

École polytechnique de Louvain

Location-aware iEEG signal analysis for epileptic patients

Author: **Maëlys DEGAND**

Supervisors: **Ron KUPERS, Benoît MACQ**

Readers: **Tim BARY, Nicolas DELINTE, Enrique GERMANY**

Academic year 2025–2026

Master [120] in Biomedical Engineering

Abstract

Epilepsy is a chronic neurological disorder affecting more than 50 million people worldwide, characterized by recurrent seizures resulting from abnormal brain electrical activity. For drug-resistant patients, surgical treatment may be considered, requiring precise localization of the epileptogenic zone. [1] In this context, intracranial electroencephalography (iEEG) provides direct, high-fidelity recordings of brain activity and constitutes a key tool for both clinical decision-making and neuroscientific research. However, the preprocessing of iEEG data remains a challenging and poorly standardized task, typically requiring significant domain expertise and manual intervention. [2] [3] [4]

This thesis presents a semi-automated pipeline for location-aware iEEG signal preprocessing in epileptic patients, integrating two prior works developed at UCLouvain [5], [6]. The proposed pipeline covers the full preprocessing workflow: data preprocessing, electrode identification, electrode localization using the ElectroLoc algorithm [6], anatomical contact labeling via an atlas, and epileptic seizure detection using a transformer-based model (MViT) [5]. A shift correction mechanism is introduced to handle mislocalized electrodes, and atlas dilatation is applied to improve anatomical labeling accuracy.

The pipeline is validated on data from two patients implanted with iEEG depth electrodes. Results demonstrate competitive performance across all modules, with particular strengths in electrode localization and contact labeling. Limitations related to dataset size, shift correction conditions, and seizure detection generalization are discussed, along with perspectives for future improvements including few-shot learning approaches for seizure detection.

The code produced for this thesis can be found in the following repository:

<https://github.com/maedegand/MasterThesis.git>

Keywords—intracranial electroencephalography, stereo electroencephalography, epileptic seizure detection, electrode localization, semi-automated pipeline, anatomical atlas, image co-registration, transformer neural network.

Acknowledgements

I am deeply grateful to **Benoit Macq** for giving me the opportunity to work on this master's thesis and discover the field of neurosciences. Through this project, I developed a strong interest in neurosciences and in their applications to engineering.

I would like to sincerely thank **Ron Kupers** for his interest in the project, for providing access to the clinical data used in this work and for sharing his clinical expertise throughout this work.

I would also like to sincerely thank **Nicolas Delinte** for his constant support, from the earliest stages of this master's thesis to its very end. His availability, feedback, suggestion and expertise constituted valuable help in the project.

I am also grateful to **Tim Bary** for his help with the seizure detection part of this work. His explanations and guidance were particularly valuable in helping me understand and build upon his previous work.

I would also like to thank **my parents** for their constant support throughout this work. I am especially grateful to my father, whose careful proofreading and thoughtful advice were particularly valuable in improving the quality of this thesis.

Finally, I would like to spare a warm thought for **my friends** who have been an important source of support and motivation throughout my studies. The many co-working sessions, shared breaks and moments of mutual encouragements made this journey far more enjoyable.

I acknowledge the use of AI tools (specially ChatGPT) to assist with the academic formulation of my ideas and to help with code documentation.

Contents

Abstract	3
Acknowledgements	5
List of Figures	11
List of Tables	12
List of Abbreviations	13
Introduction	17
I State of the Art	19
1 Epilepsy	21
1.1 Overview of Epilepsy	21
1.2 Diagnosis of Epilepsy	22
1.3 Treatment of Epilepsy	23
2 Medical Imaging in context of Epilepsy	24
2.1 Computed Tomography	24
2.2 Magnetic Resonance Imaging	25
3 Intracranial Electroencephalography	27
3.1 Basics of iEEG	27
3.2 iEEG for epileptogenic zone detection	29
4 Methodological Pipelines in Intracranial EEG Research	32
4.1 Electrode Localization	32
4.2 Electrode Identification	35
4.3 iEEG signal preprocessing	36
4.4 Epileptic Seizure Detection	37
II Pipeline Semi-Automated for Pre-Processing in iEEG Research	41
5 Materials	43
5.1 Data Acquisition	43

5.2	Preprocessing	44
5.2.1	File format	44
5.2.2	Skull stripping and brain mask extraction	44
5.2.3	Discrepancy observed in Patient 2	45
6	Methods	46
6.1	Data Preprocessing	46
6.1.1	iEEG signal preprocessing	46
6.1.2	Images preprocessing	47
6.1.3	Summary of Data preprocessing module	48
6.2	Electrode Identification	49
6.3	Electrode Localization	49
6.4	Contact Labeling	50
6.4.1	Post-processing of ElectroLoc	50
6.4.2	Atlas dilatation	53
6.4.3	Anatomical atlas labeling	54
6.4.4	Summary of contact labeling module	55
6.5	Epileptic seizure detection	56
6.5.1	Pre-processing of iEEG signal for MViT	56
6.5.2	Prediction for iEEG segments	59
6.5.3	Summary of epileptic seizure detection module	60
6.6	Algorithm summary	61
6.7	Performance Assessment	61
7	Results & Discussions	64
7.1	Performance of the Algorithm	64
7.1.1	Performance of data preprocessing	64
7.1.2	Performance of electrode localization	66
7.1.3	Performance of post-processing of ElectroLoc	68
7.1.4	Performance of contact labeling	71
7.1.5	Performance of epileptic seizure detection	75
7.1.6	Performance of overall algorithm	76
7.2	Comparison to Existing Methods	77
7.3	Limitations	78
7.3.1	Validation	78
7.3.2	Condition for shifted electrode selection	79
7.3.3	Epileptic seizure detection	79
8	Perspectives and potential improvements	81
8.1	Perspectives for contact extraction	81
8.1.1	Electrode identification	81
8.1.2	Electrode localization	81
8.1.3	Shift correction	82
8.1.4	Co-registration	82
8.1.5	Graphical User Interface	82
8.2	Perspectives for epileptic seizure detection	83
	Conclusion	85

Bibliography	89
A Condition for shift electrode	90
B Code and Repository Overview	92
B.1 Structure of the repository	92
C Few-shot Learning approaches for iEEG seizure detection	94

List of Figures

1.1	Graphical representation of the framework to classify epilepsy	22
2.1	Drawing of CT fan beam and a patient in a CT imaging system. . . .	25
2.2	Coronal slice of brain observed with T1-weighted and T2-weighted . .	26
3.1	Intraoperative photograph showing subdural grid electrodes and depth electrodes.	28
3.2	Example of intracranial EEG signal	29
3.3	Example of interictal biomarkers	30
3.4	Example of epilepsy biomarkers	31
4.1	Flowchart of Brainstorm software for electrode localization.	33
4.2	Steps for localization of electrode's contacts in Brainstorm	33
4.3	Summary of the overall algorithm of ElectroLoc	33
4.4	System 10-20 of conventional scalp EEG	35
4.5	Standardized nomenclature proposed by Stone et al. (2021)	35
4.6	AAL atlas views	36
4.7	Architecture of a MViT	38
4.8	Schematic overview of the model of T. Bary's master's thesis	39
4.9	Scalogram of EEG signal from training phase of T. Bary's model. . .	39
4.10	Morlet mother wavelet	40
5.1	A slice of Patient 1's MRI	45
5.2	A slice of Patient 1's CT	45
6.1	Summary of the proposed algorithm.	46
6.2	Electrode mask and brain's surface mask from Patient 1	47
6.3	Visualization of entry point	48
6.4	Summary of the data preprocessing module.	48
6.5	Visualization of ElectroLoc's output	50
6.6	Visualization of shifted electrode	51
6.7	Visualization of AAL atlases dilated and undilated	53
6.8	Summary of the contact labeling module	55
6.9	Illustration of preprocessing steps for epilepsy detection module . . .	57
6.10	Illustration of patches from scalogram LB13-LB14 from Patient 2. . .	59
6.11	Summary of epileptic seizure detection	60
7.1	Undetected electrodes from Patient 2	65
7.2	Validation of ElectroLoc contact localization.	66
7.3	Visualization of performance of electrode localization	67

7.4	Visualization of 3 False Negative	68
7.5	Comparison of reindexation approaches	69
7.6	Impact of shift correction on contacts	70
7.7	Shift correction from Patient 1.	71
7.8	Coregistration results for Patient 1.	72
7.9	Impact of atlas dilatation on contact labeling of Patient 1	73
7.10	Anatomical labeling along each electrode for both Patients.	74
7.11	Visualization of electrodes with abrupt change in labeling	75
7.12	Predicted seizure probabilities for 5-minute recordings.	76
7.13	Distribution of predicted seizure probabilities for 5-minute recordings.	76
7.14	Shift configuration not handled by the selected condition	79
8.1	Semi-supervised teacher-student adaptation strategy	84
8.2	Prototype-based adaptation strategy	84
A.1	Electrode configurations considered during the shift correction step.	90
A.2	Electrode configurations considered for shift correction.	91
B.1	Simplified structure of the GitHub repository.	92
C.1	Framework of semi-supervised FSL for epileptic seizure detection	95
C.2	Proposed adaptation of original transformer-based model	96

List of Tables

2.1	Mass attenuation coefficients for different materials and two different X-ray energies (keV) [11]	24
5.1	CT acquisition parameters for both patients.	43
5.2	MRI acquisition parameters for both patients.	44
6.1	Inputs and outputs of proposed data preprocessing module	48
6.2	Electrode information available after electrode identification	49
6.3	Parameters used for co-registration with the Python package <i>regis.core</i>	54
6.4	Inputs, outputs and hyperparameters of contact labeling module	56
6.5	Hyperparameters of the TUSZ model architecture from [5]	59
6.6	Inputs, outputs and hyperparameters of epileptic seizure detection module	60
7.1	Performance of the data preprocessing module with a comparison between the extracted entry point and the ground-truth.	64
7.2	Classification metrics for contact localization of ElectroLoc output.	67
7.3	Classification metrics for indexation output.	70
7.4	Classification metrics for shift correction output.	71
7.5	Dice Similarity Coefficients before and after CT-to-MRI and MNI-to-MRI co-registrations.	72
7.6	Time execution of contact extraction module	77
7.7	Time execution of epileptic seizure detection module	77
A.1	Electrode configurations classified as shifted according to each condition.	91

List of Abbreviations

Acronym	Definition
AAL	Automated Anatomical Labeling
AED	Anti-Epileptic Drug
BIP	Bipolar reference
CAR	Common Average Reference
CT	Computed Tomography
CWT	Continuous Wavelet Transform
DICOM	Digital Imaging and Communication in Medicine
DCS	Direct Cortical Stimulation
DSC	Dice Similarity Coefficient
EC	Eloquent Cortex
ECoG	Electrocorticography
EDF	European Data Format
EEG	Electroencephalogram
EMA	Exponential Moving Average
EZ	Epileptogenic Zone
fMRI	functional Magnetic Resonance Imaging
FSL	Few-Shot Learning
FLAIR	Fluid-Attenuated Inversion Recovery
GELU	Gaussian Error Linear Unit
HD	Hyperdimensional
HFB	High-Frequency Broadband
HFO	High-Frequency Oscillation
HU	Hounsfield Units
IED	Interictal Epileptiform Discharge
iEEG	intracranial Electroencephalogram
IZ	Irritative Zone
LAP	Laplacian reference
LBP	Local Binary Pattern
LFP	Local Field Potential
LVFA	Low-Voltage Fast Activity
MLP	Multi-Layer Perceptron
MNI	Montreal Neurological Institute
MRI	Magnetic Resonance Imaging
MViT	Multichannel Vision Transformer
NIfTI	Neuroimaging Informatics Technology Initiative
PCA	Principal Component Analysis
PET	Positron Emission Tomography

Acronym	Definition
SCA	Scalp reference
SEEG	Stereo-Electroencephalogram
SENSA	Standardized Electrode Nomenclature for SEEG Application
SOZ	Seizure Onset Zone
TIRDA	Temporal Intermittent Rhythmic Delta Activity
WM	White Matter reference

Introduction

Epilepsy is a neurological disorder affecting approximately 50 million people worldwide. It is characterized by abnormal brain activity that can lead to recurrent seizures. Beyond the seizures themselves, epilepsy can have major physical and social consequences imposing a significant burden on people living with the condition. Nevertheless, with appropriate diagnosis and treatment, up to 70% of people with epilepsy could become seizure-free. [1] The first-line treatment usually relies on antiepileptic drugs (AEDs). However, in some cases, seizures remain resistant to medication. For these drug-resistant patients, surgical treatment may be considered. Epilepsy surgery requires an accurate localization of the brain region responsible for seizure generation. [2] [3] [4]

In this context, intracranial electroencephalography (iEEG) is a powerful tool, as it provides direct recordings of the brain electrical activity with a high signal-to-noise ratio. These recordings are essential for clinical decision-making in epilepsy surgery. Beyond this clinical application, iEEG data constitute rich and rare data offering a valuable opportunity to better understand brain function and to support research on neurological disorders. Several studies are therefore conducted in parallel in the field of neurological research. However, iEEG remains an invasive examination and is therefore prescribed only in selected cases. As a result the available recordings are limited and intrinsically linked to the clinical context, here, epilepsy. These signals must therefore be carefully preprocessed before any subsequent analysis. [7]

However, the literature still lacks standardized approaches for iEEG signal preprocessing for epileptic patients. Existing methods often require many different inputs and substantial amount of human intervention which usually implies specific expertises in neuroimaging or signal processing.

The objective of this thesis is therefore twofold. First, this work aims to develop a semi-automated iEEG signal preprocessing pipeline, integrating two previous works conducted at UCLouvain (specifically [5] and [6]). The proposed pipeline seeks to provide a functional and reproducible algorithm while reducing as much as possible the number of required inputs and manual interventions. Second, this thesis aims to provide sufficient methodological and biomedical context to make the work accessible both to engineers with no prior background in the biomedical field and to medical experts interested in the technical aspects of the pipeline.

Part I

State of the Art

Chapter 1

Epilepsy

Epilepsy is one of the most common chronic brain disease. It affects more than 50 million people worldwide. According to World Health Organization, almost 70% of people with epilepsy could be seizure free if properly diagnosed and treated.[\[1\]](#)

1.1 Overview of Epilepsy

Epilepsy is a neurological disorder characterized by a long-term predisposition of the brain to generate epileptic seizures. An epileptic seizure corresponds to an abnormal, excessive and synchronous electrical activity in the brain. This abnormal firing disrupts normal brain signaling and may result in involuntary movements, unusual sensations, sudden emotion changes or altered/lost consciousness and awareness. [\[8\]](#), [\[9\]](#)

In the context of epilepsy, several spatial concepts are used to describe the brain regions involved in seizure generation and propagation: [\[3\]](#)

- Seizure onset zone (SOZ) refers to the cortical area that initiates clinical seizures.
- Epileptogenic zone (EZ) is the area of the cortex that is necessary and sufficient for initiating seizures and whose removal (or disconnection) is necessary for complete abolition of seizures.
- Epileptogenic lesions correspond to the macroscopic lesions causing the seizure detectable by MRI.
- Eloquent cortex (EC) is the area of cortex that if removed will result in loss of motor, sensory or language function.
- Irritative zone (IZ) refers to the cortical area that generates interictal spikes

Epilepsy can be characterized by three-level classification according 2017 ILAE classification: seizure type, epilepsy type and epilepsy syndrome. This is illustrated on Figure [1.1](#).




Figure 1.1: Graphical representation of the framework to classify epilepsy [10]

Two main categories of seizure type are distinguished: focal seizure, which originates from a precise area of the brain, and the generalized seizure, which involves both cerebral hemispheres of the brain. Focal seizures may occur with preserved or altered consciousness whereas generalized seizures usually provoke loss of consciousness.

The epilepsy type is supported by electroencephalography (EEG) findings. Generalized spike activity on EEG represents generalized epilepsy, while localized and unilateral epileptiform discharge with normal interictal EEG activity are indicative of focal epilepsy.

Epilepsy syndrome refers to a cluster of features including seizure type, EEG characteristics and imaging findings. It is often age dependent and may be associated with co-morbidities including intellectual disability, psychological and behavioral difficulties. [4], [10]

Epilepsy can arise from multiple etiologies, including structural causes (tumors, trauma, AVC), genetic mutations affection ion channels, infectious disease (meningitis, encephalitis), metabolic disorders, immune-related conditions. However, in approximately half of cases, the cause remains unknown. [4]

1.2 Diagnosis of Epilepsy

Epilepsy is typically diagnosed through EEG which records brain electrical activity, combined with medical imaging techniques. Magnetic resonance imaging (MRI) is used to detect structural abnormalities, such as epileptogenic lesions, while computed tomography (CT) scans are performed in emergency settings for rapid screening to exclude acute or life-threatening conditions such as hemorrhage, tumor or brain edema. More details about medical imaging used in the context of epilepsy is described in Section 2.

Accurate diagnosis is a critical step in determining the most appropriate treatment strategy and improving patient outcomes. Misdiagnosis may expose patients to unnecessary risks, including inappropriate treatments, potential adverse effects and social and psychological consequences. [4], [8]

1.3 Treatment of Epilepsy

The main objective of treatment is to achieve almost total control of seizures while minimizing side effects.

Anti-epileptic drugs (AEDs) constitute the first-line treatment. Approximately 60-70% of patients achieve satisfactory seizure control with a single medication (monotherapy). However, when epilepsy remains drug-resistant, typically defined as the failure of two attempts with high-dose monotherapy and one polytherapy, epilepsy surgery may be considered.[4], [8]

Epilepsy surgeries include several approaches depending on the nature and the localization of the epileptic seizures. Resection aims to remove the brain region responsible for epileptic discharges. Disconnection procedures, such as multiple subpial transections or corpus callosotomy, aim to interrupt the propagation of seizure between two brain regions. In addition, less invasive techniques have been developed, including laser thermal ablation to precisely destruct the zone, neurostimulation methods such as vagus nerve stimulation, deep brain stimulation to reduce seizures. [8]

The success of these interventions relies on achieving a highly precise balance between removing brain areas responsible for seizure generation, referred to as the EZ while preserving regions essential to normal cerebral functions, known as the EC. This challenge makes presurgical evaluation a critical step in the management of drug-resistant epilepsy. Although non-invasive methods such as scalp EEG, MRI and CT scans are used as first-line investigations, they may be insufficient to accurately delineate these critical regions. To complement these structural examinations, functional and metabolic imaging techniques are used to obtain advanced information. Positron emission tomography (PET) measures the cerebral glucose metabolism. It typically reveals hypometabolic regions corresponding to the EZ during interictal periods. Functional magnetic resonance imaging (fMRI) is also widely used in presurgical evaluation to identify and map eloquent cortex. When non-invasive investigations remain inconclusive, invasive monitoring becomes a necessary diagnostic step to precisely map these critical areas. Intracranial electroencephalography (iEEG) often combined with direct cortical stimulation (DCS) enables to establish surgical candidacy and delineate surgical margins. [2], [3], [4]

Chapter 2

Medical Imaging in Diagnosis and Presurgical Evaluation of Epilepsy

This chapter focuses on MRI and CT imaging, as these modalities are predominantly used in iEEG studies.

2.1 Computed Tomography

X-ray imaging relies on the differential absorption of ionizing radiation by biological tissues. As X-ray passes through the body, their intensity is attenuated depending on the density and atomic composition of the tissues encountered (see Table 2.1). The attenuation of X-ray in biological tissues follows the Beer-Lambert law, which describes an exponential decrease in intensity as a function of the material thickness and its attenuation coefficient:

$$I = I_0 e^{-\mu x} \quad (2.1)$$

where I_0 is the incident intensity, I is the transmitted intensity, μ is the linear attenuation coefficient of the material, and x is the thickness of the medium. [11]

Table 2.1: Mass attenuation coefficients for different materials and two different X-ray energies (keV) [11]

Material	20 keV (cm ² /g)	100 keV (cm ² /g)
Air	0.7779	0.1541
Water	0.8096	0.1707
Bone	0.4001	0.1855
Muscle	0.8205	0.1693
Lung	0.8316	0.1695
Brain	0.8281	0.1701

Dense structures such as bone attenuate a larger fraction of radiation and therefore appear brighter in the resulting image, whereas soft tissues attenuate less and appear darker. The level of attenuation depends also on the energy of the X-ray (expressed in keV): higher photon energies lead to lower coefficients, resulting in

increased penetration but reduced image contrast. Conventional radiography produces 2D projection image by capturing the transmitted X-rays on a detector placed opposite to the source.

CT extend this principle by acquiring multiple projections from different angle using motorized x-ray source. It allows the source to rotate around the patient inside a circular structure called gantry. As the patient is gradually translated through the scanner, the system continuously emits X-ray beams. Numerical detectors located on the opposite side measure the intensity of transmitted radiation of each projection angle. This is illustrated on Figure 2.1.




Figure 2.1: Drawing of CT fan beam and a patient in a CT imaging system.[12]

These measurements are then reconstructed using numerical algorithms to generate cross-sectional images of the body called "slice". The resulting slices, typically ranging from 1 to 10 millimeters in thickness, are stacked to produce a 3D representation of the scanned regions. In CT imaging, the intensity values of each voxel are expressed in Hounsfield Units (HU), a standardized scale that quantifies the X-ray attenuation of tissues. This scale is defined such that air has a value of approximately -1000 HU, water 0 HU, and dense bone can achieve 2000 HU. Soft tissues typically range from 0 to 100 HU. This representation allows for a consistent differentiation between various anatomical structures based on their radio-density. To enhance the visualization of specific anatomical structure, contrast agents such as iodine or barium are often administrated, as they increase X-ray attenuation in targeted aeras.[13]

2.2 Magnetic Resonance Imaging

MRI is a non-invasive 3D imaging technique based on nuclear magnetic resonance. A powerful magnet generates a static magnetic field that forces the hydrogen protons in the body to align along its direction. A radio-frequency (RF) pulse is then applied, which excites the protons and temporarily tips their magnetization away from equilibrium. When the RF pulse is turned off, the protons return to their equilibrium state. During this relaxations process, they release energy that is detected

by receiver coils. These signals are then processed to reconstruct detailed images of the internal structures of the bodies via Radon Transfer. [14]

The return of the proton magnetization to equilibrium is governed by two relaxation mechanisms: T1 and T2.

T1 relaxation, also known as longitudinal relaxation, corresponds to the recovery of the magnetization along the direction of the magnetic field. It reflects the ability of protons to exchange energy with their surrounding environment. In highly mobile environment such as fluids, this energy exchange is less efficient, leading to longer T1 relaxation times. Conversely, in more structure or constrained tissues, molecular interactions are stronger and facilitate energy transfer, resulting in shorter T1 relaxation times. Therefore, tissues with short T1 relaxation times recover quickly and appear bright in T1-weighted images, whereas tissues with long T1 appear darker.

T2 relaxation, also known as transverse relaxation, describes the loss of phase coherence between protons in the transverse plane due to the local interactions between neighboring protons. In highly mobile environment, protons experience more uniform magnetic fields and maintain their phase coherence for a longer time, leading to longer T2. In contrast, in more structured or heterogeneous tissues, local magnetic field variations are stronger, causing faster dephasing of protons and therefore shorter T2 relaxation times. Tissues with long T2 relaxation times retain their signal longer and appear bright in T2-weighted images, while those with short T2 appear darker.

This variability in relaxation behavior across different tissues is what ultimately generates contrast in MRI images, illustrated on Figure 2.2. [15]




Figure 2.2: Coronal slice of a brain observed with MRI with T1-weighted (left) and T2-weighted (right) .[16]

Fluid-Attenuated Inversion Recovery (FLAIR) is an MRI sequence similar to T2-weighted imaging, in which the cerebrospinal fluid (CSF) signal is suppressed. This improves the visualization of lesions and abnormalities located near the ventricles and cortical sulci.

MRI provides high soft-tissues contrast without exposing patients to ionizing radiation. For epilepsy imaging, the ILAE recommends the HARNESS-MRI protocol, designed to be applicable worldwide. It consists of specific volume sequences including 3D T1-weighted, 3D FLAIR and 2D coronal T2-weighted images to improve subtle epileptogenic lesions detection and especially mesio-temporal lesions for the last image. The use of high-field (3T) and ultra-high-field (7T) MRI further enhances the sensitivity for identifying subtle cortical abnormalities that may be associated with epilepsy.

Chapter 3

Intracranial Electroencephalography

For decades, research on human brain has largely relied on fMRI to localize functional neural network with high spatial resolution. However, this technique has important limitations. In particular, fMRI has a low temporal resolution due to the indirect and slow nature of the hemodynamic response. In addition, fMRI signals mainly capture correlations between brain regions instead of directly measuring neuronal activity, limiting the ability to establish causal relationship between neural activity and observed cognitive or behavioral processes. To overcome these limitations, iEEG arises as a powerful tool for studying neural activity with high temporal resolution. [17]

This section introduces the principles of iEEG and reviews its main application in presurgical evaluation of epilepsy.

3.1 Basics of iEEG

Intracranial electroencephalography (iEEG) is defined as an electroencephalographic recording utilizing invasive methods with electrodes surgically implanted. It overcomes several limitations of fMRI by capturing neuronal activity with millisecond temporal resolution and millimeter-scale spatial precision. This allows the dynamics of cerebral networks to be studied at the speed of cognition. Unlike non-invasive recording techniques, iEEG is less sensitive to external noise such as electromagnetic interference and physiological artifacts including muscular activity, eye movements and cardiac signals. Moreover, it achieves a signal-to-noise ratio up to one hundred times higher than conventional scalp EEG and the amplitude of the recorded signals is typically about ten times larger. This improvement results from the direct placement of electrodes within the brain tissues whereas scalp EEG recordings are strongly affected by attenuation and distortions caused by the skull and scalp.[3], [18]

iEEG recordings can be performed either using subdural grid electrodes placed on the brain surface, referred to as electrocorticography (ECoG) or using depth electrodes inserted into the brain tissue referred to as stereo-encephalography (SEEG), both illustrated on Figure 3.1. Both provide much higher spatio-temporal resolution than conventional scalp EEG. However, only depth electrodes allow the recording of neural activities in deeper brain structures such as hippocampus, amygdala and insula. There is a growing preference for SEEG, as it does not require a craniotomy, thereby

reducing the risk of hemorrhage, infection, and other complications associated with brain surgery, which are more common in ECoG. The implantation of SEEG electrodes only requires burr holes and its removal does not need a second operation. Electrodes are made of stainless steel or platinum, the latter being preferred as it is non-ferromagnetic and therefore compatible with MRI. [7], [18]




Figure 3.1: Intraoperative photograph showing a left cerebral hemisphere covered by subdural grid electrodes arrays. Two depth electrodes are also seen (see white arrows) entering the brain tissue.[7]

Both ECoG and SEEG are used during the presurgical evaluation of drug-resistant epilepsy. However, SEEG has become preferred in many epilepsy centers due to its advantage explained previously. As the latter work focuses exclusively on SEEG using depth electrodes, the remainder of this state-of-the-art review focuses on this technique. [3], [19]

iEEG records complex signal originating directly from brain tissue and corresponding to the sum of local field potentials (LFPs) generated by large cell populations. It is estimated that there are 500 000 cells underneath or surrounding these electrodes. The signal is composed of several bands of oscillatory activity, each with a functional role.

- Slow frequency oscillations : these bands are associated with different functional brain states, ranging from deep sleep (delta: 1–3 Hz) and memory-related processes (theta: 4–7 Hz) to relaxed wakefulness (alpha: 8–12 Hz), active cognition (beta: 13–20 Hz) and high-order processing such as perception and attention (gamma: 21 to 40 or 50 Hz). It reflects the collective dynamics of neural population and are associated to different functional state of the brain.
- High-frequency broadband (HFB) activity: located above 50 HZ, it represents the local engagement of the cortical tissues located near the electrode.

iEEG provides information about the directionality of the information flow where HFB activity reflects responses of local population of cells and slow frequency oscillation coordinate the dialogue between regions.[20]

3.2 iEEG for epileptogenic zone detection

In a typical surgical framework for drug-resistant epilepsy, 5 to 20 depth electrodes, each containing 4-18 contacts spaced 2-10 mm apart, are stereotactically implanted into the patient's brain. iEEG recording are conducted following a two-stage clinical approach aimed at delineating both the EZ and EC.

First, a neurosurgeon implants the depth electrodes stereotactically through small burr holes in the skull. Following implantation, patients undergo prolonged monitoring which generally lasts between 3 to 14 days. During this period, clinicians analyze the recorded electrical activity. The second stage consists of the resection of the identified epileptogenic zone. In addition, iEEG enables direct cortical stimulation for functional mapping of eloquent cortex. These capabilities make iEEG a key tool in the presurgical evaluation of epilepsy.[7]

Compared to scalp EEG, iEEG provides access to a broader frequency range, enabling the detection of subtle electrophysiological biomarkers, explained in the next paragraph. The temporal evolution of iEEG signals typically includes interictal activity, pre-ictal transition phase and ictal onset corresponding to seizure initiation. An example of these signal patterns is illustrated in Figure 3.2. [3], [18]




Figure 3.2: Example of intracranial EEG signals showing the transition from interictal activity to pre-ictal state and ictal onset.[21]

At each stage of seizure progression, characteristic electrophysiological patterns observed in iEEG recordings serve as biomarkers for the detection and characterization of epilepsy. It is worth noting that EZ is not directly measurable in clinical practice. The SOZ is therefore considered as the best clinical approximation, although the two concepts are not strictly identical.

Ictal biomarkers constitute the gold standard for SOZ delineation. One of the most characteristic features is low-voltage fast activity (LVFA), which typically consists of a suppression of background activity followed by the emergence of high-frequency activity and subsequent slow potential. This pattern allows the recognition of the seizure onset. During monitoring, clinicians analyse iEEG recordings to identify the earliest electrode contacts exhibiting ictal discharges, thereby localizing the SOZ. However, ictal recordings are not ideal for systematic evaluation due to their unpredictability, low occurrence rate and associated clinical constraints in terms of time, cost, and patient safety. [18], [22]

Interictal biomarkers, which occur between seizure and are more frequent, provide useful complementary diagnostic information. They are essential to localize the irritative zone, corresponding to cortical area that generates interictal spikes. Two main classes of interictal biomarkers are commonly described : Interictal Epileptiform Discharges (IEDs) and High-Frequency Oscillations (HFOs).

IEDs, also referred as interictal spikes, are the oldest and most established biomarkers of epilepsy in clinical practice. To be classified as a IED, a wave has to fulfill at least one of the four following criteria, illustrated on Figure 3.3: a very brief sharp waves (spikes), a rapid succession of multiple spikes (polyspikes), waves with rapid rising phase and slower falling phase (sharp waves), temporal intermittent rhythmic delta activity (TIRDA). The occurrence rate of IEDs generally increases in regions with higher epileptic activity. It is worth noting that their spatial specificity is limited, as interictal spikes can propagate beyond their site of origin. IEDs are more reliable for identifying the IZ than for precisely localizing the SOZ. [18], [23]




Figure 3.3: Example of (A) polyspikes, (B) focal spike-waves, (C) TIRDA, (D) generalized spike-waves. Red rectangle and arrows indicate EEG patterns of interest.[23]

High-frequency oscillations (HFOs) are transient, low amplitude, typically characterized by frequencies ranging from 80 to 500 Hz. Two main categories of HFO are distinguished based on their frequency range. Ripples (80-250Hz) may occur during normal physiological processes such as sleep, memory processing or other normal cognitive functions. It can also be pathological. This is illustrated on Figure 3.4b. Fast ripples (250-500Hz) are almost exclusively associated to epilepsy, illustrated on Figure 3.4c. A strong correlation has been demonstrated between the presence of pathological HFOs (mostly fast ripples) and SOZ. The complete resection of tissue generating high rate of fast ripples, is correlated to seizure-free condition postsurgery. HFOs act as a marker of disease activity, increasing, for example, during reductions in AED treatment. These biomarkers have although a main challenge: it is difficult to distinguish pathological HFOs from physiological one. [18], [22], [23]

Recent approaches increasingly rely on hybrid biomarkers combining IEDs and HFOs to refine the localization. They focus their research on spikes containing high-frequency oscillations, also referred to as spikes ripples. These events reflect neuronal populations generating synchronous pathological discharges and actively

contribute to seizure generation and propagation. Figure 3.4a illustrates intracranial EEG recordings showing IEDs associated with HFOs. [22]




Figure 3.4: Example of intracranial EEG recordings showing epilepsy biomarkers including interictal epileptiform discharges and high-frequency oscillations. (a) iEEG recordings showing iEDs with HFOs, (b) Ripples, (c) Fast ripples [22]

During the monitoring period, additional experimental protocols may also be conducted, as iEEG provides a unique opportunity to access a source of information that are otherwise not available. [7] However, these rich and complex datasets require careful preprocessing before they can be reliably used for further analysis or research applications.

Chapter 4

Methodological Pipelines in Intracranial EEG Research

Despite variability in experimental paradigms, most iEEG studies share a common methodological backbone composed of (i) Electrode localization and anatomical registration, (ii) Electrode identification and labeling (iii) iEEG signal preprocessing (iv) Epileptic seizure detection.

4.1 Electrode Localization

The electrode localization is essential in any iEEG pipeline, as it enable the localization of electrode contacts. Without it achieving spatio-temporal coherence in the results is not possible. Several approaches have been proposed in the literature to address this problem. In this review of the state of the art, we focus on Brainstorm and ElectroLoc.

Brainstorm [24] is a collaborative, free, and open-source application dedicated to the analysis of the brain recordings. The tool is developed in MATLAB and provides a graphical user interface (GUI), thus not requiring any programming knowledge. Importantly, it does not require users to own a MATLAB license. The required inputs include a pre-implantation MRI, a post-implantation CT scan and a detailed knowledge of the electrode's physical characteristics such as the number, the length and diameter of electrodes and contacts and the contact spacing. The software performs a co-registration of the CT to the MRI in order to align the electrodes with the brain anatomy. The user then manually selects the deepest (Figure 4.2 a) and the most superficial (Figure 4.2 b) contacts by clicking on them successively in the GUI. Brainstorm constructs a line segment and automatically places intermediate contacts along it. The corresponding coordinates are subsequently exported for further processing. The flowchart of this software is illustrated on Figure 4.1.





Figure 4.1: Flowchart of Brainstorm software for electrode localization.[24]

The software offers a wide range of tools for linking signal analysis with anatomical data within a single application. In addition, the tutorials make it accessible to users with no prior experience.

Despite being user-friendly and versatile, Brainstorm requires users to have some expertise in interpreting MRI scans. Moreover, it does not handle curved electrodes well. In such case, the user must manually adjust the position of the contacts. Even after correction, the visualization remains limited, as electrodes are still displayed as straight line rather than curved one. [24]



(a) Selection of deepest contact (tip) by user

(b) Selection of most superficial contact by user

Figure 4.2: Steps for localization of electrode's contacts in Brainstorm

ElectroLoc is an internal method at UCLouvain developed by Q. De Laet as part of his master's thesis in 2025 [6]. It is based on mathematical morphology, graph theory and graph search procedures. The required inputs are a post-implantation CT, a brain mask, the number of electrodes and a file .csv containing the coordinate in voxel domain of entry point of each electrode and the number of contact associated. The algorithm, illustrated on Figure 4.3 is composed of three step: centroid extraction, linear modeling and a post-processing.




Figure 4.3: Summary of the overall algorithm of ElectroLoc [6]

The centroid extraction step is based on mathematical morphology where the electrodes are first segmented via a hard threshold of CT's voxel in Hounsfield Units (HU). A threshold of 2500 HU was applied to create an electrode mask. Mathematical morphology such as dilatation and erosion were used to overcome the following issues of the electrode: electrode can be separated in several pieces and several electrodes can be close and considered connected in the mask. The dilatation with a spherical structuring element allows to guarantee that each contact from the same electrode are joint into a subgroup. The erosion's goal with a 2D cross structuring element is to separate each contact into the smallest group of voxel possible. From each group of voxel, the center of mass, also called centroid, were calculated and used to extract the voxel coordinate of the position of a contact.

The linear modeling step aims to estimate a set of K linear electrode models that best fit the centroid cloud while assigning each centroid to its corresponding electrode. Before reconstructing the electrodes, the algorithm first needs to determine two key element: the inter-contact distance and potential candidates for electrode extremities. To estimate the inter-contact distance, it computes the distances between two neighboring centroids and selects the most frequently occurring value. In parallel, the algorithm constructs a minimum spanning tree for each subgroup of centroids. From these trees, it identifies points of interest (leaf nodes, neighbors of intersections and point forming sharp angles) which are considered potential candidates for electrode extremities. Using these candidates, the algorithm search for optimal set of linear models that best represent the electrodes. For each selected linear model, centroids are assigned based on their proximity on the corresponding line and the model parameters are refined through linear regression. This process ultimately result in K line segments that model the electrodes.

The post-processing stage addresses two remaining challenges at this point in the pipeline. First, the linear model are not fully representative of real world conditions, as electrodes may bend during implantation. Second, centroids constitute useful information about the general location of contacts, they do not ensure precise positioning along the electrodes. To overcome these limitations, the post-processing step introduces curved electrode models, specifically 3D parabolic curves, obtained through a regression process similar to that used for linear models. In addition, the centroids are replaced by more accurate contact locations using additional information from the CSV file which contains the entry point coordinates and the number of contacts per electrode. Each curved model is associated with its corresponding entry point. The contacts are then generated along the curve based on the estimated inter-contact distance, ensuring a regular and realistic spacing along the electrodes.

The algorithm's output is a CSV file containing the coordinates of the contact in voxel and world space, electrode identifier and a contact identifier ordered as follow $\{vox_x, vox_y, vox_z, world_x, world_y, world_z, electrode_id, c_id\}$. Despite the strong performance of the algorithm, some limitations remain. In particular, it requires the entry point coordinates as input, which currently need to be manually extracted using external visualization software such as MRICroGL [25] or MisterI [26]. This manual step depends on the user which can introduce variability and potential inaccuracies. Moreover, contacts are assumed to be ordered in the output CSV file such that the deepest contact is assigned the lowest identifier (0), while the most superficial contact receives the highest identifier, which is not always the case in practice. [6]

4.2 Electrode Identification

In the context of conventional scalp EEG, it exists several standard to place the electrode, for example the system 10-20 illustrated on Figure 4.4. This is possible only because the scalp is standardizable and the positions are external and reproducible.




Figure 4.4: System 10-20 of conventional scalp EEG where F=frontal, Fp=pre-frontal, C=central, T=temporal, P=parietal, O=occipital, A=auricular, Z=zero [27]

In iEEG, electrode placement is guided by the suspected EZ and the EC, rather than by specific research interest. As a result, implantation is hypothesis-driven. The association between contacts and channel labels is typically performed manually by the user, which is time-consuming and may lead to inconsistencies between clinical and research teams. [28, chap.27]

Although no standardized spatial placement system exists for iEEG, electrode naming conventions are commonly used to reflect their anatomical targets. Stone et al. (2021) proposed a standardized nomenclature known as Standardized Electrode Nomenclature for SEEG Applications (SENSA), on which channel names are intrinsically linked to their anatomical location. Channel labels are composed of at least two characters and follow the structured rules summarized in Figure 4.5. Despite the very high reported inter-rate reliability (97,5 %), this convention is not universally adopted across all research centers and therefore cannot be systematically relied upon in all studies. [29]

First Character Group		Second Character Group		
Lobe of Entry (can Add ^L or ^R if Bilateral)	Gyrus/Region	Position Relative to Electrodes Within Same Gyrus/Region	Standard Distal Electrode Target	Position Relative to Electrodes Within Same Target
Frontal	1—superior	a—most anterior	A—Amygdala	a—most anterior
Parietal	2—subsequent superior	b—2nd most anterior, and so on	C—Cingulate gyrus	b—2nd most anterior, and so on
Occipital	3—most inferior (if applicable)		H—Hippocampus	
Temporal			I—Insula	
			OF—Orbitofrontal cortex	
			OT—Occipitotemporal	
			gyrus	
			PH—Parahippocampal	
			gyrus	

Figure 4.5: Standardized nomenclature proposed by Stone et al. (2021) .[29]

To facilitate comparisons across patients, electrode positions are typically pro-

jected into standard space, such as the Montreal Neurological Institute (MNI) space, which corresponds to a standardized brain template constructed from multiple subjects. This process of spatial normalization is essential to overcome individual anatomical variability by establishing a common coordinate system, despite the absence of standardized rules regarding electrode implantation and electrode identification procedures. Once mapped into MNI space, electrodes can be overlaid on standardized anatomical or functional atlases (such as Automated Anatomical Labeling (AAL) illustrated on Figure 4.6) to automatically assign precise labels to each contact. This step is necessary to enhance the comparability and reproducibility of result across different research centers and datasets. [28, chap.27]




Figure 4.6: The AAL atlas views: (A) left lateral, (B) left medial, (C) anterior, (D) superior, (E) inferior, (F) right lateral, (G) right medial, and (H) posterior. The regions are colored to identify region boundaries. Similarity of colors between spatially separated regions is not meaningful; all regions are spatially contiguous. The cortical parcellation is based on anatomical landmarks, e.g., sulci. There are 78 cortical regions. [30]

4.3 iEEG signal preprocessing

Re-referencing is a fundamental step in iEEG data preprocessing, required by the biophysical reality that any recorded signal represents potential difference between two points. The reference problem arises because, while the reference site is theoretically assumed to be neutral, it is a myth that any area of the brain or scalp is truly electrically silent. Consequently, the choice of a reference directly influences the morphology, localization and interpretation of the signal. This step is essential to increase the signal-to-noise ratio (SNR) by removing shared environmental noise and physiological artifacts. Furthermore, re-referencing is a powerful tool to transform the data to isolate specific neural components. As there is no "gold standard" approach, the choice of montage must be a decision based on the specific research question, anatomical region and the frequency range of interest. Different type of re-referencing scheme are described below [18] :

- Monopolar Reference (Scalp Rereference - SCA) : each contact is referenced to a unique point, generally an electrode located on the scalp.

- Bipolar Reference (BIP) : widely used in iEEG, this method consists of the subtraction of each channel from its neighbor from the same electrode. It improves the spatial resolution and allows to reduce the global noise.
- Laplacian Reference (LAP) : each channel is referenced to the average of its direct neighbors. Similar to BIP, this method avoids to favor a specific direction of subtraction.
- Common Average Reference (CAR) : the signal is referenced to the average of all contacts. This method preserves the global information but is still very sensitive to high-amplitude artifacts.
- Zero Reference (REF0) : it is a virtual reference without any physical point associated.
- White Matter Reference (WM) : this method uses a contact located in the white matter, supposed inactive. Although theoretically sound, this hypothesis is incorrect because white matter is not entirely electrically neutral.

Hagen and Al. (2025) compared five montage (SCA, CAR, REF0, BIP, LAP) based on 77 patients undergoing a iEEG monitoring. The main results of the research was the temporal and spatial distribution were globally similar between the five montages. They conclude that there is no unique and ideal reference for iEEG. The choice depends on the research question, the region of interest and the range of studied frequency. They recommend to avoid SCA for high-frequency analysis, to give priority to local referencing (BIP and LAP) to capture low-amplitude and response. [31]

Parish and al. (2023) explores the impact of re-referencing on two type of tasks : motor (high-frequency) and memory (low-frequency). Local method (BIP and LAP) isolate efficiently very local signals in motor cortex and suppress artifacts. For slow oscillation like theta, local method tends to cancel the signal of interest. CAR or WM wrongly introduce artificial alpha oscillation in the signal of hippocampus. [28, chap.28]

Both recommend using at least two referencing schemes (one local and one common or data-driven) to assess the robustness of results.

4.4 Epileptic Seizure Detection

Since iEEG recordings are obtained from epileptic patients, they may be contaminated by seizure-related activity that can dominate and obscure the underlying neural processes of interest. Consequently, segments corresponding to seizures are generally excluded from analyses unrelated to epileptic activity.

There are several approaches to detect epileptic seizure in iEEG recordings. For this master's thesis, the choice was made to rely on the implementation developed in a previous master's thesis from UCLouvain named "Transformers for EEG classification: architectures, pre-training, and applications to epileptic seizure forecasting" by T. Bary (2023) [5]. In that work, T. Bary implemented the Multichannel Vision Transformers (MViT) architecture [32] and proposed novel pre-training strategies

aimed at improving training efficiency. The MViT architecture itself was introduced by Hussein et al. [32] as an adaptation of vision transformers to multi-channel signals such as EEG and iEEG. The main idea is that each channel is transformed into an image (here a scalogram), then processed independently by a transformer encoder. The encoder outputs are then concatenated and then sent to a final Multi-Layer Perceptron (MLP) which combines information from all channels and makes a decision about the segment. The classification task consist in distinguish between pre-ictal and inter-ictal segments. The architecture of a MViT is illustrated on Figure 4.7.




Figure 4.7: Architecture of a MViT [32] used for EEG classification. Scalograms computed from each EEG channel were divided into patches, embedded and processed by transformer encoders. The extracted features were then aggregated and classified using a multi-layer perceptron.

The required input is a scalp EEG signal. The pipeline divided into several main steps: data preprocessing, transformation into scalograms, patches embedding, model training, illustrated on Figure 4.8, and regularization strategies.




Figure 4.8: Schematic overview of the neural network architecture proposed by T. Bary [5] based on the MViT framework. The model combines a parallel encoder with a MLP decision head for binary classification.

During the data preprocessing step, the raw EEG recording are cleaned. Depending on the dataset, unusable portions of the signals are removed, labels are inferred or corrected and the signals are converted into a common montage when necessary. The recordings are then split into segments to increase the number of training sample to overcome the large data requirements inherent to transformer-based architectures.

As MViT architecture requires an image as input, each EEG segment is converted from a one-dimensional temporal signal into a two-dimensional time-frequency scalogram (shown in Figure 4.9) using Continuous Wavelet Transform (CWT). The Morlet family is the most commonly wavelet used in EEG analysis and is defined as a complex sine waves tapered by Gaussian windows illustrated on Figure 4.10.




Figure 4.9: Scalogram of EEG signal from training phase of T. Bary's model. (a) the heatmap, (b) the 3D plot of the heatmap. [5]




Figure 4.10: Real (a) and imaginary (b) parts of the Morlet mother wavelet [5]

The patch embedding step divides these scalograms into a fixed number of non-overlapping patches. Each of them are then projected into a lower-dimensional embedding space. The dimension reduction is performed from a space

$$A \in \mathbb{R}^{n_{\text{channels}} \times n_{\text{patches}} \times n_{\text{pixels/patch}}} \quad (4.1)$$

to a space

$$B \in \mathbb{R}^{n_{\text{channels}} \times n_{\text{patches}} \times n_{\text{embedding}}}. \quad (4.2)$$

A positional encoding is also added so that the model can retain information about the spatial location of each patch within the image. The patch embeddings are represented as TensorFlow tensors of shape (B, N, D) , where B denotes the batch size, N the number of extracted patches per sample, and D the embedding dimension. The model is implemented within the same TensorFlow framework, ensuring seamless integration between data representation and learning architecture.

Once the input tensors are prepared, the model training step can be performed as follows: each EEG channel is processed independently by a transformer encoder, and the encoded representations of all channels are then concatenated and passed to a MLP for the final classification task. The Gaussian Error Linear Unit (GELU) is the activation function chosen for the MLP.

To mitigate overfitting given the limited dataset size, several regularization strategies are employed, including weight decay, dropout and early stopping.

It is worth emphasizing that T. Bary's original contribution lies specifically in the pre-training strategies he developed to accelerate and improve MViT training, while the underlying architecture and signal processing pipeline follow the framework established in [32]. This implementation demonstrates strong performance and enables efficient parallelization of computations, significantly reducing execution time. However, it also presents an important limitation for this master's thesis: the model has been trained exclusively on scalp EEG data. Therefore, its reliability on iEEG data cannot be ensured at this stage, as the model has not been trained nor validated on such signals.

Part II

Pipeline Semi-Automated for Pre-Processing in iEEG Research

Chapter 5

Materials

This chapter presents the dataset of medical images and iEEG signals used for the development and validation of the algorithm described in Chapter 6.

The MRI and CT scans were first preprocessed to obtain 3D volume suitable for subsequent analysis. A dataset correction was also performed to address discrepancy identified between the imaging and signal data.

5.1 Data Acquisition

For the development of this pipeline, a dataset collected from two patients was used. These patients, one female and one male, both had drug-resistant epilepsy and were undergoing a pre-surgical evaluation through implantation of depth electrodes. The aim of their visit was to localize the EZ for a potential later ablation.

Patient 1 was implanted with 8 electrodes, each containing from 12 to 18 contacts. Patient 2 was implanted with 9 electrodes each containing 15 and 18 contacts.

The dataset is composed of a pre-implantation MRI, a post-implantation CT scan, an iEEG recording and a SEEG surgical plan.

Tables 5.1 and 5.2 summarize the CT and MRI acquisition parameters for both patients.

Parameters	Patient 1	Patient 2
Rows	512	512
Columns	512	512
Pixel spacing (mm)	0.469×0.469	0.605×0.605
Slice thickness (mm)	0.625	1

Table 5.1: CT acquisition parameters for both patients.

Parameters	Patient 1	Patient 2
Rows	256	256
Columns	256	256
Pixel spacing (mm)	1×1	0.898×0.898
Slice thickness (mm)	1	0.900
Magnetic field (T)	3	3
TR (ms)	2300	2000
TE (ms)	2.26	2.58
Flip angle (deg)	8	8

Table 5.2: MRI acquisition parameters for both patients.

5.2 Preprocessing

To meet the requirements format file for the algorithm presents in this thesis, several preprocessing step had to be performed.

5.2.1 File format

The initial dataset is provided in the following formats: CT and MRI data in Digital Imaging and Communication in Medicine (DICOM, .dcm) file format and iEEG signals in European Data Format (EDF, .edf) format.

EDF files contain both recorded electrophysiological signals and global metadata describing the recording and channel-specific information, including signal names, sampling frequencies, and physical units.

DICOM is an international standard to transmit, store, retrieve, print, process and display medical imaging information. It ensures compatibility and interoperability between different imaging systems (CT, MRI, etc.). In addition to image data, DICOM files contains associated metadata, including spatial information, acquisition parameters and patient details. [33]

Although DICOM is the standard format for medical image acquisition, it is commonly converted into the Neuroimaging Informatics Technology initiative (NIfTI, .nii) format for further analysis. NIfTI is a widely used format in neuroimaging for storing volumetric brain imaging data such as MRI and CT scans. Unlike DICOM, which stores each slice as a separate file, NIfTI represents the entire volume in a single file. It stores both the image data and a simplified metadata than DICOM including voxel dimension, data type, affine transformation matrix, which is essential for mapping voxel coordinates to real-world coordinates. NIfTI usually has .nii.gz extension containing both the header and the data. [34]

The conversion of CT and MRI images from DICOM format to NIfTI format is performed using dcm2niix, available within the graphical user interface of MRIcroGL. [25]

5.2.2 Skull stripping and brain mask extraction

Head CT scan and MRI capture the entire cranial region, including the skull and, in the case of CT with implanted electrodes, external elements such as wires and head support. To focus the analysis on intracranial content, images were processed using

the *mri_synthstrip* tool from FreeSurfer, based on SynthStrip [35], which removes non-brain areas.

This tool supports both MRI and CT images, including cases with implanted depth electrodes, and is designed to preserve intracranial structures while removing surrounding tissues. Figures 5.1 and 5.2 illustrate the effect of skull stripping on MRI and CT images, respectively. In addition, it allows the generation of a corresponding binary brain mask from the input images. The corresponding relationship between the resulting images can be expressed as Equation 5.1:

$$\text{output} = \text{input} \times \text{mask} \quad (5.1)$$




Figure 5.1: A slice of Patient 1's MRI (gray) and the corresponding image skull-stripped (red)




Figure 5.2: A slice of Patient 1's CT scan (gray) and the brain mask generated by SynthStrip [58] (green).

5.2.3 Discrepancy observed in Patient 2

The MRI and CT images showed an acquisition date different from one month prior the acquisition of the iEEG recording. Moreover, the expected number of electrode from the headers of iEEG signal described 9 electrodes whereas post-implantation CT scan showed 13 electrodes. This discrepancy led us to a modification in the dataset for their use in the following algorithm.

The modification was achieved by creating 4 "false" electrode named FA, FB, FC and FD. Their entry points and number of contacts were extracted manually.

Chapter 6

Methods

This Chapter presents the algorithm developed as part of this thesis. The overall goal of this pipeline is to localize SEEG electrode contacts within the brain, assign them an anatomical label and flag signal segments associated with epileptic activity. To this end, the framework, illustrated in Figure 6.1 is organized into sequential modules: data preprocessing, electrode identification, electrode localization, contact labeling and epileptic seizure detection. A final section in this Chapter is dedicated to the summary of the overall algorithm.




Figure 6.1: Summary of the proposed algorithm. The workflow combines data preprocessing, seizure detection, electrode identification, electrode localization and anatomical contact labeling.

6.1 Data Preprocessing

This section describes the preprocessing steps developed in this master's thesis to prepare the three inputs required by following steps: the CT scan, the brain mask and the raw iEEG signal.

6.1.1 iEEG signal preprocessing

iEEG signals were stored in an EDF file, which included header information describing the recording. Channel names were extracted from this header and typically followed a structured naming convention such as "EEG PA1, EEG PA2,..." where the prefix denoted the electrode identifier and the suffix indicated the contact index.

Non-EEG channels were removed from the list and the remaining channels were then grouped based on their electrode identifier, enabling the counting of number of

contact per electrode. From this step, a CSV file was generated with the electrode identifier and its number contacts associated.

6.1.2 Images preprocessing

The segmentation of electrodes in the CT scan was implemented for further analysis. It consisted of the separation of electrode's voxels from voxels of the rest of the image. The technique used for this thesis was the hard-thresholding segmentation, similar to Q. De Laet's master's thesis. As explained in Chapter 2, in CT, tissues are distinguished according their mass attenuation coefficient and are expressed in HU. Human tissues are generally ranging from -100 to 2000 HU and thus from the input CT scan post-implantation, the only other tissus than brain was the electrodes. The threshold used is 2200 HU. The resulting electrode mask is illustrated on Figure 6.2.

From the brain mask generated from FreeSurfer tool explained in Chapter 5, a mask corresponding with the surface of the brain was generated. This extraction was implemented using a neighborhood analysis: a voxel was classified as part of the brain surface if at least one of its neighboring voxel lied outside the brain mask (i.e. had a value of 0), thus indicating a boundary between brain and non-brain regions. The generated brain's surface mask is illustrated on Figure 6.2




Figure 6.2: Electrode and brain surface masks obtained for Patient 1. The brain surface mask is shown in white, while the electrode mask is shown in red. (left) coronal view, (middle) 3-D rendering, (right) axial view.

From the intersection of the two masks, the coordinates of the entry points were extracted. Since each contact was represented by multiple voxels, the resulting intersection consisted of several groups of voxels rather than single points. To address this, the voxels were grouped into clusters. The number of cluster was determined based on the number of electrodes identified in the iEEG signal preprocessing, described in previous chapter. For each cluster, the center of mass was computed in order to obtain a single representative point corresponding to the entry point of each electrode. This is illustrated on Figure 6.3. The coordinates in voxel and world space of the resulting entry points were stored in two separate CSV files.




Figure 6.3: Visualization of electrode mask (grey), cluster of voxels (blue) and center of mass (red).

6.1.3 Summary of Data preprocessing module

The flowchart in Figure 6.4 summarizes the processes presented in this section. First the channel names and number of contacts were extracted. Then, an electrode mask and a surface brain mask were generated. Following that, the intersection between these two masks was computed. Finally, the center of mass of each resulting cluster was computed and used as the corresponding entry point. Table 6.1 details the input and output of the data preprocessing module.




Figure 6.4: Summary of the data preprocessing module with hyperparameters in blue

Type	Name	Description
Inputs	CT	Post-implantation 3D CT volume
	Brain mask	Binary brain mask in the CT space
	iEEG signal	EDF file containing raw iEEG signals and metadata
Output	Voxel coordinates	Contact coordinates in voxel space.
	World coordinates	Contact coordinates in millimetres.
	Electrode mask	Binary mask of the electrode in the CT space.
	Surface brain mask	Binary mask of the surface brain in the CT space.
	Contact number	Number of contacts associated with each electrode.

Table 6.1: Inputs and outputs of proposed data preprocessing module

6.2 Electrode Identification

While entry points were extracted automatically in the previous section, associating each with its electrode identifier required a manual step, as the data did not follow the SENSE naming convention proposed by Stone et al. (2021).

Using the SEEG surgical implantation plan and the extracted coordinates of entry point, a visual inspection was performed to associate each entry point to its electrode identifier. During this step, the correction with the four false electrodes were performed on Patient 2 by adding them manually in the list of electrodes. The identification was possible only if the SEEG surgical implantation plan provided enough information about the electrode's placement. For this work, MRICroGL was used. The resulting information now available is illustrated on Table 6.2

<i>Electrode ID</i>	<i># Contacts</i>	<i>(vox_x, vox_y, vox_z)</i>	<i>(world_x, world_y, world_z)</i>
OS	18	(271, 89, 159)	(−7.031, −79.698, 179.806)
DC	12	(12, 314, 151)	(−27.186, −73.135, 174.806)
OA	12	(339, 113, 125)	(−38.906, −68.448, 158.556)
...	...	...	...

Table 6.2: Electrode information available after electrode identification including number of contacts and entry point coordinates from Patient 1.

At this stage of the algorithm, two CSV file were generated. The first one contained the number of contacts and the coordinate expressed in voxel space of entry points. This file was formatted to meet the input requirements of the ElectroLoc module introduced in Section 4.1. It was therefore structured as follows: $\{nb_contacts, vox_x, vox_y, vox_z\}$.

The second one contained the number of contacts, the electrode identifier and the coordinates in world space of entry points, structured as followed: $\{electrode_id, nb_contacts, world_x, world_y, world_z\}$.

6.3 Electrode Localization

The ElectroLoc module by Q. De Laet, presented in Section 4.1 was used as a blackbox for the electrode localization in this pipeline. Previous sections detail the methodology used to generate the required inputs: a CT scan post-implantation, a brain mask and a CSV file containing the number of contacts and the coordinate expressed in voxel space of the entry points (corresponding to the first CSV generated in Section 6.2).

The output's algorithm was a CSV file containing the following information of each contact: $\{vox_x, vox_y, vox_z, world_x, world_y, world_z, electrode_id, c_id\}$. It is worth noting that the electrode identifier from this module is provisional and is ranging from 0 to n where n is the number of electrode.

The implementation of the module can be found at this following adress: https://github.com/Quent-DL/MasterThesis_ElectroLoc.git.

6.4 Contact Labeling

This section presents the steps developed to assign anatomical and labels to the detected electrode contacts. This process is essential for enabling further anatomically informed analysis of iEEG signals. It includes a post processing of the ElectroLoc output, an atlas dilatation procedure to handle unlabeled regions and the final anatomical labeling based on atlas correspondence.

6.4.1 Post-processing of ElectroLoc

Re-indexation

The numbering phase of the ElectroLoc module did not achieve perfect accuracy, as illustrated in Figure 6.5 where an incorrect contact ordering can be observed. This step, developed as part of this thesis, aimed to assign final electrode names and re-order contacts indices for each electrode. It used two input: the list of detected contacts (Section 6.3) and the list of entry points expressed in world space (Section 6.2). The objective was to ensure that contacts were consistently ordered from the deepest to the most superficial.




Figure 6.5: Visualization of the ElectroLoc's output. (pink) Brain tissue, (red) contacts, (green) electrode mask

This procedure relied on a key assumption: the detected groups of contacts were already aligned with the order of the provided entry points. In other words, the first group of contacts corresponded to the first entry point, the second group to the second entry point, and so on. It should be noted that the estimated entry point coordinates did not perfectly coincide with the actual entry points identified during electrode localization. This discrepancy could be explained by the approximation of entry points as the center of mass of voxels at the intersection between the brain surface and the electrode mask. Additionally, in practice, the first contact was

not necessarily located at the brain surface and might instead lie deeper along the electrode trajectory.

For each group of contacts, the corresponding entry point was identified based on the assumption presented earlier. A straightforward approach would have consist in ordering the contacts according to their Euclidian distance to this entry point, assuming that the farthest contact was the deepest and closest is the most superficial. However, this assumption did not hold in all cases. In some configuration, electrodes might be inserted such that not all the contacts lied within the brain. As a result, certain contacts might be located beyond this entry point, leading to larger distances to entry point despite being anatomically more superficial. This could result in incorrect ordering if distance alone was considered. To address this limitation, the ordering was instead performed by projecting each contact onto the estimated electrode axis. This axis was obtained using a principal component analysis (PCA) of the contact position. The projection provided a signed distance along the electrode trajectory relative to the entry point, enabling a consistent ordering of the contacts from the deepest to the most superficial. Positive values indicated contacts located inside the brain, deeper than the entry point along the trajectory, whereas negative values indicated contacts located on the opposite side of the entry point, corresponding to superficial or extracranial contacts.

Next, the contacts were then reindexed accordingly: the deepest was assigned an index $c_id = 0$ and the index increased progressively toward the most superficial contact. The provisional electrode identifiers from ElectroLoc algorithm were also replaced by the final identifiers provided in the entry point list (i.e., the channel names).

As a verification step, the number of detected contacts in each group was compared with the expected number of contacts for consistency checking.

Shift correction

An additional post-processing step was developed in this Section to adress the shift issue, identified by Q. De Laet as a major limitation of his algorithm. This phenomenon, illustrated on Figure 6.6, induced both false positive and false negative detections at extremities of the electrode which could lead to incorrect signal assignation for subsequent studies.




Figure 6.6: Shifted electrode (green) from electrode mask (red) with False Positive (yellow arrow) and False Negative (pink arrow)

The correction was applied after contact reconstruction and re-indexation. Its purpose was to compensate potential longitudinal shifts along the electrode trajectory, while preserving the relative spatial configuration of the contacts.

The presence of a shift was first assessed by evaluating the distance between the electrode mask and the contact with $c_id = 0$, corresponding to the deepest contact. If the distance was large (in this work, greater than 2 mm), it suggested that the reconstructed electrode was shifted along its axis and therefore required correction.

The correction was performed by applying a set of predefined range of shifts along the electrode axis. In this work, shifts were explored between -20 mm and 20 mm with a step of 0.2 mm. The optimal shift was selected by minimizing a cost function, detailed in Equation 6.1, combining the alignment with the entry point and the consistency with the electrode mask.

$$J(s) = \lambda \cdot \left\| \mathbf{x}_{\text{superficial}}^{(s)} - \mathbf{x}_{\text{entry_point}} \right\| + \frac{1}{N} \sum_{i=1}^N d(\mathbf{x}_i^{(s)}, \mathcal{M}) \quad (6.1)$$

where $\mathbf{x}_{\text{superficial}}^{(s)}$ denotes the position of the most superficial contact after applying a shift s , and $\mathbf{x}_{\text{entry_point}}$ corresponds to the associated entry point. The first term therefore represents the Euclidean distance between these two points. It penalizes shifts which the superficial contact remains far from the entry point.

The term $d(\mathbf{x}_i^{(s)}, \mathcal{M})$ denotes the distance between the i -th shifted contact and the electrode mask $\mathcal{M}$, computed using a distance transform, and N is the total number of contacts. The second term thus corresponds to the mean distance of all contacts to the electrode mask. It evaluates how well the shifted electrode remains aligned with the electrode mask.

The parameter λ is a weighting factor that controls the trade-off between alignment with the entry point and consistency with the electrode mask. This formulation was introduced to avoid solutions that would perfectly align the most superficial contact with the entry point while placing the remaining outside the electrode mask. In this work, λ was set to 0.5 .

In addition to this cost function, a second criterion was introduced to account for partially implanted electrode for which some contacts might lie outside the brain. The deepest contact corresponding to $c_id = 0$ needed to remain close to the electrode mask for every electrode. It was way more important in case of partially implanted electrode where therefore the alignment with the entry point became unreliable as the most superficial contact did not necessarily correspond to the anatomical entry point. Therefore, for these electrodes, the optimization strategy was modified: the contribution of the entry point was removed from the cost function by setting λ to 0 and the shift was selected by minimizing the distance between the deepest contact ($c_id = 0$) and the electrode mask only. This ensured that the electrode remained correctly aligned with the detected trajectory within the brain, even when part of the electrode lied outside of it. Partially implanted electrodes were identified by threshold of the distance between $c_id = 0$ and the electrode (here, 1mm).

At the end of the optimization, the selected shift was then applied once to all contact of the electrode while updating both world and voxel coordinates.

6.4.2 Atlas dilatation

An atlas dilatation was developed in this pipeline to adress the limitation of atlases such as AAL which are mainly defined on grey matter and assign anatomical labels to cortical areas. Consequently, white matter voxels are often non-labeled (typically assigned a value of 0). In the context of iEEG, depth electrode may be located outside the grey matter or even pass through white matter which can result in contacts being be assigned to an unlabeled region. To reduce this issue, a morphological dilatation of the atlas was performed, extending cortical labels into neighboring white matter regions. This allowed contacts located close to the cortical surface to be associated with a relevant anatomical regions. This step was particularly important to compensate small spatial inaccuracies arising from image registration, anatomical variability across subjects. Without dilatation, a significant number of contacts could have remained unlabeled.

The atlas dilatation was developed using a distance-based approach. First, unlabeled voxels were identified. A brain mask was used in this process to ensure that the dilatation process remained restricted to brain tissue. The identified unlabeled voxel were retrieved by excluding labeled voxels and outside brain voxels.




Figure 6.7: Visualization of AAL atlases undilated (top) and dilated (bottom) with a dilatation width of 6 mm

From this identification, for each unlabeled voxel, both the nearest labeled voxel and the distance with it were retrieved. For all unlabeled voxels located within the predefined distance threshold (i.e., dilatation width), the label of the nearest labeled voxel was assigned. An illustration of this process is in Figure 6.7. During the implementation, dilatation widths from 0 mm to 10 mm were tested in order to define the most appropriate value.

This step remained optional depending on the user's preferences and the study interest.

6.4.3 Anatomical atlas labeling

To assign anatomical labels to each electrode contact, a direct application of the atlas to the patient’s MRI was not possible due to differences in spatial reference systems. The processing pipeline therefore involved several steps.

First, the MNI space was co-registered in the patient’s MRI space. This step computed a transformation that aligns the individual anatomy with the standard template MNI. Then, the atlas was projected in the new co-registered MNI space. This ensured that the anatomical regions defined in the atlas were properly aligned with the patient-specific anatomy. Once the atlas has been mapped to the MRI space, each contact could be associated with the anatomical label. The contact coordinates, initially expressed in the CT space, were first transformed into the MRI space using the CT-to-MRI co-registration mapping. These transformed coordinates were then converted from world coordinates to MRI voxel coordinates using the inverse MRI affine matrix.

Both co-registration steps were performed using the Python package *regis.core* version 0.0.15 (available in <https://github.com/PiLAB-Medical-Imaging/registration>), specifically through the `find_transform` and `apply_transform` functions. The parameters used for each co-registration are summarized in Table 6.3. All co-registrations were performed with the parameter *level_iters* set to [1000, 100, 10]. This multi-resolution strategy allowed a high number of iterations at the coarsest level to ensure robust global alignment while progressively reducing the number of iteration at finer levels as anatomical details increased. The parameter *diffeomorph* was set to False in order to restrict the registration to linear transformations and avoid local non-linear deformations. This choice was made to preserve the geometrical consistency of the images, especially for the CT-to-MRI registration where electrode positions must not be locally distorted. The parameter *sanity_check* was enabled for the registrations involving the atlas spaces, in order to visually verify the quality of the alignment and reduce the risk of assigning incorrect anatomical labels due to a faulty registration. Finally, intensity normalization was applied for the MNI-to-MRI and atlas-to-MNI registrations, as these images have comparable anatomical intensity distributions, but was not applied for CT-to-MRI registration because CT and MRI intensities represent different physical quantities and the CT image contains high-intensity electrode artefacts.

Parameter	CT to MRI	MNI to MRI	Atlas to MNI
<i>level_iters</i>	[1000, 100, 10]	[1000, 100, 10]	[1000, 100, 10]
<i>diffeomorph</i>	False	False	False
<i>sanity_check</i>	False	True	True
<i>normalize</i>	False	True	True

Table 6.3: Parameters used for co-registration with the Python package *regis.core*.

As the atlas has been previously aligned with the patient MRI, the voxel coordinate of the contacts could be directly used to query the atlas volume. For each contact, the voxel value extracted from the atlas was interpreted as a region index and matched to its corresponding anatomical label using the atlas text file. If a contact fell outside the atlas boundaries, the label attributed was `{out_of_bounds}`.

This module finally stored, for each contact, its electrode identifier, contact index, MRI voxel coordinates, MRI world coordinates, and assigned anatomical label in a CSV file.

6.4.4 Summary of contact labeling module

The flowchart in Figure 6.8 summarizes the processes presented in this section. First, a post-processing of electrode localization was performed to re-order the indexes of contacts, re-name the electrode identifiers with the channel names extracted in Section 6.1 and to handle unexpected shift. Then, an anatomical label was assigned for each contact. An atlas dilatation was possible according to the user's preference. Table 6.4 details the input, output and hyper-parameters of the contact labeling module.




Figure 6.8: Summary of the contact labeling module including contact re-indexing, shift correction, co-registration, anatomical labeling and an optional atlas dilatation. Hyper-parameters are shown in blue.

Type	Name	Description
Inputs	Entry points	Entry point coordinates in world space
	Contacts	Contact coordinates with electrode ID and contact index
	Electrode mask	Binary mask of the electrodes in the CT scan.
	Atlas volume	3D NIfTI volume containing labels for each anatomical regions.
	Atlas labels	Text file mapping each integer label to its corresponding anatomical region.
	CT-to-MRI mapping	Transformation from CT to MRI space
	MRI affine	Transformation from MRI voxel to world space
Hyperparameters	MNI brain mask	Template brain mask defined in MNI space.
	Shift threshold	Distance used to identify shifted electrode.
	Dilatation width	Distance used to extend atlas labels.
Output	Contacts	List of contacts with associated anatomical labels

Table 6.4: Inputs, outputs and hyperparameters of contact labeling module

6.5 Epileptic seizure detection

The epileptic seizure detection module operates independently from the contact extraction pipeline. Its purpose is to classify iEEG segments as either pre-ictal or inter-ictal, using the MViT-based model implemented by T. Bary [5], whose pre-training procedure and weights are available on his GitHub repository: <https://github.com/tbary/EEGPreTrainingDatasets>. Due to the absence of labeled iEEG data, the model could not be fine-tuned on the target domain. Consequently, predictions on iEEG signals were performed using the pre-trained model originally trained on EEG data.

6.5.1 Pre-processing of iEEG signal for MViT

In order to comply with the input requirements of the pre-trained model, the iEEG signals underwent preprocessing steps similar to those described in Section 4.4. However, the hyperparameters adopted in this work were directly inherited from that study and were not specifically optimized for iEEG data, which might impact the model's performance. Figure 6.9 illustrates each step of the process.




Figure 6.9: Illustration of 5 minutes iEEG recordings (only 5 channels displayed), segmentation (5s segment with bipolar referencing) and scalogram generation for channel LB13-LB14 from Patient 2.

The data were first extracted from EDF files using the MNE Python library version 1.11.0 [36], which offered a standardized framework and built-in functions for loading, preprocessing and analyzing EEG and iEEG signal. Non-EEG channels were excluded and a downsampling step was applied to reduce data size, as high temporal resolution was not essential for seizure detection. The default downsampling frequency used for this master’s thesis was 1024 Hz.

The data were subsequently re-referenced using a common montage in order to enhance their interpretability and reduce noise or common-mode artifacts. Several referencing strategies can be applied, depending on the user’s preference and analysis requirements, including monopolar, bipolar, Laplacian and common average reference. For this work, the bipolar referencing were chosen arbitrarily.

To ensure compatibility with the input structure expected by the model, the iEEG signals were segmented into non-overlapping windows of 5 seconds. This segmentation allowed the transformation of long recordings into a set of independent samples suitable for model inference. For patient 1, this preprocessing step resulted in a total of 327 segments, each of which was processed individually by the model. For Patient 2, the same preprocessing pipeline was applied, leading to a set of 354 segments.

For each segment, the signal from each channel was then transformed into a scalogram using the PyWavelets library version 1.8.0 [37] and stored in .mat file. This transformation relied on the Continuous Wavelet Transform with a wavelet from the Morlet family.

During scalogram generation, a key step was the definition of the scale axis, which determines the range of frequencies analyzed. However, it does not directly correspond to physical frequencies. Instead, scales are inversely related to frequency up to a proportionality factor that depends on the chosen wavelet. In the case of Morlet wavelet, this relationship can be approximated as:

$$f = \frac{f_c}{s \cdot \Delta t} \quad (6.2)$$

where f denotes the frequency (Hz), s the scale, Δt the sampling period (i.e., $\Delta t = 1/f_s$), and f_c the central frequency of the wavelet.

The central frequency f_c is an intrinsic parameter of the chosen wavelet and does not depend on the analyzed frequency range. In the case of the Morlet wavelet implemented in PyWavelets, f_c is fixed (approximately 0.8125 in normalized units). The function `pywt.frequency2scale` were used in this work and the selected frequency range was $f = [0, 500Hz]$. Moreover, frequency f were normalized by multiplying by Δt before introduced into the function.

In order to preserve the multichannel structure of the iEEG signal, scalograms corresponding to the same temporal segment were subsequently grouped together. This results in a structured dataset where each sample was composed of multiple scalograms, one per channel, allowing the model to process spatial information across channels consistently.

Each group of scalograms corresponding to a segment was then loaded using a custom PyTorch dataset. This dataset ensured that the channel-wise scalograms belonging to the same temporal segment were retrieved together and formatted as a single model input. For each scalogram, the .mat file was loaded and optionally log-transformed, normalized and resized to a fixed resolution. For this work, the choice were made to log-transform, normalize and resize each scalogram to a size of 80×80 to ensure compatibility with the model input.

The preprocessed scalograms were finally stacked into a tensor representations of shape (C, H, W) where C denotes the number of channels and $H \times W$ the spatial dimensions of the scalogram.

The patch embedding stage was implemented as a custom PyTorch module by inheriting from `nn.Module`. The input scalograms were represented as tensors of shape (B, C, H, W) , where B denotes the batch size. They were divided into non-overlapping patches of size 16×16 , as illustrated in Figure 6.10. Rather than explicitly extracting each patch one by one, this operation was implemented using a 2D convolution layer (`nn.Conv2d`) with kernel size equal to the patch size and stride also set equal to the patch size. This configuration ensured that the convolutional kernel moved across the image without overlap. In addition, the convolution performed a linear projection of each patch into an embedding space of dimension 40, producing a sequence of tokens compatible with the transformer architecture. Finally, a learnable positional embedding was added to the patch embeddings in order to preserve spatial information.




Figure 6.10: Illustration of patches from scalogram LB13-LB14 from Patient 2.

The output of this preprocessing module was a sequence of embedded tokens derived from the input scalogram. The resulting tensor (PyTorch) had the shape (B, C, N, D) , where $N = 25$ is the number of patches per channel and $D = 40$ is the embedding dimension. A learnable positional embedding was added to preserve spatial information.

6.5.2 Prediction for iEEG segments

The model was re-implemented in Tensorflow following exactly the architecture described in [5] without modification. The encoder layers and the decision layer were directly reused from the original implementation without any modification. As a reminder, the model relied on a parallel transformer-based encoder designed to process multichannel inputs and a decision layer to perform the final binary classification. The parameters of the model for this work are detailed in Table 6.5 and corresponds exactly to those used during the original training phase except for the number of channel C which here was 122 and 137 for Patient 1 and 2 respectively instead of 20 for original model.

Table 6.5: Hyperparameters of the TUSZ model architecture from [5]

Hyper-parameter	Value
Input tensor shape	$(B, C, 25, 40)$
Number of parallel encoders	20
Transformer layers per encoder	8
Attention heads	4
Transformer MLP dimensions	[80, 40]
Decision head hidden layers	[512, 256]
Output dimension	1

Once initialized, the weights were loaded and the model was used to perform inference on the preprocessed iEEG segments. The model outputted logits, which

corresponded to unbounded scores for each input sample. These scores were converted into probabilities using a sigmoid activation function, mapping them to the range $[0, 1]$. A threshold of 0.5 was then applied to obtain the final binary classification. In this context, class 1 corresponds to pre-ictal activity, while class 0 corresponds to inter-ictal activity.

As mentioned previously, the model was implemented in Tensorflow whereas the preprocessed iEEG data were represented as Pytorch tensor. Therefore, an additional preprocessing step was implemented before inference to ensure the compatibility with the model. This conversion was performed by first detaching the PyTorch tensor from the computational graph, transferring it to the CPU, and converting it into a NumPy array. The resulting NumPy array was then converted into a Tensorflow tensor, which could be used as input to the model.

6.5.3 Summary of epileptic seizure detection module

The flowchart in Figure 6.11 summarizes the overall algorithm of the epileptic seizure detection module. First, the segmented iEEG signal were preprocessed into scalogram and subsequently divided into non-overlapping patches. The model was then reconstructed and the input tensor was converted from PyTorch into Tensorflow format to ensure compatibility. Finally, the model performed the prediction. The Table 6.6 details the input, output and hyperparameters of the module, excepting the model's parameters already described in Table 6.5.




Figure 6.11: Summary of epileptic seizure detection module with hyperparameters in blue

Type	Name	Description
Inputs	iEEG signal	EDF file containing the raw iEEG signal.
	Pre-trained weights	Weights from the original training model.
Hyperparameter	Sampling frequency	Sampling rate for downsampling the signal. [Hz]
	Segment duration	Duration of each signal segment. [s]
	Frequency range	Frequency interval retained for seizure detection. [Hz]
Output	Prediction	Binary classification of each segment as pre- or inter-ictal.

Table 6.6: Inputs, outputs and hyperparameters of epileptic seizure detection module

6.6 Algorithm summary

Figure 6.1 summarizes the overall pipeline of this master's thesis. The CT scan and MRI were introduced to the image preprocessing module while the iEEG signal to the signal preprocessing and epilepsy seizure detection modules. These steps produced separately the required inputs for the electrode localization module. The electrode identification module then linked these inputs, enabling electrode localization. A shift correction and re-indexation of contacts enables an accurate final localization. Finally, each contact was assigned with anatomical label and associated with the classification obtained from the epilepsy seizure detection module.

The results were stored in two CSV files. The first file contained the following information for all predicted contacts:

- electrode identifier;
- contact index;
- coordinates in voxel space;
- coordinates in world space;
- anatomical label.

The second file contained the prediction results for each segment.

6.7 Performance Assessment

The performance assessment can be divided into several parts.

Data Preprocessing The performance of the data preprocessing module was assessed by comparing the semi-automated extraction of the number of contacts and entry points with a manually established ground truth obtained by visual inspection. The precision of entry point detection was evaluated by computing the distance between each extracted entry point and its corresponding ground truth location. An entry point was considered correctly detected when this distance was below 2 mm which is higher than the spatial resolution of the CT scan.

Characterization of electrode localization The performance of the contact localization, the re-indexation and shift correction was assessed using the same three steps process. First, a comparison between the detected contacts and the electrode mask was achieved. The distance between each contact and the nearest point of the mask is computed using the Euclidean distance transform.

A visual inspection of the localized contacts relative to the electrode mask is also performed. For this purpose, a binary mask is generated from the CSV file containing the contact coordinates. This mask is created by assigning spherical region around each contact position using a predefined radius parameter corresponding to the number of neighboring voxels included around the contact coordinates. In this work, a radius of 2 mm is used, corresponding to approximately 3 voxels in the CT space. This value was selected to ensure good visibility of the contacts while maintaining a good

localization accuracy. The resulting contact mask is then overlayed on the electrode mask and visually inspected to assess the spatial consistency between the two.

A ground truth was subsequently established manually by visually inspecting the electrode mask and retrieving the contact coordinates, electrode identifiers and contact indices. Based on this ground truth, the following classification metrics are computed:

- Precision $:= \frac{TP}{TP+FP}$. This metric evaluates the proportion of true positives among all predicted positives.
- Recall $:= \frac{TP}{TP+FN}$. This metric evaluates the proportion of true positives among all actual positives.
- F1-score $:= 2 \cdot \frac{Precision \cdot Recall}{Precision+Recall}$. This metric provides the Harmonic mean of precision and recall offering a balanced measure of the model's performance.

where TP, FP and FN are defined as follows:

- True Positive (TP): predicted contact located within the selected distance threshold from the associated ground truth contact.
- False Positive (FP): predicted contact located beyond the selected distance threshold from the associated ground truth contact.
- False Negative (FN): ground truth contact for which no corresponding predicted contact with the same electrode identifier and contact index is found.

Contact labeling The impact of dilatation on atlas was assessed by evaluating of the number of labeled contacts, unlabeled contacts and unique labels obtained for different dilatation widths. A range of dilatation widths from 0 to 10 mm was tested. Contact labeling was then assessed through visual inspection to verify the spatial consistency of the assigned label, in particular the absence of abrupt change and unlabeled contacts along electrodes.

The co-registration was first assessed by visual inspection. In addition, the Dice Similarity Coefficient (DSC) was computed to quantify the spatial overlap between two binary masks:

$$DSC = \frac{2|A \cap B|}{|A| + |B|} \quad (6.3)$$

where A corresponds to the static mask, B corresponds to the moving mask transformed into the MRI space, $|A \cap B|$ represents the number of voxels shared by both masks, $|A|$ and $|B|$ represent the total number of voxels in each mask.

The brain masks of the MRI and CT scans were generated using the same process described in Chapter 5, with `synth_strip`. The CT-to-MRI mapping obtained during the co-registration step of the algorithm was then applied to the CT brain mask. The same process were performed on the MNI brain mask. The overlap between the transformed CT/MNI brain mask and the MRI brain mask therefore provided a relevant quantitative assessment of the co-registration performed by the algorithm. The Dice score ranges from 0 to 1 where the value of 1 indicates a perfect overlap and a value of 0 indicate no overlap between the two binary mask.

Extending this evaluation to the AAL parcellation was not feasible, as it is a multi-label atlas comprising 116 regions rather than a binary mask, and no manual region-wise segmentations were available for the patient data.

Epileptic seizure detection Due to lack of groundtruth annotations, the performance of the model could not be quantitatively assessed using standard classification metrics. Instead, the analysis relies on a qualitative evaluation of the predicted seizure probabilities to verify the stability and consistency of the model outputs. A temporal analysis of the predicted probabilities is performed to verify that the model outputs are consistent with the temporal dynamics of the corresponding iEEG signal. Moreover, a consistency analysis across segments is also performed to verify that the model outputs are stable across different segments of the same recordings. A histogram of the predicted probabilities is also generated to evaluate the distribution of the model outputs to assess the confidence of the model.

Chapter 7

Results & Discussions

This chapter presents the experimental results obtained to evaluate the performance of the algorithm described in Chapter 6, following the evaluation procedure introduced in Section 6.7. The analysis is structured according to the main stages of the pipeline. First the data processing module is evaluated, as it provides the inputs required by the subsequent steps. The raw ElectroLoc output is then assessed to identify the main localization errors. The post-processing is evaluated to determine the how re-indexation and shift correction improve the final contact localization. Finally, the contact labeling and epileptic seizure detection modules are assessed and discussed.

The proposed pipeline is then compared with existing methods, before discussing its main limitations.

7.1 Performance of the Algorithm

7.1.1 Performance of data preprocessing

The performance of the data preprocessing module was evaluated according to the process described in Section 6.7. Table 7.1 reports the result of the comparison between the extracted entry points and the corresponding ground-truth annotations. Accuracy was defined as the percentage of entry points located within 2 mm of their corresponding ground-truth point.

Patient	< 1 mm	1 mm < d < 2 mm	> 2 mm	Accuracy	Median distance [mm]	Execution time [s]
1	4	4	0	1.00	1.071	63.33
2	3	6	2	0.692	1.55	76.37

Table 7.1: Performance of the data preprocessing module with a comparison between the extracted entry point and the ground-truth.

The comparison between the generated entry points and the ground truth revealed differences in performance between patients. For Patient 1, half of the entry points were accurately localized, with a distance below 1 mm from the ground truth. The remaining entry points were located between 1 mm and 2 mm from the corresponding ground truth positions. For Patient 2, 12 entry points were detected whereas 13 were expected. Among these 12 detected entry points, one did corresponded to any electrode in the ground truth, meaning that two electrode were

completely missed by the algorithm at all. Overall, only three of the entry points were accurately detected while six were located between 1 mm and 2 mm from the ground truth and two were located more than 2 mm away. For both patients, the median distance is approximately 1 mm.

Since the spatial resolution of the CT scan is approximately 1 mm, distance errors below than 2 mm may corresponds to an error of only one or two voxels and therefore be considered acceptable. Moreover, the ground truth was established manually. Each entry point visible the CT scan corresponds to a cluster of voxel rather than to a single point. As a result, if the voxel selected manually as the ground truth is only two voxels away from the point selected by the algorithm, the entry point may still be classified as incorrect, although both points refer to the same electrode. In this context, the algorithm provides a more consistent representation of entry point by selecting the center of mass of the detected voxel cluster. This approach helps to reduce the subjectivity associated with manual extraction.

The two undetected entry points can be explained by the presence of plastic parts in the electrode shaft, which are not visible on the CT scan. This is illustrated in Figure 7.1 which shows the entry points of undetected electrode (named TM and TC) for Patient 2.




Figure 7.1: Undetected electrodes from Patient 2. (pink arrow) electrode TC, (green arrow) electrode TM, (red) surface brain mask.

The number of contacts extracted from the iEEG signals headers was compared with the number of contact reported in the manual ground truth. For Patient 1, 123 contacts were expected from the iEEG signal headers whereas 119 were visible on the CT scan. For Patient 2, it was 138 contacts expected from the iEEG headers compared to 182 contacts identified on the CT scan.

However, the ground truth was manually established by visual inspection of the CT scan which can be challenging due to crossing electrodes, limited image resolution and contacts located outside the brain tissue. Consequently, some contacts may have been missed or incorrectly annotated during the manual process. The algorithm arises therefore as a strong solution to generate the expected number of contacts during the electrode localization module. For Patient 2, the discrepancy may also be related to the potential mismatch between the iEEG recordings and the medical images, as discussed in Chapter 5.

To conclude, the comparison with the manually established ground truth does not always suggest a strong performance at first sight. However, these results must

be interpreted in light of the fact that the manual process used to generate the ground truth corresponds to the original process required to provide the inputs to the ElectroLoc module. The proposed preprocessing algorithm reduces several limitations of this manual procedure, including user-dependent annotation errors, missing contacts and the subjective selection of a voxel to represent an entry point. Moreover, the execution time of the proposed module is considerably shorter than that of the fully manual process. Overall, the algorithm provides a more robust, consistent and efficient approach for generating the inputs required by the subsequent steps of the pipeline.

7.1.2 Performance of electrode localization

The evaluation procedure used to assess the performance of contact coordinate extraction algorithm is detailed in Section 6.7. This analysis aims to assess the quality of the initial contact localization before any post-processing correction. It is conducted at three complementary levels: a comparison with the electrode mask, a visual inspection of the detected contacts and a quantitative comparison with the established ground-truth.

First, the comparison between the detected contacts and the electrode mask, shown in Figure 7.2, indicates that most detected contacts are located within the electrode mask, suggesting a good overall localization performance. A small proportion of detected contacts were found within 1 mm outside the mask which remains acceptable given the spatial resolution of the CT images. However, this criterion is not sufficient to fully validate the localization. Indeed, a contacts may lie within the mask while still being assigned to an incorrect signal channel. Therefore, the mask-based evaluation is mainly informative for contacts located far from the electrode mask which can be confidently considered incorrectly detected.






Figure 7.2: Validation of ElectroLoc contact localization for Patient 1 (left) and Patient 2 (right). Contacts are classified according to their position relative to the electrode mask: inside the mask, outside the mask by less than 1 mm or outside the mask by at least 1 mm.

The visual inspection, illustrated on Figure 7.3a, confirms the overall satisfactory detection of the contacts. Nevertheless, some local inaccuracies remain visible as shown in Figure 7.3b. These inaccuracies are consistent with the limitations

expected by Q. De Laet in particular incorrect contact indexing at the output of the module and the presence of unexpected electrode shift. These observations motivated the introduction of an additional post-processing stage whose performance is evaluated in Section 7.1.4.



(a) Overall localization.



(b) Local misalignment due to shift and indexing errors.

Figure 7.3: Visual comparison between the ElectroLoc output (blue) and the electrode mask (red) from Patient 1.

Finally, the classification metrics computed using the established ground truth are summarized in Table 7.2. These results further support the hypothesis that some contacts located within the electrode mask may still be incorrectly localized or incorrectly assigned. For Patient 1 and Patient 2, no corresponding ground truth was available for 4 and 14 contacts, respectively. These contacts were therefore excluded from the metric computation, which explains the lower number of contacts reported in the table. The Table 7.2 shows a low precision and F1-score as well as a high median distance for both patients. Such distances are too large to ensure a reliable anatomical and channel-wise assignment of the contacts. The low precision is mainly due to the high number of FP. In contrast, the recall is perfect for Patient 1 since all TP are correctly identified and no FN are observed. For Patient 2, the three FN are explained by the discrepancy between the iEEG recordings and the CT scan. The iEEG signal header indicates 15 contacts for the electrode LA whereas the CT scan, from which the ground truth was manually extracted, shows 18 contacts. This mismatch therefore leads to three additional contacts being present in the ground truth but absent from the generated contacts, as illustrated in Figure 7.4.

Patient	TP	FP	FN	Precision	Recall	F1-score	Median Distance
1	4	115	0	0.034	1.00	0.065	25.55 mm
2	5	173	3	0.025	0.625	0.054	24.64 mm

Table 7.2: Classification metrics for contact localization of ElectroLoc output.




Figure 7.4: Visualization of 3 False Negative from electrode LA (see red arrow). (yellow) ground truth, (blue) ElectroLoc's output

Overall, the performance of the contact localization module is consistent with the results reported by Q. De Laet. Most of contacts are correctly detected and globally localized within the electrode mask. However local inaccuracies remain mainly due to electrode shifts and incorrect contact indexing.

The indexing errors may be related to the way contact positions are computed in the ElectroLoc module. Contact coordinates are first generated sequentially before being overwritten by PCA-based recomputation step that determines the depth ordering using the mean position of all detected contacts as an approximation of the CT image center. Since this approximation can deviate from the true image center and because the contact positions are computed twice, the resulting indices may appear to be inconsistently distributed along the electrodes. In addition, the observed shifts may result from the loss function that does not penalize shifts as expected and allowing the optimization to converge to a local optimum.

7.1.3 Performance of post-processing of ElectroLoc

The previous evaluation highlighted two main sources of local inaccuracies in the raw ElectroLoc output: incorrect contact indexing and electrode shifts. To improve the reliability of the final localization, a post-processing stage was introduced. It consists of two complementary steps: a PCA-based re-indexation step, which corrects the ordering of contacts along each electrode, and a shift correction step, which adjusts contact positions when a longitudinal displacement is detected.

Re-indexation

The re-indexation step is evaluated using visual inspection and comparison with established ground-truth. Since the step only modifies the contact indices and does not change the contact coordinates, the comparison with the electrode mask was not repeated as it would provide no additional information.

The visual inspection shows that the contact indices follow the expected ordering with the deepest contact assigned $c_id = 0$ and the index increasing progressively as the contact becomes more superficial. The comparison between the selected method (the PCA-based approach) and the Euclidean distance-based method, shown in Figure 7.5, supports the approach's choice. Both methods perform well in standard configurations where the most superficial contact corresponds to the entry point

(see electrode DA in Figure 7.5). However, the PCA-based method more robustly handles partially inserted electrodes ensuring a consistent ordering along the electrode trajectory even when some contacts lie outside the brain. In contrast, the Euclidean distance-based method leads to incorrect ordering in such configurations (see electrode DB in Figure 7.5).




Figure 7.5: Comparison of ElectroLoc output, Euclidean based-distance and PCA reindexing for Patient 1.

The quantitative comparison with the ground truth, summarized in Table 7.3 confirms the improvement brought by the re-indexation step. The number of TP considerably increases compared with the raw ElectroLoc output, leading to substantial improvement in the precision and F1-score. This indicates that many contact that were previously detected at plausible spatial locations were incorrectly indexed and are therefore correctly reassigned after re-indexation. As expected, the recall remains unchanged as expected for Patient 1 since no FN were present before re-indexation. For Patient 2, the recall increases because several contact that previously considered incorrectly assigned are now correctly matched to the ground-truth. The median distance to the corresponding ground-truth is also significantly reduced after the re-indexation and becomes close to the CT image resolution. This is consistent with the visual inspection and confirms that the re-indexation step effectively corrects a large part of contact errors.

Patient	TP	FP	FN	Precision	Recall	F1-score	Median Distance
1	81	38	0	0.681	1.00	0.81	0.82 mm
2	106	72	3	0.596	0.972	0.739	1.13 mm

Table 7.3: Classification metrics for indexation output.

Shift correction

The shift correction step is evaluated through the three steps introduced in Section 6.7: comparison with electrode mask, visual inspection and quantitative evaluation with the established ground-truth.

As mentioned previously, the comparison with the electrode mask provides an indication of how strongly this step reduce the number of incorrect contacts. However, this criterion does not allow for definitive conclusions to be drawn, in particular, contacts located within the mask cannot be assumed to be correctly positioned. Figure 7.6 shows a slight reduction in the number of contacts located more than 1 mm outside the electrode mask after the shift correction. The remaining contacts remaining in this category may corresponds to contacts located outside the brain tissue, since the electrode mask is strictly defined within the brain. Therefore, their presence does not necessarily indicate an incorrect shift correction. The increase in the number of contacts located within 1 mm of the mask should also be interpreted with caution, as this distance is close to the spatial resolution of the CT images and may correspond to minor voxel-based deviations. These assumptions are therefore further assessed through comparison with the ground truth.




Figure 7.6: Validation of shift correction of Patient 1 (left) and Patient 2 (right). Contacts are classified according to their position relative to the electrode mask: inside the mask, outside the mask by less than 1 mm or outside the mask by at least 1 mm.

The visual inspection results are illustrated in Figure 7.7 and show an improvement in the alignment of the detected contacts with the electrode mask after the shift correction. The double criteria used for shift correction shows strong performance in handling different electrode configurations. In particular, the algorithm

effectively corrects shifts for standard electrodes, partially implanted electrodes, and electrodes with a contact-less shaft, ensuring a robust and adapted shift correction across various scenarios.




Figure 7.7: Shift correction from Patient 1.

The classification metrics are summarized in Table 7.4. The Precision and F1-score are considerably higher after the shift correction which is consistent with the visual inspection and validates the assumption that the majority of contacts are located within 1 mm of the electrode mask are correctly positioned. The median error distance is also reduced after shift correction, reaching a value consistent with the spatial resolution of the CT images.

Patient	TP	FP	FN	Precision	Recall	F1-score	Median Distance
1	109	10	0	0.916	1.00	0.956	0.79 mm
2	157	21	3	0.961	0.983	0.972	0.76 mm

Table 7.4: Classification metrics for shift correction output.

In conclusion, the proposed post-processing algorithm significantly improves the localization performance of the raw ElectroLoc output. Overall, the final corrected localization shows high performance with high precision and F-1 score value as well as a median error distance remaining sufficiently small. This improvement is mainly achieved through the combination of two complementary corrections. First, the re-indexation step significantly improves the contact localization by correcting the contact ordering. The PCA-based approach for re-indexation shows a robust performance across different electrode configurations, ensuring a consistent and accurate representation of contact indices along the electrode trajectory. Second, the proposed shift correction algorithm also shows strong performance. It provides a robust correction adapted to different electrode configurations including partially implanted electrode, electrode with a contact-less shaft and standard electrodes.

7.1.4 Performance of contact labeling

Contact labeling was conducted into three main steps including registration, atlas dilatation and anatomical labeling. The process used to evaluate the performance of these steps are described in Section 6.7.

Registration

Three coregistrations were performed including CT to MRI, MNI to MRI and finally, atlases to MNI. The resulting transformations are illustrated in Figure 7.8. The

visual inspection of the co-registration results showed strong alignment between the different images and atlases.




Figure 7.8: Coregistration results for Patient 1. (top) CT scan in red overlaid on the MRI. (bottom) undilated AAL atlas overlaid on the MRI.

This qualitative assessment was confirmed by the DSC value, summarized in Table 7.5 which indicate a high degree of overlap after co-registration. However, the DSC never reached 1 since a perfect voxel-by-voxel overlap is not expected when comparing images or masks derived from different source and modalities. For example, the post-implantation CT contains electrode that are absent from the pre-implantation MRI. Similarly, the MNI atlas corresponds to an average anatomical template and therefore cannot perfectly match the individual anatomy of each patient. Overall, the DSC values indicate accurate and robust co-registration.

Patient	CT to MRI		MNI to MRI	
	Before	After	Before	After
COG_011	0.000	0.982	0.661	0.963
COG_004	0.011	0.981	0.401	0.958

Table 7.5: Dice Similarity Coefficients before and after CT-to-MRI and MNI-to-MRI co-registrations.

Atlas dilatation

The impact of atlas dilatation is illustrated in Figure 7.9. As expected, the number of labeled contacts increases with the dilatation width. It reaches a plateau once all contacts located within the brain are assigned a label. However, when some contacts lie outside the brain tissue, the number of unlabeled contacts cannot reach 0 as the dilatation is constrained by the brain mask and therefore limited to the brain boundaries. Similarly, the number of unique anatomical labels slightly increases with the dilatation width before stabilizing. This behavior indicates that atlas dilatation enables the assignment of additional regions while remaining stable for

larger dilatation values. Importantly, the absence of a continuous increase suggest that the dilatation does not introduce excessive anatomical variability and thereby preserving a consistent and controlled spatial labeling.




Figure 7.9: Impact of atlas dilatation on contact labeling of Patient 1

Atlas dilatation improves the labeling coverage by capturing more contacts. However aggressive dilatation increases the risk of assigning the contact to incorrect neighboring regions. The choice of dilatation width therefore reflects a precision-recall tradeoff: increasing dilatation improves coverage but also raises the likelihood of boundary spillover.

Anatomical labeling

Figure 7.10 presents the anatomical labels assigned to each electrode using a dilated atlas with a width of 6 mm. As expected, the remaining "Unknown_0" labels corresponds to contacts with high indices which are positioned more superficially along the electrode according the ordering. These contacts are therefore likely located outside the brain tissue. Overall, the labeling appears consistent with no abrupt changes observed along individual electrodes or within the same anatomical region.

Exception are observed for Electrode OS from Patient 1, and Electrode LA, TC, TB, TA and FB from Patient 2, where sudden changes in anatomical labels are observed and may indicate a local misalignment or segmentation issue. The Figure 7.11 shows the anatomical labels overlapped on the MRI for the patient for these identified electrodes. It can be observed that the abrupt change does not correspond by a local misalignment but rather to a contact location at the boundary of a gyrus.




Figure 7.10: Anatomical labeling along each electrode for both Patients.




Figure 7.11: Visualization of Electrode LA, TB, FB, TC, TA and OA. Red arrow highlights the abrupt changes.

7.1.5 Performance of epileptic seizure detection

The process used to evaluate the performance of the epileptic seizure detection module is described in Section 6.7. Since no annotated seizure was available for the the analyzed recording, the evaluation of this module was limited to a qualitative assessment of the predicted seizure probabilities over time. Figure 7.12 shows the predicted seizure probabilities for a 5-minute recordings. The predicted probabilities are relatively stable throughout the recording for both patients with no clear transient increase suggesting pre-ictal activity. This observation is further supported by the distribution shown in Figure 7.13 which is concentrated around low probability values. These results suggest that the model consistently classified the analyzed recordings as interictal. However, in the absence of a labeled seizure segment, no quantitative evaluation could be conducted. Therefore, this analysis only provides a preliminary validation of the module behavior on non-seizure data, rather than a complete assessment of its seizure detection performance.




Figure 7.12: Predicted seizure probabilities for 5-minutes recordings from both patients. The raw and smoothed probabilities remain below the decision threshold, with no transient increase suggesting pre-ictal activity.




Figure 7.13: Distribution of predicted seizure probabilities for the 5-minutes recordings from both patients. The probabilities are concentrated at low values, supporting the absence of detected pre-ictal activity in the analyzed recordings.

7.1.6 Performance of overall algorithm

The overall algorithm can be divided into two main components: the contacts extractions and the seizure detection. Both components behave independently and their performance is assessed separately as described in the previous sections. The contact extraction component which includes the data preprocessing, the ElectroLoc module and the contact labeling module shows a strong overall performance. It significantly improves in the automation of input generation and the post-processing of the detected contacts. Moreover, even when the estimated entry point coordinates are not fully accurate, the ElectroLoc module is still able to reconstruct the correct electrode. The execution time of each subcomponent is summarized in Table 7.6. The overall execution time of this component is more than 30 minutes for both patients. The contact labeling step is the most time-consuming step representing approximately 98% of the total execution time due to the multiple coreregistrations performed. The optimization of this step of coregistration is therefore a key point to reduce the overall execution time of this part of the pipeline.

Patient	Contact extraction step	Time processing
Patient 1	Preprocessing	63.33 s
Patient 1	ElectroLoc	7.52 s
Patient 1	Contact labeling	32.988 min
Patient 2	Preprocessing	76.37 s
Patient 2	ElectroLoc	4.28 s
Patient 2	Contact labeling	47.414 min

Table 7.6: Time execution of contact extraction module

The seizure detection module was evaluated in previous section and showed a limited predictive performance. However, no quantitative validation could be performed as no ground truth was available to assess the correctness of the prediction. The execution time is summarized in Table 7.7. This high computational cost is mainly due to the input format required by the model which consists of scalograms for each channel and each time segment. This limitation is amplified in this present dataset: while the dataset used by T. Bary contained approximately 20 channels, the recordings considered in this work include more than 100 channels. Consequently, the amount of input data increases substantially resulting in much longer execution time. As a result, the module is not yet efficient for recordings of this duration, which are approximately 30 minutes in this work.

Patient	Seizure Detection	Time processing
Patient 1	Recording of 5s	276.40 s
Patient 1	Recording of 5min	120.06 min
Patient 2	Recording of 5s	217.95 s
Patient 2	Recording of 5min	106.98 min

Table 7.7: Time execution of epileptic seizure detection module

Another limitation is that the predictions are not currently linked to the anatomical location of the contacts. The model classifies each time segment as *pre-ictal* or *inter-ictal* but does not provide information about the spatial origin of the predicted seizure-related activity. Integrating the seizure detection output with the anatomical information provided by the contact extraction module represents therefore an important perspective for improving the overall pipeline.

7.2 Comparison to Existing Methods

A qualitative comparison between Brainstorm and the proposed pipeline can be conducted. Based on [24], this comparison can be focused on four main aspects: the general approach, the electrode localization method, the post-processing strategy, and the anatomical labeling procedure.

Regarding the general approach, Brainstorm provides a graphical user interface implemented in MATLAB, although it can be used without requiring a MATLAB license. This interface makes the tool accessible and user-friendly for clinical and

research applications. In contrast, the proposed pipeline is currently implemented in Python and does not include a graphical user interface. However, its main objective is to reduce manual intervention as much as possible by automating several steps of the electrode localization and labeling process.

In Brainstorm, contact localization relies on the manual selection of the deepest contact and the entry point for each electrode. The intermediate contacts are then placed along the electrode trajectory according to the inter-contact distance provided by the user. In contrast, the proposed pipeline includes an automated extraction of the entry points. Manual intervention is only required during the electrode identification step, which also remains partly user-dependent in Brainstorm. In addition, the proposed pipeline models the electrode trajectory using a 3-D parabolic representation. This could be more appropriate for implanted electrodes, since their trajectory can be curved due to bending during implantation.

The two approaches also differ in their correction strategies. In Brainstorm, correction remains primarily manual and requires expertise in MRI interpretation. In the proposed pipeline, additional post-processing steps were implemented to reduce localization errors, including shift correction and contact re-indexation when incorrect indexing is detected. These steps aim to reduce user-dependent errors and improve the consistency of the localization results.

Finally, anatomical labeling is integrated directly into the Brainstorm graphical interface. In the proposed pipeline, anatomical labeling is performed using an atlas-based approach, with an optional label dilation step constrained by the brain mask. This dilation allows unlabeled contacts located close to labeled anatomical regions to be assigned to the nearest relevant label while limiting the risk of assigning labels outside the brain.

To conclude, Brainstorm appears as a robust and user-friendly tool for electrode localization, applicable not only to stereo-EEG electrodes but also to subdural electrodes. Its graphical interface facilitates user interaction and manual verification. The proposed pipeline, on the other hand, focuses on reducing manual intervention and user-dependent variability through automation of the entry point extraction, parabolic electrode modeling, shift correction, and atlas-based anatomical labeling. Therefore, although no direct quantitative comparison was performed in this work, the proposed pipeline may be particularly relevant when the priority is to minimize manual interaction and improve the consistency of electrode localization for curved intracranial trajectories.

7.3 Limitations

7.3.1 Validation

The implementation of the pipeline and its evaluation were conducted on a limited dataset comprising two patients which introduces several limitations to the validation. First, the same dataset was used for both development and evaluation. It introduces potential bias and overfitting issues as the pipeline may be specifically

adapted to these cases. Second, the small sample size restricts the generalizability of the results and may not capture the full variability of electrode implantations and anatomical configurations. Third, for the second patient, a mismatch was observed between the number of electrodes identified on the CT scan and the number of channels available in the iEEG recordings. This discrepancy may be explained by the fact that the CT scan was acquired one month prior to the iEEG recordings raising the possibility of an incorrect correspondence between the two datasets. Consequently, the performance observed for this patient needs to be interpreted with caution may not be representative of the true performance of the pipeline. Validation on a larger and more diverse dataset is therefore necessary to confirm these results and assess the robustness of the proposed approach.

7.3.2 Condition for shifted electrode selection

During the implementation of the shift correction module, several conditions were tested to identify electrodes requiring correction (more details in Appendix A). The condition leading to the best performance was selected. The selected criterion was based on the distance between the deepest contact and the electrode mask :

$$\text{dist}(c_{id} = 0, \text{electrode mask}) < 2 \text{ mm.} \quad (7.1)$$

This criterion allows shifted electrodes to be detected while discarding electrodes for which the estimated entry point is likely incorrect because it does not rely on the entry point location. However, one specific shift configuration is not handled by this condition, as illustrated in Figure 7.14. In this configuration, the deepest contact remains close to the electrode mask even though the electrode is shifted. This is not considered a major limitation in the present study, since this type of shift was not observed in the evaluated patients. It may therefore either be uncommon in practice or simply absent from the available dataset.




Figure 7.14: Shift configuration not handled by the selected condition

7.3.3 Epileptic seizure detection

The current implementation of the epileptic seizure detection module relies on the use of a pre-trained model which was trained on an EEG signal dataset. This approach presents several limitations. As iEEG and scalp EEG signals have different characteristics (higher signal-to-noise ratio, higher amplitude) and biomarkers (LVFA, IEDs and HFOs) highlighted in Section 3.2, the model is not adapted to predict seizures from iEEG data. The elaboration of a training/fine-tuning phase using

iEEG data is therefore necessary to improve the performance and reliability of the seizure detection module. However, this process is currently hindered by the lack of available iEEG datasets with annotated seizures which are necessary for training and validating the model. A theoretical strategy to address this issue is detailed in the [Section 8.2](#).

Chapter 8

Perspectives and potential improvements

8.1 Perspectives for contact extraction

8.1.1 Electrode identification

Electrode identification remains the only module of the pipeline that requires manual intervention. The lack of standardized nomenclature during the implantation prevents full automation of this step. It therefore leads to an increased processing time, potential user-dependent errors and the need for technical expertise. In addition, this process depends on SEEG implantation plan which may not always be available or easily interpretable by non-expert users.

Future work on the pipeline should focus on automating this module. The use of standardized conventions such as SENSA proposed by Stone et al. (2021) and presented in Section 4.2 could represent a promising improvement. Such conventions need to be adopted during the implantation phase and integrated into the pipeline. A review of the main electrode-naming conventions currently used in clinical practice could be conducted to develop an algorithm capable of recognizing these conventions and automatically identifying the electrodes.

8.1.2 Electrode localization

The improvements introduced in this pipeline to address the current limitations of the ElectroLoc module showed strong performance and contributed to a more robust electrode localization process. In particular, the pre- and post-processing steps significantly enhance the overall consistency and accuracy of the results. In future implementations, these steps could be directly embedded within the ElectroLoc module itself rather than applied externally. Integrating shift correction directly into ElectroLoc would also prevent it from appearing as a separate post-processing step. This would result in a more efficient, consistent, and user-friendly pipeline.

Furthermore, the current ElectroLoc module requires the voxel coordinates of the entry point as input which are subsequently converted into world coordinates. Given the coordinate extraction process already provides both voxel and world representation, this requirement could be reconsidered. Allowing the use of world coordinates directly as a input would eliminate the need for an additional conversion

step. This modification would improve both computational efficiency and clarity of the processing chain.

8.1.3 Shift correction

The current method performs shift correction along the electrode trajectory. The improvements introduced in this step demonstrate strong performance and enable precise localizations of the contacts. However, some residual inaccuracies can still be observed.

In future work, the shift correction approach could be extended beyond longitudinal adjustments to also account for orientation errors and lateral deviations of the electrodes. Incorporating such corrections could allow a more comprehensive alignment and further improve the localization accuracy.

8.1.4 Co-registration

The parameter *diffeomorph* was kept to False for all three co-registrations (CT-to-MRI, MNI-to-MRI and atlas-to-MNI). For the CT-to-MRI, both images corresponded to the same patient. In this case the main differences between the volumes are expected to arise from position, orientation and inter-modality contrast rather than anatomical deformation. Moreover, since the CT scan was acquired post-implantation, applying a non-linear deformation could alter the relative geometry of the contacts which is undesirable for electrode localization.

In contrast, setting *diffeomorph* to True could be relevant for the MNI-to-MRI registration in future work. Indeed the MNI template and the patient MRI are both anatomical images of comparable modality, but they correspond to different anatomies. A non-linear diffeomorphic registration could therefore better account for inter-subject anatomical variability and improve the alignment between the template space and the patient-specific MRI.

Moreover, the execution time of the co-registration step should be significantly reduced. At present, this step represents the most time-consuming part of the pipeline, requiring more than 30 minutes to process a single patient. Therefore, optimizing this step would be an important perspective for future work, especially if more computationally expensive registration strategies such as diffeomorphic are considered.

8.1.5 Graphical User Interface

The creation of a graphical user interface (GUI) would represent a significant improvement to the current pipeline. Although the proposed pipeline requires limited human intervention once configured, running the algorithm and adjusting its parameters still requires some familiarity with Python, or at least the ability to execute commands from a terminal.

The comparison with Brainstorm also suggested that a GUI could significantly improve the usability of the proposed pipeline. Such an interface would make the workflow more accessible, simplify parameter selection, and reduce the technical barrier for users without programming experience. Its integration therefore facilitate the adoption of the pipeline in clinical or research environments.

8.2 Perspectives for epileptic seizure detection

The epileptic seizure detection module developed in this work is currently limited by the lack of labeled iEEG data. Although the transformer-based classification provides a promising framework for automatic seizure detection, its performance cannot be fully quantified or optimized without a sufficient number of annotated segments. A theoretical strategy based on Few-Shot Learning is introduced in this Section. A presentation of the FSL literature and a more detailed description of these strategies are provided in Appendix C.

FSL represents a promising direction to address this limitation. Unlike conventional deep learning approaches requiring large-scale annotated datasets, FSL methods are designed to adapt to new tasks or patients using only a few samples. This property is particularly relevant for epileptic seizure detection given the strong inter-subject variability and data scarcity inherent to iEEG signals. Based on FSL approaches identified in the literature [38] [39] [40], a theoretical strategy could be proposed to improve the current module. Rather than training a new model from scratch, the objective would be to transfer knowledge from the pretrained EEG-based transformer and adapt it to intracranial recordings using a limited number of labeled iEEG examples. This adaptation could be structured at three levels.

At the data and algorithmic level, the pretrained transformer encoder would be used to leverage prior knowledge acquired from previous patients. The encoder would be frozen and used as a feature extractor while only the final classification layers would be fine-tuned using a few labeled pre-ictal and interictal iEEG segments. This strategy would reduce the number of learnable parameters and limit the risk of overfitting.

At the prototype-based classification level, the original MLP decision head would be replaced by a prototype-based decision rule. The core idea is to represent each class by a single representative vector, called a prototype, computed in the latent space of the pretrained transformer encoder. Concretely, a few labeled segments from each class would be preprocessed through the current pipeline including the encoder. The resulting latent representations would then be averaged separately for each class to produce two prototypes. Classifying a new unlabeled iEEG segment would then become straightforward: Euclidean distance between the processed segment and each prototype is computed. The predicted label corresponds to the class of the closest prototype.

At the semi-supervised level, the abundant unlabeled iEEG recordings could be exploited through a teacher-student framework. Two augmented versions of each input scalogram would be fed respectively to a student and a teacher model sharing the same transformer-based architecture. The student would be trained via both a supervised classification loss on labeled data and a consistency loss against teacher predictions on unlabeled data. The teacher weights would be updated as an Exponential Moving Average (EMA) of the student weights, providing stable pseudo-labels throughout training.

Finally, a temporal post-processing step could be added to reject isolated predictions not supported by neighboring segments, improving robustness given that epileptic seizures typically span several consecutive windows.

Overall, two concrete adaptation strategies can be derived from these levels. The first one, illustrated in Figure 8.1, combines the algorithmic and semi-supervised levels: the pretrained encoder is kept as a frozen feature extractor while the MLP decision head is adapted through the teacher-student framework using both labeled and unlabeled iEEG recordings. The second one, illustrated in Figure 8.2, combines the algorithmic and prototype-based levels: the same frozen encoder feeds a prototype-based decision rule built from a small number of labeled iEEG-specific pre-ictal and interictal examples.




Figure 8.1: Semi-supervised teacher-student adaptation strategy for iEEG seizure prediction. The pretrained EEG encoder is kept frozen and used as a feature extractor, while the student model is trained using classification and consistency losses. The teacher model is updated from the student weights using an exponential moving average.




Figure 8.2: Prototype-based adaptation strategy for iEEG seizure prediction. The pretrained EEG encoder is used as a frozen feature extractor to project iEEG scalograms into a latent representation space. Class prototypes are then built from a limited number of labeled pre-ictal and interictal examples, and new segments are classified according to their distance to these prototypes.

Conclusion

The objective of this thesis was to provide an overview of the main steps involved in iEEG signal preprocessing while introducing the clinical and methodological context of epilepsy and intracranial EEG recordings. In addition, this work aimed to contribute to the development of preprocessing pipelines for iEEG signals acquired from epileptic patients.

The literature review highlighted a lack of standardized, automated and generalizable methods for iEEG signal preprocessing. Existing approaches often require several manual inputs and expert interventions which limits their reproducibility and their accessibility to non-specialist users.

The proposed approach distinguishes itself from existing methods by reducing the required inputs and human intervention, while introducing specific contributions such as an automatic entry point extraction, contact index, shift correction and anatomical labeling. The pipeline starts by generating two binary masks: an electrode mask and a brain surface mask. Their intersection is then used to compute the coordinates of the electrode entry points. The number of contacts associated with each electrode is extracted directly from the iEEG signal metadata. At this stage, electrode identification is still performed manually, in order to associate each detected electrode entry point with its corresponding clinical electrode name. Electrode localization is then performed using ElectroLoc, developed in previous work by Q. De Laet [6], as a black box localization module. The contact labeling module assigns an anatomical label to each contact with an optional atlas dilatation. It also corrects contact indices and possible shift along the electrode trajectory. Finally, an adaptation of the transformer-based classification method developed by T. Bary [5] is implemented for epileptic seizure detection. Overall, the proposed pipeline significantly reduces the time and manual effort required to localize, index and anatomically label electrode contacts.

Future work could further improve the pipeline through several enhancements. These include full-automation of the pipeline through modifications to the electrode identification module, fine-tuning the transformer-based method for iEEG signals using few-shot learning, reducing the execution time of the co-registration step and developing a graphical user interface (GUI).

Bibliography

- [1] World Health Organization, *Epilepsy*, <https://www.who.int/news-room/fact-sheets/detail/epilepsy>, Fact sheet, accessed: 16 March 2026, 2024.
- [2] K. Ivankovic, A. Principe, R. Zucca, M. Dierssen, and R. Rocamora, “Methods for Identifying Epilepsy Surgery Targets Using Invasive EEG: A Systematic Review,” *Biomedicines*, vol. 12, no. 11, p. 2597, Nov. 2024, ISSN: 2227-9059. DOI: [10.3390/biomedicines12112597](https://doi.org/10.3390/biomedicines12112597) Accessed: May 25, 2026.
- [3] S. Kovac, V. N. Vakharia, C. Scott, and B. Diehl, “Invasive epilepsy surgery evaluation,” *Seizure*, vol. 44, pp. 125–136, Jan. 2017, ISSN: 10591311. DOI: [10.1016/j.seizure.2016.10.016](https://doi.org/10.1016/j.seizure.2016.10.016) Accessed: Mar. 9, 2026.
- [4] L. Harris and H. Angus-Leppan, “Epilepsy: Diagnosis, classification and management,” *Medicine*, vol. 48, no. 8, pp. 522–528, Aug. 2020, ISSN: 13573039. DOI: [10.1016/j.mpmed.2020.05.001](https://doi.org/10.1016/j.mpmed.2020.05.001) Accessed: Mar. 3, 2026.
- [5] T. Bary and B. Macq, “Transformers for EEG classification: Architectures, pre-training, and applications to epileptic seizure forecasting,” M.S. thesis.
- [6] Q. D. Laet, B. Macq, R. Kupers, T. Bary, and H. Verhaeghe, “Semi-Automated Detection of sEEG Electrodes in Computed Tomography Scans,” M.S. thesis.
- [7] A. K. Shah and S. Mittal, “Invasive electroencephalography monitoring: Indications and presurgical planning,” *Annals of Indian Academy of Neurology*, vol. 17, no. Suppl 1, S89–S94, Mar. 2014, ISSN: 0972-2327. DOI: [10.4103/0972-2327.128668](https://doi.org/10.4103/0972-2327.128668) Accessed: Mar. 9, 2026.
- [8] *Epilepsy and Seizures — National Institute of Neurological Disorders and Stroke*.
- [9] R. S. Fisher et al., “ILAE Official Report: A practical clinical definition of epilepsy,” *Epilepsia*, vol. 55, no. 4, pp. 475–482, Apr. 2014, ISSN: 0013-9580, 1528-1167. DOI: [10.1111/epi.12550](https://doi.org/10.1111/epi.12550) Accessed: Mar. 9, 2026.
- [10] I. E. Scheffer et al., “ILAE classification of the epilepsies: Position paper of the ILAE Commission for Classification and Terminology,” *Epilepsia*, vol. 58, no. 4, pp. 512–521, Apr. 2017, ISSN: 0013-9580, 1528-1167. DOI: [10.1111/epi.13709](https://doi.org/10.1111/epi.13709) Accessed: Mar. 14, 2026.
- [11] *NIST: X-Ray Mass Attenuation Coefficients - Section 2*. Accessed: Apr. 9, 2026. [Online]. Available: <https://physics.nist.gov/PhysRefData/XrayMassCoef/chap2.html>
- [12] Advacom, *X-ray/CT Technique*. Accessed: May 25, 2026. [Online]. Available: <https://www.radiology.expert/en/modules/xrayct-technique/types-of-ct-techniques/>

- [13] *Computed Tomography (CT)*. Accessed: Apr. 9, 2026. [Online]. Available: <https://www.nibib.nih.gov/science-education/science-topics/computed-tomography-ct>
- [14] *Magnetic Resonance Imaging (MRI)*. Accessed: May 25, 2026. [Online]. Available: <https://www.nibib.nih.gov/science-education/science-topics/magnetic-resonance-imaging-mri>
- [15] L. Dricot, *Lgbio1115: Introduction aux neurosciences – thématique neuroimagerie part 1*, Course lecture, Université catholique de Louvain, accessed: March 2026, Mar. 2021.
- [16] *Figure 8.1 -Une image IRM pondéré T1 et T2*. Accessed: May 25, 2026. [Online]. Available: https://www.researchgate.net/figure/Une-image-IRM-pondere-T1-et-T2_fig43_344380023
- [17] K. E. Davis, R. K. Kana, and J. P. Szaflarski, “Intracranial Recordings in Epilepsy: An Approach to Understanding the Neural Mechanisms of Social Cognition,” *Current Neurology and Neuroscience Reports*, vol. 25, no. 1, p. 44, Dec. 2025, ISSN: 1528-4042, 1534-6293. DOI: [10.1007/s11910-025-01431-9](https://doi.org/10.1007/s11910-025-01431-9) Accessed: Dec. 17, 2025.
- [18] M. C. Tenorio, “Data Preprocessing and Machine Learning for Intracranial Electroencephalography,”
- [19] K. IIDA and H. OTSUBO, “Stereoelectroencephalography: Indication and Efficacy,” *Neurologia medico-chirurgica*, vol. 57, no. 8, pp. 375–385, Aug. 2017, ISSN: 0470-8105. DOI: [10.2176/nmc.ra.2017-0008](https://doi.org/10.2176/nmc.ra.2017-0008) Accessed: Mar. 16, 2026.
- [20] J. Parvizi and S. Kastner, “Human Intracranial EEG: Promises and Limitations,” *Nature neuroscience*, vol. 21, no. 4, pp. 474–483, Apr. 2018, ISSN: 1097-6256. DOI: [10.1038/s41593-018-0108-2](https://doi.org/10.1038/s41593-018-0108-2) Accessed: Mar. 18, 2026.
- [21] L. Zhong et al., “Epileptic seizure prediction in intracranial EEG using critical nucleus based on phase transition,” *Computer Methods and Programs in Biomedicine*, vol. 226, p. 107 091, Nov. 2022, ISSN: 01692607. DOI: [10.1016/j.cmpb.2022.107091](https://doi.org/10.1016/j.cmpb.2022.107091) Accessed: Mar. 16, 2026.
- [22] R. J. Staba, M. Stead, and G. A. Worrell, “Electrophysiological Biomarkers of Epilepsy,” *Neurotherapeutics*, vol. 11, no. 2, pp. 334–346, Apr. 2014, ISSN: 18787479. DOI: [10.1007/s13311-014-0259-0](https://doi.org/10.1007/s13311-014-0259-0) Accessed: Mar. 17, 2026.
- [23] S. Gallotto and M. Seeck, “EEG biomarker candidates for the identification of epilepsy,” *Clinical Neurophysiology Practice*, vol. 8, pp. 32–41, Dec. 2022, ISSN: 2467-981X. DOI: [10.1016/j.cnp.2022.11.004](https://doi.org/10.1016/j.cnp.2022.11.004) Accessed: Mar. 16, 2026.
- [24] F. Tadel, S. Baillet, J. C. Mosher, D. Pantazis, and R. M. Leahy, “Brainstorm: A User-Friendly Application for MEG/EEG Analysis,” *Computational Intelligence and Neuroscience*, vol. 2011, p. 879 716, 2011, ISSN: 1687-5265. DOI: [10.1155/2011/879716](https://doi.org/10.1155/2011/879716) Accessed: May 25, 2026.
- [25] *NITRC: MRICroGL: Tool/Resource Info*. Accessed: May 25, 2026. [Online]. Available: <https://www.nitrc.org/projects/mricrogl/>
- [26] *MisterI*. Accessed: May 25, 2026. [Online]. Available: <https://www.benoitscherrer.com/MisterI/index.html>

- [27] Flexikon-Pictures, *10-20-System*. Accessed: May 25, 2026. [Online]. Available: <https://www.doccheck.com/de/detail/photos/42650-10-20-system>
- [28] N. Axmacher, Ed., *Intracranial EEG: A Guide for Cognitive Neuroscientists* (Studies in Neuroscience, Psychology and Behavioral Economics). Cham: Springer International Publishing, 2023, ISBN: 978-3-031-20909-3 978-3-031-20910-9. DOI: [10.1007/978-3-031-20910-9](https://doi.org/10.1007/978-3-031-20910-9) Accessed: May 25, 2026.
- [29] S. Stone, J. R. Madsen, J. Bolton, P. L. Pearl, V. Chavakula, and E. Day, “A Standardized Electrode Nomenclature for Stereoelectroencephalography Applications,” *Journal of Clinical Neurophysiology*, vol. 38, no. 6, pp. 509–515, Nov. 2021, ISSN: 0736-0258. DOI: [10.1097/WNP.0000000000000724](https://doi.org/10.1097/WNP.0000000000000724) Accessed: Apr. 11, 2026.
- [30] *The AAL atlas views: (A) left lateral, (B) left medial, (C) anterior,...* Accessed: May 25, 2026. [Online]. Available: https://www.researchgate.net/figure/The-AAL-atlas-views-A-left-lateral-B-left-medial-C-anterior-D-superior-E_fig1_259462969
- [31] S. Hagen et al., “Intracranial EEG referencing for large-scale category-selective mapping in the human ventral occipito-temporal cortex,” *Imaging Neuroscience*, vol. 3, imag_a_00479, Feb. 2025, ISSN: 2837-6056. DOI: [10.1162/imag_a_00479](https://doi.org/10.1162/imag_a_00479) Accessed: Mar. 30, 2026.
- [32] R. Hussein, S. Lee, and R. Ward, “Multi-Channel Vision Transformer for Epileptic Seizure Prediction,” *Biomedicines*, vol. 10, no. 7, p. 1551, Jun. 2022, ISSN: 2227-9059. DOI: [10.3390/biomedicines10071551](https://doi.org/10.3390/biomedicines10071551) Accessed: May 28, 2026.
- [33] National Electrical Manufacturers Association (NEMA), *Digital Imaging and Communications in Medicine (DICOM) Standard*, Rosslyn, VA, USA: National Electrical Manufacturers Association. [Online]. Available: <https://www.dicomstandard.org/>
- [34] *NIfTI (Neuroimaging Informatics Technology Initiative) Reader/Writer*. Accessed: May 25, 2026. [Online]. Available: <https://docs.safe.com/fme/html/FME-Form-Documentation/FME-ReadersWriters/nifti/nifti.htm>
- [35] A. Hoopes, J. S. Mora, A. V. Dalca, B. Fischl, and M. Hoffmann, “SynthStrip: Skull-stripping for any brain image,” *NeuroImage*, vol. 260, p. 119474, Oct. 2022, ISSN: 10538119. DOI: [10.1016/j.neuroimage.2022.119474](https://doi.org/10.1016/j.neuroimage.2022.119474) Accessed: May 12, 2026.
- [36] E. Larson et al., *MNE-Python*, Zenodo, Apr. 2026. DOI: [10.5281/zenodo.19666955](https://doi.org/10.5281/zenodo.19666955) [Online]. Available: <https://doi.org/10.5281/zenodo.19666955>
- [37] *Continuous Wavelet Transform (CWT) — PyWavelets Documentation*. Accessed: May 25, 2026. [Online]. Available: <https://pywavelets.readthedocs.io/en/latest/ref/cwt.html>
- [38] J. Nazari, A. Motie Nasrabadi, M. B. Menhaj, and S. Raiesdana, “Epilepsy seizure prediction with few-shot learning method,” *Brain Informatics*, vol. 9, no. 1, p. 21, Dec. 2022, ISSN: 2198-4018, 2198-4026. DOI: [10.1186/s40708-022-00170-8](https://doi.org/10.1186/s40708-022-00170-8) Accessed: May 19, 2026.

-
- [39] A. Burrello, K. Schindler, L. Benini, and A. Rahimi, “One-shot Learning for iEEG Seizure Detection Using End-to-end Binary Operations: Local Binary Patterns with Hyperdimensional Computing,” arXiv:1809.01926, arXiv, Sep. 2018. DOI: [10.48550/arXiv.1809.01926](https://doi.org/10.48550/arXiv.1809.01926) arXiv: [1809.01926](https://arxiv.org/abs/1809.01926) [eess.SP]. Accessed: May 25, 2026.
- [40] Z. Zhang, X. Li, F. Geng, and K. Huang, “A Semi-Supervised Few-Shot Learning Model for Epileptic Seizure Detection,” in *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, Nov. 2021, pp. 600–603. DOI: [10.1109/EMBC46164.2021.9630363](https://doi.org/10.1109/EMBC46164.2021.9630363) Accessed: May 20, 2026.
- [41] S. An, S. Kim, P. Chikontwe, and S. H. Park, “Few-Shot Relation Learning with Attention for EEG-based Motor Imagery Classification,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Las Vegas, NV, USA: IEEE, Oct. 2020, pp. 10 933–10 938, ISBN: 978-1-7281-6212-6. DOI: [10.1109/IROS45743.2020.9340933](https://doi.org/10.1109/IROS45743.2020.9340933) Accessed: May 20, 2026.

Appendix A

Condition for shift electrode

During the implementation of the pipeline, several criteria were tested to identify electrodes requiring shift correction. The first criterion was based on the distance between the deepest contact and the electrode mask:

$$\text{dist}(c_{id} = 0, \text{electrode mask}) < 2 \text{ mm.} \quad (\text{A.1})$$

The second criterion was based on the distance between the most superficial contact and the estimated entry point:

$$\text{dist}(c_{id} = \text{max}, \text{entry point}) < 5 \text{ mm.} \quad (\text{A.2})$$

Three selection strategies were evaluated: the criterion defined in Equation A.1, the criterion defined in Equation A.2, and the union of both criteria. Based on these conditions, five electrode configurations were identified, as illustrated in Figure A.1.




Figure A.1: Electrode configurations considered during the shift correction step. The generated electrode contact are shown in green, the electrode mask in red and the scalp surface in gray. (a)-(c) show shifted electrode requiring longitudinal correction along the electrode trajectory. (d) illustrates a contact-less shaft configuration. (e) shows a partially implanted electrode for which the most superficial contact does not necessarily coincide with the scalp entry point.

The Table A.1 summarizes which electrode configurations are classified as shifted according to each tested condition. Condition 1 only classified configuration a and b as shifted. It did not identify configuration c, but correctly discarded configurations d and e, which corresponds to electrodes that do not require shift correction. Conversely, Condition 2 included four configurations and discarded configuration b. This condition therefore incorrectly classified the configuration d and e as shifted.

Condition	Config a	Config b	Config c	Config d	Config e
1	X	X			
2	X		X	X	X
$1 \cup 2$	X	X	X	X	X

Table A.1: Electrode configurations classified as shifted according to each condition.

It was observed that the criterion based on the entry point distance, defined in Equation A.2, may lead to incorrect classifications. This occurs when the estimated entry point is located more than 5 mm away from the most superficial contact, even though the electrode itself is correctly localized. This case is illustrated in Figure A.2.




Figure A.2: Electrode configurations considered for shift correction.

Since the preprocessing module did not reach 100% precision, it was preferable to select the condition that was independent of other steps of the algorithm. The selected condition was therefore the first condition defined by Equation A.1.

Appendix B

Code and Repository Overview

All the code developed in the frame of this thesis is available at the following adress:
<https://github.com/maedegand/MasterThesis.git>.

B.1 Structure of the repository

The main branch of the repository is composed of two directories. The content of these files is visualized in Figure B.1. The `src` directory contains the python scripts for the contact extraction module and seizure detection module. Each of them has their own main python file and therefore enables to run the module separately.

The `data` directory stores the AAL atlas, MNI template, the pretrained weights from [5] and the model built in this work.




Figure B.1: Simplified structure of the GitHub repository.

The `README.md` file in the repository contains all relevant information regarding the use or the development of the algorithm, particularly:

- a quick installation guide covering both the proposed pipeline and the ElectroLoc module from [6];

- details about the electrode identification module which requires an user intervention;
- a presentation of the project's organisation;
- a explanation of optionnal args for personalization of the algorithm.

For more details about the repository or the algorithm, the reader is kindly redirected to the `README.md` file.

Appendix C

Few-shot Learning approaches for iEEG seizure detection

The epileptic seizure detection module developed in this work is currently limited by the lack of labeled iEEG data. Few-Shot Learning (FSL) represents a promising direction to address this limitation. FSL aims to obtain good predictive performance from a limited number of labeled examples making it particularly relevant for epileptic seizure detection given the strong inter-subject variability and data scarcity inherent to iEEG signals. Several studies have investigated FSL strategies for EEG and iEEG-based classification tasks, four of which are presented below.

The first approach, proposed in [41], focuses on relation-based few-shot learning for EEG-based motor imagery classification. The model is composed of three main modules: an embedding module, an attention module and a relation module. The embedding module first extracts relevant features from the EEG signal using convolutional layers. The attention module then identifies the N class-representative features that are most relevant to the query sample by assigning them an attention score. Finally, the relation module computes a relation score between the query and each class representation and the class with the highest score is selected as the predicted label. As a preprocessing step, the signal is divided into frequency range to emphasize different types of information located in different frequency range.

A second approach, presented in [38], combines several strategies to reduce the need for patient-specific labeled data. At the data level, the amount of usable training data is increased through windowing with variable overlap between signal segments. At the level model, the number of learnable parameters is reduced by freezing part of the neural network during training. At algorithmic level, prior knowledge learned from other patients is exploited and the model is fine-tuned using only a few samples from the target patient. These strategies combine the advantages of a generalizable model with patient-specific adaptation.

The third approach, proposed in [39], relies on Local Binary Patterns (LBP) and Hyperdimensional (HD) computing for iEEG seizure detection. LBP first transforms iEEG signals into binary codes encoding local temporal variations. More specifically, it encodes whether the signal increases or decreases over time. During interictal periods, LBP codes are generally more evenly distributed across the possible patterns.

In contrast, ictal periods tend to be characterized by a dominant LBP pattern due to the slow and asymmetric oscillations typically observed during seizures. The LBP codes are then converted into high-dimensional binary vectors using HD computing and combined to obtain a global representation of brain activity for each temporal window. One prototype for each class is built and new unlabeled segments are classified according to their Hamming distance to these prototypes. The class associated with the closest prototype is assigned as the predicted label. A post-processing step is performed to reject isolated predictions. Instead of considering each short segment independently, the algorithm evaluates the temporal consistency of the predictions across consecutive windows.

The last approach, proposed in [40], exploits both labeled and unlabeled EEG/iEEG recordings through a semi-supervised teacher-student framework for epileptic seizure detection. Both models share the same CNN-based architecture and receive slightly different augmented versions of the same input. In this method the signals are first transformed into frequency domain using Fast Fourier transform for iEEG signals and Short-Time Fourier Transform for EEG signal. A Mel-scale power spectrum is then extracted and used as input to the classifier. Although the student and teacher models share the same architecture, their weights are updated differently. The student model is updated by backpropagation using supervised classification loss on labeled segments and the consistency loss on unlabeled segments where the teacher prediction is used as a reference target in this second case. This allows the model to learn from annotated data while also exploiting unlabeled iEEG recordings. The teacher model is not directly updated by backpropagation. Instead, its weights are updated as an Exponential Moving Average (EMA) of the student weights:

$$\theta_{\text{teacher}}^{(t)} = \alpha \theta_{\text{teacher}}^{(t-1)} + (1 - \alpha) \theta_{\text{student}}^{(t)}, \quad (\text{C.1})$$

where α controls the smoothing of the update. This strategy makes the teacher model more stable than the student model and allows it to provide more reliable pseudo-labels for unlabeled examples. The algorithm is illustrated in Figure C.1.




Figure C.1: Framework of semi-supervised FSL for epileptic seizure detection using Mean Teacher. [40]

Based on the FSL approaches identified in the literature, a theoretical strategy is proposed to adapt the pretrained transformer-based model originally trained on scalp EEG data to iEEG signals. The main core of this approach is to adapt the pretrained model to the iEEG domain at three levels: an data/algorithmic

adaptation level, prototype-based classification level and a semi-supervised learning level. The objective would not be to train a new model from scratch but rather to transfer the knowledge from scalp EEG data and adapt it to intracranial recordings using a limited number of labeled iEEG examples.

At data/algorithmic level, the pretrained transformer could be used as a general feature extractor. The current preprocessing pipeline would be kept unchanged: the iEEG signals would be first segmented, transformed into scalogram, divided into patches and embedded into tokens before being provided as input to the transformer-based model. The encoder of the pretrained model could then be frozen, meaning that its weights would not be updated during training. Only the final classification layers would be fine-tuned using a few labeled pre-ictal and inter-ictal iEEG segments. This strategy would reduce the number of trainable parameters and limit the risk of overfitting.

A second level would consist of creating iEEG-specific class prototypes in the latent space of the transformer. Inspired by the prototype-based FSL approaches, this strategy would represent each class using only a small number of labeled examples. In this proposed strategy, a few labeled pre-ictal and inter-ictal iEEG segments would first be preprocessed by the pretrained transformer encoder. The resulting latent representations would be used to define class prototypes. In this configuration, the original MLP decision head would be replaced by a prototype-based decision rule. Instead of mapping the latent representation to a binary output through the MLP classifier, each segment would be classified according to its distance to the class prototypes. An illustration of the change is displayed in Figure C.2




Figure C.2: Proposed prototype-based adaptation of the original transformer-based model for iEEG seizure prediction. The pretrained encoder is kept as a feature extractor, while the original MLP decision head is replaced by a prototype-based decision rule comparing the distance between extracted latent representation and iEEG-specific pre-ictal and inter-ictal prototypes.

More specifically, the embeddings extracted from the labeled pre-ictal segments would be averaged to create an pre-ictal prototype, while the embeddings extracted from the labeled interictal segments would be averaged to create an interictal prototype. A new unlabeled iEEG segment would then be projected into the same latent space and compare with both prototype using a distance metric such as Euclidean distance. The predicted label would correspond to the class of the closest prototype. This level could be considered as an adaptation of the prototype based idea

proposed in [39]. However, instead of constructing prototypes from LBP codes and hyperdimensional vectors, the proposed strategy would build prototypes from the embeddings extracted by the transformer encoder.

A third level would consist in exploiting the method implemented by [40] using the unlabeled iEEG recordings through semi-supervised teacher-student strategy. This approach would be relevant because long iEEG recordings are often available whereas only a limited number of segments are manually annotated. In this strategy, the iEEG signals would first undergo the same preprocessing as in the current pipeline including segmentation and scalogram generation. Then two slightly different augmented versions of each scalogram would be generated. This choice keeps the augmentation closer to the signal domain and makes it easier to ensure that seizure-related patterns are not artificially distorted. These two versions would be provided respectively to the student and teacher model which would share the same transformer-based architecture. The rest of the training procedure would follow the teacher-student scheme described previously, where the teacher provides pseudo-labels on unlabeled data while student would be trained to produce consistent predictions from the augmented inputs. The teacher weights would then be updated from the student weights, typically using an EMA.

Finally, a temporal post-processing step could be added to improve the robustness of the predictions, since epileptic seizures usually extend over several consecutive windows: predictions could be rejected if they are not supported by neighboring segments.

Overall, this strategy would allow the model to preserve the general time-frequency representation learned from EEG data while adapting the final decision process to the iEEG domain. The algorithmic level would reuse the pretrained transformer encoder as a general feature extractor for intracranial signals. Two possible adaptation strategies could then be considered. The first one would replace the original MLP decision head with a prototype-based decision rule providing an explicit few-shot mechanism based on iEEG-specific pre-ictal and inter-ictal prototypes. The second would keep the MLP decision head and adapt it through a semi-supervised teacher-student framework using both labeled and unlabeled iEEG recordings. The two theoretical frameworks proposed in this Section are illustrated in Figure 8.1 and 8.2.

UNIVERSITÉ CATHOLIQUE DE LOUVAIN

École polytechnique de Louvain

Rue Archimède, 1 bte L6.11.01, 1348 Louvain-la-Neuve, Belgique | www.uclouvain.be/epl