
Electronic Thesis and Dissertation Repository

9-20-2024 10:30 AM

Segmentation of intracranial electrode contacts using convolutional neural networks

Arun Thurairajah, *Western University*

Supervisor: Lau, Jonathan, *The University of Western Ontario*

Co-Supervisor: Khan, Ali, *The University of Western Ontario*

A thesis submitted in partial fulfillment of the requirements for the Master of Science degree in Neuroscience

© Arun Thurairajah 2024

Follow this and additional works at: <https://ir.lib.uwo.ca/etd>



Part of the [Computational Neuroscience Commons](#), [Neurosciences Commons](#), and the [Neurosurgery Commons](#)



This work is licensed under a [Creative Commons Attribution 4.0 License](#).

Recommended Citation

Thurairajah, Arun, "Segmentation of intracranial electrode contacts using convolutional neural networks" (2024). *Electronic Thesis and Dissertation Repository*. 10478.

<https://ir.lib.uwo.ca/etd/10478>

This Dissertation/Thesis is brought to you for free and open access by Scholarship@Western. It has been accepted for inclusion in Electronic Thesis and Dissertation Repository by an authorized administrator of Scholarship@Western. For more information, please contact wlsadmin@uwo.ca.

Abstract

For patients with intractable epilepsy the surgical placement of intracranial electrodes can better localize the seizure onset zone. Stereoelectroencephalography (SEEG) is one technique, where depth electrodes made of multiple contacts record activity in the brain. The precise interpretation of recordings requires the anatomical localization of each contact. Contact positions can be manually localized or determined using semi-automated algorithms. This thesis works towards the automation of SEEG contact localization with a 3D U-Net, a deep learning architecture optimized for biomedical image segmentation. The first chapter will introduce the clinical workflow for SEEG, available tools, and the potential role of deep learning. The second chapter will cover the proposed algorithm and validation methods. The last two chapters will present the accuracy of the U-Net in contrast to approaches currently employed in the clinic. Overall, the error and accuracy of the proposed method compares favorably in an independent set of clinical data. Future work will look to continue to optimize performance.

Keywords

Stereoelectroencephalography (SEEG), epilepsy, deep learning, computer vision, computed tomography, image processing, neuroimaging

Summary for Lay Audience

Epilepsy is a neurological disorder characterized by abnormal electrical activity known as seizures. A subset of patients diagnosed with epilepsy do not respond to medication. For some, surgical removal of a part of the brain determined to be the “seizure onset zone” can significantly reduce the severity and number of seizures and can be curative. Potential candidates for surgery are determined during a pre-surgical evaluation, where diagnostic tests aim to determine any structural deficits or abnormal brain activity. When non-invasive methods fail to localize the seizure onset zone, the implantation of electrodes directly within the brain can be used to record electrical activity. Stereoelectroencephalography (SEEG) is one method using depth electrodes made of multiple contacts to simultaneously study multiple deep brain regions. Analysis of recordings requires the anatomical position of each contact in the brain, as viewed in post-operative imaging. This localization process is typically done manually or with semi-automated tools, in a process that can be both time consuming and labor intensive. The thesis works towards the automated localization of SEEG contacts using a U-Net, a deep learning model that is used in most computer vision tasks with medical images. The first chapter covers the clinical workflow, the prerequisite knowledge for medical imaging, image processing, and deep learning, and the current tools for SEEG contact localization. The second chapter will outline the implementation of the U-Net model and how the model will be validated. Finally, it will compare errors to the current tool employed in the clinic. Overall, the error and accuracy of the proposed method compares favorably in an independent set of clinical data, while reducing the amount of manual intervention required. Future work will look to continue to optimize performance.

Co-Authorship Statement

Arun Thurairajah wrote the manuscript and implemented the codebase for the project. Dr. Jonathan Lau and Dr. Ali Khan supervised the project, contributed to the conception and design, and provided revisions to the thesis manuscript.

Dr. Greydon Gilmore led data acquisition and preparation, the manual annotation of SEEG contact localizations, and contributed to the conception of the study. Segments of Dr. Gilmore's code was also used to prepare figures.

Mauricio Cespedes Tenorio developed the labelling algorithm described in Section 2.4.2 and reviewed the codebase for the U-Net model training and inference.

Alaa Taha provided valuable insights to the study design, code review, and manuscript preparation. Feyi Ogunsanya reviewed code for the U-Net model training.

Dr. Jason Kai, Alaa Taha, Mauricio Cespedes Tenorio, and Brad Karat provided suggestions and edits to the content and clarity of the thesis.

Acknowledgments

This thesis would not have been possible without the help of numerous individuals. First, I would like to thank Dr. Jonathan Lau and Dr. Ali Khan for introducing me to this field and taking me on to be a part of their labs. Jon, thank you for your mentorship and guidance throughout my master's degree. I truly appreciate the care you have taken to overseeing my personal development and your unwavering support despite the ups and downs of this journey. Ali, thank you for your advice and expertise throughout the two years. Your open-door policy and positive attitude have fostered a collaborative lab environment that encourages teamwork and innovation.

I would like to thank my advisory committee Dr. Ana Suller-Marti and Dr. Lyle Muller, and the program representative Dr. Walter Rushlow, for their valuable time and insights over the past two years.

To the past and present members of the Khan and AIMS Lab, thank you for welcoming me with open arms to Western. There are too many names to thank, but I will fondly remember the friendship and memories over the past two years. I would like to mention the following individuals who played a large part in seeing this thesis to completion. Mau, Alaa, Dr. Greydon Gilmore, and Dr. Jason Kai, thank you for your mentorship over the past two years.

Finally, I would like to thank my parents and my brothers for their continuous love and support.

Table of Contents

Abstract.....	ii
Summary for Lay Audience.....	iii
Co-Authorship Statement.....	iv
Acknowledgments.....	v
Table of Contents.....	vi
List of Tables.....	ix
List of Figures.....	x
List of Appendices.....	xiii
List of Abbreviations.....	xiv
Chapter 1.....	1
1 Introduction.....	1
1.1 Stereoelectroencephalography (SEEG).....	3
1.1.1 Clinical workflow for SEEG implantation.....	4
1.1.2 The study of epilepsy with SEEG.....	6
1.2 Medical imaging for stereoelectroencephalography.....	7
1.2.1 Computed Tomography (CT).....	7
1.2.2 Magnetic Resonance Imaging (MRI).....	9
1.3 Image processing.....	10
1.3.1 Coordinate systems.....	10
1.3.2 Image resolution.....	11
1.3.3 Image registration.....	12
1.3.4 Image interpolation.....	14
1.3.5 Brain atlases and templates.....	15
1.3.6 Image segmentation.....	15

1.4	Deep learning for image segmentation	17
1.4.1	Convolutional neural networks (CNNs)	18
1.4.2	Loss functions	19
1.4.3	U-Net for medical image segmentation	21
1.5	Current approaches for SEEG contact localization	23
1.5.1	Existing challenges for contact localization	26
1.5.2	Deep learning for SEEG contact segmentation	29
1.6	Open science and reproducibility for SEEG contact segmentation	30
1.7	Thesis overview	30
	Chapter 2.....	32
2	U-Net Implementation Methods	32
2.1	Patient Data.....	32
2.1.1	CT imaging	32
2.1.2	Ground truth semi-automatic contact localization.....	32
2.2	Pre-processing workflow for U-Net.....	35
2.3	U-Net model architecture.....	37
2.4	Post-processing	37
2.4.1	Probabilistic non-max suppression	38
2.4.2	Contact labeling	39
2.5	Model validation	40
2.5.1	Traditional image processing for contact localization	41
2.5.2	Comparison to uncorrected SEEG Assistant (SEEGA) placements.....	42
	Chapter 3.....	44
3	U-Net Implementation Results.....	44
3.1	Effect of patch size on model prediction	44
3.2	Localization accuracy of U-Net approach	47

3.3	Localization error for U-Net compared to SEEG Assistant.....	50
3.3.1	Distribution of errors.....	51
3.3.2	Localization error due to electrode bending	52
3.3.3	Localization error at a subject level	54
3.3.4	Error by inter-contact spacing.....	55
3.3.5	Absolute error in x, y, and z dimensions	57
3.4	Localization error of traditional image processing approach	59
Chapter 4	62
4	Discussion	62
4.1	Limitations	65
Chapter 5	69
5	Conclusion	69
	Bibliography	70
	Appendices.....	85
	Curriculum Vitae	90

List of Tables

Table 1-1 Toolboxes for SEEG contact localization using CT. <i>Proprietary software or tools without open-access code were not hyperlinked. N/A indicates there was no reported accuracy. *N/A for iElvis from Groppe et al., 2017, indicates a reported error for ECoG contacts, but no measure for SEEG. *N/A for BrainQuake indicates error was assessed based on the degree contact positions deviated from the planned trajectory and how closely the predicted inter-contact distance matched the known spacing of the electrode. *N/A for CranialVault indicates that the tool is proprietary, with the details regarding test size, programming language, and reported accuracy not openly available.</i>	24
Table 3-1 Training Summary and Loss Metrics	44
Table 3-2 Contact localization accuracy of the U-Net model in the validation set, based on patch size (n = 1910 contacts). <i>Compared to weighted centroids of artifacts with 3-D connectivity from thresholded CT at 2500 Hu.</i>	46
Table 3-3 Contact localization accuracy of the U-Net model in the validation set, following the contact-labelling algorithm (n = 1910 contacts).	47
Table 3-4 Contact localization accuracy of the U-Net model on the test set. (n = 35 subjects, 4020 contacts, 402 electrodes). Compared to a traditional image processing approach of thresholding the CT at 2500 Hu and masking the intracranial volume.	48
Table 3-5 Average Euclidean error by contact, where n indicates number of contacts for each electrode type. <i>Matched localizations for both the U-Net and SEEGA methods, where both errors were below 2 mm.</i>	57
Table 3-6 Mean absolute error in each dimension by contact and method. <i>n = 35 subjects, 3757 contacts.</i>	58

List of Figures

Figure 1-1. Illustrative diagram of an SEEG electrode and contacts, as seen in post-operative imaging. <i>The first panel makes the distinction between the SEEG electrodes and contacts. The second shows how contact length, diameter, and spacing are defined. Captured using 3D Slicer version 5.6.1, a medical image viewer (Fedorov et al., 2012).</i>	2
Figure 1-2. Visual representation of depth electrodes and different inter-contact spacing from two manufacturers, AdTech and DIXI medical.	5
Figure 1-3. Post-operative CT image. <i>Seen on the left is the axial (top-down) view, with the right highlighting the different components and corresponding HU. Captured using 3D Slicer version 5.6.1, a medical image viewer (Fedorov et al., 2012).</i>	8
Figure 1-4. Standard 2-D U-Net Architecture, adapted from Ronneberger et al., 2015. <i>The blue boxes represent the feature maps of the image, with the width representing the number of feature maps and the height representing the spatial size.</i>	22
Figure 1-5 Example of challenging contact localizations, with the manual “ground-truth” placements in green. <i>The first panel from the left shows a vertically implanted electrode, where the contact artifacts appear to be “fused” together. The second panel shows an example of both a vertically implanted electrode and an example of where contacts are at the level of the skull. The bottom panel demonstrates an example of where electrodes curve from the planned trajectory and must be manually adjusted for most contact localization algorithms.</i>	28
Figure 2-1 Flowchart diagram for SEEG imaging pre-processing workflow.....	34
Figure 2-2. 3-D patches of various sizes taken from post-operative CT. <i>Sagittal views are shown with the voxel sizes listed above.</i>	36
Figure 2-3. CT of an electrode implanted vertically, with the U-Net model prediction in blue. <i>On the right are the “ground-truth” coordinates for each contact in green. Individual contacts cannot be separated in the electrode artifact or by the deep learning model.</i>	38
Figure 2-4 Flowchart of contact labelling for U-Net predictions.	40

Figure 3-1 Training and validation loss curve for U-Net trained with 64 voxel isotropic patch sizes. *Dice + Cross Entropy loss values shown on the y axis, with the epoch number on the x axis.* 45

Figure 3-2 Training and validation loss curve for U-Net trained with 96 voxel isotropic patch sizes. *Dice + Cross Entropy loss values shown on the y axis, with the epoch number on the x axis.* 45

Figure 3-3 U-Net model contact localizations (in red) and ground-truth manual placements (in green) for subject P126, plotted in MNI152NLin2009cSym template space (n = 120 contacts). *Side (right hemisphere), front, and top views are shown. The model achieved a perfect F1-score, indicating all contacts were found within 1 mm with no false positives or negatives.* 49

Figure 3-4 False negatives or contacts missed by the U-Net model, plotted in MNI152NLin2009cSym template space (n = 183/4020). *From left to right, side (right hemisphere), top, and front views are shown.* 50

Figure 3-5 Distribution of Euclidean error (in mm), comparing the semi-automated (SEEG Assistant) localizations to those matched by the automated U-Net method. n = 3912 contacts. *Dashed lines represent the median error, while curves indicate the cumulative frequencies. U-Net median = 0.46 mm, SEEGA median = 0.43 mm.* 51

Figure 3-6 Bland-Altman plot of Euclidean error (in mm), comparing the U-Net method to the SEEGA localizations. *Thresholded at 2 mm for both methods, n = 35 subjects, 3757 contacts. Upper and lower limits were set at $\pm 1.96 * 1$ standard deviation.* 52

Figure 3-7 Visual schematic of the perpendicular distance between the manual contact position and electrode axis used to quantify the degree of electrode bending. *a* represents the vector between the entry point and SEEG contact, while *b* represents the vector between the entry and target points..... 53

Figure 3-8 Perpendicular distance of each contact’s position from the electrode trajectory and the contact localization error for both the U-Net and SEEGA methods. *Pearson’s*

correlations are included for both methods ($r = 0.36$ for SEEGA, 0.01 for the U-Net). Linear regression lines for both methods are also plotted with a 95% CI band..... 54

Figure 3-9 Localization error by subject for the U-Net model and SEEGA, thresholded at 2 mm. Ordered by the mean error for the U-Net model (increasing from left to right). Significance was assessed using a Wilcoxon signed-rank test with Bonferroni correction. $n = 35$ subjects, 3757 contacts.

Figure 3-10 Error by electrode inter-contact spacing for U-Net and SEEGA methods. Significance was assessed using a Wilcoxon signed-rank test, with Bonferroni correction. $N = 35$ subjects.

Figure 3-11 Error in each dimension by contact for the U-Net approach compared to SEEGA. Significance was assessed using a Wilcoxon signed-rank test with Bonferroni correction. $n = 35$ subjects, 3757 contacts.

Figure 3-12 Distribution of Euclidean error (in mm), comparing the automated (U-Net) approach to the matched localizations by the traditional image processing algorithm. Thresholded for errors below 2 mm, $n = 2501$ contacts. Median U-Net = 0.45 mm (dashed red line), Median Thresholded CT = 0.14 mm (dashed green line).

Figure 3-13 Distribution of Euclidean error (in mm), comparing the semi-automated SEEG Assistant (SEEGA) tool to the matched localizations by the traditional image processing algorithm. $n = 2604$ contacts. SEEGA median = 0.44 mm (dashed red line), Thresholded CT median = 0.14 mm (dashed green line).

Figure 3-14 Contacts missed (false negatives) by the traditional image processing approach and the U-Net. In red are the false negatives by the U-Net model, with blue representing the false-negatives by the outlined image processing approach.

Figure 4-1 Examples of three bending electrodes in subject D149. In red are the U-Net model predictions, green are the manual localizations, and blue are the SEEGA predictions. The CT electrode/contact segmentation artifact can be seen in yellow (thresholded at 2500 HU). Captured with 3D Slicer, version 5.6.1 (Fedorov et al., 2012).

List of Appendices

Appendix A: Summary of localization accuracy and error for 3-D U-Net by subject, alongside the error from SEEG Assistant (SEEGA) in the test set.....	85
Appendix B: Ethics Approval.....	89

List of Abbreviations

ANN	Artificial Neural Network
BIDS	Brain Imaging Data Structure
CE	Cross Entropy
CNN	Convolutional Neural Network
CT	Computed Tomography
DBS	Deep Brain Stimulation
DOF	Degrees of Freedom
DSC	Dice-Sørensen similarity coefficient
ECoG	Electrocorticography
EEG	Electroencephalography
FOV	Field-of-View
HU	Hounsfield Unit(s)
IED	Interictal Epileptiform Discharge
iEEG	Intracranial Electroencephalography
LFP	Local Field Potential
LHSC	London Health Sciences Centre
MONAI	Medical Open Network for AI
MRI	Magnetic Resonance Imaging
SEEG	Stereoelectroencephalography

SEGA SEEG Assistant

STP Standard pressure and temperature

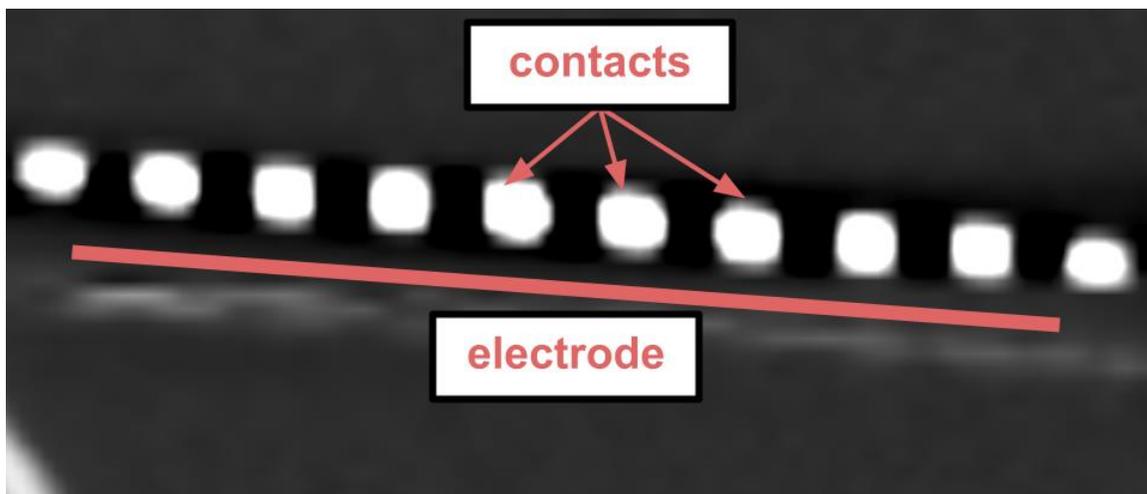
WCE Weighted Cross Entropy

Chapter 1

1 Introduction

Stereotactic neurosurgery is a subspecialty of neurosurgery concerned with precisely targeting brain structures. The most common applications involve the implantation of electrodes or probes to deliver electrical stimulation or other therapy (e.g. ablative) in specific regions to provide therapeutic relief. In other instances, electrodes serve a diagnostic role and record different areas in the brain, including stereoelectroencephalography (SEEG), where depth electrodes are implanted to localize seizures in patients with drug-resistant epilepsy.

SEEG recordings offer superior spatial and temporal resolution compared to noninvasive methods such as scalp electroencephalography. Electroencephalography uses electrodes placed on the scalp to record the activity from the cortical surface transmitted through the scalp and skull (Zijlmans et al., 2019). SEEG electrodes in contrast are implanted directly within the soft tissue of the brain. An individual SEEG electrode consists of multiple metallic contacts along a shaft between 0.5 and 2 mm in length (Figure 1-1). Contacts record the local field potential (LFP), which is the electric potential generated by the synaptic activity from populations of neurons. A single SEEG contact records the electrical activities within a 5-mm radius, sampling from about 30 mm^3 of the brain (Frauscher et al., 2024).



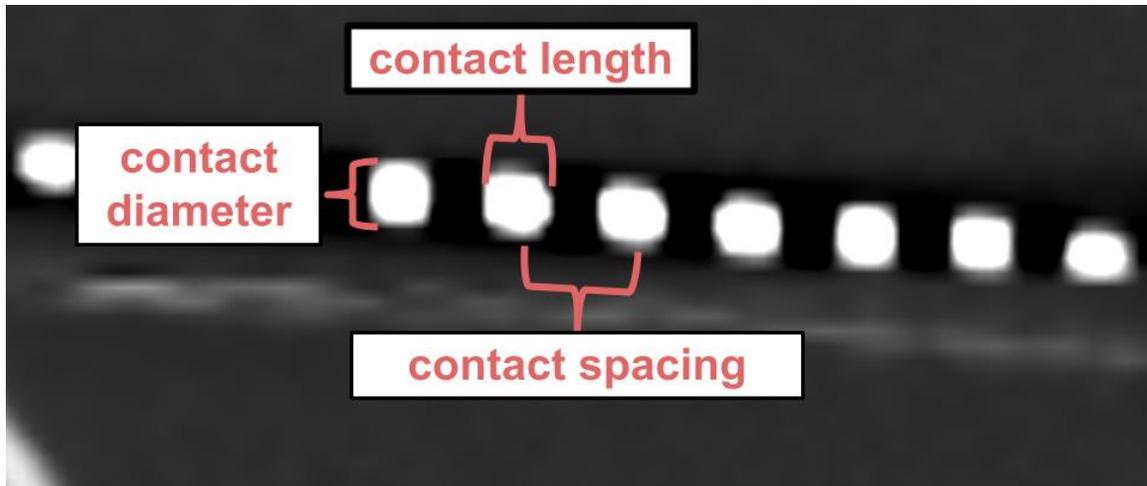


Figure 1-1. Illustrative diagram of an SEEG electrode and contacts, as seen in post-operative imaging. *The first panel makes the distinction between the SEEG electrodes and contacts. The second shows how contact length, diameter, and spacing are defined. Captured using 3D Slicer version 5.6.1, a medical image viewer (Fedorov et al., 2012).*

An accurate interpretation and understanding of these signals require the anatomical localization of each contact as viewed in the post-operative imaging. Traditionally, contacts are manually localized by annotating imaging artifacts to confirm their position. Several computational toolboxes exist to expedite this process, ranging from mathematical approaches that consider the physical properties of each electrode, image processing techniques identifying radiological contact artifacts, or machine learning algorithms (Davis et al., 2021; Granados et al., 2021; Lucas et al., 2024; Narizzano et al., 2017). These methods are all semi-automated, requiring the manual placement of at least two contacts per electrode and a thorough review of localizations. This results in an incredibly labor- and resource-intensive task for each subject, prone to human error, further making the group-wise analysis across multiple subjects untenable.

The thesis works towards an approach for the automated localization of SEEG contacts through deep learning. It addresses a gap in the literature by contrasting the accuracy with semi-automated methods used today. This introductory chapter will cover the history of stereotactic surgery and SEEG, the role of medical imaging, image processing for stereotactic surgery, an overview of deep learning, and current methods of contact localization.

1.1 Stereoelectroencephalography (SEEG)

A discussion around stereoelectroencephalography begins with the origins of stereotactic surgery. In 1908, Sir Victor Horsley and Robert Clark developed the first stereotactic frame to target and lesion deep cerebellar nuclei in animals (Horsley & Clarke, 1908). Nearly 40 years later, Ernest A. Spiegel and Henry T. Wycis designed a stereotactic frame for humans, marking the birth of human stereotactic surgery (Spiegel et al., 1947). Jean Talairach of Saint Anne Hospital in Paris, France was a contemporary of Spiegel and Wycis, who worked to refine the stereotactic methodology. A psychiatrist and neurosurgeon, his goals were to establish precise stereotactic definitions for deep-brain structures and improve the accuracy of stereotactic surgery. Talairach made several significant contributions to the field, notably defining reproducible neuroanatomical landmarks, creating the first human stereotactic *atlas* (i.e. labeling) of deep brain structures, and establishing a proportional grid system for the human brain known as the “Talairach coordinate space” (Talairach et al., 1952; Talairach, 1957; Talairach, 1967).

This cartographical exploration of the brain was the precipice for stereoelectroencephalography. Jean Bancaud, a neurologist and electroencephalographer, joined Saint Anne Hospital in 1952 and suggested that Talairach apply his stereotactic principles in treating epilepsy (Zanello et al., 2022). At this time, Wilder Penfield and Herbert Jasper had established pre-operative cortical stimulation for epilepsy surgery in Montreal. Their method of electrocorticography (ECoG) provided recordings from the cortical surface. It would use the activity observed between seizures (*interictal* discharges) to define the *epileptogenic zone*: the brain region that, once resected, would produce seizure freedom (Parrent, 2009; Penfield & Jasper, 1954). However, Bancaud and Talairach believed that understanding a seizure's origin and organization would better localize the epileptogenic focus. Therefore, analyzing network dynamics during a seizure or the *ictal* period would require precisely placing intracerebral electrodes within the structures of interest (Kahane et al., 2006). Talairach's *repérage radiologique indirect* (indirect radiological investigation) enabled the minimally invasive and accurate implantation of depth electrodes. Electrodes were implanted through burr holes made in the skull, avoiding the craniotomy required for ECoG. Combined with the patient's clinical symptoms, Bancaud's meticulous analysis of ictal recordings, and Talairach's atlases, the team could form a surgical plan tailored

to the patient's "anatomy-electro-clinical" correlations. The SEEG procedure was first published in 1962, the name referencing the three-dimensional nature and simultaneous recording of multiple brain regions (Talairach et al., 1962).

1.1.1 Clinical workflow for SEEG implantation

The primary purpose of SEEG implantation is to identify the origin of epileptic activity and define the epileptogenic zone. Removing or disconnecting this region during epilepsy surgery can be an effective treatment option to reduce seizures (Wiebe et al., 2001). However, the epileptogenic zone is a hypothetical area of the brain that is only confirmed by the post-operative outcome. Given that the actual location of the epileptogenic zone cannot be determined pre-operatively, a pre-surgical evaluation is needed to estimate the extent of this area. At this stage of the evaluation, several diagnostic tools identify different conceptual regions contributing to the epileptogenic network. Overlap of these areas indicates the likely location of the epileptogenic zone (Zijlmans et al., 2019). When there is insufficient information from non-invasive tools, intracranial electrodes are temporarily implanted to better understand the network. Also referred to as intracranial electroencephalography (iEEG), the two techniques for implantation are ECoG using subdural strip or grid electrodes and SEEG with depth electrodes.

Contacts in ECoG are organized along the grid or strip, and activity is recorded from the cortical surface (Zijlmans et al., 2019). Subdural grids offer extensive sampling of the cortical surface and are implanted through a craniotomy, which requires removing a portion of the skull. Due to this craniotomy, common complications with subdural grids include infections (5.3%) and intracranial hemorrhages (4.0%), as observed in a systematic review of implantations (Arya et al., 2013). Subdural strips can be placed using a smaller incision or burr hole in the skull, but samples a much smaller region of the brain (Wu et al., 2024). In contrast, SEEG electrodes are implanted to allow the sampling from deep brain structures across multiple layers through burr holes made in the skull. In comparison to ECoG, the method of implantations is considered less invasive and is associated with fewer intraoperative complications (Joswig, Steven, et al., 2018; Joswig et al., 2020). The number of contacts, contact size, and inter-contact spacing along an SEEG electrode can vary depending on the manufacturer. Examples include 10-contact (0.86 mm in diameter) electrodes with 3-, 4-, 5-, or 6-mm contact spacing, and 5-18-contact (0.8 mm

diameter) electrodes with 2 mm contact spacing (Dixi Medical, Besançon, France) (Cardinale et al., 2016). A visual illustration of a few electrode examples can be seen in Figure 1-2.

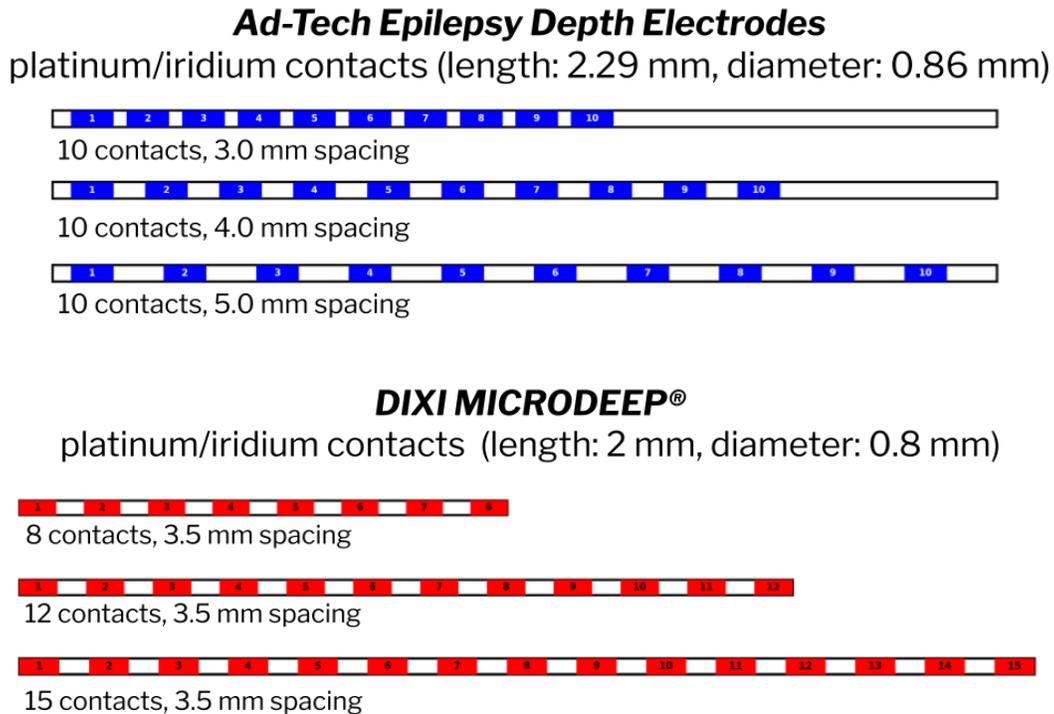


Figure 1-2. Visual representation of depth electrodes and different inter-contact spacing from two manufacturers, AdTech and DIXI medical.

Traditionally, electrode implantation for SEEG is performed in two steps with the assistance of a stereotactic frame. The first step is the *répérage* or stereotaxic localization through medical imaging, outlined by Talairach. Electrode trajectories are safely planned using different imaging modalities to visualize the vasculature and soft tissue of the brain. (see Sections 1.2 for more detail on specific imaging modalities). The second step is the implantation of electrodes, guided by the stereotactic frame. The stereotactic coordinates for each planned trajectory are set using the frame before a roughly 2 mm twist drill is used to perforate through the scalp, skull and dura mater. Once this hole is made, an anchor bolt is screwed in to avoid the overflow of cerebrospinal fluid. A needle is passed through the bolt to establish the path to the desired target before the electrode is placed along this trajectory (Cossu et al., 2005; Joswig, Benson, et al., 2018). At the end of the procedure, patients are transferred to a post-operative anesthetic care unit, where additional imaging is acquired to view electrodes and assess for potential

complications. Today, advancements in image-guided surgery and robotics have facilitated the transition to robot-assisted implantations, where the planning and surgical intervention can be performed in a single step. Neuro-navigation software allows for the seamless integration of multi-modal/planar imaging, and the robot aligns the drill with the vector of each trajectory (Cardinale et al., 2013, 2019).

1.1.2 The study of epilepsy with SEEG

The signal recorded by electrode contacts is the *local field potential* or LFP. The LFP is the electric potential generated by the flow of ions across the extracellular and intracellular space in the brain. These ionic currents are induced by the firing of neurons, with synaptic activity (i.e. communication between neurons) being the main contributor. ECoG records a smoothed version of the LFP at the cortical surface, as the signal is propagated through the soft tissue of the brain. The “true” LFP from local neuron populations is captured with multiple micro or macro electrodes placed directly within the tissue of interest, such as in SEEG (Buzsáki et al., 2012).

Intracranial recordings are separated into the ictal (during a seizure), interictal (between seizure), and postictal (after seizure) periods (Fisher et al., 2014). These segments are analyzed to identify the *irritative zone* and the *seizure onset zone*. The *irritative zone* is the origin of interictal spikes known as interictal epileptiform discharges (IEDs). Localization of the irritative zone is insufficient in predicting the epileptogenic zone or lesion (Paredes-Aragon et al., 2022). The gold standard in identifying the epileptogenic zone is the *seizure onset zone*, defined by the initiation of low-voltage fast activity during the ictal period (Jobst et al., 2020). Early research by Bancaud and Talairach revealed that seizures originated from several different structures in the brain rather than a single area or zone (J. Talairach & Bancaud, 1966). This led to the concept of epileptogenic networks, with secondary discharges deemed to be part of the *propagation zone*. Seizure propagation is a complex process dependent on the structural connectivity of cortical and subcortical structures. Defining an epileptogenic network involves the study of individual units (i.e. neurons or brain regions) to study the activation and connectivity between them (Bartolomei et al., 2017). The implantation of SEEG electrodes in carefully selected regions facilitates this investigation and can tailor the resection plan to disrupt these networks (Kreinter et al., 2024). This makes the precise anatomical localization of individual contacts essential to accurately mapping these networks and understanding seizure dynamics.

1.2 Medical imaging for stereoelectroencephalography

SEEG analysis bridges the study of electrophysiology to neuroanatomy. This requires advanced imaging technologies to visualize the implanted deep-brain structures and the locations of electrodes, in the form of Computed Tomography (CT) and Magnetic Resonance Imaging (MRI). CT uses a rotating X-ray tube to provide a 3-dimensional view of the brain and contrast high-density tissues. A CT is typically acquired post-operatively to localize electrodes and their contacts but lacks the anatomical contrast to distinguish the soft tissue of the brain. MRI uses the magnetic properties of atoms to create an image with superior tissue resolution. It is acquired pre-operatively to define the anatomical location of contacts. The following subsections will cover the principles behind each modality and their relevance to SEEG.

1.2.1 Computed Tomography (CT)

Computed Tomography (CT) was first described by Sir Godfrey Hounsfield in 1973, for which he received the Nobel Prize in 1979 (Hounsfield, 1973). A whole-brain CT scan is a three-dimensional volume, divided into a matrix of volumetric pixels or “*voxels*.” CT combines multiple X-rays at different angles to produce axial (horizontal) slices, in contrast to the anterior and lateral views from X-rays. The data from each X-ray projection is fed into a computerized algorithm that reconstructs the individual two-dimensional slices (see Natterer, 1986, for a detailed explanation of computerized tomography). The nature of X-ray projections avoids geometric distortions that could be introduced during image reconstruction, allowing CT to provide the most accurate stereotactic coordinates of structures and imaged electrodes (Gorgulho et al., 2009).

The intensity value at each voxel reflects the attenuation or absorption of the x-ray beam by the structural content. It is quantitatively scored in Hounsfield units (HU), scaled according to the radiodensity of distilled water at standard pressure and temperature (STP), defined as 0 HU, and the radiodensity of air at STP as -1000 HU. The formula to calculate the Hounsfield unit of a single voxel is as follows:

$$HU = 1000 \times \frac{\mu - \mu_{water}}{\mu_{water} - \mu_{air}}.$$

The coefficient μ describes the average attenuation at the voxel, with μ_{water} describing the average attenuation of water, and μ_{air} the average attenuation of air. Soft tissue generally ranges between 100 and 300 Hu, cortical bone between 500 and 1900 Hu, and metallic implants exceeding 2500 Hu.

The difference in attenuation for the soft tissue of the brain (i.e. the gray and white matter) is 0.5% in CT, preventing the anatomical labeling of contacts (Peters, 2001). However, CT is essential in localizing implanted electrodes and their contacts (Joswig, Benson, et al., 2018; Mercier et al., 2022). As seen in Figure 1-3, the attenuation for SEEG contacts and other metallic objects starkly contrasts against the rest of the brain. Individual electrode contacts are bright hyper-intense regions, with the centroid representing the most accurate localization. Contacts are also surrounded by low-intensity voxels in the negative Hu range, forming a “shadow” artifact. Contact localization is typically performed using this difference in voxel intensities, either by manually inspecting the image or through semi-automated algorithms (more in Section 1.5).

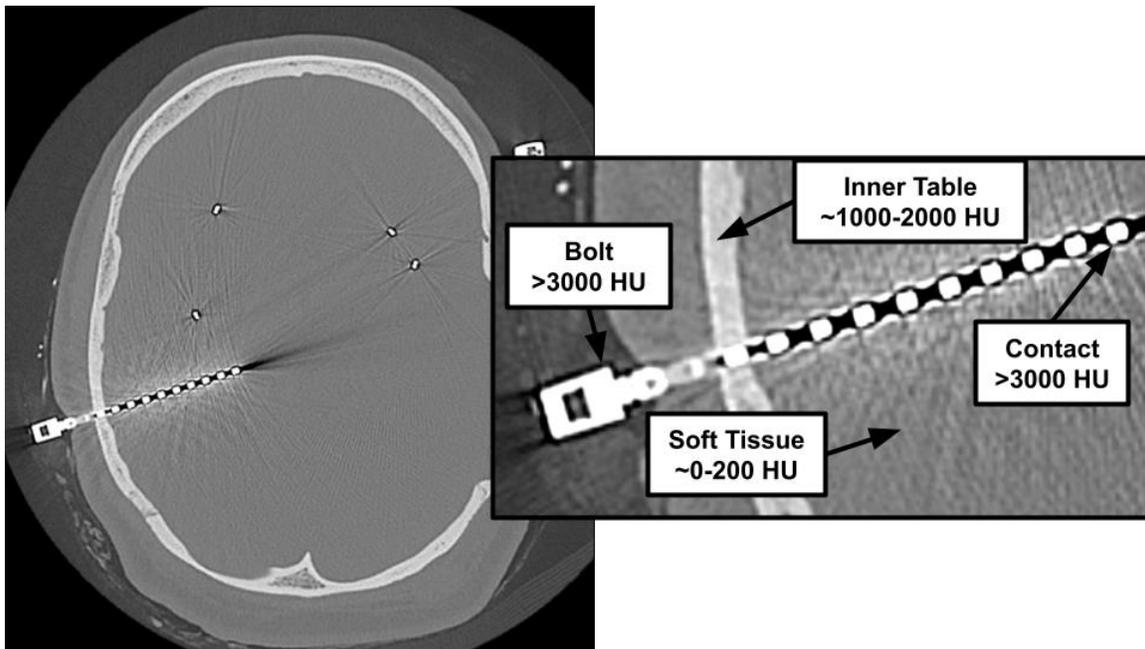


Figure 1-3. Post-operative CT image. *Seen on the left is the axial (top-down) view, with the right highlighting the different components and corresponding HU. Captured using 3D Slicer version 5.6.1, a medical image viewer (Fedorov et al., 2012).*

1.2.2 Magnetic Resonance Imaging (MRI)

The human brain primarily consists of fat and water, which are partly made up of hydrogen atoms. These atoms have a quantum *spin* property or angular momentum, analogous to tiny spinning magnets (Aston et al., 2017). In the presence of a strong magnetic field (B_0), the spin of atoms will exist in a superposition between the lower energy state, parallel to B_0 and the higher energy state anti-parallel to B_0 . At a given moment, more atoms will align to be parallel, since the low-energy state is favored. This creates a net magnetic moment along the direction of B_0 . Atomic nuclei in this state will start to precess; that is, the axis of the atom will spin around the axis of B_0 (Aston et al., 2017). This precession occurs in a random phase at the Larmor frequency (ω) defined as follows:

$$\omega = \gamma * B_0,$$

where γ is the gyromagnetic ratio, a property intrinsic to the atom, and B_0 is the magnetic field. While aligned to B_0 , applying an oscillating radio-frequency pulse at the Larmor frequency will excite these spins to a higher energy state and tip their net magnetic moment perpendicularly, where the atoms also precess in phase. After the radio-frequency pulse the system will return to equilibrium. As they return to equilibrium, nuclei emit energy “relaxing” along the longitudinal and transverse axes (Aston et al., 2017). The relaxation times vary based on the tissue composition and are responsible for MRI's superior anatomical contrast. T_1 -weighted images are the most used, reflecting the time for spins to relax to the low-energy state along the longitudinal axis. The T_1 -relaxation time for white matter is shorter in comparison to gray matter, resulting in a brighter signal. Despite the superior resolution compared to CT, the nonlinearity of the gradient field and the resonance effects from the local magnetic field of the tissue can cause geometric distortions or artifacts in the acquired MRI (Chen et al., 2008).

During SEEG surgery, MRI is acquired pre-operatively for surgical planning (unless there is some contraindication) to provide anatomical context to the implanted electrodes visualized in post-operative imaging. This is done through *image registration* (see Section 1.3.3), allowing the post-operative scan to be in the same space as the preoperative MRI. While the visualization of SEEG contacts is most commonly done with CT, it can also be performed with MRI (Hall & Khoo, 2018). In a postoperative MRI, contacts cause signal voids due to magnetic susceptibility

artifacts, appearing as “black holes” (Yang et al., 2012). Theoretically, a postoperative MRI would provide a better registration to the pre-operative imaging, as the anatomical structures are better visualized and can be easily matched. However, the electrode contacts cause significant artifacts and image distortions, limiting the anatomical resolution of nearby structures for registration. This also limits the spatial accuracy of contact localization. Post-operative CT has been shown to have a similar registration quality to post-operative MRI and provides the added benefit of an undistorted view of contact positions (Darcey & Roberts, 2010; van Rooijen et al., 2013). The choice in post-operative imaging modality ultimately varies from center to center.

1.3 Image processing

There are several aspects to imaging and image processing that are relevant for SEEG contact localization and analysis. The following section will first review the image coordinate system, how medical imaging data is represented, and define image resolution. It will then overview common pre-processing steps for analysis in image registration and image interpolation. Finally, the application of brain atlases/templates and image segmentation in relation to SEEG contact localization will be covered.

1.3.1 Coordinate systems

In geometry, the physical world can be represented as a three-dimensional *Euclidean space* where any location is described by three *Cartesian coordinates*. Three axes (x, y, z) make up the Euclidean space, intersecting a shared point or *origin*. An individual coordinate describes the distance of said point from the origin along a particular axis. The distance between two points in Euclidean space is defined as the *Euclidean distance*. In three dimensions, the Euclidean distance (ED) can be calculated as follows:

$$ED(\mathbf{a}, \mathbf{b}) = \sqrt{(\mathbf{a}_x - \mathbf{b}_x)^2 + (\mathbf{a}_y - \mathbf{b}_y)^2 + (\mathbf{a}_z - \mathbf{b}_z)^2}.$$

The concept of Cartesian coordinates is fundamental to imaging, as the raw data (i.e., the imaging voxels) and the scanner or “real-world” are two distinct Euclidean spaces or coordinate systems. As previously discussed, a digital radiological image (i.e. a CT or MRI scan) is a three-

dimensional volume composed of volumetric pixels or voxels. The image is indexed as an array of voxels, divided into three axes: i , j , and k . The i -axis extends to the right (i.e. which column), the j -axis goes to the bottom (which row), and the k -axis extends into the plane (which slice). The first voxel in the array is indexed at the coordinate (0, 0, 0). However, the coordinates or index do not provide any context of where voxels are positioned in relation to the image scanner or the imaged object. This requires a mapping from the image or voxel space to the scanner space (i.e. the axes that define the position of the imaged object relative to the scanner). In the scanner space, the x axis indicates the left-right direction, the y axis indicates the anterior-posterior direction, and the z axis the superior-inferior direction. The mapping is typically stored as a 4x4 *affine matrix* (more in Section 1.3.3), allowing us to perform a geometric transformation between two spaces. The “RAS” orientation indicates values increase as they go along the Right, Anterior, and Superior directions. Another convention is to use an “LPS” orientation, indicating the Left, Posterior, and Superior directions. This division also defines three planes anatomically:

1. The *axial* plane divides the brain from top-down or superior-inferior (x-y plane)
2. The *coronal* plane divides the brain from front-back or anterior-posterior (x-z plane)
3. The *sagittal* plane divides the brain from left-right (y-z plane).

Once in the scanner or patient space, the Cartesian coordinates reflect the distance from the scanner iso-center. This can now be used to determine the relative distances in millimeters of objects within an image, such as an SEEG contact or brain structures. Different image acquisitions or modalities are defined in their own coordinate space that must be transformed into one another using *image registration*. Templates or atlases (more in Section 1.3.5), represent another coordinate system. A classic example is the Talairach and Tournoux atlas of 1988, where the origin (0,0,0) was originally defined at the anterior commissure with the y-axis being the line connecting the anterior and posterior commissure (Brett et al., 2002; J. Talairach & Tournoux, 1988). Today, references to the Talairach space commonly refer to the origin as the midpoint of the line joining the anterior and posterior commissure (Horn et al., 2017).

1.3.2 Image resolution

The *spatial resolution* of the image is the ability to distinguish between different structures and is affected by several image acquisition parameters. This is distinct from image resolution, which is

defined by the size of the voxel. The voxel dimensions can be calculated from the field-of-view, the image matrix, and the slice thickness. The field-of-view (FOV) is the distance (in millimeters) over which the image is acquired, and the image matrix divides the FOV into a set number of pixels. The slice thickness describes the depth of the voxel in the z -dimension. Image voxels can be *isotropic* (equal-length dimensions) or *anisotropic* (at least one dimension shorter/longer than the others). For MR imaging, the matrix size is typically 256×256 with an isotropic voxel resolution of approximately one millimetre³ (Perera, 2022). Intraoperative CT for stereotactic surgery generally has a matrix size of 512×512 , with a FOV of 210 mm and a slice thickness of 0.8 mm (Cardinale et al., 2013; Furlanetti et al., 2021). This makes the initial voxel resolution anisotropic, at approximately $0.41 \times 0.41 \times 0.8$ mm. If scanner resolutions are limited or initially anisotropic, images can be resampled to produce the desired voxel dimensions (Perera, 2022). Note that the voxel dimensions of a 3D volume are not necessarily the same as the spatial resolution, since approaches such as zero-padding in K-space (effectively super-sampling the image) can be used to decrease the voxel size with no effect on spatial resolution.

1.3.3 Image registration

Following image acquisition, a critical processing step in surgical planning is *image registration* (also referred to as *image fusion*). *Image registration* is the spatial transformation of one coordinate system to another. It can be performed between imaging volumes or between a volume and a stereotactic device to guide the implantation (Kall, 2009). Registration identifies the geometric transformation that maps points in one image to corresponding points in a second image. In neuroimaging, registration is typically used to normalize individual anatomy and allow for a group-wise analysis of multiple subjects. Another application involves registering an individual image to a *brain atlas*, enabling the detailed labeling of structures. In stereotactic surgery, registration is performed for a single subject, merging images across (e.g., CT and MR) or within a modality (e.g., T1-weighted MRI and T2-weighted MRI). The stationary image is referred to as the *reference* or *fixed* volume; the second is the *moving* or *working* volume. The possible transformations that take images from one space to another can be divided into *rigid* or *non-rigid* transformations. This division is characterized by the *degrees of freedom* (DOF) applied to the image.

A *rigid* transformation is restricted to rotations or translations of the reference volume. These transformations are considered linear as they preserve the distances between points and the straightness of lines (Fitzpatrick et al., 2000). Additionally, a rigid registration preserves all non-zero angles between lines and assumes that the original images are *isotropic* in voxel size (i.e. not warped, skewed, or distorted) (Kall, 2009). When considering three-dimensional objects, a rigid registration has *six* DOF, representing three translations and three rotations. The transformation of an individual voxel of our moving image indexed by the array (x_1, x_2, x_3) to the fixed image space (y_1, y_2, y_3) is as follows:

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} = \begin{bmatrix} r_{11} & r_{12} & r_{13} \\ r_{21} & r_{22} & r_{23} \\ r_{31} & r_{32} & r_{33} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} + \begin{bmatrix} t_1 \\ t_2 \\ t_3 \end{bmatrix}$$

$$y = Rx + T$$

R and T represent the rotation and translation to be applied. For computational efficiency, we can express the transformations as a single 4x4 matrix and the image indices as a 4x1 unit vector (Hartov & Roberts, 2009):

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ 1 \end{bmatrix} = \begin{bmatrix} r_{11} & r_{12} & r_{13} & t_1 \\ r_{21} & r_{22} & r_{23} & t_2 \\ r_{31} & r_{32} & r_{33} & t_3 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ 1 \end{bmatrix}$$

The transformation can now be described as a matrix multiplication, $\mathbf{y} = \mathbf{M}(\mathbf{x})$. The benefit of describing our transformations as matrices is that several transformations can be combined through matrix multiplication. Additionally, we can obtain transformation in the opposite direction by simply by applying the matrix inverse (Friston, 2007). Rigid registrations are performed between scans acquired from the same patient where the gross anatomy is not expected to alter drastically. In stereotaxy, the pre-operative T1 typically acts as the reference volume. These matrices can be applied in a similar fashion to points within the image space, to transform them between coordinates systems.

Rigid registrations are a subset of the more general *affine* registrations. Affine registrations allow global scaling and shearing of the original image. In three dimensions, this provides 12 *DOF* (rotation, translation, shearing, and scaling along three axes). While this transformation can alter the angles between lines, it is still considered linear as it preserves parallel lines and the straightness between them.

The discussion above surrounded the transformation model, which are determined using an optimizer. The optimization for registration algorithms seeks to minimize a cost function which assesses the degree of similarity between the fixed and moving image. Cost functions can be broadly categorized according to their method of correspondence (points, surfaces, edges, voxel intensities/similarities). A more comprehensive review of registration algorithms is outside the scope of this thesis but can be found in the following references (Fitzpatrick et al., 2000; Hartov & Roberts, 2009). As part of the image processing workflow for SEEG employed at the London Health Sciences Centre (LHSC), the post-operative CT volume was rigidly registered to the pre-operative T1 image with 6 degrees of freedom, using a block-matching method to establish correspondence between the two images and a trimmed least squared optimization (Modat et al., 2010, 2014).

1.3.4 Image interpolation

During any image transformation (e.g., registration or resampling), we can imagine that the intensity values for each voxel are *pulled* from the original image to a new grid. Determining the values for the new voxel intensity requires some form of *image interpolation*. Image interpolation transforms a discrete matrix of numbers into a continuous image. There are several methods for interpolation in medical imaging analysis. The simplest method is a nearest neighbor interpolation, which assigns the new value based on the closest voxel. This preserves the original voxel intensities, but results in a blurred image. Other examples include bilinear interpolation using a weighted average of the nearest voxels and bicubic interpolation, which applies cubic polynomials. More complex methods apply a *sinc* function or linear combination of basis functions to produce a higher-quality image (Friston, 2007).

1.3.5 Brain atlases and templates

Brain atlases label anatomical structures or features within a *stereotactic space*. The labeling of structures in an atlas allows for the investigation of specific regions and the connectivity between them (Evans et al., 2012). The term *stereotactic space* is used interchangeably with *template* or *reference space* and refers to the description of imaging data using Cartesian coordinates (Lau, 2022). *Brain templates* can be derived using an average of multiple brain scans from a single subject or a population. Templates can also be defined using histological sections as seen in the original Talairach or Schaltenbrand atlases. Templates act as a standard coordinate system for the alignment or normalization of individual subjects. For SEEG, registering patient imaging with a template projects the contact coordinates to a normalized space. This allows for comparisons across multiple subjects. Contacts can also be classified according to the atlas label for a refined look at contact positions (Mercier et al., 2022).

1.3.6 Image segmentation

Image segmentation is a computer vision task where each voxel is identified as belonging to the *background* or a specified *target*. This is distinct from image classification where a label is applied to the entire image or object detection where a bounding box denotes the location of the target. Segmentation can be categorized as *semantic* or *instance* segmentation. Semantic segmentation classifies all instances of a target as one class in the image. Instance segmentation assigns a different label to each distinct instance of a class. In neuroimaging segmentation traditionally refers to tissue classification, where tools such as Freesurfer or other cortical thickness algorithms label areas belonging to the classes of gray matter, white matter, or cerebrospinal fluid (Fischl, 2012). Further distinctions can be made between healthy or pathological, different cortical and subcortical structures.

At a basic level, segmentations can be generated through thresholding, region growing, or edge-detection techniques. These approaches generate *discrete segmentations* where each label has a unique integer value. However they can perform sub-optimally in more complex scenarios, such as if there is noise present or if the image intensities are non-uniform (Withey & Koles, 2008). Advanced methods for automatic segmentation work to tackle these problems, based on image-

derived features. Examples include knowledge-driven methods, probabilistic-based methods, and deformable template-based methods (Khan et al., 2008). Knowledge-driven methods use prior anatomical knowledge to guide the segmentation. Probabilistic-based methods frame image segmentation as a classification problem, where each class has a unique statistical distribution of features constrained using a probabilistic atlas. The class for each voxel is determined by maximizing the *a posteriori* probability. Deformable template-based methods determine the affine registration between an atlas and an image using intensity-based metrics and concordant landmarks (Khan et al., 2008). Today deep learning approaches have vastly grown in popularity and shown to be quite successful in a variety of segmentation tasks (more in Section 1.4). These advanced methods generate *probabilistic segmentations*, where each voxel is assigned a probability of belonging to the label. Probabilistic segmentations are thresholded to create discrete segmentations.

The segmentation quality is assessed by determining the degree of overlap between the prediction and the ground truth. The two most common assessment methods are the Dice-Sørensen similarity coefficient (DSC) or Dice Score and the Jaccard Index (Dice, 1945; Jaccard, 1912). For a predicted segmentation A about the ground-truth segmentation B , the DSC can be calculated as follows:

$$DSC = \frac{2 |A \cap B|}{|A| + |B|}.$$

The numerator ($A \cap B$) represents the intersection or the degree of overlap between the predicted and ground-truth segmentation. Alternatively, the Jaccard Index or Jaccard similarity coefficient can be used:

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|} = \frac{|A \cap B|}{|A| + |B| - |A \cap B|}.$$

Here the denominator ($A \cup B$) indicates the union of the predicted segmentation (A) and the ground-truth (B). The Dice coefficient is related to the Jaccard index as follows:

$$DSC = \frac{2J}{1 + J}.$$

SEEG contact localization can be framed as a semantic image segmentation task, identifying which voxels belong to a contact. Low-level techniques such as thresholding or edge detection could separate contact artifacts. However, the presence of bolts and other artifacts outside the skull can lead to misclassification errors. A sophisticated deep learning approach could be better suited for the task of contact localization. An introduction to deep learning will be provided in the next section.

1.4 Deep learning for image segmentation

Machine learning describes the use of statistical models to predict outcomes or patterns in data. Most machine learning tasks are *supervised* learning tasks, where a model is trained on raw input data paired with labeled targets. The predictions are compared to the target using an *objective* or *loss function*. A model learns by adjusting a set of internal parameters known as weights to minimize the error of the loss function. After training over a set number of epochs or until reaching a minimum loss, the model is assessed on an unseen test dataset to assess the final performance.

For data such as medical images, variations to the input that do not account for changes to the target class (i.e. shift position of the head in the scanner, individual anatomical differences, etc.) can significantly change the values at the voxel level. This leads to poor mapping between the input and target, causing inaccurate predictions by most machine learning algorithms. Their performance can be improved with hand-crafted feature extractors that create representations of the input insensitive to these variations. However, the engineering of feature extractors requires a significant amount of time, domain expertise, and is biased towards the training data (LeCun et al., 2015). Automating the design of these extractors with machine learning can save time and improve their generalizability to unseen datasets (Goodfellow et al., 2016; LeCun et al., 2015).

Deep learning trains multi-layered networks (“deep” referring to the number of layers) to automatically generate feature extractors. These networks are known as Artificial Neural Networks (ANN), in reference to the structure of biological neurons. The data is vectorized at an initial set of input nodes or “neurons” before propagating through a series of interconnected hidden layers. Each layer has a set of weights that are adjusted (learned) during model training,

enabling the model to create increasingly complex and higher-level representations of the input (LeCun et al., 2015).

While generally effective, ANNs are not suited to handle the complexity of larger images or high-dimensional data types such as medical imaging. For instance, a patch from a medical image that is 64x64x64 voxels would be summarized as a vector of 262, 144 features. This can be computationally impractical as a single neuron in the network would need to train 262, 144 weights at the first layer. Convolutional neural networks (CNN) are a deep learning architecture designed to address this problem and work with data organized in arrays. The following section will provide an overview of CNNs.

1.4.1 Convolutional neural networks (CNNs)

CNNs are one of the more easily trainable and widely used models in deep learning (LeCun et al., 2015). Inspired by the visual cortex, CNNs are designed with neurons that connect only to small regions of the input, defining their “receptive field.” Neurons are also organized to match the depth of the input data, allowing for the processing of multi-dimensional arrays. For example, in volumetric images (such as in medical imaging), neurons are organized in three dimensions corresponding to the data's height, width, and depth. This local connectivity allows neurons to learn a small set of weights within their receptive fields that are shared across the depth of the image.

CNNs consist of two primary layers: *convolutional* and *pooling* layers. In convolutional layers, a set of kernels or filters are applied to the input using a convolution. A kernel acts as a feature detector to identify specific patterns in the input. Mathematically, a kernel is an array smaller in height and width than the original input but shares the same depth. *Convolution*s are the element-wise multiplication and summation as the kernel slides across the input. This operation transforms the original image into a new representation known as a *feature map*. The values of the kernel are individual weights learned during the training phase. Therefore the kernel learns to detect features at different locations of the image that minimize the loss function (O’Shea & Nash, 2015).

A CNN will consist of multiple convolutional layers that each have multiple kernels. Kernels in the earlier layers effectively have a smaller receptive field, and thus become tuned to lower-level features, such as edges or boundaries. In contrast, later layers identify more abstract or global characteristics, since they effectively have a much larger receptive field, and thus become tuned to entire shapes or specific objects. Convolutional layers reduce the parameters for the model to learn in two ways. First, the kernel size or “receptive field” restricts the connectivity of a neuron to focus on local patterns. Secondly, weight sharing (the reuse of the same kernel across the entire input) introduces *translational invariance*, allowing the model to detect features regardless of their position within the input (Chollet, 2017; LeCun et al., 2015).

Pooling layers downsample the feature maps generated by convolutional layers. Pooling is typically done after a convolution operation, using a smaller kernel to summarize the most essential features. Rather than performing a matrix multiplication, a pooling kernel can take the feature map's maximum, minimum, or average values. Typically, max-pooling or average pooling is employed in most CNNs (L. Zhao & Zhang, 2024). Pooling further reduces the size of the feature map and the number of parameters for the model to learn, preventing overfitting of the network to the training data. Pooling also provides some translational invariance, allowing features to be detected even if they are slightly shifted or translated in the input.

1.4.2 Loss functions

CNNs are highly effective for computer vision tasks such as image classification and segmentation, outperforming traditional methods (Krizhevsky et al., 2017; LeCun et al., 2015; Szegedy et al., 2015). The selection of a loss function is vital to the success of a CNN, acting as the mathematical representation of the desired task. In medical image segmentation, loss functions must address the particular issue of class imbalance since the desired segmentations are smaller in volume than the entire image (Ma et al., 2021).

Loss functions for segmentation tasks can be categorized as distribution-based, region-based, boundary-based, or combination loss functions (Jadon, 2020; Ma et al., 2021). Distribution-based loss functions look to minimize the dissimilarity between two distributions of classes. Cross entropy is fundamental to these methods and is used in many machine learning and optimization algorithms. Cross entropy is derived from statistics and for this thesis, can be described as a

measure of how one probability distribution differs from a reference. In binary cross entropy (CE), the segmentation is binarized to indicate the presence of a class label. In image segmentation, the image I can be represented as a domain of real numbers in 2 or 3 dimensions. N is the number of voxels in the image, while C is the number of classes. If we define A to be our predicted segmentation and B to be the target or ground-truth segmentation, the equation for binary cross entropy in a CNN-based segmentation task is as follows (Ma et al., 2021);

$$L_{CE} = -\frac{1}{N} \sum_{c=1}^C \sum_{i=1}^N B_i^c \log A_i^c$$

Here B_i^c and A_i^c indicate the class label c at the voxel i for the ground-truth and predicted segmentations. Weighted cross entropy (WCE) can also be used, where classes more frequent in the ground-truth segmentation are penalized in the loss. Here w_c denotes the weight for each class.

$$L_{WCE} = -\frac{1}{N} \sum_{c=1}^C \sum_{i=1}^N w_c B_i^c \log A_i^c$$

Region-based loss functions look to maximize the areas of overlap between the prediction segmentation and the ground-truth (or minimize the mismatch). The most well-known of these methods is the Dice loss function, which optimizes the Dice Coefficient. Unlike cross-entropy loss functions, Dice loss does not require class re-weighting. The equation for the Dice loss function is as follows (Ma et al., 2021);

$$L_{Dice} = 1 - \frac{2 \sum_{c=1}^C \sum_{i=1}^N B_i^c A_i^c}{\sum_{c=1}^C \sum_{i=1}^N B_i^c + \sum_{c=1}^C \sum_{i=1}^N A_i^c}$$

Boundary-based loss functions look to minimize the distance between the ground truth and predicted segmentations. It is worth noting that boundary-based loss functions determine the boundaries of segmentations in a region-based way. They are therefore closely related to other region-based loss functions such as Dice loss (Ma et al., 2021).

Finally, compound losses are functions that consider a weighted combination of the aforementioned methods. A common example is the Dice Cross Entropy (DiceCE) loss, also referred to as a combo loss, which is a sum of the cross entropy and Dice loss.

$$L_{DiceCE} = L_{CE} + L_{Dice}$$

The addition of cross entropy loss is thought to regulate the trade-off between false positive (voxels incorrectly labeled as a target) and false negative (voxels mislabeled in multi-class segmentations or incorrectly labeled as background) predictions (Taghanaki et al., 2021).

1.4.3 U-Net for medical image segmentation

The success of CNN models in computer vision can partially be attributed to the availability of training data. For example, the heavily-cited AlexNet paper in 2017 used 1 million images to train a CNN for image classification i.e. assigning a single class to the entire image (Krizhevsky et al., 2017). In medical imaging, obtaining a comparable number of curated labels for training is impractical. Furthermore, most CNNs such as AlexNet are used for image classification. Medical image segmentation requires the localization of classes, assigning a class label to each voxel (Ronneberger et al., 2015). The U-Net was developed to address these issues by combining the core elements of CNNs with other advances in deep learning. Today it is one of the more popular and best-performing architectures for medical image segmentation (Isensee et al., 2021; Ronneberger et al., 2015; Siddique et al., 2021).

The “U” refers to the encoder-decoder framework shown in Figure 1-4. The first half of the “U” is the encoder or contracting portion, where the model progressively generates more feature maps from the input data using a series of convolutional and max-pooling layers. In the original U-Net paper, each level of the encoder had two convolutional layers followed by a max-pooling operation. The number of convolutional filters and resulting feature maps double at each level, while pooling operations “downsample” these maps by halving their spatial size. This process is repeated for a specified depth, reaching the bottom portion of the “U” or the bottleneck layer. At this stage, repetitive convolutional operations are performed to compress the data and capture the most critical features.

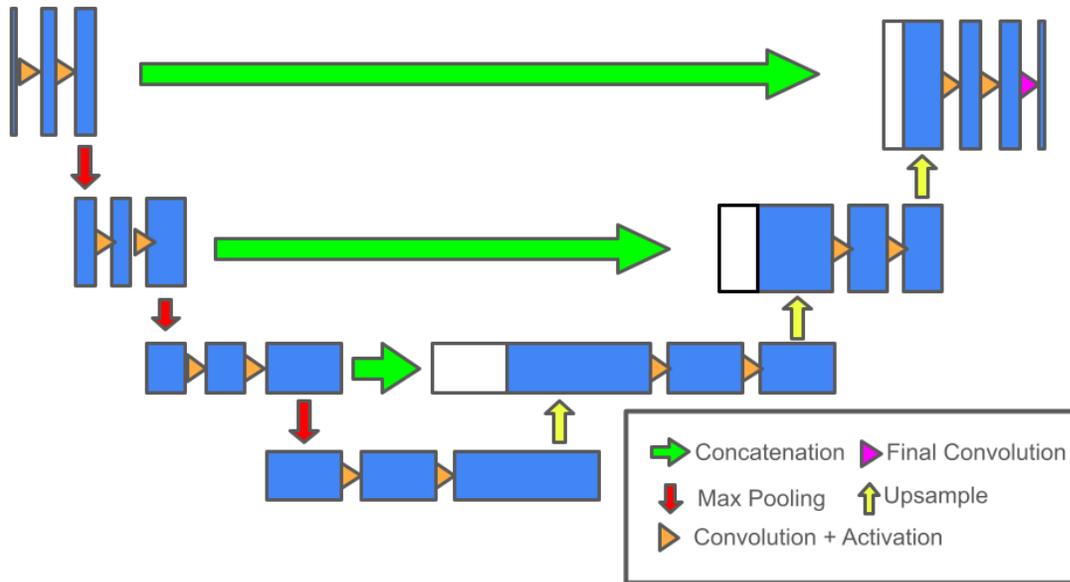


Figure 1-4. Standard 2-D U-Net Architecture, adapted from Ronneberger et al., 2015. *The blue boxes represent the feature maps of the image, with the width representing the number of feature maps and the height representing the spatial size.*

The data is now upsampled through the second half of the “U” known as the decoder or expanding path. The primary goal of the decoder is to increase the spatial size of the new feature maps to match the dimensions of the original input and determine the predicted segmentation. Upsampling is used to increase the spatial size to match the original input, by interpolating unknown voxel values. This is followed by two convolution operations to learn relevant information from the newly formed feature maps. Skip connections (in green) at each level concatenate feature maps (in white) from the encoder to the upsampled representations of the decoder. This step allows the model to “remember” high resolution feature maps of the encoder and prevent overfitting of the model. At the final layer (top-right of the “U”) a 1x1 convolution (1 indicating the size of the kernel) is applied to summarize the feature maps and create the prediction (Ronneberger et al., 2015).

The initial U-Net architecture was implemented for cell segmentation from microscopy images and used a pixel-weighted cross entropy loss. To solve the issue of limited training data, the authors applied “excessive” data augmentations in the form of elastic deformations (Ronneberger et al., 2015). The 2-D U-Net architecture was extended for volumetric imaging by

another group, with batch normalization and drop-out added as *regularization* techniques to speed up model training (Çiçek et al., 2016). Batch normalization normalizes the inputs to convolutional layers, adjusting the distribution of weights (Ioffe & Szegedy, 2015). Dropout turns off a random subset of neurons within the convolutional layers, forcing the model to learn new features and not over-rely on a particular subset (Srivastava et al., 2014).

The U-Net is a promising tool to be applied to the task of SEEG contact segmentation. As discussed in Section 1.2.1, metallic artifacts in CT are hyper-intense regions, starkly contrasting from the rest of the scan. Simple image processing techniques such as thresholding or edge detection would separate the individual contacts in the image, but the presence of bolts and wires in a similar Hu range limits the effectiveness of this method. Masking the CT to only include the intracranial volume could address this issue; however, there are instances where contacts are localized at the level of the skull. Additionally, the implantation of electrodes along the superior-inferior direction can blur the radiological artifacts of contacts, such that they appear as a single “fused” electrode (more details on the challenges for contact localization will be covered in Section 1.5.1). For these reasons, a more advanced computer vision approach such as the U-Net could encapsulate these nuances and provide an accurate segmentation of contact positions. The next sections will cover the current approaches for contact localization, the current applications of deep learning in contact segmentation, and present the motivations and objectives of the thesis.

1.5 Current approaches for SEEG contact localization

As previously discussed, the precise anatomical localization of each electrode and its contacts is crucial to interpreting SEEG data. Given their millimetric size, localization errors must exist within a small margin. Inaccurate placements can misidentify the type of tissue (gray matter vs. white matter) sampled or, on a larger scale, the anatomical subregion. Contact localization was traditionally performed by visually reviewing and annotating post-operative imaging (either CT or MRI) for each contact. CT is used in most centers, as localizations represent the most accurate stereotactic coordinates (relatively free of geometric distortions), and contacts are clearly visualized as bright regions of hyper-intensity (Darcey & Roberts, 2010).

Today, many computational toolboxes aim to automate the process for post-implantation CT data. Each follows a similar pattern. First, the CT is thresholded to identify the high-intensity voxels within the image. These voxels are associated with SEEG contacts, bolts, wires, and other metal artifacts. From here, the voxels belonging to contacts are identified by excluding any objects found outside the brain. This is typically done by creating a mask of the intracranial space (Aslam et al., 2024; Blenkmann et al., 2024; Cai et al., 2022; Granados et al., 2018; Janca et al., 2023; Medina Villalon et al., 2018; Qin et al., 2017; Wong et al., 2024; B. Zhao et al., 2023). Alternatively, artifacts can be categorized by geometric features (ex., how spherical or circular they are) or by the distances between objects (artifacts belonging to contacts would presumably be closer together than those outside the brain) (Centracchio et al., 2021; Ervin et al., 2021). Once filtered, the contact coordinates are obtained by determining the center of mass for each artifact (Arnulfo et al., 2015; Davis et al., 2021; Ervin et al., 2021; Granados et al., 2018; Monney et al., 2024; Qin et al., 2017; Wang et al., 2023; Wong et al., 2024). The extracted contacts must then be grouped by electrode. Semi-automated approaches use two points along the electrode trajectory or iteratively select the nearest contact from a manually defined point at the electrode tip (Davis et al., 2021; Ervin et al., 2021; Groppe et al., 2017; Janca et al., 2023; Monney et al., 2024; Narizzano et al., 2017; Wang et al., 2023). Other toolboxes propose clustering of identified contacts to achieve full automation (Cai et al., 2022; B. Zhao et al., 2023). Table 1.1 summarizes the currently published literature for SEEG contact localization, the programming language, and the reported accuracy either as the displacement from the manually localized coordinate or the number of correctly localized contacts.

Table 1-1 Toolboxes for SEEG contact localization using CT. *Proprietary software or tools without open-access code were not hyperlinked. N/A indicates there was no reported accuracy. *N/A for iElvis from Groppe et al., 2017, indicates a reported error for ECoG contacts, but no measure for SEEG. *N/A for BrainQuake indicates error was assessed based on the degree contact positions deviated from the planned trajectory and how closely the predicted inter-contact distance matched the known spacing of the electrode. *N/A for CranialVault indicates that the tool is proprietary, with the details regarding test size, programming language, and reported accuracy not openly available.*

Authors	Tool Availability (Hyperlinked)	Language	n for Testing	Reported Accuracy
Aslam et al., 2024	DELRecon	MATLAB	8 subjects, 1196 contacts	98% (1173/1196)
Blenkmann et al., 2024	iElectrodes	MATLAB	20 subjects, 779 contacts	0.46 ± 0.18 mm
Monney et al., 2024	voxelLoc	MATLAB	5 subjects, 812 contacts	<1 mm inter/intra-user reliability
Wong et al., 2024	SEEGLoc	Python	20 subjects, 196 electrodes	0.25 ± 0.51 mm
Lucas et al., 2024	voxTool	Both Python and MATLAB	N/A	N/A
Janca et al., 2023	SEEG contact detection and skull measurement	MATLAB	54 subjects, 8745 contacts	0.43 ± 0.84 mm
Wang et al., 2023	Your Advanced Electrode Localizer (YAEL)	R/Python	N/A	N/A
Zhao et al., 2023	DELLO	MATLAB	7 subjects, 80 electrodes, 1030 contacts	N/A
Cai et al., 2022	BrainQuake	Python	8 subjects, 743 contacts	*N/A
Rockhill et al., 2022	Intracranial Electrode Location and Analysis in MNE- Python	Python	N/A	N/A

Davis et al., 2021	LeGUI	MATLAB	38 subjects, 3726 contacts	0.23 mm (0.00–1.37)
Ervin et al., 2021	FASCILE	Python	35 subjects, 4661 contacts	0.73 ± 0.15 mm
Li et al., 2019	iEEGview	MATLAB	22 subjects, 3158 contacts	N/A
Lin et al., 2019	ContactSegmentation_SYSU	MATLAB	12 subjects, 135 electrodes, 1812 contacts	1.055 ± 0.449 mm
Medina Villalon et al., 2018	GARDEL	MATLAB	30 subjects, 4590 contacts	0.59 mm
Granados et al., 2018	N/A	C++ (ITK/MITK)	224 electrodes, 1843 contacts	0.37 mm (n = 109 contacts)
Groppe et al., 2017	iElvis	MATLAB	*N/A	*N/A
Narizzano et al., 2017	SEEG Assistant (SEEGA)	Python	40 subjects, 9626 contacts	0.50 ± 0.06 mm
D’Haese et al., 2012	CranialVault and CRAVE (CRAnialVault Explorer)	*N/A	*N/A	*N/A

1.5.1 Existing challenges for contact localization

SEEG Assistant (SEEGA) is a semi-automated toolbox validated in over 500 electrodes, with a reported error of 0.5 mm (Narizzano et al., 2017). SEEGA is currently implemented as part of the image reconstruction workflow for intracranial electrode implantation at LHSC and can be viewed as the current state-of-the-art tool for contact localization. However, SEEGA requires the manual tuning of several parameters to achieve optimal results.

1. Users must move the points defined by the electrode trajectory, to be at the start and end of the radiological electrode artifact.
2. Users manually set the inter-contact distance of each electrode
3. Users must set the number of contacts present in each electrode.

Additionally, SEEGA assumes contacts to not deviate from the electrode axis (defined by the entry and target points of the electrode trajectory) by more than 10 degrees. In actuality, the degree of bending for electrodes is dependent on the properties of the soft tissue and interactions with the physical electrode (Granados et al., 2021). This can significantly alter the expected contact positions from the trajectory and must be manually corrected to account for this deviation.

The manual interventions involved with SEEGA are common to most contact localization algorithms and performed to address the inherent challenges of automating the task. The first is that localized contacts must be grouped into their electrodes and assigned a unique label. The most intuitive method is to use the two points from the planning trajectory; that is the point of entry at the skull and the target of the brain. Contacts are iteratively labeled from the deepest contact to the point closest to the skull. This naming scheme for SEEG is in large part due to the bipolar reference scheme used to analyze recordings, where the signal represents the activity of two adjacent contacts (Mercier et al., 2022). Alternatively, some papers attempt to automate localization without the placement of two points along the trajectory. These approaches generally rely on a clustering algorithm initialized by the number of electrodes implanted (Aslam et al., 2024; B. Zhao et al., 2023). However clustering approaches are highly dependent on the initial selection of centroids and can fail if done improperly. To combat this, other groups have proposed applying a Hough transform to detect line-shaped trajectories or to initialize clustering using the planned electrode trajectories (Cai et al., 2022; Janca et al., 2023).

The key challenge for all algorithms is the presence of electrodes implanted vertically (i.e. going along the superior-inferior direction) with a smaller inter-contact spacing. Algorithms relying on traditional image processing techniques will fail since at oblique angles (perpendicular to the in-plane slice), the individual contacts are indistinguishable in the image. Generally, these tools rely on the manual correction of points from the post-implantation trajectory to be at the start and end

of the electrode artefact. The remaining contact positions are then interpolated from these two points, using the inter-contact spacing and the number of contacts in an electrode (Monney et al., 2024; Narizzano et al., 2017; Qin et al., 2017; Wong et al., 2024). Other tools such as EpiTools propose an additional erosion of these electrodes to segment contacts. However this can be an imprecise method of determining positions and likely do not accurately represent the true position of contacts (Medina Villalon et al., 2018). Figure 1-5 provides a visual example of some of these difficult electrode configurations.

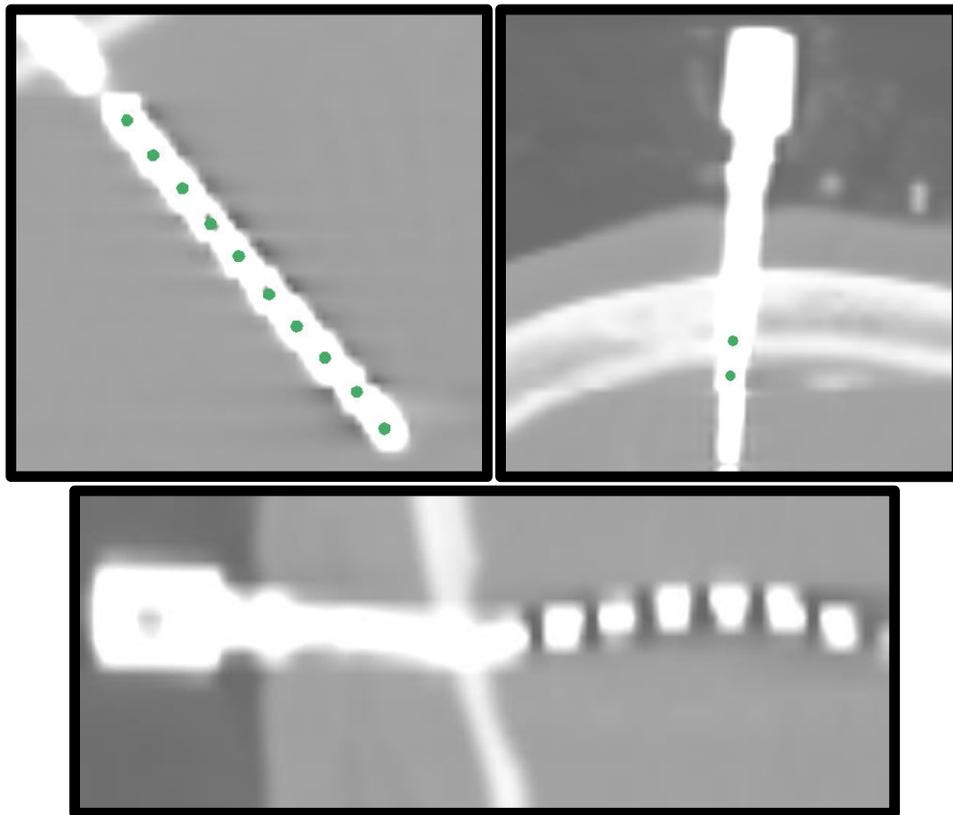


Figure 1-5 Example of challenging contact localizations, with the manual “ground-truth” placements in green. *The first panel from the left shows a vertically implanted electrode, where the contact artifacts appear to be “fused” together. The second panel shows an example of both a vertically implanted electrode and an example of where contacts are at the level of the skull. The bottom panel demonstrates an example of where electrodes curve from the planned trajectory and must be manually adjusted for most contact localization algorithms.*

1.5.2 Deep learning for SEEG contact segmentation

Two papers have trained U-Nets for SEEG contact segmentation using post-operative CT. The first paper aimed to segment the entire electrode trajectories and the intracranial space. The architecture was a 3D Cascade U-Net and Dice Cross-Entropy as the loss function (Vlasov et al., 2021). The cascade U-Net hierarchically uses two connected U-Nets; the first network aims to predict coarse or more global features and is fed into the second network to refine predictions. In their design, the first model was trained on a low-resolution whole-brain CT volume, while the second U-Net was trained on 160x160x80 voxel-size patches. To increase the training data and generalize the model to other stereotactic surgeries, the authors included postoperative scans from deep brain stimulation (DBS) surgeries. DBS is a stereotactic procedure where 2 electrodes (4-8 contacts) are permanently implanted to deliver electrical stimulus. In total, 35 DBS CT and 45 CT volumes were selected as part of their training set, randomly split into five cross-validation sets. The network was also trained on an external test set of 10 images (5 DBS, 5 SEEG) with a Dice Coefficient of 0.897 ± 0.043 .

The second paper combined outputs from a 2-D and 3-D U-Net to segment contacts (Pantovic et al., 2022). Their model was trained on scans from 18 patients and 18 augmented volumes. The augmentations included random rotation and flipping of the original CT volume. The authors employed leave-one-out cross-validation, where 1 original image and the corresponding augmentation were left out for testing. Contacts were clustered into electrodes using a Gaussian Mixture Model. The model was validated using a connected components count to identify the number of segmented instances and determine the number of true positives (correctly identified contacts), false positives (artifacts incorrectly labeled as a contact), and false negatives (missed contacts). They report an average of 183.0 true positives, 1.1 false positives, and 2.0 false negatives across all cross-validation folds, with an average Dice coefficient of 0.867. The authors mention that vertically implanted electrodes or electrodes with an uninterrupted shape could not be segmented into individual contacts and would require additional post-processing steps. It is unclear whether these contacts were included in the results.

Despite the examples of U-Nets in the literature, there are no openly available models to determine their generalizability to data from other centers or compare performance. Additionally,

they are limited in their train and test size and fail to address the challenge of localizing contacts for vertically implanted electrodes.

1.6 Open science and reproducibility for SEEG contact segmentation

The recruitment of patients implanted with intracranial electrodes (both SEEG and ECoG) for research studies offers a unique and privileged opportunity combining medical imaging with electrophysiological recordings precisely localized in different brain regions. The organization of these large datasets is essential in ensuring accessibility and reproducibility of these studies. The Brain Imaging Data Structure (BIDS) was first developed to standardize neuroimaging data (Gorgolewski et al., 2016). This was later extended to encompass intracranial EEG data, with the anatomical contact localizations being provided in an “_electrodes.tsv” file, the corresponding coordinate system in a “_coordsystem.json” and the associated imaging (Holdgraf et al., 2019). To date, only four tools have incorporated BIDS specification as part of their workflow in iElvis, VoxelLoc, iEEG-recon/voxTool, and the manual tool provided through MNE-python (Groppe et al., 2017; Lucas et al., 2024; Monney et al., 2024; Rockhill et al., 2022). This limits the reusability of potentially effective tools created by other centers and presents an additional hurdle for most contact localization algorithms.

1.7 Thesis overview

The challenges with semi-automated approaches for SEEG contact localization and the millimetric sampling radius of contacts illustrate the need for an automated tool capable of providing precise localizations. Current tools for contact localization all require a degree of manual user intervention, from the placement of two contacts per electrode, knowledge of the number of contacts in an electrode, and the inter-contact spacing. In addition, vertically implanted electrodes are unsolvable without manually inspecting the post-operative imaging and having prior information regarding the inter-contact spacing and number of contacts. For a typical SEEG implant where 12 electrodes and 150 contacts are implanted, these manual steps can add an additional 1-3 hours of labor to the clinical workflow. Deep learning models for image segmentation such as the 3-D U-Net present a potential solution for automatic localization.

The objectives of the thesis are as follows:

1. Work towards the automation of intracranial depth or SEEG electrode and contact localization using a 3-Dimensional U-Net. Specifically, the proposed approach would eliminate the need to define the inter-contact spacing, the number of contacts per electrode, and the manual placement/correction of points.
2. Compare the millimetric accuracy of the automatic contact localizations from the 3-D U-Net, the semi-automatic localizations of SEEG Assistant (the current state-of-the art tool), and a traditional image processing approach to the manually localized contact coordinates.

The thesis is the first to quantify the millimetric accuracy of deep learning approaches and compare them to traditional image processing methods for contact localization. The proposed algorithm also reduces the number of manual steps required for contact localization, reducing the cognitive load for clinicians and facilitating the group-wide analysis of iEEG datasets.

Chapter 2

2 U-Net Implementation Methods

The following chapter will cover the dataset, image pre-processing workflow, the U-Net model architecture, and how the model will be validated.

2.1 Patient Data

2.1.1 CT imaging

Postoperative CT volumes were acquired from 191 patients who underwent SEEG surgery at the London Health Sciences Centre (LHSC). The CT scans were acquired with the following parameters: tube voltage, 120 kV; tube current, 145 mA; data acquisition diameter, 1,331 mm; reconstruction diameter, 320 mm; matrix size, 512×512 voxels; pixel spacing, 0.625 × 0.625 mm²; axial slices, 96; slice thickness, 0.625 mm; gantry tilt, 0°. Many patients were implanted with 10-contact Ad-Tech Medical Instrument Corporation electrodes with 3-, 4-, 5-, or 6-mm spacing (0.86-mm diameter, 2.29-mm length). 14 patients were implanted with other SEEG electrodes. This included 3 patients with Dixi Medical electrodes (Besançon, France), 8 patients with 8 macro contacts spaced 5mm apart and 8 micro wires, and 3 patients with Behnke-Fried electrodes made up of 9 macro contacts. These 14 patients and 5 patients who had re-implantations (i.e. 2 CT scans), were held-out from the data split. From here, the remaining data from 172 patients was randomly split using *scikit-learn* to include 70% for training (120 patients), ~10% for validation (17 patients), and ~20% as a test dataset (35 patients) (Pedregosa et al., 2011).

2.1.2 Ground truth semi-automatic contact localization

For SEEG electrode and contact localization an in-house image processing pipeline which involves electrode contact detection, brain tissue segmentation and atlas fitting (<https://github.com/greydongilmore/ieegProc>). A workflow of this pipeline can be found in Figure 2-1. Semi-automatic contact localization was performed in 3D Slicer using the SEEG Assistant (SEEGA) module (Narizzano et al., 2017). The entry and target points of each

electrode were manually defined on the post-operative CT image. The CT was rigidly registered to the pre-operative T1w-MRI using NiftyReg (Modat et al., 2010, 2014). These labels were then provided to the SEEGA algorithm, for semiautomatic segmentation of the electrode contacts. Following the procedure, each contact position was visually inspected and manually corrected to be at the center of a contact by a trained clinical electrophysiologist, neurosurgeon, or technician (1-3 hours/subject). These coordinates were then retrospectively reviewed by two raters forming the *ground-truth* localizations. The localizations were later used for training the deep-learning model and to assess the final performance.



Figure 2-1 Flowchart diagram for SEEG imaging pre-processing workflow.

2.2 Pre-processing workflow for U-Net

The 3-D U-Net model was trained on patches from a resampled CT and a target image of the individual contacts. Using a bicubic interpolation, the original CT was resampled to be 0.4 mm isotropic. The target image was a mask generated from the localized contact coordinates in the space of the resampled CT. Individual contacts were represented as a 1 mm sphere dilated from the coordinate position using the `convert3d` command-line tool (Yushkevich et al., 2006). 3-D isotropic patches of 64 and 96 voxels were then taken for the resampled CT and the target contact mask. These sizes were chosen as larger patch sizes are known to improve the performance of convolutional neural networks in medical image segmentation (Hamwood et al., 2018). At 0.4 mm isotropic resolution, this corresponded to 25.6 mm and 38.6 mm isotropic patch. Considering the length of an SEEG contact at ~2 mm and a diameter of ~0.8 mm, the selected patch sizes encompass multiple contacts from different trajectories. In comparison, smaller patch sizes of 16 or 32 voxels would only include a single contact. It can be hypothesized that a model trained on smaller patches might be unable to distinguish contacts from some of the other high-intensity artifacts found in the image (ex. bolts, wire artefacts), resulting in a larger number of over-predictions. Examples of patches taken at 32, 64, and 96 voxels can be seen in Figure 2-2.

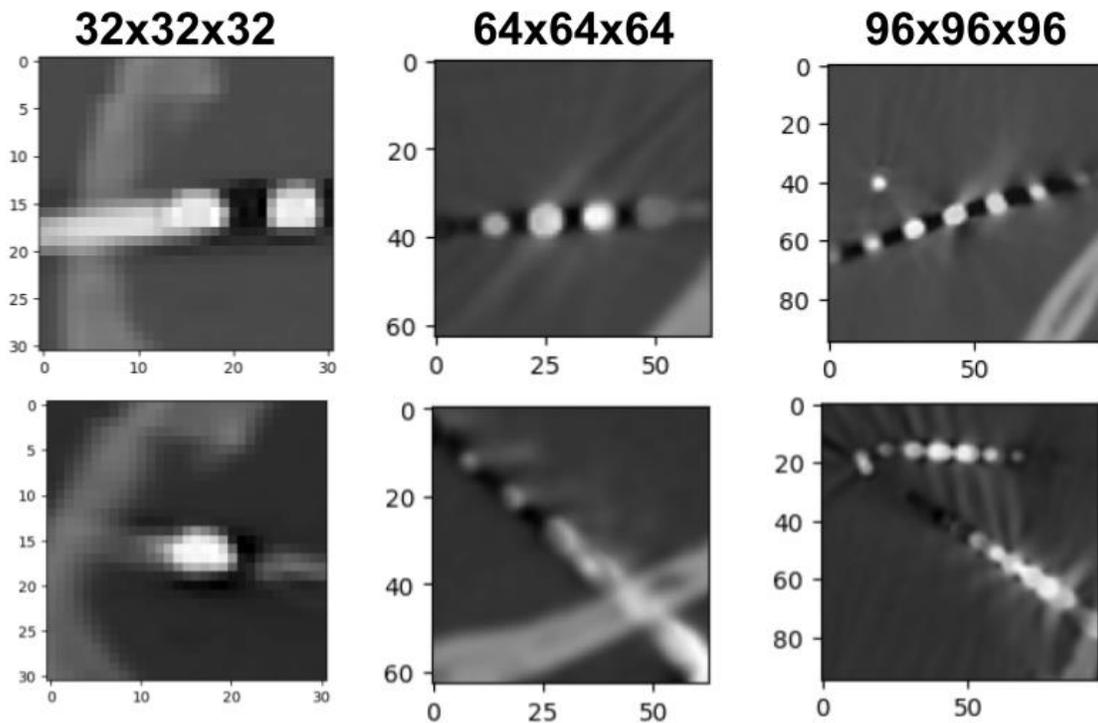


Figure 2-2. 3-D patches of various sizes taken from post-operative CT. Sagittal views are shown with the voxel sizes listed above.

Patches were made according to the overlap between the input CT, the target contact mask, and a third cylindrical mask representing the electrode trajectory. The third mask was created to ensure enough patches were sampled with contacts present. Briefly, Bresenham’s Algorithm was used to first draw a line from two contact coordinates at the entry of the electrode at the inner table of the skull and the electrode tip in the CT space. This line was then dilated to be a 3-dimensional cylinder with a radius of 4 mm and smoothed with a Gaussian kernel using the *scikit-image* and *scipy* libraries (van der Walt et al., 2014; Virtanen et al., 2020). The radius was chosen to include the soft tissue surrounding the contacts. To augment the data for training, patches were randomly rotated at angles within 30 st.dev of a normal distribution.

The number of randomly sampled patches from the training set ($n = 120$) and validation set ($n = 17$) are summarized below:

- 64x64x64 Voxel Patches: 7626 for training, 957 for validation
- 96x96x96 Voxel Patches: 6201 for training, 768 for validation

2.3 U-Net model architecture

A 3D U-Net (`monai.networks.nets.UNet()`) was implemented using the Medical Open Network for AI (MONAI) ([version 1.2, development 2316](#)). MONAI is a PyTorch based framework for deep learning in medical imaging (Cardoso et al., 2022). Training was done on two NVIDIA Tesla P100 Pascal GPUs with 12 GB of memory. Each level of the network consisted of two convolution layers followed by a pooling layer, increasing from 64, 128, 256, and 512 filters. The decoder block used the same number of filters, instead decreasing from 512, 256, 128, and 64. Strided convolutions were used in the encoder path and transpose convolutions were used in the decoder path. A kernel size of (3,3,3) was used for both convolutional and pooling filters. Regularization techniques included a dropout of 20% and batch normalization. The Adam optimizer was used, with an initial learning rate of 0.001 and default parameters (Kingma & Ba, 2017). A combination Dice + Cross-Entropy Loss was used, where the Dice and Cross Entropy losses were equally weighted. An evenly weighted Dice + Cross-Entropy loss is standard and has shown to outperform models trained with just the Dice Loss function on external test data (as assessed by the Dice Coefficient) (Galdran et al., 2023; Ma et al., 2021; Taghanaki et al., 2021). A batch size of 8 was used for training, as this was the number of 96x96x96 patches that could fit into GPU memory. The same architecture and hyperparameters were trained for the two patch sizes (64 and 96 isotropic). Early-stopping was implemented if no reduction to the average Dice CE Loss was observed in the validation set for over 50 epochs. If 50 epochs were not reached, models would be trained for a maximum of 200 epochs.

2.4 Post-processing

The U-Net predicts a probability map from the input CT data, which corresponds to the locations of SEEG contacts (represented as a binary mask of 1 mm spheres). This output is typically thresholded or transformed using an activation function to generate the final binary segmentation mask. The coordinate for each contact can then be found from the centroid of each discrete segmentation. As mentioned in Section 1.6, contact localization algorithms determine the contact coordinate from the center of mass of the radiological artifact. However, this approach will fail for electrodes implanted vertically, as the artifact encapsulates the entire electrode (see Figure 2-3). Similarly, a U-Net will predict connected segmentations in most cases rather than individual

contacts. For this reason, model outputs must undergo additional image processing (ex., erosion or additional segmentation) to extract individual contacts from the connected electrode.

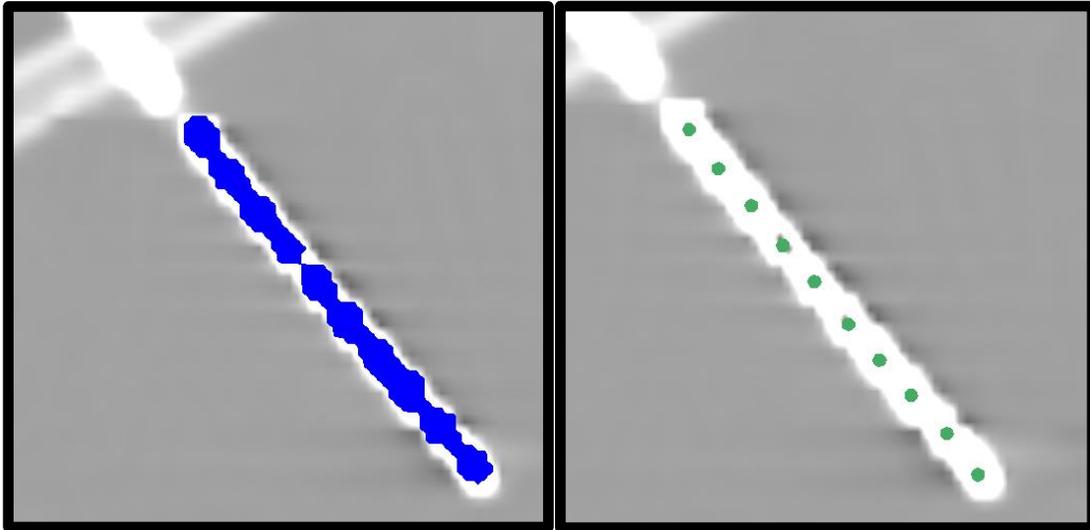


Figure 2-3. CT of an electrode implanted vertically, with the U-Net model prediction in blue. On the right are the “ground-truth” coordinates for each contact in green. Individual contacts cannot be separated in the electrode artifact or by the deep learning model.

Once the final contact locations are determined, they also must be labeled according to the target region and their position along the electrode. The traditional naming scheme in other algorithms labels electrodes by the targeted structure, with the deepest contact along the electrode (i.e. the contact farthest from the skull) as the *first* contact. Therefore, the predictions from the deep-learning model must undergo three post-processing steps: separating connected segmentations into individual contacts, extracting the coordinates of each contact, and labeling each contact with their position along the electrode. The first subsection will cover the probabilistic non-max suppression algorithm used to separate “connected” contacts and obtain the coordinate locations. The second will cover how the final coordinates were labeled and filtered for over-predictions.

2.4.1 Probabilistic non-max suppression

Non-max suppression is a technique used in object detection and computer vision tasks to reduce the number of overlapping segmentations. A probabilistic non-max suppression algorithm from the MONAI library was implemented to identify relevant contacts (Cardoso et al., 2022).

Briefly, the algorithm iteratively selected the voxel with the highest probability from the predicted map. The coordinate for that voxel was recorded before neighboring voxels in a 10x10x10 cube were suppressed to 0. This process was repeated until the remaining probabilities were below 0.5.

2.4.2 Contact labeling

During the surgical workflow for SEEG, two coordinates are placed to mark the electrode trajectory. The first is where the electrode enters the dura mater and inner table of the skull, termed the entry point. The second is at the tip of the intended target in the brain. This trajectory is first defined during the pre-surgical plan and later confirmed post-operatively to assess the accuracy of implantation (Cardinale et al., 2013). The placement of these two points was used to label predicted contacts along the electrode. This algorithm also served to filter out over-predictions or false positives in the model (i.e. instances where the probabilistic non-max suppression algorithm outputs two coordinates for a single contact). A diagram of the workflow can be found below in Figure 2-4. Unlabeled contacts provided by the U-Net and non-max suppression algorithm are iteratively labeled along the electrode trajectory. The first contact for an electrode is identified as being the point closest to the target point within an arbitrary search space of 2 mm from the target point. The next contact belonging to that electrode is selected based on a restricted search space of 35° from the electrode trajectory and 2 mm from the first contact. Once two contacts in an electrode are found, the average distance between these localizations is used to restrict the search space. When at least three contacts have been found, the trajectory is also adjusted to be the line of best fit between these points. Following a first pass of the labelling algorithm, the number of contacts identified for each electrode trajectory is recorded. The most frequently occurring contact number, was used in a second pass of the algorithm to indicate when to stop the labelling of contacts along an electrode.

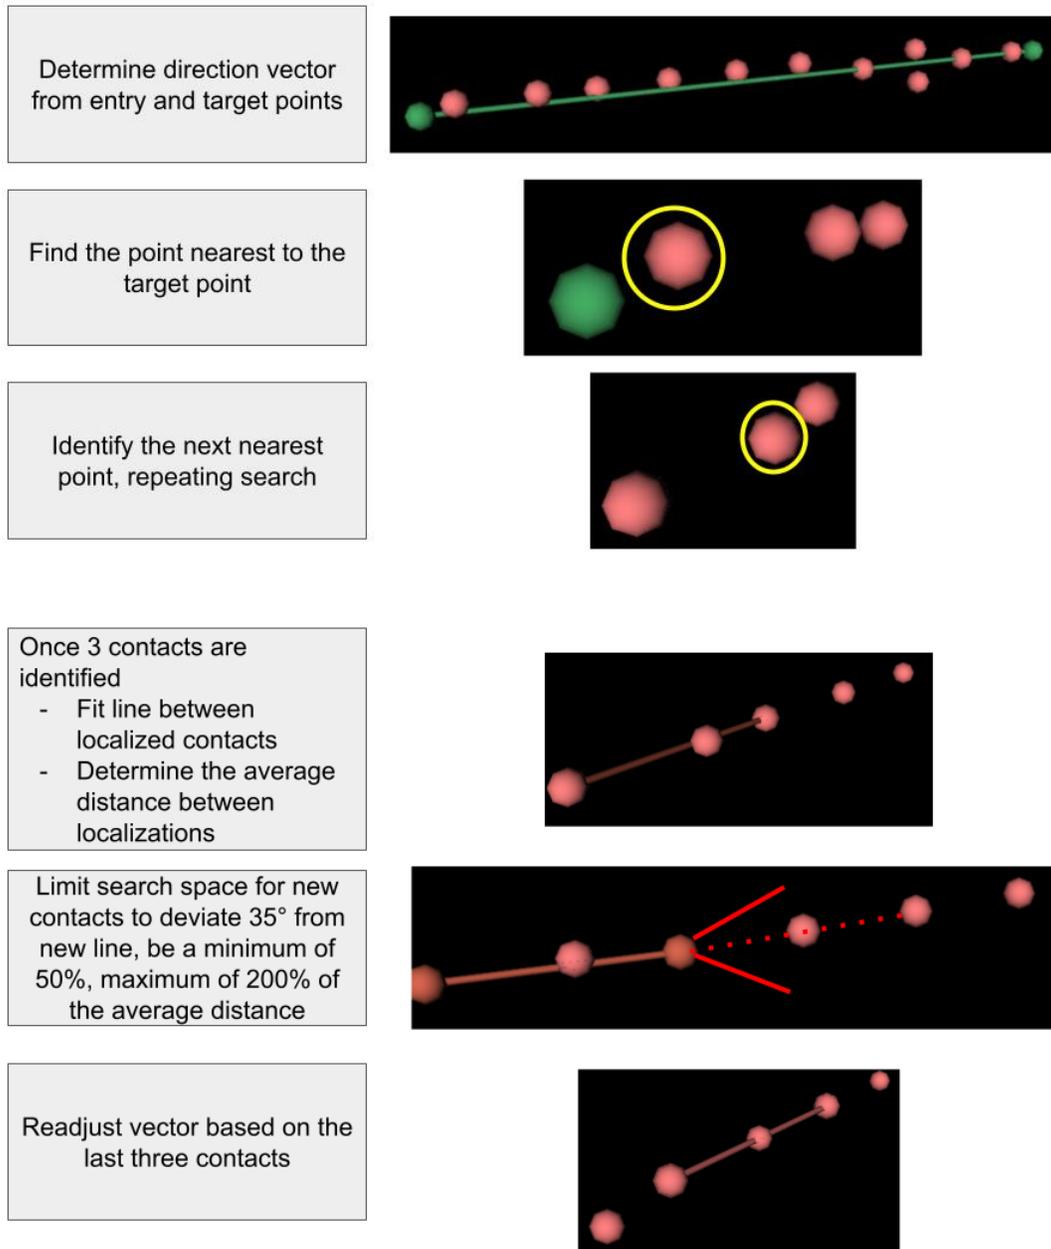


Figure 2-4 Flowchart of contact labelling for U-Net predictions.

2.5 Model validation

The accuracy was assessed in two ways; by the number of correctly localized contacts and the distance from the manually localized coordinates, which acted as our ground-truth. The model prediction at each contact was determined as a *true positive* if the prediction was within 1 mm of the manually localized coordinates. The remaining predictions were labeled *false positives* and

missed contacts as *false negatives*. This can be used to assess the accuracy of the model by defining the *precision*, *recall*, and *F1-score*, which are standard metrics in binary classification tasks. These measures were first used to assess the effect of patch size on model performance in the validation set. They were then used to assess the final localizations for the best performing U-Net model and the traditional image processing approach. The formulas for each are shown below:

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

$$Recall = \frac{TP}{TP + FN}$$

$$F1 - score = 2 * \frac{precision * recall}{precision + recall}$$

The second measure of the model's accuracy uses the Euclidean distance (described in Section 1.3.1) between the model predictions (a) and the manually localized contact coordinate (b). The Euclidean distance represents the linear distance between the two points in Cartesian space.

$$Euclidean\ Distance\ (a, b) = \sqrt{(a_x - b_x)^2 + (a_y - b_y)^2 + (a_z - b_z)^2}$$

The absolute difference along the x (lateral), y (anteroposterior), and z (superoinferior) directions was also determined for each predicted position.

$$\Delta X = |a_x - b_x| \quad \Delta Y = |a_y - b_y| \quad \Delta Z = |a_z - b_z|$$

Statistical comparisons for the Euclidean distance and absolute difference between the methods was assessed using a Wilcoxon signed-rank test using the `scipy.stats.wilcoxon()` function from the statsmodel package (Charlier et al., 2022; Seabold & Perktold, 2010).

2.5.1 Traditional image processing for contact localization

The effectiveness of the deep learning model was also compared to a traditional image processing algorithm for subjects within the test set ($n = 35$). The post-operative CT images were

thresholded at 2500 Hu to include contact, electrode and bolt artifacts, before being masked by the intracranial volume. The mask was generated using the *MNI152NLin2009cSym* template (Fonov et al., 2009). The pre-operative T1w MRI was first non-linearly registered to the template, to determine the transformation. The template was then binarized to create a mask, before applying the inverse transform with the `antsApplyTransforms` algorithm from Advanced Normalization Tools 2.2.0 (Avants et al., 2008). The number of discrete objects with 3-dimensional connectivity within this thresholded image were identified with the `skimage.morphology.label()` function. The `skimage.measure.regionprops()` function extracted the centroid and volume for each unique object (van der Walt et al., 2014). The extracted objects were filtered to include volumes between 1 mm^3 and 40 mm^3 (slightly larger than the volume of a sphere with a radius of 2.5 mm). This range eliminates objects that could be mistaken as a contact and instances where multiple contacts can appear “fused” in the post-operative imaging (ex., electrodes implanted vertically). The centroids from the remaining segmentations formed the localizations for this image processing approach. The threshold, volumetric range and a similar masking approach was used in the SEEGLoc tool (Wong et al., 2024). These points were transferred to the T1w space using the previously computed registration and labeled by electrode according to the post-operative target and entry points.

2.5.2 Comparison to uncorrected SEEG Assistant (SEEGA) placements

The SEEG Assistant (SEEGA) tool was also run on the external test set ($n = 35$) to assess performance before the localizations are manually corrected. This allowed for a comparison between the U-Net method and the tool implemented clinically. SEEGA builds off the DEETO algorithm previously published by their group in 2015, requiring two points, the inter-contact spacing, and the number of contacts for each electrode (Arnulfo et al., 2015). For this comparison, the entry and target points for each electrode from the post-operative trajectory were used. The entry points were manually corrected to be at the start of the electrode artifact closest to the inner table as viewed in 3D Slicer. SEEGA uses these points to first estimate the electrode axis. From here the algorithm identifies the image artifacts falling along the axis that match the inter-contact spacing of the electrode. This search is also restricted to artifacts that deviate 10 degrees from the original axis. Contact positions are then placed according to these parameters,

representing the unadjusted outputs from SEEGA. These were also compared to the manually corrected “ground-truth” coordinates to determine the localization error of the tool.

Chapter 3

3 U-Net Implementation Results

Chapter 3 compares the outputs of two U-Net models trained with different patch sizes on the validation set. To review, the U-Net model only required as input the CT volume resampled to 0.4mm isotropic. The contact localizations for the best-performing model will be compared to the manual annotation for each contact in the test set. Finally, the U-Net error will be compared to the semi-automated SEEG Assistant (SEEGA) tool and a traditional image processing approach.

3.1 Effect of patch size on model prediction

Two models were trained on 64 and 96 isotropic voxel patches to assess whether an increased patch size would improve the final performance. Table 3.1 summarizes the training metrics for the models using a patch size of 64 voxels and 96 voxels. Each model was trained for a maximum of 200 epochs with early stopping implemented at 50 epochs if no change was observed to the validation DiceCE loss. Figures 3-1 and 3-2 show the training curves for the 64 isotropic voxel patches and 96 isotropic voxel patches. Both models appear to have converged at a similar epoch number, with similar metrics for the training and validation loss, as well as the final Dice Coefficient for predictions in the validation set. Figure 3-1 shows the model trained with 64 isotropic voxel patches to have overfit to the training set, in comparison to the model trained with 96 voxel patches.

Table 3-1 Training Summary and Loss Metrics

Model Patch Size (Voxels)	Training Time (hh:mm)	Epochs Before Early Stopping	Training DiceCE Loss	Validation DiceCE Loss	Validation Dice Coefficient
64x64x64	07:05	190	0.181	0.248	0.725
96x96x96	17:50	200	0.206	0.243	0.736

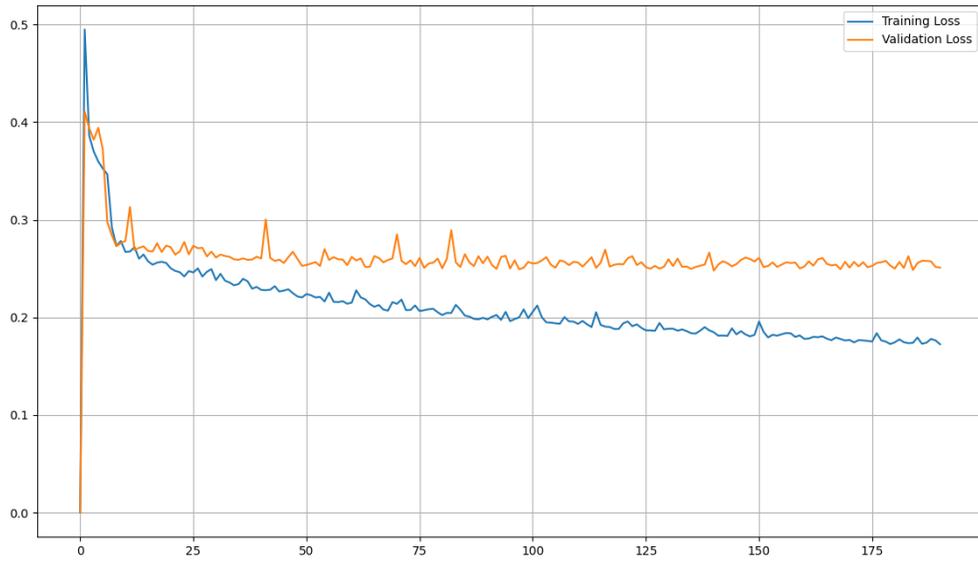


Figure 3-1 Training and validation loss curve for U-Net trained with 64 voxel isotropic patch sizes. *Dice + Cross Entropy loss values shown on the y axis, with the epoch number on the x axis.*

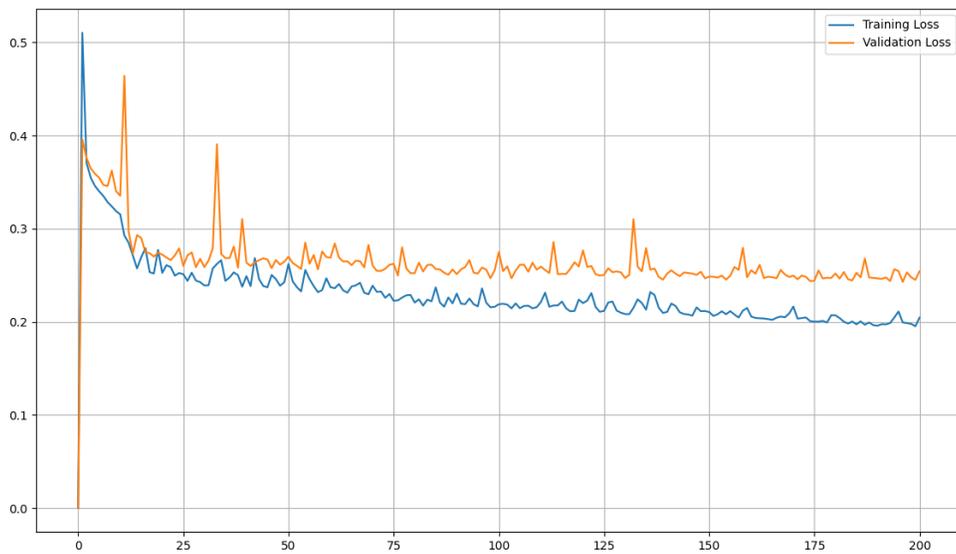


Figure 3-2 Training and validation loss curve for U-Net trained with 96 voxel isotropic patch sizes. *Dice + Cross Entropy loss values shown on the y axis, with the epoch number on the x axis.*

The effect of patch size was also compared following the use of the probabilistic non-max suppression algorithm. The coordinates provided by this post-processing step represent the model’s prediction for the most likely location of a contact (i.e. a 1 mm sphere). This was used to compare the number of contacts localized in the validation set (n = 1910 contacts). The precision, recall and F1-score were determined as described in Section 2.5 and summarized in Table 3.2. These metrics were taken before the predictions were filtered by the labeling algorithm. The goal was to assess how effective each model was in accurately detecting contact positions, without being aided by knowledge of the electrode trajectory. The precision score indicates what proportion of predictions are relevant (true localizations of contacts vs noise). The recall can be viewed as the accuracy of the model, while the F1-score is a weighted balance of both measures. The deep learning models were also compared to the number of localized contacts found by thresholding the CT image at 2500 Hu, masking for the intracranial volume, and extracting the weighted centroids from the remaining artifacts.

Table 3-2 Contact localization accuracy of the U-Net model in the validation set, based on patch size (n = 1910 contacts). *Compared to weighted centroids of artifacts with 3-D connectivity from thresholded CT at 2500 Hu.*

Model Patch Size (Voxels)	# of Contacts	Total TP	Total FP	Total FN	Precision	Recall	F1-score
64x64x64	1910	1823	569	87	0.762	0.954	0.848
96x96x96	1910	1818	634	92	0.741	0.952	0.834
<i>Thresholded CT (2500 Hu)</i>	1910	1279	139	631	0.902	0.670	0.769

Patch sizes of 64 and 96 voxels provided similar performance metrics for precision, recall and F1-score, with the 64-patch model slightly out-performing the 96-patch model. While the traditional image processing approach of thresholding the CT limited the number of false positives, it was unable to resolve a large portion of contacts positions (631/1910) resulting in lower recall and F1-score. These corresponded to electrodes implanted vertically and with a smaller inter-contact spacing.

The contact-labelling algorithm was then applied to reduce the number of false positives and label the model predictions by the electrode trajectory. The number of true positives, false positives, and false negatives in the validation set after the algorithm are summarized in Table 3-3.

Table 3-3 Contact localization accuracy of the U-Net model in the validation set, following the contact-labelling algorithm (n = 1910 contacts).

Model Patch Size (Voxels)	# of Contacts	Total TP	Total FP	Total FN	Precision	Recall	F1-score
64x64x64	1910	1816	63	94	0.966	0.951	0.959
96x96x96	1910	1809	44	101	0.976	0.947	0.961

The contact-labelling algorithm worked to filter a large portion of the “false positive” identified by the U-Net model, specifically in cases where multiple coordinates were predicted for a single contact or where the bolt/wires were incorrectly identified. Following this post-processing step, both models again achieve similar metrics as defined by the precision, recall and F1-score. The 64-patch model localized slightly more true positives as evident by the recall score when compared to model trained on 96 voxel patches. These results from the validation set appear to demonstrate that increasing the spatial context for the model from 64 to 96 voxels did not substantially change model performance. The 64-voxel model was also faster to train and required less computational resources. Considering this performance on the validation set, the 64-voxel model was used for the test set.

3.2 Localization accuracy of U-Net approach

The remaining sections of this chapter will focus on the results from the test set (n = 35 patients, 402 electrodes, 4020 contacts). To reiterate, *true positive* localizations were within 1 mm of the ground-truth contact localization, *false positive* localizations were greater than 1 mm from the ground-truth contact position, and *false negatives* were contacts missed by the model. The accuracy by subject for the U-Net is summarized in appendix A. The total accuracy of the original model predictions and those filtered by the labelling algorithm in the test set of 35 subjects is summarized in Table 3-3. This was compared to a traditional image processing

approach, where the CT was thresholded at 2500 Hu and then masked by the intracranial volume. Contact localizations were determined by extracting the centroids for the remaining objects between 1 mm³ to 40 mm³ in volume.

Table 3-4 Contact localization accuracy of the U-Net model on the test set. (n = 35 subjects, 4020 contacts, 402 electrodes). Compared to a traditional image processing approach of thresholding the CT at 2500 Hu and masking the intracranial volume.

Description	# of Contacts	TP	FP	FN	Precision	Recall	F1-score
<i>U-Net without post-processing (64 voxel patches)</i>	4020	3858	1076	162	0.782	0.960	0.862
<i>U-Net with post-processing (64 voxel patches)</i>	4020	3837	84	183	0.979	0.954	0.966
<i>Thresholded CT (2500 Hu)</i>	4020	2545	86	1475	0.967	0.633	0.765

The image processing approach achieves a high precision of 0.967, indicating the thresholded contact localizations match the manually corrected or ground-truth position. However, the recall score of 0.633 shows that a large proportion of contacts are missed using this approach. In comparison the U-Net predictions using a probabilistic non-max suppression algorithm and filtered by the electrode trajectories displayed high precision and recall scores of 0.979 and 0.954. From the test set, six subjects achieved a perfect F1-score, indicating that all contacts were within 1 mm of the manually localized coordinate, with no false positive or negative predictions. Figure 3-3 shows the contact localizations for subject 126. Figure 3-4 shows the locations of all missed contacts or false negatives by the deep learning model in the *MNI152NLin2009cSym* template space (Fonov et al., 2009). Note that contacts that are outside of the brain in the following plots represent contacts found at the level of the skull or near the bolt of the electrode, where the bony and metal artifacts obscure the contact.

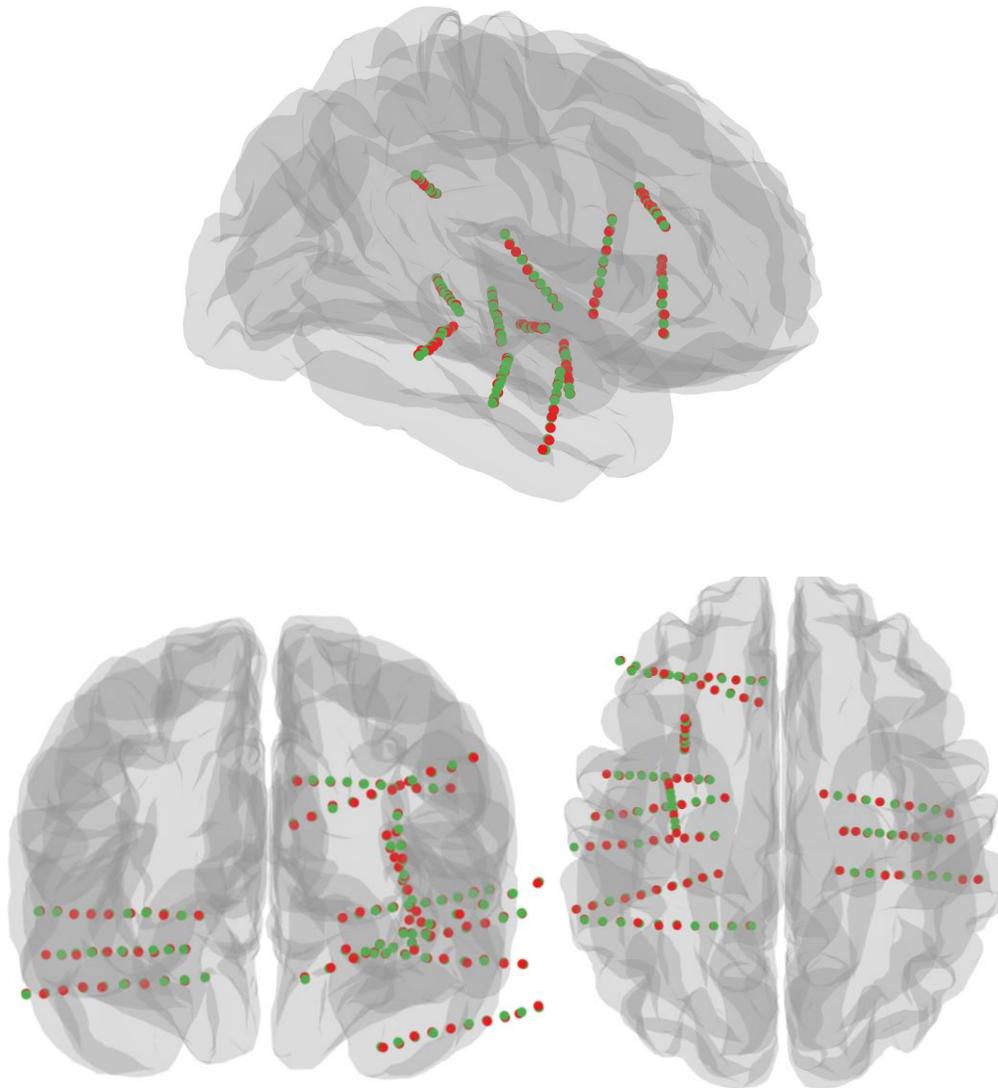


Figure 3-3 U-Net model contact localizations (in red) and ground-truth manual placements (in green) for subject P126, plotted in MNI152NLin2009cSym template space (n = 120 contacts). Side (right hemisphere), front, and top views are shown. The model achieved a perfect *F1*-score, indicating all contacts were found within 1 mm with no false positives or negatives.

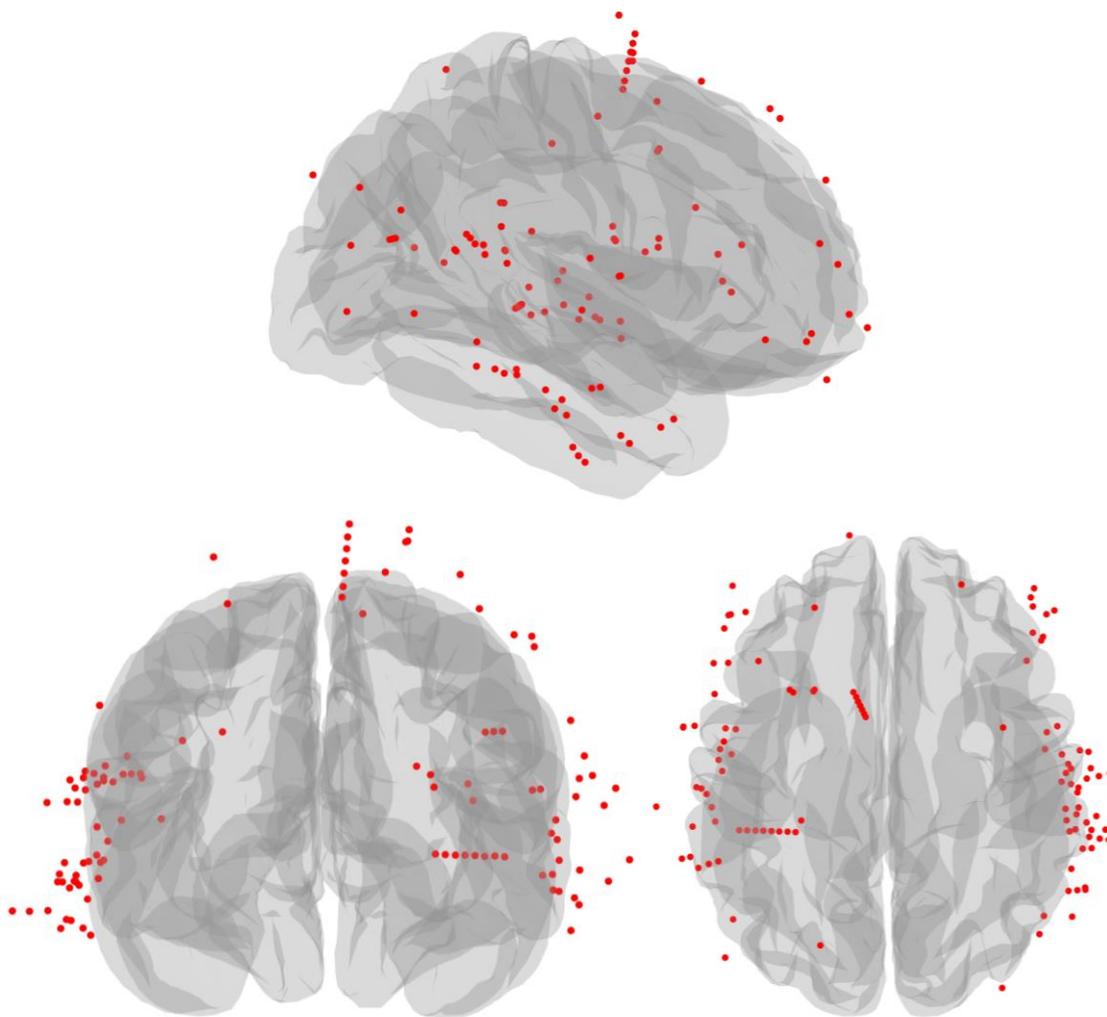


Figure 3-4 False negatives or contacts missed by the U-Net model, plotted in MNI152NLin2009cSym template space (n = 183/4020). From left to right, side (right hemisphere), top, and front views are shown.

3.3 Localization error for U-Net compared to SEEG Assistant

The localization error for each contact predicted by the U-Net model can be quantified using the Euclidean distance from the manually annotated contact, which acts as our “ground-truth”. This same error was calculated for the points provided by the semi-automated SEEGA tool (Section 2.1.2). The following subsections will compare the localization error between the two approaches by subject and contact-spacing within an electrode. It will then compare the absolute difference by each approach along the x, y, and z planes.

3.3.1 Distribution of errors

The error distribution for the U-Net model was contrasted against the unadjusted SEEGA localizations within the test set. This analysis included 3912 contacts by the U-Net that had a matching label to the SEEGA label. The mean error for the automated U-Net model was 0.54 ± 0.53 mm, with a median of 0.46 mm. For these same points, the mean error of SEEGA was 0.57 ± 0.76 mm and a median of 0.43 mm. It is worth noting that the ground-truth or manually localized points were derived from corrections to SEEGA outputs. This introduces some bias for the unadjusted SEEGA, as there is a potential for the manually placed points to be identical to the outputs from SEEGA lowering the localization error for the unadjusted points. In contrast, the U-Net errors represent an unbiased assessment or prediction of contact localizations.

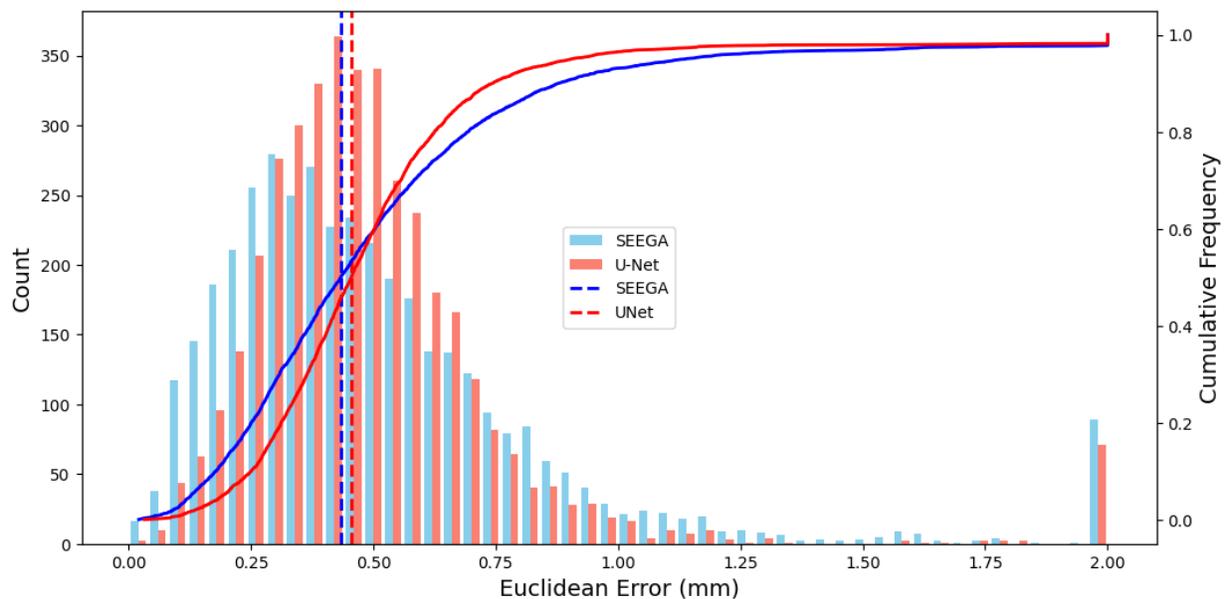


Figure 3-5 Distribution of Euclidean error (in mm), comparing the semi-automated (SEEG Assistant) localizations to those matched by the automated U-Net method. $n = 3912$ contacts. Dashed lines represent the median error, while curves indicate the cumulative frequencies. U-Net median = 0.46 mm, SEEGA median = 0.43 mm.

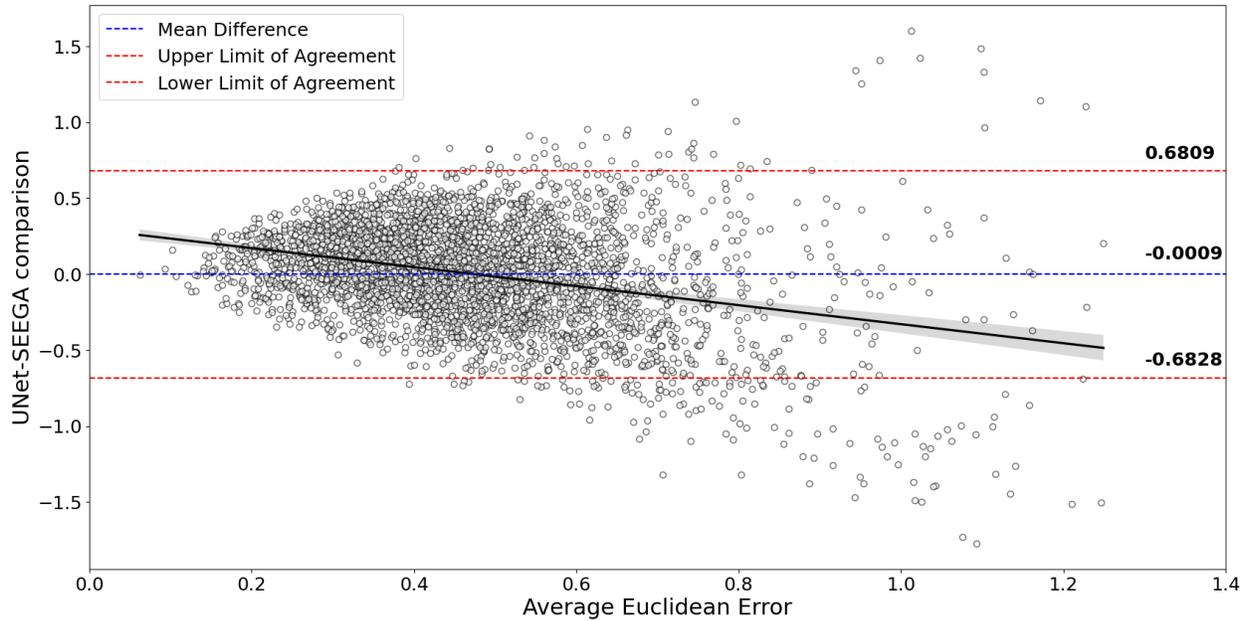


Figure 3-6 Bland-Altman plot of Euclidean error (in mm), comparing the U-Net method to the SEEGA localizations. *Thresholded at 2 mm for both methods, $n = 35$ subjects, 3757 contacts. Upper and lower limits were set at $\pm 1.96 * 1$ standard deviation.*

Figure 3-6 is a Bland-Altman plot contrasting the U-Net and SEEGA error for paired contacts. Errors for both methods greater than 2 mm were removed for this analysis, reducing the number of contacts for comparison to 3757 contacts. The U-Net matches with SEEGA for most contact positions, as indicative by the low mean difference at -0.0009 mm. Interestingly, the U-Net appears to slightly outperform SEEGA as the average error of both methods increases.

3.3.2 Localization error due to electrode bending

SEEGA restricts contact localizations to deviate 10 degrees from the electrode axis, which can arise in larger errors if contacts bend from the electrode trajectory (Narizzano et al., 2017). To assess the effect of electrode bending on localization performance, the bending of an electrode was quantified as the distance of the actual contact position (i.e. the “ground-truth contact localization) from the line joining the target and entry points (i.e. the electrode trajectory). Figure 3-7 visually shows a visual schematic of the determination of this distance. Briefly, this distance was determined from two vectors: the first (*a*) being the vector from the entry point to the individual contact and the second (*b*) being the vector from the entry to the target point. The

distance was the magnitude of the vector perpendicular to the projection vector of a onto b , seen in red and green respectively.

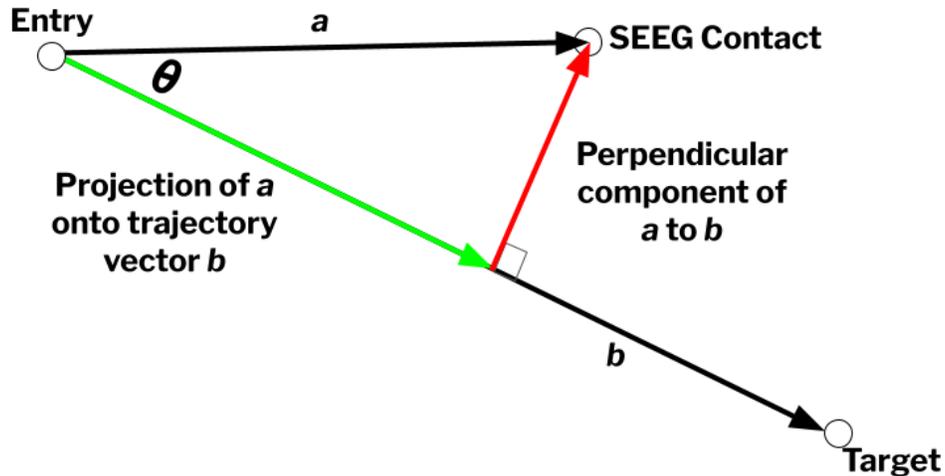


Figure 3-7 Visual schematic of the perpendicular distance between the manual contact position and electrode axis used to quantify the degree of electrode bending. a represents the vector between the entry point and SEEG contact, while b represents the vector between the entry and target points.

The larger the magnitude of this perpendicular component, the further away the contact was from the electrode axis indicative of electrode bending. Figure 3-8 examines the relation between the localization error (the Euclidean distance from the predicted contact position and the ground-truth) and this perpendicular axis distance. Interestingly as this perpendicular component increases beyond 1 mm, the localization error for SEEGA also increases. In contrast, the error for the U-Net model remained relatively consistent and at sub-millimetric values. This trend quantitatively suggests that SEEGA performs worse as contacts bend or deviate away from the electrode trajectory, while there is a weaker correlation for the U-Net localization.

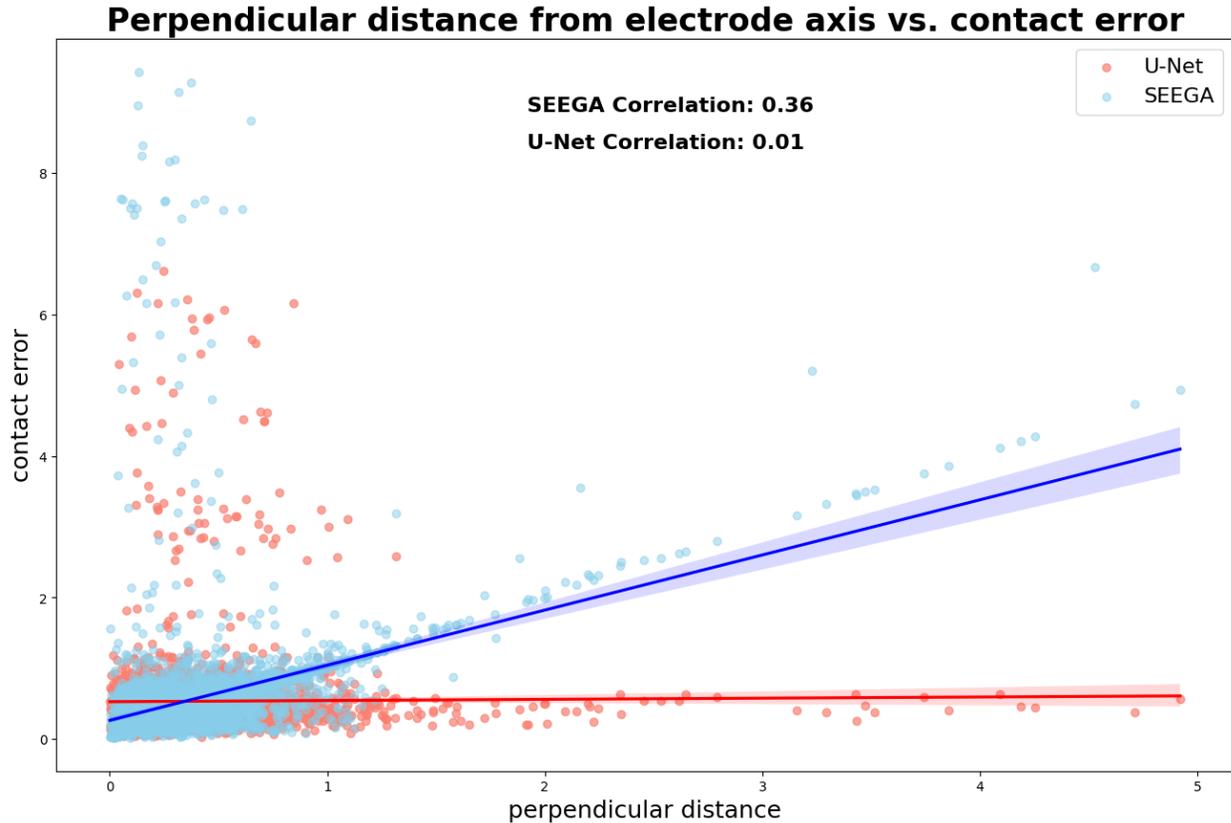


Figure 3-8 Perpendicular distance of each contact’s position from the electrode trajectory and the contact localization error for both the U-Net and SEEGA methods. *Pearson’s correlations are included for both methods ($r = 0.36$ for SEEGA, 0.01 for the U-Net). Linear regression lines for both methods are also plotted with a 95% CI band.*

3.3.3 Localization error at a subject level

Next, the localization error by subject was reviewed for both methods. Errors for both methods were thresholded at 2 mm. This reduced the number of contacts for comparison to 3757. The localization for 6 subjects by the U-Net model outperformed SEEGA and reached significance ($p < 0.05$). A summary can be found in Figure 3-9.

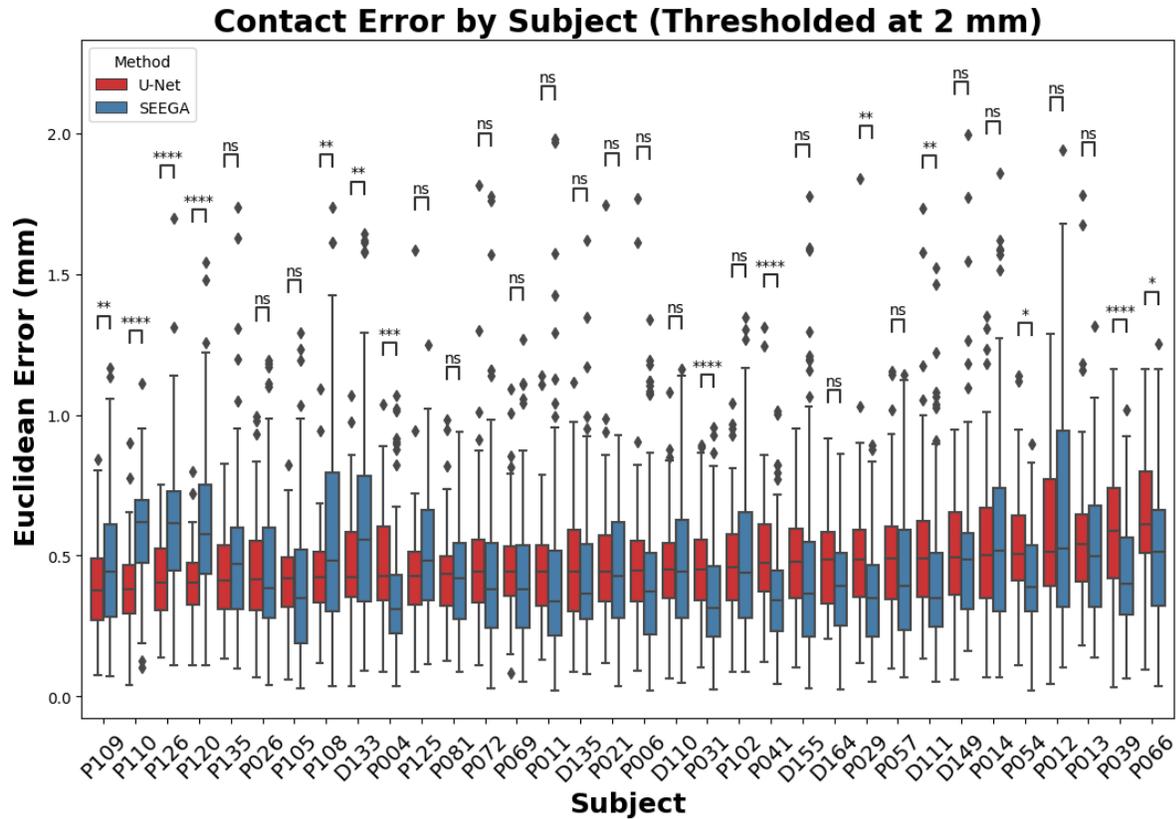


Figure 3-9 Localization error by subject for the U-Net model and SEEGA, thresholded at 2 mm. Ordered by the mean error for the U-Net model (increasing from left to right). Significance was assessed using a Wilcoxon signed-rank test with Bonferroni correction. $n = 35$ subjects, 3757 contacts.

3.3.4 Error by inter-contact spacing

The localization error for each approach was compared based on the spacing of contacts within an electrode. As mentioned in Section 2.1.1, the model was trained on patients implanted with 10-contact Ad-Tech electrodes, with 3-mm, 4-mm, 5-mm, 6-mm or 7-mm contact spacing. The localization error for the automated U-Net and SEEGA for each electrode model is shown in Figure 3-10 and Table 3-5. This was compared for errors less than 2 mm by both methods.

Error by Contact Spacing (Thresholded at 2 mm)

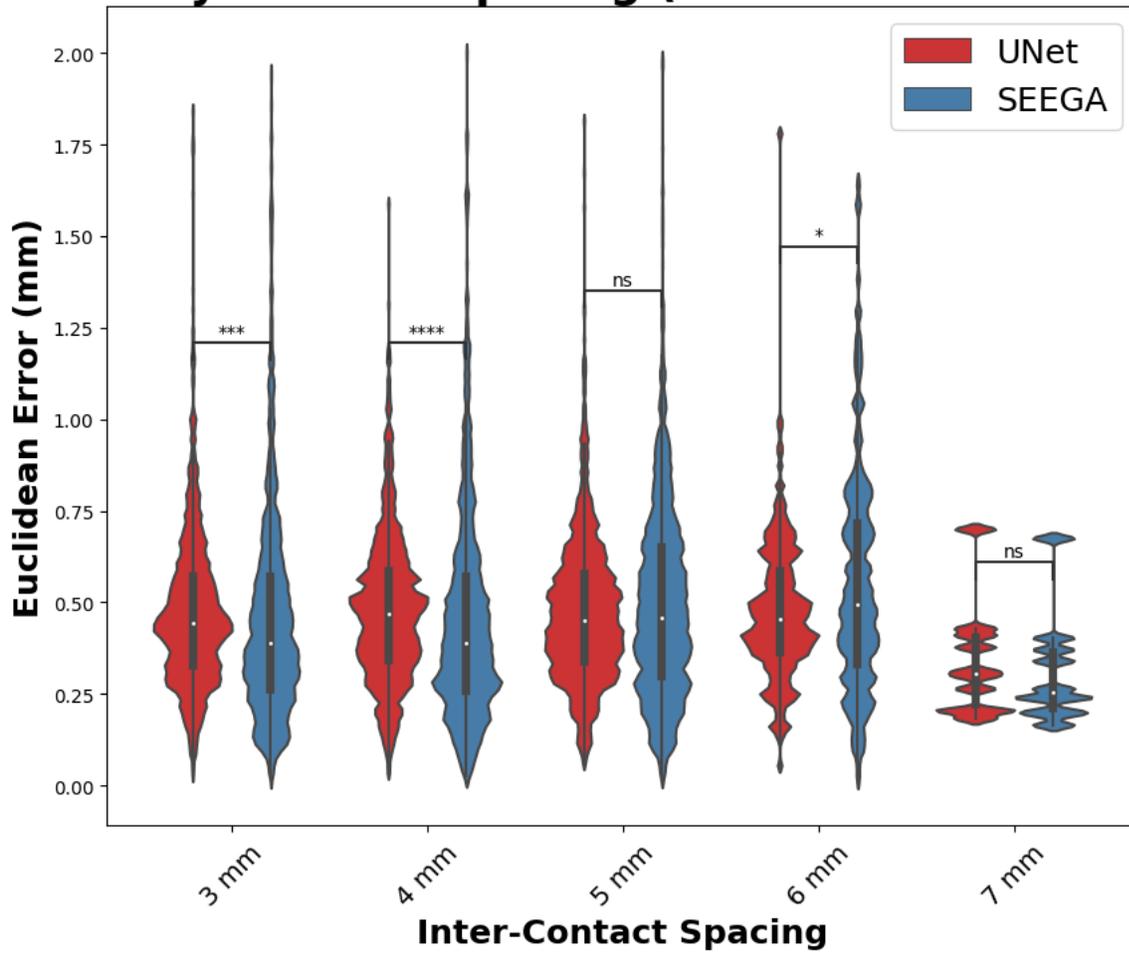


Figure 3-10 Error by electrode inter-contact spacing for U-Net and SEEGA methods. Significance was assessed using a Wilcoxon signed-rank test, with Bonferroni correction. $N = 35$ subjects.

Table 3-5 Average Euclidean error by contact, where n indicates number of contacts for each electrode type. Matched localizations for both the U-Net and SEEGA methods, where both errors were below 2 mm.

Error (Mean ± SD mm) by Electrode Inter-Contact Spacing					
Method	3 mm (n = 1404)	4 mm (n = 1094)	5 mm (n = 948)	6 mm (n = 192)	7 mm (n = 10)
U-Net	0.471 ± 0.215	0.473 ± 0.192	0.470 ± 0.198	0.593 ± 0.196	0.433 ± 0.139
SEEGA	0.439 ± 0.266	0.469 ± 0.303	0.490 ± 0.266	0.462 ± 0.340	0.520 ± 0.224

3.3.5 Absolute error in x, y, and z dimensions

The absolute error in each dimension from the manually localized contact coordinate was also calculated for both the U-Net and SEEGA tools. Figure 3-11 and Table 3-6 show the absolute error between the semi-automated SEEGA and the U-Net model. Despite reaching significance both the U-Net and SEEGA models appear to display a similar distribution of error.

Error In Each Dimension (Thresholded at 2 mm)

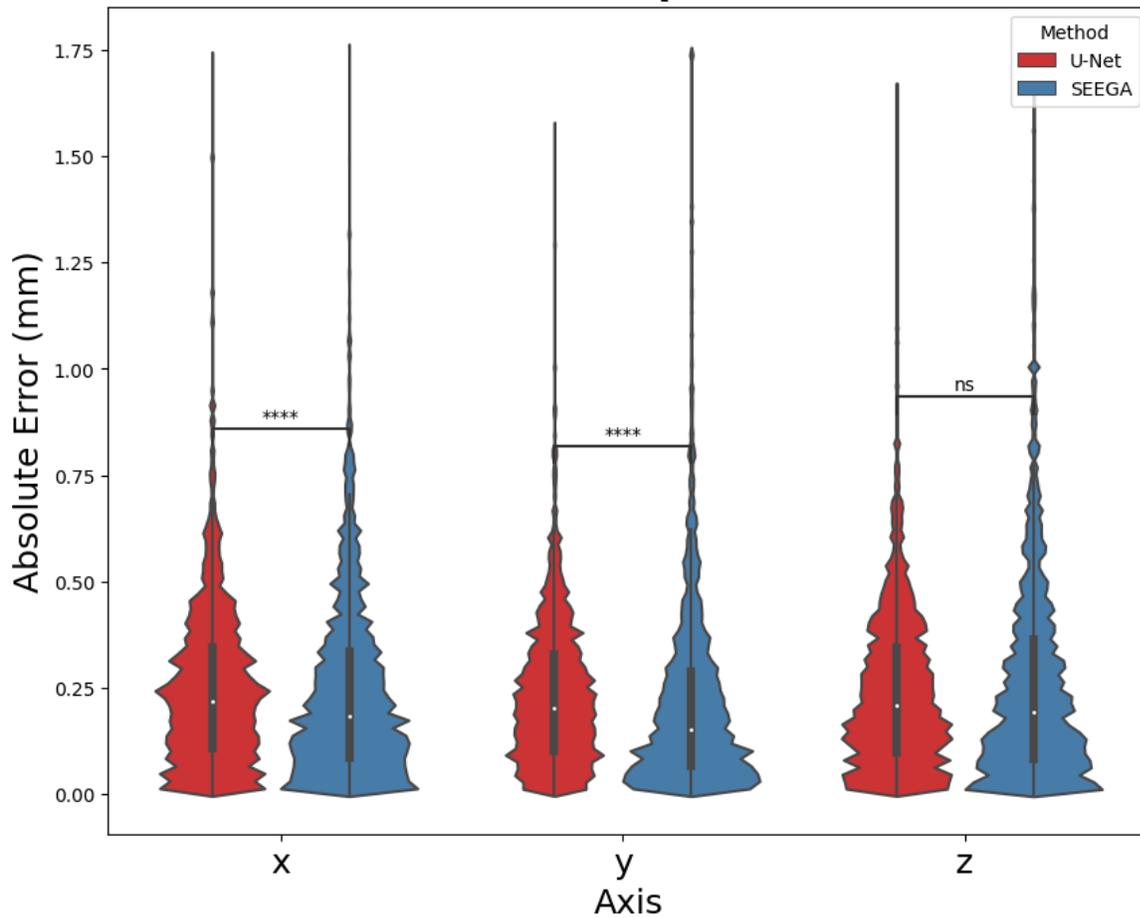


Figure 3-11 Error in each dimension by contact for the U-Net approach compared to **SEEGA**. Significance was assessed using a Wilcoxon signed-rank test with Bonferroni correction. $n = 35$ subjects, 3757 contacts.

Table 3-6 Mean absolute error in each dimension by contact and method. $n = 35$ subjects, 3757 contacts.

Mean Absolute Error \pm SD (mm)			
Method	x	y	z
U-Net	0.244 ± 0.178	0.230 ± 0.166	0.240 ± 0.180
SEEGA	0.237 ± 0.203	0.211 ± 0.203	0.255 ± 0.203

3.4 Localization error of traditional image processing approach

The following section compares the error of a traditional image processing approach that extracted the centroids from the post-operative CT thresholded at 2500 Hu. The localization error for the image processing approach was calculated using the Euclidean distance to the manually annotated contact, which acts as our “ground-truth”. This was compared to the contacts with the same label predicted by the U-Net model in Figure 3-12. A similar comparison was performed for the SEEGA tool in Figure 3-13. Despite the low number of contacts localized, the weighted centroid from the image processing achieved a lower average error of 0.14 mm when compared to corresponding localizations from the U-Net and SEEGA methods.

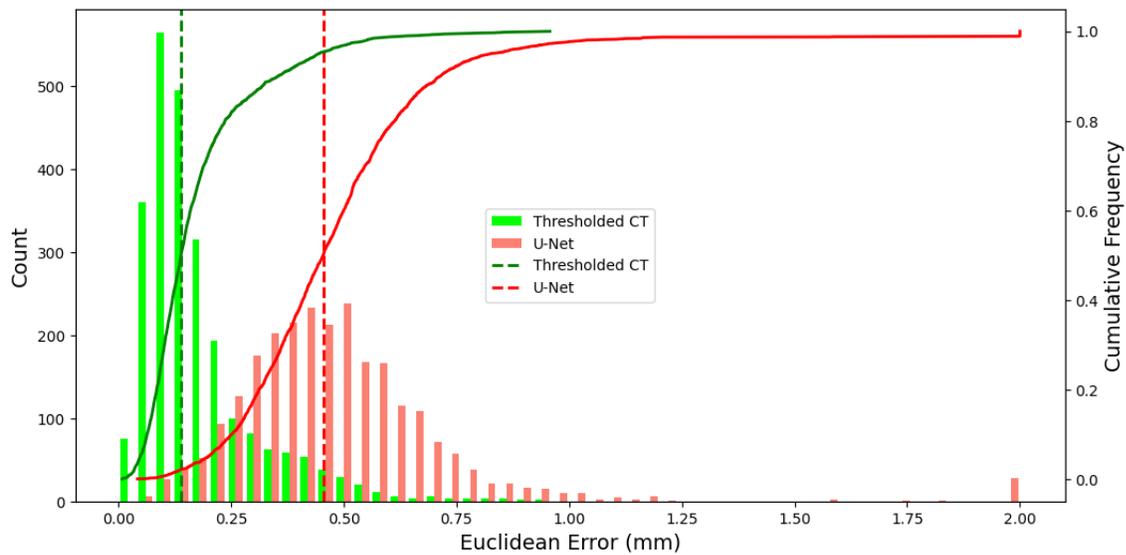


Figure 3-12 Distribution of Euclidean error (in mm), comparing the automated (U-Net) approach to the matched localizations by the traditional image processing algorithm. Thresholded for errors below 2 mm, $n = 2501$ contacts. Median U-Net = 0.45 mm (dashed red line), Median Thresholded CT = 0.14 mm (dashed green line).

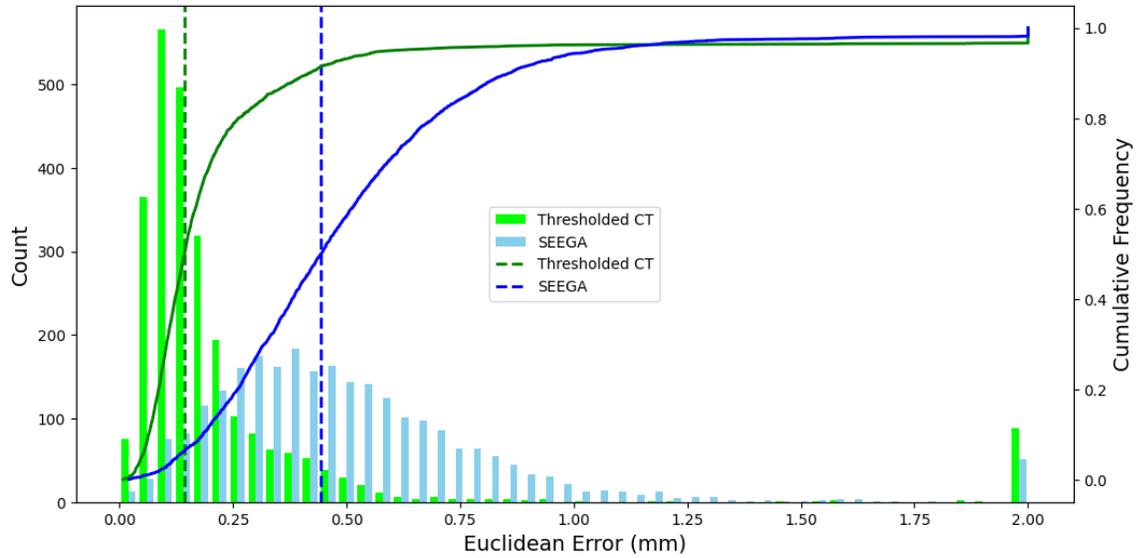


Figure 3-13 Distribution of Euclidean error (in mm), comparing the semi-automated SEEG Assistant (SEEGA) tool to the matched localizations by the traditional image processing algorithm. $n = 2604$ contacts. SEEGA median = 0.44 mm (dashed red line), Thresholded CT median = 0.14 mm (dashed green line).

Finally, a comparison of the false negatives for both the U-Net and traditional image processing approach can be seen in Figure 3-14.

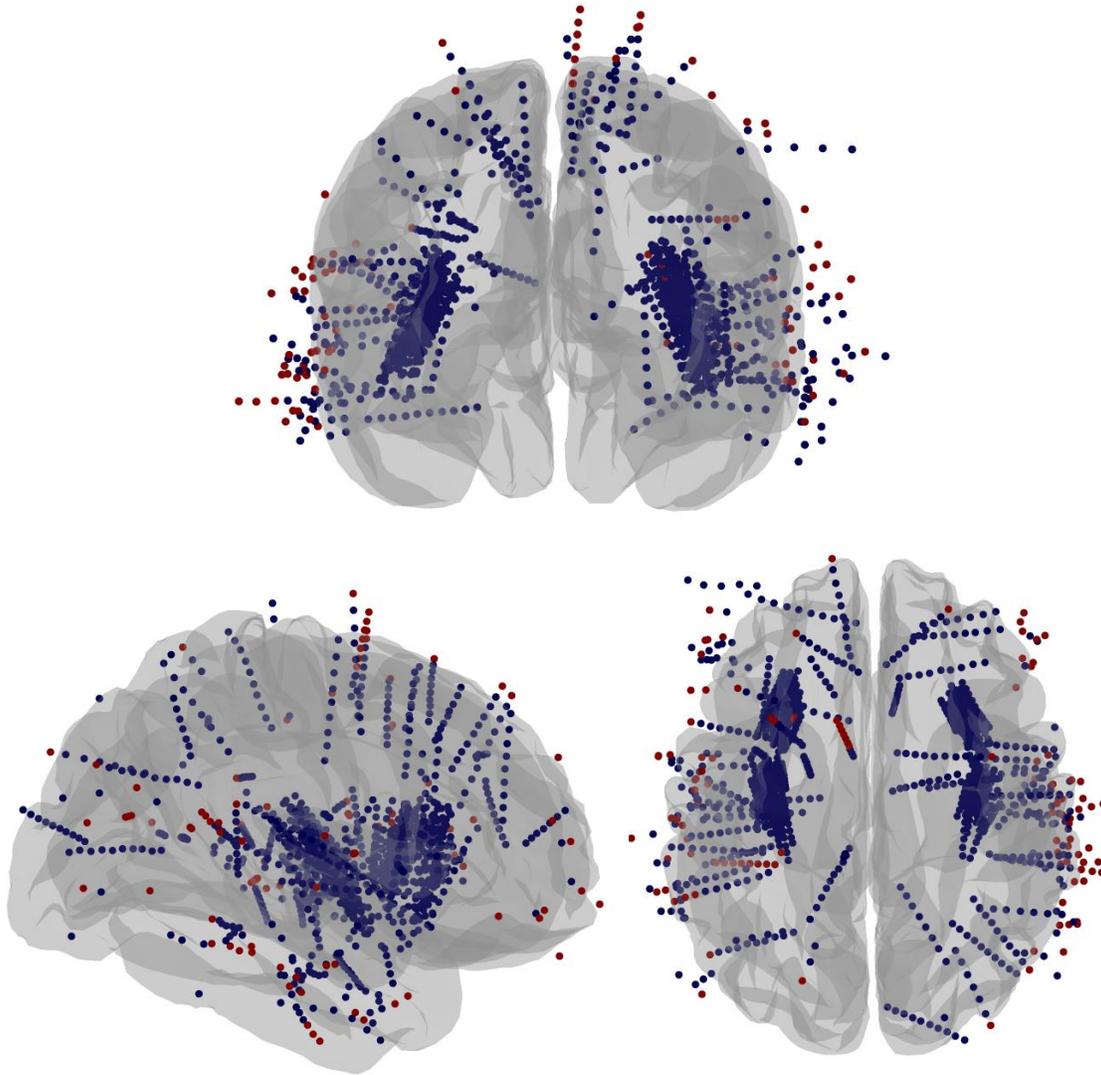


Figure 3-14 Contacts missed (false negatives) by the traditional image processing approach and the U-Net. *In red are the false negatives by the U-Net model, with blue representing the false-negatives by the outlined image processing approach.*

Chapter 4

4 Discussion

In summary, the use of a 3-D U-Net in combination with a probabilistic non-max suppression algorithm could extract most contact coordinates and perform comparably to the semi-automated method used in clinical practice (SEEGA). The approach was tested on an independent sample of 35 subjects (i.e. not used for training or validation) and achieved a similar localization accuracy to the current literature. The trained model and the post-prediction labeling algorithm were implemented in Python, building towards open-science infrastructure. Model predictions could be obtained in ~40 seconds, with the only input being a postoperative CT scan resampled to 0.4 mm isotropic. Labeling and filtering were achieved using the target and entry coordinates from the post-implantation electrode trajectory. These two points are defined in electrode trajectories as part of most SEEG workflows to assess implantation accuracy (Cardinale et al., 2013; Vakharia et al., 2017). In comparison, the SEEGA algorithm employed at LHSC required manually inspecting each image, correcting the entry and target points based on the imaging artifact, and configuring the inter-contact spacing and the number of contacts per electrode, culminating in a 1–3 hour process for each subject by a trained expert. The original SEEGA paper reports an accuracy of 87.57% (8429/9626 contacts) using the pre-implantation trajectory and these priors; with manual correction this improved to 97% (9366/9626 contacts) (Narizzano et al., 2017). In comparison, the U-Net model with the post-processing algorithm had a localization accuracy of 95.4% (3837/4020 contacts), with 6 subjects achieving a perfect F1-Score (no false positives or false negatives). Of the 183 contacts missed by the U-Net, a majority (107/183) corresponded to the most superficial contacts in the electrode, obscured by the bony artefacts of the skull or bolt artefacts of the electrode. 21 contacts responded to the 9th contact from the electrode tip, while 86 responded to the 10th contact found near the bolt of the electrode.

The histogram and Bland-Altman plot in Figures 3-5 and 3-6 demonstrate that the localization errors compare favorably for both the U-Net and SEEGA methods across the test set. There are a few notable trends that can be observed. The first is that SEEGA has a larger portion of errors above 1 mm ($n = 277$) compared to the U-Net model ($n = 137$). Figure 3-8 examines the errors in relation to electrode bending and shows that the SEEGA performance worsens as the contact

position deviates from the electrode axis. In contrast, the U-Net performance is relatively consistent as the degree of electrode bending increases. Specific examples of electrode bending within the test set include subjects D149 in Figure 4-1. The U-Net achieved a perfect F1-Score (0 FP or FN), with a mean Euclidean error of 0.498 ± 0.179 mm. In comparison the error for SEEGA was 1.492 ± 1.459 mm ($p < 0.05$).

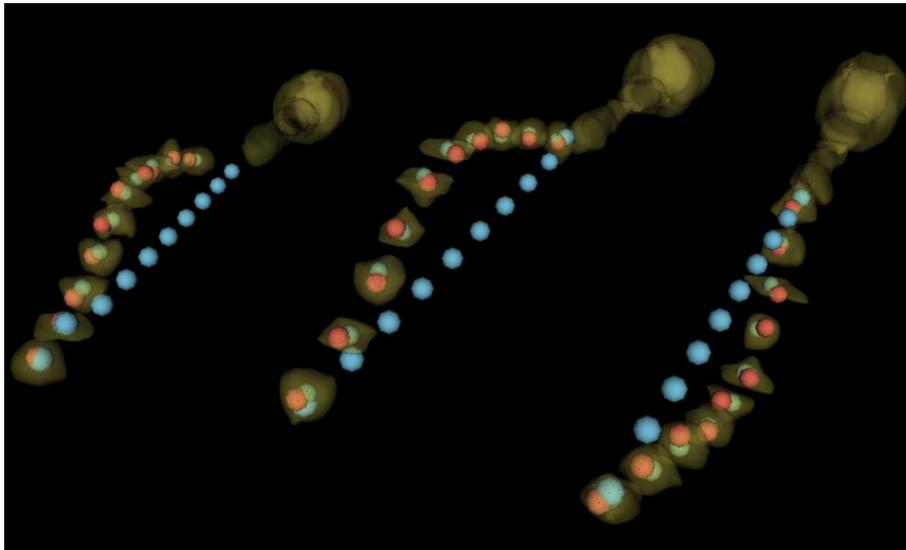


Figure 4-1 Examples of three bending electrodes in subject D149. In red are the U-Net model predictions, green are the manual localizations, and blue are the SEEGA predictions. The CT electrode/contact segmentation artifact can be seen in yellow (thresholded at 2500 HU). Captured with 3D Slicer, version 5.6.1 (Fedorov et al., 2012).

The larger errors by the U-Net model (> 1 mm) are observed when the model misses a contact along an electrode or assigns multiple coordinates to a single contact. This results in an incorrect labeling of contacts that does not reflect the manual or ground-truth position; examples include subjects D133 and P057. The distribution of errors in Figure 3-5 also shows that there is a higher proportion of errors < 0.3 mm for SEEGA. This can be attributed to the fact that the “gold-standard” manual localizations are based on the SEEGA outputs. In cases where SEEGA provides an “adequate” localization, the reviewer was unlikely to manually edit the contact position, creating an inherent bias towards the SEEGA performance.

Examining the error by the electrode model or the inter-contact spacing, the U-Net outperformed the semi-automated SEEGA for contacts with 5 mm and 6 mm spacing. The performance for the

5 mm electrode did not reach significance. In contrast, SEEGA outperformed the U-Net for electrodes with 3 and 4 mm spacing. Error by axis or dimension showed SEEGA to outperform the U-Net along both the x (0.237 mm for SEEGA, 0.244 mm for the U-Net) and y axes (0.211 mm compared to 0.230 mm), reaching significance. The U-Net had a lower absolute error along the z axis (0.240 mm), but this result did not reach significance.

The U-Net models presented here were trained with a patch size of 64 and 96 isotropic voxels, with Dice scores of 0.725 and 0.736 respectively, on the validation set. The two previously published U-Nets trained for SEEG contact segmentation in Section 1.5.2 report higher scores than our model. Pantovic et al., report an average of 0.867 from their hybrid U-Net (concatenated outputs of 2-D and 3-D models). The 2-D U-Net model (trained on slices of the CT) and 3-D U-Net model report average Dice scores of 0.850 and 0.810 respectively. The improved Dice score could be due to the use of larger patch sizes (144x144x135) for training (Pantovic et al., 2022). Vlasov et al., report an average Dice score of 0.897 ± 0.043 for 10 subjects using a 3-D Cascade U-Net mode. However, the performance of their model includes the intracranial volume and the entire electrode trajectories as part of the target segmentation.

The general application of setting an initial threshold and masking for the intracranial space to extract relevant contact artifacts is applied in a majority of contact localization algorithms (Aslam et al., 2024; Blenkmann et al., 2024; Cai et al., 2022; Davis et al., 2021; Janca et al., 2023; Medina Villalon et al., 2018; Qin et al., 2017; Wong et al., 2024; B. Zhao et al., 2023). The U-Net's performance was also compared to a similar image processing approach. The CT for each patient in the test set was thresholded at 2500 Hu and masked by the intracranial space. Contact positions were then determined from a weighted centroid including radiological artifacts between 1 mm^3 and 40 mm^3 . The average localization error from the manually annotated position appeared to outperform both U-Net and SEEGA in localization error but was limited in accuracy (2579/4020 contacts). As seen in Figure 3-14, the contacts missed by this approach were found within the midbrain, near the insular regions. These correspond to the tightly packed contacts (3-4 mm spacing) found in vertically implanted electrodes, which appear as a continuous artifact in the post-operative image (see Section 1.5.2). Determining the contact localizations for these electrodes is typically done by interpolating the remaining contacts using the inter-contact spacing, the number of contacts for that electrode, and two points placed at the

start and end of the radiological artifact (Narizzano et al., 2017; Wong et al., 2024). In comparison the use of a probabilistic non-max suppression algorithm was able to identify most contacts within said electrodes. The U-Net was agnostic to any information regarding the electrode model (i.e. the number of contacts or inter-contact spacing). A higher threshold (ex. 3000 or 3500 Hu) could improve the number of contacts localized by the image processing approach particularly for vertically implanted electrodes, but also risks reducing the number of true positives. Combining the outputs from a range of thresholds could reduce the missed contacts by this approach and should be explored for a more accurate comparison.

The application of the U-Net in this thesis still required a degree of manual intervention to label the predicted coordinates, in the form of the post-implantation electrode trajectory. However, given that the placement of coordinates at the entry and target is necessary as part of the surgical planning of each trajectory in an SEEG case, this data can reasonably be acquired when working with the neurosurgeons performing the implantations. These points also serve the dual purpose of labeling individual electrodes by the target region and are used in the referencing scheme to analyze recordings (Mercier et al., 2022). Other papers propose alternatives in the form of clustering of points into electrodes (Aslam et al., 2024; Blenkmann et al., 2024; B. Zhao et al., 2023). Unsupervised clustering provides a potential solution but requires the number of electrodes implanted as a prior and the eventual labelling of each individual electrode cluster by target region. Additionally, an improper initialization of centroids can lead to sub-optimal results or missed electrodes. For these reasons, it could be argued that the post-implantation trajectories represent the most intuitive method of labelling SEEG contacts. Future work could compare the localization accuracy of a fully automated approach through unsupervised clustering.

4.1 Limitations

The thesis presents a potential approach toward automated SEEG contact localization, but several challenges and limitations remain. A summary of these can be found below:

1. Limited to post-operative CT for SEEG contact localization

The U-Net model presented in this thesis was trained on post-operative CT scans for contact localizations, limiting the use for other centers where a post-operative MRI is

acquired to determine contact positions. Acquiring datasets with post-operative MRI would allow for the training of a multi-modal model optimized to work with said images.

2. Euclidean distance as an error metric

While the Euclidean distance provides a scalar value representing the distance between two points in 3-D space, it does not indicate the directionality of that distance. Within the thesis, the Euclidean distance was chosen as the error metric to compare the accuracy of both models (SEEGA and the U-Net model). However, given that this is a scalar quantity this metric does not represent the directionality of the error. Instead, representing these errors within a normalized point cloud would help to tease out this directionality for both the U-Net and SEEGA methods. Future work will look to explore this method.

3. Selection of patch size

As outlined in Chapter 2, the model was trained on 3D isotropic patches 64 and 96 voxels, a size that includes the contacts along an electrode as well as multiple electrode trajectories. Larger patch sizes provide more spatial context to the model and are known to improve performance for deep learning models, while requiring a larger number of computational resources to train. In the context of contact segmentation, a smaller patch size (16 or 32 voxels) would only include a single contact; it is likely that a model would overpredict the number of contacts, when considering the other high-intensity artifacts found in the image (ex. bolts, wire artefacts). A larger patch size would likely help the model differentiate between the individual contact positions and the other high-intensity artifacts or noise found in the image, such as the bolts or wires. However, the specific selection at 64 and 96 voxels (compared to 80 voxels or an increased size) is arbitrary. Future work should further explore and look to quantify the effect of more patch sizes on model performance within the validation set.

4. Thorough comparison of hyperparameters for the current U-Net model

The number of filters and depth of the current model was chosen to represent the traditional U-Net architecture. Since the segmentation task is relatively simple, a generic model was thought to be the best first approach. However, changing these specific hyperparameters, along with larger patch sizes, or different batch sizes could drive towards improved performance. Future work will look to quantitatively compare the assessment of different hyper-parameters to determine the best approach.

5. Exploring the effects of different loss functions

While the Combo loss function of an equally weighted Dice and Cross Entropy losses is effective for most segmentation tasks, the development of a loss function tailored to the specific task of contact localization (taking close care to more difficult use cases such as bending or vertically implanted electrodes) could improve performance. Additionally, different weightings of Dice and Cross-Entropy should be explored to determine the optimal combination. Examples in the literature include “soft” fine-tuning where the full weight is given to the Cross-Entropy Loss, and the weight of the Dice Loss is linearly increased at each epoch or “hard” fine-tuning where Dice Loss is only minimized for the last 10% of the total training epochs (Galdran et al., 2023). Future work could quantitatively compare the effect of the weighting of Dice and Cross-Entropy in the Combo loss, different loss functions, as well as work towards the development of a custom loss function for contact localization.

6. Testing different deep learning architectures

There is good reason to believe that a well-configured U-Net should outperform other deep learning or CNN models (Isensee et al., 2021). However, variants of the U-Net such as a cascaded U-Net and a more comprehensive look at model hyperparameters could improve performance. Recent work in the space of vision transformers should also be explored to assess any changes to performance (Shamshad et al., 2023).

7. Validation in external test sets (i.e. different electrode models and data from different groups) and model generalizability

The current model was trained on patients implanted with AdTech electrodes that had 10 contacts spaced 3, 4, 5, or 6 mm apart. As previously mentioned, other manufacturers produce electrodes that contain 5-18 electrodes and have smaller contact spacings of 2.5 mm or incorporate macro and micro electrodes. Additionally, the generalizability of the U-Net model has not been tested to data from other centers and different CT acquisition parameters. Future work will look to assess the generalizability of the current model to external datasets, along with electrodes with different inter-contact spacing and varying number of contacts.

8. Addressing instances where model fails to predict a contact

While achieving relatively high accuracy, the number of false negatives is still a concern. To address this issue, the outputs of the model could be provided additional information regarding the number of contacts present in the electrode or the inter-contact spacing to interpolate the remaining points. Previous discussion around the selection of hyperparameters, different model architectures, or loss functions could also lead to improved segmentations.

Chapter 5

5 Conclusion

The use of a U-Net and probabilistic non-max suppression algorithm localized most contact positions with a similar accuracy to the current tool, SEEG Assistant (SEEGA), employed in the clinic. The model did not require the manual correction of points from the post-implantation trajectory or information regarding the contact spacing and number of contacts present in the electrode. This provides evidence suggesting that deep learning models such as the 3-D U-Net can address some of the inherent challenges associated with contact localization and semi-automated algorithms. Future work will assess the performance of the model on different electrode models, explore alternative methods of representing the accuracy, and optimize performance, all while building this infrastructure in an open and reproducible way.

Bibliography

Arnulfo, G., Narizzano, M., Cardinale, F., Fato, M. M., & Palva, J. M. (2015). Automatic segmentation of deep intracerebral electrodes in computed tomography scans. *BMC Bioinformatics*, *16*(1), 99. <https://doi.org/10.1186/s12859-015-0511-6>

Arya, R., Mangano, F. T., Horn, P. S., Holland, K. D., Rose, D. F., & Glauser, T. A. (2013). Adverse events related to extraoperative invasive EEG monitoring with subdural grid electrodes: A systematic review and meta-analysis. *Epilepsia*, *54*(5), 828–839. <https://doi.org/10.1111/epi.12073>

Aslam, S., Chen, Q., & Zhang, J. (2024). DELRecon: Depth Electrode Reconstruction Toolbox for Stereo-EEG. In X. Xie, I. Styles, G. Powathil, & M. Ceccarelli (Eds.), *Artificial Intelligence in Healthcare* (Vol. 14976, pp. 135–148). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-67285-9_10

Aston, J., Thompson, W., Lindquist, M., & Ombao, H. (2017). *Handbook of neuroimaging data analysis*. CRC Press, Taylor & Francis Group.

Avants, B., Epstein, C., Grossman, M., & Gee, J. (2008). Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. *Medical Image Analysis*, *12*(1), 26–41. <https://doi.org/10.1016/j.media.2007.06.004>

Bartolomei, F., Lagarde, S., Wendling, F., McGonigal, A., Jirsa, V., Guye, M., & Bénar, C. (2017). Defining epileptogenic networks: Contribution of SEEG and signal analysis. *Epilepsia*, *58*(7), 1131–1147. <https://doi.org/10.1111/epi.13791>

Blenkman, A. O., Leske, S. L., Llorens, A., Lin, J. J., Chang, E. F., Brunner, P., Schalk, G., Ivanovic, J., Larsson, P. G., Knight, R. T., Endestad, T., & Solbakk, A.-K. (2024). Anatomical registration of intracranial electrodes. Robust model-based localization and deformable smooth brain-shift compensation methods. *Journal of Neuroscience Methods*, *404*, 110056. <https://doi.org/10.1016/j.jneumeth.2024.110056>

Brett, M., Johnsrude, I. S., & Owen, A. M. (2002). The problem of functional localization in the human brain. *Nature Reviews Neuroscience*, 3(3), 243–249.

<https://doi.org/10.1038/nrn756>

Buzsáki, G., Anastassiou, C. A., & Koch, C. (2012). The origin of extracellular fields and currents—EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, 13(6), 407–420.

<https://doi.org/10.1038/nrn3241>

Cai, F., Wang, K., Zhao, T., Wang, H., Zhou, W., & Hong, B. (2022). BrainQuake: An Open-Source Python Toolbox for the Stereoelectroencephalography Spatiotemporal Analysis. *Frontiers in Neuroinformatics*, 15, 773890.

<https://doi.org/10.3389/fninf.2021.773890>

Cardinale, F., Casaceli, G., Raneri, F., Miller, J., & Lo Russo, G. (2016). Implantation of Stereoelectroencephalography Electrodes: A Systematic Review. *Journal of Clinical Neurophysiology*, 33(6), 490–502. <https://doi.org/10.1097/WNP.0000000000000249>

Cardinale, F., Cossu, M., Castana, L., Casaceli, G., Schiariti, M. P., Miserocchi, A., Fuschillo, D., Moscato, A., Caborni, C., Arnulfo, G., & Russo, G. L. (2013).

Stereoelectroencephalography. *Neurosurgery*, 72(3), 353–366.

<https://doi.org/10.1227/NEU.0b013e31827d1161>

Cardinale, F., Rizzi, M., Vignati, E., Cossu, M., Castana, L., d’Orio, P., Revay, M., Costanza, M. D., Tassi, L., Mai, R., Sartori, I., Nobili, L., Gozzo, F., Pelliccia, V., Mariani, V., Lo Russo, G., & Francione, S. (2019). Stereoelectroencephalography: Retrospective analysis of 742 procedures in a single centre. *Brain*, 142(9), 2688–2704.

<https://doi.org/10.1093/brain/awz196>

Cardoso, M. J., Li, W., Brown, R., Ma, N., Kerfoot, E., Wang, Y., Murrey, B., Myronenko, A., Zhao, C., Yang, D., Nath, V., He, Y., Xu, Z., Hatamizadeh, A., Myronenko, A., Zhu, W., Liu, Y., Zheng, M., Tang, Y., ... Feng, A. (2022). *MONAI: An open-source framework for deep learning in healthcare* (arXiv:2211.02701). arXiv.

<http://arxiv.org/abs/2211.02701>

Centracchio, J., Sarno, A., Esposito, D., Andreozzi, E., Pavone, L., Di Gennaro, G., Bartolo, M., Esposito, V., Morace, R., Casciato, S., & Bifulco, P. (2021). Efficient automated localization of ECoG electrodes in CT images via shape analysis. *International Journal of Computer Assisted Radiology and Surgery*, *16*(4), 543–554. <https://doi.org/10.1007/s11548-021-02325-0>

Charlier, F., Weber, M., Izak, D., Harkin, E., Magnus, M., Lalli, J., Fresnais, L., Chan, M., Markov, N., Amsalem, O., Proost, S., Krasoulis, A., getzze, & Repplinger, S. (2022). *Statannotations* (Version v0.6) [Computer software]. Zenodo. <https://doi.org/10.5281/zenodo.7213391>

Chen, C. C., Chapman, P. H., Kooy, H., & Loeffler, J. S. (2008). Neuroimaging in Radiosurgery Treatment Planning and Follow-up Evaluation. In L. S. Chin & W. F. Regine (Eds.), *Principles and Practice of Stereotactic Radiosurgery* (pp. 9–23). Springer New York. https://doi.org/10.1007/978-0-387-71070-9_2

Chollet, F. (2017). *Deep Learning with Python*. Manning. <https://books.google.ca/books?id=wzozEAAAQBAJ>

Çiçek, Ö., Abdulkadir, A., Lienkamp, S. S., Brox, T., & Ronneberger, O. (2016). *3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation* (arXiv:1606.06650). arXiv. <http://arxiv.org/abs/1606.06650>

Cossu, M., Cardinale, F., Castana, L., Citterio, A., Francione, S., Tassi, L., Benabid, A. L., & Lo Russo, G. (2005). Stereoelectroencephalography in the Presurgical Evaluation of Focal Epilepsy: A Retrospective Analysis of 215 Procedures. *Neurosurgery*, *57*(4), 706–718. <https://doi.org/10.1227/01.NEU.0000176656.33523.1e>

Darcey, T. M., & Roberts, D. W. (2010). Technique for the localization of intracranially implanted electrodes: Clinical article. *Journal of Neurosurgery*, *113*(6), 1182–1185. <https://doi.org/10.3171/2009.12.JNS091678>

Davis, T. S., Caston, R. M., Philip, B., Charlebois, C. M., Anderson, D. N., Weaver, K. E., Smith, E. H., & Rolston, J. D. (2021). LeGUI: A Fast and Accurate Graphical User

- Interface for Automated Detection and Anatomical Localization of Intracranial Electrodes. *Frontiers in Neuroscience*, 15, 769872.
<https://doi.org/10.3389/fnins.2021.769872>
- Dice, L. R. (1945). Measures of the Amount of Ecologic Association Between Species. *Ecology*, 26(3), 297–302. <https://doi.org/10.2307/1932409>
- Ervin, B., Rozhkov, L., Buroker, J., Leach, J. L., Mangano, F. T., Greiner, H. M., Holland, K. D., & Arya, R. (2021). Fast Automated Stereo-EEG Electrode Contact Identification and Labeling Ensemble. *Stereotactic and Functional Neurosurgery*, 99(5), 393–404. <https://doi.org/10.1159/000515090>
- Evans, A. C., Janke, A. L., Collins, D. L., & Baillet, S. (2012). Brain templates and atlases. *NeuroImage*, 62(2), 911–922. <https://doi.org/10.1016/j.neuroimage.2012.01.024>
- Fedorov, A., Beichel, R., Kalpathy-Cramer, J., Finet, J., Fillion-Robin, J.-C., Pujol, S., Bauer, C., Jennings, D., Fennessy, F., Sonka, M., Buatti, J., Aylward, S., Miller, J. V., Pieper, S., & Kikinis, R. (2012). 3D Slicer as an image computing platform for the Quantitative Imaging Network. *Magnetic Resonance Imaging*, 30(9), 1323–1341. <https://doi.org/10.1016/j.mri.2012.05.001>
- Fischl, B. (2012). FreeSurfer. *Neuroimage*, 62(2), 774–781.
- Fisher, R. S., Scharfman, H. E., & deCurtis, M. (2014). How Can We Identify Ictal and Interictal Abnormal Activity? In H. E. Scharfman & P. S. Buckmaster (Eds.), *Issues in Clinical Epileptology: A View from the Bench* (Vol. 813, pp. 3–23). Springer Netherlands. https://doi.org/10.1007/978-94-017-8914-1_1
- Fitzpatrick, J., Hill, D., & Maurer, C. (2000). Image Registration. In *Handbook of Medical Imaging, Volume 2. Medical Image Processing and Analysis* (pp. 447–513). SPIE. <https://doi.org/10.1117/3.831079.ch8>
- Fonov, V., Evans, A., McKinstry, R., Almlí, C., & Collins, D. (2009). Unbiased nonlinear average age-appropriate brain templates from birth to adulthood. *Organization*

for *Human Brain Mapping 2009 Annual Meeting*, 47, S102.

[https://doi.org/10.1016/S1053-8119\(09\)70884-5](https://doi.org/10.1016/S1053-8119(09)70884-5)

Frauscher, B., Mansilla, D., Abdallah, C., Astner-Rohracher, A., Beniczky, S., Brazdil, M., Gnatkovsky, V., Jacobs, J., Kalamangalam, G., Perucca, P., Ryvlin, P., Schuele, S., Tao, J., Wang, Y., Zijlmans, M., & McGonigal, A. (2024). Learn how to interpret and use intracranial EEG findings. *Epileptic Disorders*, 26(1), 1–59.

<https://doi.org/10.1002/epd2.20190>

Friston, K. J. (Karl J.). (2007). *Statistical parametric mapping: The analysis of functional brain images* (1st ed.). Elsevier / Academic Press.

Furlanetti, L., Hasegawa, H., Oviedova, A., Raslan, A., Samuel, M., Selway, R., & Ashkan, K. (2021). O-Arm Stereotactic Imaging in Deep Brain Stimulation Surgery Workflow: A Utility and Cost-Effectiveness Analysis. *Stereotactic and Functional Neurosurgery*, 99(2), 93–106. <https://doi.org/10.1159/000510344>

Galdran, A., Carneiro, G., & Ballester, M. A. G. (2023). On the Optimal Combination of Cross-Entropy and Soft Dice Losses for Lesion Segmentation with Out-of-Distribution Robustness. In M. H. Yap, C. Kendrick, & B. Cassidy (Eds.), *Diabetic Foot Ulcers Grand Challenge* (pp. 40–51). Springer International Publishing.

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.

Gorgolewski, K. J., Auer, T., Calhoun, V. D., Craddock, R. C., Das, S., Duff, E. P., Flandin, G., Ghosh, S. S., Glatard, T., Halchenko, Y. O., Handwerker, D. A., Hanke, M., Keator, D., Li, X., Michael, Z., Maumet, C., Nichols, B. N., Nichols, T. E., Pellman, J., ... Poldrack, R. A. (2016). The brain imaging data structure, a format for organizing and describing outputs of neuroimaging experiments. *Scientific Data*, 3(1), 160044.

<https://doi.org/10.1038/sdata.2016.44>

Gorgulho, A. A., Ishida, W., & De Salles, A. A. F. (2009). General Imaging Modalities: Basic Principles. In A. M. Lozano, P. L. Gildenberg, & R. R. Tasker (Eds.), *Textbook of*

Stereotactic and Functional Neurosurgery (pp. 249–267). Springer Berlin Heidelberg.
https://doi.org/10.1007/978-3-540-69960-6_17

Granados, A., Han, Y., Lucena, O., Vakharia, V., Rodionov, R., Vos, S. B., Miserocchi, A., McEvoy, A. W., Duncan, J. S., Sparks, R., & Ourselin, S. (2021). Patient-specific prediction of SEEG electrode bending for stereotactic neurosurgical planning. *International Journal of Computer Assisted Radiology and Surgery*, *16*(5), 789–798.
<https://doi.org/10.1007/s11548-021-02347-8>

Granados, A., Vakharia, V., Rodionov, R., Schweiger, M., Vos, S. B., O’Keeffe, A. G., Li, K., Wu, C., Miserocchi, A., McEvoy, A. W., Clarkson, M. J., Duncan, J. S., Sparks, R., & Ourselin, S. (2018). Automatic segmentation of stereoelectroencephalography (SEEG) electrodes post-implantation considering bending. *International Journal of Computer Assisted Radiology and Surgery*, *13*(6), 935–946.
<https://doi.org/10.1007/s11548-018-1740-8>

Groppe, D. M., Bickel, S., Dykstra, A. R., Wang, X., Mégevand, P., Mercier, M. R., Lado, F. A., Mehta, A. D., & Honey, C. J. (2017). iELVis: An open source MATLAB toolbox for localizing and visualizing human intracranial electrode data. *Journal of Neuroscience Methods*, *281*, 40–48. <https://doi.org/10.1016/j.jneumeth.2017.01.022>

Hall, J. A., & Khoo, H. M. (2018). Robotic-Assisted and Image-Guided MRI-Compatible Stereoelectroencephalography. *Canadian Journal of Neurological Sciences / Journal Canadien Des Sciences Neurologiques*, *45*(1), 35–43.
<https://doi.org/10.1017/cjn.2017.240>

Hamwood, J., Alonso-Caneiro, D., Read, S. A., Vincent, S. J., & Collins, M. J. (2018). Effect of patch size and network architecture on a convolutional neural network approach for automatic segmentation of OCT retinal layers. *Biomedical Optics Express*, *9*(7), 3049.
<https://doi.org/10.1364/BOE.9.003049>

Hartov, A., & Roberts, D. W. (2009). Accuracy in Stereotactic and Image Guidance. In A. M. Lozano, P. L. Gildenberg, & R. R. Tasker (Eds.), *Textbook of Stereotactic and*

Functional Neurosurgery (pp. 443–452). Springer Berlin Heidelberg.
https://doi.org/10.1007/978-3-540-69960-6_28

Holdgraf, C., Appelhoff, S., Bickel, S., Bouchard, K., D’Ambrosio, S., David, O., Devinsky, O., Dichter, B., Flinker, A., Foster, B. L., Gorgolewski, K. J., Groen, I., Groppe, D., Gunduz, A., Hamilton, L., Honey, C. J., Jas, M., Knight, R., Lachaux, J.-P., ... Hermes, D. (2019). iEEG-BIDS, extending the Brain Imaging Data Structure specification to human intracranial electrophysiology. *Scientific Data*, 6(1), 102.
<https://doi.org/10.1038/s41597-019-0105-7>

Horn, A., Kühn, A. A., Merkl, A., Shih, L., Alterman, R., & Fox, M. (2017). Probabilistic conversion of neurosurgical DBS electrode coordinates into MNI space. *NeuroImage*, 150, 395–404. <https://doi.org/10.1016/j.neuroimage.2017.02.004>

Horsley, V., & Clarke, R. H. (1908). THE STRUCTURE AND FUNCTIONS OF THE CEREBELLUM EXAMINED BY A NEW METHOD. *Brain*, 31(1), 45–124.
<https://doi.org/10.1093/brain/31.1.45>

Hounsfield, G. N. (1973). Computerized transverse axial scanning (tomography): Part 1. Description of system. *British Journal of Radiology*, 46(552), 1016–1022.
<https://doi.org/10.1259/0007-1285-46-552-1016>

Ioffe, S., & Szegedy, C. (2015). *Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift* (arXiv:1502.03167). arXiv.
<http://arxiv.org/abs/1502.03167>

Isensee, F., Jaeger, P. F., Kohl, S. A. A., Petersen, J., & Maier-Hein, K. H. (2021). nnU-Net: A self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods*, 18(2), 203–211. <https://doi.org/10.1038/s41592-020-01008-z>

Jaccard, P. (1912). THE DISTRIBUTION OF THE FLORA IN THE ALPINE ZONE. *New Phytologist*, 11(2), 37–50. <https://doi.org/10.1111/j.1469-8137.1912.tb05611.x>

- Jadon, S. (2020). A survey of loss functions for semantic segmentation. *2020 IEEE Conference on Computational Intelligence in Bioinformatics and Computational Biology (CIBCB)*, 1–7. <https://doi.org/10.1109/CIBCB48159.2020.9277638>
- Janca, R., Tomasek, M., Kalina, A., Marusic, P., Krsek, P., & Lesko, R. (2023). Automated Identification of Stereoelectroencephalography Contacts and Measurement of Factors Influencing Accuracy of Frame Stereotaxy. *IEEE Journal of Biomedical and Health Informatics*, 1–10. <https://doi.org/10.1109/JBHI.2023.3271857>
- Jobst, B. C., Bartolomei, F., Diehl, B., Frauscher, B., Kahane, P., Minotti, L., Sharan, A., Tardy, N., Worrell, G., & Gotman, J. (2020). Intracranial EEG in the 21st Century. *Epilepsy Currents*, 20(4), 180–188. <https://doi.org/10.1177/1535759720934852>
- Joswig, H., Benson, C. M., Parrent, A. G., MacDougall, K. W., & Steven, D. A. (2018). Operative Nuances of Stereotactic Leksell Frame-Based Depth Electrode Implantation. *Operative Neurosurgery*, 15(3), 292–295. <https://doi.org/10.1093/ons/opx245>
- Joswig, H., Lau, J. C., Abdallat, M., Parrent, A. G., MacDougall, K. W., McLachlan, R. S., Burneo, J. G., & Steven, D. A. (2020). Stereoelectroencephalography Versus Subdural Strip Electrode Implantations: Feasibility, Complications, and Outcomes in 500 Intracranial Monitoring Cases for Drug-Resistant Epilepsy. *Neurosurgery*, 87(1), E23–E30. <https://doi.org/10.1093/neuros/nyaa112>
- Joswig, H., Steven, D. A., Parrent, A. G., MacDougall, K. W., Mirsattari, S. M., McLachlan, R. S., Diosy, D. C., & Burneo, J. G. (2018). Intracranial Electroencephalographic Monitoring: From Subdural to Depth Electrodes. *Canadian Journal of Neurological Sciences / Journal Canadien Des Sciences Neurologiques*, 45(3), 336–338. <https://doi.org/10.1017/cjn.2018.4>
- Kahane, P., Landré, E., Minotti, L., Francione, S., & Ryvlin, P. (2006). The Bancaud and Talairach view on the epileptogenic zone: A working hypothesis. *Epileptic Disorders*, 8(S2). <https://doi.org/10.1684/j.1950-6945.2006.tb00206.x>

Kall, B. A. (2009). Image Reconstruction and Fusion. In A. M. Lozano, P. L. Gildenberg, & R. R. Tasker (Eds.), *Textbook of Stereotactic and Functional Neurosurgery* (pp. 335–343). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-540-69960-6_24

Khan, A. R., Wang, L., & Beg, M. F. (2008). FreeSurfer-initiated fully-automated subcortical brain segmentation in MRI using Large Deformation Diffeomorphic Metric Mapping. *NeuroImage*, *41*(3), 735–746.
<https://doi.org/10.1016/j.neuroimage.2008.03.024>

Kingma, D. P., & Ba, J. (2017). *Adam: A Method for Stochastic Optimization* (arXiv:1412.6980). arXiv. <http://arxiv.org/abs/1412.6980>

Kreinter, H., Espino, P. H., Mejía, S., Alorabi, K., Gilmore, G., Burneo, J. G., Steven, D. A., MacDougall, K. W., Jones, M.-L., Pellegrino, G., Diosy, D., Mirsattari, S. M., Lau, J., & Suller Marti, A. (2024). Disrupting the epileptogenic network with stereoelectroencephalography-guided radiofrequency thermocoagulation. *Epilepsia*, *65*(7), e113–e118. <https://doi.org/10.1111/epi.18005>

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2017). ImageNet classification with deep convolutional neural networks. *Communications of the ACM*, *60*(6), 84–90.
<https://doi.org/10.1145/3065386>

Lau, J. C. (2022). Stereotactic spaces. In *Connectomic Deep Brain Stimulation* (pp. 169–184). Elsevier. <https://doi.org/10.1016/B978-0-12-821861-7.00023-3>

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, *521*(7553), 436–444. <https://doi.org/10.1038/nature14539>

Lucas, A., Scheid, B. H., Pattnaik, A. R., Gallagher, R., Mojena, M., Tranquille, A., Prager, B., Gleichgerricht, E., Gong, R., Litt, B., Davis, K. A., Das, S., Stein, J. M., & Sinha, N. (2024). iEEG-recon: A fast and scalable pipeline for accurate reconstruction of intracranial electrodes and implantable devices. *Epilepsia*, *65*(3), 817–829.
<https://doi.org/10.1111/epi.17863>

- Ma, J., Chen, J., Ng, M., Huang, R., Li, Y., Li, C., Yang, X., & Martel, A. L. (2021). Loss odyssey in medical image segmentation. *Medical Image Analysis*, *71*, 102035. <https://doi.org/10.1016/j.media.2021.102035>
- Medina Villalon, S., Paz, R., Roehri, N., Lagarde, S., Pizzo, F., Colombet, B., Bartolomei, F., Carron, R., & Bénar, C.-G. (2018). EpiTools, A software suite for presurgical brain mapping in epilepsy: Intracerebral EEG. *Journal of Neuroscience Methods*, *303*, 7–15. <https://doi.org/10.1016/j.jneumeth.2018.03.018>
- Mercier, M. R., Dubarry, A.-S., Tadel, F., Avanzini, P., Axmacher, N., Cellier, D., Vecchio, M. D., Hamilton, L. S., Hermes, D., Kahana, M. J., Knight, R. T., Llorens, A., Megevand, P., Melloni, L., Miller, K. J., Piai, V., Puce, A., Ramsey, N. F., Schwiedrzik, C. M., ... Oostenveld, R. (2022). Advances in human intracranial electroencephalography research, guidelines and good practices. *NeuroImage*, *260*, 119438. <https://doi.org/10.1016/j.neuroimage.2022.119438>
- Modat, M., Cash, D. M., Daga, P., Winston, G. P., Duncan, J. S., & Ourselin, S. (2014). Global image registration using a symmetric block-matching approach. *Journal of Medical Imaging*, *1*(2), 024003. <https://doi.org/10.1117/1.JMI.1.2.024003>
- Modat, M., Ridgway, G. R., Taylor, Z. A., Lehmann, M., Barnes, J., Hawkes, D. J., Fox, N. C., & Ourselin, S. (2010). Fast free-form deformation using graphics processing units. *Computer Methods and Programs in Biomedicine*, *98*(3), 278–284. <https://doi.org/10.1016/j.cmpb.2009.09.002>
- Monney, J., Dallaire, S. E., Stoutah, L., Fanda, L., & Mégevand, P. (2024). Voxeloc: A time-saving graphical user interface for localizing and visualizing stereo-EEG electrodes. *Journal of Neuroscience Methods*, 110154. <https://doi.org/10.1016/j.jneumeth.2024.110154>
- Narizzano, M., Arnulfo, G., Ricci, S., Toselli, B., Tisdall, M., Canessa, A., Fato, M. M., & Cardinale, F. (2017). SEEG assistant: A 3DSlicer extension to support epilepsy surgery. *BMC Bioinformatics*, *18*(1), 124. <https://doi.org/10.1186/s12859-017-1545-8>

- O'Shea, K., & Nash, R. (2015). *An Introduction to Convolutional Neural Networks* (arXiv:1511.08458). arXiv. <http://arxiv.org/abs/1511.08458>
- Pantovic, A., Ollivier, I., & Essert, C. (2022). Hybrid U-Net for segmentation of SEEG electrodes on post-operative CT scans. *Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization*, 1–7.
<https://doi.org/10.1080/21681163.2022.2152376>
- Paredes-Aragon, E., AlKhaldi, N. A., Ballesteros-Herrera, D., & Mirsattari, S. M. (2022). Stereo-Encephalographic Presurgical Evaluation of Temporal Lobe Epilepsy: An Evolving Science. *Frontiers in Neurology*, 13, 867458.
<https://doi.org/10.3389/fneur.2022.867458>
- Parrent, A. G. (2009). History of Stereotactic and Functional Neurosurgery in Canada. In A. M. Lozano, P. L. Gildenberg, & R. R. Tasker (Eds.), *Textbook of Stereotactic and Functional Neurosurgery* (pp. 113–123). Springer Berlin Heidelberg.
https://doi.org/10.1007/978-3-540-69960-6_10
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., & Cournapeau, D. (2011). Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12(85), 2825–2830.
- Penfield, W., & Jasper, H. (1954). *Epilepsy and the functional anatomy of the human brain*.
- Perera, T. (2022). Chapter 5—DBS imaging methods I: Preprocessing. In A. Horn (Ed.), *Connectomic Deep Brain Stimulation* (pp. 111–125). Academic Press.
<https://doi.org/10.1016/B978-0-12-821861-7.00006-3>
- Peters, T. M. (2001). Image-guided surgery: From X-rays to Virtual Reality. *Computer Methods in Biomechanics and Biomedical Engineering*, 4(1), 27–57.
<https://doi.org/10.1080/10255840008907997>

- Qin, C., Tan, Z., Pan, Y., Li, Y., Wang, L., Ren, L., Zhou, W., & Wang, L. (2017). Automatic and Precise Localization and Cortical Labeling of Subdural and Depth Intracranial Electrodes. *Frontiers in Neuroinformatics, 11*.
<https://doi.org/10.3389/fninf.2017.00010>
- Rockhill, A., Larson, E., Stedelin, B., Mantovani, A., Raslan, A., Gramfort, A., & Swann, N. (2022). Intracranial Electrode Location and Analysis in MNE-Python. *Journal of Open Source Software, 7*(70), 3897. <https://doi.org/10.21105/joss.03897>
- Ronneberger, O., Fischer, P., & Brox, T. (2015). *U-Net: Convolutional Networks for Biomedical Image Segmentation* (arXiv:1505.04597). arXiv.
<http://arxiv.org/abs/1505.04597>
- Seabold, S., & Perktold, J. (2010). Statsmodels: Econometric and Statistical Modeling with Python. In S. van der Walt & J. Millman (Eds.), *Proceedings of the 9th Python in Science Conference* (pp. 92–96). <https://doi.org/10.25080/Majora-92bf1922-011>
- Shamshad, F., Khan, S., Zamir, S. W., Khan, M. H., Hayat, M., Khan, F. S., & Fu, H. (2023). Transformers in medical imaging: A survey. *Medical Image Analysis, 88*, 102802. <https://doi.org/10.1016/j.media.2023.102802>
- Siddique, N., Paheding, S., Elkin, C. P., & Devabhaktuni, V. (2021). U-Net and Its Variants for Medical Image Segmentation: A Review of Theory and Applications. *IEEE Access, 9*, 82031–82057. <https://doi.org/10.1109/ACCESS.2021.3086020>
- Spiegel, E. A., Wycis, H. T., Marks, M., & Lee, A. J. (1947). Stereotaxic Apparatus for Operations on the Human Brain. *Science, 106*(2754), 349–350. JSTOR.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research, 15*(56), 1929–1958.

- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., & Rabinovich, A. (2015). Going Deeper with Convolutions. *Computer Vision and Pattern Recognition (CVPR)*. <http://arxiv.org/abs/1409.4842>
- Taghanaki, S. A., Zheng, Y., Zhou, S. K., Georgescu, B., Sharma, P., Xu, D., Comaniciu, D., & Hamarneh, G. (2021). *Combo Loss: Handling Input and Output Imbalance in Multi-Organ Segmentation* (arXiv:1805.02798). arXiv. <http://arxiv.org/abs/1805.02798>
- Talairach, J., & Bancaud, J. (1966). Lesion, “irritative” zone and epileptogenic focus. *Confinia Neurologica*, 27(1), 91–94. <https://doi.org/10.1159/000103937>
- TALAIRACH, J., BANCAUD, J., BONIS, A., SZIKLA, G., & TOURNOUX, P. (1962). Functional stereotaxic exploration of epilepsy. *Confinia Neurologica*, 22, 328–331. <https://doi.org/10.1159/000104378>
- TALAIRACH, J., DE ALJURIAGUERRA, J., & DAVID, M. (1952). Etudes stéréotaxiques des structures encéphaliques profondes chez l’homme; technique; intérêt physiopathologique et thérapeutique. *La Presse Médicale*, 60(28), 605–609. Scopus.
- Talairach, J. (Jean). (1967). *Atlas d’anatomie stéréotaxique du télencéphale: Études anatomo-radiologiques*. Masson et Cie. <https://cir.nii.ac.jp/crid/1130282272570629120>
- Talairach, J., & Tournoux, P. (1988). *Co-planar Stereotaxic Atlas of the Human Brain: 3-dimensional Proportional System: An Approach to Cerebral Imaging*. G. Thieme. <https://books.google.ca/books?id=pYFiQgAACAAJ>
- Vakharia, V. N., Sparks, R., O’Keeffe, A. G., Rodionov, R., Miserocchi, A., McEvoy, A., Ourselin, S., & Duncan, J. (2017). Accuracy of intracranial electrode placement for stereoelectroencephalography: A systematic review and meta-analysis. *Epilepsia*, 58(6), 921–932. <https://doi.org/10.1111/epi.13713>
- van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., Boulogne, F., Warner, J. D., Yager, N., Guillaud, E., Yu, T., & contributors, the scikit-image. (2014). scikit-image: Image processing in Python. *PeerJ*, 2, e453. <https://doi.org/10.7717/peerj.453>

van Rooijen, B. D., Backes, W. H., Schijns, O. E. M. G., Colon, A., & Hofman, P. A. M. (2013). Brain Imaging in Chronic Epilepsy Patients After Depth Electrode (Stereoelectroencephalography) Implantation. *Neurosurgery*, *73*(3), 543–549. <https://doi.org/10.1227/01.neu.0000431478.79536.68>

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., ... SciPy 1.0 Contributors. (2020). SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, *17*, 261–272. <https://doi.org/10.1038/s41592-019-0686-2>

Vlasov, V., Bofferding, M., Marx, L., Zhang, C., Goncalves, J., Husch, A., & Hertel, F. (2021). Automated Deep Learning-based Segmentation of Brain, SEEG and DBS Electrodes on CT Images. In C. Palm, T. M. Deserno, H. Handels, A. Maier, K. Maier-Hein, & T. Tolxdorff (Eds.), *Bildverarbeitung für die Medizin 2021* (pp. 92–97). Springer Fachmedien Wiesbaden. https://doi.org/10.1007/978-3-658-33198-6_22

Wang, Z., Magnotti, J. F., Zhang, X., & Beauchamp, M. S. (2023). YAEL: Your Advanced Electrode Localizer. *Eneuro*, *10*(10), ENEURO.0328-23.2023. <https://doi.org/10.1523/ENEURO.0328-23.2023>

Wiebe, S., Blume, W. T., Girvin, J. P., & Eliasziw, M. (2001). A Randomized, Controlled Trial of Surgery for Temporal-Lobe Epilepsy. *New England Journal of Medicine*, *345*(5), 311–318. <https://doi.org/10.1056/NEJM200108023450501>

Withey, D. J., & Koles, Z. J. (2008). A Review of Medical Image Segmentation: Methods and Available Software. *International Journal of Bioelectromagnetism*, *10*(3), 125–148.

Wong, S. M., Arski, O. N., & Ibrahim, G. M. (2024). An automated algorithm for stereoelectroencephalography electrode localization and labelling. *Seizure: European Journal of Epilepsy*, *117*, 293–297. <https://doi.org/10.1016/j.seizure.2024.04.002>

Wu, S., Issa, N. P., Rose, S. L., Haider, H. A., Nordli, D. R., Towle, V. L., Warnke, P. C., & Tao, J. X. (2024). Depth versus surface: A critical review of subdural and depth electrodes in intracranial electroencephalographic studies. *Epilepsia*, *65*(7), 1868–1878. <https://doi.org/10.1111/epi.18002>

Yang, A. I., Wang, X., Doyle, W. K., Halgren, E., Carlson, C., Belcher, T. L., Cash, S. S., Devinsky, O., & Thesen, T. (2012). Localization of dense intracranial electrode arrays using magnetic resonance imaging. *NeuroImage*, *63*(1), 157–165. <https://doi.org/10.1016/j.neuroimage.2012.06.039>

Yushkevich, P. A., Piven, J., Hazlett, H. C., Smith, R. G., Ho, S., Gee, J. C., & Gerig, G. (2006). User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. *NeuroImage*, *31*(3), 1116–1128. <https://doi.org/10.1016/j.neuroimage.2006.01.015>

Zanello, M., Duriez, P., Savoureux, A.-S., Vinckier, F., Chretien, F., Gavaret, M., Gorwood, P., Gaillard, R., & Pallud, J. (2022). Jean Talairach (1911–2007). An untold story of the pioneer of stereotactic and functional neurosurgery. *Neurochirurgie*, *68*(4), 398–408. <https://doi.org/10.1016/j.neuchi.2021.12.008>

Zhao, B., Zhao, X., Hu, W., Zhang, C., Wang, X., Mo, J., Shao, X., Zhang, K., & Zhang, J. (2023). Efficient volume-based localization and automatic labeling of intracranial depth electrodes. *Annals of Translational Medicine*, *11*(6), 242–242. <https://doi.org/10.21037/atm-22-3712>

Zhao, L., & Zhang, Z. (2024). A improved pooling method for convolutional neural networks. *Scientific Reports*, *14*(1), 1589. <https://doi.org/10.1038/s41598-024-51258-6>

Zijlmans, M., Zweiphenning, W., & Van Klink, N. (2019). Changing concepts in presurgical assessment for epilepsy surgery. *Nature Reviews Neurology*, *15*(10), 594–606. <https://doi.org/10.1038/s41582-019-0224-y>

Appendices

Appendix A: Summary of localization accuracy and error for 3-D U-Net by subject, alongside the error from SEEG Assistant (SEEGA) in the test set.

								Mean Euclidean Error \pm SD (mm)	
subject	Total # of Contacts	TP	FN	FP	Precision	Recall	F1-score	U-Net	SEEGA
sub-D110	90	86	4	1	0.989	0.956	0.972	0.455 \pm 0.183	0.468 \pm 0.25
sub-D111	150	136	14	4	0.971	0.907	0.938	0.516 \pm 0.245	0.43 \pm 0.307
sub-D133	140	136	4	2	0.986	0.971	0.978	0.729 \pm 0.836	0.606 \pm 0.325
sub-D135	130	120	10	2	0.984	0.923	0.952	0.497 \pm 0.419	0.438 \pm 0.259
sub-D149	60	60	0	0	1.000	1.000	1.000	0.498 \pm 0.179	1.492 \pm 1.459
sub-D155	110	108	2	1	0.991	0.982	0.986	0.474 \pm 0.165	1.036 \pm 1.961
sub-D164	70	70	0	1	0.986	1.000	0.993	0.489 \pm 0.184	0.398 \pm 0.212

sub-P004	130	126	4	2	0.984	0.969	0.977	0.46 ± 0.184	0.357 ± 0.2
sub-P006	80	77	3	2	0.975	0.962	0.969	0.635 ± 0.676	0.429 ± 0.308
sub-P011	120	115	5	4	0.966	0.958	0.962	0.47 ± 0.354	0.455 ± 0.408
sub-P012	120	103	17	7	0.936	0.858	0.896	0.575 ± 0.266	0.675 ± 0.437
sub-P013	80	75	5	4	0.949	0.938	0.943	0.561 ± 0.272	0.52 ± 0.245
sub-P014	120	108	12	7	0.939	0.900	0.919	0.534 ± 0.247	0.6 ± 0.401
sub-P017	110	105	5	5	0.955	0.955	0.955	0.599 ± 0.366	0.477 ± 0.179
sub-P021	100	95	5	1	0.990	0.950	0.969	0.475 ± 0.221	0.448 ± 0.212
sub-P026	120	118	2	1	0.992	0.983	0.987	0.607 ± 0.661	0.482 ± 0.272
sub-P029	90	82	8	4	0.953	0.911	0.932	0.575 ± 0.56	0.388 ± 0.268
sub-P031	140	130	10	1	0.992	0.929	0.959	0.487 ± 0.382	0.346 ± 0.185

sub-P039	140	127	13	7	0.948	0.907	0.927	0.595 ± 0.241	0.435 ± 0.198
sub-P041	130	119	11	3	0.975	0.915	0.944	0.525 ± 0.397	0.358 ± 0.187
sub-P054	80	77	3	2	0.975	0.962	0.969	0.53 ± 0.202	0.422 ± 0.178
sub-P057	110	106	4	4	0.964	0.964	0.964	0.859 ± 1.365	0.453 ± 0.252
sub-P066	100	89	11	7	0.927	0.890	0.908	0.819 ± 0.725	0.976 ± 1.621
sub-P069	140	135	5	3	0.978	0.964	0.971	0.455 ± 0.157	0.412 ± 0.234
sub-P072	140	132	8	3	0.978	0.943	0.960	0.487 ± 0.283	0.533 ± 0.579
sub-P081	100	89	11	0	1.000	0.890	0.942	0.429 ± 0.164	0.426 ± 0.181
sub-P102	130	125	5	2	0.984	0.962	0.973	0.892 ± 1.359	0.486 ± 0.264
sub-P105	120	120	0	0	1.000	1.000	1.000	0.405 ± 0.142	0.398 ± 0.261
sub-P108	130	129	1	2	0.985	0.992	0.989	0.423 ± 0.146	0.616 ± 0.467

sub-P109	130	130	0	0	1.000	1.000	1.000	0.378 ± 0.153	0.476 ± 0.232
sub-P110	100	100	0	0	1.000	1.000	1.000	0.387 ± 0.132	0.588 ± 0.192
sub-P120	140	140	0	1	0.993	1.000	0.996	0.403 ± 0.112	0.605 ± 0.257
sub-P125	140	139	1	1	0.993	0.993	0.993	0.685 ± 1.033	0.526 ± 0.235
sub-P126	120	120	0	0	1.000	1.000	1.000	0.412 ± 0.143	1.23 ± 1.826
sub-P135	110	110	0	0	1.000	1.000	1.000	0.416 ± 0.154	1.35 ± 2.11

Appendix B: Ethics Approval



Date: 10 October 2024

To: Assistant Professor Ana Suller Marti

Project ID: 114420

Review Reference: 2024-114420-100360

Study Title: Intracranial EEG signal Analysis: quantification of the signal frequency and topographic localization.

Application Type: Continuing Ethics Review (CER) Form

Review Type: Delegated

Date Approval Issued: 10/Oct/2024 11:44

REB Approval Expiry Date: 01/Nov/2025

Dear Assistant Professor Ana Suller Marti,

The Western University Research Ethics Board has reviewed the application. This study, including all currently approved documents, has been re-approved until the expiry date noted above.

REB members involved in the research project do not participate in the review, discussion or decision.

Western University REB operates in compliance with, and is constituted in accordance with, the requirements of the Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans (TCPS 2); the International Conference on Harmonisation Good Clinical Practice Consolidated Guideline (ICH GCP); Part C, Division 5 of the Food and Drug Regulations; Part 4 of the Natural Health Products Regulations; Part 3 of the Medical Devices Regulations and the provisions of the Ontario Personal Health Information Protection Act (PHIPA 2004) and its applicable regulations. The REB is registered with the U.S. Department of Health & Human Services under the IRB registration number IRB 00000940.

Please do not hesitate to contact us if you have any questions.

Electronically signed by:

Joshua Hatherley, Ethics Coordinator on behalf of Dr. N. Poonai, HSREB Chair 10/Oct/2024 11:44

Reason: I am approving this document

Note: This correspondence includes an electronic signature (validation and approval via an online system that is compliant with all regulations).

Curriculum Vitae

Name: Arun Thurairajah

**Post-secondary
Education and
Degrees:** McMaster University
Hamilton, Ontario, Canada
2017-2021 BSc.

University of Western Ontario
London, Ontario, Canada
2022-2024 MSc. Neuroscience
*Collaborative Specialization in Machine Learning in Health &
Biomedical Sciences*

**Honours and
Awards:** Neuroscience Travel Award
2023

Ontario Graduate Scholarship
2023-2024

Presentations:

Thurairajah, A. et al. (2024). Segmentation of stereoelectroencephalography electrode contacts using convolutional neural networks. 2024 Clinical Neurological Sciences Research Day. University of Western Ontario.

Posters:

Thurairajah, A. et al. (2023). Automatic segmentation of stereoelectroencephalography electrodes using convolutional neural networks. 2023 Clinical Neurological Sciences Research Day. University of Western Ontario.

Thurairajah, A., et al. (2023 June). Labelmerge: a BIDS app to combine different atlas parcellations. 2023 Organization for Human Brain Mapping. Montreal, Canada