



BRNO UNIVERSITY OF TECHNOLOGY

VYSOKÉ UČENÍ TECHNICKÉ V BRNĚ

FACULTY OF ELECTRICAL ENGINEERING AND COMMUNICATION

FAKULTA ELEKTROTECHNIKY
A KOMUNIKAČNÍCH TECHNOLOGIÍ

DEPARTMENT OF BIOMEDICAL ENGINEERING

ÚSTAV BIOMEDICÍNSKÉHO INŽENÝRSTVÍ

GRAPH NEURAL NETWORKS IN EPILEPSY SURGERY

GRAPH NEURAL NETWORKS IN EPILEPSY SURGERY

MASTER'S THESIS

DIPLOMOVÁ PRÁCE

AUTHOR

AUTOR PRÁCE

Bc. Valentina Hrtoňová

SUPERVISOR

VEDOUCÍ PRÁCE

Ing. Marina Filipenská, Ph.D.

BRNO 2024

Master's Thesis

Master's study program **Biomedical Engineering and Bioinformatics**

Department of Biomedical Engineering

Student: Bc. Valentina Hrtoňová

ID: 203142

**Year of
study:** 2

Academic year: 2023/24

TITLE OF THESIS:

Graph Neural Networks in Epilepsy Surgery

INSTRUCTION:

1) Familiarize yourself with the issue of epileptogenic zone (EZ) localization from stereoelectroencephalography recordings using machine learning techniques and graph connectivity measures. 2) Familiarize yourself with graph neural networks. 3) Design and implement a deep learning model for the localization of EZ based on graph neural networks. 4) Validate and test the method on clinical data (drug-resistant patients who had SEEG and subsequent surgery). 5) Evaluate and discuss the results and contributions to the current state of the art in EZ localization.

RECOMMENDED LITERATURE:

[1] CIMBALNIK, Jan, Petr KLIMES, Vladimír SLADKY, et al. Multi-feature localization of epileptic foci from interictal, intracranial EEG. *Clinical Neurophysiology* [online]. 2019, 130(10), 1945-1953 [cit. 2023-07-03]. ISSN 13882457. Dostupné z: doi:10.1016/j.clinph.2019.07.024

[2] PANZICA, Ferruccio, Giulia VAROTTO, Fabio ROTONDI, Roberto SPREAFICO a Silvana FRANCESCHETTI. Identification of the Epileptogenic Zone from Stereo-EEG Signals: A Connectivity-Graph Theory Approach. *Frontiers in Neurology* [online]. 2013, 4 [cit. 2023-07-03]. ISSN 1664-2295. Dostupné z: doi:10.3389/fneur.2013.00175

[3] ZHOU, Jie, Ganqu CUI, Shengding HU, et al. Graph neural networks: A review of methods and applications. *AI Open* [online]. 2020, 1, 57-81 [cit. 2023-07-03]. ISSN 26666510. Dostupné z: doi:10.1016/j.aiopen.2021.01.001

**Date of project
specification:** 5.2.2024

**Deadline for
submission:** 22.5.2024

Supervisor: Ing. Marina Filipenská, Ph.D.

Consultant: Ing. Petr Klimeš, Ph.D.

prof. Ing. Valentine Provazník, Ph.D.

Chair of study program board

WARNING:

The author of the Master's Thesis claims that by creating this thesis he/she did not infringe the rights of third persons and the personal and/or property rights of third persons were not subjected to derogatory treatment. The author is fully aware of the legal consequences of an infringement of provisions as per Section 11 and following of Act No 121/2000 Coll. on copyright and rights related to copyright and on amendments to some other laws (the Copyright Act) in the wording of subsequent directives including the possible criminal consequences as resulting from provisions of Part 2, Chapter VI, Article 4 of Criminal Code 40/2009 Coll.

ABSTRACT

Successful epilepsy surgery relies on precise localization of the epileptogenic zone (EZ), yet only about 60% of patients become seizure-free post-surgery, often due to inaccurate EZ identification. This thesis presents a novel method for EZ localization using Graph Neural Networks (GNNs) to analyze interictal biomarkers, specifically interictal spikes and relative entropy. The GNN models were used to localize resected seizure-onset zone electrode contacts based on interictal stereoelectroencephalography data, validated on a clinical dataset of 37 patients from two institutions. The best-performing GNN model - Graph Attention Network - scored a median Area Under the Receiver Operating Characteristic (AUROC) of 0.971 and a median Area Under the Precision-Recall Curve (AUPRC) of 0.525 across a cohort of 19 patients with a good surgical outcome, significantly outperforming a benchmark model based on spike rates (Wilcoxon Signed Rank test, $p < 0.05$) in both metrics. This is the first application of GNNs in interictal EZ localization, and the overall superior results of the GNN models compared to a Support Vector Machine and Spike Rate model references show the potential of GNNs in enhancing the precision of epilepsy surgery planning.

KEYWORDS

graph neural networks, deep learning, signal processing, intracranial EEG, epilepsy, epileptogenic zone, seizure-onset zone, postsurgical outcome prediction

ABSTRAKT

Úspěch epileptochirurgického zákroku závisí na přesné lokalizaci epileptogenní zóny (EZ), avšak pouze 60% pacientů je po operaci bez záchvatů, což je často způsobeno nepřesnou identifikací EZ. Tato práce představuje novou metodu lokalizace EZ využívající grafové neuronové sítě (GNN) k analýze interiktálních biomarkerů - konkrétně interiktálních spiků a relativní entropie. Modely GNN byly využity pro lokalizaci kontaktů elektrod v resekované zóně vzniku záchvatu na základě dat z interiktální stereoelektroencefalografie a validovány na souboru klinických dat 37 pacientů ze dvou institucí. Nejlépe hodnocený model GNN - Graph Attention Network - dosáhl mediánu Area Under the Receiver Operating Characteristic (AUROC) 0,971 a mediánu Area Under the Precision-Recall Curve (AUPRC) 0,525 v souboru 19 pacientů s dobrým pooperačním výsledkem, přičemž v obou metrikách statisticky významně překonal referenční model založený na četnosti spiků (Wilcoxon Signed Rank test, $p < 0,05$). Jedná se o první použití GNN v interiktální lokalizaci EZ a celkově lepší výsledky modelů GNN ve srovnání s referenčními modely Support Vector Machine a modelu založeného na četnosti spiků ukazují na potenciál GNN při zvyšování přesnosti plánování epileptochirurgických zákroků.

KLÍČOVÁ SLOVA

grafové neuronové sítě, deep learning, zpracování signálů, intrakraniální EEG, epilepsie, epileptogenní zóna, zóna počátku záchvatu, predikce pooperačního výsledku

ROZŠÍŘENÝ ABSTRAKT

Diplomová práce je věnována lokalizace epileptogenní zóny (EZ) pomocí grafových neuronových sítí (GNN). Cíle zahrnují zpracování přehledu o GNN a jejich aplikacích v analýze elektroencefalografických (EEG) signálů, a následný vývoj a implementace modelů hlubokého učení pro lokalizaci EZ na základě GNN. Součástí práce je také testování a validace navržené metody na klinickém datasetu, spolu s podrobnou diskusí o přínosech k současným metodám lokalizace EZ.

Teoretická část práce začíná přehledem farmakorezistentní epilepsie, chirurgické léčby epilepsie a významu lokalizace EZ v předoperačním vyšetření. Následně je poskytnut úvod do lokalizace EZ z stereo-EEG (SEEG) nahrávek, zahrnující představení klíčových interiktálních biomarkerů spojených s EZ a zkoumání metod analýzy grafů používaných pro lokalizaci EZ. Představení GNN rozšiřuje teoretický úvod o reprezentaci dat jako grafu, různé typy grafů, typy predikčních úkolů a klíčové komponenty modelů GNN. Tři konkrétní modely - Graph Sample and Aggregate (GraphSAGE), Graph Convolutional Network (GCN) a Graph Attention Network (GAT) - jsou prezentovány jako příklady GNN. Na závěr teoretické části jsou probírány aplikace GNN v analýze EEG.

V praktické části práce je popsán použitý soubor klinických dat, nastíněn proces výběru pacientů a podrobně popsány údaje o pacientech. V části věnované metodám jsou vysvětleny kroky předzpracování dat, včetně definice lokalizačních cílů a reprezentace EEG signálů ve formě grafů. Práce podrobně popisuje implementaci tří modelů GNN - GraphSAGE, GCN a GAT - spolu se dvěma referenčními modely, Spike Rate a Support Vector Machine (SVM), včetně procesu trénování a nastavení jejich hyperparametrů. Jsou popsány metody testování a vyhodnocování lokalizace EZ a predikce úspěšnosti epileptochirurgické operace. Následuje prezentace a diskuze výsledků, včetně porovnání výsledků modelů GNN s referenčními modely a porovnání s nejmodernějšími metodami v oblasti interiktální lokalizace EZ ze stereo-EEG dat. Praktickou část uzavírá rozbor limitací studie a návrhy pro budoucí výzkum.

Hodnocení modelů GNN na souboru 19 pacientů s dobrými výsledky chirurgického zákroku odhalilo výrazné zlepšení v lokalizaci epileptogenních zón oproti tradičním metodám. Modely prokázaly vynikající výsledky ve více metrikách, zejména Area Under the Receiver Operating Characteristic (AUROC) a Area Under the Precision-Recall Curve (AUPRC), kde modely GraphSAGE a GAT dosáhly mediánu skóre 0,971 AUROC a 0,563 AUPRC pro model GraphSAGE a 0,971 AUROC a 0,525 AUPRC pro model GAT, čímž statisticky výrazně překonaly tradiční modely, jako jsou SVM a Spike Rate. Tyto výsledky poukazují na robustní schopnost GNN rozlišovat mezi epileptogenními a neepileptogenními kontakty s vysokou přesností a spolehlivostí. Zejména model GAT vykazoval vynikající výsledky v preciznosti

pro identifikaci resekovaných kontakty, což zdůrazňuje jeho klinickou využitelnost při přesné identifikaci chirurgických cílů.

Další srovnání mezi pacienty s dobrými a špatnými postoperačními výsledky zvýraznilo účinnost těchto modelů v různých klinických aplikacích. Modely GAT a GraphSAGE vykazovaly významné rozdíly ve všech testovaných metrikách mezi oběma skupinami. Především model GAT prokázal výrazný rozdíl v přesnosti identifikace resekovaných kontaktů mezi skupinami dobrých a špatných postoperačních výsledků.

V oblasti predikce výsledků ukázala tato práce, že všechny modely GNN překonaly v přesnosti tradiční referenční modely, přičemž model GAT dosáhl přesnosti predikce 0,757. To naznačuje zřetelnou výhodu sítí GNN při rozpoznávání komplexních vzorů v datech SEEG, které korelují s výsledky epileptochirurgických zákroků. Model GAT zejména prokázal schopnost opravit chybné klasifikace referenčního modelu Spike Rate, což zdůrazňuje jeho potenciál pro zlepšení predikční spolehlivosti v klinických podmínkách. Statistické testy, konkrétně McNemarův test, však neukázaly žádné významné rozdíly mezi modely GNN a modelem Spike Rate, což naznačuje, že ačkoli GNN nabízejí lepší prediktivní poznatky, jejich statistická nadřazenost v predikci výsledku zákroku nad tradičními modely v této studii nebyla jednoznačně prokázána.

Diskuse v této diplomové práci zdůrazňuje několik kritických aspektů použití GNN pro plánování epileptochirurgických zákroků. Zaprvé, napříč kohortou pacientů je patrná značná variabilita ve výsledcích modelů, což je projevem přirozené komplexnosti a heterogenity klinických dat SEEG. Navzdory vysokým celkovým výsledkům v metrikách, jako jsou AUROC a AUPRC, se výkonnost u jednotlivých pacientů výrazně lišila, což poukazuje na obtížnost obecného použití těchto modelů v klinické praxi u různých pacientů. Studie dále porovnává dosažené výsledky s již publikovanými studiemi a ukazuje, že použité modely GNN (zejména GraphSAGE a GAT) obecně překonávají tradiční metody a předchozí studie, což svědčí o významném pokroku v této oblasti. Nicméně studie upozorňuje na několik omezení, včetně možnosti nadměrného přizpůsobení modelů na použitá data v důsledku malé velikosti datasetu, potřeby širší validace napříč různými institucemi a hodnocení pouze jediného experimentálního běhu, který nemusí plně zachytit variabilitu ve výkonnosti modelu. Tyto body naznačují oblasti pro další výzkum a zdokonalení s cílem zvýšit spolehlivost a použitelnost GNN v klinických podmínkách.

Závěrem lze říci, že prezentované výsledky zdůrazňují potenciál modelů GNN - konkrétně GraphSAGE a GAT - pro navádění epileptochirurgických operací pomocí interiktálních dat SEEG. Tato studie prokázala, že modely GNN dosahují signifikantně lepších výsledků při lokalizaci EZ ve srovnání s tradičními metodami. Reprezentace SEEG dat ve formě grafů, která umožňuje GNN zachytit topolog-

ické a funkční vztahy mezi signály, má mimořádný přínos při studiu komplexních epileptických sítí. Tyto modely vynikají v řadě výkonnostních metrik, což svědčí o jejich vyšší přesnosti a spolehlivosti v klinických aplikacích. Využití GNN rovněž prokázalo znatelně vyšší přesnost při predikci pooperačních výsledků, což dále zdůrazňuje jejich potenciál při napomáhání přesnějšimu a efektivnějšimu plánování epileptochirurgických zákroků. Prostřednictvím analýzy předpovědí a vizualizace výstupů modelů GNN lze potenciálně získat nové poznatky o epileptických sítích, které by mohly být podnětem pro další výzkum léčebných možností v oblasti farmakorezistentní epilepsie.

Tento výzkum představuje významný pokrok v léčbě epilepsie prostřednictvím implementace GNN pro interiktální analýzu SEEG, přičemž navrhovaná metodika otevírá cestu ke zvýšení přesnosti interiktální lokalizace EZ a tím zkrácení doby trvání invazivního SEEG monitorování.

HRTOŇOVÁ, Valentina. *Graph Neural Networks in Epilepsy Surgery*. Master's Thesis. Brno: Brno University of Technology, Faculty of Electrical Engineering and Communication, Department of Biomedical Engineering, 2024. Advised by Ing. Marina Filipenská, Ph.D.

Author's Declaration

Author: Bc. Valentina Hrtoňová
Author's ID: 203142
Paper type: Master's Thesis
Academic year: 2023/24
Topic: Graph Neural Networks in Epilepsy Surgery

I declare that I have written this paper independently, under the guidance of the advisor and using exclusively the technical references and other sources of information cited in the paper and listed in the comprehensive bibliography at the end of the paper.

As the author, I furthermore declare that, with respect to the creation of this paper, I have not infringed any copyright or violated anyone's personal and/or ownership rights. In this context, I am fully aware of the consequences of breaking Regulation § 11 of the Copyright Act No. 121/2000 Coll. of the Czech Republic, as amended, and of any breach of rights related to intellectual property or introduced within amendments to relevant Acts such as the Intellectual Property Act or the Criminal Code, Act No. 40/2009 Coll. of the Czech Republic, Section 2, Head VI, Part 4.

Brno

.....

author's signature*

*The author signs only in the printed version.

ACKNOWLEDGEMENT

I sincerely thank my thesis supervisor, Ing. Marina Filipenská, Ph.D., for her valuable comments and insight during the thesis project, and my thesis consultant, Ing. Petr Klimeš, Ph.D., for his expert guidance and support. Furthermore, I extend my gratitude to Petr Nejedlý, M.Sc., for inspiring and proposing the topic of this thesis.

Contents

Introduction	14
1 Theory	16
1.1 Drug-Resistant Epilepsy and Epilepsy Surgery	16
1.1.1 Epilepsy	16
1.1.2 The Concept of the "Epileptogenic Zone"	17
1.1.3 Presurgical Evaluation and Surgical Treatment	18
1.1.4 Surgical Outcome	19
1.2 Epileptogenic Zone Localization in Intracranial Electrophysiology . .	20
1.2.1 Stereo-Electroencephalography	21
1.2.2 Interictal Biomarkers of the Epileptogenic Zone	22
1.2.3 Graph Analysis in Epileptogenic Zone Localization	25
1.3 Graph Neural Networks	26
1.3.1 Graph Representation of Data	28
1.3.2 Graph Types	29
1.3.3 Prediction Tasks and Training	30
1.3.4 Computational Modules	30
1.3.5 Architectures	32
1.3.6 Applications in EEG Analysis	35
2 Clinical Dataset	37
2.1 Patients	37
2.2 Patient Data	37
2.2.1 SEEG Recordings	37
2.2.2 Sleep Staging	38
2.2.3 Electrode Contact Annotations	38
2.2.4 Surgical Outcomes	39
3 Methods	40
3.1 Data Pre-Processing	40
3.1.1 Localization Target Definition	40
3.1.2 EEG Signal Pre-Processing	42
3.1.3 Graph Representation of Data	45
3.2 GNN Models	48
3.2.1 GraphSAGE Network	49
3.2.2 Graph Convolutional Network	50
3.2.3 Graph Attention Network	51

3.2.4	Linear Classifier	51
3.2.5	Training	52
3.3	Reference Models	55
3.3.1	Spike Rate	55
3.3.2	Support Vector Machine	55
3.4	Hyperparameter Tuning	56
3.5	Testing and Evaluation	57
3.6	Outcome Prediction	59
4	Results	61
4.1	Epileptogenic Zone Localization	61
4.1.1	Performance for Good Outcome Patients	61
4.1.2	Performance for Good vs Poor Outcome Patients	64
4.2	Outcome Prediction	65
5	Discussion	68
5.1	Epileptogenic Zone Localization	68
5.1.1	Graph Neural Networks Outperform Reference Models	68
5.1.2	Models Show Variability in Patient Results	71
5.1.3	Models Perform Better for Good-Outcome Patients	72
5.2	Outcome Prediction	73
5.3	Comparison of GNN Models: GraphSAGE, GCN, and GAT	76
5.4	Comparison to Published Studies	78
5.5	Study Limitations	78
	Conclusion	80
	Bibliography	82
	Symbols and abbreviations	95
6	Disclosure of Artificial Intelligence Tool Usage	98
7	Electronic Attachment Content	99

List of Figures

1.1	Cortical Zones	18
1.2	Surgical Resection Outcome	20
1.3	SEEG Implantation and Recording	22
1.4	Interictal Epileptiform Discharge Patterns	23
1.5	HFOs in Intracranial EEG	24
1.6	General Framework of GNNs	27
1.7	Euclidean and Non-Euclidean Space	27
1.8	Graph and Adjacency Matrix	29
1.9	GNN Prediction Tasks	31
1.10	GraphSAGE Framework	33
1.11	Study of Grattarola et al.	36
3.1	GNN Pipeline	41
3.2	Localization Targets	42
3.3	Spike Detection	43
3.4	Functional Connectivity Matrix Before and After Thresholding	45
3.5	Relative Entropy Connectivity Matrix for Patient 723	46
3.6	GNN Model Implementation	49
3.7	GraphSAGE Training	52
3.8	GCN Training	53
3.9	GAT Training	54
4.1	Epileptogenic Zone Localization Results	62
4.2	Localization Results Good vs Poor Outcome Patients	66
4.3	McNemar's Test Results	67
5.1	Patient Case 723	70
5.2	ROC and PR Curves: GAT	71
5.3	Patient Case 74	74
5.4	Patient Case 82	75
5.5	Patient Case 1159	77

List of Tables

2.1	Engel's Postsurgical Outcome Classification	39
3.1	Electrode Contact Count Statistics	42
3.2	Clinical Dataset Data Structures	47
3.4	Hyperparameter Tuning: GNNs	56
3.5	Hyperparameter Tuning: SVMs	57
4.1	Model Localization Performance	61

Introduction

Epilepsy, one of the most prevalent and challenging brain disorders, affects around 70 million people worldwide [1]. This disorder is characterized by chronic recurrent seizures resulting from excessive neuronal cluster discharge [2]. Although numerous antiseizure medications are available, they fail to control seizures effectively in about 30-40% of epilepsy patients [3, 4]. For these individuals with drug-resistant epilepsy, epilepsy surgery is frequently the treatment of choice when other options are insufficient [5]. The goal of epilepsy surgery is to completely remove or disconnect the epileptogenic zone (EZ) defined as "the area of cortex indispensable for the generation of clinical seizures" while preserving the eloquent cortex¹ [7]. While the EZ is a theoretical concept that cannot be directly measured by any of the currently available methods, a variety of diagnostic tools, including seizure semiology, electroencephalography recordings, functional testing, and neuroimaging techniques, are currently used to localize the potential EZ and construct an individualized resection plan for each patient [7, 8]. The seizure-onset zone (SOZ), identified based on electroencephalography recordings, is currently used by clinicians to approximate the EZ [5].

Among the used diagnostic tools, stereo-electroencephalography (SEEG) has emerged as a valuable and relatively well-tolerated invasive method for correctly identifying the SOZ for patients with unclear noninvasive data [5]. Through the implantation of depth electrodes, SEEG provides a unique opportunity to capture electrical patterns directly from within the brain and to visually identify the SOZ from the recordings of seizures. This option, while effective, comes with inherent challenges. Patients undergoing SEEG must tolerate the invasive nature of the procedure, face associated risks, and endure the process of reducing antiseizure medication (ASM) to minimize seizures. This reduction process is intricate and can take 1-2 weeks, sometimes extending to 4 weeks. Moreover, traditional methods relying on visual inspection of recordings by medical professionals are time-consuming and subjective. The fact that only approximately 60% of carefully selected patients achieve seizure freedom after epilepsy surgery, often due to the challenges in accurately identifying the EZ, underscores the urgent need for advanced diagnostic tools capable of localizing the EZ quickly and precisely using SEEG data [5].

Automatic methods, particularly those analyzing interictal data², present a promising avenue to address these challenges and simplify the entire presurgical workup process. By leveraging automated approaches, the need for prolonged inva-

¹Eloquent cerebral structures are defined as areas of the brain with readily identifiable neurological function, in which injury results in disability. [6]

²Interictal data refers to the data collected away from seizures.

sive monitoring in the Epilepsy Monitoring Unit (EMU) can be minimized, along with the associated risks and costs. Consequently, there is an opportunity to enhance the efficiency of epilepsy surgery, providing a more timely and effective intervention for patients facing drug-resistant epilepsy.

Several promising automatic methods analyzing interictal data have been developed, including those studying high-frequency oscillations [9], interictal spikes [10], and connectivity measures [11, 12, 13]. While these methods have shown effectiveness, they often overlook crucial aspects, such as the specifics of each patient’s electrode implantation, and fail to consider the topology of brain networks, which are essential for accurate seizure source localization. This limitation highlights the need for more sophisticated tools that can integrate this information to enhance the understanding of epileptic activity and improve clinical outcomes.

Graph Neural Networks (GNNs) have emerged as robust deep learning tools for electroencephalography (EEG) signal analysis, demonstrating state-of-the-art performance in diverse applications such as automating sleep staging [14] and emotion recognition [15]. The distinctive strength of GNNs lies in their ability to represent EEG data as a graph, enabling them to capture implicit topological and functional relationships between signals. This unique capability is particularly relevant in examining epileptic networks, as GNNs can potentially uncover the patterns and relationships that traditional methods struggle to discern. By leveraging the graph-based nature of GNNs, it becomes possible to incorporate the patient-specific electrode implantation and brain network topology into the analysis, offering a more tailored and accurate approach to localizing the EZ.

This thesis aims to leverage GNNs to localize the EZ in drug-resistant epilepsy patients using interictal SEEG recordings. The objective is to streamline the presurgical workup process and enhance clinical insights in epilepsy diagnostics by providing a more efficient and effective approach to EZ localization for patients facing drug-resistant epilepsy. This approach not only addresses the shortcomings of existing diagnostic methods but also paves the way for advancements in personalized epilepsy treatment.

1 Theory

This Chapter begins with an overview of drug-resistant epilepsy and the role of epilepsy surgery, discussed in detail in Section 1.1. It then summarizes the methods of localizing the EZ from intracranial EEG recordings, covered in Section 1.2. Following this, Section 1.3 introduces and explains the concept of GNNs.

1.1 Drug-Resistant Epilepsy and Epilepsy Surgery

This Section provides a comprehensive overview of DRE and the surgical interventions used to treat it. Subsection 1.1.1 covers the nature of epilepsy and defines what constitutes DRE, Subsection 1.1.2 explores the concept of the EZ, Subsection 1.1.3 discusses methods for presurgical evaluation and the criteria for surgical treatment, and Subsection 1.1.4 examines surgical outcomes.

1.1.1 Epilepsy

Epilepsy is defined as a brain disorder characterized by the occurrence of repetitive and unprovoked seizures. The term "seizure" refers to a sudden and temporary change in neurological function resulting from an abnormal excessive or synchronized activity of neurons in the brain, and epileptic seizures, caused by abnormal neuronal firing, can be distinguished from non-epileptic events, such as psychogenic seizures, by epileptic discharges in the EEG. [16, 17] Epilepsy can arise from numerous causes, which can be categorized as genetic, structural, metabolic, infectious, immune, or unknown [18]. The developing brain is particularly prone to seizures, reflected by the fact that approximately 75% of epilepsy cases begin during childhood [16].

ASMs, which help 60-70% patients with any epilepsy achieve seizure control, are the mainstay of symptomatic epilepsy treatment [16]. The ASMs suppress seizures by interactions with various cellular targets, modifying the bursting properties of neurons and reducing synchronization in the neuronal system [17]. Despite the availability of more than 30 ASMs with diverse anti-seizure mechanisms, seizures may be resistant to treatment in 30-40% of all patients with epilepsy [3, 4, 19]. For patients with drug-resistant epilepsy who have failed to respond to two appropriate ASMs [4], treatment options include resective epilepsy surgery, neuromodulation with neurostimulation devices, and ketogenic diet therapy [17]. Studies have shown that 4.5% of all patients with epilepsy could potentially benefit from resective epilepsy surgery [7] whose primary goal is the complete resection or disconnection of the "epileptogenic zone".

1.1.2 The Concept of the "Epileptogenic Zone"

The theoretical concept of the "epileptogenic zone" was defined as "the area of cortex that is necessary and sufficient for initiating seizures and whose removal (or disconnection) is necessary for the complete abolition of seizures" in 1993 by Luders et al. [20]. No currently available method can be used to measure the EZ directly. However, its location can be inferred by defining the following cortical zones [7] illustrated in Figure 1.1:

- **The irritative zone** is the area of the cortex that generates interictal epileptiform discharges (IEDs), which can be captured by EEG, magnetoencephalography (MEG), or functional magnetic resonance imaging (fMRI). To give rise to clinical symptoms, IEDs must be repetitive, have sufficient strength, and invade the symptomatogenic zone.
- **The symptomatogenic zone** is the cortical area that produces the ictal symptoms when activated by an epileptiform discharge. This zone can be delineated by analyzing the patient's ictal symptoms through the analysis of seizure history and ictal video recordings or invasively by direct cortical electrical stimulation, which creates conditions similar to the activation of the cortex by an epileptiform discharge.
- **The seizure-onset zone** is the area of cortical tissue from which clinical seizures are generated, and it is usually the area of the irritative zone producing IEDs capable of inducing clinical symptoms. The SOZ is localized by visual inspection of the scalp or intracranial EEG or by ictal single photon emission computed tomography (SPECT). Although in current practice, the SOZ is often used as a proxy for the EZ [5], the extent of these areas often does not correspond, meaning that the EZ is either smaller or larger than the clinically determined SOZ.
- **The epileptogenic lesion** is the radiographic lesion that causes epileptic seizures. Brain lesions are localized by high-resolution magnetic resonance imaging (MRI), and their epileptogenicity is determined by video-EEG monitoring and seizure semiology. The resection of the MRI-visible epileptogenic lesion is often not sufficient to achieve seizure freedom, as the EZ can extend beyond lesion borders. Conversely, if the EZ covers only a part of the epileptogenic lesion, a partial resection of the lesion may result in complete seizure freedom.
- **The functional deficit zone** is the cortical area, which is functionally abnormal in the interictal period. Although the functional deficit zone's relationship to the EZ and the value of its presurgical definition is limited compared to the other zones, the correlation of the functional deficit zone with the other areas

may provide supporting information.




Fig. 1.1: Illustration of the cortical zones in the epileptic brain. Adapted from [21].

Apart from the above-mentioned cortical zones, the localization of the eloquent cortex adjacent to the seizure onset zone is necessary to construct a surgical plan. The resection of the eloquent cortex, which is the cortex related to a given neurological function, may result in the loss of that function, and the possible deficits should be discussed with the patient before the procedure. [7]

1.1.3 Presurgical Evaluation and Surgical Treatment

An extensive presurgical work-up is conducted to determine the patient's eligibility for surgery, starting with noninvasive testing including detailed medical history, neuropsychological evaluation, brain MRI with an epilepsy-dedicated protocol, and long-term video EEG [22]. Additional investigations, including positron emission tomography (PET), SPECT (both interictal and ictal), MEG, and fMRI, may be employed to investigate the EZ further [17, 22].

Curative surgery may be directly considered if the noninvasive phase is conclusive and the delineation of the epileptogenic focus from the eloquent areas is achieved. This might include cases of epilepsy derived from a well-defined focal brain lesion or cases of mesial temporal lobe epilepsy caused by mesial temporal sclerosis, which are associated with good surgical outcomes (60-90%) and do not always require invasive testing. [17, 4, 22] However, when noninvasive evaluation yields inconclusive results (25-50% of cases), long-term invasive intracranial EEG monitoring is indispensable for precise localization of the epileptogenic focus [8].

Following an exhaustive presurgical evaluation focusing on the definition of the cortical zones and postoperative deficit risk assessment, various surgical strategies are considered based on seizure type and location. Resective surgery, the conventional approach involving open-skull removal of epileptogenic tissue, is an option with broad applicability but necessitates careful consideration of potential functional deficits and recovery. Alternatively, less invasive options include SEEG-guided radiofrequency thermocoagulation, which disrupts the EZ with heat-induced lesions, and laser interstitial thermal therapy, utilizing laser-induced heat for ablation. [3, 23]

Additionally, even in cases where the multimodal evaluation provides discordant results or when the resection of the epileptogenic region is not possible, palliative surgery, which aims to reduce the frequency and severity of seizures, may be a viable option [17, 4].

1.1.4 Surgical Outcome

The success of both curative and palliative epilepsy surgery relies on accurate presurgical localization of the EZ. It is crucial to note that currently, the presurgical delineation of the EZ is not feasible, even in cases where information gained from defining the main cortical zones is concordant. This is due to the possibility that a potential SOZ exists apart from the actual SOZ where the patient's seizures are generated, which would only become clinically apparent postoperatively [7]. Examples of seizure-free and seizure-persistent outcomes are illustrated in Figure 1.2.

It is only after the surgery successfully eliminates seizures that we can be sure the EZ was either resected or disconnected. This complicates the use of the definition as a guide when developing a surgical plan and necessitates the use of postoperative outcomes in confirming diagnostic accuracy.

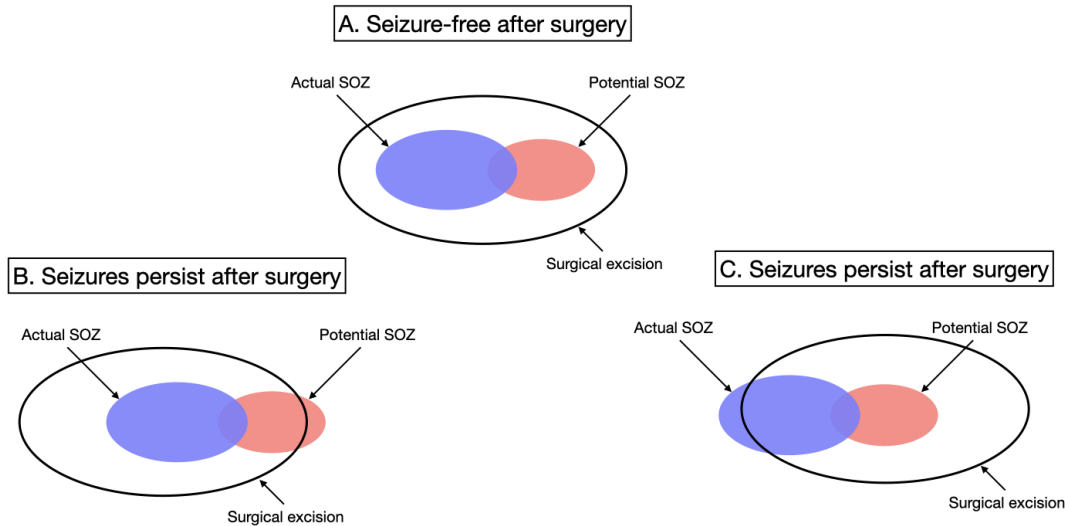


Fig. 1.2: Examples of surgical resection of the actual and potential seizure-onset zone (SOZ) with seizure-free or seizure-persistent surgical outcomes. Adapted from [24].

1.2 Epileptogenic Zone Localization in Intracranial Electrophysiology

Invasive EEG monitoring, conducted with depth electrodes (SEEG) or subdural strip or grid electrodes (electrocorticography, ECoG), directly captures the onset and propagation of seizures from the cortex. ECoG, which enables dense sampling of the subdural cortex, can be particularly helpful in localizing the irritative and EZs and identifying the eloquent cortex regions. Additionally, ECoG can be employed intraoperatively to assist the neurosurgeon during the surgical procedure. Meanwhile, SEEG, utilizing depth electrodes, enables the recording of seizure onset and spread even from deep brain structures beyond the reach of scalp, strip, or grid EEG electrodes. [22]

This Section delves into the methods for localizing the EZ using intracranial electrophysiological techniques, with a focus on the use of SEEG in monitoring interictal brain activity described in Subsection 1.2.1. It examines the key biomarkers of the EZ detectable in interictal periods in Subsection 1.2.2 and evaluates how graph analysis techniques can be applied to SEEG data to improve the accuracy of EZ localization in Subsection 1.2.3.

1.2.1 Stereo-Electroencephalography

The SEEG method was developed by Jean Talairach and Jean Bancaud in the 1950s with the aim of implementing a methodology for comprehensive analysis of the morphological and functional brain space [25]. While scalp EEG and ECoG primarily capture the electrical activity from the outer layers of the cortex, SEEG directly records the Local Field Potential (LFP), also referred to as 'micro-EEG'. The recorded signal contains action potentials and other fluctuations in membrane potential within a small group of neurons, providing a highly detailed and informative signal for studying the mechanism of cortical electrogenesis. [26] SEEG are characterized by their high amplitude, typically ranging from 50 to 1,500 μV , and high spatial resolution, generally about 3.5 mm. To capture the high-frequency changes in the signal, the recordings are typically sampled with sampling rates from 500 to 2,000 Hz. [8]

Implantation and Monitoring

Multicontact SEEG electrodes, varying in length and number of contacts, are employed to target specific brain regions. Typically, commercially available SEEG electrodes are characterized by flexibility and semi-rigidity, featuring a rounded tip as a recording contact. Each electrode typically consists of 5–18 contacts, regularly spaced at intervals of 2–5 mm or 10 mm along the electrode. [27]

Various implantation techniques, including frame-based, frameless, and robotic systems, are currently employed to insert the depth electrodes through 2.5 mm diameter drill holes under general anesthesia [25]. While frameless techniques are more time-efficient than frame-based methods, they generally exhibit lower overall accuracy. On the other hand, robotic-assisted methods, such as Neuromate and ROSA, provide exceptional accuracy with shorter implantation times compared to non-robotic techniques but require expensive robotic equipment. [8] An illustrative case of SEEG implantation, and recording is shown in Figure 1.3.

Following the surgical procedure, the patient is transferred to the EMU, where clinical and electrographic monitoring through SEEG electrode recordings is conducted. The stay in the EMU is approximately seven days but varies based on factors such as the quantity, quality, and patterns of ictal and interictal recordings. Throughout the SEEG monitoring period, the patient is not administered any AMS to increase the probability of seizures. Once a sufficient amount of ictal data has been recorded, the patient returns to the operating room for the removal of the SEEG electrodes. [28]




Fig. 1.3: Illustrative case of SEEG implantation with X-ray imaging (A) and SEEG recording (B). Adapted from [25].

Advantages and Limitations

SEEG presents several significant advantages in epilepsy monitoring. It allows for the recording of deep brain structures, enables bihemispheric recordings, serves as an alternative if subdural grids fail, and facilitates three-dimensional mapping of epileptic networks without the need for a large craniotomy [28, 22].

However, SEEG also has major limitations. The placement of depth electrodes requires an invasive neurosurgical procedure, introducing inherent risks and increasing the complexity of the monitoring process. Additionally, the spatial coverage of SEEG may be limited compared to scalp EEG and ECoG, which can capture more widespread brain activity. Moreover, prolonged stays in the EMU pose significant risks, while they may not ensure seizure occurrence. [29, 5]

1.2.2 Interictal Biomarkers of the Epileptogenic Zone

Identifying the EZ from interictal data can significantly reduce the length of EMU stays and facilitate diagnosis even for patients who do not experience seizures during the monitoring period. Several promising automatic methods of analyzing interictal data have been developed, including those studying interictal epileptiform discharges, high-frequency oscillations, and connectivity measures.

Interictal Epileptiform Discharges

Interictal epileptiform discharges (IEDs), often called 'spikes', are characterized by brief and abnormal electrical discharges resulting from the synchronous firing of hyperexcitable neurons. Traditionally, a spike is clinically defined as a sharp transient

with a duration between 20 and 70 ms, which can be clearly distinguished from background activity [30]. Figure 1.4 shows spike examples.

Spikes can be observed between seizures in patients with epilepsy, and the variability in spike rate and spatial extent with sleep state and seizure activity suggest fluctuations in spike activity over time [31]. Moreover, the relationship between spikes and seizures is complex and remains unclear, with some studies proposing that spikes may promote seizure generation while others suggest a protective benefit against seizures. [32, 33]




Fig. 1.4: Interictal epileptiform discharge (A) and a group of interictal epileptiform discharges from neocortical dysplasia (B) recorded in human partial epilepsies with intracranial EEG. Adapted from [33].

The resection of areas with a high frequency of spikes was associated with seizure-free outcome [34, 35], with spikes preceded by gamma activity proving to be especially discriminative of the EZ [36, 10]. Nevertheless, research has demonstrated that while interictal spike activity serves as a valuable tool for localizing the seizure generator, achieving seizure freedom does not always necessitate the removal of all areas generating spikes [37].

High-Frequency Oscillations

High-frequency oscillations (HFOs) are transient oscillatory events observed in high frequency ranges between 80 Hz and 500 Hz, consisting of at least four oscillations that clearly stand out from background activity [38]. HFOs can further be divided into ripples (80–250 Hz), fast ripples (250–500 Hz), and very-fast ripples with frequencies exceeding 500 Hz (17–20) [39, 40]. Ripples and fast ripples in intracranial

EEG are shown in Figure 1.5.

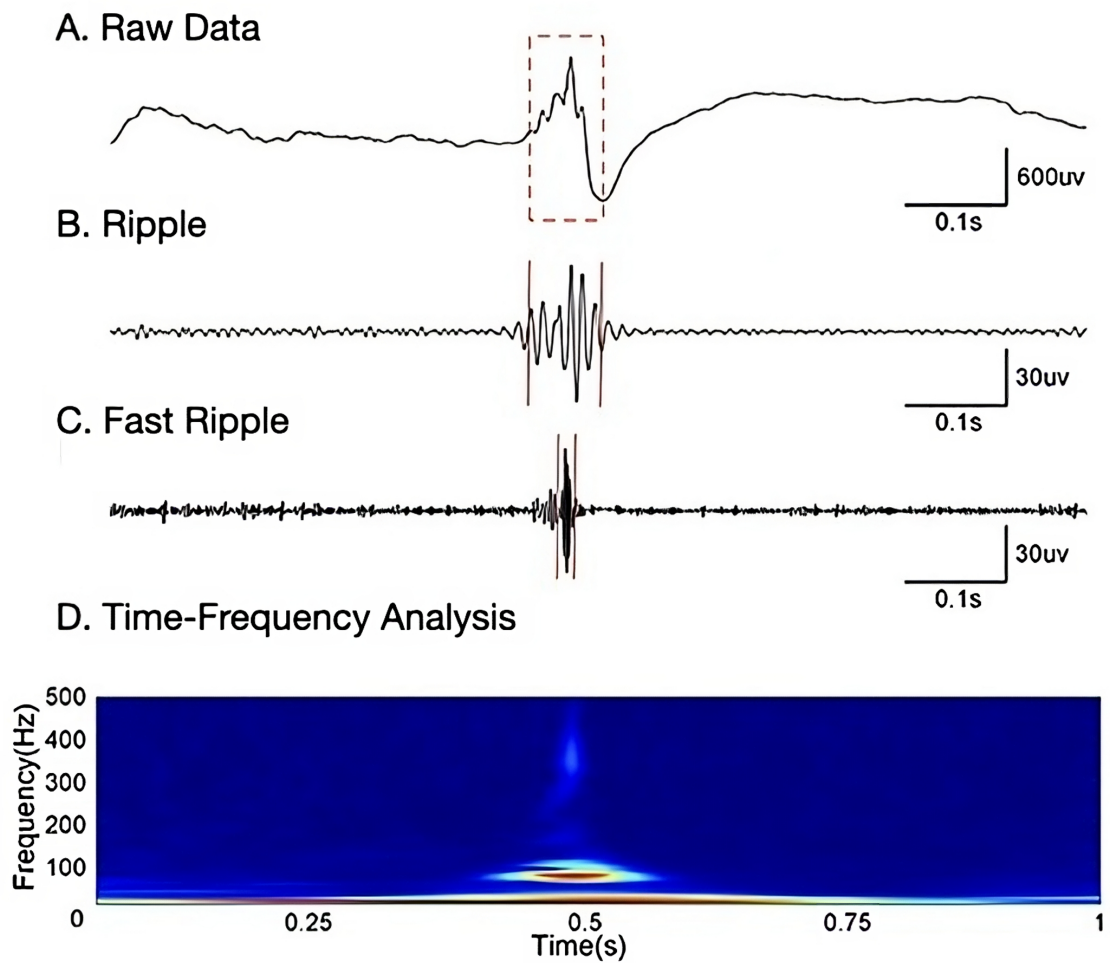


Fig. 1.5: High-frequency oscillations in the intracranial EEG recording. Data was filtered between 80-200 Hz (ripples) and 200-500 Hz (fast ripples). Adapted from [41].

Retrospective studies have demonstrated that the resection of areas characterized by high HFO rates is correlated with a more favorable seizure outcome compared to the resection of regions with low HFO rates [38, 42, 43]. However, since HFOs are also associated with normal brain function [44], distinguishing between physiological and pathological HFOs is imperative for their utility as biomarkers of the EZ. Despite this challenge, several methods based on HFOs have shown promise for the localization of the EZ with the frequency, amplitude, and duration of HFOs as main measures [9, 45, 46]. Furthermore, the combination of interictal spikes and oscillations in the ripple band improved the localization of epileptogenic tissue compared to spikes, spikes-gamma, wideband HFOs, ripples, and fast ripples [47].

Connectivity Measures

Researchers explore three types of brain connectivity: structural (anatomical), functional, and effective connectivity. Structural connectivity involves the physical interconnection of axons, detectable through noninvasive MRI. Functional connectivity assesses statistical dependencies between brain regions, typically through the correlation of time-series signals, while effective connectivity reveals directional data flow during cognitive tasks, uncovering causal brain networks. To measure functional and effective connectivity, intracranial EEG, functional MRI, and MEG are commonly used [48, 49].

Widely used functional connectivity measures include Pearson’s correlation coefficient, coherence, phase-locking value, and **relative entropy** (discussed in detail in Subsection 3.1.2). **Pearson’s correlation coefficient**, measuring the linear correlation between two signals, has been presented as a potential predictor of seizure outcome in temporal lobe epilepsy patients [50]. Additionally, epileptogenic tissue was associated with increased phase synchronization, as measured by mean phase **coherence**, during interictal periods, with changes in synchronization preceding seizures [51]. The **phase-locking value**, also measuring phase synchronization of two signals, has successfully been used to identify the area of seizure onset [52].

The main effective connectivity measures used for EZ localization are the Granger causality measures and directed information. **Granger causality** measures enable the exploration of causal relationships between brain regions by examining the directed flow of information within the brain [53], showing potential in localizing seizure networks from interictal data [54, 55]. Effective connectivity can also be inferred based on **directed information**, measuring the directed flow of information between two signals. Directed information has proven helpful for inferring the SOZ as the sole connectivity measure [56], and in combination with Granger causality [57].

1.2.3 Graph Analysis in Epileptogenic Zone Localization

As epilepsy is increasingly viewed as a network disorder, there is a growing emphasis on proposing network-based measures to characterize the connectivity patterns within the epileptic brain network [58]. Graph theory is a promising approach for analyzing networks, offering valuable measures to characterize brain networks’ topological properties and functional organization in both normal functioning and diseases. In this approach, the brain is conceptualized as a graph composed of nodes (such as SEEG contacts) and edges, representing interactions between pairs of nodes. The graph representation of data will be further discussed in Subsection 1.3.1.

In an analysis of ictal effective connectivity, the electrode contact exhibiting the highest total out-degree consistently corresponded to the clinically marked SOZ and was located within the resected brain region in postoperatively seizure-free patients [59]. In another study, high degree and betweenness centrality values at ictal-onset were correlated to the EZ, with in-degree being the most effective measure [60].

Importantly, studies indicate that abnormalities in the epileptic network are not only evident during seizures but also during rest and even independent of spike occurrence [61]. The analyzed time-varying effective connectivity networks in various non-seizure and seizure periods and found that high values of graph measures like out-degree, closeness centrality, betweenness centrality, and clustering coefficient localized the SOZ and the EZ [62]. Furthermore, greater overlap between resected and highly synchronous electrodes was associated with good surgical outcomes, suggesting that spatially-informed interictal graph measures can effectively distinguish between good and poor surgical outcomes [63].

1.3 Graph Neural Networks

Graph Neural Networks, first proposed in 2009 [64], now represent the state-of-the-art deep learning methods for analyzing graph-structured data. GNNs transform an input graph structure to generate embeddings for nodes, edges, and entire graphs. The key idea is to propagate information across the graph, allowing each node to gather and update its representation based on the information from its neighborhood. This process is typically conducted through iterative message-passing steps performed by GNN layers, where each node aggregates information from its neighbors and updates its representation. The resulting embeddings can then be used for various prediction tasks. The GNN framework is illustrated in Figure 1.6.

Essentially, GNNs can be viewed as an extension of conventional Convolutional Neural Networks (CNNs) tailored for processing data organized in graph structures instead of the processing of images. While CNNs can only operate on Euclidean data such as images (2D matrices) and texts (1D sequences), these data structures can be regarded as instances of graphs. The generalization of CNNs to graph structures extends the applicability of deep neural models beyond Euclidean domains and into non-Euclidean graph-based domains, which enables modeling relationships and dependencies in complex data. Figure 1.7 illustrates the difference between Euclidean and non-Euclidean data.

The design of a GNN model consists of four main parts described in this Section: 1) the representation of the data as a graph (Subsection 1.3.1), 2) the specification of graph type and structure (Subsection 1.3.2), 3) the design of the learning task (Subsection 1.3.3), and 4) the construction of the GNN model using computational

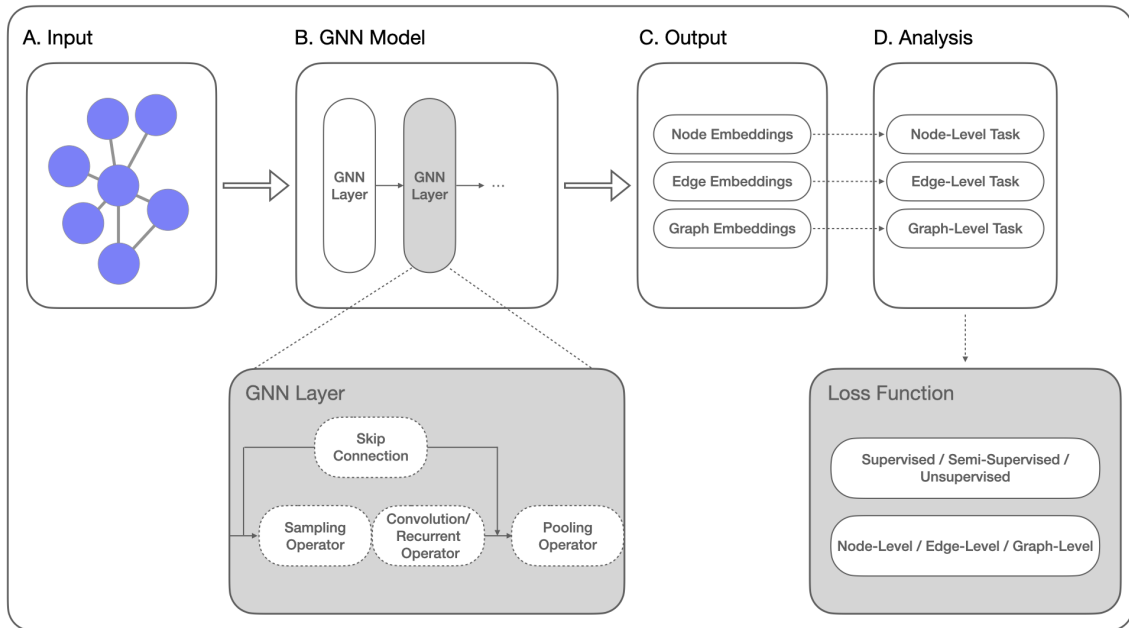



Fig. 1.6: In the general framework of GNNs, the input computational graph (A) is transformed by a series of GNN layers (B), and the resulting node, edge, and graph embeddings (C) can be used for various prediction tasks (D). The training loss function is defined based on the prediction task and training setting.

A. Euclidean Space



B. Non-Euclidean Space




Fig. 1.7: Illustration of data representation in the Euclidean (A) and non-Euclidean space (B). CNNs operate in Euclidean space, while GNNs also extend to non-Euclidean space.

modules (Subsection 1.3.4). [65] In Subsection 1.3.5, three established architectures of GNNs are described, and the applications of GNNs in EEG analysis are outlined in Subsection 1.3.6.

1.3.1 Graph Representation of Data

The graph structure must be identified in the data based on the application of the model. The structure is either explicitly given by the structural nature of the data, such as the molecular structure of analyzed proteins [66], or it must be constructed, exemplified by the design of word graphs in natural language processing [67]. The resulting graph representation then serves as an input for the GNN.

A **graph** $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ consists of a set of **nodes** \mathcal{V} , also called vertices, and a set of **edges** \mathcal{E} connecting these nodes. $|\mathcal{V}| = N$ or N^v is the number of nodes in the graph and $|\mathcal{E}| = N^e$ the number of edges. [65]

Adjacency Matrix

The **adjacency matrix** $\mathbf{A} \in \mathbb{R}^{N \times N}$ is commonly used to represent the connectivity patterns in the graph, stored as edge weights. Each edge $(i, j) \in \mathcal{E}$ represents a relationship between nodes i and j . For a weighted graph, \mathbf{A}_{ij} equals the weight of the edge (i, j) , whereas for an unweighted graph, \mathbf{A}_{ij} is defined as follows [68]:


$$\mathbf{A}_{ij} = \begin{cases} 1, & \text{if there is an edge between nodes } i \text{ and } j, \\ 0, & \text{otherwise.} \end{cases} \quad (1.1)$$

Moreover, edges in the graph structure can be directed or undirected. In the case of directed edges, a clear distinction exists between a source node, where information originates, and a destination node, representing the endpoint of the directed flow of information. Conversely, in an undirected edge, information flows bidirectionally, and its adjacency matrix is therefore symmetric ($\mathbf{A}_{ij} = \mathbf{A}_{ji}$).

Node Features Matrix

In addition to the structural information provided by the adjacency matrix, each node $v \in \mathcal{V}$ in the graph may have associated features. The features of all nodes are typically represented in a feature matrix $\mathbf{X} \in \mathbb{R}^{d \times |V|}$, where each row corresponds to a node and each column corresponds to a feature dimension. [68] For example:

A. Undirected Graph



B. Adjacency Matrix

		To				
		A	B	C	D	E
From	A	0	1	1	0	1
	B	1	0	1	0	0
	C	1	1	0	1	1
	D	0	0	1	0	0
	E	1	0	1	0	0

Fig. 1.8: Example of an undirected graph data representation (A) and a corresponding adjacency matrix (B). The undirected graph has five nodes A-E, some connected with edges. The adjacency matrix represents the undirected graph within which a value of 1 indicates the existence of an edge between nodes, while a value of 0 signifies the absence of an edge.

$$\mathbf{X} = \begin{bmatrix} x_{1,1} & x_{1,2} & \dots & x_{1,d} \\ x_{2,1} & x_{2,2} & \dots & x_{2,d} \\ \vdots & \vdots & \ddots & \vdots \\ x_{N,1} & x_{N,2} & \dots & x_{N,d} \end{bmatrix}, \quad (1.2)$$

Where N is the number of nodes and d is the number of node features.

1.3.2 Graph Types

Graph type and scale are key concepts for the construction and processing of graphs. Some of the usual categories, which are mutually orthogonal, are:

- **Directed/Undirected:** Described in Subsection 1.3.1.
- **Homogeneous/Heterogeneous:** In homogeneous graphs, nodes and edges are of the same type, while heterogeneous graphs incorporate nodes and edges of multiple types.

- **Static/Dynamic:** Dynamic graphs exhibit varying input features or graph topology, while static graphs remain unchanged.

There is no clear classification in terms of graph scale, and small and large graphs are usually distinguished based on a device’s capacity to store and process them.

1.3.3 Prediction Tasks and Training

To design a suitable loss function that will allow the optimal training of the model, the general task type and training setting must be defined. There are three main prediction tasks conducted on graphs:

- **Node-Level Tasks:** Property of each node within the graph is predicted in node-level tasks, with node embedding as the main model outputs. This category includes tasks such as node classification, where each node is assigned a label; node regression, where a numerical property for each node is predicted; and node clustering, which involves grouping nodes based on their features or connections.
- **Edge-Level Tasks:** Edge-level tasks involve predicting properties associated with the edges of a graph. Common examples include edge classification, where edges are categorized into different types, and link prediction, where the existence of an edge between two nodes is forecasted. In these tasks, the primary output is the embeddings of the graph edges.
- **Graph-Level Tasks:** GNNs are also applied to tasks that involve predicting a single property for the entire graph. This could be a classification task, where the goal is to categorize the graph as the whole, or a regression task, where a numerical property of the entire graph is predicted. The key output here is the learned embeddings of the whole graph.

Examples of the prediction tasks are illustrated in Figure 1.9.

Based on the supervision setting of the task, we differentiate:

- **Supervised models:** Labels are provided to the model for training.
- **Semi-supervised models:** Only a small amount of labeled data is provided.
- **Unsupervised models:** No labels are provided to the model, and clustering is commonly applied to find patterns in the data.

1.3.4 Computational Modules

The three commonly used computational modules to construct the GNN model are sampling, propagation, and pooling. A GNN model is typically built by combining and layering these modules.




Fig. 1.9: Examples of prediction tasks performed with GNNs. Node classification (A) is an example of a node-level task, link prediction (B) of an edge-level task, and graph classification (C) of a graph-level task.

Sampling Modules

To propagate information, GNN models aggregate information for each node’s neighborhood in the preceding layer into the node in the current layer. Therefore, the node’s neighborhood grows exponentially with the depth of the network, necessitating sampling of the neighborhood. In large graphs, sampling is also needed for storing and processing information. Different strategies for sampling based on node sampling, layer sampling, or subgraph sampling can be employed.

Node sampling involves selecting a subset of neighboring nodes for aggregation, as demonstrated by Graph Sample and Aggregate (GraphSAGE) in Subsection 1.3.5. **Layer sampling** used in models like FastGCN [69] preserves only a limited set of nodes for aggregation within each layer, thereby regulating the expansion of the node’s neighborhood. **Subgraph sampling** samples multiple subgraphs from the computational graph and confines the neighborhood search within these subgraphs instead of directly sampling nodes and edges from the entire graph.

Propagation Modules

Propagation modules transmit information between nodes through edges, aggregating information about node features and the graph topology. Convolution or recurrent operators are typically used to gather information from neighboring nodes, and skip connections can be included to incorporate information from past node representations and prevent over-smoothing.

From the propagation modules, **convolution operators** are the most popular. Convolution operators work directly with the graph topology (i.e., spatial approaches) or with the spectral representations of graphs (i.e., spectral approaches).

The choice between spatial and spectral approaches depends on the specific requirements and characteristics of the graph data. The GraphSAGE propagation framework [70] is a popular spatial convolutional module for induction and semi-supervised tasks. Attention-based spatial approaches, such as the Graph Attention Network (GAT) [71], can assign different weights to neighbors for information propagation, enabling noise reduction and improved results.

Reccurent operators, unlike convolution operators, share the same weights across layers, and they contain mainly convergence-based and gate-based methods. Convergence-based methods aim to learn state embeddings for each node, with the limitation of lower efficiency. Gate-based methods, which do not guarantee convergence, aim to enhance computational efficiency and long-term information propagation with the Gated GNN (GGNN) [72] as an example.

Pooling Modules

Propagation modules are typically followed by pooling layers, which help collect more general features from the graph. Direct and hierarchical pooling layers promote information propagation even in complex and large-scale graphs.

Direct pooling modules, also known as readout functions, directly learn graph-level representations from nodes using different node selection strategies. Simple node pooling methods include node-wise max, mean, sum, or attention operations applied to node features to obtain a global graph representation. More advanced SortPooling mechanism [73] first sorts the node embeddings according to their structural roles and then obtains their representation by passing them through CNNs. On the other hand, **hierarchical pooling modules** consider hierarchy in the graph structure and learn graph representations layer by layer.

1.3.5 Architectures

GraphSAGE

The GraphSAGE inductive framework, introduced by Hamilton et al. [70], is an example of node sampling with spatial convolution. It enables the generalization of node embeddings across graphs for node classification by training a set of sampling and aggregator functions that aggregate information from the node’s neighborhood instead of training the embedding vector for each node. The main components of the GraphSAGE framework are illustrated in Figure 1.10.

For the generation of embeddings, a fixed-size neighborhood N_i^k in the distance of k -hops from each node i is first sampled. The parameter k defines the depth of the search, and each layer k has a corresponding aggregation function with a distinct

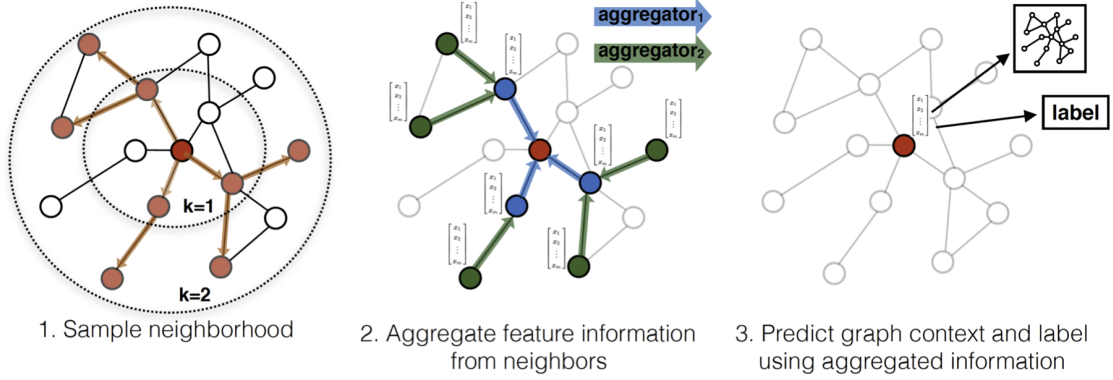


Fig. 1.10: Illustration of the GraphSAGE framework steps: (1) sampling a fixed-size neighborhood, (2) aggregating feature information from neighbors, and (3) predicting graph context and node labels using the aggregated features. Adapted from [70].

trainable matrix of weights W^k . Node features serve as initial node embeddings. Next, an aggregation function is used to combine information from the neighborhood into a single vector for each node as

$$h_i^{k'} = \text{AGGREGATE}_k(\{h_j^{k-1}, \forall j \in N_i^{k-1}\}), \quad (1.3)$$

where $h_i^{k'}$ is the aggregated representation for node i at the k -th layer, and h_j^{k-1} are the embeddings from the previous layer for each neighboring node j within the k -hop neighborhood N_i^k . Finally, the node representation is updated by the combination of the aggregated representation $h_i^{k'}$ with the existing node representation h_i^{k-1} :

$$h_i^k = \sigma(W^k \cdot \text{CONCAT}(h_i^{k'}, h_i^{k-1})), \quad (1.4)$$

where h_i^k is the new embedding of node i , W^k is a matrix of weights for the k -th layer, and σ is a non-linear activation function.

The model parameters are trained through standard stochastic gradient descent and backpropagation techniques, and the trained model is applied to the unseen nodes during testing.

Graph Convolutional Network

Graph Convolutional Networks, first introduced by Kipf and Welling [74], represent a significant advance in graph neural networks by simplifying the graph convolution process.

GCNs are designed to utilize a spectral approach involving the eigendecomposition of graph Laplacians to define graph convolutions. The key innovation of GCNs,

which simplifies the traditional spectral graph convolution for propagation, can be expressed as:

$$H^{(l+1)} = \sigma(\tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}} H^{(l)} W^{(l)}) \quad (1.5)$$

where $H^{(l)}$ is the matrix of activations in the l -th layer (with $H^{(0)} = X$), $\tilde{A} = A + I$ is the adjacency matrix of the graph A with added self-connections (identity matrix I), \tilde{D} is the diagonal degree matrix of \tilde{A} , $W^{(l)}$ is the trainable weight matrix for the l -th layer, and σ denotes the activation function such as Rectified Linear Unit (ReLU).

This approach takes advantage of the localized first-order approximation of spectral graph convolutions, effectively capturing information about immediate neighbors without excessive computational cost. By operating directly on graphs, GCNs are particularly effective for tasks such as node classification, where labels are available for at least a subset of nodes.

Graph Attention Network

Graph Attention Networks, developed by Veličković et al. [71], introduce the attention mechanism to the architecture of graph neural networks. The main contribution of GATs is the use of the attention mechanism in the propagation step, which allows the model to learn the relative weights in information aggregation from the node neighborhood.

The attention coefficients are calculated as follows:

$$\alpha_{ij} = \text{softmax}_j \left(a \left(W^k \cdot h_i^{k-1}, W^k \cdot h_j^{k-1} \right) \right), \quad (1.6)$$

where α_{ij} is the attention coefficient that signifies the importance of node j 's features to node i , W^k is a shared linear transformation applied to every node at layer k , and a is a shared attentional mechanism. The softmax function is used across all neighbours j to make the coefficients comparable.

The GAT layer specifically computes the hidden representations of each node by attending over its neighbors, following the formula:

$$h_i^k = \sigma \left(\sum_{j \in \mathcal{N}_i} \alpha_{ij} \cdot W^k \cdot h_j^{k-1} \right), \quad (1.7)$$

where h_i^k is the updated node representation, α_{ij} is the attention coefficient that signifies the importance of node j 's features to node i , W^k is a shared linear transformation for nodes at layer k , and σ denotes the nonlinear activation function.

To stabilize the learning process, GATs often employ multi-head attention mechanism defined as:

$$h_i^k = \sigma \left(\frac{1}{K} \sum_{k=1}^K \sum_{j \in \mathcal{N}_i} \alpha_{ij}^k \cdot W^k \cdot h_j^{k-1} \right), \quad (1.8)$$

where K is the number of independent attention heads, α_{ij}^k are the attention coefficients computed by the k -th attention mechanism, and W^k denotes the corresponding weight matrices.

Attention mechanisms allow GATs to focus on the most critical parts of the input graph, improving model interpretability and performance. GATs are particularly useful in node classification, graph classification, and other tasks where understanding the relational importance of nodes significantly improves performance.

1.3.6 Applications in EEG Analysis

GNNs have emerged as powerful tools for EEG signal analysis by treating EEG signals as vertices within a graph. This approach enables them to capture the implicit topological and functional relationships between signals, proving particularly valuable in the examination of epileptic networks.

GNNs have been used for the **detection of spikes** in scalp EEG [75] and for the **prediction of seizures** from intracranial EEG [76]. Several studies have also applied GNNs for the task of **seizure detection** [77, 78].

To note a few examples of the seizure detection models, a self-supervised Recurrent Graph Neural Network proposed for seizure detection and classification introduced two EEG graph structures capturing electrode geometry and dynamic brain connectivity [79]. Moreover, the use of Focal loss for GNNs was proposed as an effective tool to tackle data class imbalance in seizure detection [80] and a Graph Attention Network leveraging electrode and functional connectivity features was used to explore graph representations of intracranial EEG data [81].

Notably, Grattarola et al. [78] designed a Graph Attention Network for seizure detection and localization of the SOZ from intracranial data. Their approach, illustrated in Figure 1.11, involved generating undirected functional networks during interictal and ictal phases, employing two connectivity metrics – Pearson’s correlation coefficient and the phase-locking value. Subsequently, they applied a message-passing GNN equipped with an attention mechanism for seizure localization. The attention scores generated by the GNN were then analyzed to predict the SOZ, identifying nodes with the highest attention scores as points of high interest in the functional networks.

GNNs also found applications in other areas of EEG analysis, including:




Fig. 1.11: Schematic of the GNN-based pipeline for seizure detection and localization proposed by Grattarola et al. A functional network is computed from raw intracranial EEG data and used as input for the GNN model, composed of an edge-aware message passing operation followed by an attention-based readout. The Graph-level embeddings are classified to detect seizures and attention scores are analyzed to localize the seizures. Adapted from [78].

- **Sleep Staging:** Sleep staging involves categorizing sleep into distinct stages. An Adaptive Spatial-Temporal Graph Convolutional Network framework outperformed eight other baseline methods for sleep stage classification [14]. The model used an adaptive sleep graph learning mechanism to dynamically learn the graph structure from EEG data and spatial-temporal graph convolution with spatial-temporal attention to extract features in the spatial domain.
- **Brain Motor Imagery:** Brain motor imagery is used in brain-computer interface technology for patients with motor disorders. A Graph Convolutional Neural Network managed to precisely decode relevant features from EEG in various motor imagery tasks by leveraging the functional topological relationships of the electrodes [82]. Additionally, an Adaptive Spatiotemporal Graph Convolutional Network outperformed existing models by simultaneously considering EEG signal characteristics in the time domain and channel correlations in the spatial domain [83].
- **Emotion Recognition:** For emotion classification, a Dynamic Graph Convolutional Neural Network achieved state-of-the-art performance by dynamically learning intrinsic relationships between EEG channels [15]. Furthermore, a Graph Convolutional Broad Network, using stacked regular convolutional layers, extracted deep-level information from EEG for superior emotion recognition [84].

2 Clinical Dataset

For the practical part of the master’s thesis, a clinical dataset consisting of adult patient data from St. Anne’s University Hospital in Brno (SAUH) and the Montreal Neurological Institute & Hospital (MNI) was used. In Section 2.1, the patient eligibility criteria and the final patient cohort are described, and in Section 2.2, the available patient data is characterized.

2.1 Patients

All consecutive adult patients with drug-resistant focal epilepsy who underwent SEEG and subsequent resective surgery at SAUH between 3/2012 and 3/2022 and the MNI between 1/2010 and 12/2015 were analyzed for the study. The inclusion criteria for the patient cohort were: (i) availability of high-resolution 3D MRI datasets; (ii) availability of 24 hours of continuous SEEG recording for sleep staging; (iii) availability of scalp EEG, electro-oculography, and electromyography or video for sleep staging; and (iv) the availability of a postsurgical outcome after a minimum follow-up period of 1 year. The final patient cohort consisted of 37 patients, 16 patients from SAUH, and 21 patients from the MNI.

The study was approved by Brno Epilepsy Center - St. Anne’s University Hospital Research Ethics Committee and the Montreal Neurological Institute & Hospital Ethics Review Board. All patients granted written informed consent.

2.2 Patient Data

The clinical dataset consisted of raw SEEG recordings (Subsection 2.2.1), sleep staging annotations (Subsection 2.2.2), electrode contact annotations (Subsection 2.2.3), and patient surgical outcomes (Subsection 2.2.4).

2.2.1 SEEG Recordings

Patients participating in the study underwent stereotactic depth electrode implantation as part of their presurgical assessment for drug-resistant focal epilepsy treatment.

At SAUH, platinum depth electrodes from either DIXI (DIXI Medical; Chaudfontaine, France) or ALCIS (ALCIS, Besançon, France) were utilized. The DIXI electrodes had a diameter of 2 mm, contact spacing of 3.5 mm, and 15-18 contacts per electrode. The ALCIS electrodes featured a 0.8 mm diameter, contact spacing of 1.5 mm, a contact surface area of 5 mm², and contact length of 2 mm. Each patient

received 3–14 intracerebral electrodes, with varying numbers of individual contacts (5, 8, 10, or 15), targeting the temporal lobe and optionally other brain lobes using the Talairach stereotaxic system. A 192-channel research EEG acquisition system (M&I; Brainscope, Czech Republic) recorded the signals, with the average of all intracranial signals used as a recording reference. The sampling rate was 25 kHz, and the dynamic range was ± 25 mV with 10 nV (24 bits). Acquired EEGs underwent low-pass filtering and downsampling to 5 kHz for subsequent storage and analysis.

At the MNI, the SEEG was recorded with standard stereotactically inserted clinical SEEG depth electrodes compatible with MRI. Either commercially available DIXI electrodes, or depth electrodes manufactured onsite (diameter 0.5-1 mm, contact spacing 5 mm, 9 contacts, and contact surface 0.8 mm²) at the MNI were implanted using an image-guidance system (SSN Neuronavigation System, Mississauga, Ontario, Canada). The SEEG was recorded using the Harmonie monitoring system (Stellate, Montreal, Quebec, Canada), low-pass filtered at 500 Hz and sampled at 2 kHz.

2.2.2 Sleep Staging

Sleep staging involves categorizing sleep into five stages, as established by the American Academy of Sleep Medicine, which serves as the standardized classification system developed by sleep researchers. These stages include wakefulness (W), three phases of non-rapid eye movement (NREM) sleep—N1, N2, and N3—and rapid eye movement sleep (REM).

At both institutions, visual sleep staging was conducted in 30-second epochs, as outlined in the AASM guidelines [85]. At SAUH, the process involved using three scalp electrodes covering the frontal, central, and parietal regions, in addition to electrodes for electrooculography and electromyography of the chin. Meanwhile, at the MNI, scalp EEG for sleep staging utilized subdermal thin wire electrodes placed at F3, Fz, F4, C3, Cz, C4, P3, Pz, and P4 [86].

The data provided for each patient included annotations marking 5-minute segments of artifact-free recordings at least 1 hour away from seizures across various sleep stages based on the sleep staging. The annotations were used to select 30 minutes of artifact-free NREM sleep recordings for analysis.

2.2.3 Electrode Contact Annotations

The **seizure-onset electrode contacts** were determined based on the earliest changes at seizure onset irrespective of the fast activity content by board-certified epileptologists [87]. Pre- and post-resection MRI was used to identify areas of the brain and marking of **resected electrode contacts**. The localization of electrode

contacts was done using the SEEGAtlas plugin of the open-source software IBIS [88].

2.2.4 Surgical Outcomes

The classification of postsurgical outcomes is necessary to confirm that the EZ was successfully removed or disconnected during epilepsy surgery, as outlined in Subsection 1.1.4. In this study, surgical outcomes with respect to epileptic seizures were classified according to Engel’s classification [89] described in Table 2.1. For this study, patients with an Engel IA to ID outcome were considered as **good surgical outcome** patients ($N = 19$), and the remaining patients with Engel outcome II to IV as **poor surgical outcome** patients ($N = 17$).

Table 2.1: Engel’s classification of postoperative outcomes.

Class	Description
Class I	A. Completely seizure-free since surgery B. Nondisabling simple partial seizures only since surgery C. Some disabling seizures after surgery, but free of disabling seizures for at least 2 years D. Generalized convulsions with AED discontinuation only
Class II	A. Initially free of disabling seizures but has rare seizures now B. Rare disabling seizures since surgery C. More than rare disabling seizures since surgery, but rare seizures for the last 2 years D. Nocturnal seizures only
Class III	A. Worthwhile seizure reduction B. Prolonged seizure-free intervals amounting to greater than half the followed-up period, but not < 2 years
Class IV	A. Significant seizure reduction B. No appreciable change C. Seizures worse

3 Methods

In the practical part of the master’s thesis, three GNN models - GraphSAGE, GCN, and GAT - for the localization of the EZ from interictal SEEG data were implemented and tested along with two reference models - a Support Vector Machine (SVM) and a single-feature Spike Rate model. Models were then also used to predict postsurgical outcomes. The pipeline, illustrated in Figure 3.1, was implemented in Python, utilizing the PyTorch Geometric (PyG) library for model development, training, and testing. The data was first pre-processed as described in Section 3.1 and then used as inputs for the GNN models described in Section 3.2 and reference models described in Section 3.3. Section 3.4 details the tuning of model hyperparameters, Section 3.5 the testing and evaluation of EZ localization, and Section 3.6 the methodology and evaluation of outcome prediction.

The localization methodology builds on the multi-feature SVM approach for EZ localization. Previous research [13] demonstrated the potential of combining event-based features, such as spike rate, with connectivity features like relative entropy for distinguishing the EZ in interictal SEEG. What sets this study apart is the integration of GNNs, which have the potential to enhance the multi-feature approach by effectively leveraging topological information within the data. The capability of GNNs to process SEEG data as graphs introduces a novel perspective, potentially uncovering patterns that traditional methods might struggle to discern.

Additionally, the inspiration for representing SEEG data as graphs comes from a study by Grattarola et al. [78], which analyzed ictal recordings based on functional connectivity measures. To the best of our knowledge, there exists no prior research published on EZ localization using GNNs exclusively from interictal data.

3.1 Data Pre-Processing

The localization target was defined as described in 3.1.1. The raw EEG was first pre-processed, and features were extracted from the pre-processed signals according to Subsection 3.1.2. The data was then represented in a graph structure suitable for processing by the GNN model, described in Subsection 3.1.3.

3.1.1 Localization Target Definition

The localization target for training and testing of the model was defined, based on one of the most common approximations of the EZ, as SOZ contacts removed during surgery ('SOZ&Resected') [90, 91]. The definition of the target is illustrated in Figure 3.2. For the training of the model, only patients with a good postsurgical

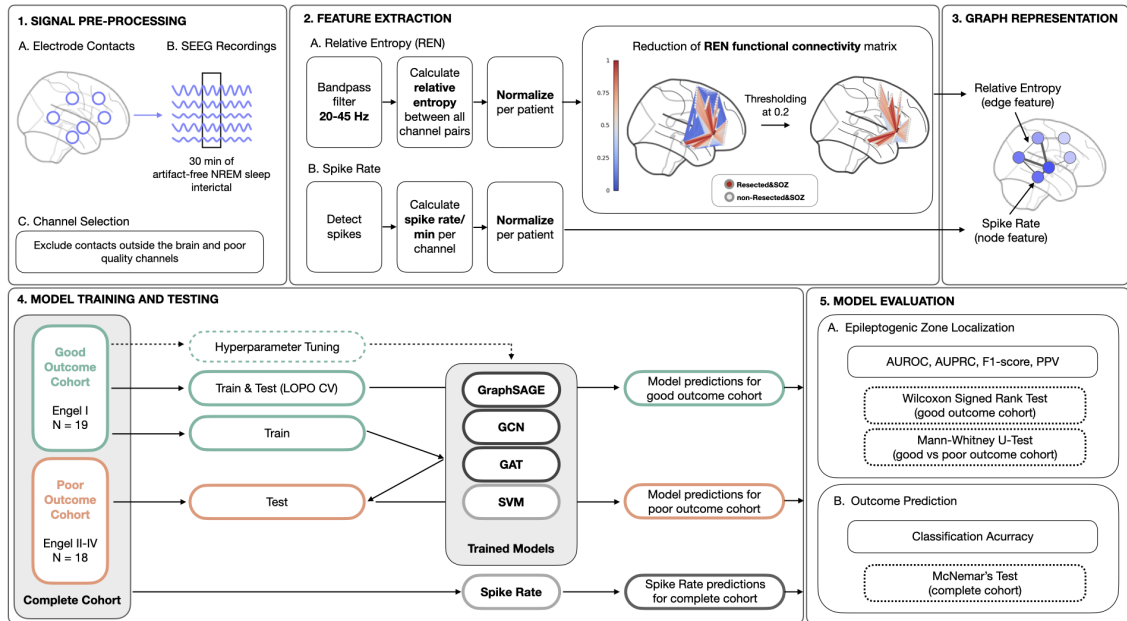


Fig. 3.1: Schematics of the implemented pipeline. The SEEG recordings were first pre-processed (1) and features (relative entropy and spike rate) were extracted (2). Next, the data of each patient was represented in a separate graph structure (3). The training and testing scheme is visualized in (4). First, the good outcome cohort was used to tune the model hyperparameters of the trained models (GraphSAGE, GCN, GAT, and SVM). Next, the good outcome cohort was tested in a leave-one-patient-out cross-validation (LOPO CV) scheme. The poor-outcome patients were tested on models trained on all good-outcome patients. The entire patient cohort was then used to produce results for the Spike Rate model. Lastly, model predictions were used to evaluate the localization of the epileptogenic zone (EZ) and the prediction of a postsurgical outcome (5).

outcome were used to ensure that the approximation of the EZ was as accurate as possible. Subsequently, the model was tested on patients with all outcomes, and the results for good and poor outcome patients were compared to evaluate the clinical utility of the model.

For the group of good outcome patients, a median of 4 electrode contacts out of 93 belonged to the 'SOZ&Resected' localization target, indicating a class imbalance in the dataset. For poor outcome patients a median of 3 electrode contacts out of 80 contacts belong to the localization target, indicating a smaller overlap between the SOZ and resected contacts compared to the good outcome patients. The complete statistics are presented in Table 3.1.




Fig. 3.2: The 'SOZ&Resected' localization target is defined as an intersection of seizure-onset zone (SOZ) electrode contacts and electrode contacts removed during surgery out of all electrode contacts.

Table 3.1: Median (IQR) counts of seizure-onset zone (SOZ), resected and 'SOZ&Resected' electrode contacts for good and poor outcome patient groups.

	SOZ contacts	Resected contacts	SOZ&Resected contacts	All contacts
Good Outcome	4.0 (6.0)	11.0 (13.0)	4.0 (2.5)	93.0 (60.0)
Poor Outcome	5.0 (5.3)	11.0 (10.8)	3.0 (1.8)	80.0 (84.3)

3.1.2 EEG Signal Pre-Processing

As a first step of the EEG signal preprocessing, segments corresponding to 30 minutes of NREM sleep were selected from the complete recording based on annotations. Nonoverlapping recordings of 15 minutes of N2 sleep, 10 minutes of N3 sleep, and 5 minutes of random NREM sleep were analyzed for each patient. NREM sleep and especially N3 sleep have been identified as the states that best identify the EZ [91, 92]. The selection of channels for analysis was performed with the exclusion of the following: (i) non-SEEG channels, (ii) SEEG channels from contacts with locations outside of the brain, and (iii) channels of poor recording quality based on visual inspection and automated detection [93]. White matter contacts were included in the analysis. An average from all SEEG contacts with a confirmed location inside the brain was used as an SEEG reference and subtracted from each SEEG signal to suppress far-field potentials caused predominantly by volume conduction.

Subsequently, the spikes in the signals were detected, and the measure of relative entropy in the low gamma band (20 – 45 Hz) was calculated between all

channel pairs for each patient. The biomarkers were selected based on the results of previous research, where they have proven effective for EZ localization [12, 13]. ElectroPhYsiology COmputational Module (EPYCOM) [94], an open-source Python library developed by the Biomedical Engineering Department of International Clinical Research Center at St. Anne’s University Hospital in Brno, was used for the extraction of features from SEEG signals.

Spike Detection

The Janca detector implemented in the EPYCOM library, was used for the detection of spikes. The detection scheme is summarized in Figure 3.3. The detector, which distinguishes the spikes from the background activity according to signal envelope modeling, achieved similar results to human reviewers in terms of spike rate. The Janca detector allows precise spike time indexing thanks to operating in the time domain, it is robust to the non-stationarity of the brain signals, and it minimizes the number of false positive detections arising from long-term interference. A limitation of the algorithm is that it is not resistant to technical artifacts, which can, however, be addressed in postprocessing. [95]




Fig. 3.3: Schematics of the Janca detector’s algorithm. Adapted from [95].

First, the signal is zero-phase filtered in the 10 – 60 Hz band by high-pass and low-pass 8th-order type II Chebyshev digital filters, and the instant envelope of the signal is calculated using the Hilbert transform. Spikes manifest in the 10 – 60 Hz band as peaks in the envelope. Next, the signal envelope is analyzed with a sliding window, statistical distribution of the envelope is calculated for each segment and approximated using the maximal likelihood algorithm (MLE). The log-normal model was the best-fitting statistical model, and the mode and median of the log-normal model were empirically chosen to discriminate spikes from the background activity. The distribution of these two parameters was used to define a threshold based on which segments with and without spikes were distinguished. A threshold curve is created by interpolating between threshold values of all segments using a cubic spline algorithm. Finally, the spikes are detected as the local maxima at the intersections of the threshold curve with the envelope, and events closer than 120 ms are merged. [95]

Spikes were detected in the 30-minute recordings, and the number of spikes per minute was calculated. The spike rates were normalized to fall between 0 and 1 for each patient.

Calculation of Relative Entropy

Relative entropy (REN), also called the Kullback-Leibler divergence, was used as a measure of functional connectivity. As a measure of how the entropy of one signal diverges from a second, REN evaluates the randomness and spectral richness between two time-series signals. The value of REN varies between 0 and ∞ , with $REN = 0$ indicating identical statistical distributions of the analyzed signals, while $REN > 0$ indicates that the two signals carry different information. REN is calculated between signals X and Y as

$$REN = \sum_{i=1}^n p(X_i) \cdot \log \left(\frac{p(X_i)}{p(Y_i)} \right), \quad (3.1)$$

where $p(X)$ and $p(Y)$ are the probability distribution functions of X and Y . Due to REN's asymmetrical properties, $REN(X, Y)$ is not equal to $REN(Y, X)$. REN was therefore calculated in two steps for both directions, and the maximum value of REN was considered the final result.

To calculate REN in the low gamma band (20 – 45 Hz), the signals were filtered using a 3rd-order Butterworth bandpass filter. REN was calculated in 1-second windows with no overlap, and the median REN value for each segment was considered. A mean REN value across segments (total of 30 minutes) was used as the final result. REN values were normalized to fall between 0 and 1, and a functional connectivity matrix was constructed for each patient.

For further computations, the functional network was sparsified. REN values below the empirically set threshold of 0.2 were set to zero, meaning that the weakest connections for each patient were discarded. Threshold values between 0.1 and 0.5 were iteratively tested, and the graph representations with connectivity matrices thresholded at 0.2 of REN were used for analysis since they maximized model results. As an example, the reduction of the functional connectivity network is visualized in Figure 3.4, and the thresholded functional connectivity matrix for patient 723 with a good surgical outcome (Engel IA) is shown in Figure 3.5. Using a sparse graph as an input for the GNN helps to prevent over-smoothing, which occurs in densely connected graphs and leads to all nodes having similar embeddings. A previous study adopted a similar approach, where the non-normalized connectivity values were thresholded at 0.1 [78].

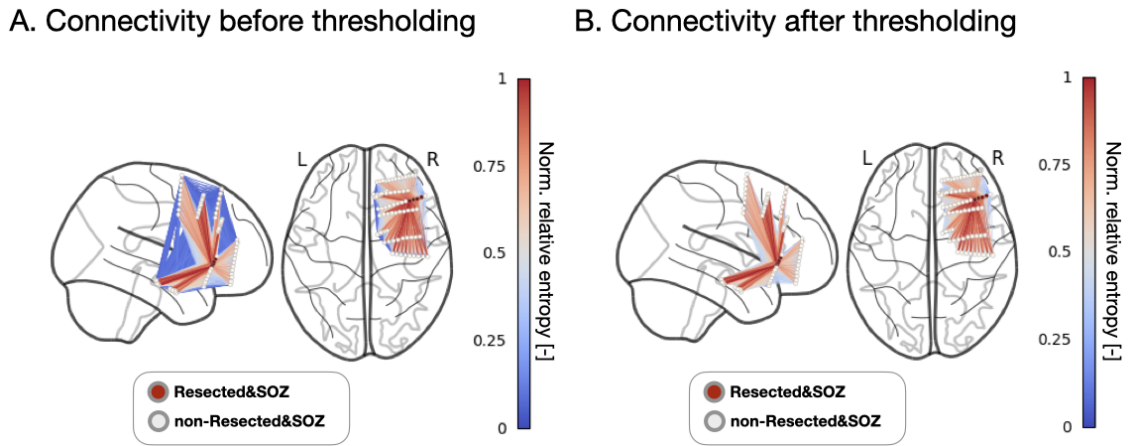


Fig. 3.4: Connectivity matrix for patient 723 with good surgical outcome (Engel IA) before (A) and after (B) thresholding. Electrode contacts are projected onto a standard brain model. Target contacts ('SOZ&Resected') are visualized in red and non-target contacts in white color. The color of the edges corresponds to the normalized relative entropy value.

3.1.3 Graph Representation of Data

The graph data was represented as 'Data' objects implemented within PyG to describe homogenous graphs. The Data object mimics the behavior of a Python dictionary and can contain node-level, edge-level, and graph-level attributes. A 'Data' object was constructed for each patient with spike rate as the node feature, REN as the edge feature, and patient ID as a graph attribute. According to the classification described in Subsection 1.3.2, the graphs are undirected, homogenous, and static.

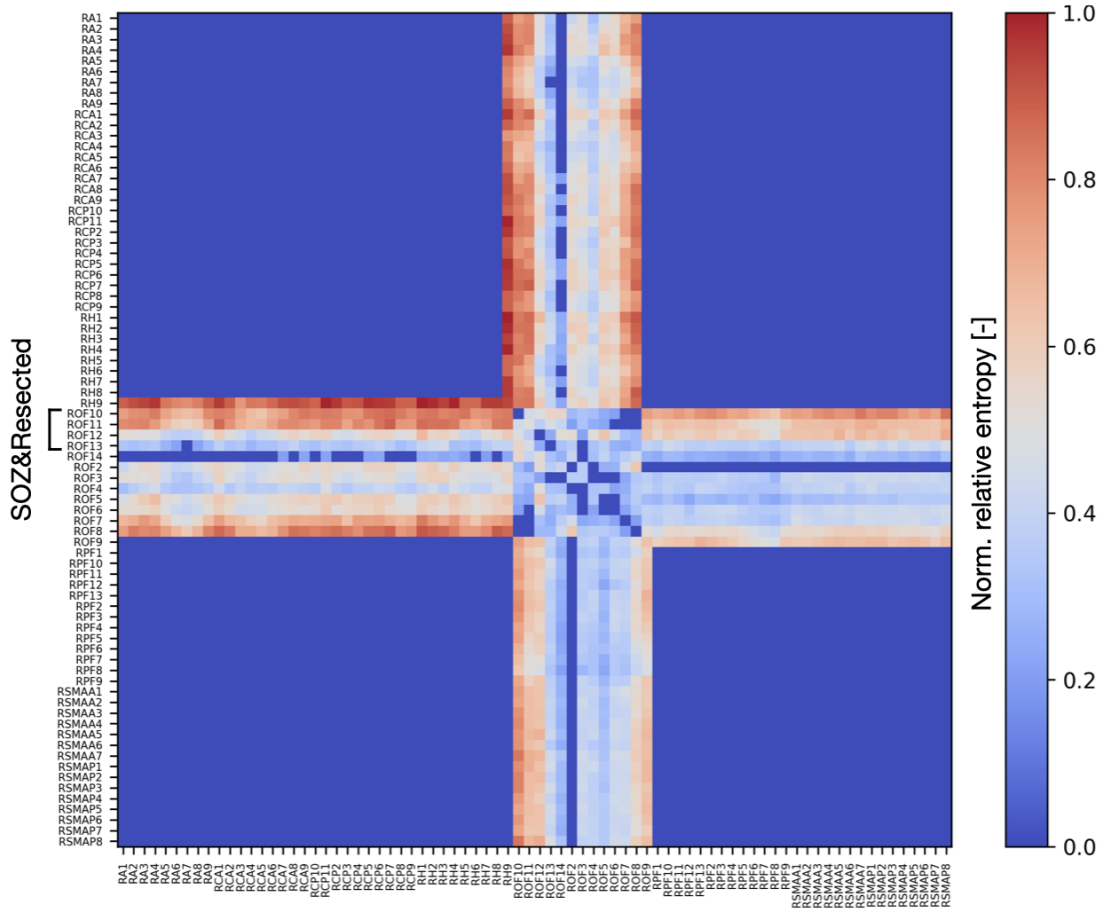


Fig. 3.5: Connectivity matrix constructed from the thresholded normalized relative entropy values of patient 723. On the x and y-axis, SEEG channels are listed with the target ('SOZ&Resected') channels highlighted on the x-axis. The color of each field corresponds to the normalized relative entropy value between the two channels according to the color bar.

Attribute Encoding

First, unique channel names were mapped to integer indices. The spike rate was encoded into a **node feature tensor** X of shape $[N, 1]$, where N is the number of nodes (i.e., channels), and X_i corresponds to the spike rate for the i -th channel. Next, the functional connectivity matrix was encoded into an **edge index tensor** E_{idx} and an **edge attribute tensor** E_{attr} . The edge index tensor E_{idx} of shape $[2, N^e]$, where N^e is the number of edges in the graph, contains the indices of all node pairs that are connected with an edge. The source and destination channel names were mapped to indices using the previously generated node mappings. The edge attribute tensor E_{attr} of shape $[N^e, 1]$ contains the functional connectivity values for

each of the existing edges, with E_{attr_i} corresponding to the weight of the i -th edge.

Ground truth labels for nodes were encoded in a **label tensor** Y of shape $[N, 1]$, where N is the number of nodes (i.e., channels), and Y_i corresponds to the ground truth for the i -th channel. According to the definition of the localization target in Subsection 3.1.1, binary labels were assigned to each channel, corresponding to 1 for channels marked as 'SOZ&Resected' and 0 for channels outside the 'SOZ&Resected'.

Subsequently, the positions of nodes were encoded into a **node position tensor** of shape $[N, 3]$. The position tensor used for visualization contained the coordinates of all nodes on the x , y , and z axis in the MNI template coordinate space. As a graph attribute, the postsurgical **outcome** was encoded into the graph with the value of 0 corresponding to a poor outcome and a value of 1 to a good outcome. Finally, the **patient ID** was encoded as a graph-level attribute.

Dataset Construction

For each patient in the dataset, a 'Data' structure is created to represent the graph, summarized in Table 3.2. Node features, edge indices, edge attributes, ground truth labels, contact positions, and patient ID are incorporated into the 'Data' object. The data structures are stored in a list and later loaded into a 'DataLoader' object for training and evaluation.

Table 3.2: Data structure shapes for patients in the dataset. X refers to the node feature tensor, E_{idx} to the edge index tensor, E_{attr} to the edge attribute tensor, Y to the label tensor, and 'Pos' to the node position tensor. 'Outcome' value of 0 refers to a poor postsurgical outcome and 1 to a good postsurgical outcome.

Patient ID	X	E_{idx}	E_{attr}	Y	Pos	Outcome
59	(121, 1)	(2, 7260)	(7260, 1)	(121)	(121, 3)	0
61	(149, 1)	(2, 11026)	(11026, 1)	(149)	(149, 3)	1
63	(120, 1)	(2, 7140)	(7140, 1)	(120)	(120, 3)	1
66	(170, 1)	(2, 14365)	(14365, 1)	(170)	(170, 3)	0
67	(153, 1)	(2, 11628)	(11628, 1)	(153)	(153, 3)	0
69	(169, 1)	(2, 14196)	(14196, 1)	(169)	(169, 3)	0
71	(155, 1)	(2, 11935)	(11935, 1)	(155)	(155, 3)	1
73	(152, 1)	(2, 11476)	(11476, 1)	(152)	(152, 3)	0
74	(174, 1)	(2, 15051)	(15051, 1)	(174)	(174, 3)	1
77	(162, 1)	(2, 13041)	(13041, 1)	(162)	(162, 3)	1
80	(142, 1)	(2, 10011)	(10011, 1)	(142)	(142, 3)	0
82	(83, 1)	(2, 3403)	(3403, 1)	(83)	(83, 3)	1

Continued on the next page

Continued from the previous page

Patient ID	X	E_{idx}	E_{attr}	Y	Pos	Outcome
81	(164, 1)	(2, 13366)	(13366, 1)	(164)	(164, 3)	0
84	(95, 1)	(2, 4465)	(4465, 1)	(95)	(95, 3)	1
88	(39, 1)	(2, 741)	(741, 1)	(39)	(39, 3)	0
90	(80, 1)	(2, 3160)	(3160, 1)	(80)	(80, 3)	0
92	(174, 1)	(2, 15051)	(15051, 1)	(174)	(174, 3)	0
93	(156, 1)	(2, 12090)	(12090, 1)	(156)	(156, 3)	1
583	(79, 1)	(2, 3081)	(3081, 1)	(79)	(79, 3)	1
657	(44, 1)	(2, 946)	(946, 1)	(44)	(44, 3)	1
705	(76, 1)	(2, 2850)	(2850, 1)	(76)	(76, 3)	0
717	(61, 1)	(2, 1830)	(1830, 1)	(61)	(61, 3)	1
723	(79, 1)	(2, 3081)	(3081, 1)	(79)	(79, 3)	1
777	(28, 1)	(2, 378)	(378, 1)	(28)	(28, 3)	0
909	(63, 1)	(2, 1953)	(1953, 1)	(63)	(63, 3)	0
953	(69, 1)	(2, 2346)	(2346, 1)	(69)	(69, 3)	1
965	(83, 1)	(2, 3403)	(3403, 1)	(83)	(83, 3)	1
1029	(91, 1)	(2, 4095)	(4095, 1)	(91)	(91, 3)	0
1041	(101, 1)	(2, 5050)	(5050, 1)	(101)	(101, 3)	1
1043	(70, 1)	(2, 2415)	(2415, 1)	(70)	(70, 3)	1
1149	(48, 1)	(2, 1128)	(1128, 1)	(48)	(48, 3)	0
1153	(93, 1)	(2, 4278)	(4278, 1)	(93)	(93, 3)	1
1159	(72, 1)	(2, 2556)	(2556, 1)	(72)	(72, 3)	0
1162	(33, 1)	(2, 528)	(528, 1)	(33)	(33, 3)	0
1233	(45, 1)	(2, 990)	(990, 1)	(45)	(45, 3)	1
1246	(103, 1)	(2, 5253)	(5253, 1)	(103)	(103, 3)	1
1247	(79, 1)	(2, 3081)	(3081, 1)	(79)	(79, 3)	0
1299	(72, 1)	(2, 2556)	(2556, 1)	(72)	(72, 3)	0
1446	(50, 1)	(2, 1225)	(1225, 1)	(50)	(50, 3)	0

3.2 GNN Models

The architecture of all GNN models was designed for a node classification task. Each of the models consisted of a sequence of GNN layers (either GraphSAGE, GCN, or GAT) followed by a linear classifier. The **GNN layers** take node features, edge indices, and edge weights as input and produce node embeddings. The transformed node embeddings are then fed into a **linear classifier** with a **sigmoid activation function** for binary classification. The output of the linear classifier is a score

between 0 and 1, indicating the likelihood of each node belonging to a particular class.

The implementation of **GraphSAGE**, **GCN** and **GAT** architectures is described in Subsections 3.2.1, 3.2.2, and 3.2.3, respectively. The linear classifier is described in Subsection 3.2.4. Subsection 3.2.5 then describes the training of the GNN models. The architecture is illustrated in Figure 3.6.

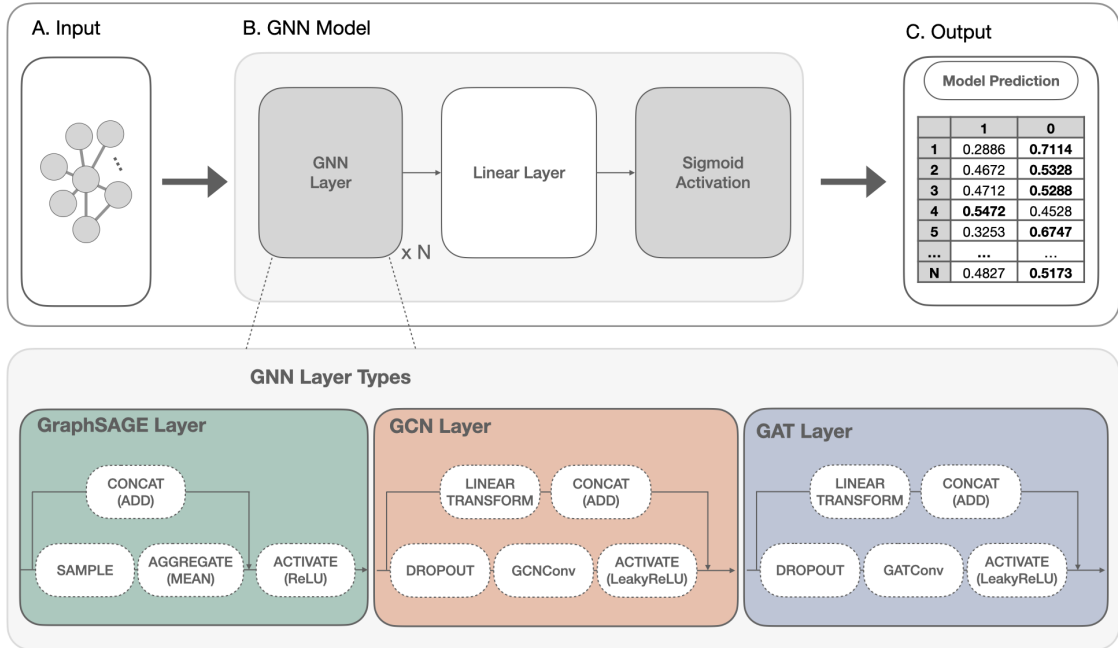


Fig. 3.6: Implementation of the GNN models. Input graph structure (A) is transformed by the GNN model (B), which contains a series of N GNN layers (GraphSAGE, GCN, or GAT) and a linear classifier with a sigmoid activation. The outputs of the model (C) are scores between 0 and 1 for each node of the graph, indicating the probability that the node is an SOZ&Resected electrode contact.

3.2.1 GraphSAGE Network

The GraphSAGE layer uses the **SAGEConv operator** for message passing with the sampling, aggregation, and concatenation steps performed as follows:

$$h_i^k = \sigma \left(W_1 \times h_i^{k-1} + W_2 \times \text{MEAN} \left(\{h_j^{k-1}\} \cup \{h_j^{k-1}, \forall j \in N_j^{k-1}\} \right) \right), \quad (3.2)$$

where h_i^k are the new embeddings for node i , h_i^{k-1} are the embeddings for node i from a previous layer, and $\{h_j^{k-1}, \forall j \in N_j^{k-1}\}$ are the node embeddings for the neighborhood N_j^{k-1} of node i . The weights W_1 and W_2 are learned during training, and σ is a non-linear activation function.

The **mean operator** was used to aggregate each node h_i^k . For every node h_i^k , it computes the mean of vectors $\{h_j^{k-1}, \forall j \in N_j^{k-1}\}$. The mean operator was chosen as the aggregator function since it performed better than other aggregation functions like max pooling and long short-term memory aggregators. The output was then transformed with a **ReLU activation function**, which introduces non-linearity by setting negative values to zero and keeping positive values unchanged. The output of the GraphSAGE layer was passed to the next GraphSAGE layer or directly to the linear classifier described in Subsection 3.2.4.

The number of GraphSAGE number message-passing layers and of hidden channels in each GraphSAGE layer, which controls the dimensionality of the node embeddings during the message-passing process, were tuned as described in Section 3.4.

3.2.2 Graph Convolutional Network

The GCN model employs a series of GCN layers that utilize the **GCNConv operator** for message passing. In the GCN layers, information is aggregated from neighboring nodes to the target node using graph convolution, mathematically represented as follows:

$$h_i^k = \sigma \left(\sum_{j \in N_i} \frac{e_{ij}}{\sqrt{d_i d_j}} W \cdot h_j^{k-1} \right), \quad (3.3)$$

where h_i^k represents the updated embedding of node i at layer k and N_i represents the neighborhood of node i consisting of nodes j . Edge weight between nodes i and j is denoted as e_{ij} , and d_i and d_j are the degrees (number of connections) of nodes i and j , respectively. The weight matrix W is learned during training, and σ is a non-linear activation function.

The **Leaky ReLU activation function** used in the GCN model allows a small non-zero gradient for negative values, helping to avoid the "dying ReLU" problem and promoting gradient flow through the graph structure. A **dropout layer** before the GCNConv operator is used to prevent overfitting by randomly setting some node embeddings to 0, and each GCN layer includes a **skip connection** that projects the input features directly to the output embedding using a separate linear projection and adds them to the activated GCNConv outputs.

The number of GCN layers, the number of output channels of each GCNConv operator, and the dropout probability were tuned as described in Section 3.4.

3.2.3 Graph Attention Network

The GAT model comprises multiple attention-based convolutional layers that utilize the **GATConv operator** for graph attention. The multi-head attention mechanism, described in Subsection 1.3.5, assigns different importance scores to each neighboring node and aggregates information based on these scores according to Equation 1.8.

The **Leaky ReLU activation function** was used as in the GCN, and **dropout** was used to prevent overfitting. **Skip connections**, same as those in the GCN layers, were used in the GAT layers, projecting linearly transformed input features to the output embedding directly.

The number of GAT layers, the number of attention heads of each GATConv operator, and the dropout probability were tuned as described in Section 3.4.

3.2.4 Linear Classifier

The linear layer with sigmoid activation serves as a binary classifier, making predictions for the node classification task based on node embeddings generated by the GNN layers. In mathematical terms, the operation performed by the linear layer with sigmoid activation can be expressed as follows:

$$\hat{y} = \sigma(W_{\text{linear}} \times h + b_{\text{linear}}), \quad (3.4)$$

where \hat{y} are the predicted probability scores, σ is the sigmoid activation function, W_{linear} represents the weight matrix for the linear layer, h denotes the node embeddings generated by the GNN layers, and b_{linear} is the bias for the linear layer. The sigmoid activation function is defined as:

$$\sigma(x) = \frac{1}{1 + e^{-x}}, \quad (3.5)$$

where $\sigma(x)$ is the sigmoid function, and e is the base of the natural logarithm.

The linear layer applies a linear transformation to the input data, and the sigmoid activation function maps the linear layer output to a range between 0 and 1, producing a probability score. This score indicates the likelihood of each node belonging to a specific class. To obtain a class assignment for each node, we can predict class 1 if the probability is greater than or equal to 0.5, and class 0 otherwise.

During model training, the sigmoid activation is applied as part of the Binary Cross-Entropy with Logits loss function, and during testing, it is added as the final layer.

3.2.5 Training

Each model was trained separately for good and poor outcome patient cohorts. First, it was trained and tested on the good outcome patient cohort with leave-one-patient-out cross-validation, meaning that it was trained on the data from all patients except one and then tested on the withheld patient’s data.

For each cross-validation fold, a new model was trained for a given number of training epochs. A suitable number of epochs for a given model architecture and optimizer learning rate was first estimated using hyperparameter optimization described in Section 3.4. Next, each of the GNN models was trained for 200 training epochs while metrics were monitored, and the number of epochs that yielded the best balance of AUROC and AUPRC across cross-validation folds for each model was used for final model training. Training of GraphSAGE, GCN, and GAT models is visualized in Figures 3.7, 3.8, and 3.9, respectively.




Fig. 3.7: Evolution of training loss, validation loss, and performance metrics (Median AUROC, Median AUPRC, and Median Mean AUROC and AUPRC) for the Graph Sample and Aggregate (GraphSAGE) model across 200 training epochs. The median mean AUROC and AUPRC is highlighted at the maximum value during Epoch 105.

In each training epoch, the data of all patients in the training loader were fed into the model for training by mini-batches of size 4 (4 patients at a time), and the model predicted the output based on the input graphs. Next, the loss was calculated using the Binary Entropy with Logits loss function and backpropagated through the

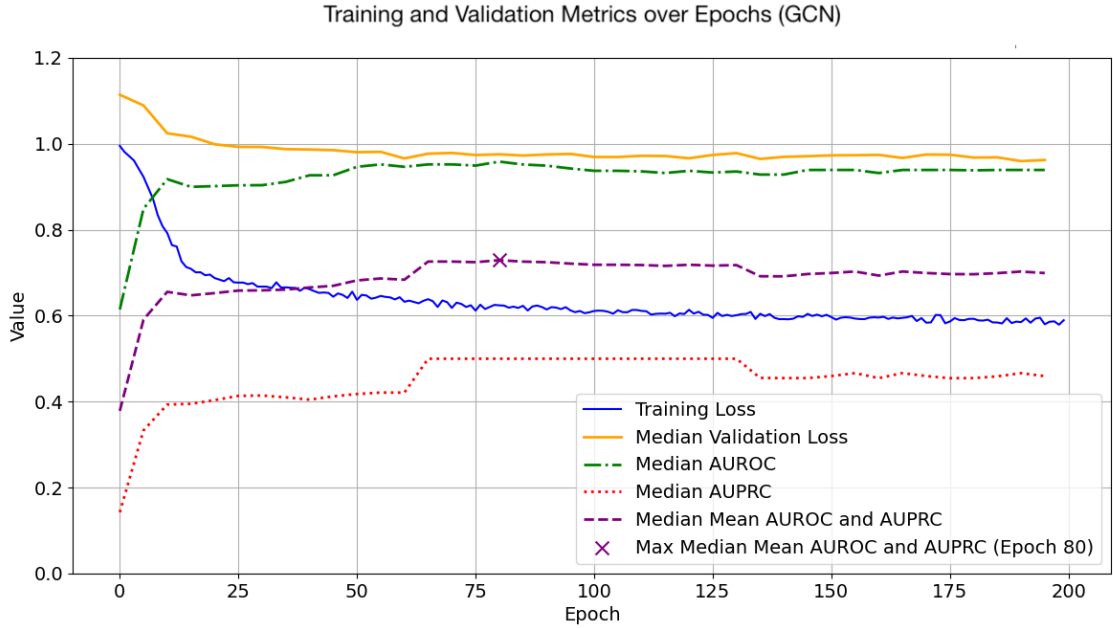


Fig. 3.8: Evolution of training loss, validation loss, and performance metrics (Median AUROC, Median AUPRC, and Median Mean AUROC and AUPRC) for the Graph Convolutional Network (GCN) model across 200 training epochs. The median mean AUROC and AUPRC is highlighted at the maximum value during Epoch 80.

model. Then, the optimizer, in this case, the Adam optimizer, updated the model's parameters to minimize the loss.

Subsequently, the model was retrained with the same hyperparameters on all good-outcome patients and tested on all poor-outcome patients.

Binary Cross Entropy with Logits

The Binary Cross Entropy with Logits (BCE with Logits) loss combines the Binary Cross Entropy (BCE) loss and the Sigmoid activation function into a single computation. The loss function measures the dissimilarity between predicted values and ground truth labels while ensuring numerical stability. The formula of BCE with Logits is:

$$\text{BCE with Logits Loss} = -\frac{1}{N} \sum_{i=1}^N [y_i \log(\sigma(z_i)) + (1 - y_i) \log(1 - \sigma(z_i))], \quad (3.6)$$

where N is the number of samples, y_i is the ground truth label, z_i is the predicted value (also called "logit"), and σ is the Sigmoid activation function.

Custom class weights for the loss were defined to address imbalances in the distribution of positive and negative samples within the dataset during the training process.

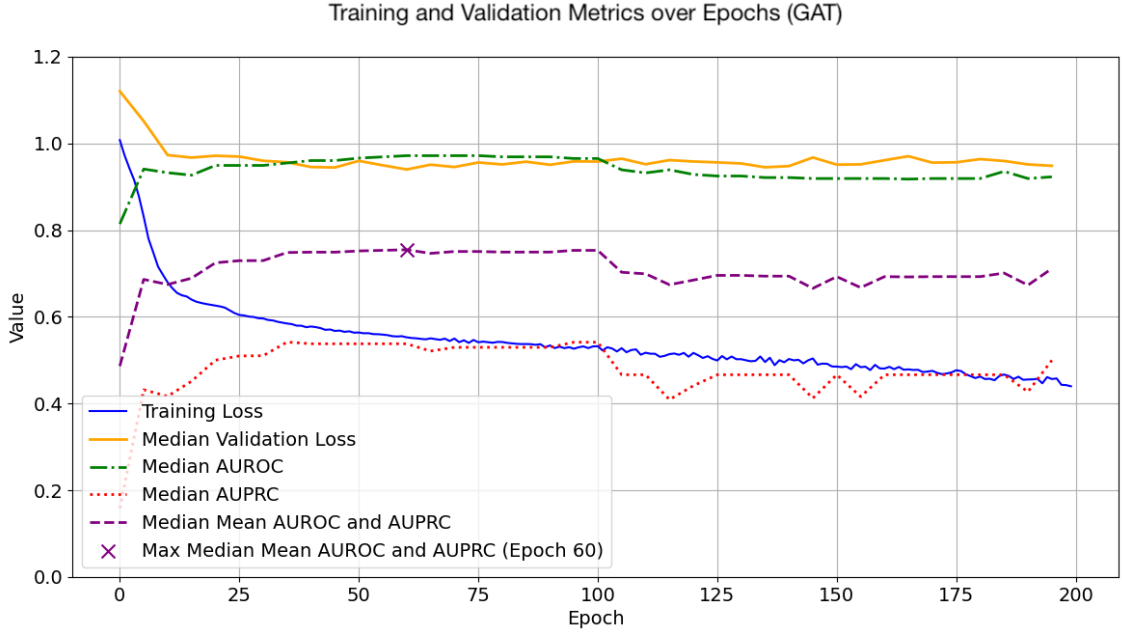


Fig. 3.9: Evolution of training loss, validation loss, and performance metrics (Median AUROC, Median AUPRC, and Median Mean AUROC and AUPRC) for the Graph Attention Network (GAT) model across 200 training epochs. The median mean AUROC and AUPRC is highlighted at the maximum value during Epoch 60.

The weights were calculated based on the numbers of positive and negative samples in the entire dataset:

$$w_{classes} = \frac{N}{n_{classes} \times \text{bincount}(y)}, \quad (3.7)$$

where $w_{classes}$ are the class weights, N is the number of all samples, $n_{classes}$ is 2 in binary classification, and $\text{bincount}(y)$ contains the counts of negative and positive samples in the ground truth label vector y .

The BCE with Logits loss was tested alongside the original BCE loss and the Focal loss, which is specifically designed for imbalanced datasets, and it produced the most favorable results. In practice, the BCE with Logits loss is more computationally efficient and numerically stable than the BCE loss due to its ability to pass the logits directly into the loss function. Compared to the Focal loss, the BCE with Logits loss also demonstrated superior performance, and significant improvement in results after the application of class weights confirms that BCE with Logits can effectively be applied to imbalanced datasets.

Adam Optimizer

The Adam optimizer [96] was employed to update the model’s trainable parameters during training. It combines adaptive learning rates for each parameter and incorporates momentum to achieve efficient optimization. The key parameters of the Adam optimizer are the learning rate (α) and the exponential decay rates β_1 and β_2 . The learning rate controls the size of the steps taken during optimization, and its value was optimized with other hyperparameters as described in Section 3.4. Values of the exponential decay rates were fixed at $\beta_1 = 0.9$ and $\beta_2 = 0.999$. Additionally, weight decay penalizing large weights was applied with the value of 0.0005 for regularization.

3.3 Reference Models

To evaluate the validity and utility of the proposed GNN models, two reference models were implemented and tested using the same clinical dataset. First, the raw normalized spike rate values were used (termed as the ‘Spike Rate’ model) as described in Subsection 3.3.1, and second, an SVM was trained and tested using the same features as the GNNs as described in Subsection 3.3.2, i.e. normalized spike rate and normalized REN.

3.3.1 Spike Rate

In the Spike Rate model, the normalized spike rate values were used the same way as scores output by the GNN models. Therefore, there was no training, and the normalized spike rate values were directly evaluated according to Subsection 3.5 for each patient in the clinical dataset.

3.3.2 Support Vector Machine

The SVM model was trained and tested using an identical scheme to the GNN models. Specifically, the SVM utilized two features: the normalized spike rate and the normalized REN. The REN values were computed between each SEEG channel and its nearest neighboring channel. They were then normalized for each patient individually to lie within the range of 0 to 1.

The hyperparameters of the SVM model, i.e. the value of a regularization parameter C and kernel type, were optimized according to Section 3.4 to match the optimization of GNN hyperparameters. The kernel coefficient γ was automatically scaled by the algorithm, and class weights were automatically balanced and set as inversely proportional to class frequencies in the data.

3.4 Hyperparameter Tuning

Hyperparameters of the GNN models and the reference SVM model were tuned using the Python library 'Tune' of an open-source framework 'Ray'. Ray Tune enables the execution of experiments and the tuning of hyperparameters, and the Population Based Training (PBT) algorithm [97] implemented in Ray Tune was used in this pipeline. Inspired by genetic algorithms, PBT combines the strengths of random search and hand-tuning while reducing computational demands. Initially, multiple neural networks are trained in parallel with random hyperparameters. Unlike conventional random search, PBT allows these networks to periodically exchange information, leveraging successful configurations discovered by other models in the population. By balancing the exploration and exploitation of hyperparameters, PBT focuses computational resources on promising models and dynamically adapts hyperparameters during training.

For all models, the hyperparameters were tuned in a leave-one-patient-out cross-validation scheme on the cohort of good-outcome patients. For each hyperparameter setting, 19 separate models (corresponding to 19 good-outcome patients) were trained and tested on the left-out patients. The median mean of AUROC and AUPRC was calculated across the cross-validation folds and reported back to Ray Tune. Ray Tune was configured to find the parameters that yielded the highest score of mean between AUROC and AUPRC to balance these two performance metrics. For each model, 50 different sets of hyperparameters were evaluated using all available computational resources (28 CPU cores and 2 GPU cores). Searched and best parameter values are reported in Table 3.4 for the GNN models, and in Table 3.5 for the SVM model.

Table 3.4: Searched parameter values and best configurations for the Graph Sample and Aggregate (GraphSAGE), Graph Convolutional Network (GCN), and Graph Attention Network (GAT). The parameter 'Hidden Channels' refers to the number of attention heads for the Graph Attention Model.

Parameter	Hidden Channels	Layers	Dropout	Learning Rate	Epochs
Searched Values	4, 8, 16, 32	2, 3, 4, 5	0.0, 0.1, 0.2, 0.4	0.01, 0.005, 0.001, 0.0005	50, 75, 100, 150
Best Config.	Hidden Channels	Layers	Dropout	Learning Rate	Epochs
GraphSAGE	32	3	0.4	0.01	75
GCN	32	5	0.4	0.001	150
GAT	32	5	0.0	0.001	50

Table 3.5: Searched parameter values and the best configuration for the Support Vector Machine (SVM) model. The value of regularization parameter C and kernel type was optimized.

Parameter	C	Kernel
Searched Values	log-uniform $[1e^{-3}, 1e^3]$	linear, poly, rbf, sigmoid
Best Configuration	0.1767	linear

3.5 Testing and Evaluation

Model testing was performed separately for the good outcome and poor outcome patient cohorts. First, the models were trained and tested for the good outcome patient cohort using a leave-one-patient-out cross-validation scheme, yielding the output data for the good outcome patients. Next, the model was trained on all good-outcome patients and tested on all poor-outcome patients to yield the output data for poor-outcome patients. The outputs were then evaluated separately for each patient and aggregated for each patient cohort, i.e., good and poor-outcome patients.

Evaluation Metrics

The true positives (TP) were defined as correctly classified SOZ&Resected contacts, true negatives (TN) as correctly classified non-SOZ&Resected contacts, false positives (FP) as non-SOZ&Resected contacts misclassified as SOZ&Resected, and false negatives (FN) as SOZ&Resected contacts misclassified as non-SOZ&Resected contacts. The following metrics can be calculated from the TP, TN, FP, and FN values:

$$\text{Sensitivity (Recall)} = \frac{\text{TP}}{\text{TP} + \text{FN}}, \quad (3.8)$$

$$\text{Specificity} = \frac{\text{TN}}{\text{TN} + \text{FP}}, \quad (3.9)$$

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \quad (3.10)$$

For a comprehensive evaluation of model performance across a range of classification thresholds, the trade-off between these metrics was evaluated using the Receiver Operating Characteristic (ROC) and the Precision-Recall (PR) curve and their respective quantifications, Area Under ROC (AUROC) and Area Under PR Curve (AUPRC).

The AUROC evaluates the trade-off between specificity and sensitivity, providing a model-wide assessment of the model’s ability to discriminate between the positive and negative class samples (i.e. SOZ&Resected and non-SOZ&Resected contacts). The AUROC was calculated using the trapezoidal rule [98] from the ROC curve, which plots the sensitivity against the false positive rate (calculated as 1 - specificity) across all classification thresholds. The chance level for AUROC is 0.5 in a binary classification task.

The AUPRC evaluates the trade-off of precision and recall for all possible thresholds. AUPRC focuses on the positive class, which makes it suitable for imbalanced datasets. The average precision method was used to estimate the AUPRC [99], constructed by plotting the precision against recall across all classification thresholds. Average precision AP is calculated as the mean precision P achieved for n thresholds, weighed by the increase in recall R from the previous threshold:

$$AP = \sum_n (R_n - R_{n-1}) \cdot P_n, \quad (3.11)$$

where R_n and P_n denote the recall and precision at the n -th threshold, respectively. The always positive classifier defines the chance level for AUPRC. The AUPRC of such a classifier equals the precision that would be achieved if all samples were classified as positive (chance level precision), corresponding to the proportion of positive samples in the dataset.

To evaluate the models’ ability to find a suitable classification threshold and assign correct class labels to SEEG contacts, the F1-score was used. F1-score is the harmonic mean of precision and recall at a specific threshold, reflecting the model’s capability to identify an optimal threshold that achieves a balanced trade-off between recall and precision, and it is calculated as:

$$F_1\text{-score} = \frac{2}{\frac{1}{\text{PREC}} + \frac{1}{\text{REC}}} = \frac{TP}{TP + \frac{1}{2}(FP + FN)} \quad (3.12)$$

When calculating the chance level for the F1-score, recall is set to 1, indicating perfect recall, and the precision is set to the chance level precision (p). The chance level F1-score is then calculated using the F1-score formula as $(2 \times p)/(p+1)$.

Statistical Analysis

The Wilcoxon Signed Rank test [100] was employed to compare the localization results between models. Specifically, the test assessed the distributions of metric values across the good outcome patient cohort, comparing each GNN model against the main reference Spike Rate model. The Wilcoxon Signed Rank test is a non-parametric statistical test for paired data that evaluates whether two related samples

originate from the same distribution, which allows direct comparisons between the performance of the GNN models and the reference model across the identical patient cohort.

To evaluate the statistical differences between the good and poor outcome patient groups, the Mann-Whitney U-test [101] was performed. Similar to the Wilcoxon Signed Rank test, the Mann-Whitney U-test is a non-parametric method, but it compares two independent groups to determine whether their distributions differ significantly, suitable for two non-overlapping patient cohorts.

In both statistical tests, results with a p-value below 0.05 were considered significant.

3.6 Outcome Prediction

The prediction of postsurgical outcome was done as in a published study by Nevalainen et al. [102]. For each patient, model outputs were used to calculate precision in identifying resected electrode contacts. Precision was calculated for each patient according to Equation 3.10, where TP was defined as correctly identified resected contacts, TN as correctly classified non-resected contacts, FP as non-resected contacts misclassified as resected, and FN as resected contacts misclassified as non-resected. The value of precision for each patient reflects how many from all contacts predictive as positive (i.e. epileptogenic) were resected during surgery. Subsequently, a good postsurgical outcome was predicted for patients with a precision value higher than 0.5 (meaning that more than 50% of predicted positive contacts were resected).

Evaluation Metrics

Outcome prediction was evaluated for each model using classification accuracy as:

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}}, \quad (3.13)$$

where TP corresponds to patients correctly classified as good-outcome patients, TN to patients correctly classified as poor-outcome patients, FP to poor-outcome patients misclassified as good-outcome patients, and FN to good-outcome patients misclassified as poor-outcome patients. Accuracy is suitable in this case since the dataset is balanced in terms of patient outcomes, with 19 good-outcome patients and 18 poor-outcome patients.

Statistical Analysis

The outcome prediction results were also evaluated statistically, using the McNemar test [103]. The McNemar test is a non-parametric test for paired data that can be

used to compare the performance of classification models. It is applied to a 2x2 contingency table which contains the counts of samples (i.e. patients): (i) whose outcomes were correctly predicted by both models, (ii) whose outcomes were correctly predicted only by model A, (iii) whose outcomes were correctly predicted only by model B, and (iv) whose outcomes were incorrectly predicted by both models. The McNemar test was performed between the three GNN models and the main reference model (Spike Rate model) to compare the clinical utility of the proposed GNN models. Differences in the McNemar test with a p-value below 0.05 were considered significant.

4 Results

This chapter presents the findings from the practical application of the GNNs in localizing the EZ in Section 4.1 and predicting surgical outcomes in Section 4.2.

4.1 Epileptogenic Zone Localization

This section details the results from the application of GNNs in localizing the EZ, with a specific focus on performance for good-outcome patients in Subsection 4.1.1 and comparison of performance for good and poor-outcome patients in Subsection 4.1.2.d

4.1.1 Performance for Good Outcome Patients

The GNN models for localization of the EZ were first evaluated on the cohort of 19 good-outcome patients. Evaluation metrics described in Subsection 3.5, namely AUROC, AUPRC, and F1-score, were calculated to evaluate the model’s ability to localize EZ contacts defined as SOZ&Resected electrode contacts. Next, the models’ ability to localize resected contacts was evaluated using precision as an indicator of clinical utility. The distribution of results across the patient cohort is visualized in Figure 4.1, and median and interquartile range (IQR) values are presented in Table 4.1. The Wilcoxon Signed Rank test was used to compare results between models.

Table 4.1: Performance of models across the cohort of 19 good outcome patients (median and IQR). Values in bold indicate that the model outperformed both reference models in terms of a given evaluation metric.

Model	AUROC	AUPRC	F1-score	PPV (Res)
GraphSAGE	0.971 (0.184)	0.563 (0.582)	0.364 (0.411)	0.500 (0.441)
GCN	0.939 (0.205)	0.500 (0.476)	0.400 (0.415)	0.429 (0.541)
GAT	0.971 (0.184)	0.525 (0.557)	0.471 (0.392)	0.714 (0.626)
SVM	0.954 (0.236)	0.355 (0.456)	0.316 (0.289)	0.429 (0.476)
Spike Rate	0.921 (0.237)	0.407 (0.476)	0.400 (0.470)	0.615 (0.646)
Chance	0.500	0.049	0.047	0.119

A) Area Under Receiver Operating Characteristic (AUROC)

The GAT and GraphSAGE models achieved the highest median AUROC scores, both at 0.971 (IQR 0.184), outperforming both reference models regarding AUROC

Epileptogenic Zone Localization Results

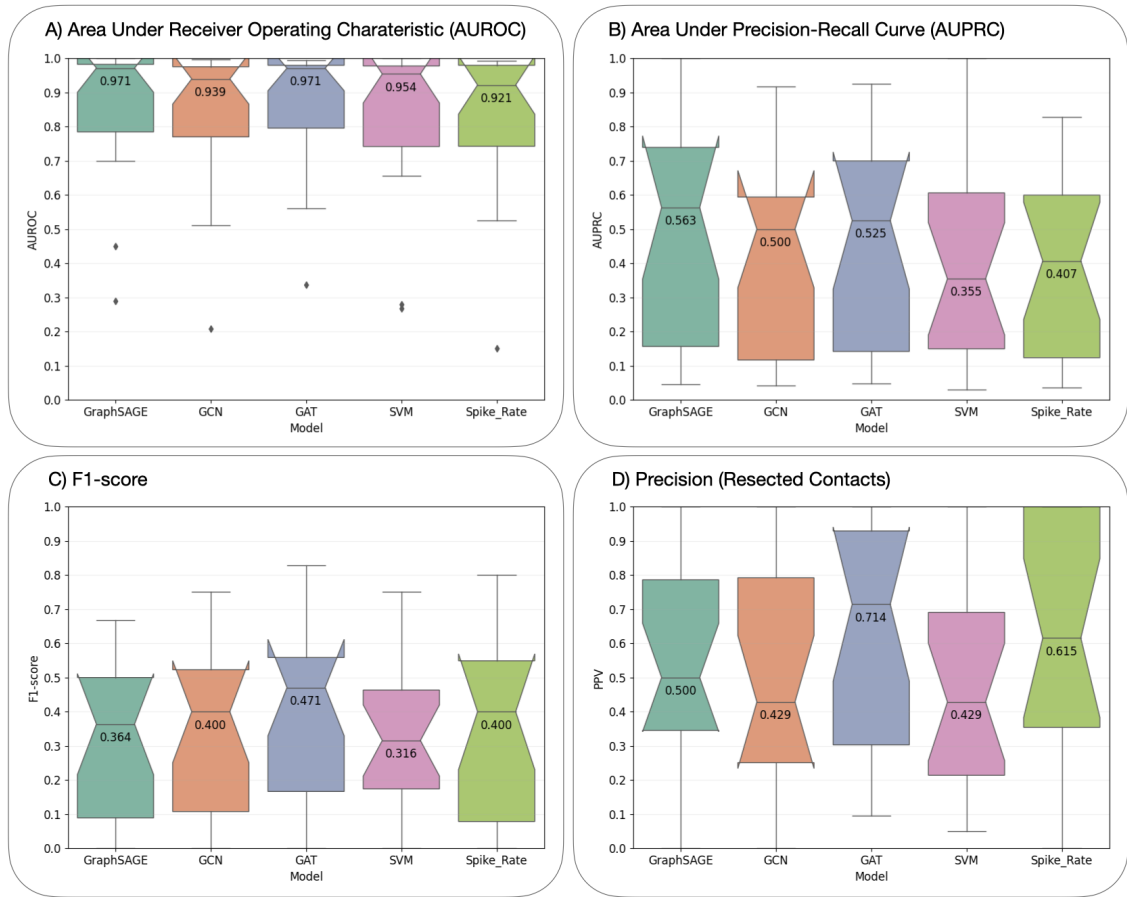


Fig. 4.1: Performance of models across the cohort of 19 good outcome patients for A) Area Under Receiver Operating Characteristic (AUROC), B) Area Under Precision-Recall Curve (AUPRC), C) F1-score, and D) precision (PPV) for resected contacts. For each distribution, the box represents the interquartile range with the horizontal line corresponding to the median and the notch to the 95% confidence interval. Whiskers extend from the box to the minimum and maximum values within a range specified by 1.5 times the interquartile range. The median value is reported below the median line, and diamond markers represent outliers.

performance. The GCN model had a median AUROC of 0.939 (IQR 0.205), surpassing only the Spike Rate reference model. The reference models, SVM and Spike Rate showed slightly lower median AUROC scores of 0.954 (IQR 0.236) and 0.921 (IQR 0.237), respectively.

In a statistical comparison of the AUROC performance of GNN models to the main reference model, Spike Rate, the GAT model showed significantly better performance ($p = 0.016$). The GraphSAGE and GCN were not significantly better ($p > 0.05$).

B) Area Under Precision-Recall Curve (AUPRC)

The GraphSAGE model achieved the highest median AUPRC score of 0.563, with an IQR of 0.582. The GAT and GCN models followed, with medians of 0.525 (IQR 0.557) and 0.500 (IQR 0.476), respectively. The Spike Rate model scored 0.407 (IQR 0.476), while the SVM model demonstrated the lowest median AUPRC of 0.355 (IQR 0.456). These results indicate that GraphSAGE consistently outperformed other models in identifying the SOZ&Resected regions with precision and recall. Moreover, all GNN models showed superior performance in terms of AUPRC compared to references. Notably, all models exhibited high IQR values across the patient cohort, indicating high variability in patient results.

In terms of AUPRC, both the GAT and the GraphSAGE significantly outperformed Spike Rate ($p = 0.017$ for both), while the comparison between GCN and Spike Rate was not significant ($p > 0.05$).

C) F1-Score

The GAT model had the highest median F1-score of 0.4706, with an IQR of 0.392. The GCN and Spike Rate models achieved a median F1-score of 0.400, with IQRs of 0.415 and 0.470, respectively. The GraphSAGE model attained a median of 0.364 (IQR 0.411), while the SVM model had the lowest median F1-score of 0.316 (IQR 0.289). The wider IQRs and presence of outliers in the F1-scores indicate varying model performance.

For the F1-score, none of the GNN models had a significantly higher median than Spike Rate ($p > 0.05$). The GAT model demonstrated a borderline significant difference in the F1-score compared to the Spike Rate ($p = 0.084$).

D) Precision for Resected Contacts (PPV)

In the analysis focusing on precision for the localization of resected contacts, the GAT model had the highest median precision of 0.714 (IQR 0.626). The Spike Rate model demonstrated a median precision of 0.615 (IQR 0.646). The GraphSAGE model had a median precision of 0.500 (IQR 0.441), while both GCN and SVM models scored the lowest with medians of 0.429 and IQRs of 0.541 and 0.476, respectively.

None of the GNN models had a significantly higher median precision in the localization of resected contacts compared to Spike Rate ($p > 0.05$).

4.1.2 Performance for Good vs Poor Outcome Patients

Comparative results of EZ localization for good versus poor outcome patients across models are shown in Figure 4.2. Mann-Whitney U-test was employed to assess statistical differences between patient groups.

A) AUROC

The GraphSAGE ($p = 0.028$) and GAT ($p = 0.014$) models observed significant differences between good and poor outcome groups. The difference for the GCN model was borderline significant ($p = 0.098$). The GAT model showed the strongest differentiation, achieving a median AUROC of 0.971 in the good outcome group versus 0.764 in the poor outcome group. Similarly, the GraphSAGE model demonstrated a significant decline from 0.971 to 0.812. These results indicate that the GAT and GraphSAGE models effectively distinguish between good and poor outcomes.

B) AUPRC

The GraphSAGE ($p = 0.002$), GAT ($p = 0.002$), and GCN ($p = 0.007$) models displayed significant differences in AUPRC between the two patient groups, suggesting these models are particularly effective at correctly identifying true positives while minimizing false positives in the good outcome group. The SVM model also demonstrated significant differences ($p = 0.028$), and the results for Spike Rate were borderline significant ($p = 0.061$). All models showed a considerable decline in performance for poor outcome patients.

C) F1-Score

All models indicated significant differences in F1-score between the two groups, highlighting the reduction in balanced precision and recall in poor-outcome patients. The GAT model exhibited the strongest differentiation ($p = 0.001$), dropping from 0.471 in the good-outcome group to 0.087 in the poor-outcome group.

D) Precision for Resected Contacts (PPV)

All models achieved significant differences in precision for resected contacts between good and poor outcome groups. The GAT model maintained high performance in the good outcome group, with a PPV of 0.714, significantly dropping to 0.116 in the poor outcome group ($p < 0.001$). Other models followed similar patterns.

4.2 Outcome Prediction

The outcome prediction accuracy of all GNN models was superior to the reference models. The GAT model showed the highest outcome prediction accuracy at 0.757, followed by GraphSAGE at 0.730 and GCN at 0.703. The reference Spike Rate model achieved an accuracy of 0.649, while the SVM model attained an accuracy of 0.568. Figure 4.3 presents contingency tables comparing the GNN models (GraphSAGE, GCN, and GAT) to the Spike Rate model in terms of correct and incorrect classifications of patient outcomes.

The McNemar test was conducted to determine the statistical significance of model performance differences. None of the GNN models displayed statistically significant improvements when compared to the Spike Rate model based on the p-values from the McNemar test: GraphSAGE ($p = 0.450$), GCN ($p = 0.617$), and GAT ($p = 0.221$).

The GraphSAGE model correctly predicted outcomes for five patients that the Spike Rate model misclassified but misclassified two patients that the Spike Rate model had predicted correctly. The GCN model was correct for three patients compared to Spike Rate but wrong for one patient that Spike Rate classified correctly. The best-performing model, GAT, correctly classified five patients that the Spike Rate model misclassified but predicted the wrong outcome for one patient the Spike Rate model classified correctly.

Epileptogenic Zone Localization Results (Good vs. Poor Outcome Patients)

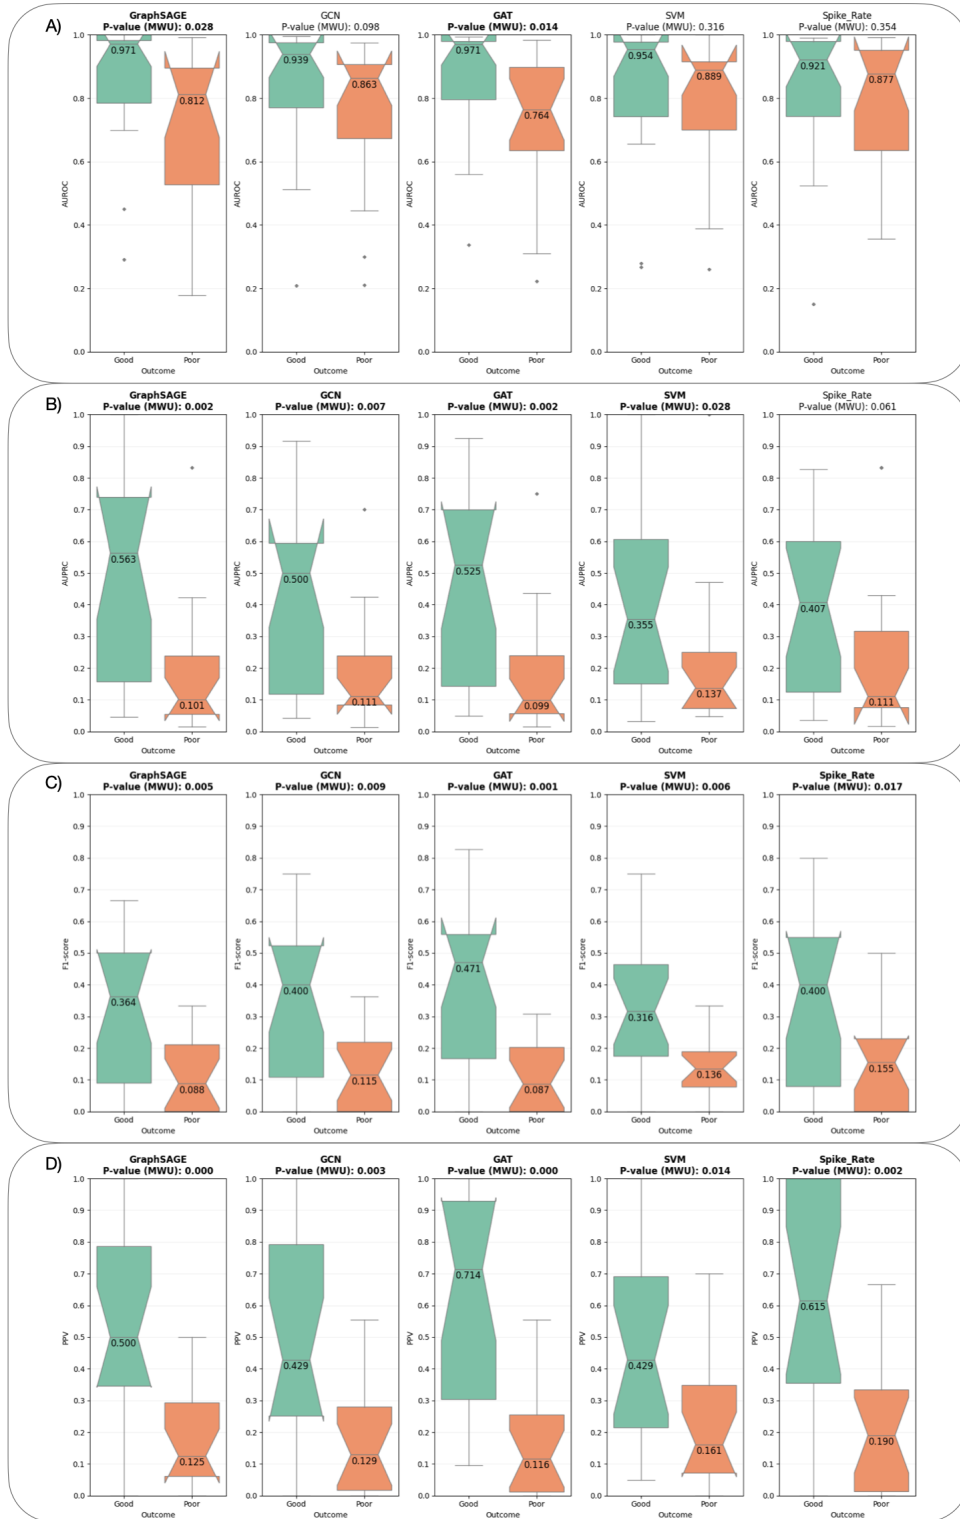


Fig. 4.2: Performance of models across the cohort of 19 good outcome patients vs cohort of 18 poor outcome patients for A) Area Under Receiver Operating Characteristic (AUROC), B) Area Under Precision-Recall Curve (AUPRC), C) F1-score, and D) precision (PPV) for resected contacts. Mann-Whitney U-test p-values are reported with significant results in bold.

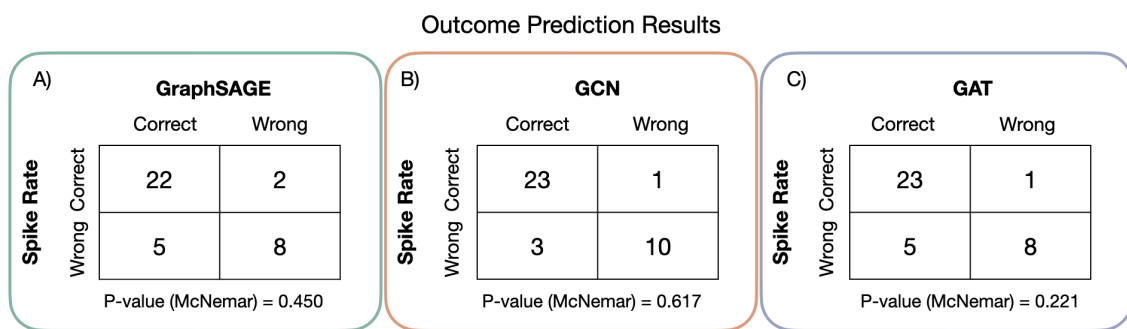


Fig. 4.3: Comparison of outcome prediction results between the reference Spike Rate model and three GNN models: GraphSAGE (A), GCN (B), and GAT (C). Each sub-figure displays a contingency table showing the number of correct and incorrect predictions for each model compared to the Spike Rate model. McNemar's test p-values are provided for each comparison to assess statistical differences between the models.

5 Discussion

In the practical part of the master’s thesis, three GNN models — GraphSAGE, GCN, and GAT — were used to localize the EZ based on interictal SEEG data. The analysis was performed on a clinical dataset from 37 patients across two institutions. From the SEEG recordings, 30 minutes of artifact-free NREM sleep, recorded at least 1 hour away from seizures, were used for the analysis. The models processed two features - spike rate and relative entropy- represented in the graph structure as node and edge features, respectively.

In a supervised learning framework, these models classified electrode contacts as either SOZ&Resected or non-SOZ&Resected. The models were evaluated in three phases: initially assessing performance using data from patients with good post-surgical outcomes, then comparing performance across cohorts with both good and poor postsurgical outcomes, and finally using the models to predict postsurgical outcomes for all patients in the dataset. Throughout these stages, the GNN models’ effectiveness was benchmarked against a reference SVM model, which utilized the same features as GNNs - spike rate and relative entropy - as features and a main reference Spike Rate that employed a single feature for localizing the EZ. The results from the practical application of GNN models show the potential of these models in epilepsy treatment planning.

This Chapter interprets the results of EZ localization in Section 5.1 and outcome prediction in Section 5.2, discussing their implications for the field of epilepsy treatment, comparing GNN model architectures in Section 5.3, and comparing the results to existing studies in Section 5.4. It also addresses the limitations of the study and suggests directions for future research in Section 5.5.

5.1 Epileptogenic Zone Localization

This Section discusses the results from the GNN models used for EZ localization, with a focus on how GNN models compared to traditional models in terms of performance metrics in Subsection 5.1.1, an exploration of the variability in model performance across different patients in Subsection 5.1.2, and discussion on why models perform better for good-outcome patients in Subsection 5.1.3.

5.1.1 Graph Neural Networks Outperform Reference Models

The results of the practical application of GNN models in localizing the EZ in 19 patients with good post-surgical outcomes indicate a promising enhancement over

traditional methods. Specifically, the GAT and GraphSAGE models demonstrated superior ability in differentiating between epileptic and non-epileptic contacts.

The GAT and GraphSAGE models demonstrated superior performance in terms of AUROC, achieving the highest scores of 0.971, clearly distinguishing themselves from the chance level of 0.5 for AUROC. This indicates a robust ability to identify true epileptic (SOZ&Resected) versus non-epileptic contacts, with low false positive rates. Both models outperformed the SVM and Spike Rate reference models, with the GAT model showing statistically significantly better performance than Spike Rate ($p = 0.016$).

In terms of the AUPRC, which considers both the precision and recall of the model, the GraphSAGE model excelled with a score of 0.563, surpassing the SVM and Spike Rate models, which scored substantially lower at 0.355 and 0.407, respectively. This superior performance indicates a better balance in the trade-off between precision and recall, which is crucial for clinical applicability, where both identifying true positives and minimizing false positives are important. Both GraphSAGE and GAT models significantly surpassed the Spike Rate model ($p = 0.017$ for both), further validating the efficacy of GNNs in this application.

The potential clinical utility of models such as GAT is highlighted by the F1-score and precision for resected contacts results, where the GAT model scored highest with a median score of 0.4706 and 0.714, respectively. The GAT surpassed both established localization methods, SVM and Spike Rate, and showed potential in predicting the surgical target with higher precision. Such performance is essential for optimizing surgical outcomes and minimizing unnecessary resections.

GNN models outperformed the chance levels in all cases, indicating a robust capability in accurately identifying the surgical targets. Furthermore, traditional models, such as SVM and the single-feature Spike Rate model, showed lower median scores across all evaluation metrics, highlighting the advanced capability of GNN models to handle complex patterns in SEEG data. The GAT model, in particular, showed exceptional results across various metrics, including F1-score and precision for resected contacts, which are both critical for clinical decision-making.

While GNN models like GAT and GraphSAGE have demonstrated their superiority over traditional models in terms of performance, it is important to note that their computational demands are significantly higher. This aspect was not evaluated in this study, and future research should consider the computational efficiency and the practical deployment challenges of using GNNs in clinical settings.

Case Study: Patient 723

The GAT model's performance for Patient 723, who underwent epilepsy surgery with a good postsurgical outcome classified as Engel IA, is illustrated in Figure 5.1. This case exemplifies the effectiveness of the GAT model, demonstrating its superior capability in discerning critical patterns in SEEG data. The GAT model outperformed the SVM model in all of the measured metrics, underscoring its higher accuracy and reliability.

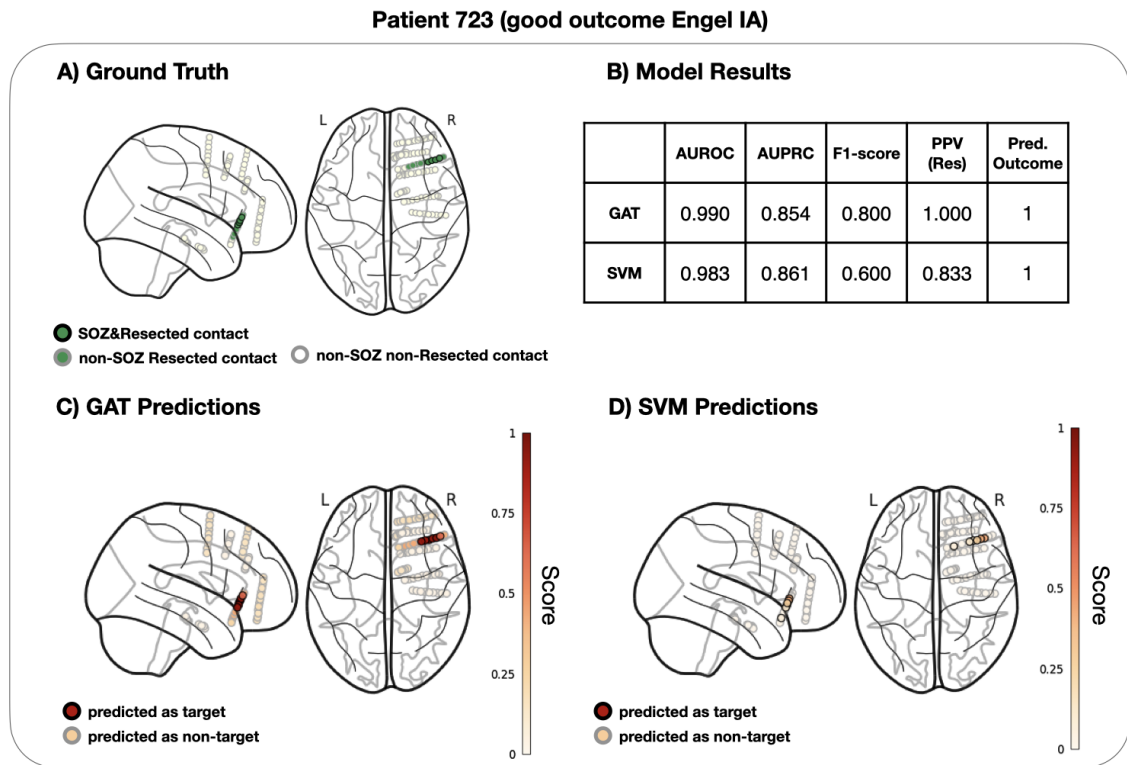


Fig. 5.1: Classification results for patient 723. In (A), the ground truth labels are visualized on a standard brain model, where green circles represent resected electrode contacts, green circles with black edges SOZ&Resected contacts, and white circles remaining contacts. In (B), results for the Graph Attention Network (GAT) and Support Vector Machine (SVM) models are presented. A predicted outcome value of 1 corresponds to a predicted good outcome and a value of 0 to a predicted poor outcome. In (C) and (D), the predictions of GAT and SVM, respectively, are visualized on a standard brain model. Contacts with black edges were classified as positive by the model, while contacts with gray edges were classified as negative. The color gradient symbolizes the scores assigned to each contact by the classifier normalized to a range between 0 and 1.

Importantly, the GAT model achieved a perfect precision score for resected con-

tacts, meaning that every contact identified as epileptogenic was indeed resected during successful epilepsy surgery. This outcome highlights the practical benefits of employing advanced graph-based neural networks for surgical planning. By representing SEEG data as graphs, the GAT model leverages complex interconnections between data points, leading to enhanced precision and clinical outcomes.

5.1.2 Models Show Variability in Patient Results

The variability in patient results, as demonstrated by the graphs for the GAT model in Figure 5.2, underscores an important aspect of using GNN models in a clinical setting. The ROC and PR curves show significant variability in results across the patient cohort, which reflects the underlying complexity and heterogeneity in clinical data used for localizing the EZ.

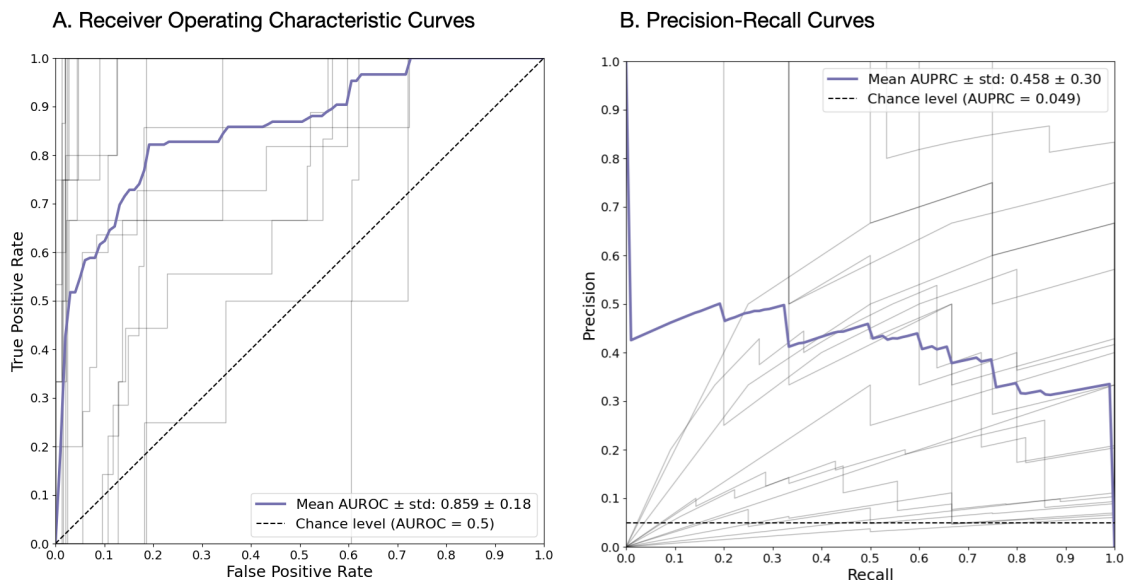


Fig. 5.2: Receiver Operating Characteristic (ROC) (A) and Precision-Recall (PR) curves (B) for the Graph Attention Network (GAT) model. Curves for testing on each patient are visualized in gray, along with the average curve in purple. The chance level in the ROC space lies on the minor diagonal, while in the PR space, the chance level is defined by the class imbalance in the data.

While the mean AUROC score is relatively high at 0.859 ± 0.18 standard deviation (std), the individual patient curves range considerably. Some approach the upper left corner, indicating excellent model performance with high true positive rates and low false positive rates, while others are closer to the diagonal line, representing performance near chance level (AUROC = 0.5). Similar variability is observed

in the Precision-Recall curves, with the mean AUPRC at 0.458 ± 0.30 std. Some curves demonstrate high precision across various recall levels, whereas others decline rapidly, showing significant challenges in maintaining precision as recall increases. This variability, however, is not unique to GNN models. Similar patterns are evident in results from the SVM and Spike Rate model, suggesting that the issue may not be inherent to GNN models specifically but rather to the size of the dataset and the heterogeneity between patients.

To address these challenges, further studies involving larger datasets are essential to better understand the influence of patient-specific factors that lead to suboptimal performance. Furthermore, future studies could explore different interictal features and their combinations, alternative data representations, and advanced GNN architectures, which could lead to more robust models performing better across different patients and more effective processing of diverse SEEG data.

For further refinement of the graph structure in GNNs for epilepsy surgery, it may be beneficial to incorporate the Euclidean distances between electrode contacts. This would allow the model to better account for patient-specific variability in electrode implantation, reflecting true anatomical spacing and potentially increasing sensitivity to spatial patterns of epileptic activity. In addition, the integration of multimodal data into the graph structure, such as imaging data, can enrich the input features and thus facilitate the generalization of the model to new patients. Among the promising directions in GNN architecture, the Graph Isomorphism Network (GIN) shows potential due to its powerful ability to distinguish between different graph structures, making it suitable to capture subtle differences in interictal brain activity. These solutions could significantly improve the predictive accuracy and clinical utility of GNNs for planning epilepsy surgery.

5.1.3 Models Perform Better for Good-Outcome Patients

The comparative analysis of EZ localization models shows significant differences in performance between good and poor-outcome patient groups. The GAT and GraphSAGE models notably demonstrated strong discriminative power, with significant AUROC differences, indicating their effectiveness in distinguishing patient outcomes. Similarly, the AUPRC results highlight these models' capability to accurately identify true positives, particularly in the good-outcome group.

The F1-score analysis further shows a significant reduction in both precision and recall in patients with poor outcomes, with the GAT model showing the best differentiation. Notably, the precision for resected contacts was significantly higher in the good outcome group, particularly for the GAT model, but decreased significantly in the poor outcome group. This pattern suggests that the GAT model was more

effective at identifying epileptogenic contacts within the resected zone in patients with good outcomes. Conversely, for patients with poor outcomes, the model detected more epileptogenic contacts outside the resected zone, potentially explaining why these patients continue to experience seizures post-surgery.

These results underscore the potential of advanced methods like GNNs in improving EZ localization and in retrospective analysis of poor-outcome cases.

5.2 Outcome Prediction

The GNN models demonstrated improved performance over traditional reference models in some patient cases in predicting outcomes following epilepsy surgery. Notably, the GAT model achieved the best results, accurately predicting outcomes for 28 out of 37 patients. This model correctly identified five patients for whom the Spike Rate model had misclassified, whereas the Spike Rate model only outperformed GAT in one instance. Similarly, the GraphSAGE model accurately predicted outcomes for 27 out of 37 patients, correcting five misclassifications by the Spike Rate model, which, however, outperformed GraphSAGE in two cases. The GCN model also performed well, accurately predicting outcomes for 26 out of 37 patients and correcting three misclassifications by the Spike Rate model, which surpassed GCN in just one case.

These results highlight the GNN models' superior capability to uncover complex patterns within SEEG data, even though they were originally designed and trained for localizing the EZ, not for predicting surgical outcomes. This fundamental aspect of the models' application suggests that their potential in outcome prediction could be further enhanced if they were directly trained for this purpose.

These results underscore the GNN models' enhanced ability to discern complex patterns within SEEG data more effectively than traditional models. However, despite their higher accuracy rates, McNemar's test revealed no statistically significant improvements in the GNN models over the Spike Rate model, with p-values indicating that the observed differences could be due to random variation. This suggests that while GNN models hold great promise, a more targeted approach in model training specifically for outcome prediction and further investigation with larger datasets are required to definitively establish their superiority over traditional models in predicting surgical outcomes.

Case Study: Patient 74

Performance for Patient 74, who underwent epilepsy surgery with a good postsurgical outcome classified as Engel IC, is illustrated in Figure 5.3. The GAT model

consistently scored higher in all metrics than the Spike Rate, showcasing GAT’s superior performance in both EZ localization and outcome prediction. Since the Spike Rate model made a large number of false positive predictions and only 38.1% of the contacts predicted as epileptogenic were inside the resected zone, the patient was incorrectly predicted to have a poor outcome by the reference model. In contrast, the GAT model was much more precise, with 85.7% of the contacts in the resected zone and a correctly predicted good postsurgical outcome.

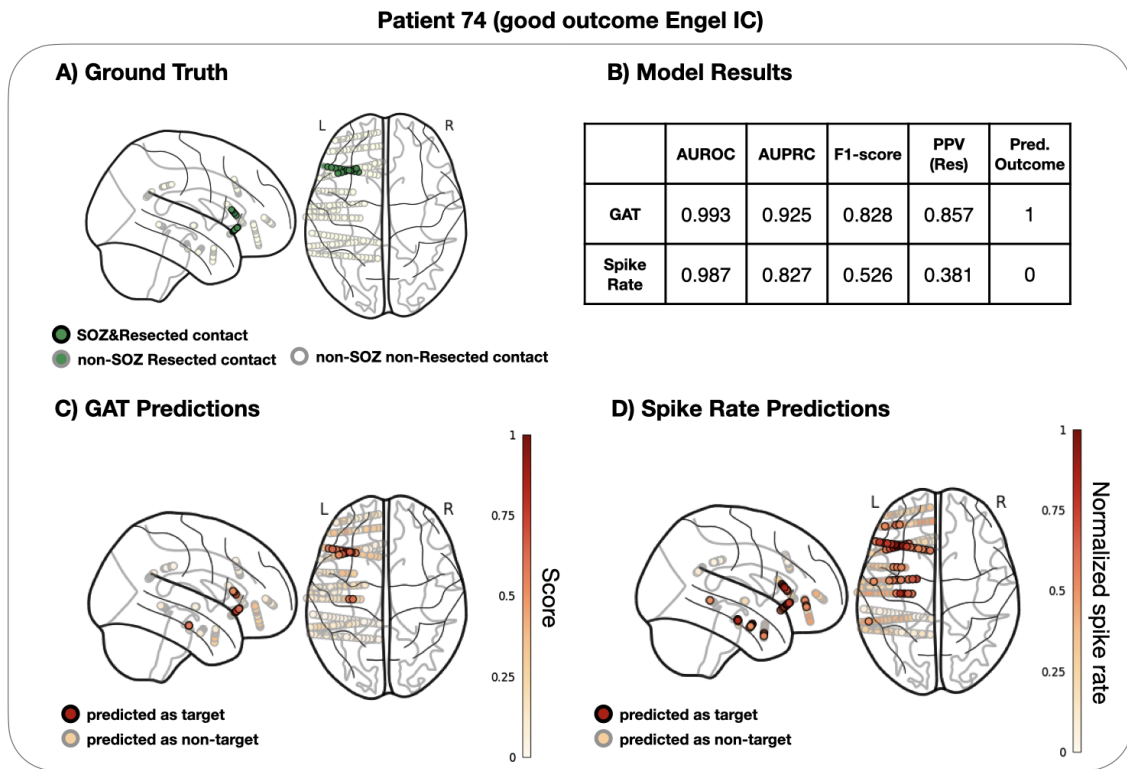


Fig. 5.3: Classification results for patient 74. In (A), the ground truth labels are visualized on a standard brain model, where green circles represent resected electrode contacts, green circles with black edges SOZ&Resected contacts, and white circles remaining contacts. In (B), results for the Graph Attention Network (GAT) and Spike Rate are presented. A predicted outcome value of 1 corresponds to a predicted good outcome and a value of 0 to a predicted poor outcome. In (C) and (D), the predictions of GAT and Spike Rate, respectively, are visualized on a standard brain model. Contacts with black edges were classified as positive by the model, while contacts with gray edges were classified as negative. For the GAT, the color gradient symbolizes the scores assigned to each contact by the classifier normalized to a range between 0 and 1. For Spike Rate, the gradient symbolizes the values of normalized spike rate.

Case Study: Patient 82

Results for Patient 82, who underwent epilepsy surgery with a good postsurgical outcome classified as Engel IA, are illustrated in Figure 5.4. Interestingly, while the GAT model underperformed relative to the Spike Rate model in terms of AUROC and AUPRC metrics, which assess model performance over a range of potential classification thresholds, it excelled in more clinically relevant metrics.

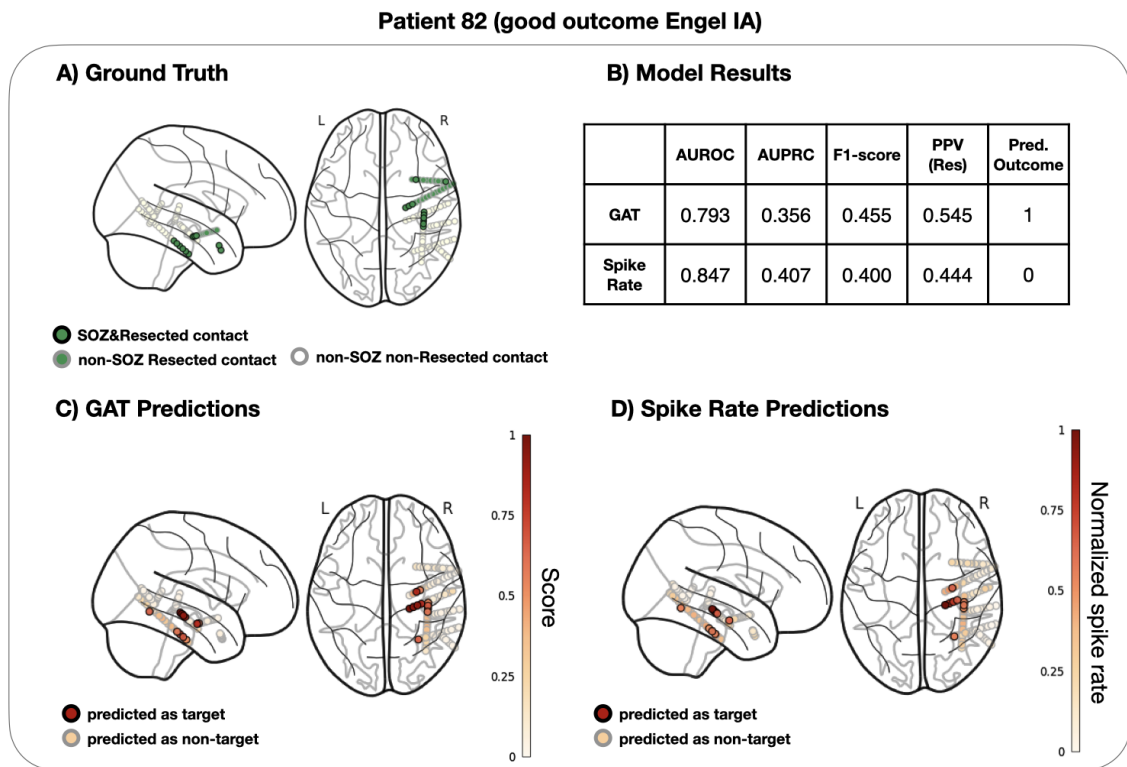


Fig. 5.4: Classification results for patient 82. In (A), the ground truth labels are visualized on a standard brain model, where green circles represent resected electrode contacts, green circles with black edges SOZ&Resected contacts, and white circles remaining contacts. In (B), results for the Graph Attention Network (GAT) and Spike Rate are presented. A predicted outcome value of 1 corresponds to a predicted good outcome and a value of 0 to a predicted poor outcome. In (C) and (D), the predictions of GAT and Spike Rate, respectively, are visualized on a standard brain model. Contacts with black edges were classified as positive by the model, while contacts with gray edges were classified as negative. For the GAT, the color gradient symbolizes the scores assigned to each contact by the classifier normalized to a range between 0 and 1. For Spike Rate, the gradient symbolizes the values of normalized spike rate.

Specifically, the GAT model outperformed the reference in terms of F1-score and

precision for resected contacts. This indicates a superior capacity to accurately categorize contacts into two distinct classes: epileptogenic and non-epileptogenic, using a single threshold. Importantly, this enhanced precision in identifying resected contacts led to a more accurate prediction of a postsurgical outcome, distinguishing itself from the Spike Rate model, which incorrectly predicted a poor outcome. This case shows the superior performance of GAT in areas critical for clinical decision-making, especially in predicting surgical success in epilepsy treatment.

Case Study: Patient 1159

A case of Patient 1159, who underwent epilepsy surgery with a poor postsurgical outcome classified as Engel IIIC, is illustrated in Figure 5.5. The postsurgical outcome means that the patient continued to experience seizures post-surgery, indicating that not all epileptogenic contacts were successfully resected.

The Spike Rate model demonstrated high scores in all metrics, including precision for resected contacts, which led to the misclassification of patient outcomes as good. Conversely, the GAT model, with lower metric scores, accurately predicted a poor outcome. This was due to its identification of the majority of epileptogenic contacts outside the resected zone, aligning better with the patient’s ongoing seizures post-surgery.

This case underlines the GAT model’s ability to identify epileptogenic contacts outside the ground truth localization target that is not guaranteed to cover the EZ for poor outcome patients. While the Spike Rate model excelled in identifying the ground truth target contacts, it potentially failed to localize additional epileptogenic contacts, crucial for the success of epilepsy surgery. In contrast, the GAT model’s broader analysis showed a more realistic prediction, possibly identifying a secondary seizure generator outside the resected area. This case demonstrates the GAT’s utility in complex clinical scenarios where comprehensive epilepsy zone identification is crucial for outcome accuracy.

5.3 Comparison of GNN Models: GraphSAGE, GCN, and GAT

In evaluating GNNs for EZ localization and outcome prediction, three models were assessed: GraphSAGE, GCN, and GAT. GraphSAGE used a simpler approach focusing only on spike rate and the graph structure, omitting edge weights corresponding to relative entropy values. This simplicity might limit its depth of analysis compared to GCN and GAT, which incorporate relative entropy values as edge weights, allowing them to capture more detailed interdependencies within the data.

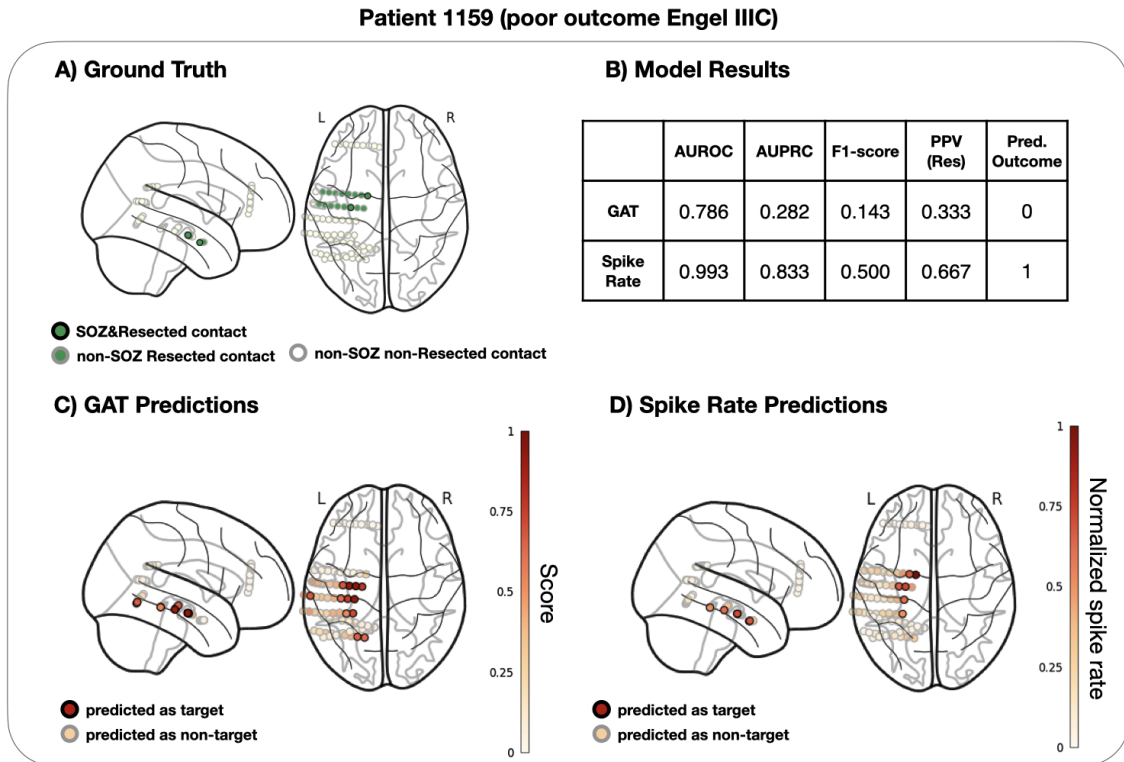


Fig. 5.5: Classification results for patient 1159. In (A), the ground truth labels are visualized on a standard brain model, where green circles represent resected electrode contacts, green circles with black edges SOZ&Resected contacts, and white circles remaining contacts. In (B), results for the Graph Attention Network (GAT) and Spike Rate are presented. A predicted outcome value of 1 corresponds to a predicted good outcome and a value of 0 to a predicted poor outcome. In (C) and (D), the predictions of GAT and Spike Rate, respectively, are visualized on a standard brain model. Contacts with black edges were classified as positive by the model, while contacts with gray edges were classified as negative. For the GAT, the color gradient symbolizes the scores assigned to each contact by the classifier normalized to a range between 0 and 1. For Spike Rate, the gradient symbolizes the values of normalized spike rate.

Among these, GAT showed the most stable and superior performance across all metrics due to its attention mechanism. This feature enables GAT to dynamically prioritize important signals, enhancing accuracy and interpretability. The attention scores from GAT can offer deeper insights into epileptic networks, highlighting key nodes critical in seizure generation. This makes GAT particularly promising for advancing our understanding of epileptic networks and improving surgical outcomes in epilepsy treatment.

5.4 Comparison to Published Studies

The GNN models, particularly the GraphSAGE and GAT, implemented in this thesis have shown superior performance compared to previous efforts with similar patient cohorts. Notably, these models surpass the multi-feature SVM approach used in localizing epileptogenic tissue from interictal SEEG data, which reported an AUROC of 0.838 in the study by Cimbalnik et al. [13]. Similarly, in a comparable cohort, Chybowski et al. [92] recorded an AUPRC of 0.521 with the best-performing SVM model, comparable to the results of the GAT model (AUPRC = 0.525), but lower to the GraphSAGE model (AUPRC = 0.563). Furthermore, the proposed GAT and GraphSAGE models achieved higher AUROC scores (AUROC = 0.971) than a benchmark study by Jiang et al. [104] from 2022, who utilized interictal SEEG connectivity to localize the SOZ, achieving an AUROC of 0.94.

This comparative analysis underscores the significant advantages of employing GNNs to analyze the networked nature of epileptic signals. GNNs' enhanced performance in localizing EZs represents a considerable improvement over traditional models.

The novel application of GNNs for EZ localization, particularly using only interictal data, marks a significant distinction from prior research, such as Grattarola et al.'s work [78], which focused on ictal intracranial EEG recordings using an attention-based GNN network. The method proposed in this thesis avoids the complexities and limitations associated with seizure recording and detection by employing only interictal segments, specifically 30-minute recordings from NREM sleep for each patient, suggesting a pathway to significantly reduce the duration of presurgical monitoring. To the best of the author's knowledge, this is the first study to explore EZ localization using GNNs exclusively with interictal data, potentially setting the stage for a more efficient and less invasive presurgical evaluation process for epilepsy patients.

5.5 Study Limitations

The primary limitation is the training and testing scheme, particularly the hyperparameter tuning methodology. Hyperparameters were tuned using a leave-one-patient-out cross-validation scheme across all good-outcome patients, potentially exposing the tuning process to data from the patient left out during testing. Ideally, hyperparameter tuning should be conducted on a separate validation set derived from the training cohort, ensuring that the test patient remains completely unseen during both the training and validation phases. This would provide a more rigorous and unbiased evaluation of model performance.

Additionally, the models were tested using data from the same two institutions from

which the training data was sourced, which might not fully demonstrate the models' capacity to generalize. Validating the models on an entirely separate cohort from a different institution would better assess their adaptability to variations in data collection, patient demographics, and other institution-specific factors.

The significant limitation is the size of the patient cohort. The small number of patients may contribute to statistical insignificance in several cases, limiting the confidence in the generalizability and robustness of the findings. A larger cohort would not only provide a stronger statistical basis for the results but would also allow for more conclusive assessments of the models' performance.

Moreover, each model was evaluated in only a single experimental run, which does not account for the inherent randomness in the training and inference processes of machine learning models. Multiple runs would provide a more reliable measure of model performance, capturing potential variability in results due to initialization and stochastic optimization processes.

Lastly, the models rely on approximations of the EZ since the exact extent and location cannot be definitively measured with currently available methods. This study used information about the SOZ, the resected region, and the postsurgical outcome to approximate the EZ. This method of approximation assumes that a good surgical outcome is indicative of a complete resection of the EZ, further refining the identification of the EZ to only those contacts also recognized as part of the SOZ. This could potentially overlook other epileptogenic contacts that were resected but not identified as seizure-onset by clinicians. Consequently, the ground truth labels used for training and testing the models may not be entirely accurate. To minimize the effect during model training, a noise-tolerant approach to classification could be implemented.

Addressing these limitations in future studies could significantly enhance the reliability and applicability of GNNs in clinical settings.

Conclusion

The thesis aims to enhance the localization of the EZ using SEEG recordings by employing advanced machine learning techniques, specifically GNNs. The objectives are to develop and implement a GNN model that effectively identifies the EZ from SEEG data, validate this model on clinical data from drug-resistant epilepsy patients who underwent SEEG and subsequent surgery, and finally, evaluate the model's performance. This work seeks to contribute to the state of the art in EZ localization by demonstrating the potential of GNNs to improve the accuracy and reliability of surgical planning in epilepsy treatment.

In the theoretical section, the introduction begins with an overview of drug-resistant epilepsy, epilepsy surgery, and the significance of EZ localization in presurgical evaluation. Following this, a detailed introduction to EZ localization from SEEG recordings is provided, incorporating a review of key interictal biomarkers associated with the EZ and an examination of graph analysis methods employed for EZ localization. The subsequent section introduces the concept of GNNs, covering the representation of data as graphs, various graph types, prediction task types, and the components integral to GNN models. GraphSAGE, GCN, and GAT architecture are presented as examples of GNN models, and the applications of GNNs in EEG analysis are reviewed.

In the practical part, the thesis describes the clinical dataset used, outlines the patient selection process, and details the patient data. The Methods section explains the data pre-processing steps, including the definition of localization targets and the representation of EEG signals as graphs. The section details the implementation of three GNN models—GraphSAGE, GCN, and GAT—alongside two reference models, Spike Rate and SVM, including their hyperparameter tuning. Testing and evaluation methods for EZ localization and outcome prediction are outlined, followed by a discussion of the results, comparing the GNN models' performance to reference models, comparing the three GNN architectures, and benchmarking against state-of-the-art methods in interictal EZ localization. Four patient case studies are included in the discussion to illustrate model performance in specific patients. The practical part concludes with an acknowledgment of the limitations of the study and suggestions for future research.

The presented results highlight the potential of GNN models—specifically GraphSAGE and GAT—in guiding epilepsy surgery using interictal SEEG data. This study demonstrated that GNN models when compared to traditional methods such as SVM and simple spike rate models, provide a substantial improvement in classifying epileptogenic versus non-epileptogenic contacts. The representation of SEEG data as graphs, which enables the GNNs to capture the topological and functional

relationships between the signals, is especially relevant in the study of complex epileptic networks. These models excel in various performance metrics, indicating their superior precision and reliability in clinical settings. The application of GNNs has also shown notable success in predicting patient postsurgical outcomes, further underscoring their potential to promote more accurate and effective epilepsy surgery planning. Through analysis of model predictions and visualization of explanations produced by the GNN, new insight into epileptic networks could potentially be gained, driving further research into epilepsy treatment options.

Despite these advancements, the study also acknowledges several limitations, including the potential for overfitting due to the small sample size and the methodological constraints of the training and testing framework. These issues highlight the need for further research, particularly involving larger datasets and more rigorous validation methods, to ensure the models' generalizability and reliability across different clinical settings.

In conclusion, this research represents a significant advancement in epilepsy treatment by introducing GNNs to interictal SEEG analysis, as the proposed methodology paves the way for enhancing the precision of EZ localization and reducing SEEG monitoring durations.

Bibliography

- [1] Roland D Thijs, Rainer Surges, Terence J O'Brien, and Josemir W Sander. Epilepsy in adults. *The Lancet*, 393(10172):689–701, February 2019. doi: 10.1016/S0140-6736(18)32596-0.
- [2] M. J. Brodie, S. D. Shorvon, R. Canger, P. Halász, S. Johannessen, P. Thompson, H. G. Wieser, and P. Wolf. Commission on European Affairs: Appropriate Standards of Epilepsy Care Across Europe. *Epilepsia*, 38(11):1245–1250, November 1997. doi:10.1111/j.1528-1157.1997.tb01224.x.
- [3] Jerome Engel. The current place of epilepsy surgery. *Current Opinion in Neurology*, 31(2):192–197, April 2018. doi:10.1097/WCO.0000000000000528.
- [4] Patrick Kwan, Steven C. Schachter, and Martin J. Brodie. Drug-Resistant Epilepsy. *New England Journal of Medicine*, 365(10):919–926, September 2011. doi:10.1056/NEJMr1004418.
- [5] Birgit Frauscher. Localizing the epileptogenic zone. *Current Opinion in Neurology*, 33(2):198–206, April 2020. doi:10.1097/WCO.0000000000000790.
- [6] Shawn L Hervey-Jumper and Mitchel S Berger. Introduction: Surgical Management of Eloquent Area Tumors. *Neurosurgery*, 87(6):1076–1077, December 2020. doi:10.1093/neuros/nyaa358.
- [7] F. Rosenow. Presurgical evaluation of epilepsy. *Brain*, 124(9):1683–1700, September 2001. doi:10.1093/brain/124.9.1683.
- [8] Debayan Dasgupta, Anna Miserocchi, Andrew W. McEvoy, and John S. Duncan. Previous, current, and future stereotactic EEG techniques for localising epileptic foci. *Expert Review of Medical Devices*, 19(7):571–580, July 2022. doi:10.1080/17434440.2022.2114830.
- [9] Jan Cimbalnik, Benjamin Brinkmann, Vaclav Kremen, Pavel Jurak, Brent Berry, Jamie Van Gompel, Matt Stead, and Greg Worrell. Physiological and pathological high frequency oscillations in focal epilepsy. *Annals of Clinical and Translational Neurology*, 5(9):1062–1076, September 2018. doi:10.1002/acn3.618.
- [10] John Thomas, Philippe Kahane, Chifaou Abdallah, Tamir Avigdor, Willemiek J. E. M. Zweiphenning, Stephan Chabardes, Kassem Jaber, Véronique Latreille, Lorella Minotti, Jeff Hall, François Dubeau, Jean Gotman, and Birgit Frauscher. A Subpopulation of Spikes Predicts Successful Epilepsy Surgery

- Outcome. *Annals of Neurology*, 93(3):522–535, March 2023. doi:10.1002/ana.26548.
- [11] Kristin M Gunnarsdottir, Adam Li, Rachel J Smith, Joon-Yi Kang, Anna Korzeniewska, Nathan E Crone, Adam G Rouse, Jennifer J Cheng, Michael J Kinsman, Patrick Landazuri, Utku Uysal, Carol M Ulloa, Nathaniel Cameron, Iahn Cajigas, Jonathan Jagid, Andres Kanner, Turki Elarjani, Manuel Melo Bicchi, Sara Inati, Kareem A Zaghoul, Varina L Boerwinkle, Sarah Wyckoff, Niravkumar Barot, Jorge Gonzalez-Martinez, and Sridevi V Sarma. Source-sink connectivity: a novel interictal EEG marker for seizure localization. *Brain*, 145(11):3901–3915, November 2022. doi:10.1093/brain/awac300.
- [12] Vojtech Travnıcek, Petr Klimes, Jan Cimbalnik, Josef Halamek, Pavel Jurak, Benjamin Brinkmann, Irena Balzekas, Chifaou Abdallah, Franois Dubeau, Birgit Frauscher, Greg Worrell, and Milan Brazdil. Relative entropy is an easy-to-use invasive electroencephalographic biomarker of the epileptogenic zone. *Epilepsia*, 64(4):962–972, April 2023. doi:10.1111/epi.17539.
- [13] Jan Cimbalnik, Petr Klimes, Vladimir Sladky, Petr Nejedly, Pavel Jurak, Martin Pail, Robert Roman, Pavel Daniel, Hari Guragain, Benjamin Brinkmann, Milan Brazdil, and Greg Worrell. Multi-feature localization of epileptic foci from interictal, intracranial EEG. *Clinical Neurophysiology*, 130(10):1945–1953, October 2019. doi:10.1016/j.clinph.2019.07.024.
- [14] Ziyu Jia, Youfang Lin, Jing Wang, Ronghao Zhou, Xiaojun Ning, Yuanlai He, and Yaoshuai Zhao. GraphSleepNet: Adaptive Spatial-Temporal Graph Convolutional Networks for Sleep Stage Classification. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, pages 1324–1330, Yokohama, Japan, July 2020. International Joint Conferences on Artificial Intelligence Organization. doi:10.24963/ijcai.2020/184.
- [15] Tengfei Song, Wenming Zheng, Peng Song, and Zhen Cui. EEG Emotion Recognition Using Dynamical Graph Convolutional Neural Networks. *IEEE Transactions on Affective Computing*, 11(3):532–541, July 2020. doi:10.1109/TAFFC.2018.2817622.
- [16] C. E. Stafstrom and L. Carmant. Seizures and Epilepsy: An Overview for Neuroscientists. *Cold Spring Harbor Perspectives in Medicine*, 5(6):a022426–a022426, June 2015. doi:10.1101/cshperspect.a022426.

- [17] Katsuhiko Kobayashi, Tomoyuki Akiyama, and Cristina Go. Epilepsy. In Toshitaka Oohashi, Hirokazu Tsukahara, Francesco Ramirez, Chad L. Barber, and Fumio Otsuka, editors, *Human Pathobiochemistry*, pages 295–310. Springer Singapore, Singapore, 2019. doi:10.1007/978-981-13-2977-7_26.
- [18] Ingrid E. Scheffer, Samuel Berkovic, Giuseppe Capovilla, Mary B. Connolly, Jacqueline French, Laura Guilhoto, Edouard Hirsch, Satish Jain, Gary W. Mathern, Solomon L. Moshé, Douglas R. Nordli, Emilio Perucca, Torbjörn Tomson, Samuel Wiebe, Yue-Hua Zhang, and Sameer M. Zuberi. ILAE classification of the epilepsies: Position paper of the ILAE Commission for Classification and Terminology. *Epilepsia*, 58(4):512–521, April 2017. doi:10.1111/epi.13709.
- [19] Wolfgang Löscher and Pavel Klein. The Pharmacology and Clinical Efficacy of Antiseizure Medications: From Bromide Salts to Cenobamate and Beyond. *CNS Drugs*, 35(9):935–963, September 2021. doi:10.1007/s40263-021-00827-8.
- [20] Lara Jehi. The Epileptogenic Zone: Concept and Definition. *Epilepsy Currents*, 18(1):12–16, February 2018. doi:10.5698/1535-7597.18.1.12.
- [21] Anatomical Images: Free Human Anatomy Images and Pictures. URL: <https://www.injurymap.com/free-human-anatomy-illustrations>.
- [22] Maeike Zijlmans, Willemiek Zweiphenning, and Nicole Van Klink. Changing concepts in presurgical assessment for epilepsy surgery. *Nature Reviews Neurology*, 15(10):594–606, October 2019. doi:10.1038/s41582-019-0224-y.
- [23] Christoph Baumgartner, Johannes P. Koren, Martha Britto-Arias, Lea Zoche, and Susanne Pirker. Presurgical epilepsy evaluation and epilepsy surgery. *F1000Research*, 8:1818, October 2019. doi:10.12688/f1000research.17714.1.
- [24] Md. Rabiul Islam, Xuyang Zhao, Yao Miao, Hidenori Sugano, and Toshihisa Tanaka. Epileptic seizure focus detection from interictal electroencephalogram: a survey. *Cognitive Neurodynamics*, 17(1):1–23, February 2023. doi:10.1007/s11571-022-09816-z.
- [25] Jorge A. Gonzalez-Martinez. The Stereo-Electroencephalography: The Epileptogenic Zone. *Journal of Clinical Neurophysiology*, 33(6):522–529, December 2016. doi:10.1097/WNP.0000000000000327.

- [26] György Buzsáki, Costas A. Anastassiou, and Christof Koch. The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, 13(6):407–420, June 2012. doi:10.1038/nrn3241.
- [27] Hui Ming Khoo, Jeffery A. Hall, Francois Dubeau, Naoki Tani, Satoru Os-hino, Yuya Fujita, Jean Gotman, and Haruhiko Kishima. Technical Aspects of SEEG and Its Interpretation in the Delineation of the Epileptogenic Zone. *Neurologia medico-chirurgica*, 60(12):565–580, 2020. URL: https://www.jstage.jst.go.jp/article/nmc/60/12/60_st.2020-0176/_article, doi:10.2176/nmc.st.2020-0176.
- [28] Jeffrey P. Mullin, Saksith Smithason, and Jorge Gonzalez-Martinez. Stereo-Electro-Encephalo-Graphy (SEEG) With Robotic Assistance in the Presurgical Evaluation of Medical Refractory Epilepsy: A Technical Note. *Journal of Visualized Experiments*, (112):53206, June 2016. doi:10.3791/53206.
- [29] Stjepana Kovac, Vejay N. Vakharia, Catherine Scott, and Beate Diehl. Invasive epilepsy surgery evaluation. *Seizure*, 44:125–136, January 2017. doi:10.1016/j.seizure.2016.10.016.
- [30] Donald L. Schomer and Fernando H. Lopes Da Silva, editors. *Niedermeyer’s Electroencephalography*, volume 1. Oxford University Press, November 2017. doi:10.1093/med/9780190228484.001.0001.
- [31] Eric D. Marsh, Bradley Peltzer, Merritt W. Brown Iii, Courtney Wusthoff, Phillip B. Storm Jr, Brian Litt, and Brenda E. Porter. Interictal EEG spikes identify the region of electrographic seizure onset in some, but not all, pediatric epilepsy patients. *Epilepsia*, 51(4):592–601, April 2010. doi:10.1111/j.1528-1167.2009.02306.x.
- [32] Philippa J. Karoly, Dean R. Freestone, Ray Boston, David B. Grayden, David Himes, Kent Leyde, Udaya Seneviratne, Samuel Berkovic, Terence O’Brien, and Mark J. Cook. Interictal spikes and epileptic seizures: their relationship and underlying rhythmicity. *Brain*, 139(4):1066–1078, April 2016. doi:10.1093/brain/aww019.
- [33] Marco de Curtis, John G. R. Jefferys, and Massimo Avoli. Interictal Epileptiform Discharges in Partial Epilepsy: Complex Neurobiological Mechanisms Based on Experimental and Clinical Evidence. In Jeffrey L. Noebels, Massimo Avoli, Michael A. Rogawski, Richard W. Olsen, and Antonio V. Delgado-Escueta, editors, *Jasper’s Basic Mechanisms of the Epilepsies*. National Center for Biotechnology Information (US), Bethesda (MD), 4th edition, 2012.

- [34] Dong Wook Kim, Hyun Kyung Kim, Sang Kun Lee, Kon Chu, and Chun Kee Chung. Extent of neocortical resection and surgical outcome of epilepsy: Intracranial EEG analysis. *Epilepsia*, 51(6):1010–1017, June 2010. doi:10.1111/j.1528-1167.2010.02567.x.
- [35] Abdullah Azeem, Nicolas Von Ellenrieder, Jeffery Hall, Francois Dubeau, Birgit Frauscher, and Jean Gotman. Interictal spike networks predict surgical outcome in patients with drug-resistant focal epilepsy. *Annals of Clinical and Translational Neurology*, 8(6):1212–1223, June 2021. doi:10.1002/acn3.51337.
- [36] Liankun Ren, Michal T. Kucewicz, Jan Cimbalnik, Joseph Y. Matsumoto, Benjamin H. Brinkmann, Wei Hu, W. Richard Marsh, Fredric B. Meyer, S. Matthew Stead, and Gregory A. Worrell. Gamma oscillations precede interictal epileptiform spikes in the seizure onset zone. *Neurology*, 84(6):602–608, February 2015. doi:10.1212/WNL.0000000000001234.
- [37] A. Hufnagel, M. Dümpelmann, J. Zentner, O. Schijns, and C. E. Elger. Clinical Relevance of Quantified Intracranial Interictal Spike Activity in Presurgical Evaluation of Epilepsy. *Epilepsia*, 41(4):467–478, April 2000. doi:10.1111/j.1528-1157.2000.tb00191.x.
- [38] J. Jacobs, R. Staba, E. Asano, H. Otsubo, J.Y. Wu, M. Zijlmans, I. Mohamed, P. Kahane, F. Dubeau, V. Navarro, and J. Gotman. High-frequency oscillations (HFOs) in clinical epilepsy. *Progress in Neurobiology*, 98(3):302–315, September 2012. doi:10.1016/j.pneurobio.2012.03.001.
- [39] Anatol Bragin, Jerome Engel, Charles L. Wilson, Itzhak Fried, and Gyorgy Buzski. High-frequency oscillations in human brain. *Hippocampus*, 9(2):137–142, 1999. doi:10.1002/(SICI)1098-1063(1999)9:2<137::AID-HIP05>3.0.CO;2-0.
- [40] Milan Brázdil, Martin Pail, Josef Halámek, Filip Plešinger, Jan Cimbálník, Robert Roman, Petr Klimeš, Pavel Daniel, Jan Chrastina, Eva Brichtová, Ivan Rektor, Gregory A. Worrell, and Pavel Jurák. Very high-frequency oscillations: Novel biomarkers of the epileptogenic zone. *Annals of Neurology*, 82(2):299–310, August 2017. doi:10.1002/ana.25006.
- [41] Rui Zuo, Jing Wei, Xiaonan Li, Chunlin Li, Cui Zhao, Zhaohui Ren, Ying Liang, Xinling Geng, Chenxi Jiang, Xiaofeng Yang, and Xu Zhang. Automated Detection of High-Frequency Oscillations in Epilepsy Based on a Convolutional

- Neural Network. *Frontiers in Computational Neuroscience*, 13:6, February 2019. doi:10.3389/fncom.2019.00006.
- [42] Hisako Fujiwara, Hansel M. Greiner, Ki Hyeong Lee, Katherine D. Holland-Bouley, Joo Hee Seo, Todd Arthur, Francesco T. Mangano, James L. Leach, and Douglas F. Rose. Resection of ictal high-frequency oscillations leads to favorable surgical outcome in pediatric epilepsy. *Epilepsia*, 53(9):1607–1617, September 2012. doi:10.1111/j.1528-1167.2012.03629.x.
- [43] Claire Haegelen, Piero Perucca, Claude-Edouard Châtillon, Luciana Andrade-Valença, Rina Zelmann, Julia Jacobs, D. Louis Collins, François Dubeau, André Olivier, and Jean Gotman. High-frequency oscillations, extent of surgical resection, and surgical outcome in drug-resistant focal epilepsy. *Epilepsia*, 54(5):848–857, May 2013. doi:10.1111/epi.12075.
- [44] György Buzsáki and Fernando Lopes Da Silva. High frequency oscillations in the intact brain. *Progress in Neurobiology*, 98(3):241–249, September 2012. doi:10.1016/j.pneurobio.2012.02.004.
- [45] Andrew Matsumoto, Benjamin H. Brinkmann, S. Matthew Stead, Joseph Matsumoto, Michal T. Kucewicz, W. Richard Marsh, Frederic Meyer, and Gregory Worrell. Pathological and physiological high-frequency oscillations in focal human epilepsy. *Journal of Neurophysiology*, 110(8):1958–1964, October 2013. doi:10.1152/jn.00341.2013.
- [46] Krit Charupanit, Indranil Sen-Gupta, Jack J. Lin, and Beth A. Lopour. Detection of anomalous high-frequency events in human intracranial EEG. *Epilepsia Open*, 5(2):263–273, June 2020. doi:10.1002/epi4.12397.
- [47] Wen Shi, Dana Shaw, Katherine G Walsh, Xue Han, Uri T Eden, Robert M Richardson, Stephen V Gliske, Julia Jacobs, Benjamin H Brinkmann, Gregory A Worrell, William C Stacey, Birgit Frauscher, John Thomas, Mark A Kramer, and Catherine J Chu. Spike ripples localize the epileptogenic zone best: an international intracranial study. *Brain*, page awae037, February 2024. doi:10.1093/brain/awae037.
- [48] Sai Sanjay Balaji and Keshab K. Parhi. Seizure Onset Zone Identification From iEEG: A Review. *IEEE Access*, 10:62535–62547, 2022. doi:10.1109/ACCESS.2022.3182716.
- [49] Stefano Gallotto and Margitta Seeck. EEG biomarker candidates for the identification of epilepsy. *Clinical Neurophysiology Practice*, 8:32–41, 2023. doi:10.1016/j.cnp.2022.11.004.

- [50] Arun R. Antony, Andreas V. Alexopoulos, Jorge A. González-Martínez, John C. Mosher, Lara Jehi, Richard C. Burgess, Norman K. So, and Roberto F. Galán. Functional Connectivity Estimated from Intracranial EEG Predicts Surgical Outcome in Intractable Temporal Lobe Epilepsy. *PLoS ONE*, 8(10):e77916, October 2013. doi:10.1371/journal.pone.0077916.
- [51] Florian Mormann, Klaus Lehnertz, Peter David, and Christian E. Elger. Mean phase coherence as a measure for phase synchronization and its application to the EEG of epilepsy patients. *Physica D: Nonlinear Phenomena*, 144(3-4):358–369, October 2000. doi:10.1016/S0167-2789(00)00087-7.
- [52] Bahareh Elahian, Mohammed Yeasin, Basanagoud Mudigoudar, James W. Wheless, and Abbas Babajani-Feremi. Identifying seizure onset zone from electrocorticographic recordings: A machine learning approach based on phase locking value. *Seizure*, 51:35–42, October 2017. doi:10.1016/j.seizure.2017.07.010.
- [53] Robert Coben and Iman Mohammad-Rezazadeh. Neural Connectivity in Epilepsy as Measured by Granger Causality. *Frontiers in Human Neuroscience*, 9, July 2015. doi:10.3389/fnhum.2015.00194.
- [54] Eun-Hyoung Park and Joseph R Madsen. Granger Causality Analysis of Interictal iEEG Predicts Seizure Focus and Ultimate Resection. *Neurosurgery*, 82(1):99–109, January 2018. doi:10.1093/neuros/nyx195.
- [55] Mona Hejazi and Ali Motie Nasrabadi. Prediction of epilepsy seizure from multi-channel electroencephalogram by effective connectivity analysis using Granger causality and directed transfer function methods. *Cognitive Neurodynamics*, 13(5):461–473, October 2019. doi:10.1007/s11571-019-09534-z.
- [56] Rakesh Malladi, Giridhar Kalamangalam, Nitin Tandon, and Behnaam Aazhang. Identifying Seizure Onset Zone From the Causal Connectivity Inferred Using Directed Information. *IEEE Journal of Selected Topics in Signal Processing*, 10(7):1267–1283, October 2016. doi:10.1109/JSTSP.2016.2601485.
- [57] Yonathan Murin, Jeremy Kim, Josef Parvizi, and Andrea Goldsmith. SozRank: A new approach for localizing the epileptic seizure onset zone. *PLOS Computational Biology*, 14(1):e1005953, January 2018. doi:10.1371/journal.pcbi.1005953.
- [58] William Stacey, Mark Kramer, Kristin Gunnarsdottir, Jorge Gonzalez-Martinez, Kareem Zaghloul, Sara Inati, Sridevi Sarma, Jennifer Stiso,

- Ankit N. Khambhati, Danielle S. Bassett, Rachel J. Smith, Virginia B. Liu, Beth A. Lopour, and Richard Staba. Emerging roles of network analysis for epilepsy. *Epilepsy Research*, 159:106255, January 2020. doi:10.1016/j.eplepsyres.2019.106255.
- [59] Pieter Van Mierlo, Evelien Carrette, Hans Hallez, Robrecht Raedt, Alfred Meurs, Stefaan Vandenberghe, Dirk Van Roost, Paul Boon, Steven Staelens, and Kristl Vonck. Ictal-onset localization through connectivity analysis of intracranial EEG signals in patients with refractory epilepsy. *Epilepsia*, 54(8):1409–1418, August 2013. doi:10.1111/epi.12206.
- [60] Yong-Hua Li, Xiao-Lai Ye, Qiang-Qiang Liu, Jun-Wei Mao, Pei-Ji Liang, Ji-Wen Xu, and Pu-Ming Zhang. Localization of epileptogenic zone based on graph analysis of stereo-EEG. *Epilepsy Research*, 128:149–157, December 2016. doi:10.1016/j.eplepsyres.2016.10.021.
- [61] Margherita Carboni, Pia De Stefano, Bernd J. Vorderwülbecke, Sebastien Tourbier, Emeline Mullier, Maria Rubega, Shahan Momjian, Karl Schaller, Patric Hagmann, Margitta Seeck, Christoph M. Michel, Pieter Van Mierlo, and Serge Vulliemoz. Abnormal directed connectivity of resting state networks in focal epilepsy. *NeuroImage: Clinical*, 27:102336, 2020. doi:10.1016/j.nicl.2020.102336.
- [62] Ye Ren, Fengyu Cong, Tapani Ristaniemi, Yuping Wang, Xiaoli Li, and Ruihua Zhang. Transient seizure onset network for localization of epileptogenic zone: effective connectivity and graph theory-based analyses of ECoG data in temporal lobe epilepsy. *Journal of Neurology*, 266(4):844–859, April 2019. doi:10.1007/s00415-019-09204-4.
- [63] Preya Shah, John M. Bernabei, Lohith G. Kini, Arian Ashourvan, Jacqueline Boccanfuso, Ryan Archer, Kelly Oechsel, Sandhitsu R. Das, Joel M. Stein, Timothy H. Lucas, Danielle S. Bassett, Kathryn A. Davis, and Brian Litt. High interictal connectivity within the resection zone is associated with favorable post-surgical outcomes in focal epilepsy patients. *NeuroImage: Clinical*, 23:101908, 2019. doi:10.1016/j.nicl.2019.101908.
- [64] F. Scarselli, M. Gori, Ah Chung Tsoi, M. Hagenbuchner, and G. Monfardini. The Graph Neural Network Model. *IEEE Transactions on Neural Networks*, 20(1):61–80, January 2009. doi:10.1109/TNN.2008.2005605.

- [65] Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applications. *AI Open*, 1:57–81, 2020. doi:10.1016/j.aiopen.2021.01.001.
- [66] Kanchan Jha, Sriparna Saha, and Hiteshi Singh. Prediction of protein–protein interaction using graph neural networks. *Scientific Reports*, 12(1):8360, May 2022. doi:10.1038/s41598-022-12201-9.
- [67] Lingfei Wu, Yu Chen, Kai Shen, Xiaojie Guo, Hanning Gao, Shucheng Li, Jian Pei, and Bo Long. Graph Neural Networks for Natural Language Processing: A Survey. *Foundations and Trends® in Machine Learning*, 16(2):119–328, 2023. doi:10.1561/22000000096.
- [68] Si Zhang, Hanghang Tong, Jiejun Xu, and Ross Maciejewski. Graph convolutional networks: a comprehensive review. *Computational Social Networks*, 6(1):11, December 2019. doi:10.1186/s40649-019-0069-y.
- [69] Jie Chen, Tengfei Ma, and Cao Xiao. FastGCN: Fast Learning with Graph Convolutional Networks via Importance Sampling. 2018. Publisher: arXiv Version Number: 1. doi:10.48550/ARXIV.1801.10247.
- [70] William L. Hamilton, Rex Ying, and Jure Leskovec. Inductive Representation Learning on Large Graphs, September 2018. arXiv:1706.02216 [cs, stat]. URL: <http://arxiv.org/abs/1706.02216>.
- [71] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph Attention Networks, February 2018. arXiv:1710.10903 [cs, stat]. URL: <http://arxiv.org/abs/1710.10903>.
- [72] Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. Gated Graph Sequence Neural Networks. 2015. Publisher: arXiv Version Number: 4. doi:10.48550/ARXIV.1511.05493.
- [73] Muhan Zhang, Zhicheng Cui, Marion Neumann, and Yixin Chen. An End-to-End Deep Learning Architecture for Graph Classification. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1), April 2018. doi:10.1609/aaai.v32i1.11782.
- [74] Thomas N. Kipf and Max Welling. Semi-Supervised Classification with Graph Convolutional Networks, February 2017. arXiv:1609.02907 [cs, stat].

- [75] Ahmed Hossam Mohammed, Mercedes Cabrerizo, Alberto Pinzon, Ilker Yaylali, Prasanna Jayakar, and Malek Adjouadi. Graph neural networks in EEG spike detection. *Artificial Intelligence in Medicine*, 145:102663, November 2023. doi:10.1016/j.artmed.2023.102663.
- [76] Qi Lian, Yu Qi, Gang Pan, and Yueming Wang. Learning graph in graph convolutional neural networks for robust seizure prediction. *Journal of Neural Engineering*, 17(3):035004, June 2020. doi:10.1088/1741-2552/ab909d.
- [77] Zhengdao Li, Kai Hwang, Keqin Li, Jie Wu, and Tongkai Ji. Graph-generative neural network for EEG-based epileptic seizure detection via discovery of dynamic brain functional connectivity. *Scientific Reports*, 12(1):18998, November 2022. doi:10.1038/s41598-022-23656-1.
- [78] Daniele Grattarola, Lorenzo Livi, Cesare Alippi, Richard Wennberg, and Taufik A. Valiante. Seizure localisation with attention-based graph neural networks. *Expert Systems with Applications*, 203:117330, October 2022. doi:10.1016/j.eswa.2022.117330.
- [79] Siyi Tang, Jared A. Dunnmon, Khaled Saab, Xuan Zhang, Qianying Huang, Florian Dubost, Daniel L. Rubin, and Christopher Lee-Messer. Self-Supervised Graph Neural Networks for Improved Electroencephalographic Seizure Analysis. 2021. Publisher: arXiv Version Number: 2. doi:10.48550/ARXIV.2104.08336.
- [80] Yanna Zhao, Gaobo Zhang, Changxu Dong, Qi Yuan, Fangzhou Xu, and Yuanjie Zheng. Graph Attention Network with Focal Loss for Seizure Detection on Electroencephalography Signals. *International Journal of Neural Systems*, 31(07):2150027, July 2021. doi:10.1142/S0129065721500271.
- [81] Alan A. Díaz-Montiel and Milad Lankarany. Graph representations of iEEG data for seizure detection with graph neural networks. preprint, Neuroscience, June 2023. doi:10.1101/2023.06.02.543277.
- [82] Yimin Hou, Shuyue Jia, Xiangmin Lun, Ziqian Hao, Yan Shi, Yang Li, Rui Zeng, and Jinglei Lv. GCNs-Net: A Graph Convolutional Neural Network Approach for Decoding Time-resolved EEG Motor Imagery Signals. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–12, 2022. arXiv:2006.08924 [cs, eess]. doi:10.1109/TNNLS.2022.3202569.
- [83] Biao Sun, Han Zhang, Zexu Wu, Yunyan Zhang, and Ting Li. Adaptive Spatiotemporal Graph Convolutional Networks for Motor Imagery Classification.

- IEEE Signal Processing Letters*, 28:219–223, 2021. doi:10.1109/LSP.2021.3049683.
- [84] Tong Zhang, Xuehan Wang, Xiangmin Xu, and C. L. Philip Chen. GCB-Net: Graph Convolutional Broad Network and Its Application in Emotion Recognition. *IEEE Transactions on Affective Computing*, 13(1):379–388, January 2022. doi:10.1109/TAFFC.2019.2937768.
- [85] Richard B. Berry, Rita Brooks, Charlene Gamaldo, Susan M. Harding, Robin M. Lloyd, Stuart F. Quan, Matthew T. Troester, and Bradley V. Vaughn. AASM Scoring Manual Updates for 2017 (Version 2.4). *Journal of Clinical Sleep Medicine*, 13(05):665–666, May 2017. doi:10.5664/jcsm.6576.
- [86] John R. Ives. New Chronic EEG Electrode for Critical/Intensive Care Unit Monitoring:. *Journal of Clinical Neurophysiology*, 22(2):119–123, April 2005. doi:10.1097/01.WNP.0000152659.30753.47.
- [87] F. Spanedda, F. Cendes, and J. Gotman. Relations Between EEG Seizure Morphology, Interhemispheric Spread, and Mesial Temporal Atrophy in Bitemporal Epilepsy. *Epilepsia*, 38(12):1300–1314, December 1997. doi:10.1111/j.1528-1157.1997.tb00068.x.
- [88] Rina Zelman, Birgit Frauscher, Renzo Phellan Aro, Houssein-Eddine Gueziri, and D Louis Collins. SEEGAtlas: A framework for the identification and classification of depth electrodes using clinical images. *Journal of Neural Engineering*, 20(3):036021, June 2023. doi:10.1088/1741-2552/acd6bd.
- [89] H. G. Wieser, W. T. Blume, D. Fish, E. Goldensohn, A. Hufnagel, D. King, M. R. Sperling, H. Lüders, T. A. Pedley, and Commission on Neurosurgery of the International League Against Epilepsy (ILAE). ILAE Commission Report. Proposal for a new classification of outcome with respect to epileptic seizures following epilepsy surgery. *Epilepsia*, 42(2):282–286, February 2001.
- [90] John M Bernabei, Nishant Sinha, T Campbell Arnold, Erin Conrad, Ian Ong, Akash R Pattnaik, Joel M Stein, Russell T Shinohara, Timothy H Lucas, Dani S Bassett, Kathryn A Davis, and Brian Litt. Normative intracranial EEG maps epileptogenic tissues in focal epilepsy. *Brain*, 145(6):1949–1961, June 2022. doi:10.1093/brain/awab480.
- [91] Petr Klimes, Jan Cimbalnik, Milan Brazdil, Jeffery Hall, François Dubeau, Jean Gotman, and Birgit Frauscher. NREM sleep is the state of vigilance that best identifies the epileptogenic zone in the interictal electroencephalogram. *Epilepsia*, 60(12):2404–2415, December 2019. doi:10.1111/epi.16377.

- [92] Bartłomiej Chybowski, Petr Klimes, Jan Cimbalnik, Vojtech Travnicek, Petr Nejedly, Martin Pail, Laure Peter-Derex, Jeff Hall, François Dubeau, Pavel Jurak, Milan Brazdil, and Birgit Frauscher. Timing matters for accurate identification of the epileptogenic zone. *Clinical Neurophysiology*, 161:1–9, May 2024. doi:10.1016/j.clinph.2024.01.007.
- [93] Petr Nejedly, Jan Cimbalnik, Petr Klimes, Filip Plesinger, Josef Halamek, Vaclav Kremen, Ivo Viscor, Benjamin H. Brinkmann, Martin Pail, Milan Brazdil, Gregory Worrell, and Pavel Jurak. Intracerebral EEG Artifact Identification Using Convolutional Neural Networks. *Neuroinformatics*, 17(2):225–234, April 2019. doi:10.1007/s12021-018-9397-6.
- [94] J. Cimbalnik, P. Klimes, P. Nejedly, and V. Sladky. Epycom: ElectroPhYsiology COmputational Module, 2019. URL: <https://github.com/ICRC-BME/epycm>.
- [95] Radek Janca, Petr Jezdik, Roman Cmejla, Martin Tomasek, Gregory A. Worrell, Matt Stead, Joost Wagenaar, John G. R. Jefferys, Pavel Krsek, Vladimir Komarek, Premysl Jiruska, and Petr Marusic. Detection of interictal epileptiform discharges using signal envelope distribution modelling: application to epileptic and non-epileptic intracranial recordings. *Brain Topography*, 28(1):172–183, January 2015. doi:10.1007/s10548-014-0379-1.
- [96] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization, January 2017. arXiv:1412.6980 [cs]. URL: <http://arxiv.org/abs/1412.6980>.
- [97] Max Jaderberg, Valentin Dalibard, Simon Osindero, Wojciech M. Czarnecki, Jeff Donahue, Ali Razavi, Oriol Vinyals, Tim Green, Iain Dunning, Karen Simonyan, Chrisantha Fernando, and Koray Kavukcuoglu. Population Based Training of Neural Networks. 2017. Publisher: [object Object] Version Number: 2. doi:10.48550/ARXIV.1711.09846.
- [98] Tom Fawcett. An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8):861–874, June 2006. doi:10.1016/j.patrec.2005.10.010.
- [99] Kendrick Boyd, Kevin H. Eng, and C. David Page. Erratum: Area under the Precision-Recall Curve: Point Estimates and Confidence Intervals. In David Hutchison, Takeo Kanade, Josef Kittler, Jon M. Kleinberg, Friedemann Mattern, John C. Mitchell, Moni Naor, Oscar Nierstrasz, C. Pandu Rangan, Bernhard Steffen, Madhu Sudan, Demetri Terzopoulos, Doug Tygar, Moshe Y. Vardi, Gerhard Weikum, Camille Salinesi, Moira C. Norrie, and Óscar Pastor,

- editors, *Advanced Information Systems Engineering*, volume 7908, pages E1–E1. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013. Series Title: Lecture Notes in Computer Science. doi:10.1007/978-3-642-40994-3_55.
- [100] Frank Wilcoxon. Individual Comparisons by Ranking Methods. *Biometrics Bulletin*, 1(6):80, December 1945. URL: <https://www.jstor.org/stable/10.2307/3001968?origin=crossref>, doi:10.2307/3001968.
- [101] H. B. Mann and D. R. Whitney. On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other. *The Annals of Mathematical Statistics*, 18(1):50–60, March 1947. doi:10.1214/aoms/1177730491.
- [102] Päivi Nevalainen, Nicolás Von Ellenrieder, Petr Klimeš, François Dubeau, Birgit Frauscher, and Jean Gotman. Association of fast ripples on intracranial EEG and outcomes after epilepsy surgery. *Neurology*, 95(16), October 2020. doi:10.1212/WNL.0000000000010468.
- [103] Quinn McNemar. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*, 12(2):153–157, June 1947. doi:10.1007/BF02295996.
- [104] Haiteng Jiang, Vasileios Kokkinos, Shuai Ye, Alexandra Urban, Anto Bagić, Mark Richardson, and Bin He. Interictal SEEG Resting-State Connectivity Localizes the Seizure Onset Zone and Predicts Seizure Outcome. *Advanced Science*, 9(18):2200887, June 2022. doi:10.1002/advs.202200887.
- [105] OpenAI. ChatGPT. URL: <https://chat.openai.com/chat>.
- [106] DeepL SE. DeepL. URL: <https://www.deepl.com/translator>.
- [107] Grammarly Inc. Grammarly. URL: <https://www.grammarly.com>.

Symbols and abbreviations

DRE	Drug-Resistant Epilepsy
EZ	Epileptogenic Zone
SOZ	Seizure-Onset Zone
SEEG	Stereo-Electroencephalography
ASM	Antiseizure Medication
EMU	Epilepsy Monitoring Unit
GNN	Graph Neural Network
EEG	Electroencephalography
MEG	Magnetoencephalography
SOZ	Seizure-Onset Zone
fMRI	Functional Magnetic Resonance Imaging
EMU	Epilepsy Monitoring Unit
SPECT	Single Photon Emission Computed Tomography
MRI	Magnetic Resonance Imaging
PET	Positron Emission Tomography
ECoG	Electro-Corticography
LFP	Local Field Potential
HFO	High-Frequency Oscillation
IED	Interictal Epileptiform Discharge
CNN	Convolutional Neural Network
ASM	Antiseizure Medication
GraphSAGE	Graph Sample and Aggregate
GAT	Graph Attention Network
GGNN	Gated Graph Neural Network

GCN	Graph Convolutional Network
SAUH	St. Anne's University Hospital
MNI	Montreal Neurological Institute & Hospital
NREM	Non-Rapid Eye Movement
REM	Rapid Eye Movement
SVM	Support Vector Machine
PyG	PyTorch Geometric
LOPO CV	Leave-One-Patient-Out Cross-Validation
EPYCOM	ElectroPhYsiology COmputational Module
MLE	Maximal Likelihood Algorithm
REN	Relative Entropy
ReLU	Rectified Linear Unit
BCE	Binary Cross Entropy
TP	True Positive
TN	True Negative
FP	False Positive
FN	False Negative
ROC	Receiver Operating Characteristic
PR	Precision-Recall
AUROC	Area Under Receiver Operating Characteristic
AUPRC	Area Under Precision-Recall Curve
AP	Average Precision
PPV	Positive Predictive Value
IQR	Interquartile Range
STD	Standard Deviation

GIN

Graph Isomorphism Network

6 Disclosure of Artificial Intelligence Tool Usage

Several digital artificial intelligence tools were used during the development of this thesis to ensure its clarity and linguistic precision. ChatGPT-4 [105], DeepL [106] and Grammarly [107] tools were used for linguistic support, primarily for translation, paraphrasing, and the revision of grammar and style.

7 Electronic Attachment Content

The electronic attachment includes scripts describing the preprocessing of data, the implementation, training and testing of the GNN and reference models, and evaluation and visualization of results. Additionally, a directory containing the evaluation results is included, along with a file with general information about the project and the requirements file.

The scripts were developed using the PyCharm IDE and are compatible with Python version 3.8.18. The computations were performed on a server with an Intel Core i9 processor, which includes 14 cores and 28 threads. The server is equipped with 125 GB of RAM and includes two NVIDIA GeForce RTX 2080 Ti graphics cards, each with 11.2 GB of dedicated graphics memory.

The directory structure of the Graph Neural Networks project is as follows:

```
/.....root directory
├── feature_extraction/.....feature extraction scripts
│   ├── bivariate_FNUSA.py.....relative entropy FNUSA
│   ├── bivariate_MNI.py.....relative entropy MNI
│   └── detector_release13022022.py.....spike detector
├── epycomdevelop/.....feature extraction library
│   ├── spike_detection_FNUSA.py.....spike detection FNUSA
│   └── spike_detection_MNI.py.....spike detection MNI
├── GNNs_main/.....main GNN implementation scripts
│   ├── data_loader.py.....loads and preprocesses data
│   ├── dataset_prep.py.....prepares datasets for processing
│   ├── dataset_stats.py.....computes dataset statistics
│   ├── GNN_main.py.....main GNN script
│   ├── graph_data_loader.py.....loads graph data for processing
│   ├── model_eval.py.....evaluates model outputs
│   ├── model_pipeline.py.....model training and evaluation pipeline
│   ├── model.py.....defines and trains the GNNk models
│   ├── patient_results_plotting.py.....visualizes patient results
│   ├── reference_models.py.....runs reference models
│   ├── results_evaluator.py.....evaluates model predictions
│   └── results_visualizer.py.....visualizes the results as ROC and PR curves
├── results/.....model evaluation results
│   ├── AUPRC_evaluations.csv.....evaluation results for AUPRC
│   ├── AUROC_evaluations.csv.....evaluation results for AUROC
│   ├── F1-score_evaluations.csv.....evaluation results for F1-score
│   ├── model_outcomes.csv.....outcome predictions
│   ├── model_stats_outcomes.csv.....statistical analysis of model performances
│   ├── outcome_prediction_accuracy.csv.....statistics for outcome prediction
│   └── PPV_evaluations.csv.....evaluation results for PPV
├── README.md.....overview and general information about the project
└── requirements.txt.....lists all Python libraries required to run the scripts
```