nevertheless, this constraint relies on a priori knowledge of the histological block's approximate anatomical origin, representing a fundamental trade-off in our registration pipeline.

The current pipeline cannot be considered fully autonomous for a completely blind, whole-brain search, as it requires this initial localization to help the optimization process.

An additional physical limitation specifically affected the dataset of Patient 854. Prior to the MRI acquisition, the brain sample was sectioned, which induced a macroscopic shift between the two parts of the brain. This discontinuity breaks the global anatomical coherence of the 3D volume. Because the matching phase developed relies on global affine transformations, it was incapable of correcting this physical misalignment. This is an additional physical issue that should be accounted for further algorithms.

7.4.4 Histological data quality

Additionally, the pipeline's performance is intrinsically bounded by the physical integrity of the histological sections themselves. The mechanical process of sectioning, chemical staining, and slide mounting could have introduced several tissue deformations, such as micro-tears, local folding, or missing tissue fragments. Since the global affine optimization models the section as a continuous and linearly transformable plane, it cannot mathematically account for these localized, discontinuous physical damages. The optimizer may erroneously rotate or translate the entire slice off-axis in a mathematically attempt to compensate for missing or folded brain tissue, leading to a misaligned global target space. Beyond these localized structural damages, the embedding process itself induces systemic changes. Chemical dehydration during paraffin embedding could result in tissue shrinkage. While the affine optimizer incorporates global scaling parameters (S_x, S_y) to compensate for this size reduction, it remains limited by its linear nature and cannot resolve localized, non-uniform shrinkage. As demonstrated in Patient 854 (subsection 6.2.2), when bias field correction is insufficient, the algorithm may align the brain with incorrect intensity gradients (as shown in Figure 6.13.b). In such cases, the optimizer ends on these structural or intensity-based artifacts rather than the true anatomical features, inevitably decreasing

the registration accuracy.

7.5 Future work and perspectives

While the current pipeline establishes a solid foundation for automated 3D-to-2D cross-modal registration, the limitations identified in the previous sections should be tackled in future methodological developments.

7.5.1 Machine Learning-Driven Segmentation

To address the vulnerability of the 3D MRI masking step to ultra-high-field artifacts and physical embedding cages, future iterations of this pipeline should transition from classical intensity-based thresholding to machine-learning methodologies. By training a model specifically on *ex vivo* 7T datasets, the network could learn the high-level the global shape of the cortical structures, allowing it to ignore the artificial intensity gradients, background noise or even the cage in which the brain is placed, providing a good geometric initialization for the registration. However, this technique would need to have sufficient data in order to train it, which we know is not easy to acquire.

7.5.2 Feature-Based Elastic Registration

To resolve the issue of internal tissue distortion during the non-linear alignment phase, the B-spline optimization must be updated to incorporate internal tissue properties. Future frameworks could recognize and extract the different landmarks such as the ventricles or specific multiple sclerosis lesions and use them as anchor points during the elastic deformation aiming to keep their original morphology. Additionally, by assigning specific elasticity coefficients to those different regions (like gray or white matter) depending on their properties, the algorithm would prevent the unconstrained internal warping that currently degrades the histological quality.

7.5.3 Stain adaptation of the pipeline

Finally, another area for future improvement concerns the pipeline's adaptability to varying microscopic histological stains. While the current algorithm demonstrates robust performance on standardized macroscopic tissue sections, true histological slices exhibit radically different contrast profiles depending on the specific chemical stain applied (e.g., myelin-specific versus cellular stains). Certain stains selectively highlight distinct biological structures while causing a complete loss of signal in other regions that are otherwise highly visible in the MRI volumes.

This non-linear contrast discrepancy brings bias to the Mutual Information metric. When the tissue representation changes, the statistical co-occurrences within the joint histogram break down. The optimizer struggles to find meaningful intensity correlations, frequently

driving the algorithm toward a complete spatial misalignment.

To overcome this challenge, future frameworks must integrate stain-invariant matching strategies. By artificially creating an MRI-like contrast directly from the specifically stained histological slice prior to the optimization phase, the type of pipeline could overcome the loss of intensity of certain zones, neutralizing the bias introduced by the chemical staining process.

Chapter 8

Conclusion

This master’s thesis aimed to address a major challenge in multiple sclerosis research: bridging the dimensional and morphological gap between whole-brain 7T ex vivo MRI and microscopic histological sections. This was achieved by designing a custom co-registration pipeline that progresses from the preprocessing of raw images to 3D affine alignment, ending with 2D elastic non-linear warping.

Moreover, this work serves as promising proof of concept. The resulting matching pipeline is intuitive, user-friendly, and relatively easy to handle, providing a promising mathematical framework to deal with the complexities of multimodal alignment. By successfully retrieving macroscopic cutting planes and matching external cortical boundaries on optimal datasets, it demonstrates the viability of the proposed hybrid objective function and the overall registration architecture.

However, the current framework is still far from perfect and requires future improvements before it can be deployed on a larger scale. While the pipeline performs well on high-quality data, it currently lacks the necessary robustness to handle lower-quality, noisy, or artifacted scans. A major limitation lies in the inconsistency of the pipeline’s behavior across different datasets. At this stage, the process is not fully automated, several steps still require parameters tuning depending on the specific input images. Most notably, the thresholding parameters for the masking processes and the adjustment of the α weight in the hybrid metric demand manual tuning to achieve optimal results.

Moving forward, future research should focus on minimizing this need for user intervention to standardize the process across diverse datasets. Integrating more robust methods or other approaches such as deep learning algorithms (if the data available is enough) for automated, intensity-independent skull-stripping and tissue segmentation could resolve the current masking inconsistencies. By addressing these limitations and improving the regularization of internal elastic deformations, this pipeline has the potential to become a highly reliable and fully automated tool, further paving the way for the in vivo translation of ultra-high-field MRI signatures.

AI Usage Declaration

During the preparation of this thesis, I used generative artificial intelligence tools to assist in specific technical tasks. The use of these tools was strictly limited to the following applications:

- **Language Editing:** To correct the manuscript's language, correct grammatical errors, and refine the fluency of the English phrasing.
- **Code Assistance:** To assist in debugging the Python scripts developed for the registration pipeline and to resolve programmatic issues.
- **Formatting and visuals:** To generate the LaTeX syntax for complex tables and assist in structuring graphics, ensuring a clear and professional presentation of the results without contributing to the actual writing or interpretation of those results.

These tools were used exclusively as technical assistants. They were not used to generate scientific concepts, formulate the research methodology, or interpret the core data and results. I meticulously reviewed, edited, and validated all AI-assisted outputs and take full responsibility for the final content, findings, and conclusions presented in this manuscript.

- [1] R. Nistri, A. Ianniello, V. Pozzilli, C. Gianni, and C. Pozzilli, "Advanced MRI Techniques: Diagnosis and Follow-Up of Multiple Sclerosis," en, *Diagnostics*, vol. 14, no. 11, p. 1120, May 2024, issn: 2075-4418. doi: 10.3390/diagnostics14111120.
- [2] J. Annese, "The importance of combining MRI and large-scale digital histology in neuroimaging studies of brain connectivity and disease," en, *Frontiers in Neuroinformatics*, vol. 6, 2012, issn: 1662-5196. doi: 10.3389/fninf.2012.00013.
- [3] R. Galbusera, M. Weigel, E. Bahn, S. Schaedelin, A. Cagol, P.-J. Lu, M. Barakovic, L. Melie-Garcia, J. Franz, P. Dechent, G. Nair, L. Kappos, W. Brück, C. Stadelmann, and C. Granziera, "Quantitative T1 is sensitive to cortical remyelination in multiple sclerosis: A postmortem MRI study," en, *Brain Pathology*, vol. 35, no. 5, e70010, Sep. 2025, issn: 1015-6305, 1750-3639. doi: 10.1111/bpa.70010.
- [4] S. Hametner, V. Endmayr, A. Deistung, P. Palmrich, M. Prihoda, E. Haimburger, C. Menard, X. Feng, T. Haider, M. Leisser, U. Köck, A. Kaider, R. Höftberger, S. Robinson, J. R. Reichenbach, H. Lassmann, H. Traxler, S. Trattnig, and G. Grabner, "The influence of brain iron and myelin on magnetic susceptibility and effective transverse relaxation - A biochemical and histological validation study," en, *NeuroImage*, vol. 179, pp. 117–133, Oct. 2018, issn: 10538119. doi: 10.1016/j.neuroimage.2018.06.007.
- [5] V. Wiggermann, V. Endmayr, E. Hernández-Torres, R. Höftberger, G. Kasprian, S. Hametner, and A. Rauscher, "Quantitative magnetic resonance imaging reflects different levels of histologically determined myelin densities in multiple sclerosis, including remyelination in inactive multiple sclerosis lesions," en, *Brain Pathology*, vol. 33, no. 6, e13150, Nov. 2023, issn: 1015-6305, 1750-3639. doi: 10.1111/bpa.13150.
- [6] C. Stadelmann, M. Albert, C. Wegner, and W. Brück, "Cortical pathology in multiple sclerosis," en, *Current Opinion in Neurology*, vol. 21, no. 3, pp. 229–234, Jun. 2008, issn: 1350-7540. doi: 10.1097/01.wco.0000318863.65635.9a.
- [7] V. Barletta, E. Herranz, C. A. Treaba, A. Mehndiratta, R. Ouellette, G. Mangeat, T. Granberg, J. A. Sloane, E. C. Klawiter, J. Cohen-Adad, and C. Mainero, "Quantitative 7-Tesla Imaging of Cortical Myelin Changes in Early Multiple Sclerosis," en, *Frontiers in Neurology*, vol. 12, p. 714820, Sep. 2021, issn: 1664-2295. doi: 10.3389/fneur.2021.714820.
- [8] A. Alkemade, P.-L. Bazin, E. Kirilina, K. Pine, A. Herrler, R. L. A. W. Bleijs, J. M. Kros, L. Groenewegen, S. M. M. Vermorgen, L. E. Jonkman, D. F. Swaab, R. Balesar, N. Weiskopf, and B. U. Forstmann, "A unified 3D reconstruction of microscopy

- and MRI in a brain showing Alzheimer's disease-related neuropathology," en, *Brain Pathology*, e70039, Sep. 2025, issn: 1015-6305, 1750-3639. doi: 10.1111/bpa.70039.
- [9] N. Weiskopf, S. Mohammadi, A. Lutti, and M. F. Callaghan, "Advances in MRI-based computational neuroanatomy: From morphometry to in-vivo histology," en, *Current Opinion in Neurology*, vol. 28, no. 4, pp. 313–322, Aug. 2015, issn: 1350-7540. doi: 10.1097/WCO.0000000000000222.
- [10] N. Weiskopf, J. Suckling, G. Williams, M. M. Correia, B. Inkster, R. Tait, C. Ooi, E. T. Bullmore, and A. Lutti, "Quantitative multi-parameter mapping of R1, PD*, MT, and R2* at 3T: A multi-center validation," en, *Frontiers in Neuroscience*, vol. 7, 2013, issn: 1662-453X. doi: 10.3389/fnins.2013.00095.
- [11] J. Wang, L. He, H. Zheng, and Z.-L. Lu, "Improving structural brain images acquired with the 3D FLASH sequence," en, *Magnetic Resonance Imaging*, vol. 38, pp. 224–232, May 2017, issn: 0730725X. doi: 10.1016/j.mri.2017.01.014.
- [12] K. Tabelow, E. Balteau, J. Ashburner, M. F. Callaghan, B. Draganski, G. Helms, F. Kherif, T. Leutritz, A. Lutti, C. Phillips, E. Reimer, L. Ruthotto, M. Seif, N. Weiskopf, G. Ziegler, and S. Mohammadi, "hMRI – A toolbox for quantitative MRI in neuroscience and clinical research," en, *NeuroImage*, vol. 194, pp. 191–210, Jul. 2019, issn: 10538119. doi: 10.1016/j.neuroimage.2019.01.029.
- [13] A. Kutzelnigg, C. F. Lucchinetti, C. Stadelmann, W. Brück, H. Rauschka, M. Bergmann, M. Schmidbauer, J. E. Parisi, and H. Lassmann, "Cortical demyelination and diffuse white matter injury in multiple sclerosis," en, *Brain*, vol. 128, no. 11, pp. 2705–2712, Nov. 2005, issn: 1460-2156, 0006-8950. doi: 10.1093/brain/awh641.
- [14] L. Luijckink, M. Rodriguez, and R. Machaalani, "Quantifying GFAP immunohistochemistry in the brain – Introduction of the Reactivity score (R-score) and how it compares to other methodologies," en, *Journal of Neuroscience Methods*, vol. 402, p. 110 025, Feb. 2024, issn: 01650270. doi: 10.1016/j.jneumeth.2023.110025.
- [15] N. D. Scharf, S. A. Herr, L. Madsen, S. J. Butts, C. Torres, L. B. Mendez, and A. L. Brewster, "Spatiotemporal profile of Map2 and microglial changes in the hippocampal CA1 region following pilocarpine-induced status epilepticus," en, *Scientific Reports*, vol. 6, no. 1, p. 24 988, May 2016, issn: 2045-2322. doi: 10.1038/srep24988.
- [16] S. Hametner, I. Wimmer, L. Haider, S. Pfeifenbring, W. Brück, and H. Lassmann, "Iron and neurodegeneration in the multiple sclerosis brain," en, *Annals of Neurology*,

- vol. 74, no. 6, pp. 848–861, Dec. 2013, issn: 0364-5134, 1531-8249. doi: 10.1002/ana.23974.
- [17] M. K. Fard, F. Van Der Meer, P. Sánchez, L. Cantuti-Castelvetri, S. Mandad, S. Jäkel, E. F. Fornasiero, S. Schmitt, M. Ehrlich, L. Starost, T. Kuhlmann, C. Sergiou, V. Schultz, C. Wrzos, W. Brück, H. Urlaub, L. Dimou, C. Stadelmann, and M. Simons, “BCAS1 expression defines a population of early myelinating oligodendrocytes in multiple sclerosis lesions,” en, *Science Translational Medicine*, vol. 9, no. 419, eaam7816, Dec. 2017, issn: 1946-6234, 1946-6242. doi: 10.1126/scitranslmed.aam7816.
- [18] Y. Wang, W. Ho, I. N. Huszar, P. DiGiacomo, H. M. Taghavi, L. Tao, M. Choi, N. Nguyen, S. Leventis, D. B. Camarillo, P. Schlömer, M. Axer, W. Shao, M. Rusu, I. Cobos, J. Nirschl, M. Georgiadis, and M. Zeineh, “Precise MRI-histology coregistration of paraffin-embedded tissue with blockface imaging,” en, *Imaging Neuroscience*, vol. 3, IMAG.a.106, Aug. 2025, issn: 2837-6056. doi: 10.1162/IMAG.a.106.
- [19] I. N. Huszar, M. Pallegage-Gamarallage, S. Bangerter-Christensen, H. Brooks, S. Fitzgibbon, S. Foxley, M. Hiemstra, A. F. Howard, S. Jbabdi, D. Z. Kor, A. Leonte, J. Mollink, A. Smart, B. C. Tendler, M. R. Turner, O. Ansorge, K. L. Miller, and M. Jenkinson, “Tensor image registration library: Deformable registration of stand-alone histology images to whole-brain post-mortem MRI data,” en, *NeuroImage*, vol. 265, p. 119 792, Jan. 2023, issn: 10538119. doi: 10.1016/j.neuroimage.2022.119792.
- [20] M. Jung, B. Bogner, T. D. Diallo, S. Kim, P. Arnold, H. Füllgraf, K. Kurowski, P. Bronsert, P. M. Jungmann, J. Kiefer, D. Kraus, P. Rovedo, M. Reisert, S. U. Eisenhardt, F. Bamberg, M. Benndorf, and A. Runkel, “Multiparametric magnetic resonance imaging for radiation therapy response monitoring in soft tissue sarcomas: A histology and MRI co-registration algorithm,” en, *Theranostics*, vol. 13, no. 5, pp. 1594–1606, 2023, issn: 1838-7640. doi: 10.7150/thno.81938.
- [21] M. Absinta, G. Nair, M. Filippi, A. Ray-Chaudhury, M. I. Reyes-Mantilla, C. A. Pardo, and D. S. Reich, “Postmortem Magnetic Resonance Imaging to Guide the Pathologic Cut: Individualized, 3-Dimensionally Printed Cutting Boxes for Fixed Brains,” en, *Journal of Neuropathology & Experimental Neurology*, vol. 73, no. 8, pp. 780–788, Aug. 2014, issn: 0022-3069, 1554-6578. doi: 10.1097/NEN.0000000000000096.
- [22] D. Lowe, “Object recognition from local scale-invariant features,” en, in *Proceedings of the Seventh IEEE International Conference on Computer Vision*, Kerkyra, Greece: IEEE, 1999, 1150–1157 vol.2, isbn: 978-0-7695-0164-2. doi: 10.1109/ICCV.1999.790410.

- [23] D. G. Lowe, "Distinctive Image Features from Scale-Invariant Keypoints," en, *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, Nov. 2004, issn: 0920-5691, 1573-1405. doi: 10.1023/B:VISI.0000029664.99615.94.
- [24] B. Rister, M. A. Horowitz, and D. L. Rubin, "Volumetric Image Registration From Invariant Keypoints," en, *IEEE Transactions on Image Processing*, vol. 26, no. 10, pp. 4900–4910, Oct. 2017, issn: 1057-7149, 1941-0042. doi: 10.1109/TIP.2017.2722689.
- [25] P. Viola and W. Wells, "Alignment by Maximization of Mutual Information," en, *International journal of computer vision*, vol. 24, pp. 137–154, Sep. 1997. doi: <https://doi.org/10.1023/A:1007958904918>.
- [26] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multimodality Image Registration by Maximization of Mutual Information," en, *IEEE Transactions on Medical Imaging*, vol. 16, no. 2, pp. 187–198, Apr. 1997. doi: <https://doi.org/10.1109/42.563664>.
- [27] C. R. Meyer, B. A. Moffat, K. K. Kuszpit, P. L. Bland, P. E. Mckeever, T. D. Johnson, T. L. Chenevert, A. Rehemtulla, and B. D. Ross, "A Methodology for Registration of a Histological Slide and In Vivo MRI Volume Based on Optimizing Mutual Information," en, *Molecular Imaging*, vol. 5, no. 1, p. 7290.2006.00002, Jan. 2006, issn: 1535-3508, 1536-0121. doi: 10.2310/7290.2006.00002.
- [28] D. Aquino, R. Garbelli, L. Rossini, D. De Santis, R. Spreafico, P. d'Orio, L. Tassi, and F. Padelli, "A semi-automatic registration protocol to match ex-vivo high-field 7T MR images and histological slices in surgical samples from patients with drug-resistant epilepsy," en, *Journal of Neuroscience Methods*, vol. 367, p. 109439, Feb. 2022, issn: 01650270. doi: 10.1016/j.jneumeth.2021.109439.
- [29] P.-F. Van De Moortele, C. Akgun, G. Adriany, S. Moeller, J. Ritter, C. M. Collins, M. B. Smith, J. T. Vaughan, and K. Uğurbil, " B_1 destructive interferences and spatial phase patterns at 7 T with a head transceiver array coil," en, *Magnetic Resonance in Medicine*, vol. 54, no. 6, pp. 1503–1518, Dec. 2005, issn: 0740-3194, 1522-2594. doi: 10.1002/mrm.20708.
- [30] S. Roy and P. Maji, "A simple skull stripping algorithm for brain MRI," en, in *2015 Eighth International Conference on Advances in Pattern Recognition (ICAPR)*, Kolkata, India: IEEE, Jan. 2015, pp. 1–6, isbn: 978-1-4799-7458-0. doi: 10.1109/ICAPR.2015.7050671.

- [31] N. J. Tustison, B. B. Avants, P. A. Cook, Yuanjie Zheng, A. Egan, P. A. Yushkevich, and J. C. Gee, "N4ITK: Improved N3 Bias Correction," en, *IEEE Transactions on Medical Imaging*, vol. 29, no. 6, pp. 1310–1320, Jun. 2010, issn: 0278-0062, 1558-254X. doi: 10.1109/TMI.2010.2046908.
- [32] N. J. Tustison and J. Gee, "N4ITK: Nick's N3 ITK Implementation For MRI Bias Field Correction," en, *The Insight Journal*, Dec. 2010, issn: 2327-770X. doi: 10.54294/jculxw.
- [33] F. Bookstein, "Principal warps: Thin-plate splines and the decomposition of deformations," en, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 6, pp. 567–585, Jun. 1989, issn: 01628828. doi: 10.1109/34.24792.
- [34] Y. Liu, B. R. Sajja, M. G. Uberti, H. E. Gendelman, T. Kielian, and M. D. Boska, "Landmark Optimization Using Local Curvature for Point-Based Nonlinear Rodent Brain Image Registration," en, *International Journal of Biomedical Imaging*, vol. 2012, pp. 1–8, 2012, issn: 1687-4188, 1687-4196. doi: 10.1155/2012/635207.