

Università degli Studi di Salerno

.DIEM

Dipartimento di Ingegneria dell'Informazione ed Elettrica e Matematica Applicata

Dottorato di Ricerca in Ingegneria dell'Informazione
Ciclo 38

TESI DI DOTTORATO / PH.D. THESIS

Lightweight and Interpretable Deep Learning Framework for Personalized EEG-Based Epilepsy Monitoring

ROSANNA FERRARA

SUPERVISOR: **PROF. FRANCESCO TORTORELLA**
CO-SUPERVISOR: **PROF. MARTINO GIAQUINTO**

PHD PROGRAM DIRECTOR: **PROF. PASQUALE FOGGIA**

Anno 2026

Abstract

Deep Learning (DL) models make it possible to automate biosignal analysis. When dealing with Electroencephalography (EEG), a widely used tool to monitor brain activity, this translates into accurate detection of neurological events and scalable systems suitable for real-time applications. Within this framework, the present PhD work aims to design and optimize DL architectures tailored to patient-specific seizure detection, specifically addressing two main aspects of the problem. On the one hand, the proper management of the spatial complexity of EEG signals, which requires extracting meaningful information from multichannel data distributed across different brain regions; on the other hand, the reduction of the computational complexity of DL algorithms, in order to make their use feasible in a wearable context. This latter aspect is highly relevant, as it aligns with the growing interest in *personalized medicine* and remote healthcare, where the analysis of brain activity through continuous monitoring can help clinicians track the progression of neurological conditions more accurately.

The research activity is divided into two lines of investigation and is evaluated using two public datasets, i.e., the Children's Hospital Boston-Massachusetts Institute of Technology (CHB-MIT) database, widely adopted in the literature, and the Siena Scalp EEG dataset. The first study investigates how to effectively reduce the spatial dimensionality of EEG signals and exploit the resulting reduced input in a lightweight DL model trained on raw EEG data for supervised detection of epileptic seizures. To this end, we propose a data-driven selection algorithm that identifies the most informative scalp regions involved in seizure detection and a computationally efficient Convolutional Neural Network (CNN) that processes data

from only two EEG channels in a patient-specific setting, where both training and testing are performed on data from the same subject.

The two selected channels are representative of the two main phases of epileptic activity, i.e. ictal activity corresponding to labeled seizure events and no-seizure activity corresponding to the absence of ictal activity. The results show high performance compared to the state of the art. On the CHB-MIT dataset, the model achieves a balanced accuracy of 83% averaged across patients and an overall accuracy of 99%, correctly detecting 152 of 181 epileptic events and outperforming the same architecture applied to the four-channel fixed configuration commonly used in wearable EEG applications. Even on the Siena Scalp, the model achieves a balanced accuracy of 83% (with 31 of 36 epileptic events) compared to 78% obtained in the case of the fixed four-channel configuration.

Based on these results, the proposed method is further evaluated in a single-channel framework, in order to show the impact of choosing informative versus non-informative locations, and across various implementation strategies that address the patient-independent scenario, assessing how the model could generalize across patients. In the first evaluation, each of the two selected channels is tested individually, together with a non-informative EEG channel not highlighted by the channel selection procedure. The analysis demonstrates that using a non-informative channel leads to a 22% decrease in balanced accuracy on both datasets, confirming that not all channels carry the information required for seizure detection. Moreover, although the dual-channel configuration provides higher performance by combining complementary information from both ictal and no-seizure representative channels, the single-channel setup still offers a compelling trade-off between performance and simplicity. On the CHB-MIT dataset, the ictal-representative single channel achieves a balanced accuracy of 77% and detects 136 out of 181 epileptic events. Under the same training conditions, the no-seizure single channel reaches a balanced accuracy of 69% and detects 98 out of 173 events. Similar results are replicated on the Siena dataset, with a balanced accuracy of 81% for the ictal-representative channel and

77% for the no-seizure channel.

In the second evaluation, we investigate different implementation strategies for the patient-independent scenario. Specifically, we compare a fully personalized configuration, in which both the channel subset and the lightweight DL model are customized using only data acquired from a single patient; a fully generalized configuration, where both channel selection and model training are derived from a multi-patient dataset, and two hybrid approaches. In the first hybrid strategy, channel selection is performed on a large multi-patient dataset, while the model is trained on data from a specific patient. In the second hybrid strategy, the model is trained on a multi-patient dataset, whereas EEG channels are selected based on data from a specific patient. This analysis suggests that, although the fully customized approach achieves the best performance, an intermediate strategy, which consists of fixing the electrode layout while customizing the detection model, offers an effective balance between simplicity and accuracy, with a small decrease in balanced accuracy (80% vs. 83% on CHB-MIT and 73% vs. 83% on Siena Scalp).

The second study analyzes how to best exploit the spatial relationships between EEG electrodes in scenarios where the goal is to use information from all channels, rather than selecting only a few. Traditional CNNs, which assume input organized on a regular grid, are not ideal for this purpose, as EEG electrodes are arranged on a curved and irregular surface (the scalp) and do not follow a structured two-dimensional grid layout. For this reason, the use of Graph Neural Networks (GNNs) represents a more natural and effective approach to modeling EEG data, in which each channel is described as a node in the graph and the edges encode the functional connection (FC) between the electrodes. To identify the FC measure that best highlights differences between seizure phases while keeping the graph as sparse as possible, thereby enabling a lighter graph-based network fed by more selective and informative inputs, a data-driven statistical comparison of two widely used measures, i.e., phase locking value (PLV) and Pearson correlation coefficient (PCC), is carried out. As a result, PLV emerges as the most effective FC measure, and the developed architec-

ture, which combines Graph Convolutional Networks (GCN) for spatial feature extraction and a Long Short-Term Memory (LSTM) module for representing temporal dynamics, achieves a balanced accuracy of 97.7%, comparable to the state of the art while using more than 50% fewer parameters. Using PCC, performance is reduced to a balanced accuracy of 93.13% with the same proposed GNN model. Similar results are observed on the Siena Scalp-EEG dataset, where a balanced accuracy of 96.32% is obtained using PLV and 93.68% using PCC.

In conclusion, even if the main result of the entire PhD work is based on the implementation of lightweight DL frameworks for the extraction of meaningful descriptors of brain functioning, particular attention is paid to result visualization, which is essential for the clinical adoption of the models and for effective collaboration between human experts and Artificial Intelligence (AI) systems during the diagnostic process. Indeed, one of the main limitations to the adoption of AI in healthcare is the "black box" nature of state-of-the-art models. In this perspective, both proposed approaches generate intermediate outputs. The first returns EEG topographic maps, which allow the spatial distribution of the most relevant activity in distinguishing between the phases of the epileptic seizure to be visualized. The second, through the use of GNNExplainer, an interpretability tool that identifies the most influential nodes for model predictions, obtains relevance scores at the node level, maintaining a clear focus on the regions of the brain involved. These intermediate readings show strong consistency between the two approaches: graph-based explanation tools highlight regions of interest in accordance with the channels selected in the first study, reinforcing the reliability and cross-validation of the methodological paths adopted.

Abstract in italiano

I modelli di Deep Learning (DL) consentono di automatizzare l'analisi dei biosegnali. Nel caso dell'elettroencefalografia (EEG), uno strumento ampiamente utilizzato per monitorare l'attività cerebrale, ciò si traduce in un rilevamento accurato degli eventi neurologici e in sistemi scalabili adatti ad applicazioni in tempo reale. In questo contesto si inserisce il presente lavoro di dottorato che mira a progettare e ottimizzare architetture DL su misura per il rilevamento delle crisi epilettiche specifiche per ciascun paziente, affrontando in particolare due aspetti principali del problema. Da un lato, la corretta gestione della complessità spaziale dei segnali EEG, che richiede l'estrazione di informazioni significative da dati multicanale distribuiti in diverse regioni del cervello; dall'altro, la riduzione della complessità computazionale degli algoritmi DL, al fine di renderne possibile l'uso in un contesto indossabile. Quest'ultimo aspetto è molto rilevante, in quanto in linea con il crescente interesse per la *medicina personalizzata* ed il monitoraggio a distanza, dove l'analisi dell'attività cerebrale attraverso il monitoraggio continuo può aiutare i medici a seguire con maggiore precisione la progressione delle condizioni neurologiche.

L'attività di ricerca è suddivisa in due linee di indagine ed è valutata utilizzando due set di dati pubblici, ovvero il database del Children's Hospital Boston-Massachusetts Institute of Technology (CHB-MIT), ampiamente adottato in letteratura, ed il dataset Siena Scalp. Il primo studio indaga come ridurre efficacemente la dimensionalità spaziale dei segnali EEG e sfruttare l'input ridotto risultante in un modello DL leggero addestrato su dati EEG grezzi per il rilevamento supervisionato delle crisi epilettiche. A tal fine, si propone un algoritmo di selezione basato sui dati che identifica le regioni dello scal-

po più informative coinvolte nel rilevamento delle crisi e una rete neurale convoluzionale (CNN), efficiente dal punto di vista computazionale, che elabora i dati provenienti da soli due canali EEG in un contesto specifico per ciascun paziente, in cui sia l'addestramento che il test vengono eseguiti su dati provenienti dallo stesso soggetto. I due canali selezionati sono rappresentativi delle due fasi principali dell'attività epilettica, ovvero la fase ictale corrispondente agli eventi epilettici etichettati e la fase no-seizure corrispondente all'assenza di attività epilettica. I risultati mostrano prestazioni elevate rispetto allo stato dell'arte. Sul set di dati CHB-MIT, il modello raggiunge un'accuratezza bilanciata dell'83% in media tra i pazienti e un'accuratezza complessiva del 99%, rilevando correttamente 152 dei 181 eventi epilettici e superando la stessa architettura applicata alla configurazione fissa a quattro canali proposta in letteratura per le applicazioni EEG indossabili. Anche sul Siena Scalp, il modello raggiunge un'accuratezza bilanciata dell'83% (con 31 eventi epilettici su 36 rilevati) rispetto al 78% ottenuto nel caso della configurazione fissa a quattro canali.

Sulla base di questi risultati, il metodo proposto è stato ulteriormente valutato in un framework a canale singolo, al fine di mostrare l'impatto della scelta di posizioni informative rispetto a quelle non informative, e attraverso varie strategie di implementazione che affrontano lo scenario indipendente dal paziente, valutando come il modello potrebbe generalizzare tra i pazienti. Nella prima valutazione, ciascuno dei due canali selezionati viene testato individualmente, insieme a un canale EEG non informativo non emerso dalla procedura di selezione dei canali. L'analisi dimostra che l'utilizzo di un canale non informativo porta a una diminuzione del 22% dell'accuratezza bilanciata su entrambi i dataset, confermando che non tutti i canali trasportano l'informazione necessaria alla rilevazione dell'attacco. Inoltre, sebbene la configurazione a doppio canale fornisca prestazioni più elevate combinando informazioni complementari provenienti da canali rappresentativi sia ictali che no-seizure, la configurazione a canale singolo offre comunque un compromesso interessante tra prestazioni e semplicità. Sul dataset CHB-MIT, il canale singolo rappresentativo della fase ictale raggiunge un'accuratezza bilanciata del 77% e rileva

136 dei 181 eventi epilettici. Nelle stesse condizioni di addestramento, il canale singolo non epilettico raggiunge un'accuratezza bilanciata del 69% e rileva 98 dei 173 eventi. Risultati simili sono stati replicati sul dataset Siena Scalp, con un'accuratezza bilanciata del 81% per il canale rappresentativo della fase ictale e del 77% per il canale no-seizure.

Nella seconda valutazione, sono esaminate diverse strategie di implementazione per lo scenario indipendente dal paziente. Nello specifico, si confronta una configurazione completamente personalizzata, in cui sia la selezione dei canali che il modello DL leggero sono personalizzati utilizzando solo i dati acquisiti da un singolo paziente; una configurazione completamente generalizzata, in cui sia la selezione dei canali che l'addestramento del modello derivano da un set di dati multi-paziente; e due approcci ibridi. Nella prima strategia ibrida, la selezione dei canali viene eseguita su un ampio set di dati multi-paziente, mentre il modello viene addestrato sui dati di un paziente specifico. Nella seconda strategia ibrida, il modello viene addestrato su un set di dati multi-paziente, mentre i canali EEG vengono selezionati per un paziente specifico. Questa analisi suggerisce che, sebbene l'approccio completamente personalizzato raggiunga le migliori prestazioni, una strategia intermedia, che consiste nel fissare la disposizione degli elettrodi personalizzando il modello di rilevamento, offre un equilibrio efficace tra semplicità e accuratezza, con una leggera diminuzione dell'accuratezza bilanciata (80% contro 83% sul CHB-MIT e 73% contro 83% sul Siena Scalp).

Il secondo studio analizza come sfruttare al meglio le relazioni spaziali tra gli elettrodi EEG in scenari in cui l'obiettivo è quello di utilizzare le informazioni provenienti da tutti i canali, piuttosto che selezionarne solo alcuni. Le CNN tradizionali, che presuppongono un input organizzato su una griglia regolare, non sono ideali per questo scopo, poiché gli elettrodi EEG sono disposti su una superficie curva ed irregolare (lo scalpo) e non seguono una disposizione strutturata su una griglia bidimensionale. Per questo motivo, l'uso delle reti neurali grafiche (GNN) si presta come un approccio più naturale ed efficace per modellare i dati EEG, in cui ogni canale è descritto come un nodo nel grafo ed i bordi codificano la connessione funzionale

(FC) tra gli elettrodi. Per identificare la misura FC che meglio evidenzia le differenze tra le fasi delle crisi epilettiche mantenendo il grafo il più sparso possibile, consentendo così una rete basata su grafici più leggera alimentata da input più selettivi ed informativi, si effettua un confronto statistico basato sui dati di due misure ampiamente utilizzate, ovvero il valore di Phase locking value (PLV) ed il Coefficiente di Correlazione di Pearson (PCC). Per quanto riguarda i risultati, il PLV emerge come la misura FC più efficace e l'architettura sviluppata, che combina reti convoluzionali a grafo (GCN) per l'estrazione di caratteristiche spaziali ed un modulo di Long Short-Term Memory per rappresentare le dinamiche temporali, raggiunge un'accuratezza bilanciata del 97.7%, paragonabile allo stato dell'arte, utilizzando oltre il 50% in meno di parametri. Utilizzando il PCC, le prestazioni scendono ad un'accuratezza bilanciata del 93.13% con lo stesso modello GNN. Risultati simili si osservano sul dataset Siena Scalp, dove si ottiene un'accuratezza bilanciata del 96.32% con PLV e del 93.68% con PCC.

In conclusione, sebbene il risultato principale del lavoro di dottorato si basi sull'implementazione di framework DL leggeri per l'estrazione di descrittori significativi del funzionamento del cervello, particolare attenzione è rivolta alla visualizzazione dei risultati, essenziale per l'adozione clinica dei modelli e per un'efficace collaborazione tra esperti e sistemi di Intelligenza Artificiale (AI) nel processo diagnostico. Infatti, uno dei principali limiti all'adozione dell'IA in ambito sanitario è la natura "scatola nera" dei modelli. In questa prospettiva, entrambi gli approcci proposti generano output intermedi. Il primo restituisce mappe topografiche che consentono di visualizzare la distribuzione spaziale dell'attività più rilevante nel distinguere le fasi della crisi. Il secondo, attraverso l'uso di GNNExplainer, uno strumento di interpretabilità che identifica i nodi più influenti per le previsioni del modello, ottiene punteggi di rilevanza a livello di nodo, mantenendo una chiara attenzione sulle aree cerebrali coinvolte. Queste letture intermedie mostrano una forte coerenza tra i due approcci: gli strumenti di spiegazione basati su grafo evidenziano regioni di interesse in accordo con i canali evidenziati nel primo studio, rafforzando l'affidabilità e la validazione incrociata dei percorsi metodologici adottati.

Acknowledgment

I am pleased to have had the opportunity to pursue my Ph.D. in Computer Science at the University of Salerno. I wish to express my deepest and most sincere gratitude to the entire MIVIA group for welcoming and encouraging me, especially in overcoming the initial challenges arising from my different biomedical background. This experience has been, for me, not only a pivotal stage in my academic development but also a journey of personal growth that I will always cherish and value.

First and foremost, I would like to thank Prof. Mario Vento, Prof. Francesco Tortorella, Prof. Gennaro Percannella, and Prof. Pasquale Foggia for giving me the opportunity to join their research group and for patiently guiding me in both my scientific work and my personal growth. Prof. Francesco Tortorella, in his role as supervisor, taught me the importance of responsibility and critical thinking, providing me with guidance, assistance, and support throughout my entire PhD journey. I am also grateful to Prof. Alessia Saggese for constantly offering me thought-provoking insights and a different perspective, which proved invaluable in broadening my research horizons and improving the quality of my work.

I would like to extend my special thanks to my co-supervisor, Dr. Martino Giaquinto, for his meticulous guidance, unwavering support, and attentive presence throughout every stage of my research journey. Without his generous mentorship, my path would have been far more challenging and demanding. My sincere thanks also go to Prof. Benoit Gauzere and Prof. Florian Yger, together with their research group at INSA Rouen, for the valuable scientific and personal support I received during my research stay abroad as part of my PhD program.

Finally, I would like to express my heartfelt gratitude to my family and, in particular, dedicate a few words to someone who was my point of reference and my greatest certainty: my grandmother. A strong woman, remarkable in her humility, she taught me, through her quiet strength, the deepest meaning of freedom. She was an example of a life lived with commitment, sacrifice, dedication, and extraordinary endurance in the face of pain. This journey was shared with her, through every joy and every hardship. Every evening, when I came home, she was ready to listen to the story of my day; every morning, with her affection and encouragement, she wished me great things. She was my anchor, my refuge, and my constant source of motivation. The fact that I defended this PhD thesis on the day of her name day made me feel even more strongly that, although she passed away only recently, she did not want to miss this milestone of ours. I felt her presence beside me, as if she were still there, supporting me, encouraging me, and sharing in the joy of this journey.
I feel immensely fortunate to have had her by my side.

List of Acronyms

ACC - Accuracy
AI - Artificial Intelligence
BCI - Brain–Computer Interface
bAcc - Balanced Accuracy
CHB-MIT - Children's Hospital Boston–Massachusetts Institute of Technology
CNN - Convolutional Neural Network
CWC - Channel Weight Coefficient
DGNNs - Dynamic Graph Neural Networks
DL - Deep Learning
DWT - Discrete Wavelet Transform
EC - Effective Connectivity
ECG - Electrocardiography
EEG - Electroencephalography
FC - Functional Connectivity
FGS - Fully Generalized System
FN - False Negative
FP - False Positive
FP/h - False Positives per Hour
FPS - Fully Personalized System
GAT - Graph Attention Network
GCN - Graph Convolutional Network
GNN - Graph Neural Network
GRU - Gated Recurrent Unit
HSSC - Hybrid System-Specific Channels
HSSM - Hybrid System-Specific Model
ILAE - International League Against Epilepsy
LORO - Leave-One-Record-Out
LOSO - Leave-One-Subject-Out
LSTM - Long Short-Term Memory

ML - Machine Learning
NAS - Neural Architecture Search
PCA - Principal Component Analysis
PCC - Pearson's Correlation Coefficient
PCs - Principal Components
PLV - Phase-Locking Value
RF - Random Forest
RNNs - Recurrent Neural Networks
SC - Structural Connectivity
SN - Sensitivity
SP - Specificity
STFT - Short-Time Fourier Transform
STGNN - Spatio-Temporal Graph Neural Network
SUDEP - Sudden Unexpected Death in Epilepsy
SVM - Support Vector Machine
TN - True Negative
TP - True Positive
TUSZ - Temple University Hospital Seizure Corpus

Contents

1 Introduction **3**
1.1 Motivation . . . 3
1.2 Contributions of the thesis . . . 8
1.3 Organization of the thesis . . . 10

2 Background **15**
2.1 Electroencephalography . . . 15
2.1.1 Conventional brain rhythms . . . 16
2.1.2 Electrode placement system . . . 18
2.1.3 EEG montage . . . 19
2.1.4 Artifacts . . . 21
2.2 Epilepsy . . . 21
2.2.1 Introduction . . . 22
2.2.2 Seizure-related phases . . . 23
2.2.3 Seizure classification problems . . . 24
2.3 Graph-based EEG Analysis . . . 29
2.3.1 Brain Connectivity . . . 29
2.3.2 Graph construction . . . 31
2.3.3 Graph Neural Networks . . . 36

3 Literature review **41**
3.1 Introduction . . . 41
3.2 Dataset . . . 48
3.3 Lightweight approaches . . . 51
3.3.1 Channel selection . . . 58
3.3.2 EEG wearable devices . . . 60
3.4 GNN-based approaches . . . 63
3.5 Summary . . . 67

4 Proposed methods to manage complexity **73**
4.1 Reduced-Channel Framework 73
4.1.1 Introduction 75
4.1.2 Related works 76
4.1.3 Methods 79
4.1.3.1 Channel selection algorithm . . 80
4.1.3.2 Classification model 88
4.2 Full-Channel Framework 91
4.2.1 Introduction 91
4.2.2 Related work 94
4.2.3 Method 96
4.2.3.1 Connectivity evaluation 97
4.2.3.2 Graph computation 103
4.2.3.3 Graph-based model 104
4.2.3.4 GNN-explainer 106
4.3 Evaluation metrics 107

5 Validation of reduced-channel approach **111**
5.1 Optimized dual channel analysis 111
5.1.1 Implementation details 112
5.1.2 Results 115
5.1.3 Discussion 122
5.2 Single-Channel Analysis 134
5.2.1 Implementation details 134
5.2.2 Results 136
5.2.3 Discussion 138
5.3 Generalization Analysis 144
5.3.1 Implementation details 145
5.3.2 Results 148
5.3.3 Discussion 152
5.4 Conclusion 157

6 Validation of GNN-based approach **161**
6.1 Implementation details 161
6.2 Results 164
6.2.1 Connectivity results 164
6.2.2 Classification performance 172
6.3 Discussion 174
6.4 Conclusion 184

7 Conclusion **187**
7.1 Main contributions 188
7.2 Limitations 193
7.3 Future directions 195

Bibliography **197**

List of Figures **220**

List of Tables **225**

In loving memory of my mother, whose guidance and unwavering support have shaped the person I am today. Though she is no longer with me, her spirit and wisdom live on in every page of this thesis. I dedicate this work to her memory, as a tribute to the profound impact she had on my growth and academic journey.

To my dear husband, whose strength and love have been a constant source of inspiration, this milestone is a testament to the values we are building together. Thank you for being the pillar of my life: even in the face of loss, you have given me the strength to move forward.

Rosanna

Chapter 1

Introduction

1.1 Motivation

Electroencephalography (EEG) is a technique used to record the electrical activity of the brain through electrodes placed on the scalp of the patient. It is widely adopted by researchers and practitioners in neuroscience, neurology and psychology due to its non-invasive nature and its ability to provide valuable information that offers insights into a wide range of cognitive processes and neurological disorders. However, despite its high potential, the EEG signal is more challenging to interpret, not only because it evolves both in time and space, but also because its waveform is less stereotypical than other biosignals, such as electrocardiographic (ECG) or photoplethysmographic signals, which have well-defined patterns.

Consequently, although traditional methods based on spectral analysis have the advantage of being reliable and well established in the current clinical practice, EEG analysis still presents significant analytical challenges. First, EEG signals are highly variable and non-stationary, meaning that patterns of brain activity change over time and depend strongly on the experimental conditions. Moreover, brain activity differs considerably between individuals and this explain why traditional approaches are often limited in capturing the full richness of the information contained in the signal, making it necessary to adopt more advanced techniques for effective interpretation.

In response to these limitations, in recent years, machine

learning (ML) methods and, in particular, deep learning (DL) techniques are emerging as powerful tools for learning informative representations from EEG data, improving their robustness in analysis and interpretation. In the literature, these methodologies have proven particularly effective in the automated processing of EEG signals, the identification of complex patterns, and even the prediction of neurological disorders and cognitive states, although important questions remain open. Indeed, when applying artificial intelligence (AI) models to EEG signals, the accurate representation of spatial characteristics, the ability to generalize to different subjects and experimental conditions, and the need to improve the interpretability of results remain critical. This last aspect is particularly important in medical and healthcare contexts. Indeed, although traditional DL models are known for their accuracy, they often operate as "black boxes". This lack of transparency in their outputs and decision-making processes is particularly critical, as understanding the reasoning behind predictions is essential to strengthen trust in AI-based tools and to promote their adoption in clinical practice.

In parallel with the application of AI models to EEG, recent developments in wearable technologies have also expanded the use of EEG for *telemedicine* and remote patient monitoring. EEG devices with dry electrodes are capable of detecting conditions such as epilepsy and sleep disorders, or monitoring post-stroke rehabilitation, without the need for frequent hospital visits. However, despite their well-known potential, wearable EEG systems have not yet reached the level of integration and usefulness of wearable sensor patches and smartwatches, which are already widely used to monitor respiratory, motor, and cardiovascular biosignals in both clinical and everyday-life contexts. Bringing EEG to wearable devices requires particular attention to the accurate management of the spatial information held in the signals. On the one hand, one possibility is to reduce the number of electrodes and explicitly study the trade-off between information loss and greater system comfort. On the other hand, when all channels are maintained, the main challenge becomes optimizing the representation of their spatial relationships in order to preserve the most informative pat-

terns while reducing the redundancy and computational complexity of the system.

Both strategies described above play a central role in this research, which arises from the need to integrate EEG acquisition systems with AI-based diagnostic tools for signal analysis, opening up important opportunities in the clinical field. Such integration could, for example, enable the automatic identification of anomalies such as epileptic seizures, thus supporting both neurologists and patients in clinical care. The overall objective is to develop automatic detection models capable of operating in real time on systems optimized for resource-limited environments, and therefore subject to low power and computational constraints. A further purpose is the evaluation of compact and highly discriminative input representations, able to effectively capture and manage the spatial complexity of EEG signals for the classification task, before they are fed into the DL models. Identifying the right data representation is also useful for improving the interpretability of data and results, which is essential in healthcare contexts since systems must ensure that their outputs can be understood even by users who are not AI experts. This line of research is also supported by recent technological advances, which are moving in the same direction by bringing wearable EEG devices [1] to the market, although current solutions still focus primarily on signal acquisition alone, without integrating advanced analysis processes or real-time automatic analysis.

The study is divided into two lines of investigation. The first concerns reducing the computational complexity of the network by decreasing the number of input EEG channels, with the aim of adopting lighter models by selecting the minimum set of channels that provide the most information. The second line, which is more exploratory in nature, starts from the idea that brain function is intrinsically distributed and cannot be studied by analyzing individual units in isolation. From this perspective, we investigate the possibility of reducing the computational complexity of Graph Neural Networks (GNN) through a careful study of the most suitable measure for modeling, in graph representation, the relationships between different brain regions. In this model, electrodes are represented

as nodes and their interactions as edges, typically defined by functional connectivity (FC) measures, which capture statistical dependencies between brain regions.

Additionally, to provide effective support for the discussion of results, particular attention is also paid to the correct visualization of outputs, an essential aspect in the healthcare context towards effective collaboration between human experts and AI-based systems during the diagnostic process. In this perspective, to represent the spatial distribution of brain activity, the first approach produces EEG topographic maps, while the second uses GNNExplainer, an interpretability tool that identifies the most influential EEG electrodes for model predictions.

The starting point for the entire project is the identification of a suitable case study. To this end, a review of the literature covering the last five years carried out using the search keywords ("Artificial Intelligence" OR "Machine Learning" OR "Deep Learning") AND "EEG" reveals several key areas of interest. These areas include brain-computer interfaces (BCI) [2], sleep stages [3], monitoring of epileptic seizures [4], emotion recognition [5], and neuro-degenerative disorder detection [6, 7], among others. These typical applications are illustrated in Figure 1.1. Each area typically involves classification tasks but differs in terms of the number of classes and data generation mechanisms. For instance, sleep monitoring pertains to biological events (event level), whereas pathology screening focuses on individual patients in relation to population-level data (subject level).

Within this range of possible applications, epilepsy is identified as the main focus of this study for several reasons. Firstly, automatic seizure detection systems have not yet reached the level of reliability required to replace manual annotation by experts. Consequently, a fundamental clinical need for accurate and reliable seizure detection systems clearly emerges. Additionally, the growing body of research on wearable EEG systems for epilepsy further underlines this need. The development of wearable EEG devices specifically designed for epilepsy monitoring has become a key objective of recent technological advances, and the literature emphasizes the importance of these solutions due to the growing demand for long-term

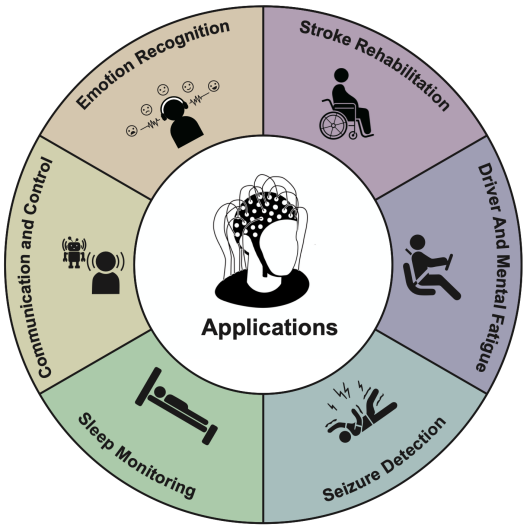


Figure 1.1: The six most typical EEG applications [8].

monitoring [9]. This demand is supported by a recent online survey of 40 healthcare professionals, who emphasizes the critical attributes of high diagnostic performance, data quality, patient mobility, and comfort in long-term EEG monitoring systems [10]. Notable benefits mentioned include more natural monitoring environments, reduced application times, earlier diagnosis of epilepsy, and decreased support needs. Innovative EEG devices are regarded as essential for diagnosing epilepsy, monitoring seizures, documenting seizures automatically, raising alarms, and predicting seizures. Consequently, there is significant interest in integrating them into clinical practice, particularly for managing drug-resistant epilepsy and reducing the risk of sudden unexpected death in epilepsy (*SUDEP*) [10], also because, at present, the most common method for real-time home monitoring of epileptic patients is video-EEG telemetry [11]. In this approach, the patient or caregiver supervises the process, which requires the simultaneous acquisition of video signals and EEG data. However, this method does not provide the same practicality and portability offered by next-generation wearable technologies.

Furthermore, another important aspect is that there is a growing interest within the research community in applying

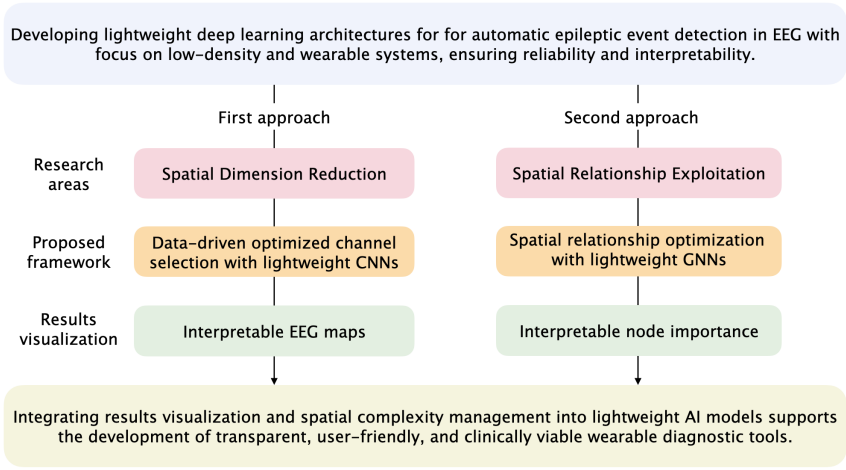


Figure 1.2: An overview of the targets and contributions of the thesis, including the objectives, research areas, proposed methods, interpretability aspects, and conclusion.

AI algorithms to EEG data for epilepsy detection and monitoring. Numerous recent competitions are arising in this area highlighting advancements in AI-driven solutions for epilepsy, such as the *AI-based Seizure Detection With Wearable EEG competition* organized by the IEEE Signal Processing Society [12]. This trend is evident not only in traditional AI applications on EEG but also in the integration of AI within wearable devices, indicating a substantial interest in advancing the field.

1.2 Contributions of the thesis

Figure 1.2 summarizes the contributions and objectives of this thesis, highlighting the research areas explored, the proposed methods, and the strategies adopted to visualize the results. The thesis begins with an in-depth literature review and then presents two complementary methods designed to tackle the crucial challenge of reducing model computational complexity while effectively managing spatial complexity in EEG data analysis.

The first study proposes the use of a data-driven channel selection algorithm to reduce the number of channels for the

task of detecting epileptic seizures [13]. The framework adopts a two-stage strategy: (i) selection of the two most informative EEG channels for identifying the epileptic phase, guided by interpretability criteria; (ii) experimental validation on two public datasets, demonstrating that the developed method significantly improves the performance of personalized seizure detection compared to the four-channel fixed configuration proposed in literature for wearable purposes. The contribution includes:

- channel visualization: performed using topographic maps to highlight the EEG channels with the highest contribution in the two epileptic phases for the specific patient, i.e. presence and absence of seizure events.
- lightweight architecture application: a DL architecture optimized for embedded systems with resource constraints, combining three temporal convolutional layers and one spatial layer to enhance the accuracy and timeliness of epilepsy diagnosis from two selected channels.

Based on the results obtained, the impact of using a single channel on seizure detection performance, as well as the system's flexibility in moving away from a patient-specific approach to adapt to a broader range of patients, is also investigated. In particular, applying the same lightweight model in a single-channel configuration allows us to assess whether channel selection has a tangible impact on seizure detection [14]. In this case, the analysis includes a comparative channel analysis among three settings: a single channel associated with epileptic activity, a single channel associated with the absence of epileptic activity, and a channel not involved in either phase. On the other hand, since the channel selection algorithm operates independently of the downstream lightweight detection system, exploring different levels of personalization and generalization allows us to investigate the practicality of the proposed detection system [15]. In this case, the analysis includes a comparative channel analysis among four different implementation strategies. Starting from the *fully personalized* system of the first study, in which both the model and

the channel selection are patient-specific, the other strategies considered include the *fully generalized* methodology in which both channel selection and model training are derived from a multi-patient dataset, and two *hybrid solutions*, the first focused on multi-patient channel selection with a patient-specific DL model, and the second on patient-specific channel selection with a multi-patient generalized model.

The second study investigates the optimal graph-based representation of EEG signals to enable the exploitation of full spatial information for capturing underlying brain activity in seizure detection. The framework adopts a two-stage strategy: (i) selection of the most appropriate functional connectivity (FC) measure, chosen among the most commonly used in the literature to represent spatial relationships in graph-based EEG, guided by statistical evaluations; and (ii) experimental validation on two public datasets, demonstrating performance comparable to state-of-the-art baseline approaches while employing a significantly lighter model, obtained by feeding it input graphs that compactly represent EEG activity across different epileptic phases. The contribution includes:

- channel visualization: performed using GNNExplainer, a dedicated interpretability tool that identifies which channels the model focuses on during the detection process.
- lightweight architecture application: a DL architecture optimized for reduced resource constraints, combining a spectral-based Graph Convolutional Network (GCN) to model spatial relationships among channels and a Long Short-Term Memory (LSTM) module to capture temporal dependencies.

1.3 Organization of the thesis

The dissertation consists of seven chapters, structured as illustrated in Figure 1.3.

Chapter 2, Background, provides the basic knowledge and fundamental concepts necessary to acquire the background in-

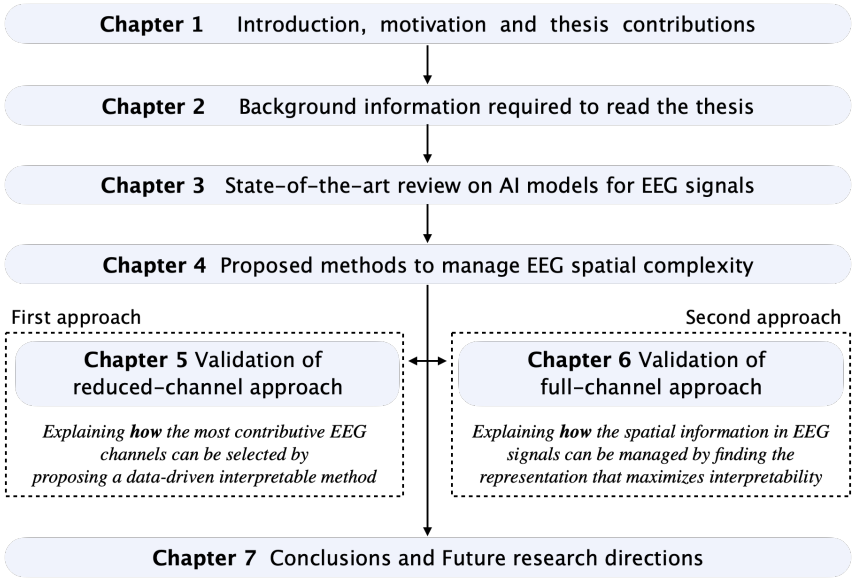


Figure 1.3: Thesis structure.

formation required to read this dissertation. First, the neurological condition of epilepsy is described. EEG as a recording technique and the resulting signal is also outlined. Finally, how advancements in technology are reinforcing the trend towards smarter EEG wearable devices are discussed. The Chapter also introduces the concept of brain connectivity and presents the main methods found in the literature for representing brain graphs from EEG.

Chapter 3, Literature review, presents a systematic review of the state of the art across all major lines of research on AI models for EEG analysis in the field of epilepsy. It begins by describing the datasets identified in the literature and used in this thesis. The Chapter is then structured into two parts, to reflect the two parallel studies on which the thesis focuses. The first part reviews lightweight models for resource-constrained scenarios and the most widely used EEG channel selection methods, while the second part critically assesses the use of GNNs for EEG analysis. Finally, the Chapter highlights the gaps that emerged from the review and clarifies how this thesis proposes to address them, thus precisely framing the

challenges addressed.

Chapter 4, Proposed methods to manage complexity, introduces the methodological aspects of the two proposed approaches. The first focuses on a method that identifies the two EEG channels that best capture the variance associated with the dynamics of epileptic seizures. These channels are then fed into a lightweight Convolutional Neural Network optimized for a small number of channels and designed for potential use on embedded systems. The second proposes a graph-based modeling of the full scalp EEG in which edges encode functional connectivity. Through a statistical comparison between phase synchronization and Pearson's linear correlation measures, the most suitable connectivity metric is selected for an interpretable distinction between epileptic and non-epileptic activity, thus allowing the use of a less complex Graph Neural Network. The proposed architecture combines Graph Convolutional Networks with a dynamic recurrent module and includes an explainability component based on GNNExplainer, which learns masks over node features and edges to highlight the electrodes that are most decisive for the prediction made by the network on a given graph.

Chapter 5, Validation of reduced-channel approach, presents the experimental validations of the reduced-channel framework. It starts from the optimized two-channel analysis which demonstrates performance comparable to, and in some cases superior to, state-of-the-art systems based on four channels in fixed positions. Based on this, two follow-up studies are described to further strengthen the framework. The first applies the channel selection method to quantify the impact of using a single EEG channel for detecting epileptic seizures, thereby further evaluating the efficiency and robustness of the algorithm. The second explores the transition to a subject-independent configuration, exploiting the decoupling between channel selection and model training introduced in Chapter 4.1. The purpose is to evaluate a practical trade-off between detection performance and the design of wearable devices targeted for multi-patient use rather than single-patient personalization.

Chapter 6, Validation of full-channel approach, presents the experimental validation of the graph-based multichannel framework. It starts from the optimized selection of the connectivity measure and a justifiable sparsification of the graph, demonstrating that this yields a lightweight network matching state-of-the-art performance with significantly fewer parameters. Then, GNNExplainer results are estimated to provide valuable insights into neural dynamics, deepening the understanding of epileptic activity and allowing the validation of the channel selection procedure proposed in Chapter 4.1.

Chapter 7, Conclusion, concludes the thesis and discusses further plans and road maps for the field of research. The concluding remarks include a summary of the results achieved, a critical analysis of the difficulties and limitations of the study, the identification of open questions and the definition of possible future developments.

Chapter 2

Background

This section will outline the problem and background information required to read this dissertation. First, Electroencephalography (EEG) as a recording technique and the characteristics of the resulting signal are described in Section 2.1. An overview of the main functional aspects of this widely used method for studying brain activity is provided, including electrode placement according to the international 10–20 system and the main EEG montage and referencing configurations. Then, the neurological condition of epilepsy is described in Section 2.2, while Section 2.3 discusses the main graph-based approaches for modeling EEG signals.

2.1 Electroencephalography

Scalp EEG is a non-invasive technique used to record the brain's electrical activity via electrodes placed on the scalp. It is often favored over more invasive approaches, such as intracranial EEG or stereotactic EEG (sEEG). Intracranial recordings, typically employed in neurosurgical investigations, allow for the measurement of neuronal firing patterns directly from within the brain [16]. In contrast, sEEG involves implanting electrodes into deep brain structures, enabling the monitoring of electrical activity that cannot be effectively captured from the scalp surface [17]. In this section, we provide an overview of the EEG paradigms, including the clinical EEG recording technique, the conventional EEG frequency rhythms and the 10-20

electrode placement system.

Scalp EEG was first introduced in 1929 by the German psychiatrist Hans Berger [18]. His key discovery was the so-called "alpha rhythm", an oscillatory component observed in the occipital regions at approximately 10 Hz, which remains one of the most extensively studied EEG rhythms today. Scalp EEG plays a crucial role in capturing electrical signals that reflect both neurophysiological processes and neuropathological conditions of the human brain. Thanks to its non-invasive nature and excellent temporal resolution, it is a valuable tool in clinical applications, enabling real-time monitoring of brain activity.

A typical clinical EEG acquisition involves 256 sensors, called electrodes arranged in a cap that can be easily positioned on the human scalp. Each single electrode, accurately positioned and secured using conductive gel, measures the extracellular currents generated by synchronized neuronal activity. These currents propagate through various cortical regions, revealing the underlying mechanisms of information processing in the brain and giving rise to electrical potentials large enough to be measured on the scalp. However, since the electrical potentials propagate through heterogeneous media, including the cerebrospinal fluid, skull, and scalp, the resulting EEG signals become spatially distorted. Consequently, the information picked up at a specific location is actually the summation of multiple electric fields originating from different regions of the brain.

EEG signals recorded at the scalp are usually on the order of 10 μV to 100 μV and are typically digitized at a sampling frequency that should be twice as high as the highest frequency component of interest of the signal. Common values fall within the range of 128 to 1024 Hz. The requirement follows the Nyquist-Shannon sampling theorem, which ensures that all relevant information is captured and that the recorded signal does not suffer from aliasing artifacts.

2.1.1 Conventional brain rhythms

Historically, EEG has been analyzed in the frequency domain, quantifying the total power distributed across different bands.

Table 2.1: A detailed analysis of conventional EEG frequency bands and their associated cognitive and physiological states. The bands typically fall within a range of 0.5 Hz to 80 Hz [19].

Band	**Range (Hz)**	**Cognitive/physiological states**
delta (δ)	0.5–4	Deep sleep, low-frequency activity
theta (θ)	4–8	Relaxation, drowsiness, meditation
alpha (α)	8–13	Relaxed wakefulness, eyes closed, mental inactivity
beta (β)	13–30	Cognitive activity, concentration, mental effort
gamma (γ)	30–80	Perception, conscious awareness, complex sensory processing

This approach is still the most widely used method today, as it plays an important role in clinical practice and can be interpreted immediately by doctors and practitioners. Over the years, specific frequency bands have been conventionally defined to categorize brain electrical activity in both physiological and pathological conditions. These frequency bands are therefore commonly referred to as *brain rhythms*. The five most commonly used brain rhythms are shown in Table 2.1 and described below [19].

Delta waves occur in a frequency range between 0.5 and 4 Hz and are usually observed during deep sleep. Being low-frequency waves, they can easily be confused with movement artifacts, given their similar nature. Delta waves have also been associated with sustained attention tasks. Theta waves have a range between 4 and 8 Hz and are linked to states of drowsiness and deep meditation.

Alpha waves are the most studied in EEG. They range between 8 and 14 Hz and are generally visible in the occipital lobe. They are the most common indicator of a relaxed mental state and are associated with closed eyes. States of anxiety or attention significantly reduce the presence of alpha waves. Beta waves are in the frequency range between 14 and 30 Hz and are associated with active thinking, anxiety, alertness and

concentration.

Finally, gamma waves are the highest frequency waves, ranging from approximately 30 Hz upwards. They are believed to play a complex role in brain functions, such as the integration of information from different sensory inputs. They are also used as markers in certain brain disorders.

2.1.2 Electrode placement system

The electrodes are placed on the scalp according to the international 10-20 system, which ensures standardized placement and, consequently, reproducibility of EEG measurements. This system, recommended by the International Federation of Societies for Electroencephalography and Clinical Neurophysiology (IFCN) [20], establishes that the distances between adjacent electrodes correspond to 10% or 20% of the total cranial distance, calculated between specific *anatomical landmarks* : i.e, nasion, inion and pre-auricular points. The nasion is the point of intersection between the frontal bone and the nasal bones, while the inion is the most prominent protrusion of the occipital bone.

Once the nasion-inion distance has been measured, the head is divided into segments as follows. The first one located 10% from the nasion, the last 10% from the inion, and the four central points are located at successive intervals of 20% of the nasion-inion distance. The same criterion applies to the distance between the left and right pre-auricular points passing through the vertex (Cz). The result of this mapping is shown in Figure 2.1.

Following this measurements, the positions of the electrodes on the scalp are labeled accordingly. Each label consists of one or two letters, followed by a number. The letters represent the areas of the brain, namely frontal (F), temporal (T), parietal (P), central (C) and occipital (O). Additional areas are frontal-polar (Fp) and fronto-temporal (FT). As for the numbers, odd values are located on the left side of the brain, while even values are located on the right hemisphere. Z electrodes lie along the midline of the scalp.

As can be seen in the diagram of the 10-20 electrode place-

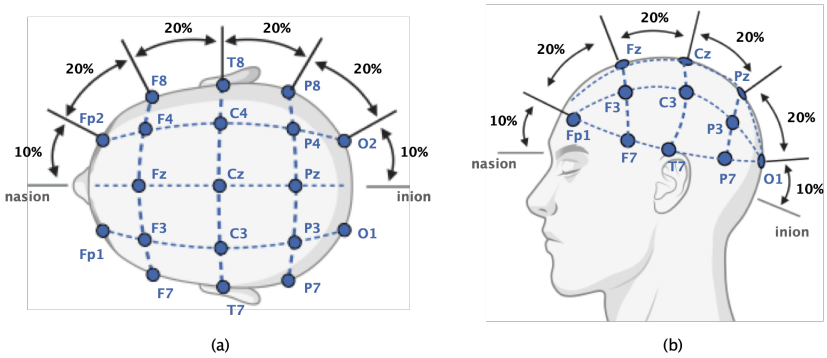


Figure 2.1: The 10-20 electrode placement system based on electrode distances. (Created with app *BioRender* [21])

ment system shown in Figure 2.1, all electrodes on the sagittal line are labeled with number 7 if they are located in the left hemisphere and with number 8 if they are located in the right hemisphere. The only exceptions are the Fp1/Fp2 and O1/O2 electrodes. This representation corresponds to the new nomenclature introduced in accordance with the guidelines of the American Clinical Neurophysiology Society [22]. Previously, the mid-temporal electrodes T7 and T8 were referred to as T3 and T4, respectively, while the parietal electrodes P7 and P8 were denoted as T5 and T6. One disadvantage of the new labeling is that the letter "P" could suggest a parietal location, whereas P7 and P8 are actually located above the posterior temporal lobe [23].

2.1.3 EEG montage

The term EEG montage refers to a specific arrangement of electrodes on the scalp, as well as the way in which the electrodes are connected to each other. In a *monopolar montage*, each electrode is connected to a common reference (see Figure 2.2 on the right). The reference electrode can be placed on the scalp or on the body. A common example is the auricular electrode placed on the ear. Alternatively, the reference can be obtained by calculating the average of the potential recorded by all electrodes, called the *average reference* [24]. In a *bipolar*

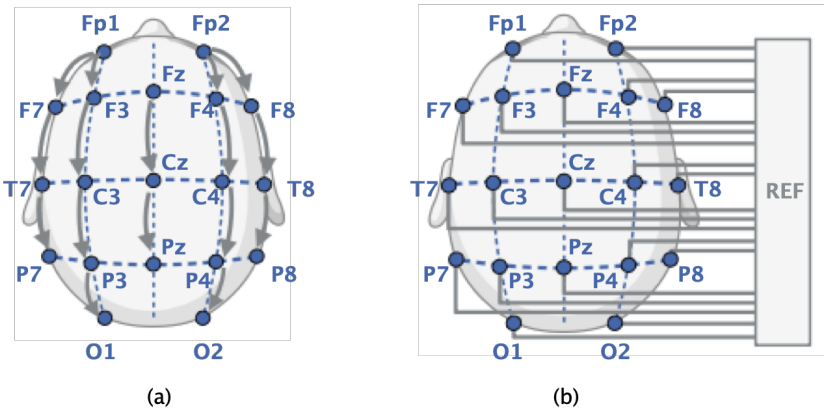


Figure 2.2: (a) Bipolar and (b) monopolar EEG montage. (Created with app *BioRender* [21])

montage, each electrode is connected to one or two neighboring electrodes [25] (see Figure 2.2 on the left). Sometimes it is possible to recalculate the reference of the data even after recording by means of simple arithmetic steps, which can be useful for highlighting new dynamics of interest. If necessary, it is possible to switch from a monopolar to a bipolar montage by eliminating the common reference. This can be done through simple arithmetic steps and can be useful for highlighting new dynamics of interest or for standardizing different montages between data sets.

When the EEG signal is recorded between a pair of electrodes, it is referred to as a *channel*. Thus, each channel records the voltage difference between two electrodes: one positioned near the brain region of interest and the other, known as the reference electrode, placed in an area with minimal neuronal activity, such as the mastoid [25]. The basic electrode array according to bipolar montage channel configurations is: F_{P1}-F_7, F_7-T_7, T_7-P_7, P_7-O_1, F_{P1}-F_3, F_3-C_3, C_3-P_3, P_3-O_1, F_{P2}-F_4, F_4-C_4, C_4-P_4, P_4-O_2, F_{P2}-F_8, F_8-T_8, T_8-P_8, P_8-O_2, Fz-Cz, Cz-Pz, T_7-F_{T9}, F_{T9}-F_{T10}, and F_{T10}-T_8. This configuration is used in this thesis.

2.1.4 Artifacts

EEG amplitudes are very small, between 10 μV to 100 μ V, so the recordings are very sensitive to artifacts. The most common types of artifacts are eye blinks, muscle artifact, electrode movement, and electrocardiogram (ECG) contamination.

Electrical and environmental interference is also a source of artifacts. External sources of electrical disturbance, such as nearby electronic devices or intense electromagnetic fields (i.e. 50 or 60 Hz mains noise), commonly affect EEG recordings.

Eye blinks are represented by a big voltage drops in the Fp channels. The frontal area is the one most susceptible to this type of artifact due to its close proximity to the eyes. Muscle artifact in an EEG recording consists of high-frequency, high-amplitude noise in the temporal T_7 electrode. The frequency content of muscle artifact can vary between 20 Hz to 300 Hz, which overlaps with EEG frequency range. Another type of artifact is electrode movement that usually causes large voltage swings that can saturate the channel.

ECG signals can also show up in EEG and are usually indicated by horizontal lines in F_7-T_7 channel. All the artifacts described in this section are common and appear with varying degree. Electrode movements and muscle artifacts occur even more frequently in ambulatory systems, as the patient is free to move around. These phenomena can also occur in the context of wearable devices. It is therefore essential to minimize their impact on the recording system in order to ensure that the actual EEG data can be analyzed correctly.

2.2 Epilepsy

This section provides an overview of the case study analyzed in this thesis, namely epilepsy. Introducing some concepts related to this condition is necessary to understand and correctly evaluate the results and discussions presented in the thesis. In particular, the section first defines what is meant by epilepsy, then presents the main classification tasks associated with it, and finally describes the main characteristics of the publicly available datasets considered in this thesis.

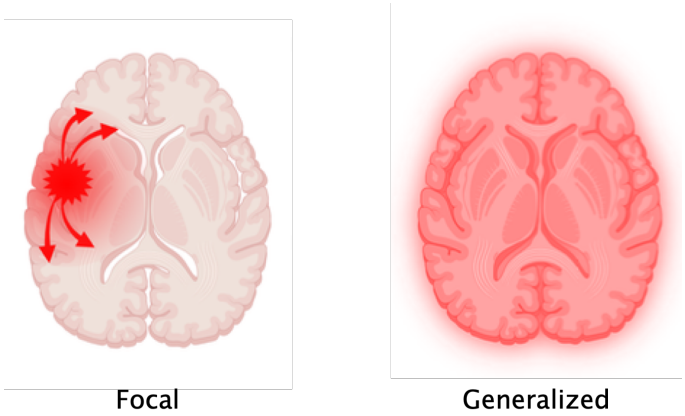


Figure 2.3: Distinction between focal and generalized seizures.

2.2.1 Introduction

Epilepsy is a chronic neurological disorder that affects about 50 million people world-wide, which is approximately 1% of the world population [26], and it is classified as the second most serious neurological condition known to man, after stroke.

Excessive and abnormal electrical discharge activity in a region of the brain causes a disturbance in normal brain activity, generating seizures known as *ictal events*. If the abnormal neuronal activity is localized in a specific part of the brain confined to one hemisphere, it is called a *focal* - or partial - *seizure*. By contrast, seizures that affect large portions of the brain involving both hemispheres are named *generalized*. Symptoms vary depending on the part of the brain affected. For example, a focal seizure on the auditory cortex will probably cause the patient to experience hearing sensations [27]. Figure 2.3 shows the main difference between focal and generalized seizures. The former begins in a specific region of the brain in one hemisphere and progressively spreads to other cortical areas. The latter presents widespread activity in both hemispheres.

In addition to the distinction between focal and generalized seizures, epileptic seizures can be further classified according to clinical, behavioural and electrographic manifestations [28, 29]. In particular, a distinction is made between simple partial seizures (SPSZ), i.e. complex partial seizures (CPSZ), which involve a loss of consciousness, generalized tonic-clonic

seizures (TCSZ), characterized by convulsions and finally absence seizures (ABSZ), generalized seizures with altered consciousness but without convulsions. However, most publicly available datasets do not include the labeling of seizure subtypes; one exception is the TUH Seizure Corpus, which provides this categorisation in one of its subsets [30].

2.2.2 Seizure-related phases

The human brain comprises approximately 86 billion cells called neurons which transfer information by propagating electrochemical currents along their main axis [26]. Neurons are capable of generating, propagating, and processing electric signals. They connect to other neurons to form functional networks: thus, the brain can be viewed as a collection of interacting neural networks. The inputs to a neural network can be excitatory if they promote activity among neurons or inhibitory if they suppress it. Epileptic seizures are transient periods involving hyperactivity and hyper-synchronizations of a large number of neurons within one or more neural networks [31]. They are characterized by four different temporal dynamics of brain activity: inter-ictal, pre-ictal, ictal and post-ictal.

Inter-ictal is the baseline EEG trace, i.e. the time interval when no epileptic seizures occur. It cannot be compared to the EEG trace of a healthy subject, as it may contain abnormalities related to the brain disorder caused by the disease, such as inter-ictal epileptiform discharges (IEDs). Such discharges may occur between convulsive events, but their study is not included in this work as their relationship with convulsive seizures remains weakly investigated [32].

Pre-ictal refers to the interval preceding the epileptic phase. During this phase, the patient may experience warning symptoms that precede the onset of the seizure. This interval is of considerable interest in research, as its correct identification would make it possible to predict epileptic seizures. Anticipating the onset of the subsequent ictal phase is a major goal in the clinical setting. However, the design of predictive systems is

severely limited by the fact that there is no clinical definition of the pre-ictal interval [33]. Indeed, different studies report that this interval not only varies across patients but also between seizures within the same patient [34].

Ictal refers to the period of seizure activity. It corresponds to the time interval annotated in the ground truth, which specifies the onset and offset of each seizure. Such labeling is performed by clinical experts through visual inspection of the EEG, leveraging their ability to recognize ictal patterns [35]. The specific EEG characteristics observed during this phase depend on the brain region affected by the epileptic activity.

Post-ictal refers to the interval right after the ictal interval. Its duration is variable, ranging from a few seconds to several minutes, during which the patient may appear confused or drowsy [36]. For this reason, this phase is often discarded from epilepsy classification tasks, such as prediction and detection [37]. Neural activity in this stage tends to gradually rebalance and return to baseline after the ictal event.

2.2.3 Seizure classification problems

Currently, EEG is considered the gold standard for detecting epileptic activity [38], even though it presents critical issues, such as delays in diagnosis or even diagnostic errors that can occur due to various factors intrinsic to the assessment process. For this reason, the introduction of artificial intelligence can provide more precise mappings of epileptic networks, thereby supporting diagnostic planning and potentially reducing the risk of misdiagnosis [39].

In the study of epileptic seizure, three different tasks are distinguished and categorized depending on their final objective, specifically: seizure detection, seizure prediction, seizure classification.

Seizure detection represents the process by which a model identifies the presence or absence of epileptic seizures or abnormal activity. The detection of epileptic seizures can be carried out when it is necessary to review EEG recordings and perform

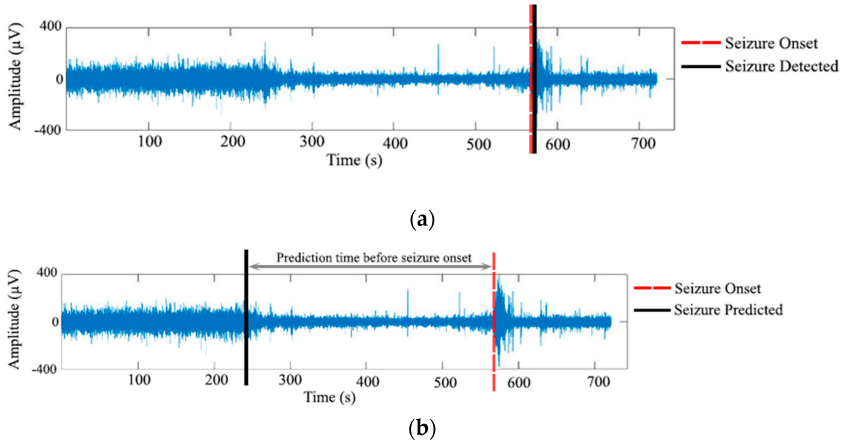


Figure 2.4: Different classification tasks for diagnosis of epilepsy: (a) seizure detection; (b) seizure prediction [41].

a retrospective analysis of critical episodes, but it is also useful as a clinical tool in cases of non-convulsive status epilepticus (NCSE) [40]. In these conditions, in fact, epileptic activity may not have observable external correlates, such as convulsions. Consequently, the detection task applied to EEGs is an essential aid in the diagnosis of NCSE [40]. In Figure 2.4a, it is possible to see when the ideal case of seizure detected verifies, as soon as the seizure onset happens.

Seizure prediction represents the process by which a model predict the likelihood of the occurrences of imminent epileptic seizures early on, by identifying the patient's pre-ictal state. Seizure prediction is usually deployed to take into account safety precautions: the objective of this task is, indeed, to spot upcoming seizures with some forewarning and alert the patient or the health care professionals. In Figure 2.4b, it is possible to see a seizure being identified before the seizure onset.

Seizure classification represents the process by which a model can categorize different types of seizures or seizure phases. Indeed, as proposed in 1981 by International League Against Epilepsy (ILAE), epileptic seizures are broadly classified according to their cerebral site of origin and spread. The first main division is between focal and generalized seizures. The

first arises from a localized region of the brain's cortex and have clinical manifestations that reflect that region of the brain. As an example, a focal originating in the part of the brain that processes emotions and short-term memory, i.e. the temporal lobe, may result in feelings such as euphoria, fear, and deja-vu or hallucinations of taste or smell. The second one affects the whole brain with abnormal electrical activity. The manifestations in this case often include the loss of consciousness. Additionally categorizations consider motor manifestations of the seizures, which may include whole-body rigidity and jerking (difference between tonic-clonic seizure) or whole-body loss of muscle tone (atonic seizure) [42]. Seizure classification may be a crucial aspect for neurologists to take the appropriate medical decision. However, the scarcity of publicly available data for the different categories of epilepsy limits the possibility of further investigating this task.

Figure 2.4 shows an overview of the first two classification tasks, namely detection and prediction [41], which represent the most established lines of research in the field of epilepsy. Although each of them is of considerable scientific and clinical interest, this thesis focuses specifically on the task of seizure detection, which still represents a significant challenge in the literature [43].

Traditional processing pipeline

To clarify how a classification problem aimed at detecting an epileptic seizure is translated operationally in the field of AI, the main phases of a traditional processing pipeline are presented schematically below. The problem analyzed concerns the distinction between ictal and non-ictal windows, with the latter including both inter-ictal and pre-ictal phases. Figure 2.5 summarizes this scheme.

The process starts with the acquisition of raw EEG data, which are then pre-processed to reduce noise and artifacts. Typical pre-processing steps include: (*i*) filtering, used either to remove power-line interference (e.g., 50 Hz) and improve signal quality or to decompose the signal into the standard frequency bands of interest; (*ii*) normalization, such as z-score normalization, to reduce variability across recordings and en-

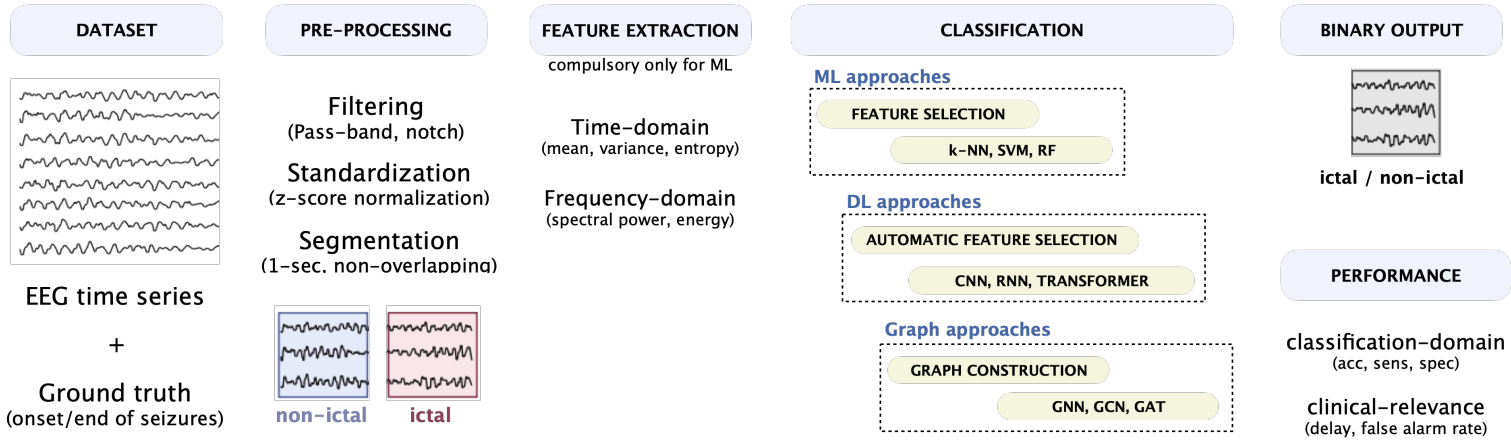


Figure 2.5: Flowchart of a common seizure detection method.

sure more consistent background activity levels; and (*iii*) segmentation through a sliding-window procedure, in which both window length and overlap must be defined.

After pre-processing, the pipeline proceeds to representation learning and classification. In traditional ML approaches, informative descriptors are manually extracted from each EEG segment in the time, frequency, or time-frequency domain, and may then be further refined through feature selection before being provided to a classifier such as Support Vector Machine (SVM) or Random Forest (RF). By contrast, DL approaches aim to automatically learn discriminative representations directly from the data, thereby reducing the reliance on manual feature engineering and enabling end-to-end optimization.

Among DL models, Convolutional Neural Networks (CNNs) have been widely adopted for EEG analysis due to their ability to effectively capture the nonlinear and non-stationary nature of EEG signals, particularly when the input is expressed in transformed forms such as spectrograms. To further model temporal dynamics, CNN-based architectures are often combined with recurrent modules such as Long Short-Term Memory (LSTM) networks, or alternatively replaced by Temporal Convolutional Networks (TCNs) [44]. While these architectures improve the modeling of sequential information, more recent seizure detection approaches increasingly integrate Transformer modules into CNN-based pipelines to better capture global spatial dependencies and overcome the memory limitations typically associated with recurrent networks. In such designs, EEG trials are first compressed into patch embeddings through convolutional blocks and subsequently processed by Transformer layers to model long-range dependencies across the signal [45]. However, these architectures generally rely on Euclidean representations of the signal and therefore do not explicitly account for the complex inter-channel relationships that characterize full-scalp EEG recordings.

In this context, graph-based approaches have emerged as a promising alternative, as they enable EEG channels to be modeled as interconnected nodes and their functional interactions to be encoded as edges. Such a formulation is better suited to capturing the complex spatial organization of multichannel

EEG data and provides the methodological basis for the adoption of Graph Neural Networks (GNNs), including Graph Convolutional Networks (GCNs) and Graph Attention Networks (GATs), with the aim of more effectively modeling the relational nature of EEG signals.

The pipeline concludes with a binary output (ictal vs. non-ictal) and an evaluation of model performance. In addition to standard metrics such as accuracy, sensitivity, and specificity, clinical applications often require more task-oriented indicators, including latency, false alarm rate, and the number of correctly detected events.

2.3 Graph-based EEG Analysis

Graph-based approaches for modeling EEG signals aim to leverage information from all available electrodes by considering their complex pairwise interactions. This section reviews the core steps in the pipeline of graph-based EEG analysis. First, the concept of brain connectivity is introduced in Section 2.3.1. Then, the methods for constructing brain graph structures from EEG data, i.e., functional connectivity measures, are presented in Section 2.3.2. Finally, Section 2.3.3 reviews the most commonly used graph neural network architectures for characterizing brain graphs.

2.3.1 Brain Connectivity

The term *brain connectivity* refers to the network of connections between different regions of the brain. Widely employed to gain a deeper understanding of complex neurological processes, connectivity analysis offers unique insights into the organization and dynamics of brain networks. The connectivity measure is calculated between all possible pairs of signals recorded on the EEG channels. This ensures that the degree of connectivity measured for a particular pair of signals is completely isolated from the rest of the EEG channels.

In general, connectivity in the brain can be described at three complementary levels: structural connectivity, which refers

to the anatomical links between regions; functional connectivity, which quantifies statistical dependencies such as correlations or synchronizations between neural signals; and effective connectivity, which characterizes directed causal influences and the dynamic flow of information between regions.

Structural connectivity (SC) represents the most intuitive notion of connectivity as it is usually understood as the pattern of fiber pathways that traverse broad brain regions and correspond to known anatomical organization [46]. By analogy with structural neuro-imaging, early EEG graph studies modeled SC by defining edges according to the physical distance between scalp sensors, for instance using Euclidean distances between electrode coordinates. While this is conceptually simple and well aligned with structural imaging, such as Magnetic Resonance Imaging (MRI) and Diffusion Tensor Imaging (DTI) [47], such a choice yields a static connectivity that does not reflect the dynamic interactions between EEG channels. In structural imaging, connectivity may evolve slowly due to brain plasticity, but in scalp EEG the inter-electrode distance is fixed and does not encode dynamic coordination.

Functional connectivity (FC) is defined as the synchronization between spatially remote neurophysiological events occurring in the regions of the cortex that reflects statistical dependencies between distinct and distant regions of information-processing neuronal populations. Thus, FC does not relate to the structure of the brain but it is purely about the presence or absence of statistical dependence between observed neural responses acquired across sensors and time. Functional graphs are often dense, with many edges representing dynamic interactions that fluctuate over time, depending on cognitive states, tasks, or external stimuli. This dynamic nature makes FC particularly well suited for EEG data analysis.

Effective connectivity (EC) is defined as the causal influence that one neural unit exerts on another and therefore requires directed edges and directed graphs [48]. However, at the sensor level, EEG is strongly affected by volume conduction

and montage referencing, which makes EC estimates difficult to interpret. These issues are further amplified in bipolar montages with shared electrodes, where apparent interactions can be inflated and spurious directionality can be introduced. For these reasons, in this study EC is not pursued and analyses focus on undirected graphs.

2.3.2 Graph construction

Graph theory is a mathematical and computational framework used to represent complex systems as sets of *nodes* (vertices) connected by *edges* (links) [49]. Although this approach has been successfully applied in several fields, including social networks, biological networks and knowledge graphs, in the context of EEG the construction of efficient and meaningful brain graphs that faithfully reflect the underlying neural activity still represents a central challenge in the literature [50].

The widely adopted method involves structuring the data in the form:

$$G = (V, E, A) \tag{2.1}$$

where $V = v_0, v_1, v_2, \ldots, v_n$ denotes the set of nodes, corresponding to the electrodes (or channels) placed on the scalp in EEG applications, and E represents the set of edges, i.e., the connections between nodes. These edges typically capture relationships such as correlation, coherence, or other measures of FC between the EEG signals recorded from two electrodes.

More specifically, E is the set of pairs of vertices. Thus, a singular edge e is represented as $e = \{x, y\}$, where x and y are nodes of the edge and are said to be *neighbors* or *adjacent nodes* in the graph G.

A refers to the weights assigned to each edge, indicating the strength of the connections, and is usually called the adjacency matrix. For example, a higher weight might indicate a stronger functional relationship between two brain regions [50].

Another representation of the graph is the Laplacian matrix, defined as:

$$\mathbf{L} = \mathbf{D} - \mathbf{A} \tag{2.2}$$

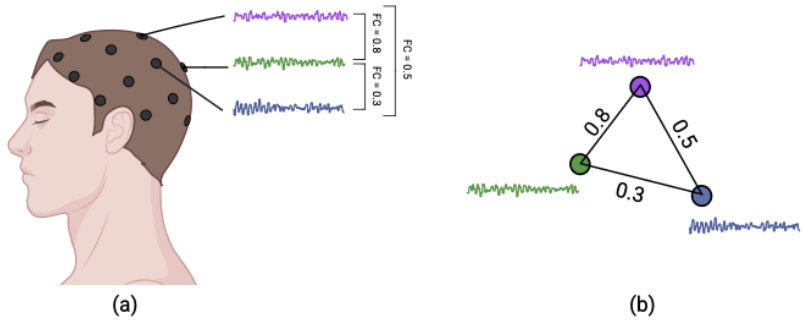


Figure 2.6: Example of graph computation from EEG recordings. (a) Calculation of FC between EEG channels. (b) Graph representation: raw EEG data represent feature nodes, FC values represent edges between nodes. (Created with app *BioRender* [21])

where $\mathbf{A}$ is the adjacency matrix while $\mathbf{D}$ is the *degree matrix*, that is, a diagonal matrix in which each entry corresponds to the number of connections of a given node. This matrix is especially relevant in spectral analysis within GNNs, as it facilitates the characterization of smoothness over the graph structure [51].

Figure 2.6 shows an example of graph computation based on EEG signals recorded on the scalp. For simplicity, only three channels in monopolar configuration are considered, indicated in purple, blue and green. Figure 2.6a shows the electrodes on the scalp and the corresponding recorded time series. For each pair of signals, the estimated FC value is reported. Figure 2.6b illustrates the corresponding mapping in an undirected graph: each electrode constitutes a node and the weighted edges represent the intensity of the connection, which is equal to the FC values. The vector of features for each node corresponds to the raw time series of the relevant EEG channel.

The connections between vertices v_i, represented by edges e_k, are encoded in the adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$. The element a_{ij}, with $i, j = 0, 1, \ldots, N-1$, indicates the presence of a connection from v_j to v_i. Thus, $a_{ij} = 1$ if the edge exists and $a_{ij} = 0$ otherwise. In the case of a weighted graph, a_{ij} takes the value of the edge weight (i.e., the measure of FC)

and is 0 in the absence of a connection.

FC can be studied both in terms of non-direct coupling between the time series of a pair of electrodes and in terms of causality, i.e. as a time-delayed cause-effect relationship between two brain signals [52]. Based on this distinction, the resulting graphs can be direct or non-direct. In *directed graphs*, the connectivity matrix is not symmetrical and may have bidirectional links, allowing for the detection of oriented and reciprocal influences; in undirected graphs, on the other hand, symmetrical connectivity measures are used to quantify the simultaneous occurrence of spatially distant neurophysiological events [52].

Graph nodes can be represented using:

- *raw data*: the raw sampled signal values from each EEG channel are directly used as node attributes. This representation is often adopted to evaluate the potential of GNNs in enabling an end-to-end learning framework without relying on intermediate feature extraction.

- *Power Spectral Density (PSD)*: given the relevance of frequency-domain information in brain state classification, the average power within specific frequency bands is computed. The common frequency bands were considered: Gamma (35-80 Hz), Beta (13-35 Hz), Alpha (8-13 Hz), Theta (4-8 Hz), and Delta (0.5-4 Hz). To estimate the PSD, both the Welch method and the Burg method were employed.

- *Band statistics*: for each of the aforementioned frequency bands, six statistical descriptors were extracted: mean, standard deviation, entropy, variance, skewness, and kurtosis. These metrics provide a compact yet informative characterization of the signal's distribution within each frequency range.

- *1D Convolutional Neural Networks (1D-CNNs)*: can be applied to the raw EEG signals from each electrode to further enhance feature representation. These networks automatically learn temporal patterns and hierarchical

features by applying convolutional filters along the time dimension, enabling the model to extract local and task-relevant features directly from the signal prior to graph modeling [53].

Graph edges, which capture interactions or correlations between nodes, are typically defined using FC measures. This type of connectivity is fundamentally different from SC, which is more commonly used in brain imaging analyses. Functional graphs are often dense, with many edges representing dynamic interactions that fluctuate over time, depending on cognitive states, tasks, or external stimuli. This dynamic nature makes FC particularly well suited to EEG data analysis. In this context, edges can be computed using:

- *Pearson correlation coefficient (PCC)*: which quantifies the linear relationship between two signals with values ranging from -1 to 1. Let $x_1(t)$ and $x_2(t)$, where $t = 1, \ldots, T$ indexes the samples, represent two EEG time series from electrodes (nodes) 1 and 2, respectively. PCC provides a more accurate representation of the degree of interaction between nodes with:

$$PCC = \frac{\sum_{t=1}^{T}(x_1(t) - \bar{x}_1)(x_2(t) - \bar{x}_2)}{\sqrt{\sum_{t=1}^{T}(x_1(t) - \bar{x}_1)^2}\sqrt{\sum_{t=1}^{T}(x_2(t) - \bar{x}_2)^2}} \tag{2.3}$$

  where $\bar{x}$ represents mean values of observed series. The sign and magnitude of PCC measure the direction and degree of linear association between two signals with zero time lag, respectively. In extreme cases, $r = -1$ reveals a complete linear inverse relationship whereas $r = 1$ indicates a perfect linear relationship. A complete absence of linear dependency manifests in $r = 0$ [54].

- *Phase-locking value* (PLV): which assesses the strength of FC in the network. It quantifies the absolute value of the average phase difference between two time-series

$x_1(t)$ and $x_2(t)$:

$$PLV(t,f) = \left| \frac{1}{T} \sum_{t=1}^{T} e^{j(\phi_{x_1}(t,f) - \phi_{x_2}(t,f))} \right| \tag{2.4}$$

where $\phi_{x_1}(t,f)$ is the instantaneous phase of the time series $x_1(t)$ obtained via Hilbert transform (and similarly for $\phi_{x_2}(t,f)$) and (x_1, x_2) are two different nodes. The value of PLV is between 0 and 1 and is used to evaluate the phase synchronization between brain regions. Values close to 1 indicate strong synchronization, values close to 0 indicate a lack of synchronization [54].

- *Phase Lag Index* (PLI) which measures phase synchronization between two signals $x_1(t)$ and $x_2(t)$, defined as the relative phase between them. With values closer to 1 indicating stronger synchronization, PLI is given by:

$$PLI(t,f) = \left| \frac{1}{T} \sum_{t=1}^{T} \text{sign}\big[\Delta\phi(t,f)\big] \right| \tag{2.5}$$

where $\Delta\phi(t,f)$ is the instantaneous phase difference between the instantaneous phases of each pair of nodes (x_1, x_2):

$$\Delta\phi(t,f) = \phi_{x_1}(t,f) - \phi_{x_2}(t,f) \tag{2.6}$$

This measure is analogous to PLV. It describes the relative distribution of phase lags and was introduced to overcome a limitation of PLV, i.e. sensitivity of PLV to volume conduction. By neglecting phase lags close to 0 mod π, PLI is more robust to common sources and zero-delay couplings, reducing spurious connections compared to PLV [54].

- *Coherence measure* which quantifies the degree of linear correlation between two signals across different frequency bands. The coherence between two EEG channels, x_1 and x_2, is defined as:

$$C(f) = \frac{S_{x_1x_2}(f)}{\sqrt{S_{x_1}(f)\, S_{x_2}(f)}} \tag{2.7}$$

where $S_{x_1x_2}(f)$ is the cross-spectral density between x_1 and x_2, and $S_{x_1}(f)$ and $S_{x_2}(f)$ are the auto-spectral densities (power spectral densities) of x_1 and x_2, respectively [54]. Coherence values range between 0 and 1, with higher values indicating stronger linear coupling at frequency f [54].

2.3.3 Graph Neural Networks

Graph Neural Networks (GNNs) are a class of deep learning architectures specifically designed to process graph-structured data. Unlike traditional neural networks, which operate on grid-like structures (such as images), GNNs are able to learn patterns from data characterized by arbitrary graph topologies [55]. For this reason, they represent a promising approach for capturing the spatial complexity of EEG signals, as they extend the notion of convolution from Euclidean domains to data organized in graph form.

A GNN consists of multiple layers in which each node aggregates the features of its neighbors and updates its own representation, exploiting information about the edges [51]. This scheme, commonly referred to as *message passing*, typically involves three key operations: message computation, aggregation, and node update. In practice, GNNs operate through iterative steps that refine the representation of each node by incorporating both the features of its neighboring nodes and the information carried by the connecting edges.

Among the most popular architectures is the Graph Convolutional Network (GCN) [56], which generalizes convolution to the graph domain. To this end, each node updates its own features by aggregating those of its immediate neighborhood, capturing both structural and attribute-based dependencies. GCNs can be broadly categorized into two main types: spatial and spectral GCNs.

Spatial GCNs, facilitate efficient processing and have gained popularity due to their simplicity and scalability. They oper-

ate directly on the graph structure using the adjacency matrix **A**. Starting from the characteristics of the nodes, the networks calculate new representations using the neighbor aggregation mechanism of message passing. The basic idea is that nodes connected by an edge should share similar representations. For this reason, at each update step, the embedding of a node is enriched with information from its surroundings. The update of the embedding of node i at level $l+1$ in a spatial GNN with L layers can be described as [56]:

$$h_i^{(l+1)} = \sigma\left(W_1^{(l)} h_i^{(l)} + \sum_{j \in \mathcal{N}(v_i)} W_2^{(l)} h_j^{(l)} e_{ji}\right) \tag{2.8}$$

where $h_i^{(l)}$ is the embedding of node i at layer l, σ is a nonlinear activation function, and $\mathcal{N}(v_i)$ denotes the neighborhood of node v_i. The learnable weight matrices are $W_1^{(l)}, W_2^{(l)} \in \mathbb{R}^{d_1 \times d_2}$, projecting node embeddings from input dimension d_1 to hidden dimension d_2. Specifically, $W_1^{(l)}$ transforms the features of the center node i (self-message), whereas $W_2^{(l)}$ transforms the features of the neighbor nodes j (neighbor-messages) prior to the weighted aggregation via e_{ji}. Finally, e_{ji} denotes the edge weight.

Each layer of a spatial GCN only takes into account the 1-hop neighborhood. By superimposing multiple layers, the model is able to capture information from increasingly distant neighbors (L-hop). However, when the network becomes very deep, a problem known as over-smoothing emerges. In this situation, node representations progressively reduce specificity, tending to become uniform until they are almost indistinguishable, since the propagated messages already contain aggregated information from the entire graph.

A special case of spatial GNN is the graph attention network (GAT). GAT allows adjusting the graph by re-weighting the edges using an attention mechanism. Generally, the attention mechanism for computing the new softmax-normalized edge weight e_{ji} is defined as follows [56]:

$$e_{i,j} = \frac{\exp\left(\sigma\left(\mathbf{w}^\top [W\mathbf{h}_i \,\|\, W\mathbf{h}_j]\right)\right)}{\sum_{k \in \mathcal{N}(i)} \exp(\sigma(\mathbf{w}^\top [W\mathbf{h}_i \,\|\, W\mathbf{h}_k]))}, \tag{2.9}$$

where w and W are the learnable parameters of the model, σ is an activation function, h is the node feature and $N(i)$ is the set of nodes connected to node i.

Spectral GCNs, in contrast, are heavily influenced by spectral graph theory and rely on the eigen decomposition of the graph Laplacian to perform convolutions in the spectral domain [57]. In other words, given the graph structured data x with adjacency matrix A and degree matrix D, the normalized Laplacian is:

$$L = I - D^{-\frac{1}{2}} A D^{-\frac{1}{2}} = U \Lambda U^{\top}. \tag{2.10}$$

where $U = [u_1, \ldots, u_n]$ contains the orthonormal eigenvectors and $\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_n)$ the eigenvalues. Then, the graph Fourier transform (GFT) and its inverse are defined as:

$$\mathcal{F}(x) \;=\; U^{\top} x, \qquad \mathcal{F}^{-1}(z) \;=\; U z.$$

Convolution between a graph signal x and a spectral filter g is a point-wise multiplication in the spectral domain followed by the inverse GFT:

$$g * x \;=\; \mathcal{F}^{-1}\big(\mathcal{F}(g) \odot \mathcal{F}(x)\big) \;=\; U\big((U^{\top} g) \odot (U^{\top} x)\big),$$

which is equivalently written by parameterizing the filter as a function of the spectrum:

$$g * x \;=\; U\, g(\Lambda)\, U^{\top} x, \qquad g(\Lambda) \;=\; \mathrm{diag}\big(g(\lambda_1), \ldots, g(\lambda_n)\big).$$

While theoretically powerful, this approach can be computationally intensive and difficult to scale to large graphs. To balance expressiveness and efficiency, recent developments in GCNs have introduced Chebyshev polynomial approximations as a solution to the computational limitations of traditional spectral methods. Instead of relying on explicit eigen decomposition, this approach approximates the graph convolution operator with a truncated Chebyshev expansion [57].

Let L (Eq. 2.10) be the normalized Laplacian and let $\lambda_{\max}$ denote its largest eigenvalue. Defined the scaled Laplacian as:

$$\tilde{L} \;=\; \frac{2}{\lambda_{\max}} L \;-\; I,$$

whose spectrum lies in $[-1, 1]$, the spectral filter $g(\Lambda)$ is then approximated as:

$$g_\theta(\Lambda) \;\approx\; \sum_{k=0}^{K} \theta_k \, T_k(\tilde{\Lambda}). \tag{2.11}$$

where T_k are the Chebyshev polynomials of the first kind and $\theta = \{\theta_k\}_{k=0}^{K}$ are learnable coefficients. Using $U\Lambda U^\top$ for the eigen decomposition, the convolution can be evaluated *without* eigen decomposition as:

$$g_\theta * x \;=\; U \, g_\theta(\Lambda) \, U^\top x \;\approx\; \sum_{k=0}^{K} \theta_k \, T_k(\tilde{L}) \, x \tag{2.12}$$

Thus, the Chebyshev polynomials satisfy the stable recurrence: $T_0(\tilde{L}) = I, T_1(\tilde{L}) = \tilde{L}, T_k(\tilde{L}) = 2\,\tilde{L}\,T_{k-1}(\tilde{L}) - T_{k-2}(\tilde{L})$ for $k \geq 2$. Each Chebyshev term $T_k(\tilde{L})x$ of 2.12 can be computed iteratively with sparse matrix-vector multiplications. This significantly reduces computational efforts yielding an overall complexity of $\mathcal{O}(K|E|)$ per layer (and $\mathcal{O}(K|V|)$ memory). The filter is K-localized, meaning that it depends on K-hop neighborhoods, scales to large graphs, and still captures both local and global structural patterns in the EEG connectivity graph-making it particularly well suited to our seizure detection task.

Dynamic Graph Neural Networks

Recently, significant progress has been made in the field of GNNs with the emergence of Dynamic Graph Neural Networks (DGNNs). These models are specifically designed to address the limitations of conventional GNNs when applied to time-varying signals such as EEG. Indeed, traditional GNNs typically rely on fixed graph topologies derived from prior data, which prevents the model from adapting to temporal changes or evolving patterns within the signal. This static structure can be particularly restrictive in the context of EEG, where brain connectivity may vary dynamically over time [58].

Among the various DGNN approaches, the most suitable for EEG signals, characterized by fixed nodes (electrodes) and dynamic edges (functional connectivity), is the Spatio-Temporal Graph Neural Network (STGNN). The core idea of STGNNs is to simultaneously model spatial relationships and temporal dependencies through a tight integration of spatial and temporal modules. The spatial module typically employs GCNs to learn the structural dependencies between nodes at each time step. Meanwhile, the temporal module captures the evolution of node features over time using sequential models such as Recurrent Neural Networks (RNNs) or Long Short-Term Memory (LSTM) networks.

Chapter 3

Literature review

The chapter provides a comprehensive literature review of the current state of the art regarding AI methods for EEG seizure detection. In particular, the current literature is reviewed from two perspectives, namely lightweight AI-based approaches and GNN-based approaches. The former focus on reducing computational complexity and model size while maintaining high performance, thereby enabling deployment on low-power wearable devices. The latter represent an innovative strategy for appropriately modeling a non-Euclidean signal, such as EEG recorded over the entire electrode configuration, effectively capturing and managing the complexity of the spatial relationships between channels. The chapter is structured as follows. Section 3.1 provides an introduction of the evolution of EEG research with the introduction of AI-based approaches. Section 3.2 reviews the main datasets employed for the task under analysis. Section 3.3 presents AI-lightweight approaches for EEG-based seizure detection and describes the problem of EEG channel reduction, while Section 3.4 describes methods based on graph neural networks. At the end of the chapter, Section 3.5 summarizes the research gap identified and the research direction proposed in this work.

3.1 Introduction

The evolution of EEG research has undergone a significant change with the introduction of AI-based approaches, partic-

ularly in the field of seizure detection. The complexity of the task, combined with the highly nonlinear nature of EEG signals, makes ML and DL valuable tools for extracting relevant information and identifying meaningful patterns, thereby improving classification accuracy [59, 60].

For several decades, traditional ML represented the dominant paradigm for automatic seizure detection, mainly because publicly available EEG datasets were generally characterized by limited sample size and high input dimensionality [59]. Under these conditions, methods based on handcrafted feature extraction became a natural choice and achieved promising results in different EEG-related tasks [61]. Among the most widely adopted classifiers, Support Vector Machines (SVMs) and Random Forests (RFs) have been extensively documented in the literature on seizure detection [59, 62]. SVM draws complex decision boundaries via representative samples of the positive and negative classes by comparing each test sample to specific algorithmically selected data points called support vectors. On the other hand, RF constructs an ensemble of decision trees, such that each tree is trained with a random subset of the input examples, and each node is optimized using a random subset of the input features. These models robust performance for small amounts of data. For instance, on the CHB-MIT data Raghu *et al.* achieved a 94.21% with SVM [63], while Xiong *et al.* reported an accuracy of 94.04% using a RF classifier [64].

Despite their effectiveness, low computational requirements, and fast training, these approaches present some important limitations. In particular, they rely on a preliminary feature-engineering stage, which requires both domain expertise and careful design choices and may still fail to capture the full complexity of EEG dynamics. In this regard, the literature offers a wide range of features for EEG analysis, the most common of which are summarized in Table 3.1. These features capture different aspects of the EEG signal, including the time domain, frequency domain, spatial domain (i.e., across different channels), or a combination of these [4]. Their aim is to provide an efficient representation of brain activity and to highlight relevant information for seizure detection. However, the performance of traditional ML methods strongly depends on the

Table 3.1: Main feature extraction domains adopted in EEG analysis for seizure detection, together with representative features commonly used in the literature.

Extraction method	**Relevant features**
Time domain	Mean, variance, median, mode, skewness, kurtosis, min, max, standard deviation, zero-crossing rate, line length, energy, power, Shannon entropy, sample entropy, approximate entropy, fuzzy entropy, Hurst exponent
Frequency domain	Spectral power, spectral entropy, energy, peak frequency, median frequency
Time-frequency domain	Line length, min, max, Shannon entropy, approximate entropy, standard deviation, energy, median, root mean square
Discrete Wavelet Transform (DWT)	Bounded variation, coefficients, energy, entropy, relative power, relative scale energy, variance, standard deviation
Continuous Wavelet Transform (CWT)	Standard deviation of energy, energy, coefficient z-score, entropy
Fourier Transform (FT)	Median frequency, power, peak frequency, spectral entropy, spectral edge frequency, total spectral power

quality and discriminative power of the selected features, which may limit both generalization and adaptability to real-world scenarios [65]. These limitations have progressively motivated the shift toward neural network architectures, which enable automatic representation learning directly from raw or minimally processed EEG signals. By reducing the dependence on handcrafted features, such models provide a more flexible framework for capturing the complex temporal and spatial patterns of EEG data [66].

Although many approaches have been proposed in the literature for processing time-series data in EEG-based seizure detection, as illustrated in Figure 3.1 which summarizes the prevalence of neural network architectures employed in epilepsy research, the present review focuses primarily on the model

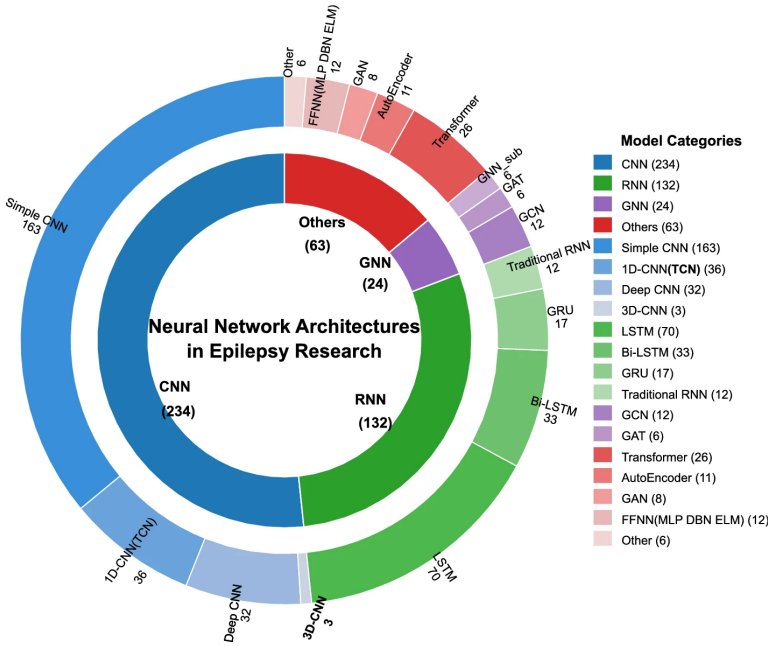


Figure 3.1: Nested donut chart of neural networks applied in epilepsy detection and prediction [67].

categories that are most widely adopted in the field [67]. This choice is motivated not only by their methodological relevance, but also by their central role in the evolution of AI methods for seizure detection.

Convolutional Neural Networks (CNNs) are the most widely used DL technique [68, 60]. Their effectiveness mainly derives from their ability to capture local patterns in the signal through convolutional filters, making them particularly suitable for identifying discriminative temporal structures in windowed EEG recordings. Since CNNs were originally designed for image-like or regularly gridded two-dimensional data, many early applications in EEG analysis relied on transforming the signals into two-dimensional time-frequency representations, such as the Short-Time Fourier Transform (STFT) or the Discrete Wavelet Transform (DWT) [61]. These representations can be more naturally processed by 2D-CNNs, allowing the models to capture spatially distributed patterns in the transformed feature maps. However, later studies have shown that

robust performance can also be achieved by directly analyzing raw EEG signals provided that the convolutional kernels are properly designed to match the signal structure [69]. This has progressively encouraged the adoption of CNN architectures operating on raw EEG, together with more advanced collaborative designs, such as multi-branch, multi-feature, and multi-scale frameworks, which have further improved recognition performance [67]

Another widely used family is that of Recurrent Neural Networks (RNNs), especially Long Short-Term Memory (LSTM) models, which are specifically designed to process sequential data and preserve information over extended temporal contexts. For example, Pandey *et al.* developed an automatic seizure detection system based on LSTM, achieving an accuracy of 89.88% on the CHB-MIT dataset and 92.59% on the Siena Scalp dataset [70]. However, these results appear to be lower than those achieved by hybrid architectures that combine CNN-based models for feature extraction with temporal modules. Such combinations are particularly effective because they jointly exploit the ability of CNNs to extract local discriminative patterns and the capability of recurrent architectures to model the long-term temporal dependencies needed to distinguish seizure (ictal) from non-seizure (inter-ictal) events over time. In this regard, Li *et al.* proposed a seizure detector based on a hybrid CNN-LSTM network applied to multichannel time series, achieving an average sensitivity of 95.29% [71].

While these extensions improve sequential modeling, they may still be less effective in capturing global spatial and temporal dependencies in a flexible and computationally efficient manner, especially over long EEG sequences. For this reason, recent seizure detection architectures have increasingly incorporated Transformer modules into CNN-based pipelines [67]. In most cases, these models adopt a serial design, in which EEG trials are first compressed into patch embeddings through convolutional blocks and subsequently processed by Transformer layers to model long-range dependencies [66]. Such frameworks still rely on a sequential processing strategy that constrains information flow to a single pathway, with local and global features being handled in successive stages rather than jointly [67].

As a result, the spatial inter-channel relationships extracted by the initial convolutional layers may be progressively attenuated in later blocks, limiting the model's ability to fully exploit the spatial complexity of EEG signals.
This aspect indicates only one of the limitations of Transformer-based models in spatio-temporal signal analysis. In the context of EEG, their use is currently associated with high computational cost, which moves away from the objectives of the present study, focused instead on lightweight and resource-efficient solutions. Moreover, these models may suffer from inefficiencies and overfitting, particularly in scenarios where large annotated datasets are not available, since they typically require substantial amounts of training data and a considerable number of epochs to achieve stable convergence [66].

It is worth stressing, however, that although Transformer-based architectures are not included among the methodological choices of this research, their potential should not be underestimated. Indeed, their integration into the emerging concept of *brain foundation models* may open new opportunities for exploiting large-scale pre-trained architectures. In this scenario, pre-training on large and heterogeneous EEG datasets could provide a robust backbone capable of supporting a variety of downstream tasks, including epileptic seizure detection. Such an approach could facilitate efficient transfer learning across different applications, reducing the overall computational burden and requiring only limited fine-tuning for adaptation to specific tasks [72].

What has been described provides an overview of the work carried out in the field of DL-based EEG analysis and highlights the significant progress achieved in this area, although several challenges remain unresolved.

Spatial complexity.
EEGs are characterized by both temporal and spatial information. Therefore, it is necessary to carefully evaluate how to format the data before it is used by an AI model. The works proposed in the literature mainly adopt two strategies. A first line of research consists of designing architectures capable of processing temporal and spatial information separately,

treating each channel individually and analyzing them one at a time [73]. On the other hand, an alternative approach involves a remodeling procedure, in which multichannel EEG signals are organized into an image-like format, allowing neural networks to simultaneously process temporal and spatial dimensions through two-dimensional convolution operations [74]. Although more common, the latter representation often reduces the actual inter-dependencies between channels, leading them away from the underlying neurophysiological complexity. Indeed, the spatial correlations resulting from the propagation of cortical potentials are not linearly separable and are easily distorted by phenomena such as volume conduction or reference choice [75]. In recent years, a graph-based representation of the EEG signal has emerged, which is a promising line of research for managing such complexities. In this approach, the topology of the signal is preserved by capturing nonlinear relationships between brain regions [56].

Inter- and intra-patient variability.

Both inter- and intra-patient variability constitute a significant and widely discussed challenge in EEG analysis [61]. Like other biosignals, EEG is significantly affected by individual anatomical, physiological, and clinical differences, which limit the definition of unique and easily generalizable predictive models. Furthermore, even in the same subject, the signal may exhibit variations over time due to patho-physiological changes, which accentuate its non-stationary nature and lack of recurring patterns [76].

With regard to this problem, the work proposed in the literature mainly adopts two AI application strategies. Intra-subject models are trained and tested on EEG data acquired from a single patient. In these cases, performance tends to be higher, as the model has to deal with reduced variability and evaluate a limited number of recordings [61]. Inter-subject models include data from multiple patients, requiring larger and potentially more representative datasets for each patient. These models report lower performance because inter-individual variability is added to the physiological variability of the individual patient, making generalization considerably

more complex [66].

What significantly affects the results obtained are the validation techniques for these two types of models. Among the most common strategies are leave-N-subject-out, in which data from different subjects are used for training and testing, and leave-N-record-out, in which different EEG recordings from a single subject are used for training and testing, respectively.

Although several reviews show that models generally perform better in intra-subject tasks than in inter-subject tasks [4, 73, 68, 66], intra-subject variability is also highly regarded by the scientific community in the epilepsy task. In fact, the morphology, amplitude, and frequency of seizures change from one epileptic episode to another in EEG recordings acquired from the same individual [77]. In addition, physiological temporal dynamics, such as the sleep-wake cycle, alter the background in which the epileptic seizure occurs, changing the pattern between events in the same patient [78].

In conclusion, the main challenges discussed represent significant obstacles to the clinical adoption of EEG-based AI systems. In this thesis, they are addressed with particular attention through two emerging lines of research that are attracting growing interest. On the one hand, the development of lightweight AI models optimized for implementation on low-power wearable devices. On the other hand, the use of graph-based architectures capable of preserving the non-Euclidean nature of the relationships between EEG signals and more accurately capturing the complexity of brain dynamics. The following sections present the state of the art in these two lines of research, in order to understand how the scientific community is addressing these challenges, which solutions are emerging as the most promising, and what gaps still remain.

3.2 Dataset

From the literature, three publicly available datasets related to epilepsy were identified, while the remaining eight were private. These datasets vary in characteristics such as the patient information, sampling frequency, total number of ictal events,

Table 3.2: Overview of the two datasets used in this study for seizure detection.

Characteristic	**CHB-MIT** [79]	**Siena Scalp** [80]
Number of patients	23	14
Age of subjects (years)	1.5–22	20–71
Number of seizures	181	47
Type of recordings	bipolar	monopolar
Total recording time (h)	950	128
Sampling frequency	256 Hz	512 Hz

and whether a monopolar or bipolar montage was used. Table 3.2 details them.

The **Children's Hospital Boston-Massachusetts Institute of Technology** (CHB-MIT) dataset is the most widely used in the literature among scientific studies [79]. It includes 5 males aged 3-22 years, 18 females aged 1.5-19 years and 1 individual with unknown age and gender. The recordings were originally sampled at 256 Hz and are provided in a bipolar montage (see Section 2.1). The database includes both focal and generalized seizures, although the specific classification of each seizure type is not provided. The seizure intervals are precisely annotated, indicating the exact start and end time of each epileptic event.

The **Siena Scalp EEG** dataset contains EEG recordings from 14 patients, collected at the Unit of Neurology and Neurophysiology at the University of Siena [81, 80]. The database include 9 males (ages 25-71) and 5 females (ages 20-58). The recordings were acquired at a sampling rate of 256 Hz. However, since some patients experienced only one or two seizures, we included only those with more than two seizures in our analysis (i.e., eleven patients).

Both datasets are based on the international 10-20 electrode configuration [20]. To ensure cross-dataset consistency,

we ensured the analysis to the 21 EEG channels common to all recordings, removing any additional channels. Moreover, among the selected datasets, CHB-MIT is provided in a bipolar montage, whereas Siena is distributed in referential (monopolar) configurations. Thus, since bipolar signals cannot be converted back to referential without the original reference (whereas monopolar recordings can be re-referenced to bipolar by differencing appropriate pairs), we derived a bipolar montage for Siena. This enabled us to validate the results obtained on the primary dataset (CHB-MIT) under a uniform montage across all three datasets. Furthermore, re-referencing the data to a bipolar montage minimizes the volume conduction issue since it eliminates the common reference. The reduced volume conduction will be particularly useful in the second study presented in this work.

Several seizure detection studies also use EEG data from the University of Bonn EEG Database [82] and Temple University Hospital (TUSZ) [83]. However, the first, while including surface EEGs from healthy subjects, its seizure recordings are intracranial rather than scalp EEG and since this work focuses on non-invasive scalp EEG recordings, we do not include Bonn data in our analysis. The second represents a subset of a Temple University Hospital (TUH) database that is attracting growing interest and is among the largest public EEG databases in epilepsy research. According to the official website (accessed August 30, 2025), it comprises 26,846 clinical EEGs. The corpus is structured into multiple annotated subsets, including TUSZ, each focusing on different phenomena of interest to research (i.e., epileptic seizures, artifacts, eye movements). The other subsets do not include aspects covered in this thesis.

Despite its size, the adoption of TUSZ is not present in this thesis, due to some practical issues: the annotations are relatively sparse compared to the available signal volume; moreover, many recordings are short and the amount of data per patient is reduced. These factors complicate the design of studies and the development of patient-specific models, which is the central objective of this work. However, although not used in this thesis, the relevance of TUH is well known in the liter-

Table 3.3: Accuracy results and computational complexity of ML algorithms [84].

ML algorithm	**Accuracy (%)**	**Computational complexity**
SVM	97.31	$O(n^3)$
RF	97.08	$O(n \log n)$
NN	93.53	$O(mn)$
k-NN	90.01	space $O(n)$

ature: TUH is increasingly used as a pre-training dataset for foundation models, as discussed in Section 7.3.

3.3 Lightweight approaches

As described in the previous section, numerous ML and DL algorithms have been proposed for application to EEG physiological signals. However, these methods require a large set of learnable parameters and large hardware resources to be implemented [74]. This means that, in order to be able to make a step forward implementing ML or DL algorithms on compact, low-power wearable devices, we have to make some considerations about the computational complexity or the potential applicability of the proposed architectures in-devices. In this regard, this section presents the main challenges and related solutions proposed in the literature regarding the application of AI algorithms on EEG time series in the context of wearable devices.

With regard to ML algorithms, Resque *et al.* [84] studied the performance of those most commonly used in the context of epilepsy detection. Focusing on their computational cost, the authors applied all models to the same EEG dataset and compared their performance, as summarized in Table 3.3.

In particular, the authors found that both SVM and RF achieved high accuracy in the classification task, with comparable results (97.31% for SVM and 97.08% for RF), but RF showed lower computational complexity than SVM. For this reason, they suggest that RF may be a better choice for appli-

cations in wearable contexts, where computational efficiency is a critical factor [84]. ML algorithms were also compared in terms of the trade-off between performance and computational cost in [85]. From the evaluation of SVM, RF, Extra Trees, and AdaBoost, Ingolfsson *et al.* achieved the same conclusions as Resque *et al.* [84]. In addition, they also implemented the RF model on an integrated 8-core platform called *bioWolf*, demonstrating that RF is promising for future wearable epilepsy monitoring systems with minimal energy requirements (approximately 1μJ) for inference.

Consistent results were also reported by Sopic *et al.* [86] and, subsequently, by Zanetti *et al.* [87], who adopted RF as the internal processing algorithm in their wearable device, further emphasizing its suitability for low-complexity real-time applications in EEG analysis (see Section 3.3.2 for further details). In both cases, the results are expressed in terms of the geometric mean, defined as the square root of the product between sensitivity and specificity, reporting 96.36% on a limited subset of 10 patients [86] and 94.5% across all 23 subjects [87] of the CHB-MIT dataset, respectively.

In contrast to these results, with regard to DL algorithms, the study by Daoud *et al.* [88] focused on evaluating the computational complexity of various models, taking into account both the feature extraction and classification phases. This represents a key difference from the work of Resque *et al.* [84], which did not account for the computational cost associated with the feature extraction process performed prior to the application of classification algorithms. Indeed, in the work of Daoud *et al.* [88], the authors evaluated four distinct DL models applied to raw EEG data from the publicly available CHB-MIT dataset [79]. The models considered were: (1) a Multi-Layer Perceptron (MLP) handling both FE and classification, (2) a CNN for FE combined with an MLP for classification, (3) a CNN for FE coupled with a Bidirectional LSTM (Bi-LSTM) network for classification, and (4) Autoencoders (AEs) for FE integrated with a Bi-LSTM for classification. The outcomes of this evaluation, reported in Table 3.4, summarize both the performance and the computational complexity of each model, with complexity measured in terms of the total number of train-

Table 3.4: Accuracy results and computational complexity (in terms of number of learnable parameters) of DL algorithms [88, 89] on the CHB-MIT [79] and Bonn [82] datasets, respectively.

Authors	DL algorithm	Accuracy (%)	Number of parameters
[88]	MLP	83.63	8,870,291
	CNN + MLP	94.10	520,477
	CNN + Bi-LSTM	99.66	27,657
	AE + Bi-LSTM	99.66	27,658
[89]	CNN	96.23	21,366
	P-1D-CNN	96.27	8,326

able parameters.

As shown in Table 3.4, the combination of CNN and Bi-LSTM networks proved highly effective in jointly capturing relevant spatial and temporal features from raw EEG data, achieving an accuracy of 99.66% in the epilepsy classification task with a total of 27,657 trainable parameters. Comparable performance and parameter count were also obtained when replacing the CNN with a pre-trained AE model for feature extraction. Building on these promising results, the authors suggested further investigation into reducing the number of learnable parameters, for instance through the adoption of a more accurate EEG channel selection procedure [88].

What was published by Daoud et al. [88] appears to build upon the work of Ullah *et al.* [89] and Lawhern *et al.* [69], who, the year before, had already addressed the trade-off between achieving strong performance and maintaining reduced complexity in DL models.

Specifically, the study by Ullah *et al.* compared the performance of two different networks [89]. The first is a standard 1D-CNN, while the second, termed the pyramidal 1D-CNN (P-1D-CNN), is designed with a significantly reduced number of trainable parameters, 61% fewer parameters. Both networks include convolutional layers, batch normalization, fully connected layers, and dropout layers, with softmax used as the

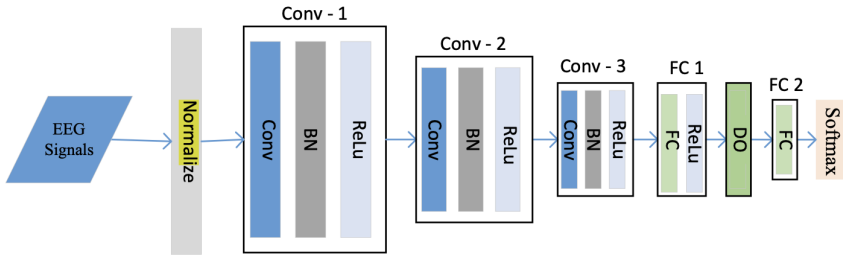


Figure 3.2: Deep Pyramidal 1D-CNN architecture (P-1D-CNN) proposed by Ullah *et al.* [89].

classifier. The key distinction among the architectures is that the P-1D-CNN model does not incorporate any pooling layers and this helps to minimize redundant or unnecessary features by using larger strides in the convolutional layers. The proposed system, P-1D-CNN model, achieved an accuracy of 96.27% (with 8326 parameters), while the traditional CNN 96.23% (21366 parameters) on the same University of Bonn dataset [82]. Figure 3.2 presents the architectural structure of proposed P-1D-CNN architecture.

On the other hand, Lawhern *et al.* [69] proposed a compact CNN architecture, named EEGNet, specifically designed for EEG analysis. Unlike the work of Daoud *et al.* [88], where multiple models were evaluated on the same dataset, EEGNet focused on identifying a single model with the smallest number of trainable parameters. The proposed network, with fewer than 2,500 parameters, was tested across four different datasets. The network, shown in Figure 3.3, was able to reduce the number of trainable parameters by considering (i) depthwise convolution, i.e. convolutions that are not fully-connected to all previous feature maps, (ii) separable convolution which decouple the relationship within and across feature maps. In this approach, the model first learns a kernel that individually summarizes each feature map, then optimally combines the resulting feature map [69]. Initially, EEGNet was applied in the context of general EEG-based brain-computer interface (BCI) tasks [69], and it was later extended to applications such as seizure detection [90] and emotion recognition [91]. Thanks to its compactness and generalization capability, EEGNet has

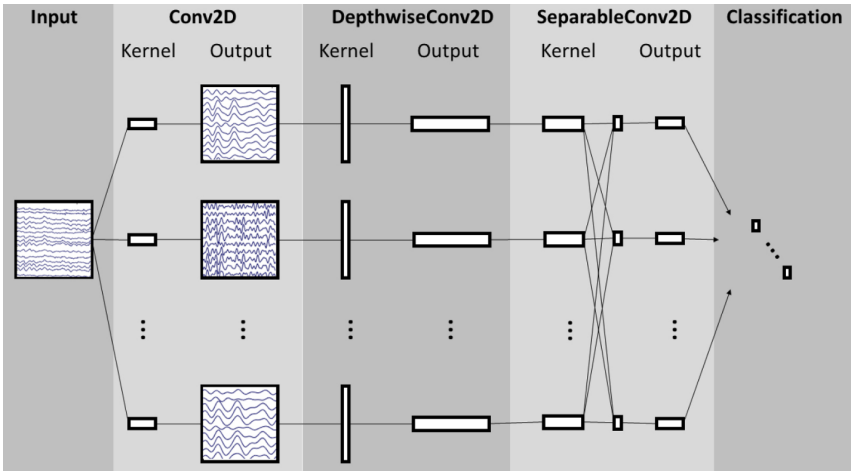


Figure 3.3: EEGNet architecture proposed by Lawhern *et al.* [69].

established itself as a state-of-the-art solution in the field of lightweight DL models for EEG analysis.

An even more remarkable result was reported by [74], where a DL model with a footprint of only 45 kB was obtained. Unlike the studies by Ullah *et al.* [89] and Lawhern *et al.* [69], which focused on architectural modifications, the work by Zhao *et al.* [74] addressed the challenge of reducing model complexity in terms of learnable parameters through a technique known as Neural Architecture Search (NAS). NAS automatically designs an optimal neural network for a given classification task by operating in three main phases: (*i*) defining the target model, (*ii*) selecting the search strategy to reach the target model, and (*iii*) evaluating its performance.

In [74], the NAS algorithm was applied to the task of epileptic seizure prediction. The pre-defined model, shown in Figure 3.4, was composed of blocks divided into two groups-one focused on temporal resolution and the other on spatial convolution. For both block groups, parameters such as the number of filters, filter dimensions, and pooling layer configurations (indicated by '?' in Figure 3.4) are determined using a reinforcement learning-based method. This approach employs a controller that generates the network architecture, trains it, and then evaluates its performance. The results of these evalu-

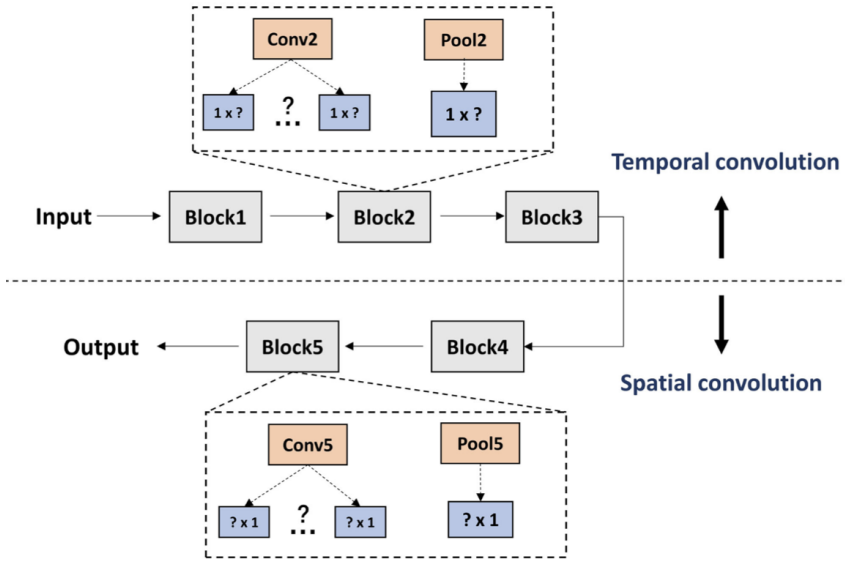


Figure 3.4: Target architecture on which the NAS algorithm has been applied [74].

ations serve as reward signals to refine the controller. After applying NAS, the optimal architecture is found and compressed using techniques like network pruning, quantization, and compact design. This process successfully reduced the model size to under 50 kB without sacrificing performance, making it suitable for low-power wearable devices.

What has been described highlights that it is possible to reduce model complexity while still maintaining good performance in DL, similar to what is achievable with ML. To do this, we refer again to the work of Zhao *et al.* [74], which shows that the most recent hardware for EEG processing predominantly uses ML-based approaches, which typically consume less than μJ/inference and are less than 1 kB in size (see Figure 3.5. However, while ensuring low power and memory requirements, these methods have lower performance than DL models, with sensitivity often not exceeding 95%. This is attributed to their dependence on manual stages of raw data pre-processing and feature extraction. In wearable applications, where automation is a key requirement, such pre-processing may not be feasible and therefore result in a drop in performance compared

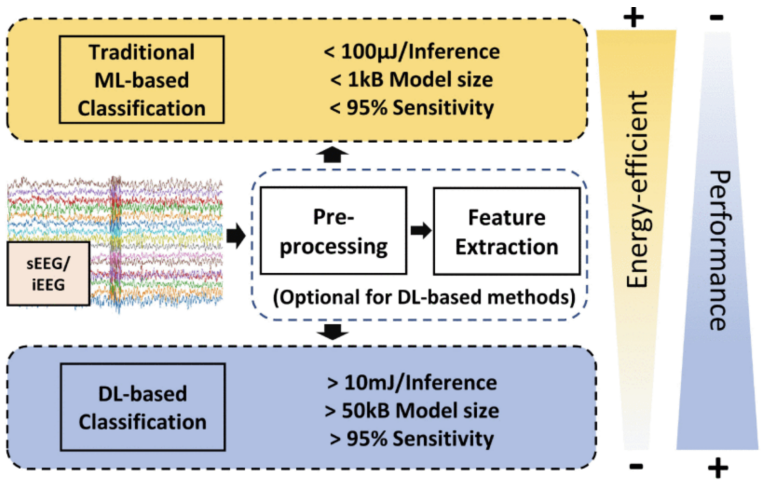


Figure 3.5: Comparison between ML and DL in terms of energy-efficiency and performance [74].

to non-wearable systems, where these operations are typically performed by specialized technicians.

On the other hand, regarding DL models, the authors found that they generally achieve higher sensitivity, although they are computationally intensive and consume significantly more memory. Indeed, these models are typically larger than 50 kB and require over 10 mJ per inference. Moreover, they also reduce the *explainability* of the data, as it becomes difficult to understand which features are involved in the classification task being analyzed. This lack of transparency makes it harder to interpret the decision-making process of the model, which is a notable drawback when compared to traditional ML methods, where the selected features are more explicit and interpretable.

However, these limitations do not preclude the development of lightweight DL solutions, as shown in the study by Zhao *et al.* [74]. On the contrary, by leveraging NAS and model compression, a valuable starting point for adapting DL architectures to the constraints of wearable EEG applications can be provided. In fact, Ingolfsson *et al.* further refined Zhao's architecture by shifting the focus from predicting epileptic seizures to detecting them. By adapting the network to work with only four input channels, they were able to further reduce the computational load and power consumption, thus moving closer to

a practical, real-time implementation in wearable devices with limited resources [92].

3.3.1 Channel selection

The number of available EEG channels generates a trade-off between the performance of the algorithm and the *wearability* of the device. In the literature, various techniques have been proposed to determine the optimal number of channels starting from the traditional 10-20 configuration [20]. The general goal is to minimize the number of electrodes (or sensors) while preserving good performance in AI-based classification tasks.

For clarity, channel selection strategies are generally categorized into filter methods, wrapper methods, and embedded methods [93]. *Filter methods* rank channels based on statistical or information-theoretic criteria, independently of the learning algorithm. They are computationally efficient and straightforward, making them suitable for feature selection. Examples include correlation-based measures, such as the Pearson-Fisher (PF) correlation, which applies Fisher's z transformation to assess inter-channel relationships [94], as well as variance analysis and entropic features that quantify EEG signal variability and uncertainty [95, 96, 97]. For instance, the work of Duun-Henriksen *et al.* [94] compared the variance of EEG signals across different epileptic events for each channel in the standard montage [20], ranking channels by maximum variance and selecting those above a threshold. This approach helped narrow down the number of EEG channels to three for reliable seizure detection, and to the best of our knowledge, it was the first to achieve the lowest number of EEG channels through methods that do not rely on visual inspection or expert opinion. For this reason, it has been cited by future works, such as [97, 96].

Despite the results achieved, filter methods are more effective with ML models that rely on a limited number of features extracted from pre-processed EEG. In contrast, raw EEG signals, which are inherently noisy, require alternative approaches. In this perspective, *wrapper methods* were proposed to evaluate subsets by iteratively training and validating an AI model,

selecting those that potentially achieve higher accuracy at the expense of increased computational load [98, 99]. For example, in [98], seizure detection began with 22 channels. Then, subsets of channels were progressively removed through an iterative procedure, continuing until the reduction in the number of EEG channels resulted in a corresponding decrease in accuracy. Similarly, in the work of Yildiz *et al.* [95], authors eliminated non-discriminative channels iteratively while optimizing accuracy evaluated through training and validation of a Linear Discriminant Analysis (LDA) model. In other cases, the iterative process was combined with filter strategies, as in the study of Daoud *et al.* [88], where the channel set for each iteration was selected based on specific features (like the highest variance-entropy product). They began with all channels [20] and selected those with the highest variance-entropy product which yielded prediction accuracy. At each step, they evaluated whether the accuracy of the model improved with the addition of a channel. A similar approach was employed in [96], where a RF model was used to identify the three most informative channels by ranking feature contributions.

Embedded methods integrate channel selection into the training process itself, often through regularization techniques such as LASSO (Least Absolute Shrinkage and Selection Operator). For example, Peng *et al.* used LDA to rank channels and select the top K channels. From each of these channels, M features were extracted to train a SVM classifier. LASSO was then employed to jointly optimize the K-M pair for each patient based on classification performance [100]. These methods strike a balance between the computational efficiency of filter methods and the accuracy of wrapper methods [93].

Beyond these algorithmic strategies, there are also expert-driven visual approaches that involve professionals, such as neurologists, and alternative techniques based on pre-defined electrode positions depending on the specific wearable device being implemented. In the first case, as proposed by Chung *et al.*, neurologists visually identified four specific channels on the scalp for seizure detection based on their knowledge of distinctive seizure locations for each patient [35]. On the other hand, the idea of using for instance glasses for EEG signal ac-

quisition, named *e-glasses*, led other researchers to select predefined fronto-temporal channels based on the specific electrode positions required by the wearable technology [86].

3.3.2 EEG wearable devices

This section provides information on the main efforts currently in progress to develop wearable EEG devices for home monitoring, reflecting the scientific community's growing interest in the technological aspects of this field. This trend not only justifies but also highlights the importance and relevance of the work carried out in this thesis. Despite the progress made, many wearable EEG systems still have critical issues related to usability and practicality. Within this scenario, the relevance of the present work becomes clear, as described in section 1.1.

As seen, traditional ambulatory EEG relies on a centralized system where all electrodes are connected to a central recorder. This setup requires wiring between the electrodes and central recording unit, which introduces artifacts and significantly reduces wearability (Figure 3.6A). Moreover, due to high costs, not comfortable nature and long waiting lists [101], non-hospital-based approaches have been developed, referred to as home video-EEG telemetry (Figure 3.6B). This monitoring integrates multiple data streams, including acquisition of audio signal via microphone, acquisition of video signal via camera, EEG and ECG (i.e. electrocardiography, the electrical activity of the heart) acquisition and seizure diaries kept by the patient (or caregiver). Although such systems enable ambulatory EEG recordings in a home-setting and are considered the accepted gold-standard for seizure detection and diagnostic evaluation in clinical practice [43], they remain bulky and impractical for unobtrusive everyday use [11].

For this reason, recently, several sophisticated, lightweight, and accurate EEG recording devices with wireless transmission have been developed, such as the Emotiv Epoc headset [102]. These devices (Figure 3.6C) enable continuous or periodic EEG monitoring, allowing patients to freely move inside or even outside the home during data acquisition. The headset provides 14 high-resolution wireless channels, with a common mode sense



Figure 3.6: (A) Traditional ambulatory acquisition system. (B) Home-setting video-EEG monitoring. (C) *Emotiv Epoc* devices. (D) *Byteflies* Sensor Dot with four electrodes placed on the mastoid bone behind the ears of the patient and wired connected to the device attached to the upper back. (E) *cEEGrid* sensor characterized by 10 electrodes arranged in a C-shape and attached around the ear. (F) *Epilog* sensor attached through "stickers" that serve as the interface between Epilog and the scalp when used below hairline

and a driven right leg electrode in place of conventional ground electrode, which are located at the third and fourth parietal nodes (P3/P4), respectively. Signal acquisition is performed using a multiplexed single-channel sequential analog-to-digital converter (ADC). The on-board ADC is actually 16 bits but 2 bits are below the instrumental noise floor and are discarded. The sampling rate is a fixed 2048 Hz, which is multi-channel sampled and yields 128 samples per second (SPS) per channel. The resolution is 14 bits where the least significant bit (LSB) represents 0.51 μ V. The input referred dynamic range per channel is 8400 μ V peak-to-peak and all channels are AC coupled. The buffered data are transmitted to a USB receiver via wireless communication across a 2.4 GHz band. In addition, the headset integrates a real-time impedance monitoring system to assess contact quality, ensuring stable referencing of the differential amplifiers to the ADC's reference voltage.

More recently, efforts to miniaturize EEG-devices have mainly focused either on the electronics of the recording device or on the electrode placement setup. In this regard, companies such as Byteflies have developed a device, named Byteflies Sensor

Dot (SD) (Figure 3.6D), characterized by four electrodes placed on the mastoid bone behind the ears of the patient. The electrodes are connected to a small SD of $24.5 \times 33.5 \times 7.73$ mm and 6.3 g, attached to the upper back via an adhesive [103]. Battery life and memory storage were 24 h and 2 Gb, respectively. Unlike the Emotiv Epoc device, in this case the acquired data are extracted by placing the SD in a docking station, which automatically uploads EEG data to the cloud via a Wi-Fi connection, allowing also a recharging of the battery.

In addition to these commercially available products, research has focused on different aspects of the miniaturization of EEG devices, particularly on electrode technology and the possibility of self-application. In this regard, ear-EEG systems have shown promising results, offering both good signal quality and high wearing comfort. These systems integrate all necessary EEG components (recording electrode, ground, and reference) into the outer ear canal, the concha, or the area around the ear (Figure 3.6E). By directly comparing simultaneously acquired scalp-EEG and ear-EEG signals, several independent research works have shown that ear-EEG can capture brain signals that are closely related to those recorded with scalp-EEG, as summarized in [104].

Starting from these considerations, cEEGrids devices can be considered as wearable-EEG devices. cEEGrids are made from conductive Ag/AgCl based polymer thick film ink and are characterized by 10 electrodes arranged in a C-shape and attached around the ear through a double-sided adhesive tape. The conductive surface of the contacts is circular with a diameter of 3 mm and the distance between electrodes is either 12 or 18 mm (center to center). As shown in Figure 3.6E, cEEGrids are typically used in combination with a small, wireless EEG amplifier that is placed at the back of the head with a headband.

To address the need for a wireless, easily applied, and wearable EEG system, the company Epitel developed the Epilog miniaturized EEG sensor designed to record high-fidelity EEG during the daily activities of the patient. The device uses a differential electrode pair positioned approximately at T5-T6 and F7-F8 according to the standard International 10-20 sys-

tem [20] (see Section 2.1.2, with an 18 mm center-to-center spacing, to acquire a single EEG channel. In addition to reliable recording performance, Epilog sensors are robust, water-resistant, and engineered to withstand the rigors of daily use. Each sensor connects directly to a Remote EEG Monitoring (REMI) platform, developed again by the company Epitel, which manages the signals acquired by sensors. In practice, the EEG data from four Epilog sensors are uploaded directly to a compatible database and converted into a 10-channel REMI montage, as detailed in [105].

The above-described devices, mainly focused on EEG signal acquisition, demonstrate that it is now possible to record EEG signals using dry electrodes and transmit them via wireless to a base station for subsequent processing. The next step forward is the integration of real-time signal processing directly on an embedded system.

3.4 GNN-based approaches

Recently, research progress on the application of GNNs to EEG signal analysis in the context of seizure detection has attracted growing interest. As shown in Fig. 3.7, the results of research conducted on Scopus using the keywords graph neural network or GNN, together with EEG and the terms seizure or epilepsy, show a strong upward trend. This trend highlights how GNNs are emerging as one of the most promising approaches for EEG analysis, including in the field of epilepsy [50].

The traditional and lightweight models widely used for feature extraction from EEG signals discussed in Sections 3.1 and 3.3 assume that the electrodes were arranged on a regular, equidistant grid, similar to the pixels of an image. However, this assumption neglects the complex functional relationships that exist between different brain areas [55]. Although extensions of convolution and pooling operations (i.e., 2D operations) have been proposed to better capture the underlying network structure between electrode sites, these strategies remain constrained to a grid-based paradigm.

To overcome this limitation, attention has shifted to models

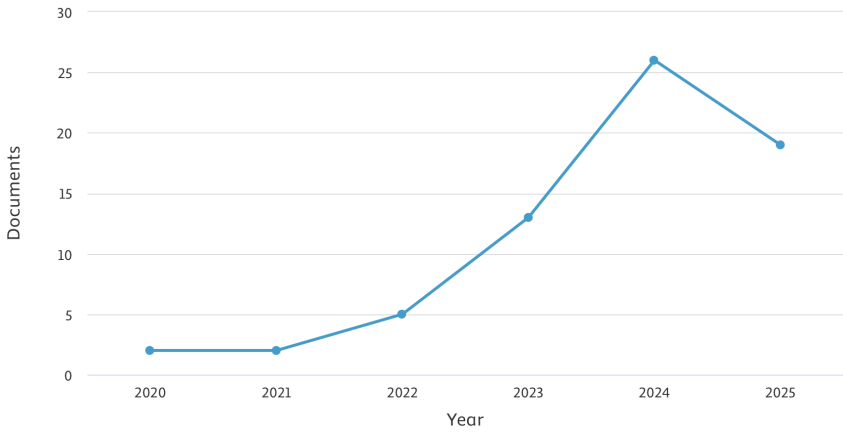


Figure 3.7: Recent trends in EEG-GNN applications.

capable of more accurately representing the inherently non-Euclidean nature of EEG data [53]. This approach is part of a broader trend in ML, where representing data as graphs has proven to be effective in several domains, from pattern recognition to molecular biology, chemistry, computer vision, and data mining [56]. In line with this, GNNs have also found increasing application in neuroscience [53], modeling electrodes as nodes and functional or anatomical relationships as edges, so as to capture the complex patterns of EEG signals through brain connectivity (see Section 2.3.1).

The most common way to apply GNNs in this context is to use static architectures based on models such as GCNs [56]. An example is the Epilepsy EEG Graph Convolutional Network (E-GCN) model proposed by Chen *et al.*, structured on five layers of graph convolution to exploit the correlations between EEG channels and distinguish normal from epileptic tracings. Convolution on the graph is implemented in approximate form using Chebyshev polynomials, a solution that reduces computational costs and allows information from non-directly adjacent nodes to be included [106]. However, a limitation of these approaches is the absence of a universally accepted method for defining the edges of the graph.

To address this problem, Lian *et al.* proposed the Joint Graph Representation Network (JGRN) model, which jointly learns both the graph structure and node representations. The

network estimates patient-specific graph structures directly from iEEG signals, simultaneously optimizing the graph and features with respect to the goal of predicting epileptic seizures [107]. On the same side of architectures with self-learning connectivity, Zhao *et al.* developed a model based on an optimized GAT, capable of dynamically learning the attention weights between electrodes and therefore the spatial connectivity. This model achieved a sensitivity of 98.33% on the CHB-MIT dataset for seizure/non-seizure classification, demonstrating the effectiveness of GATs in adapting to the specificities of the EEG signal [108].

Although the approaches described have shown promising results, they do not explicitly model the temporal dynamics of EEG signals. To overcome this limitation, dynamic GNN models have been developed, combining a spatial backbone with a temporal module. In this regard, He et al. proposed an architecture that uses a GAT for spatial feature extraction and a BiLSTM as a temporal back-end, achieving 98.52% accuracy (97.75% sensitivity, 94.34% specificity) on the CHB-MIT dataset in detecting seizures from raw EEG. Furthermore, Yan *et al.* introduced the DTS-GAN model, which integrates an LSTM encoder to capture long-term dependencies and a dynamic graph attention with probabilistic Gaussian connectivity to learn functional interactions between electrodes [109]. On the TUSZ dataset, this approach reported accuracies between 89% and 91%, demonstrating the effectiveness of dynamic representation in capturing the spatio-temporal relationships underlying EEG signals.

Similar pipelines have been implemented by using a GCN instead of a GAT while keeping the LSTM as the temporal component. In this direction, Liu *et al.* further improved epilepsy detection performance by combining a GCN with a BiLSTM for the joint extraction of contextual and temporal features from EEG signals, then integrating an extensive learning system to automate the diagnostic process [110]. Kuang *et al.* also proposed a pipeline based on GCN and LSTM. In this case, the graph was obtained using Pearson's correlation between EEG channels to reflect the functional inter-dependencies between electrodes. Starting from 20-second EEG segments, features

were extracted in the time and frequency domains and fed into the GCN. The sequences thus generated were then processed by an LSTM module for temporal modeling and finally classified using a fully connected layer followed by softmax.

The idea of concatenating an RNN component to a static GNN network is therefore the most common strategy for building dynamic models. Indeed, few studies explored alternative approaches, such as the use of temporal transformers. In this direction, Xiang *et al.* proposed SGSTAN, which combines a GAT for spatial processing on synchronized graphs with a Transformer for temporal modeling. On CHB-MIT, the model achieved an accuracy of approximately 98.2% and a sensitivity close to 97.9%. However, the computational cost of such an architecture remains high, limiting its transferability to real-time clinical scenarios. This highlights a significant gap in the literature related to the fact that most of the proposed models, while showing high performance, are too heavy to be applied in continuous monitoring contexts, such as those required in clinical practice or wearable devices [111].

This need has led to growing interest in the development of lightweight variants of GNNs, capable of balancing accuracy and computational sustainability. An example is provided by Jia *et al.*, who propose a simplified version of spectral ChebNet/GCN for seizure prediction on CHB-MIT. The model has approximately 15.5k parameters and achieves competitive performance (sensitivity of 96.5%), discussing its suitability for low-power wearable scenarios [57]. However, this is a static approach, with the graph defined based on the distance between electrodes and/or derived features such as band energy and Hjorth parameters. To combine lightweight and efficient modeling of temporal dynamics, Xu *et al.* proposed a GCN with BiGRU network, designed to be implementable in resource-limited scenarios. The GCN allows spatial dependencies to be captured, while the BiGRU efficiently models temporal dependencies. The results obtained on CHB-MIT show competitive performance with significantly reduced complexity, confirming the possibility of balancing accuracy and computational costs through lightweight but hybrid architectures [112].

3.5 Summary

This literature review pioneers a comprehensive overview of AI-based EEG systems, highlighting their growing importance and urgency. From what has been described, it is clear that the use of automatic learning systems has represented a turning point in the analysis of EEG signals [59, 66].

We have seen how, over the years, overcoming the limitations of ML methods based on manually extracted features, CNN and RNN models have established themselves as valid approaches in the analysis of EEG signals in various applications, as they are able to capture the dynamic behavior of the signal [60]. Additionally, it is worth noting that more recently attention has shifted to Transformers, architectures that, thanks to the self-attention mechanism, allow non-local relationships and long-range dependencies to be captured, showing promising results also in the detection of epileptic seizures [66]. However, in this context, Transformers appear to be computationally expensive despite their strengths. This is because self-attention has quadratic complexity in time and memory relative to sequence length, and can be inefficient in real-time scenarios or with limited resources [113].

In summary, despite the progress made over the years, applying AI methodologies to EEG signals still presents significant challenges. First, the management of spatial complexity linked to the high number of electrodes and to the need to accurately characterize the interactions between them. Second, the presence of intra- and inter-subject variability constrains the generalizability of models across patients and across repeated sessions for the same patient. Finally, the limited interpretability of many DL-based models, which hinders understanding of the underlying neural mechanisms and slows down their adoption in clinical or applied settings.

Within this scenario, this thesis proposes two approaches. The first concerns the reduction of the number of EEG channels, with the aim of limiting the dimensionality of the data without compromising the discriminating information. The second involves the use of graph-based models, able of explicitly representing the topological and functional relationships

between channels. The relevance of this work, i.e. the management of the spatial complexity of the EEG signal, is further emphasized by a parallel trend: the development of wearable EEG devices (see Section 3.3.2), which require lightweight and efficient AI models suitable for on-device execution. From this perspective, the two lines of research are complementary: channel-reduction strategies and simplified graph-based models, both targeting computational efficiency and real-world applicability. Both approaches are also developed with the aim of addressing the issue of inter- and intra-subject variability. Specifically, in this study we focus our attention on a patient-specific strategy, comparing subject-independent and patient-specific configurations in light of the variability observed within and between patients.

The first line of research concerns the design of computationally efficient DL models, with the aim of reducing computational costs while maintaining high performance. Techniques such as pruning and quantization have been used to lighten complex models. For example, LightSeizureNet introduces dilated 1D-CNN, global average pooling, and kernel-wise pruning for real-time seizure detection [114]. At the same time, architectures designed from the outset for resource-constrained scenarios have emerged, such as CNNs with separable depthwise convolutions from EEGNet (Lawhern *et al.* [69]), which drastically reduce parameters and MAC (multiply-accumulate operations) while preserving discriminatory power in various EEG tasks, and pyramidal 1D-CNNs, proposed to contain parameters in EEG time series with good performance in epilepsy detection [89]. These approaches aim to make on-device deployment (embedded/wearable) a reality, where energy and memory constraints are stringent. In parallel, the literature on EEG channel selection highlights that, despite numerous proposals, from filter, wrapper and embedded methods to expert-guided strategies, there is still no universally accepted, fully automatic method capable of reliably identifying the optimal subset of channels for epilepsy-related tasks.

Although it is clear from the above that both model compression and channel selection are extensively addressed in the literature, these aspects are mostly dealt with separately. In-

deed, the joint reduction of the number of electrodes and the use of an explicitly "*lightweight*" network for seizure detection remain relatively unexplored to date. This area requires co-optimization between compact architectures and wearable constraints and represents a research direction with a strong application impact and worthy of further structured investigation.

To the best of our knowledge, a line of work that is particularly consistent with this goal is that of a scientific group that has proposed both a compact Transformer for microcontrollers (EEGformer) and a lightweight CNN network (EpiDeNet) [92] for on-device scenarios. Both models are evaluated on the public CHB-MIT dataset [79] and are designed for temporal channel (i.e., F_7-T_7, F_8-T_8, T_7-P_7, T_8-P_8). This choice is constrained by the electrode layout of a prototype e-glasses[115]. EEGformer reports a sensitivity of 65% on a subset of 8 patients. EpiDeNet, obtained by combining the network optimization proposed by Zhao *et al.* [74] with the exclusive use of the four-fixed temporal channels mentioned above, achieves 91.16% detection of ictal events, with 68.73% sensitivity and 99.75% specificity.

Since this direction appears promising, in the present work we take these contributions as a baseline and replicate their results. In particular, we adopt EpiDeNet as the main baseline, because EEGformer was presented only on a subset of subjects and reports lower performance than EpiDeNet. Next, instead of fixing the channels a priori based on the design constraints of the wearable device, we propose an adaptive and patient-specific selection methodology that can jointly optimize the computational efficiency of the lightweight model and the portability of the EEG acquisition system.

The second line of research focuses on graph-based models, which have proven particularly suitable for handling the spatial complexity of EEG signals. In this paradigm, each electrode is represented as a node in a graph, and the connections (edges) model structural or functional relationships between pairs of channels. Therefore, GNNs allow us to learn representations that take into account brain topology and dynamics, capturing patterns that traditionally sequential or convolutional models

might overlook.

Graph construction is a central aspect of these approaches, particularly with regard to the definition of spatial relationships. In this regard, the literature proposes several metrics for defining connectivity between channels. Among the most widely used are PCC, frequency coherence, PLV, and various measures based on entropy or mutual information. The choice of connectivity measure directly affects the model's ability to represent meaningful relationships and is therefore crucial for obtaining informative graphs without resorting to overly complex models [116]. This choice cannot be random but must be carefully considered. Indeed, although comparative studies suggest that FC tends to yield better results than distance-metric approaches, no FC metric has been identified as universally superior to others [46]. Furthermore, the selection should also take into account the configuration of the EEG signal, i.e., whether it is recorded in monopolar or bipolar montage. Although many studies do not consider this aspect, it is far from marginal since bipolar reference, while improving the detection of local cortical activity, can introduce artificial dependencies between channels due to shared electrodes [117].

It is worth remembering that choosing the most appropriate connectivity measure also depends on the classification task to be pursued. For example, in the context of EEG synchronization, phase-lag only measures reduce spurious zero-lag connections due to volume conduction, producing cleaner and often sparser graphs, at the same threshold, compared to classical correlation indices [118, 54]. On the other hand, in emotion recognition [119] or motor imagery [120] PCC has repeatedly been reported as the most effective measure, sometimes further improved by metric fusion strategies [121]. All these studies highlight that an appropriate choice of connectivity measure enables more stable and discriminative performance, although in the epilepsy context there is still no systematic evidence quantifying whether PLV or PCC is more effective. This also suggests the possibility of employing models with reduced complexity, since more informative inputs can compensate for the need of highly sophisticated architectures.

Another factor influencing graph-based model complexity

is the density of the adjacency matrix. Since the computational cost of graph convolutions increases with the number of edges [116], sparsification strategies have been investigated showing that sparse graphs, when built from robust connectivity measures, can preserve performance comparable to dense graphs while reducing memory footprint and training time [122, 55]. These findings suggest that the combined use of task-appropriate connectivity measures and controlled sparsification may reduce the need for highly complex GNN architectures. However, a systematic analysis of this aspect in the specific field of seizure detection with GNNs is still missing.

In parallel, the increasing interest in wearable devices and real-time applications has driven also the graph-based research towards lightweight architectures. Variants that reduce the number of parameters or the complexity of aggregation operations have been proposed, for example through edge dropout or simplified operators such as low-order Chebyshev-polynomial (ChebNet) operators. In this direction, Jia *et al.* designed an efficient GCN for epileptic seizure prediction from scalp EEG, based on three Chebyshev graph convolutions, global pooling, and a compact fully connected head [57]. Similarly, Barbera *et al.* introduced a minimal GCN pipeline for dementia assessment, employing a few low-order GCN layers, node normalization, and graph pooling before a very small MLP classifier [123]. Both works aim to optimize the ratio between accuracy and number of parameters, favoring implementation on resource-limited devices.

Starting from the gaps highlighted in the literature, i.e. the lack of consensus on the most appropriate connectivity metric, the lack of shared criteria for the graph sparsification threshold, and the lack of GNN networks that are both lightweight and dynamic, in the second part of this thesis we propose an exploratory study that combines these three aspects. In particular, we conduct an analysis focused on seizure detection, evaluating in detail the two most widely used connectivity metrics in the literature, PCC and PLV, with the aim of identifying a compromise between connectivity measure selection and graph sparsification to improve the informativeness of the graph while reducing its complexity. Subsequently, we adopt GCN layers

based on ChebNet to reduce the complexity of the GNN model, enriched with recurrent components (like LSTM) capable of capturing the temporal dimension. In this way, we can propose a lightweight but dynamic network, which moves in the direction of real applicability in wearable contexts.

Another important aspect of this second line of research is the interpretability of models, one of the main challenges in applying GNNs to brain research. In a clinical setting, however, interpretability is crucial, since in addition to diagnostic accuracy, it is necessary to provide insights into the neurophysiological mechanisms underlying epileptic activity. In recent years, some explainable GNN approaches have begun to bridge this gap. Mazurek *et al.*, for example, proposed a GNN with attention mechanisms capable of highlighting the features and graph structure relevant to classification, showing consistency with clinical interpretation [124]. However, since attention mechanisms can be computationally expensive, lighter alternatives have been proposed, such as GNN-Explainer, which highlights the nodes and edges of the brain network that are most relevant to a given prediction or classification, thereby increasing the transparency of the decision-making process [125]. Although such tools have been successfully applied in other domains, they have not yet been explored in the context of seizure detection. In this study, we address this gap by applying GNN-Explainer to EEG data, with the aim of identifying the nodes and edges of the brain graph that are most relevant for seizure classification.

Chapter 4

Proposed methods to manage complexity

In this Chapter, we present the main contributions of the thesis: an optimized data-driven algorithm for EEG channel selection and an optimized graph representation of multi-channel EEG. As indicated in Section 1.2, the central challenge addressed in this thesis is to reduce spatial complexity, making it clinically interpretable and functional for the detection of epileptic seizures. This enhances the discriminatory power of the input, enabling lighter architectures with lower computational costs. To tackle the issue, two complementary methods are proposed. First, a minimal yet informative subset of EEG channels is selected via data-driven criteria, which can be visualized on topographic maps (Section 4.1). Second, a graph model exploits task-specific functional connectivity to capture inter-channel dependencies while preserving a compact, explainable architecture, with interpretability provided through GNNExplainer (Section 4.2). The experiments are described in the following chapters, while the metrics adopted for evaluating the proposed methods are reported in Section 4.3.

4.1 Reduced-Channel Framework

This section describes the proposed method for selecting EEG channels and the related framework, which enables the deploy-

ment of a fully automated seizure detection algorithm, also suitable for operation on micro-controllers. The approach computes channel weight coefficients to maximize detection performance, identifying the most representative channels for each phase. These coefficients are derived through Principal Component Analysis (PCA), where the loadings are calculated and subsequently weighted by the explained variance, thereby emphasizing the channels most relevant for discriminating between ictal and non-ictal phases.

The choice of PCA was motivated by both methodological and practical considerations. As discussed in Section 3.3.1, several channel selection strategies have been proposed in the literature. However, many of these approaches either require extensive model training or depend strongly on the classifier used. In contrast, PCA was selected because it provides a computationally efficient, stable, and interpretable solution, which is particularly suitable when only limited clinical data are available. By jointly modeling the variance structure of multichannel EEG, PCA captures the dominant inter-channel correlation patterns while reducing noise and redundancy. Furthermore, in the proposed framework, PCA is not employed for conventional dimensionality reduction, i.e., to replace the original variables with principal components. Instead, it is used to derive channel-wise importance coefficients from the loadings weighted by the explained variance. In this way, the method preserves direct interpretability in the original channel space, while exploiting the global covariance structure of the multichannel EEG. This makes PCA especially suitable for identifying the most informative channels in a compact and robust manner, compared with approaches based only on variance measures or simple pairwise channel comparisons proposed in literature.

The selected channels are then provided as input to a lightweight Convolutional Neural Network (CNN), specifically designed for efficient deployment on low-power hardware, enabling real-time operation. The use of a CNN for time-series processing was also a deliberate choice. Although CNNs were originally developed for image analysis and have sometimes been regarded as suboptimal for generic time-series modeling,

in EEG-based seizure detection they have proven highly effective in extracting local temporal patterns from short signal windows, while maintaining low computational cost and efficient parallelization. These characteristics are particularly relevant in the present study, whose main objective is the design of a resource-efficient seizure detection framework. Moreover, once the channel-selection stage has reduced the input dimensionality, the classification task mainly consists in identifying discriminative local temporal structures within a few selected channels, rather than modeling very long-range dependencies. Under these conditions, a lightweight CNN offers a more suitable trade-off between performance, interpretability, and computational efficiency than more complex recurrent or attention-based architectures.

4.1.1 Introduction

Epileptic seizure detection and management represent crucial research areas, with EEG widely recognized as a powerful diagnostic tool [126]. In clinical practice, raw EEG data are acquired through multiple electrodes placed on specific areas of the scalp, continuously recording the brain's electrical activity. As a result, EEG data are high-dimensional and typically contain noise, artifacts, and task-irrelevant information. Existing deep neural network approaches in the literature typically use the full set of EEG channels distributed across the scalp as input data, achieving high performance but resulting in systems that are impractical for everyday life monitoring with wearable EEG devices. The opportunity of recording EEG outside clinical settings allows long-term recordings,thus increasing the probability of detecting epileptic episodes [127]. Furthermore, seizures that would not occur in the hospital can be recorded and the neurologist has access to a more representative recording of the patient's seizures. However, to ensure the adoption and extended use of these wearable EEG devices by individuals with epilepsy, they must be designed to be comfortable, compact, and discreet [128].

To address these challenges, different research groups have focused their efforts on reducing the number of electrodes re-

quired, making the system less invasive and more feasible for daily use. Several solutions have been proposed in the field of epilepsy, including wearable devices integrated into eyeglasses [86] or positioned behind the ear [129, 130]. Following this trend, four temporal channels were recently used to define various algorithms tailored for microcontroller units, including a transformer-based model [131] and a Convolutional Neural Network (CNN) [92] with performance comparable to that obtained by a standard 10-20 system based on 21 channels [74]. However, seizures can be generalized or focal onset, with focal seizures often originating in specific regions of the brain, but highly variable and dependent on the specific patient [132]. This variability complicates optimal channel selection, especially for focal seizures, which account for approximately 60% of cases.

Therefore, although these methods based on four fixed channels allow for generalizing the technical solution based on implementation needs, while also improving patient comfort, they may result in sub-optimal channel selection from an accuracy point of view, especially when focal seizures occur in regions not covered at all by the selected channels. This has motivated my research aimed at developing a subject-specific approach, in which the selected channels vary for each subject and may not be suitable for others, in line with the new vision of *personalized medicine* [133].

4.1.2 Related works

Lightweight models. CNNs are widely used for seizure detection in EEG signals due to their ability to automatically learn hierarchical features directly from raw data [73], thus reducing the need for manual feature engineering. However, the growing interest in real-time and resource-constrained applications, such as wearable EEG devices, has highlighted the need for lightweight architectures capable of maintaining high detection performance while reducing computational cost and energy consumption.

As discussed in Section 3.3 of the literature review chapter, the evolution of this research area has progressively led to

the development of increasingly compact and efficient neural networks. In this context, several lightweight seizure detection approaches have been proposed in recent years, often relying on the full set of EEG channels provided by standard acquisition systems, such as the 10–20 electrode placement system [20]. Within this line of research, the work of Qiu *et al.* represents a relevant contribution. The authors introduced LightSeizureNet [114], a lightweight architecture that applies dilated 1D convolutions to each of the 21 channels of a standard EEG system. The model further incorporates global average pooling and kernel-wise pruning to compress the network and reduce its computational complexity. Despite its compact design, the architecture maintains high detection accuracy and provides fine-grained information, such as the brain regions and frequency bands activated during seizure events, thereby supporting physicians in clinical diagnosis.

In the same year, Zhao *et al.* proposed a Neural Architecture Search (NAS) approach to design a compact CNN architecture with approximately 11,000 parameters, achieving high sensitivity with energy consumption in the millijoule (mJ) range per inference [74]. The resulting network consists of five convolutional blocks, each including a convolutional layer followed by a pooling layer, designed to capture the essential temporal and spatial features of EEG signals. This architecture was later refined by T. M. Ingolfsson *et al.* to demonstrate its feasibility for implementation on microcontrollers [92]. Their proposed network, named EpiDeNet, further improved efficiency by adopting a four-channel input configuration, thereby reducing computational load and power consumption. However, the choice of channels was mainly driven by hardware and wearable design constraints, rather than by an assessment of the most informative EEG locations. For this reason, building upon this line of research, the present work investigates a further optimization of the input representation by reducing the dimensionality to only two EEG channels selected according to their informativeness for seizure detection.

The works discussed in this section represent a selected subset of studies chosen for their relevance to the objectives of the present work. In particular, priority was given to approaches

developed and validated on the same public datasets considered in this thesis, so as to enable a more meaningful and consistent comparison of both quantitative results and the related clinical considerations with respect to the proposed framework.

Channel selection. In previous studies, several data-driven approaches have been proposed to identify EEG channel subsets that achieve detection performance comparable to full-channel systems. These methods are generally classified into (*i*) filtering, (*ii*) wrapper, and (*iii*) embedded methods [93]. *Filtering methods* rank channels using statistical or information-theoretic criteria, such as correlation, variance, or entropy, and are computationally efficient, classifier-independent, and highly interpretable. This makes them well suited for clinical applications where selected features must be easily understood by medical experts [97, 134]. However, they may fail to capture complex inter-dependencies between channels.

Wrapper methods, instead, iteratively evaluate channel subsets by training classification models (i.e., Random Forests (RFs), K-nearest neighbors (KNN)) often yielding higher accuracy than filtering approaches but at the expense of greater computational cost and an increased risk of overfitting, particularly when working with small datasets [135].

Embedded methods integrate feature selection into the model training process, using regularization techniques such as the multi-head attention mechanism [136] or LASSO [137], thus balancing performance and computational demands. While more efficient than wrapper approaches, they are limited to models that inherently support integrated selection, reducing their general applicability [100]. Despite these advances, no single objective, reliable, and fully automated method has emerged as the definitive solution for EEG channel selection.

For this reason, the expertise of clinicians remains crucial in practical applications. In fact, studies aiming to reduce the number of channels to fewer than three [135, 97, 136, 134] frequently adopt expert-guided visual methods, where neurologists leverage domain knowledge to optimize channel placement based on each patient's unique seizure onset locations [35].

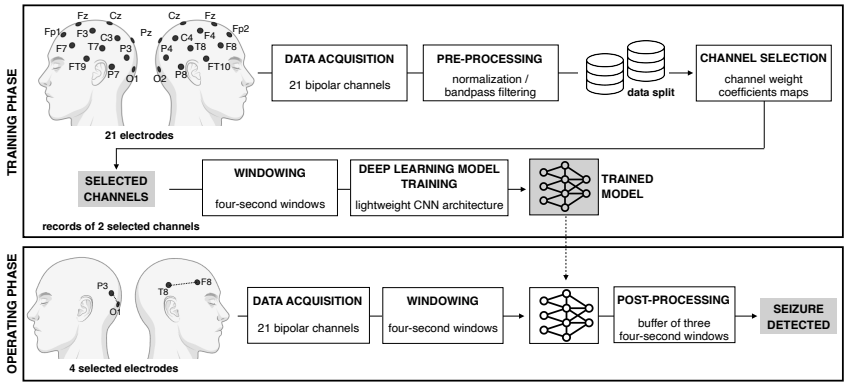


Figure 4.1: Block diagram of the proposed reduced-channel framework. Training phase: EEG channels are selected and used to train a lightweight CNN. Operating phase: the test records from the selected channels are processed, and the model output is post-processed to detect seizure events.

4.1.3 Methods

An overview of the proposed framework is provided in Figure 4.1. The method consists of two main phases: a training phase and an operating phase. In the first phase, EEG signals are recorded from 21 bipolar channels, preprocessed, and then split according to the adopted validation scheme. Channel selection is subsequently performed to identify the most relevant channels for each individual patient, and a lightweight CNN is trained on 4-second windowed segments extracted from the selected channels. In the second phase, EEG data are acquired from only two patient-specific bipolar channels, segmented into 4-second windows, and processed by the trained model to detect seizure onsets. Post-processing techniques are finally applied to refine the detection results and reduce false positives.

In this perspective, the section is organized as follows. First, Section 4.1.3.1 presents the proposed channel selection algorithm. Next, Section 4.1.3.2 describes the implemented classification model, including its architectural design and the related considerations on parameters and latency. Finally, Section 5.1 outlines the experimental setup, including the validation protocol, training procedure, and post-processing details.

4.1.3.1 Channel selection algorithm

The first goal of our analysis, which can also be considered an important contribution of this work, is to identify the scalp regions where to retrieve the most informative content during different phases of epileptic seizures. In turns, this allows to determine the optimal placement of EEG electrodes for a compact wearable device. This task is challenging due to the transient nature of EEG patterns, their high susceptibility to noise and artifacts, and the limited size of most EEG datasets, which increases the risk of overfitting when applying complex algorithms. To address these difficulties, we propose a method based on temporal Principal Component Analysis (PCA) to rank EEG channels based on their contribution to seizure-related activity. Each channel corresponds to a specific scalp region according to the $10-20$ standard system. By decomposing the data into orthogonal, uncorrelated new variables, this approach highlights the most informative channels while suppressing noise and redundancy, making it particularly suitable for small EEG datasets, typically acquired in clinical practice. The procedure adopted is defined in four main steps, as described in the following.

1. Data pre-processing. All EEG records were cleaned of the 60 Hz component using a 1 Hz-wide notch filter (two-order Butterworth filter) and then normalized using the z-score method, i.e., by subtracting the mean and dividing by the standard deviation to obtain zero mean and unitary standard deviation. After normalization, the EEG signals were filtered to extract the following standard frequency ranges: $\delta[0.5, 4]Hz$, $\theta(4, 8]Hz$, $\alpha(8, 13]Hz$, $\beta(13, 35]Hz$ and $\gamma(35, 80]Hz$, using a two-order Butterworth band-pass filter.

2. Channel weight coefficient definition. The 21-channel EEG signal was segmented into 1-second time windows for each frequency band considered. Since the signal was sampled at 256 Hz, each time window was represented as a matrix $\mathbf{X}$ of size $N \times C$, being $N = 256$ the number of samples per second, and $C = 21$ the number of bipolar EEG channels. This segmentation step allowed for a more localized analysis of neu-

ral activity over time, ensuring that transient seizure patterns could be captured.

The generic matrix $\mathbf{X}$ was then decomposed, as follows, into principal components (PCs) that capture the main patterns of variation in the data:

$$\mathbf{Y} = \mathbf{XL} \tag{4.1}$$

$\mathbf{L} \in \mathbb{R}^{C \times C}$ is the loading matrix, whose columns are the eigenvectors of the covariance matrix $\mathbf{S_X}$ of the original data $\mathbf{X}$ [138]. Each j-th column of $\mathbf{L}$ represents the orientation of $\mathbf{PC_j}$ relative to the original EEG channels. Its elements, l_{ij}, known as *loadings*, quantify the contribution of the i-th EEG channel to the corresponding $\mathbf{PC_j}$.

The loading matrix was used here to rank channels based on their contribution to the PCs. However, loadings alone do not indicate how much of the total variance each component explains. Instead, this information is provided by the covariance matrix of the transformed data $\mathbf{Y}$, given by:

$$\mathbf{S_Y} = \mathbf{L}^\top \mathbf{S_X} \mathbf{L} \tag{4.2}$$

$\mathbf{S_Y}$ is a diagonal matrix whose elements s_{Yj} are the *eigenvalues*, corresponding to the variance explained by the j-th PC. PCs with larger s_{Yj} values correspond to dominant neural activity patterns, whereas those with smaller values mainly capture noise or minor fluctuations. As a result, seizure-related activity, which is often localized to specific brain regions, are expected to be reflected by PCs that account for the greatest portion of the data variance. To integrate the explained variance into the channel ranking process, we weighted the absolute values of the loadings by the corresponding eigenvalues:

$$l_{ij}^{\text{weighted}} = |l_{ij}| \cdot s_{Yj}, \quad \text{for } i, j = 1, 2, \ldots, C. \tag{4.3}$$

This approach ensures that the most relevant seizure signatures are prioritized over less informative background activity, which is crucial for diagnostic purposes. Finally, to obtain a compact representation for each 1-second window, we summed the weighted loadings for each EEG channel to produce a single

weighted mean coefficient for each channel, hereinafter defined as the Channel Weight Coefficient (CWC):

$$\mathrm{CWC}_i = \sum_{j=1}^{C} l_{ij}^{\mathrm{weighted}}, \text{ for } i = 1, 2, \ldots, C. \tag{4.4}$$

This transformation reduces each 1-second multichannel segment to a compact vector of CWC values, with each value representing the weight of a specific channel in expressing the relevant information content. This provides us with a tool to select the most informative channels within different phases, rather than reducing the data dimensionality by exploiting the linear combination of original channels, as done in previous work employing the PCA for EEG data [139, 140]. It is important to remark that z-score normalization ensures that all channels are on a comparable scale, preventing variables with larger numerical ranges from disproportionately affecting the PCs. Moreover, being this procedure explicitly based on the eigendecomposition of the covariance matrix, it inherently captures the linear correlation structure among the analyzed variables (i.e., the original EEG channels). Therefore, this method allows for the selection of the most statistically informative and mutually uncorrelated EEG channels, giving the same importance to the informative content within the clinical frequency bands.

An example of the results of the step described is shown in Figure 4.2. Without loss of generality, it specifically refers to the 1-second time segment of record 3 of patient chb01, extracted during the ictal phase in δ band.

In more detail, the map in Figure 4.2a graphically represents the 21×21 loading matrix (with its elements in absolute value) evaluated by means of the PCA, showing the 21 PCs along the x-axis, while the 21 EEG channel labels are displayed along the y-axis. To simplify the notation, the bipolar channels of the standard $10-20$ system were numbered sequentially from 1 (corresponding to the left frontal region) to 21, (corresponding to the right occipital region) coherently with the topographic map of Figure 4.2d [25]. Specifically, numbering follows the scheme: 1: $\mathrm{Fp_1}$-$\mathrm{F_7}$; 2: $\mathrm{Fp_1}$-$\mathrm{F_3}$; 3: $\mathrm{Fp_2}$-$\mathrm{F_8}$;

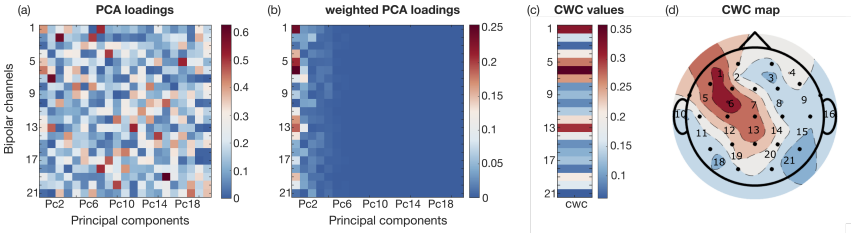


Figure 4.2: Channel weight coefficient calculation steps. (a) PCA loadings matrix computed in a 1-second time window. (b) PCA loadings matrix multiplied by the explained variance associated with each PC. (c) Weighted mean of the PCA loadings, obtained by summing up the values in the corresponding row computed in (b). (d) Topographic map representation of (c).

4: Fp_2-F_4; 5: F_7-T_7; 6: F_3-C_3; 7: Fz-Cz; 8: F_4-C_4; 9: F_8-T_8; 10: T_7-FT_9; 11: T_7-P_7; 12: C_3-P_3; 13: Cz-Pz; 14: C_4-P_4; 15: T_8-P_8; 16: FT_{10}-T_8; 17: FT_9-FT_{10}; 18: P_7-O_1; 19: P_3-O_1; 20: P_4-O_2; 21: P_8-O_2.

The loading values associated with each channel for all the evaluated PCs are visualized using a colormap ranging from blue (lowest values) to red (highest values). Following this approach, the 21×21 points map highlights with red points the channels that contribute most significantly to each PC, which stand out from those making a negligible contribution, represented instead by the blue points. Figure 4.2b shows the updated map obtained by scaling the loading matrix by the explained variance. In this revised map, the colors shift toward blue for higher PCs, indicating that only the first few PCs are most effective in capturing the most relevant information [141, 142]. Finally, Figure 4.2c shows the CWC calculated by summing up the weighted loadings for each EEG channel across all the PCs. The same values are represented in the map of Figure 4.2d, which also serves as a legend for Figure 4.2 a,b,c. In this specific example, it is evident that the analyzed δ activity at this specific instant is localized in region 6, which is covered by the F_3-C_3 channel.

3. Channel weight coefficients evolution over time. The procedure described in the previous step was repeated for all 1-second segments within a specific record, resulting in the temporal evolutions of the CWCs represented in Figure 4.3. Tracking the evolution of CWCs over time allows us to monitor changes in neural activity across seizure phases, identifying scalp regions where activity is concentrated during ictal events. This provides valuable insight into seizure propagation patterns, also serving as an effective tool for the definition of a channel selection algorithm.

Analyzing the temporal evolution of the CWCs in Figure 4.3 for the five standard frequency ranges, one in each of the sub-figures, from (a) to (e), we can notice that higher frequencies capture more efficiently the rapid neuronal discharges that occur in different regions of the scalp during the seizure's phases. Indeed, as the frequency increases (see Figure 4.3d, e), two phenomena become evident: some channels tend to dominate others and trend discontinuities emerge around the 1730th second and the 2996th second, indicated as t_0 and t_1 in the plots, respectively. The instant t_0 exactly corresponds to the onset of the *ictal* phase, as annotated in the dataset, which concludes at the instant t_2, marking the beginning of a *post-ictal* phase. The discontinuity at instant t_1, instead, although not annotated in the dataset, could be representative of the onset of a pre-ictal phase. The portion of the record preceding the instant t_0 will hereinafter be referred to as the *no-seizure* phase while the segment between t_0 and t_1 will be defined as the *pre-ictal* phase [31].

Looking at the temporal trends in more details, it can be observed that the δ and θ waves, shown in Figure 4.3a and 4.3b, exhibit a relatively homogeneous distribution across channels throughout all the records. By contrast, the α band in Figure 4.3c reveals more structured spatial patterns. During the no-seizure phase, α activity is mainly concentrated in frontal regions (channels 1-4, i.e. Fp_1-F_7, Fp_1-F_3, Fp_2-F_4, Fp_2-F_8), whereas during the pre-ictal phase it shifts toward parieto-occipital regions (mainly channels 18-20, i.e. around the P_7-O_1 area). This trend is illustrated in Figure 4.3c between t_0 and t_1. During the ictal phase (between t_1 and t_2), α waves are instead

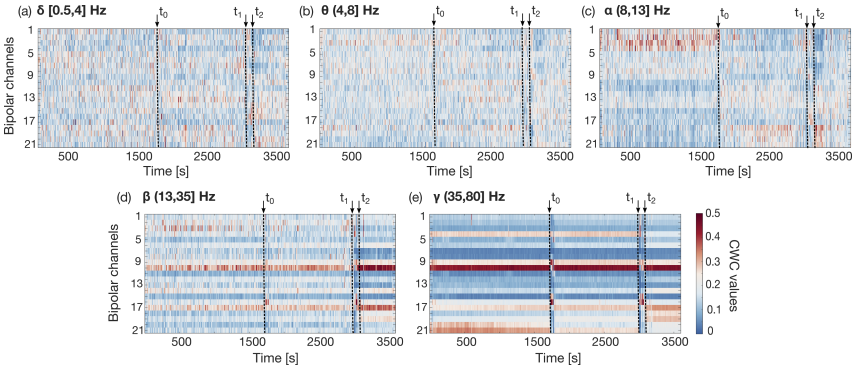


Figure 4.3: Temporal evolution of CWC decomposed into five clinical frequency ranges (a-e). x-axes represent time in seconds over 1 hour, while the y-axes show the 21 bipolar channels (defined in the text). The intensity of the CWCs is indicated by the colorbar in (e). Dotted t_1 and t_2 lines denote the start and end of the ictal event, respectively, according to the ground truth. The t_0 dotted line marks the beginning of the pre-ictal phase, characterized by trend discontinuities.

predominant in the right temporal region 16 and 17, covered by FT_9-FT_{10} and FT_{10}-T_8. Finally, in the post-ictal phase, a configuration similar to the pre-ictal tends to be reestablished.

The β waves (Figure 4.3d) and the γ waves (Figure 4.3e) reveal higher channel selectivity during different seizure phases, with γ waves localized more precisely than β. During no-seizure phases, i.e. before the instant t_0, these high-frequency activities are mainly in the left temporal region 10 (T_7-FT_9) with notable spots in the right parieto-occipital (21: P_8-O_2) and temporal-frontal regions covered by the electrode F_8, especially in the γ band. Around the $t_0 = 1730s$, marking the beginning of a pre-ictal phase, the left temporal regions 16 and 9 around T8 electrode (FT_{10}-T_8, F_8-T_8) becomes predominant, causing a trend discontinuity in Figure 4.3d and 4.3e. After this discontinuity, the γ waves tend to become dominant in the same regions that characterize the no-seizure phase.

The ictal phase again shows a strong involvement of the right temporal-frontal region (FT_{10}-T_8, F_8-T_8), also extending to the right frontal region (Fp_2-F_8, Fp_2-F_4) and right temporal-parietal region (T_8-P_8). In the post-ictal phase, after the in-

stant t_2, high-frequency activity tends to shift again toward the region 10 covered by T_7-FT_9. We can observe that γ waves are dominant on the FT_9 channel throughout the whole considered record, becoming even more pronounced in the post-ictal phase. This effect could be due to individual variations in brain anatomy and functionality, reflecting involvement in cognitive and sensorimotor functions, or latent hyper-excitability, which, while not sufficient to trigger a seizure, could be present at a subclinical level as higher baseline activity [143].

An interesting aspect to observe is that within the individual phases, the trend of CWCs is quite constant. This helps us to effectively compact the information obtained by deriving, for each phase, single values of the CWCs associated with the individual channels. Actually, each temporal phase corresponds to a spatial region that is more or less extended, involving a limited number of channels.

Consequently, each phase could be effectively represented by one channel, i.e., the one with the highest CWC. Given that our task addresses a binary classification problem between ictal and inter-ictal, we specifically focus on the subset of channels that are most representative of these two phases. To further enhance the distinction between the classes, we limit the inter-ictal phase to a reduced subclass, as will be better described in Section 5.1.2. Therefore, such a temporal analysis enables the identification of the most informative channels at each phase of interest aiming at improving the detection performance.

4. Automated selection algorithm. To define a fully automated channel selection algorithm, it was useful to compress the information into a single coefficient per channel for each phase merging information from the different bands, yielding 21 values for the inter-ictal and 21 for the ictal phase. With this aim, the CWCs representative of each phase were first normalized into the $[0, 1]$ range by using a min-max normalization within each frequency band, thus preserving information from all frequency bands, ensuring equal contribution from all of them.

These normalized frequency-dependent CWCs were then averaged across frequency bands (denoted by the set $\mathcal{B}$), yield-

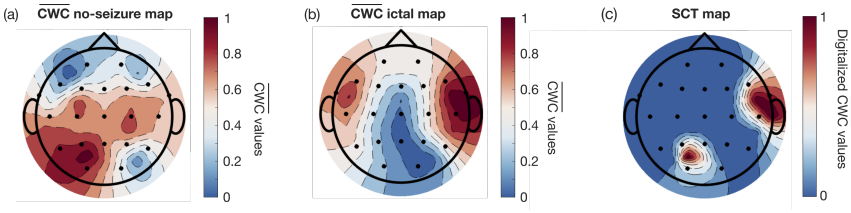


Figure 4.4: Channel selection method. EEG topographic maps showing the $\overline{\text{CWC}}$ values (normalized to the range [0,1]) for each channel, represented in an easily interpretable format. The maps were obtained from all the records of a single patient. **(a)** $\overline{\text{CWC}}$ values during the no-seizure phase. **(b)** $\overline{\text{CWC}}$ values during the ictal phase. In both maps, areas in red represent higher CWC values according to the colormap. **(c)** XOR map of the digitalized values (with a threshold at 0.9) from (a) and (b) highlighting the two most significant channels.

ing average CWCs (represented as $\overline{\text{CWC}}$) for each phase:

$$\overline{\text{CWC}}_i = \frac{1}{|\mathcal{B}|} \cdot \sum_{b \in \mathcal{B}} \text{CWC}_i^{(b)}, \text{ with } i = 1, 2, \ldots, C. \tag{4.5}$$

In our case, we consider the five standard frequency bands $\mathcal{B} = \{\delta, \theta, \alpha, \beta, \gamma\}$. Finally, since each patient dataset contains multiple records, this procedure was repeated for each of them, and the results were averaged and visualized in the EEG topographic maps of Figure 4.4a and 4.4b for no-seizure and ictal phases, respectively.

As a final step, we identified a pair of channels as the most informative in each of the two phases (ictal and inter-ictal), specifically those with the highest average weight. To maximize the difference in information content between the selected channels and minimize redundancy, we ensured that the channels were not dominant in both maps, allowing each to contribute with unique and non-overlapping information.

For instance, the map in Figure 4.4a, corresponding to the no-seizure phase, shows that the channel P_3-O_1 carries most of the information, while the map in Figure 4.4b, pertaining to the ictal phase, highlights the channel F_8-T_8.

This selection procedure can be fully automated using the following two-step algorithm: (i) digitize the average maps by

normalizing the values into a $[0, 1]$ scale, assigning a value of 1 to values above a threshold (experimentally set to 0.9) and 0 to others; and (*ii*) perform a XOR) operation on the two digitized values sets to generate the final *Selected Channels Topological (SCT) Map*, which highlights the selected channels (as shown in Figure 4.4c).

It should be noted that the selection algorithm may identify multiple representative channels for each phase if several channels exceed the $\overline{\text{CWC}}$ threshold. While this may provide insights into seizure characteristics, it also necessitates selecting an optimal channel pair for classification. To refine this selection, we introduce a stronger criterion based on inter-channel correlation. In particular, we define an average correlation coefficient ($\overline{\rho}$) in the range $[0, 1]$, based on the absolute Pearson correlation coefficient (PCC) [144]. PCC values were computed over 1-second sliding windows across frequency bands and averaged to obtain $\overline{\rho}$.

In the example discussed, the parameter $\overline{\rho}$ calculated between the two identified channels is equal to 0.32, thus indicating a low correlation degree. This coefficient helps discriminate between channel pairs when multiple options arise by selecting the least correlated pair and also serves as an interpretative tool in the experimental phase.

4.1.3.2 Classification model

Seizure detection was carried out using EpiDeNet, a lightweight deep neural network developed by Ingolfsson *et al.* [92]. The choice of this architecture was motivated by several considerations. First, although recent studies highlighted that CNNs are not always the most suitable solution for modeling multichannel EEG signals, especially when a large number of electrodes gives rise to complex and irregular spatial relationships, this limitation is much less relevant in the present framework. Indeed, the proposed classification stage operates on only two selected EEG channels, which greatly simplifies the spatial structure of the input and shifts the modeling focus toward the extraction of temporal patterns. Under these conditions, a compact CNN represents an effective trade-off between repre-

sentational capability, computational efficiency, and suitability for embedded deployment.

Second, the main objective of this work was not the development of a novel DL architecture, but rather the careful and coherent integration of a data-driven EEG channel selection strategy into a lightweight seizure detection pipeline. In this respect, adopting an existing compact and validated model made it possible to assess the contribution of the proposed channel selection method more clearly, without introducing additional variability related to architectural design choices. Finally, EpiDeNet was originally designed for resource-constrained environments and therefore features a compact structure, a limited number of trainable parameters, and low computational complexity [92]. These properties make it particularly suitable for implementation on embedded or wearable devices, which represent the target application scenario of this study.

Although the network architecture was inherited from the literature, several modifications were introduced to improve its suitability for the proposed framework. In particular, batch normalization (BN) layers were added after each convolutional block in order to stabilize the training process and facilitate faster convergence [145]. Each BN layer operates on the feature maps generated by the preceding convolution and introduces a number of parameters equal to the number of output channels of that layer. In addition, the network was adapted to accommodate the reduced input dimensionality by reshaping the spatial convolutional filter accordingly.

The resulting architecture, summarized in Table 4.1, consists of five convolutional blocks aimed at extracting spectral features through frequency-selective filters and identifying spatial patterns associated with specific frequency bands through spatial filtering. Each block includes a 2D convolutional layer with appropriate padding to preserve the input dimensions, followed by batch normalization, a ReLU activation, and a max-pooling operation. The first pooling layers progressively reduce the temporal dimension, using kernels with height equal to 1 and widths of 8, 4, and 4, respectively. The final max-pooling layer has width equal to 1 and height equal to the number of input channels to the network, denoted by M.

Table 4.1: Architecture of the proposed deep learning model.

Layer	Input	Output	Filters	Kernel	Params
Conv2D	$(1, M, 1024)$	$(4, M, 1024)$	4	1×4	20
BNorm2D	$(4, M, 1024)$	$(4, M, 1024)$			8
MaxPool	$(4, M, 1024)$	$(4, M, 128)$		1×8	
Conv2D	$(4, M, 128)$	$(16, M, 128)$	16	1×16	1040
BNorm2D	$(16, M, 128)$	$(16, M, 128)$			32
MaxPool	$(16, M, 128)$	$(16, M, 32)$		1×4	
Conv2D	$(16, M, 32)$	$(16, M, 32)$	16	1×8	2064
BNorm2D	$(16, M, 32)$	$(16, M, 32)$			32
MaxPool	$(16, M, 32)$	$(16, M, 8)$		1×4	
Conv2D	$(16, M, 8)$	$(16, M, 8)$	16	16×1	4112
BNorm2D	$(16, M, 8)$	$(16, M, 8)$			32
MaxPool	$(16, M, 8)$	$(16, 1, 8)$		$M \times 1$	
Conv2D	$(16, 1, 8)$	$(16, 1, 8)$	16	8×1	2064
BNorm2D	$(16, 1, 8)$	$(16, 1, 8)$		8×1	32
AvgPool	$(16, 1, 8)$	$(16, 1, 1)$			
Flatten	$(16, 1, 8)$	16			
Dense	16	2			34

M = number of EEG channels (2 in the *2-selected* configuration, input size 2×1024; 4 in the *4-temporal* configuration, input size 4×1024). T = number of time samples.

Specifically, M was set to 2 for the proposed two-channel configuration, corresponding to an input size of 2×1024, and to 4 for the four-channel configuration, corresponding to an input size of 4×1024. The final part of the model consists of a global average pooling layer followed by a fully connected layer with softmax activation.

The network was implemented in PyTorch (version 2.1.2) with CUDA 12.1 support. All training and testing procedures were carried out on a workstation equipped with an NVIDIA Tesla V100 DGX GPU. The source code is publicly available on GitHub: `https://github.com/MiviaLab/INBIT/tree/main/Seizure-detection`.

4.2 Full-Channel Framework

This Section presents the proposed method for modeling full-scalp EEG from a graph perspective, which is easy to use and resource-efficient, thanks to the rigorous quantification of brain connectivity changes induced by epileptic seizures. The approach compares two measures of functional connectivity (FC), Pearson's correlation coefficient (PCC) in time domain and Phase-Locking Value (PLV) in frequency domain, to identify, through statistical evaluations, the most representative measure for the seizure detection task. The selected FC measure is then evaluated to extract clinically significant spatial patterns, maintaining a compact architecture thanks to a study aimed at graph sparsification. The resulting graph is then evaluated in a framework that uses a lightweight and dynamic spatiotemporal graph network, consisting of two modules: (i) a Graph Convolutional Network (GCN) to model spatial relationships between channels and (ii) a Long Short-Term Memory (LSTM) to capture temporal dependencies. Interpretability is ensured by the integration of GNNExplainer, which highlights the most relevant channels and connections during detection, allowing clinicians to visually evaluate and validate the model's focus.

4.2.1 Introduction

In recent years, ML and DL techniques have been widely applied to the analysis of EEG signals for the detection of epileptic seizures. Traditional ML methods [63, 64] require manual feature extraction, with the risk of introducing significant human error. In contrast, DL-based approaches have demonstrated a strong ability to automatically learn meaningful representations from temporal data [70, 71, 60, 66]. Among these, CNNs have been particularly successful due to their ability to transform EEG signals into image-like representations, such as time-frequency spectra or spatiotemporal maps, making it possible to extract patterns relevant for seizure detection [60]. However, their theoretical framework assumes that the electrodes are arranged on a regular grid, thus neglecting the complex functional connections between different brain areas [55].

To overcome this limitation, there has been growing interest in recent years in Graph Neural Networks (GNNs), which allow brain connectivity to be modeled directly by representing electrodes as nodes and functional or anatomical relationships as edges. Therefore, GNNs offer a more accurate picture of the non-Euclidean nature of EEG data, integrating the local characteristics of the nodes with the overall relational structure. Numerous studies have shown performance improvements over CNN methods, highlighting how leveraging the geometric and topological properties of the EEG signal can lead to more accurate, interpretable, and generalizable models [55, 122, 46].

In this study, we aim to explore the use of lightweight GNN models for detecting epileptic seizures, with a particular focus on choosing the most suitable FC measure to represent brain states. Selecting the FC metric is no trivial task. Over the years, several options have been proposed, both in the time domain and in the frequency domain, which have specific advantages and limitations. Despite this, in several studies this choice does not receive the attention it deserves, while an accurate justification can be crucial. An informed selection, in fact, not only allows the construction of a coherent and meaningful graph-based representation of brain connections, but can also help reduce computational costs, making GNN-based models more practical and sustainable for the task at hand [116].

Once the most suitable FC measure has been identified to simplify the graph-based representation of the EEG and facilitate the use of lightweight models with reduced computational requirements, the next step is to design a lightweight GNN. In fact, a crucial aspect that has emerged from the literature concerns the possibility of translating these models into real-time, low-energy-consumption clinical applications [50], which implies that GNN-based seizure detection algorithms must be able to operate under conditions of limited computational resources.

In recent years, several studies have shown that lightweight versions of GCNs, sometimes combined with compression strategies, can drastically reduce the number of parameters while maintaining competitive performance. For example, Jia *et al.* designed an efficient GCN for predicting epileptic seizures from

scalp EEG, based on three Chebyshev convolutions, global pooling, and a compact fully connected head [57]. Similarly, Barbera *et al.* introduced a minimal GCN pipeline for dementia assessment, employing a few low-order GCN layers, node normalization, and graph pooling, followed by a small MLP classifier [123].

However, these studies do not consider the dynamic nature of EEG signals, instead treating the signal through static graphs. In contrast, Xu *et al.* proposed a spatio-temporal architecture for detecting epileptic seizures based on GCN and Bidirectional Gated Recurrent Units (BiGRU), capable of efficiently modeling both the spatial dependencies and temporal dynamics of the EEG. The idea of combining a recurrent network with a GCN is typical of spatio-temporal GNN models, applied to graphs characterized by fixed nodes and dynamic edges, as it allows spatial and temporal modules to be integrated using a relatively simple approach [112]. Starting from this study as a *baseline*, the present work aims to design a lighter ST-GNN than the reference model. In line with the first part of the thesis, this line of research is continuing in the direction of personalized medicine and wearable applications.

One final open question concerns the issue of model explainability, as GNNs remain largely black boxes. In a clinical context, it is not enough to obtain accurate predictions: it is essential to understand the reasons behind the model's decisions. To this end, the literature has introduced specific tools such as GNNExplainer, which can identify the nodes, edges, or subgraphs most responsible for classification [125]. Although such approaches have been successfully applied in other fields, their use in detecting epileptic seizures has not yet been explored.

In this perspective, a further objective of this work is to evaluate the ability of the proposed ST-GNN model to provide interpretative information. To this end, we apply GNNExplainer to the feature extraction module (i.e., the GCN component) in order to estimate the relative importance of nodes in the correct discrimination of critical events and obtain a more transparent mapping of the brain areas involved. Specifically, we propose a quantitative analysis of the model's interpretability, based on the frequency of occurrence of the nodes identified

by GNNExplainer. The results show that GNN methods confirm what emerged in the first study of this thesis, namely that some brain regions are more discriminative and informative than others. This approach not only reinforces the experimental evidence but also contributes to supporting more effective and interpretable diagnoses in real clinical contexts.

4.2.2 Related work

Li *et al.* [122] proposed a graph-generative neural network model for epileptic seizure detection based on EEG signals and brain functional connectivity. Their work utilized the Temple University Hospital Seizure Corpus (TUSZ), which includes EEG recordings from 3,047 epilepsy cases sampled at 250 Hz. The model performs seven-class classification and achieved an accuracy of 91%. In [146], the authors introduced a two-layer Linear Graph Convolutional Network (LGCN) to enhance the feature representation of raw EEG signals during both seizure and non-seizure periods. They computed the Pearson correlation matrix of EEG signals to construct the graph input for the GCN, where matrix coefficients reflect spatial relationships among EEG channels. On the CHB-MIT dataset, the model achieved 95.89% accuracy, 95.30% specificity, and 93.80% sensitivity.

Nie *et al.* [147] presented a hybrid architecture that combines Graph Convolutional Networks (GCNs) with Transformer models for epilepsy detection using EEG data. The approach showed strong performance, with an average sensitivity of 92.97 %, specificity of 94.60%, and a balanced accuracy of 93.78%.

Jibon *et al.* [148] employed a Graph Attention Network (GAT) to extract spatial features from preprocessed EEG segments and their corresponding graph structures. The GAT was integrated with a Radial Basis Function Neural Network (RBFN), which complements the GAT by leveraging linear optimization techniques to achieve a globally optimal adjustment of the network's weights. By assigning attention scores to neighboring nodes, the GAT effectively highlights the most significant inter-channel relationships within the EEG data. Operating within a subject-specific framework on the CHB-MIT

dataset, the proposed model achieved sensitivity, specificity, and balanced accuracy scores of 94.74%, 97.89%, and 96.32%, respectively.

A major limitation of most of these approaches is that they rely on static graphs, which cannot adapt to the dynamic nature of brain connectivity during different brain states. In contrast, Spatio-Temporal Graph Neural Networks (STGNNs) simultaneously capture spatial structures via Graph Neural Networks (GNNs) and model temporal dependencies through Recurrent Neural Networks (RNNs) [149]. This integrated modeling framework offers a powerful means for analyzing the complex spatiotemporal dynamics of EEG data. However, only a few studies have addressed dynamic structure learning, which is essential for extracting richer and more adaptive features from EEG signals [122]. He *et al.* [150] proposed a hybrid framework combining a GAT with a Bi-directional Long Short-Term Memory (BiLSTM) network for seizure detection. GAT captures spatial dependencies by modeling EEG channel topology and highlighting crucial inter-channel connections, while BiLSTM models temporal patterns in the EEG signals. This spatiotemporal approach achieved a balanced accuracy of 96% on the CHB-MIT dataset. In 2024, Jibon *et al.* [151] proposed a hybrid framework combining a Sequential Graph Convolutional Network (Seq-GCN) and a Deep Recurrent Neural Network (DeepRNN). The Seq-GCN captures spatial dependencies by modeling the relationships between EEG electrodes, while the DeepRNN captures the temporal dynamics of brain activity. Using 16 EEG channels and adopting a subject-specific approach across 20 patients, their model achieved a balanced accuracy of 96.54% for seizure detection.

Xu *et al.* [112] conducted a subject-dependent seizure detection study to examine computational complexity variations between ictal and inter-ictal phases. They developed several GCN-based models and evaluated them on the CHB-MIT dataset. A simple three-layer GCN achieved 95.23% accuracy, while hybrid models that combined GCN with LSTM and Gated Recurrent Unit (GRU) achieved 85.58% and 97.35% accuracy, respectively. Awais *et al.* [152] further proposed a five-layer GCN model combined with LSTM, achieving an average accuracy of

96.54% and sensitivity of 93.89% across 23 CHB-MIT patients.

4.2.3 Method

Figure 4.5 provides an overview of the experimental setup proposed in this study. The pipeline consists of several integrated stages. First, the acquired signals undergo a pre-processing procedure aimed at removing noise and artifacts, in order to obtain clean and standardized data. Then, an initial phase focuses on selecting the most suitable FC metric for seizure detection. By calculating the relationships between pairs of EEG channels and extracting the corresponding connectivity matrices, a comparative analysis is conducted to identify the most discriminating measurement following the procedure described in Section 4.2.3.1.

Once this phase is complete, the representative and appropriately distributed graphs are ready to be processed by a GNN-based model. The proposed architecture integrates a GCN-based feature extraction phase with a classification phase implemented using LSTM. Thanks to the relatively small number of trainable parameters, the model guarantees high computational efficiency while maintaining strong representational capacity. The final output of the pipeline is a binary classification that distinguishes between 'ictal' and 'non-ictal' states. In this work, the label 'non-ictal' refers to all windows assigned to class 0 that do not belong to a ground truth ictal event. This choice will be clarified in Section 4.2.3.1, at step point 1, where we explain how the issue of dataset imbalance was addressed.

Additionally, it is worth noting that the method is evaluated in a specific scenario for each patient, with the aim of developing a personalized model for each of them. Furthermore, in order to improve the decision-making phase, a GNN Explainer module has been added that allows the interpretation of the proposed model, highlighting in particular the nodes most relevant to the prediction of the output and thus providing an important interpretability tool useful in both clinical and research settings.

In this perspective, this section presents the details of the proposed methodology. First, we outline all the steps of the

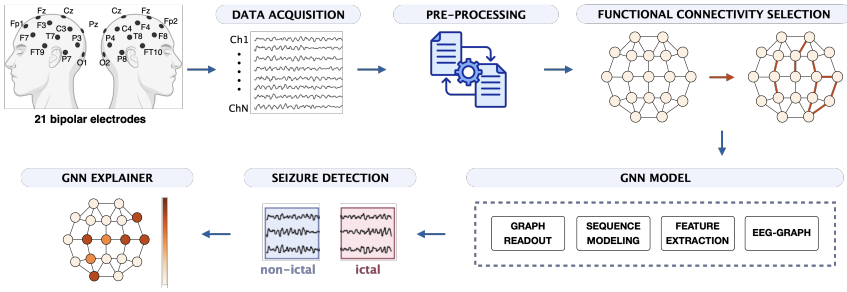


Figure 4.5: Schematic view of the proposed GNN-based pipeline for the detection of epileptic seizures.

procedure implemented to compare the considered connectivity measures in Section 4.2.3.1. Next, we describe the graph computation in Section 4.2.3.2 and the GNN-based architecture in Section 4.2.3.3. Finally, the implementation of the GNN-Explainer tool on the pre-trained graph-based models is presented 4.2.3.4.

4.2.3.1 Connectivity evaluation

The first aspect considered in this study, which is also the most relevant contribution, is the evaluation of the FC measure best suited to represent graphs able to distinguish between ictal and non-ictal states. In this perspective, we propose a comparison between two metrics, one in the time domain and one in the frequency domain, respectively Pearson correlation coefficient (PCC) and phase lag index (PLV). Both measures are widely used in EEG applications, but their effectiveness varies significantly depending on the task at hand. For example, in the context of EEG synchronization, measures based solely on phase delay produce cleaner and often more sparse graphs, at the same threshold, compared to classic correlation indices [118, 54]. Conversely, in tasks such as emotion recognition [119] or motor imagery [120], PCC has repeatedly been shown to be the most effective metric, sometimes further improved by metric fusion strategies [121]. However, in the specific case of detecting epileptic seizures there is still no systematic study directly comparing the performance of the two measures.

To address this difficulty, we propose a comparative study

based on a statistical and data-driven approach to quantify differences in FC between ictal and non-ictal phases. By combining the estimation of FC matrices, the descriptive evaluation of distributions, and non-parametric statistical analysis, the approach makes it possible non only to identify the most discriminating FC measure but also to optimize the sparsification threshold. Consequently, this reduces edge redundancy in the EEG graph-based representation and lowers the computational costs of the applied GNN.

The procedure adopted to compare FC measures can be summarized in the following steps.

1. Data pre-processing. All EEG record was filtered with a 0.5 Hz high-pass filter and an 80 Hz low-pass filter, both second-order Butterworth filters. The high-pass filter removes direct current (DC) trends, while the low-pass filter eliminates physiological artifacts that could interfere with seizure recognition. To ensure cleaner signals, a notch filter was also applied as part of the network interference removal procedure. Specifically, the power line noise was removed using a zero-phase 4th order Butterworth stop-band filter. The filter was centered at 50 Hz for the Siena Scalp dataset [80] and to 60 Hz for the CHB-MIT dataset [79], in order to take into account the different frequencies of the electrical network in the two different recording contexts. Additionally, amplitudes greater than two standard deviations from the average intensity of the signals were removed to eliminate high-intensity artifacts such as muscle activity or electrode detachment. Finally, the EEG signals were also normalized to ensure zero mean and unit variance for all recordings.

It should be noted that the pre-processing adopted in this section differs from the one described previously because the two studies had different objectives and, consequently, different pre-processing requirements. In the first study, pre-processing was required only for the channel selection stage, whereas the lightweight CNN received the raw EEG signal as input. In this case, the task of extracting relevant temporal and spectral patterns was delegated to the convolutional kernels of the network. Moreover, since the channel selection procedure was based on

a PCA-based method that is relatively robust to noise, only basic pre-processing steps were applied, namely notch filtering and band-pass filtering. By contrast, in the connectivity-based study, an additional pre-processing step was necessary because connectivity measures are known to be strongly affected by noisy EEG components. For this reason, artifact removal was introduced to reduce the influence of disturbances such as muscle activity or electrode detachment, using the amplitude-reduction criterion based on 2 times the standard deviation. Concerning the frequency range of interest, this remained unchanged with respect to the first study, and the signal was analyzed in the 0.5-80 Hz band.

From all one-hour recordings, non-overlapping one-second windows were extracted and used as input for our GNN model. This is consistent with previous GNN applications to EEG data [153]. Epileptic activity in each recording was annotated by ground truth files, which report the start and end of all seizure events. In this way, any one-second window overlapping this interval was considered a positive instance of the "ictal" class (class 1), while windows with no epileptic activity were labeled "non-ictal" (class 0). This approach is consistent with the channel selection study presented in the previous chapters of this thesis. To mitigate class imbalance while fully exploiting all available ictal events, we fixed the number of non-ictal windows to an approximately 4:1 ratio relative to ictal windows [112]. For instance, in patient chb01 from the CHB-MIT dataset [79], 442 ictal and 1,680 non-ictal windows were extracted, with non-ictal windows randomly sampled from inter-ictal and pre-ictal periods.

2. Estimation of FC matrices. One-second EEG windows comprising all electrodes were used to estimate connectivity matrices. Since the graphs were constructed considering $n = 20$ nodes, corresponding to the number of electrodes from which the data were collected, the matrices have a size of 20×20. Figure 4.6 shows the connectivity metrics computed on one-second windows for both FC metrics under analysis, i.e PLV (on the top) and PCC (on the bottom). The degree of connectivity is illustrated between each pair of bipolar channels in the

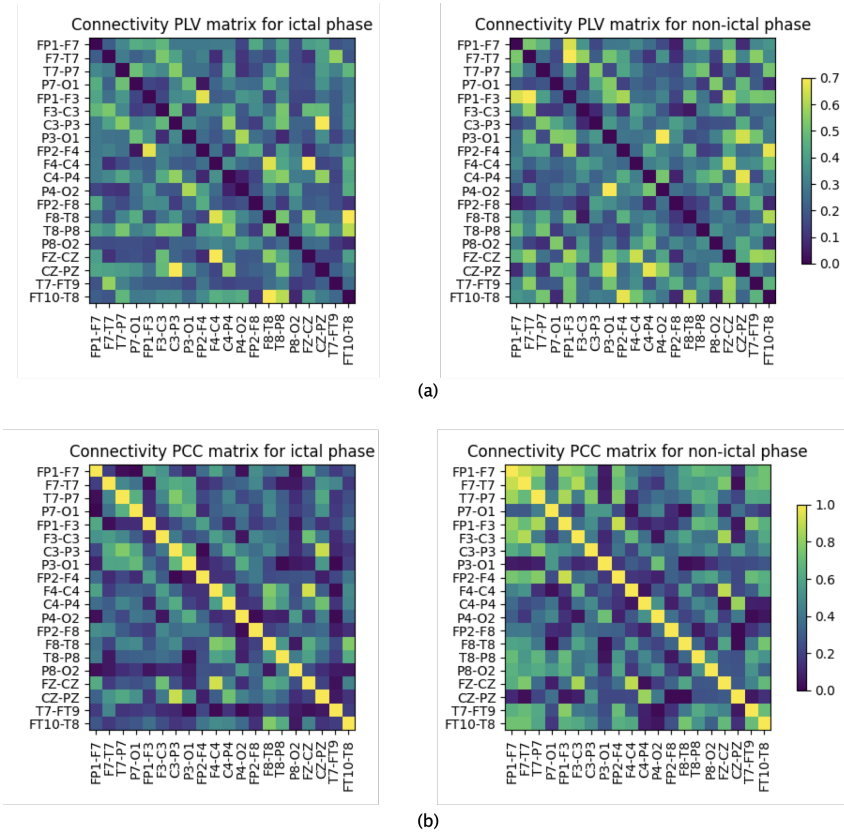


Figure 4.6: FC matrices between ictal (left) and non-ictal (right) for (a) PLV measure and (b) PCC measure. The x and y axes represent bipolar channels, while each pixel in the matrices indicates FC value between each pair of channels.

two phases, i.e. ictal (on the left) and non-ictal (on the right). Blue indicates lower connectivity, while yellow indicates higher connectivity.

All matrices are symmetric, with values ranging from 0 (no connectivity) to 1 (connectivity). The diagonal elements represent self-connections. In the case of the PCC, the diagonal elements take the value 1, as they correspond to the perfect correlation of each channel with itself. On the contrary, for the PLV, these values are set to 0, since the phase synchronization measure is not defined for a signal with respect to itself. As these elements do not provide relevant information for the

analysis, they are not considered [112, 154, 149]. Moreover, to facilitate comparison, the absolute values of the PCC were considered, and not the interval $[-1\ ,\ 1]$, since the aim is not to distinguish between positive and negative correlations, but only to verify the presence or absence of a connection [46].

3. Evaluation of FC distributions. To compare the two FCs, two different methodologies were adopted. First, connectivity matrices (20×20) were computed on all one-second windows relating to a single patient, thus obtaining averaged FC matrices. Next, each matrix was transformed into vectors following the procedure illustrated in Figure 4.7.

This procedure involves considering only the upper (or lower) triangle of the matrices, excluding the diagonal, in order to retain only distinct connections and avoid redundancies. The selected values were then linearized into vectors, $v \in R^m$ with $m = n(n-1)/2$ elements (for $n = 20$, $m = 190$), which describes the distribution of connections. This representation allowed a direct comparison of the connectivity distributions between the two classes (ictal and non-ictal), using boxplots and statistical tests, as already proposed in other contexts in the literature [155, 156].

In addition to this quantitative assessment of FC distributions, a second method of analysis was adopted. This method is based on the distributions of connection strength for each of the two classes. To this end, starting from the averaged connectivity matrices, the relative probability of connection strength was estimated using normalized histograms. This approach was repeated for each patient and made it possible to describe how FC values were distributed and how frequently they appear in the two conditions.

4. Statistical analysis. To validate both methods, we used the Mann-Whitney U test (or Wilcoxon-Mann-Whitney test), which is appropriate for independent groups and does not rely on assumptions of normality or equality of variances [158]. The test was applied between the two phases (ictal vs. non-ictal) and operates on ranks (not raw values), making it robust even with small or unbalanced samples. In our case, since the 'ictal'

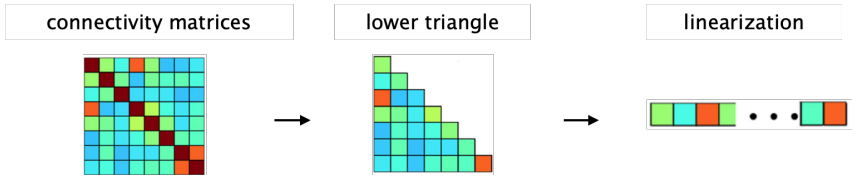


Figure 4.7: FC matrix linearization procedure [157]. Linearization is useful for both statistical analysis and comparison between distributions.

and 'non-ictal' windows are not paired $1:1$ but follow an approximate $1:4$ ratio, the assumption of independence between groups holds, and the Mann-Whitney test emerges as the appropriate choice over the Wilcoxon test, which is designed for paired data [158].

Since connectivity matrices involve multiple comparisons, we additionally considered the Benjamini-Hochberg false discovery rate [159] to correct for multiple comparison in the network analysis and the Cliff's delta δ to assess the effect size for pairwise interactions [160]. More precisely, Cliff's δ is used to quantify how much the values of one distribution are systematically greater or less than those of another distribution. Formally, it is defined as:

$$\delta = \frac{\#(x_1 > x_2) - \#(x_1 < x_2)}{n_1 n_2}$$

where x_1 and x_2 represent the values of the two distributions being compared, n_1 and n_2 are the respective sample sizes, and the symbol $\#$ indicates the count of pairs of observations that satisfy the condition [160].

Both FDR and Cliff's δ can be applied without assuming a normal distribution of data and are commonly used to analyze FC differences between groups in non-parametric contexts. All statistical analyses were performed in Python using SciPy version 1.10.1.

The significance level denoted as the p-value was set to 0.05. The *p-value* indicates the probability, under the null hypothesis (the two vectors arise from the same distribution), of observing a test statistic at least as extreme as the one measured [158]. In our case, smaller p-values indicate stronger evidence

of differences between the connectivity matrices in the ictal and non-ictal conditions.

The statistical test is applied according to the procedure described in the work by Xu *et al.* [157] and shown in Figure 4.7. For each one-second window, we estimated the connectivity matrix (20×20) and transformed it into the corresponding linear vector, following the procedure described at step point 2 of the Section 4.2.3.1. Using the same order for all windows ensures a one-to-one correspondence between each position in the vector and a specific pair of channels (i, j) of the matrix, thus allowing also consistent multiple-comparison testing.

We apply the same procedure to both PLV-based and PCC-based matrices and, to ensure a fair comparison between them, we do not apply any thresholding to sparsify the matrices prior to testing.

4.2.3.2 Graph computation

Once the FC metrics for each channel pair were computed, the connectivity matrix was used to define the edges of the graph. More precisely, we considered a single graph for each one-second window analyzed in this study (as shown in Figure 4.6). To properly represent the edges, we first derived the adjacency matrix **A** from the FC matrix, where each element a_{ij} is defined as:

$$a_{ij} = \begin{cases} 1, & \text{if } FC_{ij} \geq \tau \\ 0, & \text{otherwise} \end{cases} \tag{4.6}$$

where τ is a threshold and FC_{ij} denotes the connectivity relationship between nodes i and j [161]. As for the characteristics of the nodes, these were represented by raw EEG signals. Considering a time window of one second, each node is therefore described by a vector of 256 samples.

Then, we defined the threshold value needed to "sparsify" the graph by eliminating the edges that represent weak connections [153]. This operation is very useful for reducing the complexity of the graph and, consequently, the number of model parameters, simplifying the overall architecture. However, the

threshold represents a critical hyperparameter in the analysis of graph-based networks. Indeed, the literature frequently describes an approach in which the GNN network is trained on graphs constructed with different threshold values, ultimately selecting the one that guarantees the best performance [161, 108].

In this study, the threshold was optimized by evaluating its ability to discriminate between the ictal and non-ictal phases. To this end, starting from the normalized histograms described at step point 3 of the Section 4.2.3.1, the threshold value was selected by analyzing the distributions of connection strengths under the two conditions and identifying the interval that yielded the maximum separation between them. The rationale for focusing on the point of maximum separation lies in the observation that, during the ictal phase, connectivity increases not only in terms of strength but also in the number of connections, as a consequence of the widespread synchronization occurring at the brain level. As a result, the ictal graphs become denser, whereas the non-ictal graphs appear sparser, thus providing a more structurally discriminative representation for the GNN in seizure detection [161].

It is also worth noting that the threshold value was optimized individually for each patient, as inter-individual variability in the EEG signal prevents the definition of a fixed threshold that is valid for all patients

4.2.3.3 Graph-based model

Once the FC analysis confirmed that PLV is more suitable for representing connectivity in distinguishing between ictal and non-ictal phases, and after defining the threshold to obtain a sparse graph representation, we proceeded to evaluate the patient-specific classification performance. The proposed ST-GCN model starts from a raw EEG signal and integrates a feature extraction phase based on GCN with a classification phase implemented using LSTM. The training process was organized in two distinct phases. In the first phase, a network was trained with the aim of obtaining an independent feature extractor. In the second phase, the pre-trained extractor was

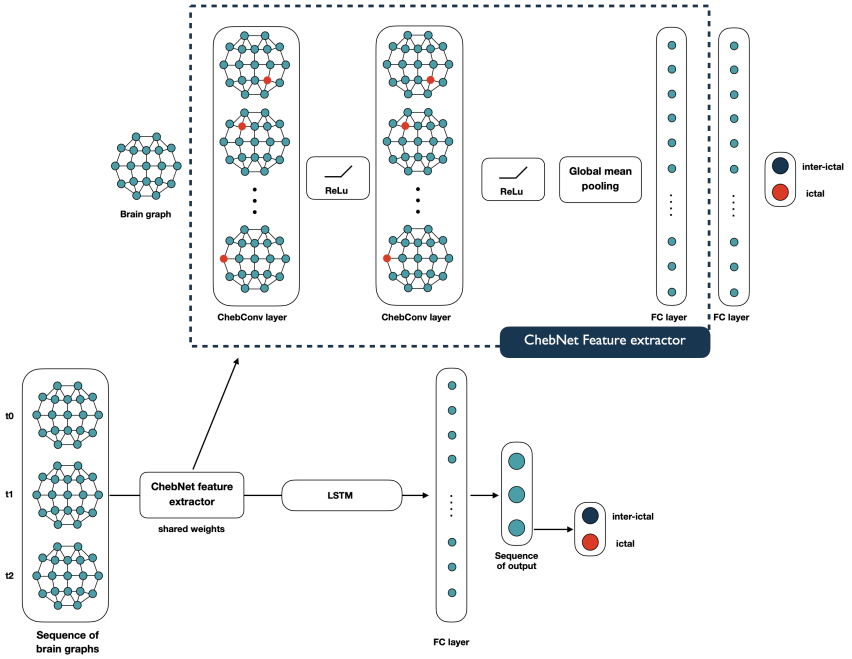


Figure 4.8: Proposed model for graph-based EEG signal analysis. (Top) ChebNet feature extractor, trained to encode the spatial properties of EEG graphs. (Bottom) Full ST-GCN architecture, in which the pre-trained ChebNet extractor is combined with an LSTM to model temporal dependencies between sequences of three consecutive graphs.

embedded in a larger architecture, with the addition of a recurrent module designed to model temporal dependencies.

Feature extraction is performed using a Chebyshev GCN (Cheb-GCN), shown at the top of Figure 4.8, applied simultaneously to all EEG channels. Specifically, two layers of ChebConv with 256 filters and order K=1 are used. Each convolution is followed by a Batch Normalization operation and a ReLU activation function, in order to stabilize learning and introduce non-linearity. To prevent overfitting, a dropout mechanism with a probability of 0.3 has also been included. At the end of the convolutional phase, global mean pooling is applied, which aggregates the information from the nodes into a single vector representing the graph. This representation is then processed by two fully connected layers. The first reduces the di-

mensionality to 64 units, applying a ReLU activation function and dropout, while the second projects into the 2-dimensional decision space (ictal vs. non-ictal) for classification.

The network described represents the first training phase and constitutes a robust and effective extractor, capable of encoding the topological and functional properties of EEG graphs. This extractor provides an output that can be used both for direct classification and, as described in the second phase, as input for the LSTM, which is responsible for modeling the temporal dependencies between graph sequences.

In the second phase (bottom of Figure 4.8), the pre-trained ChebNet extractor was reused within a more complex architecture, featuring an LSTM designed to model the temporal evolution of graphs. In this scenario, the model processes sequences of three consecutive graphs, corresponding to three-second time windows, in order to capture the dynamic evolution of the EEG signal. Each graph is processed by the ChebNet feature extractor in frozen mode (without updating the weights), producing a 64-dimensional feature vector. The sequence of the three vectors thus obtained constitutes the input for the LSTM, equipped with a hidden state equal to 32, which is responsible for modeling the temporal dependencies between the graphs. The output of the LSTM is finally projected, through a fully connected layer, into the decision space, returning a binary classification between ictal and non-ictal states.

4.2.3.4 GNN-explainer

Ying *et al.* introduced GNNExplainer [125], which defines the explanation as a compact subgraph and a subset of features that are most informative for the model's prediction. In this study, we used it to extract a frequency map of the nodes most involved in seizure detection.

Starting from the trained Chebyshev GCN, for each EEG window we applied GNNExplainer in single-instance mode at the graph level, obtaining an importance mask on the edges and extracting an explanatory subgraph G_S aimed at maximizing the information between the explanation and the model output.

Then, to move from an edge-centric explanation to a node-centric measure, we assigned each node v a score by aggregating the weights of the incident edges:

$$s(v) = \sum_{(v,u)\in E} m_e(v,u)$$

where v and u represent two nodes, m_e is the importance weight of the edge provided by GNNExplainer, $s(v)$ is the score of the single node v. Next, for each instance, we sorted the nodes by score $s(v)$ and selected the top-k most relevant ones (in our study, k = 3). Repeating the process on all ictal instances, we obtained the frequency of occurrence in the top-k for each node as:

$$f(v) = \frac{\#\{t \in \Omega : \, v \in \text{top-}k_t\}}{|\Omega|}$$

where Ω is the set of explained instances and $top - k_t$ indicates the k nodes with the highest score for instance t. Once $f(v)$ was estimated for all nodes, we represented the map of the most influential nodes in the Chebyshev GCN model training process by constructing a reference graph that included all channels in the bipolar assembly and represented each node with a color proportional to the normalized frequency $f(v)$. Note that only nodes with an occurrence > 0 were shown alongside the absolute count. This visualization made it clearly visible which channels appeared most often among the top k explanations from GNNExplainer in the analyzed ictal windows.

4.3 Evaluation metrics

This section presents the metrics used to evaluate the seizure detection performance of both models, i.e. the CNN-based model of the first study and the GNN-based model of the second study.

In details, we consider as true positives (TP) the ictal samples correctly classified by the model and as false negatives (FN) the ictal samples that are not recognized by the model. Under these definitions, we compute the sensitivity (SN) as the

ratio between the number of TP samples and the total number of actual ictal samples:

$$\mathrm{SN} = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FN}}, \tag{4.7}$$

Similarly, we define as true negatives (TN) the non-ictal samples correctly classified by the model and as false positives (FP) the non-ictal samples misclassified as ictal. In this way, we compute the specificity (SP), which measures the proportion of true non-ictal samples correctly identified among all actual non-ictal instances:

$$\mathrm{SP} = \frac{\mathrm{TN}}{\mathrm{TN} + \mathrm{FP}}, \tag{4.8}$$

In addition to SN and SP, we also consider the overall accuracy (Acc), defined as:

$$\mathrm{Acc} = \frac{\mathrm{TP} + \mathrm{TN}}{\mathrm{TP} + \mathrm{TN} + \mathrm{FP} + \mathrm{FN}}, \tag{4.9}$$

and the balanced Accuracy (bAcc), which better accounts for the imbalanced nature of the dataset by computing the arithmetic mean of the sensitivity and specificity values:

$$\mathrm{bAcc} = \frac{1}{2}(\mathrm{SN} + \mathrm{SP}). \tag{4.10}$$

For the sake of clarity, it is important to specify how the above metrics are used in the two studies. In the CNN-based channel selection framework, since the output of the binary classifier is evaluated on each 1-hour test through a post-processing phase (further details in Section 5.1.1), the metrics are computed at both the *segment level* and the *event level*. The segment level refers to individual segments, each defined as a buffer of three consecutive four-second sliding windows, with a two-second overlap, classified as either ictal or non-ictal [134].

This post-processing strategy is more restrictive than the majority-voting method applied over three consecutive windows adopted in previous studies [92] and helps to mitigate the effect of FP. The event level, instead, considers entire ictal events, where an event corresponds to the whole ictal phase

detected by the model, i.e. a sequence of contiguous ictal segments, and is used to further assess the usability of the system.

In this setting, all the metrics defined in equations 4.8–4.7–4.9–4.10 are computed at the segment level, while additional performance metrics are defined to evaluate performance at the event level. First, to gain a clearer understanding of the system's ability to detect ictal events, we consider the proportion of ictal events that are correctly identified with respect to the total number of ictal events. Moreover, we also analyze the detection latency, defined as the time difference between the annotated ictal onset and the model's detection instant, hereinafter referred to as *delay*, expressed in seconds. Finally, since a system that generates frequent false alarms is likely to lose user trust, it is crucial to quantify how false positives affect reliability. To this end, we define the false positive rate as the average number of false positive events detected per hour (FP/h) [92].

In the GNN-based full-channel framework, the output of the binary classifier is evaluated on one-second windows of the continuous EEG stream. The model includes a recurrent module, i.e. a Long Short-Term Memory (LSTM) module, that operates on sequences of three consecutive one-second windows with a one-second stride, thereby capturing very short-term temporal dynamics. For this reason, buffer-based post-processing is not required, and performance is evaluated only at the window level, without the need to define additional metrics at the event level. In this setting, all the metrics defined in equations 4.8–4.7–4.9–4.10 are computed directly on the model outputs, i.e., on individual one-second windows.

In order to assess the generalization capability of the models by evaluating their performance on unseen data, a cross-validation strategy is adopted in both frameworks. In the CNN-based approach, a leave-one-out cross-validation (LOO-CV) scheme is employed, which can be seen as a special case of k-fold cross-validation where the number of folds is equal to the number of elements in the dataset [162]. In particular, two configurations are considered: leave-one-record-out (LORO), where each EEG recording is used once as the test set while the remaining recordings are used as the training set,

and leave-one-subject-out (LOSO), where at each iteration the data of one patient are held out as the test set while the data of all remaining patients are used for training. In both cases, the process is repeated until every recording (LORO) or every subject (LOSO) has been used as the test set, resulting in a comprehensive evaluation across all patients [162]. In the GNN-based approach, a classical k-fold cross-validation is applied with $k = 5$. This choice offers a compromise between statistical reliability and computational cost, ensuring that each fold contains a sufficient number of graphs for stable training. In this case, the LORO procedure is not considered, both to maintain consistency with other works in the literature that adopt a traditional k-fold scheme [112] and because it would result in a larger number of folds, requiring the graph structure to be rebuilt at each iteration, which is computationally expensive.

Chapter 5

Validation of reduced-channel approach

This Chapter details the complete experimental protocol and results obtained for the first framework of the thesis (see Section 5.1), aimed at selecting the most informative EEG channels for a compact model suitable for contexts with limited computational resources. Specifically, Section 5.1 presents the results obtained with the *optimized dual channel* selection.

The following Sections introduce two follow-up studies, developed on the basis of evaluations carried out on the dual-channel results. The first, described in Section 5.2, validates the proposed channel-optimization strategy by assessing the tangible contribution of each of the two selected channels to seizure detection and demonstrating that choosing an uninformative channel yields insufficient performance. The second, described in Section 5.3, explores the system's flexibility in moving away from a patient-specific approach to adapt to a broader range of patients. Extension of the fully personalized framework aims at assessing its feasibility for designing a wearable device usable by multiple patients rather than customized for each individual.

5.1 Optimized dual channel analysis

The reduced-channel framework presented in Section 4.1 provides an efficient and lightweight deep learning approach for

designing a new generation of wearable EEG systems, capable of integrating personalization and advanced data processing in a low-complexity architecture. The main idea is to select the most informative channels for each patient and use them in a compact model, suitable for contexts with limited computational resources.

We perform the seizure detection experiments on the CHB-MIT public EEG dataset [79] and additionally validate the method also on the Siena Scalp dataset [80]. Both dataset specifications can be found in Section 3.2. The data include, for each patient, records acquired during their normal state as well as records obtained during seizure events. The latter are annotated with the start and end of ictal events, providing the ground truth for a classification problem with two classes: ictal, representing the period of the ictal event, and inter-ictal, representing the period between the end of one ictal phase and the onset of the next. Notably, the datasets contain data from patients with both generalized and focal seizures, even if the specific seizure type is not specified.

5.1.1 Implementation details

Within the framework of personalized medicine, this study aims to identify the optimal wearable device configuration tailored to each individual patient. This involves not only determining the most effective electrode placement but also selecting a model specifically trained to the patient's unique characteristics. To accomplish this, a leave-one-record-out (LORO) cross-validation scheme was adopted. For each patient, the model was trained on all available records except one, which was reserved for testing (see Figure 5.1 for a schematic representation of the procedure).

Given the considerable intra-patient variability, this process was repeated for each record, so that every available recording was used once as the test set. In this way, the model was trained and evaluated across the full spectrum of patient-specific EEG patterns, while also ensuring a strict record-wise separation between training and test data. As a result, the risk of information leakage due to correlations among segments ex-

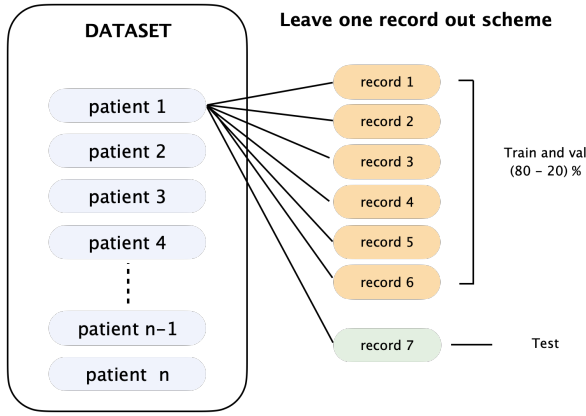


Figure 5.1: Schematic representation of the leave-one-record-out (LORO) cross-validation procedure. At each iteration, one record is excluded from the dataset and used as test sample, while all the remaining records are used to train the model. The training set is further split into training (80%) and validation (20%) subsets. The process is repeated until each record has been used once as test.

tracted from the same recording was reduced. In addition, both the normalization parameters and the channel selection procedure were computed exclusively on the training set, preventing any information from the test record from influencing model development.

We trained a specific model for each patient at most epochs and the best models were saved according to the loss registered on the validation set. In order to avoid overfitting on the training data, we stop the training if the performance on the validation set does not improve for ten consecutive epochs. The validation set was created by randomly splitting the training data in an 8:2 ratio, where 80% of the samples were used for training and the remaining 20% for validation. We apply the cross-entropy loss function, with the Adam optimizer characterized by $\beta_1 = 0.9$, $\beta_2 = 0.999$ and a learning rate of 10^{-3}.

Performance was evaluated for each test record using the post-processing procedure described in the following paragraph, and the results are reported according to the evaluation metrics as the mean and standard deviation across all records avail-

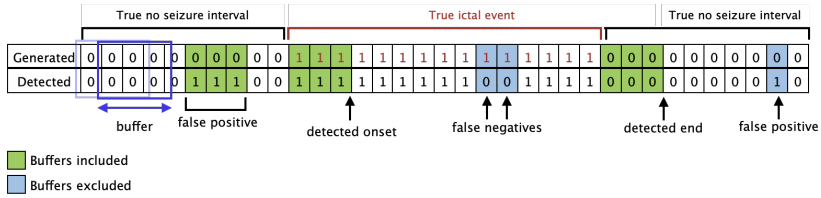


Figure 5.2: Post-processing procedure for detecting epileptic seizures based on consecutive four-second segments with a two-second overlap, using a three-window buffer rule to define ictal and no-seizure phases.

able for each patient, including those without ictal annotations, in order to assess the robustness of the system against false alarms. In this setting, temporal causality was preserved during inference, since each 4-s EEG window was processed only after its acquisition, yielding one prediction per window without relying on future samples, and the subsequent post-processing step was based exclusively on the current and preceding windows.

Post-processing.

The performance achieved on each 1-hour test record was evaluated through a post-processing phase by considering a buffer of three consecutive four-seconds windows with two-seconds sliding approach. The ictal onset is identified when three consecutive windows are classified as ictal following a long stretch of no seizure. The end of the event is marked when three consecutive windows are classified as no seizure following a stretch of ictal. The latency value is estimated based on the time difference between the annotated ictal onset and the start of the epileptic event detected by the model.

Figure 5.2 illustrates this procedure through an example comparing the ground truth (top row) and the model predictions (bottom row) over time. Each box represents a four-second signal window, extracted with a two-second overlap. A value of '*1*' indicates an ictal classification, whereas a value of '*0*' indicates a no-seizure classification. The violet rectangle highlights the buffer of three consecutive windows sliding across the sequence. Green boxes indicate buffers of three win-

dows included in the performance metrics, while blue boxes represent buffers excluded from the evaluation since they do not contain all values equal to '*0*' or all equal to '*1*'.

In the example shown, whenever three consecutive windows in the ground truth were all labeled as '*0*' or all labeled as '*1*', the corresponding interval was considered, respectively, as a no-seizure or ictal phase (both shown in green) and was therefore included in the evaluation. If, in a ground-truth no-seizure interval, the model classified all three windows as ictal, a false positive was counted. Conversely, if, in a ground-truth ictal interval, the model classified all three windows as no seizure, a false negative was recorded.

The proposed post-processing strategy was essential to obtain more robust detection decisions. Although the overlap between consecutive windows inherently introduced statistical dependence among adjacent samples, its effect on the final evaluation was mitigated by aggregating temporally correlated predictions rather than treating each window as an independent observation.

5.1.2 Results

The experiments were carried out using a patient-specific approach to assess the value of a personalized wearable device. Performance was first assessed separately for each patient and then summarized across the entire cohort. The proposed *two-selected channels* configuration was compared with the *four-temporal channels* setup not only because the latter has been proposed in the literature for reduced-electrode wearable scenarios [87, 92, 131, 163], often for practical design reasons rather than on the basis of a clearly stated methodological rationale, but also because temporal lobe epilepsy is recognized as the most common form of focal epilepsy in adults [164].

Channels optimization

As a first step, we implemented the proposed channel-selection method to identify, for each patient, the optimal scalp region for electrode placement according to the 10-20 system, thereby personalizing the wearable EEG device. To motivate and de-

Table 5.1: Selected channel pairs with their corresponding average correlation $\overline{\rho}$ for the CHB-MIT dataset. Channel A corresponds to the inter-ictal phase and Channel B to the ictal phase, respectively.

Patient	Channel A	Channel B	$\overline{\rho}$	Patient	Channel A	Channel B	$\overline{\rho}$
chb01	P_3-O_1	F_8-T_8	0.32	chb13	P_4-O_2	Fp_1-F_7	0.29
chb02	T_7-P_7	T_8-P_8	0.32	chb14	T_8-P_8	Fp_1-F_7	0.24
chb03	P_8-O_2	T_7-P_7	0.29	chb15	Cz-Pz	T_7-P_7	0.30
chb04	Cz-Pz	T_8-P_8	0.33	chb16	P_7-O_1	C_3-P_3	0.25
chb05	Cz-Pz	P_8-O_2	0.28	chb17	C_4-P_4	P_4-O_2	0.31
chb06	F_8-T_8	T_8-P_8	0.47	chb18	C_4-P_4	T_8-P_8	0.32
chb07	Cz-Pz	P_8-O_2	0.32	chb19	Cz-Pz	T_8-FT_{10}	0.26
chb08	P_8-O_2	F_4-C_4	0.27	chb20	P_8-O_2	T_7-FT_9	0.21
chb09	P_7-O_1	T_8-P_8	0.29	chb21	T_8-P_8	C_3-P_3	0.43
chb10	T_8-P_8	T_7-P_7	0.29	chb22	F_4-C_4	T_7-P_7	0.27
chb11	Fz-Cz	T_7-P_7	0.29	chb23	P_8-O_2	T_7-P_7	0.23
chb12	F_3-C_3	F_4-C_4	0.23	chb24	C_4-P_4	F_8-T_8	0.31

scribe the phases of the algorithm detailed in Section 4.1.3.1, we illustrated the procedure using data from patient chb01, but similar temporal patterns are also observed across all patients in the CHB-MIT dataset (see Figure 5.3).

Therefore, the same procedure was implemented for all the patients, and the results are shown in Table 5.1 for CHB-MIT. Bipolar channels, corresponding to specific scalp regions, identified during the inter-ictal and ictal phases are specified as Channel A and Channel B, respectively. The $\overline{\rho}$ column reports the correlation values evaluated as described at the point 4 of the Section 4.1.3.1.

It is important to remark that during the implementation of the selection algorithm, careful consideration was given to the fact that the selected channels are used as input to a classifier trained under a LORO cross-validation scheme. Specifically, it was verified that when applying the algorithm to a subset of data obtained by excluding one record from the total num-

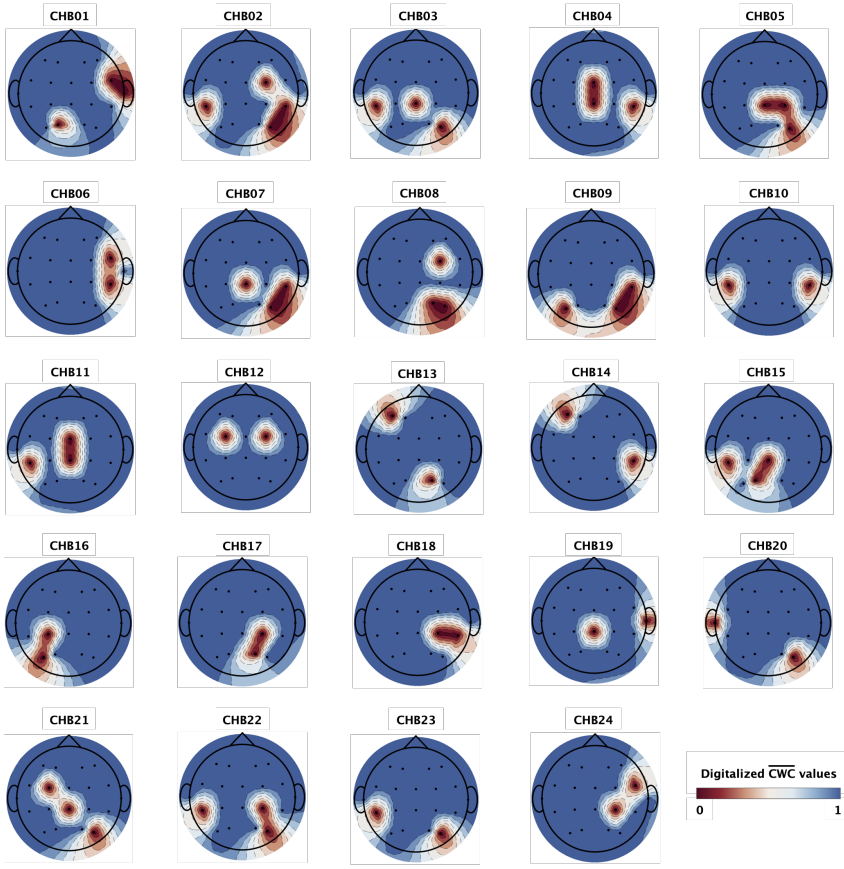


Figure 5.3: Maps of the digitalized $\overline{CWC}$ values highlighting the most significant channels across all patients of the CHB-MIT dataset.

ber of records for a given patient, the selected pair of channels remained the same across all iterations, regardless of the excluded record. This demonstrates that the selected channel pair does not depend heavily on individual records, ensuring that the classification performance does not drop due to overfitting. The observed consistency in channel selection also indicates that only a few records are sufficient to achieve a reliable and stable selection, thereby simplifying the application of the described workflow.

However, for some specific patients, variability in channel selection was observed across the excluded records. In these

cases, a majority voting strategy was applied to determine the final channel pair for the subsequent training and testing phases, although the selection was not unanimous.

The exceptions mainly involved patients characterized by a limited number of records (fewer than four) containing ictal events, such as chb04, chb07, chb17, chb19, and chb21, making more challenging the selection of the channel 2. This variability also occurs when the algorithm penalizes channel pairs that exhibit high correlation or that are not exclusively selected from the two distinct classes, as in the case of chb06, coherently with findings reported in [114].

It is interesting to observe that the presence of multiple active channels for each phase in certain patients may potentially reflect the type of seizure, focal or generalized. However, due to the lack of specific annotations for seizure types within the dataset, a deeper investigation in this regard is not feasible.

We can notice that the $\overline{\rho}$ values, computed between the channels selected for different patients, remain below 0.33 in most cases, indicating minimal overlap in information between the two channels, which is in turn advantageous at the classification stage. The only exception is related to the chb06 and chb21 patients, exhibiting high correlation values (0.47 and 0.43, respectively). This in turn reflects on poor performance at classification stage, as discussed in next sections.

Performance evaluation

Once the channel pairs were selected for each patient, we evaluated the patient-specific classification performance, reporting the obtained results as follows. Table 5.2 lists patients where the two-selected channels configuration outperformed the four-temporal channels configuration, while Table 5.3 includes patients with similar performance across both configurations. Table 5.4 shows patients for whom the two-channel configuration performed slightly worse than the four-channel setup. Finally, Table 5.5 summarizes the results across all CHB patients. To assess the significance of these paired comparisons (four-temporal channels vs. two-selected channels), we also performed a paired non-parametric Wilcoxon signed-rank test [158], considering differences significant for p-value < 0.05.

Table 5.2: Performance comparison for patients in whom the two-selected-channel configuration outperformed the four-temporal configuration.

Patient	Config.	Segment-level			Event-level		
		SP	SN*	bAcc*	Delay [s]	FP/h	Detected Seizures*
chb04	4-temp.	1.00±0.00	0.18±0.21	0.59±0.11	40.0±25.5	0.33	2/4
	2-sel.	1.00±0.00	0.71±0.27	0.85±0.13	28.8±32.9	0.11	4/4
chb13	4-temp.	1.00±0.01	0.55±0.44	0.77±0.22	9.7±3.0	0.45	9/10
	2-sel.	0.99±0.01	0.65±0.31	0.82±0.16	9.7±2.5	0.36	10/10
chb14	4-temp.	1.00±0.00	0.32±0.41	0.66±0.20	6.5±0.6	0.04	4/8
	2-sel.	1.00±0.00	0.53±0.31	0.76±0.16	6.6±1.0	0.00	7/8
chb15	4-temp.	1.00±0.02	0.73±0.19	0.87±0.09	11.5±8.2	0.15	19/20
	2-sel.	1.00±0.02	0.82±0.16	0.91±0.08	6.7±5.0	0.21	20/20
chb17	4-temp.	1.00±0.00	0.28±0.25	0.64±0.13	39.5±7.8	0.05	2/3
	2-sel.	1.00±0.00	0.46±0.10	0.73±0.05	26.3±6.0	0.00	3/3
chb18	4-temp.	1.00±0.02	0.45±0.38	0.72±0.18	13.5±6.6	0.46	4/6
	2-sel.	0.99±0.04	0.55±0.31	0.76±0.15	16.2±13.1	0.11	5/6

Metrics marked with * are statistically significant under the paired non-parametric Wilcoxon signed-rank test ($p < 0.05$).

To conduct a comparative analysis with the state-of-the-art and demonstrate the effectiveness of the proposed approach, we also included in the same tables the results obtained using the configuration that takes as input four fixed channels located in the temporal region (F_7-T_7, T_7-P_7, F_8-T_8, and T_8-P_8) [92], as described in Section 4.1.3.2. Both configurations were trained and tested following the same procedure and evaluated using identical metrics, ensuring a fair comparison. This enables us to fairly compare the performance of our two-selected-channels configuration with a four-temporal-channels configuration similar to those adopted in recent works [165, 163, 86, 134].

For both channel configurations, the Table 5.5 shows the average results for each metric across all the 24 patients in the CHB dataset. In particular, the last column provides cumulative results, reflecting the total number of detected ictal events with respect to the total number of annotated events across all patients (in the form detected/total number).

Table 5.3: Performance comparison between the four-temp. and two-selected channels configurations for each patient. The final row shows the results (mean ± standard deviation) across all CHB patients.

Patient	Config.	Segment-level			Event-level		
		SP	SN	bAcc	Delay [s]	FP/h	Detected Seizures
chb01	4-temp.	1.00±0.00	0.86±0.19	0.93±0.10	9.0±6.9	0.00	7/7
	2-sel.	1.00±0.00	0.91±0.10	0.95±0.05	6.4±3.4	0.00	7/7
chb02	4-temp.	1.00±0.00	0.93±0.08	0.96±0.04	11.0±6.6	0.00	3/3
	2-sel.	1.00±0.00	0.93±0.09	0.96±0.04	8.3±2.9	0.08	3/3
chb03	4-temp.	1.00±0.01	0.77±0.14	0.89±0.07	17.4±6.5	0.03	7/7
	2-sel.	1.00±0.01	0.92±0.11	0.96±0.05	8.0±5.8	0.08	7/7
chb05	4-temp.	1.00±0.00	0.88±0.12	0.94±0.06	9.8±4.0	0.08	5/5
	2-sel.	1.00±0.00	0.83±0.14	0.92±0.07	21.8±16.3	0.00	5/5
chb07	4-temp.	1.00±0.00	0.57±0.22	0.79±0.11	14.3±5.7	0.00	3/3
	2-sel.	1.00±0.00	0.59±0.31	0.79±0.16	28.3±12.3	0.03	3/3
chb08	4-temp.	1.00±0.00	0.83±0.15	0.91±0.07	11.2±3.3	0.20	5/5
	2-sel.	0.99±0.02	0.85±0.07	0.91±0.05	15.2±3.1	0.03	5/5
chb09	4-temp.	1.00±0.00	0.94±0.03	0.97±0.01	7.3±1.7	0.00	4/4
	2-sel.	1.00±0.00	0.93±0.03	0.97±0.01	8.3±2.5	0.00	4/4
chb10	4-temp.	1.00±0.00	0.96±0.07	0.98±0.03	6.0±2.3	0.00	7/7
	2-sel.	1.00±0.00	0.94±0.08	0.97±0.04	4.9±3.0	0.00	7/7
chb11	4-temp.	1.00±0.00	0.83±0.25	0.91±0.12	3.0±1.7	0.03	3/3
	2-sel.	1.00±0.00	0.82±0.08	0.91±0.04	5.0±5.2	0.17	3/3
chb19	4-temp.	1.00±0.00	0.94±0.03	0.97±0.01	9.3±3.8	0.10	3/3
	2-sel.	1.00±0.00	0.83±0.02	0.91±0.00	16.7±1.5	0.14	3/3
chb20	4-temp.	1.00±0.00	0.59±0.42	0.80±0.21	11.9±6.9	0.03	7/8
	2-sel.	1.00±0.00	0.57±0.21	0.78±0.10	13.9±5.0	0.07	7/8
chb22	4-temp.	1.00±0.00	0.89±0.17	0.94±0.08	12.3±11.0	0.00	3/3
	2-sel.	1.00±0.00	0.82±0.12	0.91±0.06	16.3±7.0	0.00	3/3
chb23	4-temp.	1.00±0.01	0.86±0.13	0.93±0.06	12.0±5.9	0.18	7/7
	2-sel.	1.00±0.01	0.86±0.05	0.93±0.03	9.1±2.4	0.18	7/7
chb24	4-temp.	1.00±0.00	0.52±0.27	0.76±0.13	9.8±3.2	0.14	13/16
	2-sel.	1.00±0.00	0.50±0.30	0.75±0.15	9.1±2.5	0.14	13/16

Table 5.4: Performance comparison for patients in whom the two-channel configuration performed slightly worse than the four-channel setup.

Patient	Config.	Segment-level			Event-level		
		SP	SN	bAcc	Delay [s]	FP/h	Detected Seizures
chb06	4-temp.	1.00±0.00	0.63±0.35	0.81±0.17	7.9±1.7	0.07	8/10
	2-sel.	1.00±0.00	0.24±0.32	0.62±0.16	9.0±2.2	0.00	4/10
chb12	4-temp.	1.00±0.01	0.42±0.31	0.71±0.15	11.4±5.9	0.29	20/27
	2-sel.	1.00±0.01	0.38±0.20	0.69±0.13	11.6±7.9	0.33	19/27
chb21	4-temp.	1.00±0.00	0.65±0.24	0.82±0.12	14.3±14.4	0.03	4/4
	2-sel.	1.00±0.00	0.29±0.28	0.64±0.14	26.0±18.5	0.03	3/4
Edge case							
chb16	4-temp.	0.99±0.05	0.00	0.50±0.00	–	0.88	0/8
	2-sel.	1.00±0.00	0.00	0.50±0.00	–	0.00	0/8

Table 5.5: Performance comparison between the four-temporal and two-selected channels configurations across all CHB patients reported as mean ± standard deviation.

Config.	Segment-level			Event-level		
	SP	SN	bAcc	Delay [s]	FP/h	Detected Seizures
4-temp.	1±0.01	0.66±0.32	0.83±0.16	11.38±8.02	0.15±0.21	149/181
2-sel.	1±0.01	0.67±0.31	0.83±0.16	11.47±9.75	0.10±0.11	152/181

We can notice that our two-selected channels based system successfully detected 152 out of 181 ictal events, with an average delay of approximately 11 seconds, that is about one-third of the average duration of a seizure calculated within the analyzed dataset. The classification performance yielded an average SP of 1.00 and SN of 0.67, achieving a remarkable bAcc of 0.83. Moreover, the system demonstrated high robustness against false positives, with a low false positive rate of only 0.10 FP/h. The model obtained for each patient has a size of 51 KB, with approximately 9.5k model parameters, thus suitable

for its deployment on low-cost microcontroller units.

5.1.3 Discussion

Overall analysis

The proposed method enables the calculation of CWCs, offering a quantifiable measure of the relative importance of each channel in datasets acquired through standard EEG systems. This approach has been crucial in identifying the most relevant scalp regions during seizures, thereby facilitating optimized electrode placement for continuous seizure monitoring in a personalized manner. By focusing on channels with the highest CWC values, we effectively isolate those that best capture seizure-related activity while suppressing redundant and noisy signals. Based on linear inter-relationships and orthonormal transformations among the different channels, our method is particularly advantageous for low-dimensional EEG datasets, where conventional analysis techniques often struggle to distinguish between meaningful patterns and irrelevant noise. To the best of our knowledge, no previous study has utilized PCA in this specific, time-aware, frequency-specific, and channel-focused manner. For instance, Li *et al.* proposed PCA for dimensionality reduction of the features extracted from EEG signals, transforming the original high-dimensional feature space into a smaller set of uncorrelated PCs [139]. Chakrabarti *et al.* used PCA to select the subset of channels that have the highest correlation with the PC. However, without a thorough analysis across different frequency bands and without a time-aware approach, they were only able to achieve acceptable performance by reducing the number of channels from 23 to 18 [166]. On the other hand, Mateo *et al.* applied PCA for noise elimination in EEG data [140], demonstrating effectiveness when the noise level is sufficiently low.

Intriguingly, the scalp regions identified by our algorithm are largely consistent with those highlighted by a more complex method described in [114], which involves training a CNN architecture that takes 21 channels as input and produces EEG Channel Activation Maps to identify the brain regions activated during seizures. In addition, leveraging its low computa-

tional cost, the method processes data on a second-by-second basis, producing dynamic maps that capture the temporal evolution of the observed parameters. These maps offer an effective and intuitive visualization tool for clinical experts, facilitating data interpretation and analysis.

This practical applicability should, however, be interpreted in light of the different roles played by the components of the pipeline. Specifically, the PCA-based computation of channel importance coefficients is performed only during the offline personalization stage, where the most informative patient-specific electrodes are identified. After this selection, inference is carried out directly on the reduced EEG input, without recomputing PCA. Therefore, the computational cost relevant to runtime deployment is primarily associated with the lightweight CNN operating on the selected two-channel input, rather than with the channel-selection stage.

Overall, thanks to the entire optimization strategy, we developed a system working with only two selected channels, demonstrating to be more advantageous than recently analyzed four-channel systems. Selecting more than two channels (i.e., three or four) would introduce increased system complexity and reduced wearability, while introducing negligible gain in performance for this specific detection task. On the other hand, the feasibility of a single-channel solution is investigated separately in Section 5.2, where it is shown that using only one channel leads to a noticeable performance drop with respect to the proposed dual-channel configuration.

The results obtained with the two-channel configuration are comparable to those of the four fixed-channel setup in terms of sensitivity and detection delay, while reducing the system complexity and footprint and improving the false-positive rate. Specifically, we observed no differences in terms of SP, which can be attributed to the inherent imbalance between ictal and inter-ictal classes in the training and validation datasets. This imbalance naturally results in a negligible number of FP compared to the high number of TN. When considering SN, the two-channel configuration achieved a slight improvement of about 1%, reflecting an increased number of detected ictal events (152 vs. 149) while maintaining a comparable aver-

age detection delay of approximately 11 seconds. We explicitly remark that the balanced accuracy value reported in Table 5.5 is penalized compared to the traditional accuracy calculation, which is instead based on the number of correct predictions divided by the total number of predictions. By calculating the accuracy, we obtain a value of 0.99, strongly influenced by the high number of TN correctly identified.

More importantly, our two-channel setup demonstrated increased robustness against FP, due to the reduction of noise and redundancy. The average FP rate in the state of the art four-temporal channels configuration was higher at 0.15 FP/h compared to 0.10 FP/h in the proposed two-channel system, with a non-negligible improvement of 30% for our system, that further confirms the advantage of the proposed optimized approach. Reducing the FP rate is crucial for improving the reliability of seizure detection systems, minimizing false alarms that can lead to unnecessary interventions, thus in turn improving clinical efficiency and patient quality of life by ensuring a more trustworthy and user-friendly solution. Frequent false alarms may also cause alarm fatigue, reducing responsiveness to actual seizures [131].

This aspect is particularly relevant in light of the strong class imbalance that characterizes real-world monitoring scenarios, a condition that is already partially reflected in our evaluation setting. In the patient-specific experiments, the leave-one-record-out protocol uses full one-hour recordings as test samples, which may or may not contain ictal events. Moreover, performance was averaged over all available records, including seizure-free ones, specifically to assess robustness against false alarms. As a result, the reported FP/h values already capture model behavior during prolonged non-ictal periods. Nevertheless, in long-term deployment, even a stable FP/h may accumulate into a clinically relevant number of false alarms over days or weeks.

For this reason, a comprehensive assessment of real-world applicability should not rely on FP performance alone, but must also take system sensitivity into account, since missed detections may limit the effectiveness of the alarm function in practical use. From this perspective, the sensitivity values

reported in 5.5 suggest that the proposed system should currently be regarded as a supportive monitoring tool rather than a fully autonomous safety-critical device. Although the cumulative number of detected seizures is high, a mean sensitivity of about 0.67 indicates that a non-negligible fraction of ictal events may still be missed in some patients. In a real-world setting, these false negatives could reduce the reliability of the alarm function and limit the system's effectiveness when immediate warning is required. On the other hand, the low FP rate and the short detection delay remain highly relevant for wearable applications, since excessive false alarms would strongly compromise usability and long-term adherence. Overall, these results support the feasibility of the proposed reduced-channel solution for continuous monitoring, while also indicating that further improvements in patient-specific adaptation and longitudinal updating are necessary before considering fully autonomous deployment.

Performance comparison with the state-of-the-art

The purpose of this section is to position the proposed method with respect to recent literature, rather than to provide a strict head-to-head benchmark, especially since only a limited number of studies jointly address both EEG channel reduction and lightweight network design. Accordingly, unless explicitly stated otherwise, the results shown were taken from the original publications, as a full re-implementation and re-evaluation of all competing methods within our experimental pipeline was not always feasible given the methodological details available in the corresponding studies.

The methods summarized in the Table 5.6 can be broadly divided into two groups. The first includes lightweight seizure detection approaches that do not apply any channel selection strategy, while the second can be considered fully lightweight, both in terms of network architecture and EEG setup, making them more suitable for wearable applications. Several of these studies rely on configurations that adopt all the channels provided by standard headsets (for example following the 10-20 standard). For instance, LightSeizureNet [114] achieved outstanding performance with an SN of 0.99, while other recent

Table 5.6: Comparison against existing literature approaches in terms of characteristics and achieved performance.

Ref.	# Ch.	Selection method	Classification model	SN	Acc
[114]	23	not impl.	CNN	0.99	N/A
[167]	23	not impl.	CNN	N/A	0.96
[168]	23	not impl.	VGGNet (CNN)	0.99	0.98
[169]	23	not impl.	PCNN–Bi-LSTM	0.98	0.99
[90]	21	not impl.	EEGNET-8.2 (CNN)	0.81	0.96
			EEGWaveNet (CNN)	0.69	0.98
This work	21	not impl.	CNN	0.82	0.91 (bAcc)
[166]	18	PCA	MLP	N/A	0.87
[170]	5-opt.	MI and RF	SVM	0.87	0.98
[134]	4-fixed	PCC	CNN	N/A	0.99
[92]	4-fixed	not impl.	EpiDeNet (CNN)	0.69	–
[97]	3-fixed	highest variance	SVM	N/A	0.83
[136]	3-opt.	LAFS	Multi-Head Self-Attention	0.65	0.85
[171]	2-opt.	Ch. Attention + MLP	CNN+GRU	N/A	0.72
This work	2-opt.	temporal PCA	CNN	0.67	0.99 0.83 (bAcc)
	21	not impl.	CNN	0.82	0.91 (bAcc)

fixed: fixed channels; opt.: optimized channels; N/A: Not available in the original publication; not impl.: not implemented.
All acronyms are reported in the Abbreviations list at the end of the manuscript.

studies such as those reported by Wang *et al.* [169] and Ke *et al.* [168] reported SN values above 0.97, confirming the effectiveness of compact architectures when full-channel data are available.

Anyway, it is important to point out that the concept of a lightweight approach extends beyond algorithmic compactness alone. Our focus, in fact, is on the development of a fully lightweight pipeline, which also includes a minimal EEG setup

using at most two channels, so as to be easily integrated into a wearable device. In this regard, several strategies have been recently proposed to identify optimized subsets of EEG channels. These include, for instance, techniques based on Locally-Adaptive Feature Selection (LAFS) [136], Mutual Information (MI) and RFs [170], Pearson Correlation [134], and variance-based ranking [97]. Chakrabarti *et al.* employed the PCA to select a number of channels varying from a minimum of 4 to a maximum of 22. Their best performance was achieved using an 18-channel configuration, reaching an accuracy of 0.86, compared to our reported accuracy of 0.99 using only two channels. As the number of channels decreased, performance declined, with accuracy dropping to 0.78 when using four channels [166].

In general, methods reporting the highest performance typically rely on a higher number of EEG channels, which compromises their feasibility for wearable applications. For instance, MI and RF-based approaches reported an SN of 0.86 using five channels [170], a configuration that, although effective, limits applicability in daily use scenarios. Conversely, when the number of channels is further reduced, a performance drop is often observed. As an example, Gifford *et al.* [136] used LAFS to select three channels (thus one more than the proposed approach), achieving an SN of 0.65 (vs 0.67 of the proposed approach), despite using one more channel. Similarly, Shoka *et al.* [97] adopted a variance-based selection method to identify three channels, yet their reported Acc did not exceed 0.83 (vs 0.99 of the proposed approach). In this context, our method offers a compelling alternative. By combining patient-specific optimization with an efficient channel selection strategy, we demonstrate that noteworthy performance can still be achieved using two EEG channels alone. For instance, our method allows us to significantly outperform the 0.72 Acc reported by Affes *et al.* [171] with a system based on two channels. This comparison highlights the effectiveness of our patient-specific two-channel configuration, which not only reduces complexity and enhances suitability for wearable systems but also maintains, and in some cases also improves, detection performance compared to several multi-channel solutions.

It is worth noting that we also evaluated the proposed

lightweight CNN using the full 21-channel montage, obtaining an average SN of 0.82 across patients. This closely matches the performance reported for EEGNet in Thuwajit et al. [90], which we regard as the current state of the art among lightweight models operating with all channels. A comparison in terms of bAcc is not possible because the authors do not report SP. However, our 21-channel SN is about 7 % points lower than LightSeizureNet [114], which performs patient-specific pruning, tuning the pruning level for each patient based on performance degradation. By contrast, EEGNet employs a single optimization of the architecture across all patients, and we follow the same cross-patient optimization strategy in this work.

Taken together, these observations further suggest that the comparison in Table 5.6 should be interpreted mainly as an indication of how the proposed method relates to the current literature. In particular, they show that the proposed approach remains competitive with recent lightweight solutions, while addressing a more challenging and less explored setting in which both architectural simplicity and aggressive channel reduction are simultaneously required.

CHB-MIT Patient-wise analysis

While the overall results confirm the validity of our approach, a more detailed analysis at the patient level reveals important insights into individual variability and the factors influencing performance. It is worth highlighting that most studies published in recent years on the same topic have limited their experimental evaluation to a small subset of patients from the same dataset [35, 74, 131]. In contrast, our study extends the analysis to all available patients, providing a more comprehensive overview of the system's performance and evaluating its robustness across a diverse range of seizure patterns and patient-specific characteristics.

In this regard, Table 5.2 highlights patients for whom the two-channel configuration detects more epileptic seizures than the four-temporal channels setup, with ictal events detected compared to 40, out of a total of 51. This improvement corresponds to an average increase in SN from 0.50 to 0.66, which deserves particular attention. The state-of-the-art four-temporal-

channel approach relies on fixed electrode positions, which may not capture signals from scalp regions most relevant to the epileptic activity. As a result, the recorded data may be weakly informative, leading to a higher number of FN and a consequent drop in SN. In contrast, our method selects the two most informative scalp regions for each patient, enabling the classification model to better detect ictal patterns and thus improve SN. For example, in the case of patient chb04, SN increases remarkably from 0.18 to 0.71 when using the optimized, patient-specific two-channel configuration. For all these patients, the channel selection algorithm identified a single pair of uncorrelated bipolar channels (one emerged from the ictal phase, and the other from the no-seizure phase), coherently with observation based on activation maps provided in [114]. In two of these six patients (chb13, chb17), the selected channels do not overlap with those in the four-channel configuration; in the other four cases, only one channel is shared between the configurations. This results subset underscores even more sharply the importance of selecting an optimal channel pair for effective seizure detection, often outperforming systems with a higher number of channels.

Table 5.3 focuses on patients where the two-channel configuration performs comparably to the four-temporal setup in terms of event detection. In most of these cases, the four-channel configuration identified all ictal events; however, reducing the number of channels did not adversely affect detection capability or sensitivity. For the majority of these patients, channel selection was not critical, as the proposed algorithm naturally identified a single optimal pair of channels. Conversely, in seven cases (chb02, chb03, chb11, chb19, chb20, chb22, and chb23), multiple channels were identified in both ictal and no-seizure phases, requiring correlation analysis to select the optimal channel pair with the lowest correlation coefficient. The identification of multiple potential channel pairs suggests that these patients may experience generalized seizures, although the lack of specific annotations for this type of seizure in the dataset limits further investigation. Nevertheless, even for this subset of patients, an optimized two-channel configuration was identified without compromis-

ing performance.

It should also be noted that, in some entries of Tables 5.2 and 5.3, the standard deviation of the delay is comparable to its mean value. This should not necessarily be interpreted as instability of the proposed method; rather, it reflects the variability of detection latency across held-out test records and seizures. In the LORO framework, each record is evaluated separately, while channel selection and model training are performed using only the remaining records of the same patient. As a consequence, at each LORO fold, the selected channel pair may not be equally optimal for the excluded test record, since seizure dynamics and onset patterns can vary across different ictal events even within the same patient. This may lead to earlier detections in some records and later detections in others, thereby increasing the dispersion of delay values. This effect is particularly evident in patients with a limited number of ictal records, for whom the channel-selection process is inherently more sensitive to record-specific variability. In addition, the post-processing stage identifies seizure onset only when three consecutive overlapping windows are classified as ictal. Although this rule improves robustness against isolated misclassifications, it can also delay onset detection when one or more false-negative windows occur during the transition to the ictal phase, or when seizure-related EEG patterns become clearly detectable only after the annotated clinical onset. Therefore, delay should be interpreted together with the number of detected seizures and, when relevant, in light of seizure duration.

Table 5.4 reports the few patients (3 out of 24) for whom the two-channel configuration performed slightly worse than the four-channel setup. For these patients, the model detected fewer seizures and exhibited a lower SN (0.31 vs. 0.54 on average). This reduced performance can be attributed to the complexity of the selection process, which was not straightforward for all patients. In some cases, it was challenging to identify distinct and uncorrelated channels for the ictal and no-seizure phases. For patient chb06, the algorithm identified the two channels as significant in only one phase due to the XOR operation employed, with a high correlation ($\overline{\rho} = 0.47$) be-

tween the selected channels. This resulted in a reduced number of detected ictal events and lower performance metrics, with SN dropping from 0.63 to 0.24, bAcc decreasing from 0.81 to 0.62. Similarly, patient chb21 also exhibited a high correlation ($\overline{\rho} = 0.43$) between the selected channels. These results confirm that selecting channels conveying distinct and non-overlapping information from the ictal and inter-ictal phases is crucial for optimizing classification performance.

Finally, it is important to consider that the complexity of the detected signals, in terms of the number and duration of ictal events, could affect detection performance. For instance, patient chb12 frequently experienced multiple ictal events (two to five) within the same record. This phenomenon was even more evident for patient chb16, where seizure detection was problematic in both configurations due to the very short duration of ictal phases (average 9.00 ± 3.00 seconds). These two specific cases have been non-trivial also for neurologists as discussed in previous works [35]. Notably, our approach successfully detected all the seizures occurring in the subset of patients considered in [35], also in challenging ones, outperforming earlier efforts involving expert-driven channel selection and dataset re-annotation.

Table 5.7: Selected channel pairs with their corresponding average correlation $\overline{\rho}$ for Siena Scalp dataset. Channel A refers to the inter-ictal phase, while Channel B refers to the ictal phase.

Patient	Channel A	Channel B	$\overline{\rho}$	Patient	Channel A	Channel B	$\overline{\rho}$
PN00	F_3-C_3	F_8-T_8	0.20	PN12	T_8-P_8	P_7-O_1	0.25
PN01	Cz-Pz	F_7-T_7	0.28	PN13	F_8-T_8	T_7-P_7	0.30
PN03	C_4-P_4	F_8-T_8	0.31	PN14	P_8-O_2	C_3-P_3	0.22
PN05	T_7-P_7	T_7-FT_9	0.53	PN16	F_7-T_7	T_8-P_8	0.23
PN06	F_3-C_3	T_7-P_7	0.34	PN17	T_7-P_7	T_8-P_8	0.33
PN09	F_3-C_3	T_8-P_8	0.23				

Additional validations on the Siena Scalp Dataset

To assess the robustness and generalizability of the proposed pipeline, the entire experimental protocol was replicated on the Siena Scalp dataset [80], without modifying pre-processing, training, or evaluation settings. The same channel-selection procedure was adopted providing the results shown in Table 5.7. For completeness, Figure 5.4 shows the corresponding maps of the digitized values of $\overline{\text{CWC}}$ with a threshold of 0.9. Notably, in this case, most patients do not present only two channels in the final map, but three.

Excluding the exceeding channel was straightforward, as it was strongly correlated with one of the other two (see the last paragraph of Section 4.1.3.1 for more details on the impact of correlation in channel selection ranking). Furthermore, the channel selected by the algorithm for the ictal phase corresponds to the lateralization of the seizure indicated in the dataset specifications [80].

The same model configuration and detection metrics were adopted to allow a direct comparison with the CHB-MIT analysis. The overall results for all Siena Scalp patients are summarized in Table 5.8. As can be seen, the findings on this second dataset closely mirror those obtained on CHB-MIT: the optimized two-channel configuration maintains SN and detection delay comparable to multi-channel baselines, while further reducing FP/h and preserving a minimal hardware footprint.

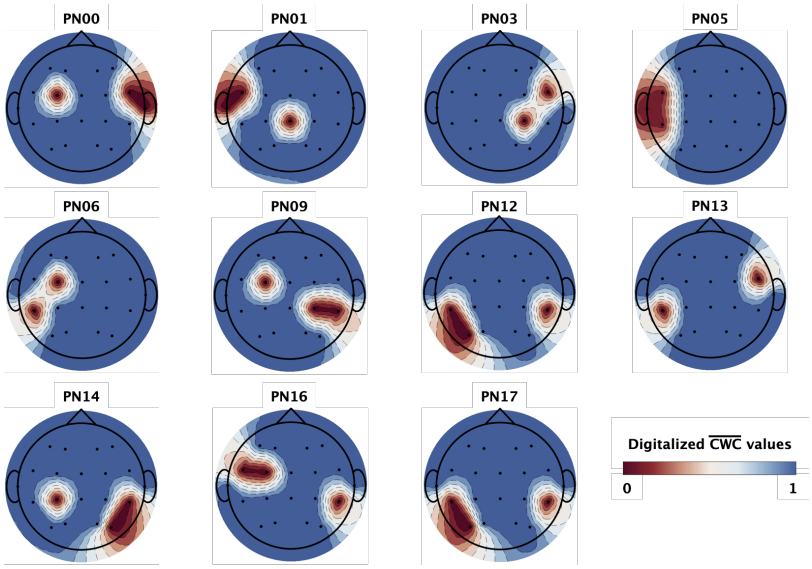


Figure 5.4: Maps of the digitalized $\overline{CWC}$ values highlighting the most significant channels across patients of the Siena dataset.

Table 5.8: Performance comparison between the four-temporal and two-selected channels configurations across all Siena scalp patients reported as mean ± standard deviation.

Config.	**Segment-level**			**Event-level**		
	SP	**SN**	**bAcc**	**Delay [s]**	**FP/h**	**Detected Seizures**
4-temp.	0.98±0.07	0.60±0.40	0.78±0.19	13.46±7.69	0.55±0.12	28/36
2-sel.	0.99±0.03	0.67±0.32	0.83±0.16	13.5±6.40	0.39±0.42	31/36

5.2 Single-Channel Analysis

Section 5.1 analyzed the application of the *Channel Weight Coefficient* (CWC), an interpretable metric proposed to quantify the importance of each channel in the standard system for detecting epileptic seizures. As described in Section 4.1.3.1, during the definition of the CWC calculation procedure, it emerged that, in the four main phases of a seizure, no-seizure, pre-ictal, ictal, and post-ictal [31], a single channel, or a small subset of channels, tends to consistently have higher weights. Based on this evidence, two channels were selected, one representative of the *ictal* phase and one of the inter-ictal phase [13]. This two-channel configuration allowed for the detection of a higher number of ictal events than the four-channel configuration, as shown in the previous Section.

Based on these findings, in this section we aim to further reduce the complexity of the system by investigating the feasibility of an approach based on a single EEG channel. This investigation represents a step forward in the development of compact, low-power, patient-customized wearable EEG devices for epilepsy management. Furthermore, it reinforces the two-channel study, as it highlights not only the need for accurate channel selection, but also the effectiveness of the proposed algorithm.

To this end, three configurations are compared: (i) a single channel associated with the ictal phase; (ii) a single channel associated with the non-ictal phase; (iii) a channel not involved in either phase. The results obtained were presented in the paper "Wearable EEG systems: toward Single-Channel personalized configurations for seizure detection" presented at the 9th National Congress of Bioengineering (GNB) [14].

5.2.1 Implementation details

Figure 5.5 shows the modified channel selection procedure used for this analysis. The scheme follows the one proposed in the first study (see section 5.1), but without the normalization, thresholding and XOR steps. In detail, the procedure comprises the following three steps:

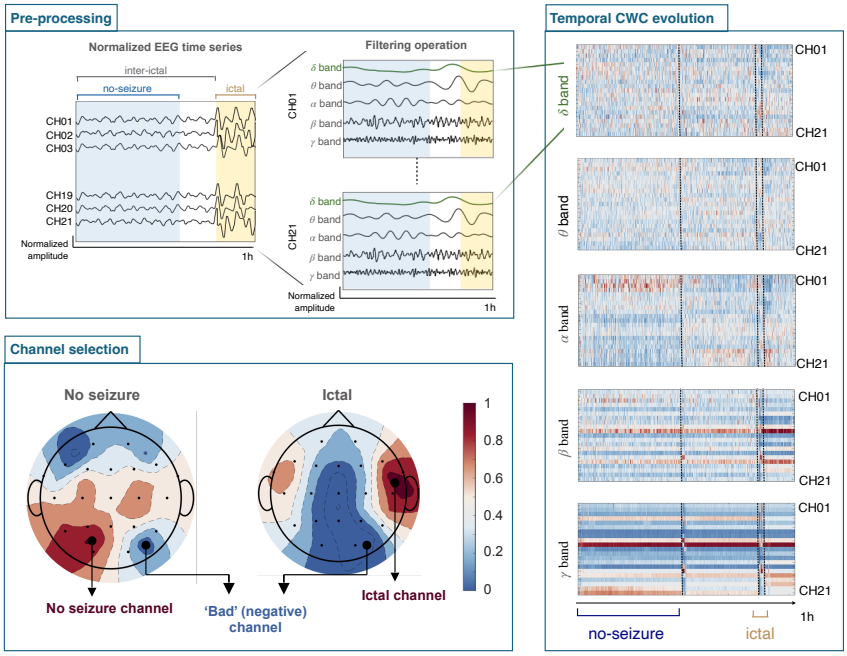


Figure 5.5: Channel selection framework for extracting a single channel in three configurations: (i) ictal, (ii) no-seizure (inter-ictal), and (iii) bad-channel.

- The pre-processing of raw EEG data. The raw EEG signals of each patient were filtered into the five standard frequency bands commonly used in clinical practice, namely $\delta[0.5, 4]Hz$, $\theta(4, 8]Hz$, $\alpha(8, 13]Hz$, $\beta(13, 35]Hz$ and $\gamma(35, 80]Hz$. The CWC coefficients were then computed over one-second time windows.

- The evaluation of the temporal evolution of the CWC. By analyzing the temporal evolution of the coefficients divided by band, it was possible to observe how the contribution of each EEG channel to the overall information fluctuates over time during critical events. This analysis provided a solid basis for the development of the channel selection algorithm we described in Section 4.1.3.1.

- The selection of the three individual channels of interest. Focusing on the ictal and non-ictal phases, the CWC values were normalized and averaged across all frequency

bands, obtaining representative maps for each phase. Three individual channels were extracted from these. One strongly representative of the non-seizure phase (red area in the non-seizure map), one representative of the ictal phase (red area in the ictal map) and one not representative of either phase (blue area in both maps).

Once the individual channels were identified, they were placed at the input of the lightweight network discussed in Section 4.1.3.2, in Section 4.1.3.2. Originally designed to process four input channels by Ingolfsson *et al.* [92], then adapted to receive two-channel EEG signals [13], the network has been applied here to Single-Channel EEG signals. The main modification concerned adjustments to the kernel size of the final max pooling layer, in order to adapt the architecture to the different input size.

5.2.2 Results

In this section, we summarize the findings for the three single-channel configurations on both the CHB-MIT [79] and Siena Scalp [80] datasets to ensure consistency across heterogeneous cohorts and acquisition protocols. We apply the modified EEG channel-selection procedure described in Section 5.2.1 to identify the optimal single channels. We then evaluate personalized seizure-detection performance on those channels using the leave-one-record-out scheme detailed in Section 5.1.1.

Channel selection.
Tables 5.9 and 5.10 report the channels selected for all patients in the CHB-MIT [79] and Siena Scalp [81] datasets, respectively. The tables highlight the three configurations considered in this study: (i) the channel active during the no-seizure phase, (ii) the channel active during the ictal phase, and (iii) the channel inactive in both ictal and non-seizure phases.

Table 5.9: EEG channels selected for each patients of the CHB-MIT dataset.

Patient	no-seizure	ictal	negative	Patient	no-seizure	ictal	negative
chb01	P_3-O_1	F_8-T_8	P_8-O_2	chb13	P_4-O_2	Fp_1-F_7	Fp_2-F_8
chb02	T_7-P_7	T_8-P_8	Fp_2-F_8	chb14	T_8-P_8	Fp_1-F_7	Fz-Cz
chb03	P_8-O_2	T_7-P_7	Fp_2-F_8	chb15	Cz-Pz	T_7-P_7	Fp_1-F_7
chb04	Cz-Pz	T_8-P_8	Fp_2-F_8	chb17	C_4-P_4	P_4-O_2	Fp_2-F_4
chb05	Cz-Pz	P_8-O_2	Fp_2-F_8	chb18	C_4-P_4	T_8-P_8	Fp_2-F_8
chb06	F_8-T_8	T_8-P_8	Fp_2-F_4	chb19	Cz-Pz	T_8-FT_{10}	Fp_2-F_4
chb07	Cz-Pz	P_8-O_2	Fp_1-F_7	chb20	P_8-O_2	T_7-FT_9	Fz-Cz
chb08	P_8-O_2	F_4-C_4	Fp_1-F_7	chb21	T_8-P_8	C_3-P_3	Fp_2-F_8
chb09	P_7-O_1	T_8-P_8	Fp_2-F_8	chb22	F_4-C_4	T_7-P_7	Fp_2-F_8
chb10	T_8-P_8	T_7-P_7	Fp_2-F_4	chb23	P_8-O_2	T_7-P_7	Fp_2-F_8
chb11	Fz-Cz	T_7-P_7	T_8-P_8	chb24	C_4-P_4	F_8-T_8	Fp_1-F_7
chb12	F_3-C_3	F_4-C_4	P_7-O_1				

Table 5.10: EEG channels selected for each patient of the Siena Scalp dataset.

Patient	no-seizure	ictal	negative	Patient	no-seizure	ictal	negative
PN00	F_3-C_3	F_8-T_8	P_7-O_1	PN12	T_8-P_8	P_7-O_1	Fp_2-F_4
PN01	Cz-Pz	F_7-T_7	Fp_2-F_4	PN13	F_8-T_8	T_7-P_7	C_3-P_3
PN03	C_4-P_4	F_8-T_8	P_8-O_2	PN14	P_8-O_2	C_3-P_3	Fp_2-F_4
PN05	T_7-P_7	T_7-FT_9	P_4-O_2	PN16	F_7-T_7	T_7-P_7	P_7-O_1
PN06	F_3-C_3	T_7-P_7	Fp_2-F_8	PN17	T_7-P_7	T_8-P_8	Fp_1-F_3
PN09	F_3-C_3	T_8-P_8	Fp_1-F_7				

Table 5.11: System performance across all configurations on the CHB-MIT dataset.

Metrics	Single-Channel			Dual-Channel
	ictal	no-seizure	negative	
SP	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.01
SN	0.54 ± 0.34	0.36 ± 0.35	0.28 ± 0.32	0.67 ± 0.30
bAcc	0.77 ± 0.17	0.69 ± 0.18	0.64 ± 0.16	0.83 ± 0.15
delay [s]	13.6 ± 11.1	17.7 ± 15.1	18.2 ± 13.1	11.5 ± 9.75
Fp/h	0.21 ± 0.27	0.09 ± 0.13	0.13 ± 0.18	0.15 ± 0.18
seizures	136 / 173	98 / 173	73 / 173	152 / 173

Performance evaluation.
Tables 5.11 and 5.12 report performance for epileptic-seizure detection across the considered single-channel configurations. Performance is provided at both the *segment level*, in terms of sensitivity (SN), specificity (SP), and balanced accuracy (bAcc), and the *event level*, in terms of average detection delay, false positives per hour (FP/h), and the number of correctly identified ictal events. See Section 4.3 for metric definitions.

In both tables, the first column shows the performance obtained using only the channel selected in the ictal phase; the second column shows the results for the channel identified during the no-seizure phase; the third column shows the performance of the single "bad" channel, i.e. the channel that did not emerge in either the ictal or no-seizure phase. Finally, the last column on the right summarizes the results obtained with the two-channel configuration, already discussed in Section 5.1.2, reported here to allow a direct comparison with the single-channel configurations.

5.2.3 Discussion

The selection of channels in this work is guided by the temporal trend of CWC coefficients in the different phases of the epileptic seizure, i.e. inter-ictal, pre-ictal, ictal, and post-ictal

Table 5.12: System performance across all configurations on the Siena Scalp dataset.

Metrics	Single-Channel			Dual-Channel
	ictal	no-seizure	negative	
SP	0.99 ± 0.03	0.99 ± 0.07	0.93 ± 0.17	0.99 ± 0.03
SN	0.64 ± 0.35	0.55 ± 0.39	0.40 ± 0.35	0.67 ± 0.32
bAcc	0.81 ± 0.18	0.77 ± 0.19	0.66 ± 0.16	0.83 ± 0.16
delay [s]	14.0 ± 8.08	13.83 ± 8.49	22.2 ± 8.9	13.5 ± 6.40
Fp/h	0.84 ± 0.8	0.39 ± 0.47	0.48 ± 0.53	0.39 ± 0.42
seizures	29/36	23/36	20/36	31/36

[31]. Since neural activity tends to concentrate in different regions of the scalp depending on the phase, visualizing the dynamics of CWC coefficients is an effective means of identifying the most informative channels at any given moment. By exploiting this property, it was possible to identify the relevant channels specifically for each patient, using a data-driven and interpretable approach.

As an example, Figure 5.6 shows the topographic maps for patient chb01 from the CHB-MIT dataset [79]. From these maps, it can be intuitively observed that the most significant bipolar channels are P_3-O_1 during the non-seizure phase (Figure 5.6a) and F_8-T_8 during the ictal phase (Figure 5.6c), according to the international 10-20 standard [20]. This approach not only allows the most relevant channels to be automatically identified for personalized medicine, but also provides clinicians with a valuable visualization tool. The maps obtained are intuitive and interpretable, promoting a better understanding of neural dynamics and supporting informed decisions.

In the specific case of patient chb01, it can be observed that neural activity spreads from the left parieto-occipital region (Figure 5.6a) to the right fronto-temporal area (Figure 5.6c) with the onset of the seizure, crossing a large fronto-central region in the pre-ictal phase (Figure 5.6b). In the post-ictal phase (Figure 5.6d), activity spreads to the posterior area of

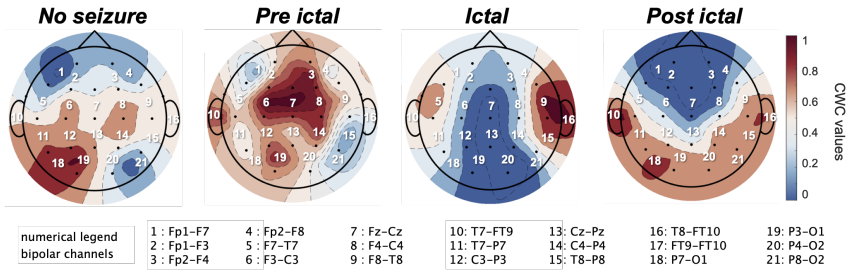


Figure 5.6: Average CWCs maps derived from all the records of patient chb01 highlighting how the informative area shifts across different regions of the scalp during a seizure.

the scalp, with greater involvement of the temporal and occipital regions, probably reflecting compensatory mechanisms of neuronal recovery.

For the other patients in both the CHB-MIT dataset (Table 5.9) and the Siena Scalp dataset (Table 5.10), a qualitative analysis of the selected channels reveals some recurrent spatial tendencies. In particular, channels associated with the ictal phase are more frequently located over temporal or fronto-temporal regions, whereas channels associated with the no-seizure/inter-ictal phase tend to be more often found over central or posterior regions. This trend is also consistent with the generalized analysis (see Section 5.3), in which Cz-Pz emerges as the representative inter-ictal channel for both datasets, while the ictal channel is lateralized over temporal regions, namely T8-P8 in CHB-MIT and T7-P7 in Siena Scalp. Nevertheless, these tendencies should not be interpreted as a universal rule, since substantial inter-patient variability is still observed, further confirming the importance of a personalized channel-selection strategy.

It is important to note that the channel selection procedure was applied using a leave-one-record-out approach (see Section 5.1.1 for further details). In most cases, this strategy led to consistent identification of the same channel for each phase in the different records, with only rare exceptions. These exceptions typically involve patients with a limited number of ictal events, or cases requiring additional discrimination based on correlation. In such instances, majority voting was applied to

select the optimal channel pair. This consistency highlights the robustness and reliability of the algorithm against intra-patient variability, regardless of the excluded record.

Results achieved on the CHB-MIT [79] dataset and reported in Table 5.11, reveal that the Dual-Channel configuration provides the best performance across all key metrics evaluated at the segment level, achieving SN of 0.67 and SP of 1.00, which result in a considerable bAcc of 0.83. Interestingly, these performances are comparable, and in some cases even superior, to those of recently proposed systems that rely on a large number of channels [136]. Our configuration detected 152 out of 173 total ictal events, with an average detection delay of 11.5 seconds, the shortest among the analyzed configurations. 0.15 FP/h indicates that the system generates one false alarm approximately every 7 hours.

The Single-Channel configurations, while simpler and more resource-efficient, exhibit reduced sensitivity, impacting their capability to identify seizures. A noticeable performance difference emerges between the two analyzed Single-Channel configurations, i.e. ictal and no seizure (second and third column of Table 5.11, respectively). The first yields a suboptimal SN value of 0.54, outperforming the 0.36 achieved with the second one. SP, instead, is 1.00 in both cases (this aspect will be discussed later). As a result, bAcc in the first case stands at 0.77, which is 6% lower than the best result obtained with two channels. In contrast, the other configuration has a lower bAcc (0.69), demonstrating a more significant performance gap.

Regarding detection delay, the ictal Single-Channel configuration remains relatively fast, with a latency of 13.6 seconds, only 2 seconds slower than the Dual-Channel benchmark. However, the delay increases substantially to 17.7 seconds in the no-seizure Single-Channel configuration. This trend is mirrored in the number of detected ictal events: the Single-Channel ictal configuration correctly identified 136 events, 16 fewer than the Dual-Channel configuration but 38 more than the no-seizure channel setup, which detected only 98 events.

The impact of FP deserves particular attention. Due to the heavy imbalance between the ictal and inter-ictal classes (skewed toward the latter), specificity alone fails to provide

meaningful insight into system behavior across different configurations, as it remains at SP = 1.00 in all cases. However, the evaluation of performance at the event level, specifically through the FP/h metric, allows for meaningful comparisons. Interestingly, selecting the ictal channel results in the highest FP/h of 0.21, doubling the rate observed in the no-seizure channel configuration. Meanwhile, the no-seizure Single-Channel setup achieves the lowest FP/h value (0.09), even outperforming the Dual-Channel configuration. This highlights the importance of accurately representing the inter-ictal class through a dedicated channel in order to maximize performance in terms of reliability. Finally, the performance obtained with the negative channel are the worst among the analyzed configurations (third column of Table 5.11).

The results obtained on the Siena Scalp dataset and reported in Table 5.12 confirm what was observed on CHB-MIT for both types of configuration. In particular, the Dual-Channel configuration provides the best performance, with SN equal to 0.67 and an overall bAcc of 0.83. Notably, the differences between the ictal single-channel configuration and the dual-channel configuration are not statistically significant. However, the main purpose of this comparison is not to demonstrate a large performance gap between these two settings, but rather to emphasize the importance of selecting an informative channel. Indeed, the fact that the ictal single-channel configuration achieves results close to those of the dual-channel setup suggests that the proposed selection strategy is able to identify a highly informative channel for seizure detection.

Also in this case, as seen for CHB-MIT, the performance is comparable and slightly superior to four-channel solutions (see the Table 5.8 of the Section 5.1.3), confirming the effectiveness of the proposed channel selection approach. We recall that the four fixed-channel configuration includes the temporal channels F_7-T_7, T_7-P_7, F_8-T_8, and T_8-P_8. Compared to FP/h equal to 0.39 with 31/36 correctly identified ictal events in the two-optimized channels configuration, the four-channels configuration detects 28/36 events with FP/h = 0.55, exhibiting lower reliability (more false alarms) and a 7% - point decrease in SN (0.60 vs. 0.67).

Although Single-Channel configurations show lower SN than Dual-Channel configurations, they present significant differences between ictal and no-seizure channels. In fact, the ictal configuration reaches SN = 0.64 with bAcc = 0.81, while the no-seizure configuration stops at SN = 0.55 with bAcc = 0.77. In both cases, SP remains high (SP = 0.99), which should be interpreted in light of the strong imbalance between classes, as with CHB-MIT.

In terms of detection delay, the two main single-channel configurations are very close to the Dual-Channel configuration: 14.0 s (ictal) and 13.83 s (no-seizure), compared to 13.5 s for both for both multi-channel setups. In terms of recognized events, the ictal channel detects 29/36 seizures (two fewer than Dual-channel configuration), while the no-seizure channel drops to 23/36.

The performance with respect to false alarms highlights a significant trade-off. On the one hand, the choice of the ictal channel produces the highest FP/h value (0.84), while the no-seizure channel is much lower (0.39). On the other hand, the Dual-Channel configuration is more reliable overall, with an FP/h of 0.32. This result reinforces the concept highlighted above, namely that the explicit representation of the inter-ictal class through a dedicated channel is crucial for maximizing the reliability of the system.

Also for the Siena Scalp dataset, the negative Single-Channel configuration is clearly the worst among those analyzed, with SN = 0.40, bAcc = 0.66, SP = 0.93, delay = 22.2 s, and 20/36 ictal episodes detected. Furthermore, this configuration is also the only one to provide an SP not close to 1.00 despite the imbalance between classes, suggesting that the choice of the 'negative' channel, and therefore not truly representative of the no-seizure phase, fails to adequately model the absence of ictal patterns.

In conclusion, the results on both datasets confirm the effectiveness of our channel selection algorithm. Choosing signals that well represent both the non-seizure and ictal phases, avoiding redundant or less informative channels, allows the Dual-Channel configuration to offer competitive performance even compared to configurations with multiple channels.

5.3 Generalization Analysis

In the Section 5.1, we demonstrated that the use of fixed subsets of channels may not be sufficient to capture the specific location of seizures in each individual, particularly in the case of focal seizures, and that when the number of channels is reduced, their positioning must be optimized to improve detection accuracy. Therefore, adopting the paradigm of personalized medicine [133], we have designed a methodology that personalizes both the selection of the subset of channels and the training of the deep learning model, using only data acquired from a single patient [13, 14]. This strategy gives rise to a system that can be defined as '*fully personalized*'.

Although this approach guarantees a high level of optimization, it can nevertheless introduce significant challenges during the production phase. In particular, the need to adapt the configuration of the electrodes to the specific characteristics of each patient makes it clear that the device must be designed and manufactured to measure for each individual. This entails further complexities in terms of time, costs and production processes.

To systematically evaluate this problem, we considered several strategies to compare with the fully personalized approach. The proposed channel selection algorithm (see Section 4.1.3.1) operates independently of the downstream detection system. This independence allows us to explore different levels of personalization and generalisation, modulating both channel selection and detection model training.

A first option is a '*fully generalized*' approach, in which both channel selection and model training are derived from a multi-patient dataset. This solution simplifies deployment but can reduce accuracy [114]. Also intermediate solutions are possible, in which only one of the two components is customized. For example, channel selection can be based on a large multi-patient dataset, while the deep learning model is trained on a single patient's data. This would allow the use of a generic acquisition system (i.e., a standard electrode configuration), but at the same time distinct from state-of-the-art approaches [92].

The latter, in fact, involve a fixed electrode configuration,

determined mainly by the characteristics and constraints of the chosen wearable device. In our case, the approach is optimized for the specific problem to be addressed, while exploiting a model trained on patient-specific data to increase performance. We define this configuration as a '*hybrid system-specific model*'. Conversely, it would also be possible to train the model on a multiple-patient dataset while selecting EEG channels based on data from a specific patient, referred to as '*hybrid system-specific channels*'. However, this method is less practical from an implementation perspective, as electrode configuration customization is presumably the most resource-intensive.

The experiments presented in this study were the subject of the presentation "Personalized vs non-personalized medicine in wearable device-based seizure detection" at the 47th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC) [15].

5.3.1 Implementation details

In order to explore the different possibilities we may have in terms of both electrode setup (which in turn means channel selection) and model training, we analyzed the four different possible configurations, summarized in Figure 5.7.

1. **Fully Generalized System (FGS)**: both the channel selection and model training are derived from multiple-patient dataset.

2. **Hybrid System-Specific Model (HSSM)**: it is based on a generalized channel selection, derived from multiple-patient dataset, combined with a patient-specific model training.

3. **Hybrid System-Specific Channels (HSSC)**: it is based on a patient-specific channel selection, combined with a detection model trained on a multiple-patient dataset.

4. **Fully Personalized System (FPS)**: both the channel selection and model training are tailored to a single patient's data.

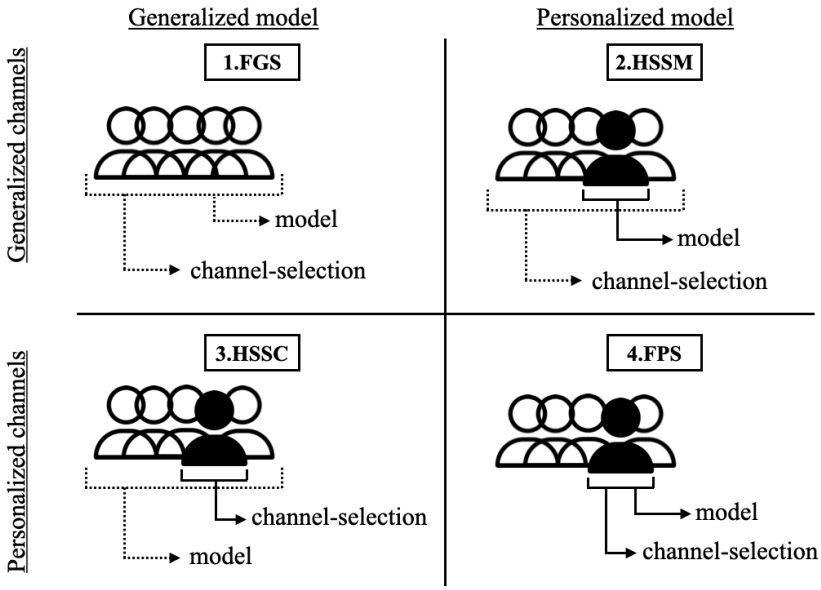


Figure 5.7: Schematic view of the four analyzed configurations, including Fully Generalized System (FGS), Hybrid System-Specific Model (HSSM), Hybrid System-Specific Channels (HSSC), Fully Personalized System (FPS).

It is important to remark that in all the cases the channel selection was carried out by using the method described before.

To select the EEG channels carrying the most relevant information for discriminating between the two defined classes, we considered the algorithm proposed in Section 4.1.3.1. The approach ensures that channels influenced by less meaningful details of brain activity are penalized in the selection algorithm.

The procedure described step-by-step in Section 4.1.3.1 enables a ranking of channels based on their total informational contribution, highlighting those that capture the most relevant variance in the EEG data. The algorithm is applied to each individual patient independently of the others.

To extend the channel selection to a more generalized approach, a variant called '*multiple patients*' was implemented. In this case, all recordings from $n-1$ patients in the dataset were combined to perform the selection process.

Validation strategy.
The datasets used have some limitations, related both to the number of patients included and to the total hours of recording available for each subject. In particular, in the case of the CHB-MIT dataset [79], adopted as the initial reference for the study, only a small number of recordings per patient contained epileptic events (up to a maximum of 14), while most of the recordings (over 100) were acquired during periods of normal behavior.

The same is valid also for the Siena Scalp dataset, in which patients present a maximum of five epileptic seizures. Even within the recordings labeled as containing epileptic seizures, the number of observable events remains relatively small. In light of these limitations, the model training, validation and testing procedures, regardless of the configuration considered, were defined taking these considerations into account.

A leave-one-out validation strategy was therefore adopted in order to maximize the use of available data and ensure a more robust evaluation of the model's performance. This is a rigorous validation strategy in the field of seizure detection as it prevents overfitting both in patient-specific and patient-independent approaches. Concretely, the models were trained on $n-1$ examples and testing on the n-th example, iterating over all n examples.
By adapting this approach for generalized and patient-specific models, we will have:

- For generalized models (i.e. those trained using data from all patients), a leave-one-subject-out (LOSO) cross-validation approach was applied. In this setup, each example corresponded to the dataset of an individual patient, meaning the model was trained and validated on $n_S - 1$ subjects and tested on the n_S-th subject.

- For patient-specific models, a leave-one-record-out (LORO) cross-validation approach was employed. Here, each example corresponded to a one-hour record, which might include or not seizure events. The model was trained and validated on $n_R - 1$ records from a single patient and tested on the n_R-th recording.

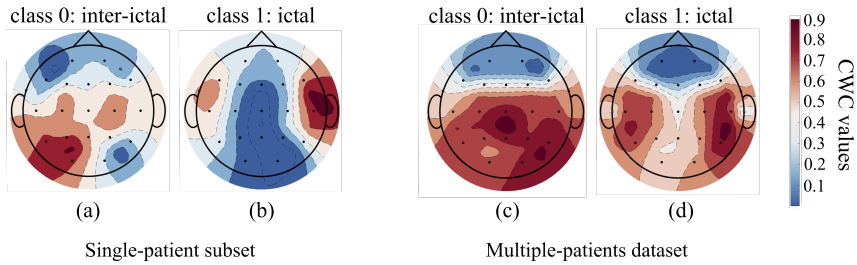


Figure 5.8: Average CWCs maps highlighting the neural activity localization during the inter-ictal (a,c) and ictal phases (b,d). In particular, (a,b) are the CWC maps derived from all the records of patient chb01, taken as example, while (c,d) are the CWC maps derived from all the records of the dataset of the CHB-MIT dataset, i.e. across all the patient

5.3.2 Results

In this section, we present the outcomes of the analyses carried out on the four analyzed configurations. Focusing on both EEG channel selection and seizure detection performance, we first describe how optimal channels were identified both for individual patients and in a generalized framework across both the CHB-MIT [79] and the Siena Scalp [80] datasets. Then, we report the performance of seizure detection models using the selected channels, comparing fully personalized, hybrid, and fully generalized configurations.

Channel selection.

Figure 5.8 shows the topographic maps relating to the two classes of interest, namely inter-ictal and ictal extracted from the CHB-MIT dataset. The maps represent the distribution of the channels according to the 10-20 system and show the average value of the CWC. The blue color indicates the lowest values, while red indicates the highest values, as indicated by the color bar shown in the right part of the figure. Specifically, Figures 5.8a and 5.8b show the mean CWC obtained for patient chb01, taken as an example, respectively in the inter-ictal and ictal phases. Figures 5.8c and 5.8d, on the other hand, show the maps in the two phases, inter-ictal and ictal, respectively, but refer to the 'generic patient', i.e., obtained by

Table 5.13: Selected channel pairs for individual patients and for the dataset including all patients (CHB-MIT dataset).

Patient	ch1	ch2	Patient	ch1	ch2
chb01	P_3-O_1	F_8-T_8	chb13	P_4-O_2	Fp_1-F_7
chb02	T_7-P_7	T_8-P_8	chb14	T_8-P_8	Fp_1-F_7
chb03	P_8-O_2	T_7-P_7	chb15	Cz-Pz	T_7-P_7
chb04	Cz-Pz	T_8-P_8	chb17	C_4-P_4	P_4-O_2
chb05	Cz-Pz	P_8-O_2	chb18	C_4-P_4	T_8-P_8
chb06	F_8-T_8	T_8-P_8	chb19	Cz-Pz	T_8-FT_{10}
chb07	Cz-Pz	P_8-O_2	chb20	P_8-O_2	T_7-FT_9
chb08	P_8-O_2	F_4-C_4	chb21	T_8-P_8	C_3-P_3
chb09	P_7-O_1	T_8-P_8	chb22	F_4-C_4	T_7-P_7
chb10	T_8-P_8	T_7-P_7	chb23	P_8-O_2	T_7-P_7
chb11	Fz-Cz	T_7-P_7	chb24	C_4-P_4	F_8-T_8
chb12	F_3-C_3	F_4-C_4	Generic	Cz-Pz	T_8-P_8

combining the data of all the patients available in the dataset.

As already described in the previous chapters, in the reported patient-specific case study it can be observed that in the inter-ictal phase (see Figure 5.8a), neural activity is concentrated in the parieto-occipital area, specifically around the region covered by the P_3-O_1 channel. During the ictal phase, however, the activity shifts to the opposite side of the brain, focusing on the region covered by the F_8-T_8 channel (see Figure 5.8b). Extending the analysis to all patients, when the algorithm is applied to each patient using the LORO approach, the patient-specific optimized channel pairs are reported in Table 5.13 for CHB-MIT and in Table 5.14 for the Siena Scalp dataset.

A key observation from the application of the LORO scheme is the consistent identification of the same channel for each phase, regardless of the record excluded from the computation, with only a few specific exceptions. These exceptions involved patients with a limited number of records containing ictal events or those for whom the algorithm identified multiple

Table 5.14: Selected channel pairs for individual patients and for the dataset including all patients (Siena Scalp dataset).

Patient	ch1	ch2	Patient	ch1	ch2
PN00	F_3-C_3	F_8-T_8	PN12	T_8-P_8	P_7-O_1
PN01	Cz-Pz	F_7-T_7	PN13	F_8-T_8	T_7-P_7
PN03	C_4-P_4	F_8-T_8	PN14	P_8-O_2	C_3-P_3
PN05	T_7-P_7	T_7-FT_9	PN16	F_7-T_7	T_7-P_7
PN06	F_3-C_3	T_7-P_7	PN17	T_7-P_7	T_8-P_8
PN09	F_3-C_3	T_8-P_8	Generic	Cz-Pz	T_7-P_7

channels per phase, necessitating stronger discrimination based on correlation. In such cases, majority voting was employed to determine the best pair of channels.

On the other hand, the generalized approach, based on the data of all patients, leads to an inter-ictal map (Figure 5.8c) that highlights widespread activity over most of the scalp, at exception of frontal regions, with a notable peak in the central region covered by channel Cz-Pz.

The same pattern is observed for the Siena Scalp dataset. Indeed, Figure 5.9, which shows only the multi-patient configuration to provide an alternative example of the generalized map on a different dataset, clearly highlights (in the left head map (a)) the same area with a central peak at Cz-Pz. Although the area is more localized, probably due to the smaller number of patients involved, the fact that the same region emerges in inter-ictal conditions in both datasets allows for a solid clinical interpretation. It is known, in fact, that the central region of the scalp overlies the sensorimotor cortex and the posterior parietal cortex, areas that belong to the central networks of the resting state. These regions contribute to the continuous and reproducible background activity and therefore represent a robust generalized baseline condition [172].

Processing the data collected during the ictal phases of all patients results in a map (Figures 5.8d for CHB-MIT dataset and 5.9b for Siena Scalp dataset) that shows strong neural activity in the left and right temporal regions, consistently

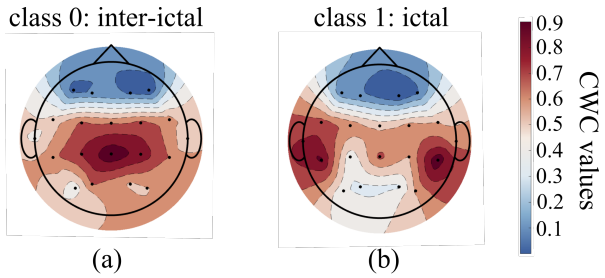


Figure 5.9: Average CWCs maps highlighting the neural activity localization during the inter-ictal (a) and ictal phases (b). Both CWC maps derive from all the records of the Siena Scalp dataset, i.e. across all the patient.

with previous studies that associate temporal lobe activity with seizure onset and propagation [173]. The asymmetry with more pronounced activity on the right for CHB-MIT dataset and on the left for the Siena Scalp, might reflect lateralization patterns typical of specific seizure dynamics within the dataset.

It is important to underline that the LOSO scheme also exhibited consistent channel selection across iterations through the subjects. This consistency underscores the robustness and reliability of the algorithm, ensuring that individual patient variability does not impact the channel selection. Clearly, increasing the dataset size and including a more diverse patient population could help generalize these findings further and reduce any potential bias.

In this generalized-channel setting, as also reported in the last row in the right column of both Tables, the selected pair results in Cz-Pz (inter-ictal) and T_8-P_8 (ictal) for the CHB-MIT dataset (see Table 5.13), and in Cz-Pz (inter-ictal) and T_7-P_7 (ictal) for the Siena Scalp dataset (see Table 5.14).

Performance evaluation

Once the optimized channel pairs were identified on both a generalized and patient-specific basis, they were fed into the detection model, and the performance was evaluated using the metrics introduced in Section 4.3.

Table 5.15: System performance for the two-optimized channel configuration on the CHB-MIT dataset.

Metrics	FGS	HSSM	HSSC	FPS
SP	0.99 ± 0.03	1.00 ± 0.01	0.94 ± 0.19	1.00 ± 0.01
SN	0.27 ± 0.33	0.60 ± 0.34	0.40 ± 0.38	0.67 ± 0.30
bAcc	0.63 ± 0.18	0.80 ± 0.17	0.67 ± 0.19	0.83 ± 0.15
Delay [s]	23.1 ± 15.6	11.9 ± 8.74	18.6 ± 16.4	11.5 ± 9.7
FP/h	1.17 ± 2.70	0.16 ± 0.19	1.77 ± 4.27	0.15 ± 0.18
Seizures	86 / 173	143 / 173	110 / 173	152 / 173

Table 5.16: System performance for the two-optimized channel configuration on the Siena Scalp dataset.

Metrics	FGS	HSSM	HSSC	FPS
SP	0.94 ± 0.12	0.97 ± 0.08	0.95 ± 0.15	0.99 ± 0.03
SN	0.30 ± 0.32	0.50 ± 0.38	0.39 ± 0.29	0.67 ± 0.32
bAcc	0.62 ± 0.16	0.73 ± 0.18	0.67 ± 0.15	0.83 ± 0.16
Delay [s]	21.1 ± 13.9	19.1 ± 9.57	24.2 ± 13.5	13.5 ± 6.40
FP/h	0.48 ± 1.53	0.34 ± 0.13	0.68 ± 1.74	0.39 ± 0.42
Seizures	21 / 36	29 / 36	26 / 36	31 / 36

The analysis, as shown in Tables 5.15 and 5.16, revealed notable differences across the evaluated configurations, particularly in terms of sensitivity (SN), balanced accuracy (bAcc), detection delay and false positive per hour (FP/h), as better discussed in the Section 6.3.

5.3.3 Discussion

The results presented in the previous section highlight clear differences between generalized, hybrid, and fully personalized approaches to EEG channel selection and seizure detection.

First of all, it is important to highlight that while the generalized approach captures common patterns, it may not fully

account for patient-specific variations. As can be noted from both Tables 5.13 and 5.14, some subjects exhibit ictal or interictal activity distributions that differ from the generalized ones. Analyzing the influence of these factors is indeed one of the main aims of this work, leading us to critically analyze and compare different generalized, hybrid, and personalized configurations.

In the following discussion, we analyze how these variations impact detection performance, the trade-offs between model complexity and wearable implementation, and the practical implications for real-time, minimal-channel EEG monitoring across both datasets used in this study.

For the CHB-MIT dataset, consistent with the above observations, the FPS (fourth column of Table 5.15) outperformed other approaches, achieving the highest SN and bAcc, alongside the lowest FP/h and detection delay. This highlights the advantage of tailoring both the channel selection and the detection model to individual patients, as it allows the model to better capture unique signal patterns associated with ictal events. Specifically, the FPS configuration achieved a SN of 67% and a bAcc of 83%, with a low FP/h of 0.15, demonstrating its reliability in accurately detecting seizures with minimal false alarms. The test phase demonstrated the capability of detecting 152 out of 173 ictal events.

In contrast, the FGS (in the first column of Table 5.15) even achieving perfect SP, due to the imbalance of the dataset, it struggled with SN, which was notably lower at 27%, reflecting in the detection of 86/173 ictal events, i.e. almost half with respect to the FPS. This indicates that FGS fails to detect a significant portion of seizures, due to the lack of patient-specific optimization. Moreover, the average detection delay is about 23 seconds, the highest among the analyzed configurations, while the FP/h is about 1.17, i.e. slightly more than one false alarm per hour.

As expected, the hybrid configurations (HSSM and HSSC) demonstrated intermediate performance. HSSM showed better SN and a lower FP/h compared to HSSC, indicating that a fixed pair of optimized channels, combined with a patient-specific detection model, may provide a better trade-off be-

tween seizure detection performance and ease of deployment. Notably, HSSM exhibited a detection delay of approximately 13 seconds, which is close to that of FPS. In terms of false alarms, HSSM delivered promising results with a FP/h of 0.17, close to the 0.15 achieved with FPS. Interestingly, HSSM detected 130 ictal events, slightly fewer than twice the events captured by FGS, and approaching the 152 events detected by FPS. Conversely, HSSC demonstrated the poorest performance in both SP (0.94) and FP/h (reaching almost 2), with unsatisfactory results in terms of bAcc (0.7) reflecting in a low number of detected ictal events (110/173). The detection delay of 18s, although lower than that obtained with the FGS, is still non-negligible. In any case, it is worth noting that the HSSC configuration, although included in this analysis for completeness, is the least favorable from an implementation perspective. When patient-specific data are available, allowing for channel personalization, it becomes less meaningful to rely on a generic model. Utilizing individualized data to tailor channel selection and detection models offers a significant advantage, ensuring more accurate and efficient seizure detection, as demonstrated by FPS.

The same described trends hold also on Siena Scalp as shown in Table 5.16. The FPS again delivers the best overall performance, with the highest SN (0.67) and bAcc (0.83), while keeping a short detection delay (13.5 s) and a moderate FP/h (0.39). It correctly detected 31/36 ictal events. The hybrid approach with subject-specific model (HSSM) provides a good compromise, since it improves markedly over generalized channels (SN of 0.50 vs. 0.30 for FGS), achieving bAcc of 0.73 and keeping delay close to FPS (19.1 s). This is important for real-time use also because it shows the lowest false-alarm rate among all configurations (FP/h of 0.34) and detect 29/36 events. In contrast, the generalized configuration (FGS) attains lower SP (0.94) and SN (0.30), capturing only 21/36 seizures, while the hybrid configuration with a generic classifier (HSSC) remains the least attractive. Indeed, despite acceptable SP (0.95), it shows lower bAcc (0.67) and the longest delay (24.2 s) with 26/36 ictal detections.

To conclude, the distinctive feature of the channel general-

Table 5.17: System performance for the four-fixed channel configuration on the CHB-MIT dataset.

Metrics	FGS	HSSM
SP	0.98 ± 0.11	1.00 ± 0.01
SN	0.26 ± 0.35	0.66 ± 0.32
bAcc	0.65 ± 0.19	0.83 ± 0.16
Delay [s]	21.5 ± 17.2	11.4 ± 8.0
FP/h	1.18 ± 3.42	0.14 ± 0.25
Seizures	62 / 173	149 / 173

ization proposed in this study lies in the fact that, although the same channels are used for all patients, as is commonly done in most existing studies, these channels are selected on the basis of a straightforward and interpretable data-driven method. In particular, many studies focus on temporal channels, as these are more easily accessible through wearable devices such as e-glasses [86, 103]. Specifically, several recent works employ a fixed set of four channels, namely T_7, T_7-P_7, F_8-T_8 and T_8-P_8 [92, 163, 85]. This configuration includes T_8-P_8, which our algorithm selects for the ictal phase on CHB-MIT, and T_7-P_7, selected for the ictal phase on Siena Scalp; however, it excludes the central region Cz-Pz that instead emerges in the inter-ictal phase in both datasets.

Performance comparison with the state-of-the-art

In order to enable a comparison with existing literature approaches, Table 5.17 reports the results obtained using a configuration based on the four fixed temporal channels F_7-T_7, T_7-P_7, F_8-T_8, T_8-P_8 introduced in Section 5.1.2. Consistently with the previous analysis, the results are presented for the model trained on a multiple-patient dataset (LOSO approach), referred to as FGS, and for the model trained on patient-specific datasets (LORO approach), referred to as HSSM. Since the set of channels is fixed *a priori*, the FPS and HSSC configurations, which involve channel personalization, were not considered in this experiment.

Table 5.18: System performance for the two HSSM configurations, one temporal pair and one central-temporal pair, on the CHB-MIT dataset.

Metrics	HSSM	
	config. T_8-P_8, Cz-Pz	config. T_7-P_7, T_8-P_8
SP	1.00 ± 0.01	1.00 ± 0.01
SN	0.60 ± 0.34	0.55 ± 0.34
bAcc	0.80 ± 0.17	0.77 ± 0.17
Delay [s]	11.9 ± 8.74	12.33 ± 8.31
FP/h	0.16 ± 0.19	0.18 ± 0.22
Seizures	143 / 173	132 / 173

Interestingly, the proposed two-optimized channels FGS outperform the four temporal channel configuration in all considered metrics, despite using fewer signals, except for the detection delay, which is approximately 2 seconds lower in the four-channel case. The improvement is particularly evident in the number of detected seizures, which decreases from 86/173 detected with two channels to 62/173 detected with four channels. This result highlights the importance of optimizing the channel subset, even in a generalized approach.

On the other hand, increasing the number of channels in the temporal regions leads to improved results when the model is tailored to a specific patient. Indeed, the resulting hybrid system based on four temporal channels outperforms the hybrid system based on two generalized channels (one central and one temporal). This can be explained by the fact that, as highlighted by our algorithm, the channels that emerge most prominently during ictal phases are mainly temporal (see Figure 5.8d.), meaning that increased coverage in these regions enhances the likelihood of seizure detection. However, this improvement comes at the cost of a higher number of electrodes, and consequently, a larger wearable systems, and in any case does not provide a significant performance gain compared to the FPS based on two channels.

For the sake of completeness, to verify whether two temporal channels could match the generalized pairing, we also evaluated the pair (T_7-P_7, T_8-P_8), i.e., the two channels that emerged in the ictal phase of the 'generic patient' model for both datasets. As shown in Table 5.18, this temporal pair performs worse than our optimized central-temporal pair, highlighting the benefit of the proposed coupling. This result reinforces the idea that the Cz-Pz channel provides a robust and subject-independent baseline, capable of integrating the most representative temporal channel appeared during the ictal phase. On the contrary, replacing it with a second temporal channel introduces redundancy and reduces cross-subject stability.

5.4 Conclusion

This chapter presents an in-depth analysis of all the experiments carried out to evaluate the robustness of the EEG channel selection algorithm introduced in Section 4.1 and to verify the feasibility of wearable systems for seizure detection with minimal hardware, balancing accuracy and practicality. Unlike conventional multi-channel EEG systems, which are often impractical for continuous monitoring, our system identifies just two personalized channels optimized for each patient's unique seizure patterns. This advancement aligns with the ongoing trend of minimizing the number of electrodes in non-invasive EEG systems, providing a robust, data-driven framework for optimized electrode placement.

The performance of the proposed system was validated through extensive experimentation on two publicly available datasets, demonstrating its effectiveness. With a balanced accuracy of 0.83 (and an impressive accuracy of 0.99) on the CHB-MIT dataset evaluated at the segment level, our system performs comparably to or even exceeds the performance of existing multi-channel systems based on four-channel configuration designed for wearable devices. Notably, our two-channel system detected 152 out of 181 ictal events, outperforming the four-channel configuration trained and tested with the same

method, which detected 149 ictal events. The same can be observed for the Siena Scalp dataset, where the two-channel configuration model detects 31/36 seizures compared to 28/36 with four channels and improves balanced accuracy by 5 percentage points. In light of these elements, two follow-up studies were also validated. One aimed to show that not all channels are equally informative, while the other aimed to evaluate the generalizability of the system in multi-patient scenarios, analyzing the feasibility of wearable devices that can be used by multiple patients.

As for the Single-Channel study, the analysis identifies for each patient the most informative single channels for the ictal and non-seizure phases, as well as the least informative negative channel across both phases. The results obtained on both datasets confirm the robustness and the meaningfulness of the EEG channel-selection algorithm, showing that among the Single-Channel configurations, the negative channel delivers the weakest performance across all evaluation metrics. This outcome, expected by design, highlights the effectiveness of the selection algorithm: not all channels carry informative content, and channels dominated by artifacts or noise should be avoided. The Single-Channel (no-seizure) setup yields the lowest false-positive rate, even outperforming the Dual-Channel configuration on this metric, though it does so at the cost of reduced sensitivity. This is plausibly because a channel representative of baseline activity triggers fewer detections overall, both true and false. Especially relevant is the Single-Channel configuration, which selects, for each patient, the most informative ictal-phase channel. Although it achieves lower performance than the Dual-Channel setup, which combines one channel representative of the ictal phase and another of the non-seizure phase, it offers a solid trade-off between performance and complexity, suggesting that the number of EEG channels can be further reduced without materially compromising accuracy.

With regard to the generalization study, the analysis carried out reinforces the importance of patient-specific optimization in seizure detection systems. The experiments showed that, although the fully personalized system configuration represents the solution with the best performance in terms of detection

accuracy and reliability, making it highly promising for clinical implementation, hybrid methods are a practical alternative. In particular, systems that maintain a fixed electrode configuration but customize the detection model provide an effective balance between ease of implementation and performance quality. While not achieving the accuracy levels of fully customized systems, these approaches still offer significant results in terms of reducing false positives per hour and latency, making them particularly suitable for daily monitoring applications.

Taken together, these results validate the channel selection framework guided by interpretability. The channels chosen to represent the ictal and non-ictal phases are confirmed as the most informative for discriminating epileptic seizures, while negative channels (inactive in both phases) are systematically uninformative. The principle of personalization remains central, even if the second follow-up study shows that, in generalized contexts, including a channel representative of the non-ictal phase is crucial. The study also discuss how an hybrid setup (HSSM: Hybrid System-Specific Model) with two fixed "generalized" channels and a patient-specific trained model represents a solid compromise when channel optimization valid for multiple patients is desired without sacrificing good operational performance. This latest result opens the path to future developments aimed at further reducing the gap between generalization and personalization, for example through transfer learning and domain adaptation (between patients and/or datasets).

Chapter 6

Validation of GNN-based approach

The second proposed framework of this thesis is to appropriately quantify the changes in brain connectivity caused by seizure disorders and to model full-scalp EEG from an ease-of-use and resource-constrained graph perspective. In particular, we propose a dynamic spatio-temporal graph convolutional network consisting of two main components: a graph convolutional network (GCN) to model spatial relationships among channels, and a Long Short-Term Memory (LSTM) module to capture temporal dependencies, jointly enabling accurate detection of seizure windows. The approach explicitly models the functional connectivity (FC) among channels, estimated through phase synchrony, to extract robust and clinically meaningful spatial patterns, while maintaining a compact architecture suitable for deployment in low-resource scenarios.

Furthermore, interpretability is ensured through the integration of GNN-Explainer, which highlights the most relevant channels during seizure detection, thereby enabling clinicians to visually assess and validate the model's focus.

6.1 Implementation details

Similar to the results described for the validation of reduced-channel framework (see Chapter 5), we perform GNN-based seizure detection experiments on the CHB-MIT public EEG

dataset [79] and additionally validate the method on the Siena Scalp dataset [80]. The specifications of both datasets are reported in Section 3.2. For each patient, recordings are available under baseline conditions and during critical events. The latter are annotated with the start and end times of the ictal event, providing the ground truth for the binary classification problem performed in this study.

The training procedure was divided into two phases. In the first phase, the architecture shown at the top of Figure 4.8 was considered, with the aim of classifying individual EEG windows as ictal (label 1) or 'non ictal' (label 0). To ensure reproducibility, the seed was set to 42 in order to standardize all training procedures. Optimization was performed using the Adam algorithm, with a learning rate of 5×10^{-5} and L2 penalty (weight decay) of 10^{-4}. The loss function chosen was Cross Entropy Loss. The model was trained for a maximum of 250 epochs, with an early stopping mechanism characterized by a patience of 10 epochs and a minimum improvement threshold of 0.0001. At this stage, the data loader was built from single one-second windows, with a batch size of 256. The dataset was divided into training, validation, and test sets using the cross-validation scheme.

In the second phase, three consecutive graphs (each corresponding to a one-second time window) were provided as input to the network. Each graph was processed by the ChebNet backbone, which produced a vector of 64 features. The sequence of the three vectors thus obtained constituted the input for an LSTM network with 32 hidden units. The LSTM layer modeled the temporal dependencies between the three graphs and provided a compact representation of the sequence as output. This representation was then projected, via a fully connected layer, into the binary decision space.

In this second stage, the model was trained for a maximum of 300 epochs, using the Cross Entropy Loss function and the Adam optimizer, with a learning rate of 10^{-3} and weight decay of 10^{-4}. In this case, an early stopping mechanism was also applied with a patience of 10 epochs and a minimum improvement threshold of 0.001. The dataloader used in the second phase was different from the first, as it was constructed from

sequences of three consecutive windows. In this case, a batch size of 256 was also adopted, maintaining the same strategy of division into training, validation, and test sets described in the first phase.

It should be noted that since the ChebNet backbone was reused within the second network, the mean and standard deviation calculated on the training set of the dataset used for pre-training the feature extractor were subsequently used to normalize the training, validation, and test sets of the sequence-based dataset. The neural architecture adopted was implemented using the PyTorch framework (version 2.1.2) with CUDA 12.1 support. All training and validation procedures were performed on a workstation equipped with an NVIDIA Tesla V100 DGX GPU. In addition to PyTorch, the Networkx libraries and PyTorch Geometric were used to manipulate, create, and study the dynamics, structure, and functions of complex networks.

Cross-validation scheme.

The following classification metrics were used to evaluate the proposed ST-GCN model: accuracy, sensitivity, and specificity. Since the dataset is unbalanced, despite applying a balance to ensure a 1:4 ratio between ictal and non-ictal windows, the main evaluation metric adopted is balanced accuracy.

Again in the context of personalized medicine [133], the aim is to obtain a specific model for each patient. To this end, the K-fold cross-validation scheme has been adopted, which is widely used in the literature, particularly in AI applications to EEG, where the available data are often limited. In this scheme, the dataset is divided into K partitions (folds) in turn, each partition is used as a test set, while the remaining K-1 constitute the training set [174]. The experiments are then repeated K times, changing the training and test partitions each time. At each iteration, the validation set was extracted from the training set, preserving the balance between classes, with a distribution of 70% for training and 30% for validation.

This procedure reduces bias, as performance is estimated on K different subdivisions of the dataset rather than on a single partition. At the end, the final metrics used to evaluate the model's performance were calculated as the average

across the K folds. The choice of this validation procedure, as an alternative to the leave-one-record-out scheme adopted in previous chapters, is motivated by the fact that the works in the literature with which this study compares itself use K-fold cross-validation [112, 150, 149]. In this study, we set $K = 5$.

6.2 Results

All experiments carried out in this study adopted a patient-specific approach, with the aim of reinforcing the concept of personalized medicine [133]. Therefore, performance was evaluated for each patient and subsequently averaged to ensure the overall validity of the study. The first results concern the comparison between PCC and PLV, while the results obtained to validate the application of the lightweight ST-GNN network are presented below, highlighting its performance against state-of-the-art methods.

6.2.1 Connectivity results

As a first step, we compared FC measures following the procedure described in Section 4.2.3.1, in order to determine which metric is most suitable for building graphs in the context of seizure detection. Without loss of generality, we present in detail the results for patient chb01 from the CHB-MIT dataset [79], as an illustrative example of the procedure used to evaluate the FC measure for seizure detection. Comparable results were consistently observed across all patients in both datasets.

The boxplots shown in Figure 6.1 compared the strength of connections estimated with PLV and PCC in non-ictal and ictal conditions for subject chb01. With PLV, there is a clear shift towards higher values during the ictal phase (higher median with 0.44 vs 0.36 values) compared to the non-ictal phase, with limited overlap between the two distributions. In the case of the PCC measure, the increase during the ictal phase is less visible and the distributions overlap extensively, as well as showing greater dispersion and numerous outliers.

Additionally, Table 6.1 reports the main descriptive statistics of both connectivity matrices averaged over the time win-

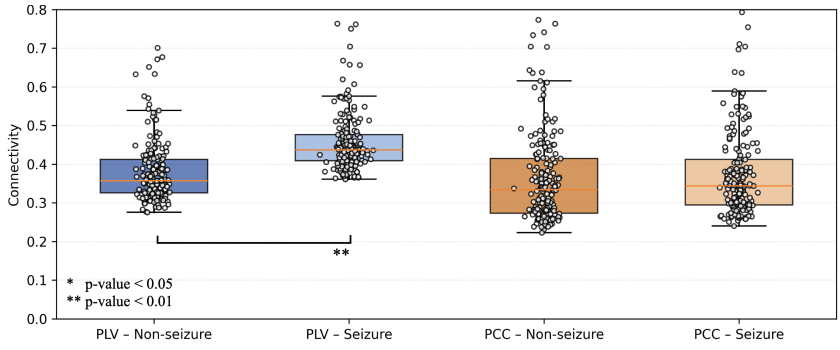


Figure 6.1: Boxplots of PLV and PCC values during seizure and non-seizure phases for patient chb01 from the CHB-MIT dataset. Dots represent the single FC values computed between each electrode pair. Symbol * denotes statistical significance.

Table 6.1: Descriptive statistics of average PLV and PCC connectivities over time in ictal and non-ictal conditions.

Condition	Minimum	Maximum	Mean	Median	Std. Dev.
PLV ictal	0.361	0.764	0.456	0.437	0.073
PLV non-ictal	0.276	0.702	0.379	0.357	0.080
PCC ictal	0.240	0.809	0.375	0.344	0.114
PCC non-ictal	0.223	0.774	0.362	0.334	0.116

dows and used for the computation of the boxplots. From a statistical point of view, to establish the ability of each FC measure to discriminate ictal states, we performed comparisons with the Mann-Whitney U test (or Wilcoxon signed-rank test). The results showed a significant difference between PLV values in the ictal and non-ictal phases (p-value < 0.05), while no significance emerged between PCC measures in the same context of application (p-value ≥ 0.05). This confirms that PLV better discriminates between the two phases with more pronounced synchronization in EEG signals during the ictal event. A summary of the statistical results for all patients in the CHB-MIT dataset is provided in Table 6.2 while the corresponding results for the Siena dataset are reported in Table 6.3. This approach is appropriate for capturing the average trend between the two distributions.

Table 6.2: Statistical comparison of PLV and PCC between ictal and non-ictal conditions across patients from the CHB-MIT dataset. The table reports the p-values obtained from the distributional comparisons and the number of significant connections detected for each FC measure using the FDR-corrected Mann-Whitney U test to control for multiple comparisons.

Patient	p-value PLV	p-value PCC	# Sig. Conn.	
			PLV	PCC
chb01	2.534e-27	6.352e-02	179	127
chb02	1.148e-16	6.631e-02	145	103
chb03	1.011e-03	3.128e-02	121	112
chb04	7.485e-16	7.725e-01	150	111
chb05	9.529e-03	6.078e-01	132	115
chb06	1.889e-17	4.009e-15	132	120
chb07	1.531e-23	7.784e-02	175	140
chb08	3.775e-49	8.546e-02	190	175
chb09	2.228e-28	6.153e-01	170	129
chb10	2.596e-11	4.423e-02	145	104
chb11	4.822e-31	3.155e-04	177	151
chb12	1.070e-16	3.994e-02	170	164
chb13	9.061e-30	1.759e-01	173	140
chb14	1.587e-14	9.880e-02	165	144
chb15	4.921e-01	7.266e-05	158	132
chb16	6.073e-02	1.406e-03	151	125
chb17	5.328e-04	1.289e-01	148	118
chb18	2.189e-11	7.431e-04	168	139
chb19	8.761e-09	5.569e-03	162	137
chb20	6.825e-01	2.047e-02	142	128
chb21	1.126e-05	3.482e-02	159	136
chb22	4.719e-08	2.409e-05	167	142
chb23	3.589e-04	8.763e-01	147	129
chb24	1.805e-02	5.377e-01	141	126

Table 6.3: Statistical comparison of PLV and PCC between ictal and non-ictal conditions across patients from the Siena Scalp dataset. The table reports the p-values obtained from the distributional comparisons and the number of significant connections detected for each FC measure using the FDR-corrected Mann-Whitney U test to control for multiple comparisons.

Patient	p-value PLV	p-value PCC	# Sig. Conn.	
			PLV	PCC
PN01	1.868e-29	5.734e-09	157	132
PN03	1.796e-03	2.015e-02	143	135
PN05	1.235e-02	2.771e-01	119	119
PN06	1.048e-26	2.313e-19	176	166
PN09	2.111e-10	3.781e-01	133	132
PN10	1.650e-07	6.031e-02	143	136
PN12	1.367e-03	1.014e-01	135	131
PN13	6.483e-08	1.969e-01	118	109
PN14	1.271e-16	4.886e-01	140	106
PN16	1.293e-07	3.080e-01	160	122
PN17	7.331e-13	7.694e-02	135	125

To directly compare ictal and non-ictal values across all time windows for each channel pair, we employed the Mann-Whitney U test with FDR correction to account for multiple comparisons. In detail, for patient chb01, the analysis revealed 179 connections that were significantly different between ictal and non-ictal conditions for the PLV measure, whereas only 127 significant differences were observed for the PCC measure. A higher number of significant connections for PLV compared to PCC was consistently found across all patients analyzed in both the CHB-MIT (see Table 6.2) and the Siena Scalp datasets (see Table 6.3). These findings confirm that PLV provides a more effective discrimination between the two phases.

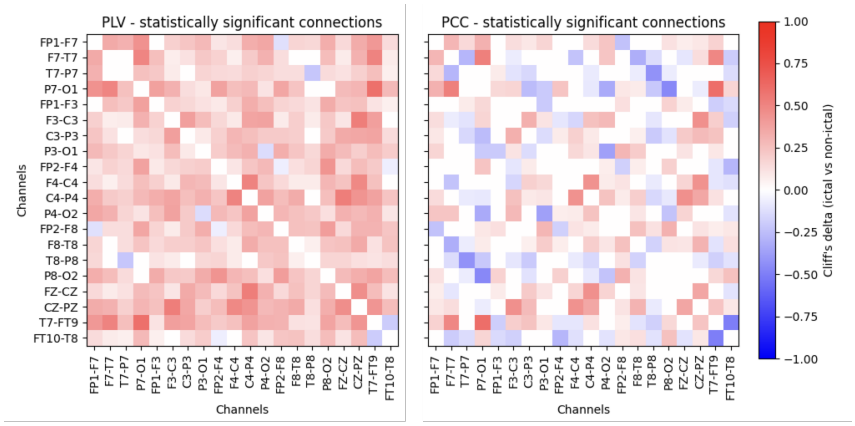


Figure 6.2: Significant FC connections for PLV (a) and PCC (b) according to Cliff's effect size (δ) between ictal and no-seizure phases. Only channel pairs with statistically significant differences are shown. The color red ($\delta > 0$) indicates stronger connectivity in ictal; the color blue ($\delta < 0$) indicates stronger connectivity in no-seizure.

We also estimated the effect size using Cliff's δ, which provides an additional aid in understanding the differences between ictal and no-seizure connectivity patterns. In this regard, Figure 6.2 shows the insights achieved by the analysis. Cliff's δ ranges between -1 and 1. A value of $\delta = -1$ indicates that, across all windows included in the statistical analysis, all ictal values are lower than non-ictal values, while $\delta = 1$ corresponds to the opposite situation, where all ictal values are higher than non-ictal values. A value of $\delta = 0$ indicates a complete overlap of the distributions, i.e., no differences between the FC measures in the ictal and non-ictal conditions. Only those channel pairs that exhibit statistically significant differences between the two phases (with p-value < 0.05), as determined by the Mann-Whitney U test with FDR correction for multiple comparisons, are shown with colors different from white. Positive values (red) indicate stronger connectivity during the ictal phase, whereas negative values (blue) indicate stronger connectivity in the non-ictal phase.

Actually, Figure 6.2 confirms what has already been observed in the boxplots (Figure 6.1), i.e., that FC increases dur-

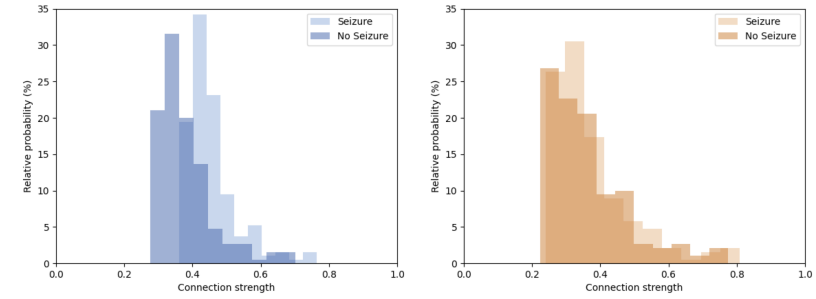


Figure 6.3: Relative probability distributions of connection strength for the two FC measures, i.e. PLV (a) and PCC (b). The histograms show the percentage of connections falling within each bin interval under ictal (light color) and non-ictal (dark color) conditions.

ing the ictal phases. In the PLV matrix, red values are prevalent, indicating greater connectivity during the ictal phase. The PCC matrix, on the other hand, shows a more heterogeneous trend, with both red and blue values, indicating the lower discriminatory power of this measure compared to PLV. Furthermore, fewer significant connections appear in the PCC matrix, as evidenced by the greater presence of white values.

Threshold results

The two FC measures were also compared considering the distribution of connection strength in ictal and non-ictal conditions. Figures 6.3a and 6.3b show the normalized histograms obtained with PLV and PCC, respectively. These histograms represent the relative probability (%) of connection strengths falling within each bin interval. In detail, Figure 6.3a shows that in ictal conditions (light blue) the distribution is shifted towards slightly higher values than in the non-ictal condition (dark blue). Connections tend to concentrate between 0.4 and 0.5 during ictal phase, while in the non-ictal state they are mainly between 0.3 and 0.4. This trend reinforces the idea of an increase in phase synchronization during epileptic seizures. On the other hand, in Figure 6.3b, the two conditions show greater overlap and are less distinguishable.

Table 6.4: Threshold τ value selected for each patient of the CHB-MIT dataset.

Patient	Threshold τ	Patient	Threshold τ	Patient	Threshold τ
chb01	0.4	chb09	0.3	chb17	0.3
chb02	0.3	chb10	0.3	chb18	0.3
chb03	0.3	chb11	0.3	chb19	0.3
chb04	0.3	chb12	0.3	chb20	0.3
chb05	0.2	chb13	0.3	chb21	0.3
chb06	0.2	chb14	0.3	chb22	0.3
chb07	0.3	chb15	0.3	chb23	0.3
chb08	0.4	chb16	0.3	chb24	0.4

Table 6.5: Threshold τ value selected for each patient of the Siena Scalp dataset.

Patient	Threshold τ	Patient	Threshold τ
PN00	0.4	PN12	0.3
PN01	0.3	PN13	0.3
PN03	0.3	PN14	0.3
PN05	0.2	PN16	0.3
PN06	0.3	PN17	0.3
PN09	0.3		

Both distributions are concentrated at low values (0.2-0.4), with longer tails towards high values, but without a clear separation between ictal and non-ictal states. The relative probability in this context is useful as it provides information on the probability of observing a given range of connection strength during the two different phases. For example, setting a threshold at 0.4 for PLV measure, we observe that ictal connections exceed this value more frequently than non-ictal connections. To demonstrate this, we applied ST-GNN to sparse graphs obtained with different PLV threshold values and observed that

the performance of the model, in terms of bAcc, varies depending on the selected threshold. In particular, as shown in Figure 6.4, the threshold value τ of 0.4 allows us to obtain the highest bAcc (0.988), confirming the effectiveness of this value in maximizing the model's discriminatory power. The patient-specific τ used in all patients are summarized in Table 6.4.

This result is explainable. In the case of PLV, the distributions for the ictal and non-ictal conditions are more distinct but concentrated in a range of medium-high values. This implies that a threshold that is too low tends to include almost all connections, reducing discriminatory power, while a threshold that is too high risks eliminating important functional edges. Conversely, in the case of PCC, the distributions of the two conditions appear more similar and shifted towards low values, making the construction of the graph less sensitive to the chosen threshold. For this reason, it becomes necessary to increase the complexity of the GNN network to be applied, in order to compensate for the lower discriminative power and improve performance.

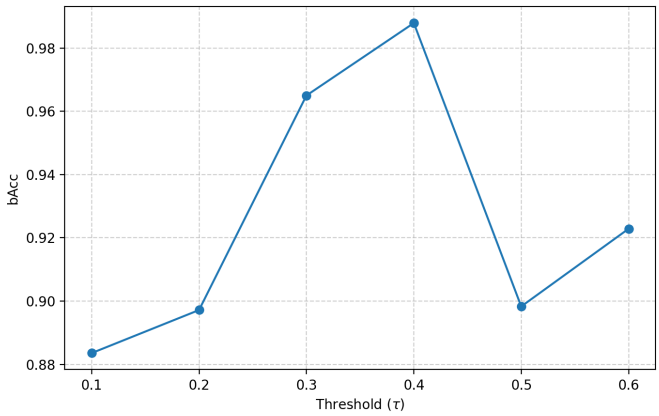


Figure 6.4: Balanced accuracy of the proposed ST-GCN model for patient chb01 as a function of the threshold τ, evaluated for values ranging from 0.1 to 0.6.

6.2.2 Classification performance

Once we evaluated the most suitable connectivity measure and an appropriate threshold for each patient, we evaluated the performance of our lightweight ST-GCN in a subject-dependent setting. Table 6.6 presents the results obtained using the evaluation metrics described in Section 4.3. The last two rows show the two variants of our model, in which the graph is constructed from FC matrices estimated with PCC or PLV. This comparison is to verify, in light of the a priori study, that the performance of the ST-GCN network confirms the greater effectiveness of PLV compared to PCC. In fact, with PCC, the model achieves Acc of 96.16%, SP of 98.45% and SN of 87.90%, with a bAcc of 93.13% and approximately 161K parameters.

On the other side, with PLV, SN increases significantly (up to 96.56%) while SP remains high (98.77%), leading to an overall improvement with Acc of 98.49% and bAcc of 97.66%, with again 161K parameters. To ensure a fair and consistent comparison, for both FC measures, we applied the patient-specific thresholds described in Table 6.4. In particular, in the case of PCC, the results proved to be insensitive to the threshold value (i.e., varying the threshold yielded the same result).

In the same Table 6.6, we also compare our model with the state-of-the-art GNN-based approaches proposed in the literature. The performance of existing methods is generally high, with Acc of 95% but they differ both in the balance between SP and SN and, above all, in the number of parameters required. Compared to the GCN-BiGRU of Xu *et al.* [112] (bAcc 97.34%, 730K parameters), our PLV-based model attains a slightly higher bAcc (97.66%) with about 78% reduction of parameters. It also delivers higher SP (98.77% vs 95.83%) but with lower SN (96.56% vs 98.85%). On the other hand, lighter methods [147, 150, 112] in the first rows of the Table reduce complexity but rely on static graphs, limiting their ability to capture dynamic variations in connectivity. It is worth pointing out that, for a comprehensive comparison that takes into account the class imbalance of the classification task, we calculated the bAcc for all state-of-the-art methods based on the SN and SP values reported in the original articles.

Table 6.6: Comparison of GNN-based models for EEG seizure detection on the CHB-MIT dataset, including the current state-of-the-art.

Work	Method	Graph	Accuracy	Specificity	Sensitivity	Balanced Acc.	# Parameters
Zhao *et al.* [146]	Linear GCN	Pearson	95.89%	95.30%	93.80%	94.55%	–
Nie *et al.* [147]	GCNTrans	Pearson	–	94.60%	92.97%	93.78%	≈ 66.8 K
Xu *et al.* [112]	GCN	Pearson	95.23%	91.83%	94.35%	92.59%	≈ 18 K
He *et al.* [150]	GAT	Learnable	96.28%	93.03%	92.75%	92.89%	≈ 8.5 K
Xu *et al.* [112]	GCN-BiGRU	Pearson	97.35%	95.83%	98.85%	97.34%	≈ 730 K
He *et al.* [150]	GAT-BiLSTM	Learnable	98.52%	97.75%	94.34%	96.05%	≈ 428 K
This work	GCN-LSTM	Pearson	96.16%	98.45%	87.90%	93.13%	≈ 161 K
on CHB-MIT dataset	GCN-LSTM	PLV	98.49%	98.77%	96.56%	97.66%	≈ 161 K
This work	GCN-LSTM	Pearson	94.11%	97.95%	89.42%	93.68%	≈ 161 K
on Siena dataset	GCN-LSTM	PLV	96.65%	98.63%	94.00%	96.32%	≈ 161 K

Thus, in accordance with the definition in Section 4.3, bAcc was derived as the arithmetic mean of SN and SP.

GNN-Explainer results

Besides the two-stage networks, the interpretation technique provided by the GNN-Explainer tool, plays a crucial role in our framework. Figure 6.5 shows the occurrences of the most relevant nodes identified by GNN-explainer during the training of the ChebNet extractor. The nodes represent the positions of the electrodes according to the international 10-20 system [20], while the color reflect the normalized frequency of appearance in the top 3. It can be seen that the darker and larger nodes appear more frequently, indicating that the model assigns them a predominant role in the classification process. Conversely, the lighter and smaller nodes are selected less frequently, suggesting marginal relevance. For completeness, the absolute frequency of occurrence (i.e., how many times each channel appeared in the top-3) is also reported inside to each node.

In the case of patient chb01, shown in Figure 6.5, it can be seen that the most informative area identified by the model lies in the central region, encompassing the bipolar channels F_3-C_3, Fz-Cz, F_4-C_4, C_3-P_3, Cz-Pz, and C_4-P_4. In particular, these same central channels consistently emerge in all patients, indicating a recurring area that is probably considered more informative by the model. Additional recurrent nodes for the patient under analysis emerge posteriorly (P_3-O_1 and P_7-O_1), which are strongly correlated with each other, and frontally at F_{P1}-F_7, as well as in the right frontal lobe at F_8-T_8.

6.3 Discussion

In this study, we proposed a ST-GCN for the task of seizure detection, designed to distinguish between ictal and non-ictal activity in 3-second segments of continuous EEG recordings. Each node feature corresponds to a 1-second raw EEG signal, from which ChebNet-based graph convolutions directly extract discriminative hidden representations. These spatial features are subsequently fed into a LSTM network, which captures the

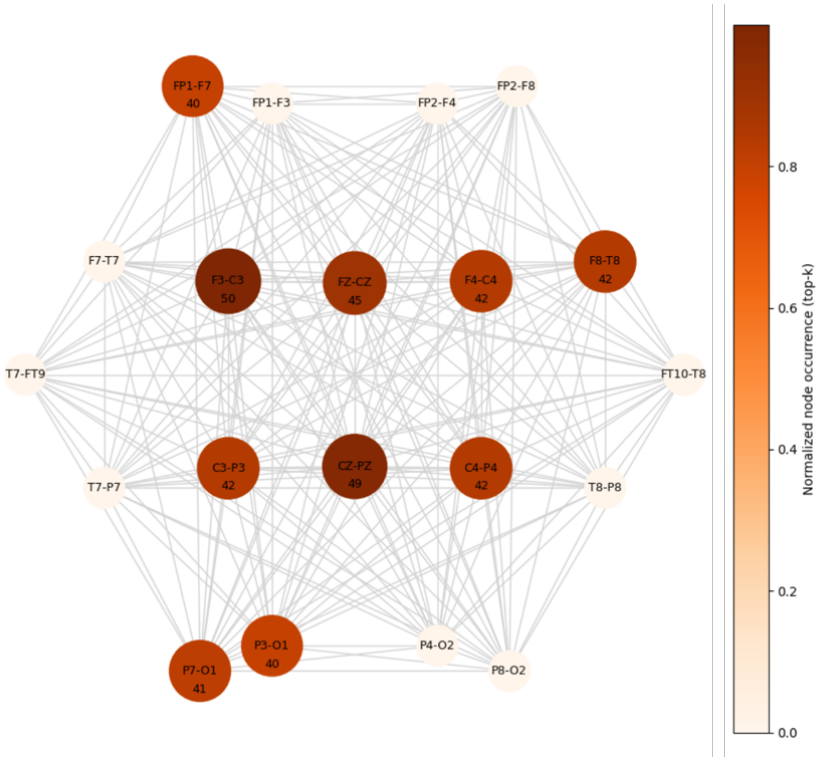


Figure 6.5: Graph representation showing the channels that appeared most frequently among the top-3 nodes identified by GNNExplainer in the analyzed ictal windows. The color intensity reflects the normalized occurrence of each channel, with darker nodes indicating higher frequency.

temporal dynamics of seizure evolution by integrating information from both past and future windows.

The integration of spatial graph representations and temporal modeling results in a detection framework that achieves high SN and low error rates. We validate the effectiveness of the proposed model on two publicly available datasets, CHB-MIT [79] and Siena Scalp EEG [80], adopting a 5-fold cross-validation protocol to ensure robust performance and assess its potential clinical applicability [112].

To better understand the effectiveness of the two different connectivity measures in seizure detection, we made a comparison based on a priori-analysis. We computed the connectivity

matrices for both ictal and non-ictal windows for each patient. Then, we selected PLV as the more appropriate metric to represent the connections that each EEG electrode has between them as a network.

Finally, the developed model outperformed the state-of-the-art models as shown in Table 6.6, achieving a bAcc of 97.66%. In contrast, the baseline of Xu *et al.* [112], which represents the prior state of the art for subject-dependent seizure detection using a GNN, achieves 97.34% accuracy. Although the accuracy gain is low, our model attains it with 161K parameters versus 730K in Xu *et al.*, a reduction that makes the improvement particularly meaningful. At the same time, it should be noted that some approaches, such as the single GCN model in [112] and the single GAT model in [150], achieve competitive balanced accuracy with fewer learnable parameters. However, their lower complexity is mainly due to the use of static architectures, which do not explicitly model the temporal evolution of EEG connectivity patterns. Indeed, when a temporal module is introduced, the number of parameters naturally increases. This further strengthens the value of the proposed approach, which achieves state-of-the-art performance while preserving a favorable trade-off between model complexity and spatio-temporal representational power.

From a computational perspective, however, this trade-off should be interpreted by distinguishing between the offline evaluation of the connectivity metric and the online graph-construction stage. The comparative analysis between PCC- and PLV-based connectivity was performed offline to determine the most suitable FC measure for seizure discrimination. However, once PLV was selected, its estimation became part of the operational pipeline, since a connectivity matrix must be computed for each one-second EEG window in order to construct the corresponding graph. Therefore, although the adopted sparsification strategy reduces graph complexity and downstream model burden, the computation of phase synchronization still represents an online pre-processing cost in the proposed full-channel framework.

Connectivity analysis.
Electrical discharges in different areas of the brain do not always occur in a perfectly synchronized mode. For this reason, understanding in which phase (ictal or non-ictal) the correlation is weakest can help identify the onset of an epileptic seizure [161]. Under normal conditions, EEG activity shows predominantly random discharges. However, during the epileptic state, the discharges become progressively more rhythmic, increasing the correlation between brain areas until reaching a peak in the hyper-synchrony phase, which marks the approaching end of the seizure [161].

Based on these considerations, we have shown that the choice of FC measure to represent the edges of the graph cannot be arbitrary and must reflect the properties of the task in question. We compared two complementary indices of channel relationships: (i) PCC, which quantifies the degree of correlation between two signals, taking into account both amplitude and phase information; (ii) PLV, which quantifies frequency-specific phase synchronization between two signals, regardless of amplitude.

In our study, both metrics showed higher values during the ictal phase than during the non-ictal phase (see Table 6.1 and Figure 6.1), confirming the increased neuronal synchronization typical of seizures. However, statistical analysis revealed a clear difference. While PLV showed a significant increase in the ictal phase ($p < 0.01$), PCC did not show the same discriminatory power [124]. This result highlights how phase-based indices are able to more effectively capture the dynamic changes in brain activity associated with the onset and propagation of seizures, offering superior discriminative properties compared to amplitude-based measures.

During the ictal phase, the increase observed in PLV is also reflected in the literature [175, 176]. In particular, the recent study by Dallmer-Zerbe *et al.* showed that an increase in connectivity, indicative of greater synchronization between brain areas, promotes the propagation of the seizure. This aspect makes PLV particularly suitable for modeling brain graphs, as it captures the phase dynamics that characterize the transition to the ictal state. Conversely, measures based purely on sig-

nal amplitude, such as PCC, tend not to reflect with the same sensitivity the synchronization mechanisms that are crucial for describing the epileptic phenomenon.

The difference observed between PCC and PLV is also related to an important methodological aspect, i.e., the influence of the reference EEG montage (see Section 2.1.3). In datasets with bipolar leads, such as the CHB-MIT used in this work, this factor is often underestimated, despite the fact that it is well known that the reference scheme can significantly alter the topology of the inferred networks [75]. Indeed, both measures are sensitive to the phenomenon known as volume conduction, which is characteristic of monopolar montages. This effect, due to the passive propagation of the electrical signal through brain and skull tissue, causes artificial overlap between activity originating in different areas of the brain. In other words, it can make two EEG channels appear to be correlated even in the absence of real functional interaction between the underlying regions, introducing spurious correlations and compromising the correct estimation of FC [177]. The use of bipolar leads partially reduces this effect thanks to differential recording, which attenuates the common component associated with signal propagation and better preserves local activity.

Several studies have confirmed the impact of montage on connectivity measures. In particular, Anastasiadou *et al.* demonstrated that the use of metrics less sensitive to reference and volume conduction allows for more stable and discriminative networks [75]. Therefore, the choice of reference scheme significantly modifies the networks inferred by correlation, altering the topology of the functional graph. Pastor and Vega-Zelaya observed that synchronization patterns derived from bipolar leads vary systematically according to brain region and age [178], while Mitsuhashi *et al.* highlighted how bipolar montage effectively suppresses extra-cortical activity, preserving that of gray matter and improving the relevance of local connectivity [179]. In this context, PLV, based on phase synchronization and consequently less sensitive to variations due to montage than purely amplitude measures such as PCC, is confirmed as the most robust and effective choice for representing functional graphs.

In summary, the above description shows that estimating FC from EEG requires careful selection of both the metric and the montage, as both substantially influence conclusions about brain interactions. However, despite empirical evidence, many studies continue to favor PCC, probably because it was the first metric introduced and remains the most commonly known and widely used in the literature. Other studies are also drawing attention to the fact that the current graph-based representation process for EEG data is mainly relying on PCC analysis to construct the adjacency matrix, which cannot show the information specific to EEG [161]. On the other hand, PLV has proven to be more discriminating not only in epilepsy but also in other clinical contexts, such as schizophrenia [149].

Threshold evaluation

The comparative study not only identified which FC metric was most suitable for discriminating between ictal and non-ictal phases, but also allowed the relative probability of connections between EEG channels to be assessed, paving the way for a more targeted analysis aimed at graph sparsification. The elimination of less influential connections proved crucial in highlighting the most relevant functional interactions in the detection of epileptic seizures, favoring the use of a less complex network.

In the literature, the threshold value is often defined based on the performance of the model validated on the validation set [161]. Some more recent studies have instead proposed data-driven thresholding approaches, designed to isolate the most significant connections and verify their robustness at the intra-subject level, comparing the functional networks obtained with null models (random and lattice graphs). A significant contribution in this direction is that of Adebisi *et al.*, who introduce the concept of Brain Functional Networks (BFNs) constructed with a subject-specific approach. These networks, derived from EEG recordings of both healthy subjects and patients with dementia, aim to capture the FC patterns that differentiate the two populations. Once the BFNs were constructed, the authors applied GCNs to support the differential diagnosis between dementigenic disorders, in particular Mild Cognitive Impairment,

Alzheimer's Disease, and Vascular Dementia [154].

In the same way, our work shows how the relative probability distributions of connection strength, shown in Figure 6.3a for PLV and in Figure 6.3b for PCC, show how the information that helps to discriminate between ictal and non-ictal phases emerges with different intensity depending on the considered FC metric. In the case of PLV, the two distributions show a more marked separation, with a central interval in which the relative probability of connections differs significantly between the two conditions. This observation allows us to identify an optimal sparsification threshold, corresponding to the interval in which the separation is maximum. For example, for patient chb01, the point of maximum separation was identified around the value 0.4, which was then selected as the threshold. The adoption of a data-driven threshold not only eliminates the least informative connections but also preserves those most relevant to the characterization of the brain state, reducing the complexity of the graph without compromising its discriminatory power. In fact, with a-posteriori analysis, we verified that choosing a threshold of 0.4 allows the model to achieve the best performance in terms of bAcc, as shown in Figure 6.4, with a maximum value of 0.988 for the patient considered.

On the contrary, in the case of PCC, the distributions relating to the ictal and non-ictal phases are more overlapping, with less evidence of clearly separable intervals. This made it more difficult to identify an optimal threshold value, making the choice necessarily more empirical.

Model evaluation

To evaluate the performance of the proposed model, we compared our results with those of recent studies on the detection of epileptic seizures from EEG, selecting studies that performed subject-specific experiments on the CHB-MIT dataset and reported standardized metrics, such as Acc, SN, and SP. The results are shown in Table 6.6.

The studies shown highlight the tendency in the literature to reduce the complexity of models to facilitate their use in real-time scenarios, a particularly pressing and widely recognized need in the scientific community [50, 53]. For example, Zhao *et*

al. achieve a bAcc of 94.55% with a Linear GCN [108], while Nie *et al.* combine GCN and spatial Transformer to achieve a bAcc of 93.78% with 66.8K parameters [147]. An even lighter approach is that of He *et al.*, who achieve a bAcc of 92.89% with 8.5K parameters. Their strength, compared to the previously mentioned works, lies in the use of learnable adjacency graphs through spatial layers equipped with attention. Furthermore, to overcome the static nature of the graph, the same authors also propose GAT-BiLSTM, bringing the bAcc to 96.05% at the cost of 428K parameters. This combination of edge learning and network dynamism represented an important reference in the state of the art.

It was only a few years later that Xu *et al.* took a further step, replacing GAT with GCN and maintaining a recurrent component to model the dynamics of the EEG [112]. To justify this replacement, the authors point out that the training time required by GAT is longer than that required by GCN, indicating the greater efficiency of their approach. In detail, the authors initially adopted GCN-LSTM, but found a clear drop in performance (bAcc of 80.6%), indicating difficulties in effectively capturing temporal dynamics with this architecture. Then, they proposed GCN-BiGRU, achieving very solid results (bAcc of 97.34%) and establishing themselves as the state-of-the-art baseline.

The choice of our GCN-LSTM architecture arises from the comparison between the last two described approaches. Considering that [150]. propose GAT-BiLSTM, while [112] adopt GCN-BiGRU, obtaining approximately 1% more bAcc than GAT-BiLSTM, we opted for GCN as the spatial backbone. Notably, GCN is also proposed in other contexts as a lightweight architecture [123]. Indeed, a main advantage of using GCN is that it can be further optimized from a computational point of view using Chebyshev operators (as suggested by Jia *et al.* [57]), reducing parameters and costs without compromising the propagation capacity on the graph. For the temporal component, we chose LSTM instead of GRU because, with the same configuration, LSTM is more stable on long sequences thanks to the cell state, which mitigates vanishing gradients and preserves long-term dependencies-useful features in the presence of

slow dynamics and class imbalance typical of EEG. The GCN-LSTM combination is also a recurring lightweight and effective network in several tasks [149], reinforcing its choice as the basis for our model.

Moreover, compared to GAT, GCN offers explicit control over connections (not self-learned), resulting in increased interpretability of the model [154]. This allowed for the integration of an a priori study of FC and the design of more informative graphs, with the potential to reduce the number of layers needed to extract representative features. In this regard, the last two rows of the Table 6.6 show the two variants of our model, in which the graph is constructed from connectivity matrices estimated with PCC or PLV.

The comparison, designed in light of the a priori study on connectivity, shows the impact of the choice of FC measure on the same architecture. Specifically, with PCC we obtain Acc of 96.16%, SP of 98.45% and SN of 87.90% with bAcc of 93.13%. On the contrary, with PLV, SN increases up to 96.56% compared to an still high SP (98.77%), with an overall improvement given by bAcc 97.66%. In both models, the number of parameters is the same (approximately 161K), highlighting that the gain does not derive from a greater capacity of the model but from a more informative representation of the graph. When compared to the state-of-the-art, our variant with PLV slightly exceeds the bAcc of the baseline GCN-BiGRU by Xu *et al.* (97.66% vs. 97.34%) while requiring far fewer parameters (161K vs. 730K), thus offering a better compromise between accuracy and complexity.

GNN Explainability

The results obtained using GNNExplainer (see Figure 6.5) further reinforce the results of our previous study (see Section 5.1.2), in which we demonstrated that a reduced set of patient-specific channels, optimized according to individual seizure patterns, is sufficient to maintain, or even improve, diagnostic accuracy. In particular, the nodes most frequently highlighted by GNNExplainer correspond to the channels that emerged as the most discriminating in the personalized selection approach for seizure detection. This confirms that the channel selection algorithm proposed in the previous study effectively identifies the cortical areas most involved in epileptic dynamics.

For patient chb01, what was found is that the central channels (F_3-C_3, Fz-Cz, F_4-C_4, C_3-P_3, Cz-Pz, C_4-P_4) were particularly prominent, and this trend was confirmed for all patients analyzed in this study. Interestingly, the systematic emergence of the central brain region is also consistent with the results obtained in the study described in Section 5.3.2, where we evaluated the generalizability of the channel selection algorithm in a multi-patient setting (i.e., by looking for two representative channels common to all subjects) and we found the central channels emerged as the most shared among patients. From the clinical point of view, the recurrence of the central area can be interpreted as an expression of the fact that the central regions of the cortex, being more connected and involved in the propagation of epileptic seizures, contribute more consistently to the discriminative patterns captured by the model.

Other recurring nodes for the patient chb01 emerge at the posterior level, in particular P_3-O_1 and P_7-O_1, which are strongly correlated with each other as they cover the same area of the scalp, at the frontal level with F_{P1}-F_7, and in the right frontal lobe with F_8-T_8. It is interesting to note that the nodes P_3-O_1 and F_8-T_8 had already been highlighted in the channel selection study presented in Section 4.1.3.1 of this thesis. This correspondence, on the one hand, confirms that some nodes are indeed more informative than others in seizure detection. On the other hand, it underlines the robustness of the approach adopted, as different analysis methods converge in identifying the same cortical areas as particularly relevant.

6.4 Conclusion

In this chapter, we proposed a novel approach employing a lightweight spatio-temporal Graph convolutional network (ST-GCN) for detecting epileptic seizures. The proposed GNN methodology combines an LSTM module, which captures the temporal evolution of signals, with a ChebNet-GCN, which extracts their spatial features.

To build truly informative brain graphs, we performed an in-depth, data-driven comparison between two measures of functional connectivity, selecting the one that best captures the inter-dependencies between nodes for the seizure detection task. This choice allowed us to build customized graphs for each patient and for each condition (no-seizure/ictal). By informatively modeling the graphs at different stages, our model was lighter in computational terms and achieved better or comparable performance compared to the baselines. From a graph construction perspective, the targeted selection of the most discriminative functional connectivity measure between the ictal and no-seizure phases allows us to construct the graph from a single informative index, reducing computational overhead and improving the clinical interpretation of the graph. Furthermore, the threshold value τ was examined and set according to a justified criterion, so as to generate a sparse graph that reduces input complexity and allows the ST-GCN to effectively discriminate between the two conditions.

The final result is a more compact network, with lower processing times and memory usage. In clinical practice, this translates into better scalability on long-term recordings, reduced inference latency, and the feasibility of implementation on hardware with limited resources (i.e., bedside monitoring), without loss of performance in detecting epileptic seizures. The clinical implications are further supported by the level of explainability introduced in this work. In addition to functional connectivity analysis, we adapted GNNExplainer to interpret the model results by analyzing the recurrence of the most relevant EEG nodes. Although a first level of robustness was already guaranteed by the choice of the most discriminative connectivity measure and the achievement of competitive per-

formance, it was necessary to add a level of understanding, i.e., insight into which channels the model focuses on. To this end, we used GNNExplainer to extract critical subgraphs and construct a node frequency map, indicative of the most relevant cortical areas. The analyses show that, during the ictal event, specific brain regions systematically emerge and that these regions coincide with the channels identified in the selection study presented in the first part of the thesis, thus reinforcing the consistency and robustness of both approaches.

However, our model is subject to some limitations. The first is the relatively small sample size, which reduces the generalizability of the best classifier. This is a common problem in DL models, which generally require large training datasets to achieve high accuracy but are limited by the limited availability of online databases. Another limitation is that our classifier relies on spectral GCNs obtained from the Laplacian of the adjacency matrix, making it less robust to variations in the number of EEG channels. These drawbacks suggest two possible extensions of the work. On one side, the fusion of functional and structural connectivity to obtain more informative graphs; on the other side, the exploration of alternative topologies, including space-based GCNNs, which are better suited to managing variable sensory layouts.

Chapter 7

Conclusion

This PhD work is based on the idea that applying Artificial Intelligence (AI) to biosignals can support the automated detection of events, reducing the time required by traditional analyses and enabling scalable signal analysis in practical, real-world settings. However, in the context of EEG, despite extensive research on AI-based analysis, clinical interpretation is still largely visual and carried out manually by expert neurologists. Therefore, the main purpose of this work is to address some of the key challenges that currently limit the adoption of automated AI methods. Specifically, the work aims to tackle the spatial complexity of EEG and the appropriate representation of EEG input, in order to better support the AI model while reducing its computational complexity, as well as the effective visualization of the output, to facilitate the interpretation of results even by users who are not AI experts.

These aspects are explored with a particular focus on low-density application systems, prompted by the scientific community's growing interest in EEG wearable solutions. Indeed, unlike other biosignals, for which dedicated devices for on-device acquisition and on-board analysis are already available, their application to EEG remains limited, despite clear technological advances in the field of EEG acquisition. This has led to the need for lightweight AI models that are suitable for operating in contexts with limited computational resources and for supporting long-term monitoring within telemedicine.

Epilepsy detection is the case study investigated. It was se-

lected not only because it is one of the most widely studied topics in the application of AI to EEG, along with brain-computer interfaces (BCI) and sleep stage classification, but also because it exhibits high variability in spatial patterns across different EEG channels. Unlike BCI, where the electrodes of interest are typically restricted to the motor cortex area (i.e., C_3, C_4, Cz), epilepsy detection involves a broader range of brain regions, making it particularly suitable for investigating spatial complexity in EEG analysis.

This chapter summarizes the main contributions of the thesis, highlighting the innovative approaches and advances made with respect to the state of the art in Section 7.1. Section 7.2 addresses the limitations of the study, while Section 7.3 proposes future directions for research in this field.

7.1 Main contributions

This study addresses the issue of the spatial complexity of EEG in epilepsy with an explicitly data-driven and interpretable approach, introducing two complementary methodological frameworks: channel selection for reduced-channel EEG analysis and graph-based modeling for analysis including the full set of EEG channels. Both frameworks are designed to balance accuracy and interpretability while respecting computational load constraints.

The main innovation lies in the in-depth analysis of the EEG signal in its full temporal and spatial complexity, with the aim of optimizing a data-light deep learning (DL) model, rather than forcing the data to adapt to a predefined architecture. In other words, the modeling process itself becomes a tool for interpreting the signal, rather than a simple pre-processing step that feeds a final "black-box" classifier. In this framework, each modeling choice is guided by intermediate evidence, such as topographic maps, node-level explanations, and connectivity analyses, which inform the subsequent stages of DL model design and support their impact on the final results.

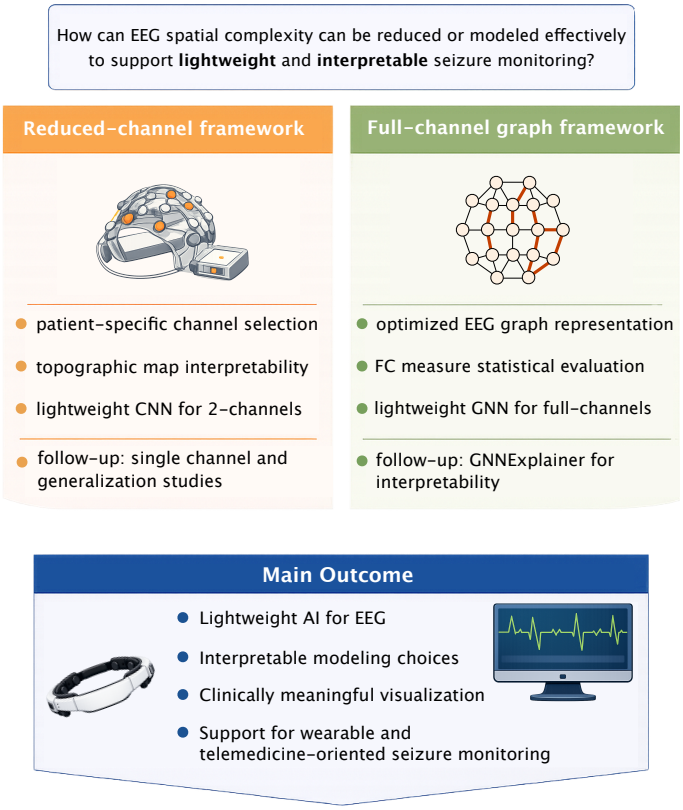


Figure 7.1: Visual synthesis of the main contributions of the thesis.

Moving on to the specific contributions proposed in this work, Figure 7.1 provides a summary of the two main research directions developed in the thesis. The first concerns the introduction of a lightweight model optimized for a reduced number of EEG channels. The basic idea is to select only the most informative EEG channels for detecting epileptic seizures. Channel selection is patient-specific, leveraging detailed information to support the development of personalized clinical monitoring and treatment strategies. Although the use of a reduced set of electrodes may lead to a deterioration in seizure detection performance, our studies (see Sections 5.1.2, 5.2.2, and 5.3.2) demonstrate that accurate and adaptive channel selection can minimize this loss. This is because some electrodes do not provide information relevant to seizure detection and mainly

introduce redundancy.

Compared to the state of the art, which tends to rely on wrapper strategies for channel reduction, for example by adding or removing channels until performance drops, we propose a selection framework that starts from interpretable visualizations and proceeds through steps guided by intermediate results. Indeed, the topographic maps obtained from the analysis made it possible to: (i) highlight the cortical areas that are more informative for the task; (ii) clinically support the seizure detection task; and (iii) show that the ictal/non-ictal combination provides complementary information and facilitates discrimination. It follows that a significant reduction in channels remains effective when choices are explained and customized to the spatial dynamics of the individual patient. The analysis also clarifies why and when performance degrades as reduction increases, offering concrete indications for future improvements. For example, one possible insight concerns the dependence on seizure phases. Since different areas of the brain emerge at different times during epileptic events, including pre-ictal and post-ictal phases, a possible extension of the work could involve the use of three or four channels, with one channel selected in each phase, shifting the focus from finding the *minimum number* of electrodes to finding the *minimum informative set*, depending on the phase and location. This insight was not explored in the present work, as publicly labeled data for these phases are not available; therefore, a dedicated study in close collaboration with clinical experts would be necessary to define appropriate annotations and validate such a phase-specific strategy.

Further valuable insights regarding reduced-channel design, explored in this work, concern both the qualitative role of channel selection and the feasibility of patient-independent configurations. With respect to the first aspect, the comparison with the reference work, selected because it introduced a lightweight DL architecture suitable for embedded implementation, highlights that channel reduction should be viewed not only in quantitative but also in qualitative terms. In that study, the configuration relies on four fixed temporal channels (F_7-T_7, T_7-P_7, F_8-T_8, and T_8-P_8) that are the same for all patients. In

contrast, our approach shows that, with patient-specific data-driven selection, performance can remain stable even with only two channels, whereas if the ictal channel is not included among the four fixed baseline channels, the performance of the patient-specific model can degrade significantly. This indicates that what truly matters is not just how many channels are retained, but which channels are selected and how effectively they capture the most informative phases of the signal.

With respect to the second aspect, in order to mitigate the intrinsic limitation of a fully customized solution, which in practical applications would require extensive customization of the wearable system for each patient, we explored the development of a patient-independent model based on two EEG channels shared across subjects. The results indicate that generalization with two fixed channels can lead to a performance degradation compared to the four fixed channels configuration assumed as the baseline. Despite this limitation, the investigation provides several interesting observations. First, when our selection algorithm is applied in a multi-patient context, the channels that emerge in the non-ictal phase tend to systematically correspond to the central area of the scalp (Cz-Pz), which is clinically meaningful, as these EEG derivations capture more widespread patterns of cortical synchronization and desynchronization. Second, the combined choice of one ictal channel and one central non-ictal channel achieves higher performance than using two of the four fixed temporal channels proposed in the literature. This result outlines a promising direction for the development of a more robust generalized model, in which fine-tuning or transfer learning techniques could handle high inter-subject variability and improve performance in the proposed two-fixed-channel configuration.

The second main contribution of this work is a complementary lightweight approach optimized for the full set of EEG channels. Instead of focusing on selecting a subset of electrodes, this study shifts the attention to modeling the spatio-relational structure of the signal through a graph representation. This choice is motivated by the growing interest in applying graph neural networks (GNNs) to EEG data, but also by the observation that the current EEG–GNN literature is

still limited, especially with regard to the definition of graph connections. On the one hand, several state-of-the-art approaches rely on attention-based mechanisms, such as graph attention networks (GATs), to learn connectivity; these models are powerful but often heavy and not fully aligned with the interpretability and efficiency objectives of this thesis. On the other hand, many works still adopt static connectivity based solely on the spatial distance between electrodes, which does not capture the dynamic nature of cortical interactions. To address this gap, we propose a systematic, data-driven analysis of functional connectivity (FC) choices in the specific context of seizure detection. In particular, we compare two of the most widely used FC measures, i.e. Pearson Correlation Coefficient (PCC) and Phase-Locking Value (PLV), and study how both the chosen metric and the sparsification thresholds shape the resulting EEG graph and, consequently, the performance of GNNs. The key message is that FC design, including the choice of the measure, the normalization strategy, and the criteria used to prune edges, can influence performance as much as architectural factors such as network depth or type of GNN layers. In other words, we show that a well-founded and explicitly justified graph construction makes it possible to use lighter GNNs while still achieving competitive performance, reducing complexity without compromising discriminatory performance.

The proposed framework is built to satisfy the same interpretability and lightweight requirements as the reduced-channel study. We adopt graph convolutional networks (GCNs) to construct graph-based EEG representations, since they are computationally lighter than GATs and offer a more transparent processing pipeline. To preserve the dynamic nature of the model, the FC between channels is modeled as a time-varying graph, updated through a Long Short-Term Memory (LSTM) recurrent module. Consistency with the first framework is further ensured by employing GNNExplainer to quantify which electrodes contribute most to the classification of ictal versus non-ictal states. The analysis shows that, even when all channels are available, the model tends to assign higher importance to the same electrodes highlighted by the earlier channel selection study across different seizure phases, providing

cross-validation between methods and reinforcing the clinical plausibility of the learned representations. In terms of impact on the state of the art, this work demonstrates that performance gains in EEG–GNN models do not necessarily require increasingly complex architectures. While recent studies often report high and very similar accuracy values by stacking deeper GNNs, these improvements typically come at the cost of higher computational burden and reduced transparency. In contrast, our results indicate that the most effective way to reduce complexity while preserving accuracy and interpretability lies in a conscious, task-driven design of FC, combined with appropriate normalization and sparsification steps that structure the graph before the network even comes into play.

7.2 Limitations

In this section, we outline the main limitations encountered during the research and explicitly use them to motivate the future developments discussed in Section 7.3. The most relevant limitations can be grouped into three categories: (i) data availability and statistical constraints; (ii) inter- and intra-subject variability and limited generalization; and (iii) the restricted clinical scope, focused on seizure detection rather than seizure prediction.

First, the major challenge in applying AI models to EEG signals is the limited availability of large public datasets. This is true not only in the clinical domain, but also in other application areas such as brain–computer interfaces (BCI), sleep, and neurodegenerative diseases, where access to high-quality, labeled signals remains restricted. This scarcity directly impacts the training and validation of complex AI architectures, as it can lead to overfitting and reduced generalization. In our specific case, the datasets used comprise only 24 subjects for CHB-MIT and 11 for Siena Scalp. To partially mitigate this issue, we increased the number of training samples by splitting continuous signals into smaller windows, which allowed us to train lightweight models effectively. However, this solution cannot be transferred to statistical analyses, because windows

coming from the same participant share variance and therefore violate the independence assumptions of most statistical tests.

Second, EEG signals exhibit substantial variability both between different patients (inter-subject variability) and within the same patient across sessions (intra-subject variability). This heterogeneity makes it difficult to develop models that generalize well across subjects and recording conditions. For this reason, the work primarily focuses on patient-specific model customization, in line with the paradigm of personalized medicine. This choice is supported by evidence from the first study (see Section 5.1), which showed that the electrophysiological "history" of each patient differs significantly from that of others, and that this effect is particularly marked in epilepsy. Indeed, unlike other tasks such as sleep staging or BCI, where the relevant cortical areas are relatively stable, epileptic phenomenology depends on seizure type (focal vs. generalized) and can involve highly variable spatial and temporal patterns. At the same time, the practical integration of personalization into a clinical workflow requires careful consideration, since patient-specific approaches may involve additional calibration effort, greater dependence on patient-specific data, and the need for model updating over time. In real-world settings, these aspects must be balanced against workflow efficiency and limited specialist time. For these reasons, we also explored inter-patient generalization (Section 5.3). However, performance was lower in the reduced-channel setting, confirming how difficult it is to move away from a fully personalized framework. These limitations motivate future research on data-efficient models and large-scale pre-training strategies, as discussed in Section 7.3.

Third, on the clinical side, this thesis focuses exclusively on seizure detection, without extending the analysis to seizure prediction. Prediction is a substantially more challenging task, mainly because publicly available datasets typically do not provide reliable labels for pre-ictal periods and because a rigorous definition of these intervals would require close collaboration with clinical experts to ensure consistency with clinical practice. As a consequence, we could not systematically address seizure prediction within the present work. However, the results obtained in the first study, in particular the different

phases highlighted by the channel selection algorithm around seizure onsets, suggest that distinct groups of EEG channels become more informative at different times during epileptic events. Since different brain areas emerge at different stages of a seizure, these observations directly motivate future work on patient-specific seizure prediction models. A natural extension would be to enrich the available annotations with pre-ictal labels and move from a binary detection task to a three-class classification problem (inter-ictal, pre-ictal, ictal), as discussed in Section 7.3.

7.3 Future directions

Building on the limitations discussed in Section 7.2, we identify several complementary research directions that can strengthen and extend the contributions of this thesis. Each of these directions is designed to address one or more of the constraints highlighted above.

A first line of development concerns seizure prediction. While the present work focused on seizure detection, prediction remains a greater challenge in both clinical practice and algorithm design. It is usually framed as a classification problem involving two brain states, such as inter-ictal vs. pre-ictal or ictal vs. pre-ictal. However, as discussed in Section 2.2.2 and in the literature, existing studies that attempt to separate the pre-ictal state from other phases still lack a clear and consistent methodology for defining its duration. There is no universally accepted clinical definition of the pre-ictal interval [33], and different works adopt markedly different time windows, ranging from tens of seconds [180] to tens of minutes [69, 74, 88]. This makes it difficult to compare results and to derive clinically meaningful guidelines.

These elements indicate that a central objective for automatic seizure prediction algorithms is the robust identification of the pre-ictal state within a predefined time window. The "ideal" case would be to define a fixed pre-ictal duration valid for an entire group of patients [33]. However, the intrinsic variability of EEG signals and dynamic brain models suggests

that each subject exhibits unique neural patterns, making the adoption of a single pre-ictal interval inappropriate. Based on the results obtained in this thesis, in particular the channel selection analysis showing that differences in the temporal evolution of weights between ictal and inter-ictal phases reveal distinct neural dynamics across inter-ictal, pre-ictal, and ictal states, a natural extension will be to transfer the proposed methodology to the predictive setting. The goal will be to systematically analyze the duration and dynamics of the pre-ictal phase across different patients, using the same interpretable channel selection and connectivity analysis tools adopted here. This would help derive new evidence for defining personalized prediction models and could lead to explicit guidelines linking channel selection strategies to specific epilepsy tasks (detection vs. prediction), directly addressing the limitations related to generalization and variability discussed in Section 7.2.

A second complementary direction directly targets these issues of data scarcity and statistical constraints described in Section 7.2. Large models have shown outstanding robustness and performance in computer vision and natural language processing, mainly due to training on extensive datasets. Integrating these strategies with the emerging concept of *brain foundation models* opens new opportunities for leveraging large-scale pre-trained architectures. By adapting such architectures to the specific characteristics and bi-dimensional complexity of EEG data, it may be possible to obtain systems that are more reliable, generalizable, and interpretable, while reducing the amount of labeled data required for downstream tasks such as seizure detection and prediction.

In the literature, promising examples of brain foundation models applied to EEG are emerging, indicating significant potential for these approaches in the neurophysiological domain as well. Among them are *BIOT* (Biosignal Transformer), a pre-trained model designed for EEG signal processing that uses self-attention mechanisms to capture temporal dependencies in sequences [181]; *LaBraM* (Large Brain Model), an advanced neural architecture trained on large EEG datasets, whose multi-head self-attention structures enable the identification of the most relevant patterns, improving classification accuracy [182];

EEGPT, a transformer-based model inspired by GPT architectures and adapted for EEG signal processing, equipped with temporal encoding techniques to more effectively handle the sequential nature of the data [183]; and *Neuro-GPT*, a large-scale model specialized and fine-tuned on neural data, which integrates domain-specific embeddings to enhance interpretability and generalization across different EEG datasets [72].

However, the adoption of these models in EEG systems remains limited due to the scarcity of available data, although EEG-specific data augmentation techniques, based on signal processing or adversarial examples, offer a promising path to mitigate this issue. Brain foundation models, in particular, hold remarkable potential for advancing EEG-based artificial intelligence. Their goal is to create universal representations of EEG signals through pre-training on large and diverse datasets, providing a solid backbone for a wide range of downstream tasks, including emotion recognition, sleep stage classification, and epilepsy detection. This approach enables efficient transfer learning across applications, requiring minimal fine-tuning.

A key element in developing such models is the use of self-supervised learning techniques, which make it possible to extract rich and transferable EEG features from unlabeled data. Techniques such as contrastive learning, masked signal modeling, and predictive coding have shown strong potential in capturing meaningful representations of the underlying structure of EEG signals without the need for large amounts of annotated data. When these representations are combined with targeted fine-tuning, they can substantially improve both the robustness and adaptability of EEG systems.

Bibliography

[1] G. Niso, E. Romero, J. T. Moreau, A. Araujo, and L. R. Krol, "Wireless eeg: A survey of systems and studies," *NeuroImage*, vol. 269, p. 119774, 2023.

[2] F. Lotte, L. Bougrain, A. Cichocki, M. Clerc, M. Congedo, A. Rakotomamonjy, and F. Yger, "A review of classification algorithms for eeg-based brain-computer interfaces: a 10 year update." *Journal of neural engineering*, vol. 15, no. 3, p. 031005, 2018.

[3] S. Xu, O. Faust, S. Seoni, S. Chakraborty, P. D. Barua, H. W. Loh, H. Elphick, F. Molinari, and U. R. Acharya, "A review of automated sleep disorder detection," *Computers in Biology and Medicine*, vol. 150, 2022.

[4] M. K. Siddiqui, R. Morales-Menendez, X. Huang, and N. Hussain, "A review of epileptic seizure detection using machine learning classifiers," *Brain Informatics*, vol. 7, no. 1, 2020.

[5] C. Yu and M. Wang, "Survey of emotion recognition methods using eeg information," *Cognitive Robotics*, vol. 2, pp. 132 – 146, 2022.

[6] M. Aviles, L. M. Sanchez-Reyes, J. M. Alvarez-Alvarado, and J. Rodriguez-Resendiz, "Machine and deep learning trends in eeg-based detection and diagnosis of alzheimer's disease: A systematic review," *Eng*, vol. 5, no. 3, pp. 1464 – 1484, 2024.

[7] H. W. Loh, W. Hong, C. P. Ooi, S. Chakraborty, P. D. Barua, R. C. Deo, J. Soar, E. E. Palmer, and U. R. Acharya, "Application of deep learning models for automated identification of parkinson's disease: A review (2011-2021)," *Sensors*, vol. 21, no. 21, 2021.

[8] X. Zhou, C. Liu, Z. Wang, L. Zhai, Z. Jia, C. Guan, and Y. Liu, "Interpretable and robust ai in eeg systems: A survey," 2024.

[9] W.-H. Wang and W.-S. Hsu, "Integrating artificial intelligence and wearable iot system in long-term care environments," *Sensors*, vol. 23, no. 13, 2023.

[10] A. Biondi, E. Dursun, P. F. Viana, P. Laiou, and M. P. Richardson, "New wearable and portable eeg modalities in epilepsy: The views of hospital-based healthcare professionals," *Epilepsy and Behavior*, vol. 159, p. 109990, 2024.

[11] F. Brunnhuber, J. Slater, S. Goyal, D. Amin, G. Thorvardsson, D. R. Freestone, and M. P. Richardson, "Past, present and future of home video-electroencephalographic telemetry: A review of the development of in-home video-electroencephalographic recordings." *Epilepsia*, vol. 61, 2020.

[12] M. De Vos, C. Chazachristos, and M. Bhagubai, "Biomed seizure detection challenge," 2024, iCASSP 2023 Grand-challenge. [Online]. Available: https://biomedepi.github.io/seizure_detection_challenge

[13] R. Ferrara, M. Giaquinto, G. Percannella, L. Rundo, and A. Saggese, "Personalizing seizure detection for individual patients by optimal selection of eeg signals," *Sensors*, vol. 25, no. 9, 2025.

[14] R. Ferrara, M. Giaquinto, G. Percannella, A. Saggese, and M. Vento, "Wearable eeg systems: toward single-channel personalized configurations for seizure detection," in *Proceedings of the 9th National Congress of Bioengineering (GNB 2025)*, Palermo, Italy, Jun. 2025, june 16–18, 2025.

[15] ——, "Personalized vs unpersonalized medicine in wearable based seizure detection," in *47th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2025, [Accepted one-page abstract].

[16] A. Chari, R. C. Thornton, M. M. Tisdall, and R. C. Scott, "Microelectrode recordings in human epilepsy: a case for clinical translation," *Brain Communications*, vol. 2, no. 2, p. fcaa082, 2020.

[17] C. Maher, Y. Yang, N. D. Truong, C. Wang, A. Nikpour, and O. Kavehei, "Seizure detection with reduced electroencephalogram channels: research trends and outlook," *Royal Society open science*, vol. 10, no. 5, p. 230022, 2023.

[18] B. H. Uber, "Uber das elektroenkephalogramm des menschen," *Arch Psychiatr Nervenk*, vol. 87, no. 1, pp. 527–70, 1929.

[19] N. Bajaj, "Wavelets for eeg analysis," in *Wavelet Theory*, S. Mohammady, Ed. Rijeka: IntechOpen, 2020, ch. 5.

[20] G. Klem, H. Luders, H. Jasper, and C. Elger, "The ten-twenty electrode system of the international federation. the international federation of clinical neurophysiology." *Electroencephalography and clinical neurophysiology. Supplement*, vol. 52, pp. 3–6, 1999.

[21] BioRender, "Biorender," 2025, created with app BioRender BioRender.com. [Online]. Available: https://www.biorender.com

[22] J. N. Acharya, A. J. Hani, P. Thirumala, and T. N. Tsuchida, "American clinical neurophysiology society guideline 3: a proposal for standard montages to be used in clinical eeg," *The Neurodiagnostic Journal*, vol. 56, no. 4, pp. 253–260, 2016.

[23] M. Seeck, L. Koessler, T. Bast, F. Leijten, C. Michel, C. Baumgartner, B. He, and S. Beniczky, "The standardized eeg electrode array of the ifcn," *Clinical Neurophysiology*, vol. 128, no. 10, pp. 2070–2077, 2017.

[24] S. Hu, Y. Lai, P. A. Valdes-Sosa, M. L. Bringas-Vega, and D. Yao, "How do reference montage and electrodes setup affect the measured scalp eeg potentials?" *Journal of Neural Engineering*, vol. 15, no. 2, p. 026013, 2018.

[25] J. W. Britton, L. C. Frey, J. L. Hopp, P. Korb, M. Z. Koubeissi, W. E. Lievens, E. M. Pestana-Knight, and E. K. St. Louis, *Electroencephalography (EEG): An Introductory Text and Atlas of Normal and Abnormal Findings in Adults, Children, and Infants*. American Epilepsy Society, 2016.

[26] V. Rao and D. Lowenstein, "Epilepsy," *Current Biology*, vol. 25, no. 17, pp. R742–R746, 2015.

[27] F. Chowdhury *et al.*, "Localizing and lateralizing value of auditory phenomena in seizures," *Practical Neurology*, vol. 21, no. 6, pp. 481–489, 2021.

[28] I. E. Scheffer, S. Berkovic, G. Capovilla, M. B. Connolly, J. French, L. Guilhoto, E. Hirsch, S. Jain, G. W. Mathern, S. L. Moshé, D. R. Nordli, E. Perucca, T. Tomson, S. Wiebe, Y. H. Zhang, and S. M. Zuberi, "Ilae classification of the epilepsies: Position paper of the ilae commission for classification and terminology," *Epilepsia*, vol. 58, no. 4, pp. 512–521, 2017.

[29] H. Lüders, J. Acharya, C. Baumgartner, S. Benbadis, A. Bleasel, R. Burgess, D. Dinner, A. Ebner, N. Foldvary, E. Geller, H. Hamer, H. Holthausen, P. Kotagal, D. Lachhwani, R. Maciunas, S. Noachtar, F. Rosenow, A. Sakamoto, B. Steinhoff, I. Tuxhorn, and E. Wyllie, "Semiological seizure classification," *Epilepsia*, vol. 39, no. 9, pp. 1006–1013, 1998.

[30] I. Obeid and J. Picone, "The temple university hospital eeg data corpus," *Frontiers in Neuroscience*, vol. 10, p. 196, 2016.

[31] R. S. Fisher, W. van Emde Boas, W. Blume, C. Elger, P. Genton, P. Lee, and J. J. Engel, "Epileptic seizures and epilepsy: Definitions proposed by the international league against epilepsy (ilae) and the international bureau for epilepsy (ibe)," *Epilepsia*, vol. 46, no. 4, pp. 470–472, 2005.

[32] E. H. Smith, J.-y. Liou, E. M. Merricks, T. Davis, K. Thomson, B. Greger, P. House, R. G. Emerson, R. Goodman, G. M. McKhann, S. Sheth, C. Schevon, and J. D. Rolston, "Human interictal epileptiform discharges are bidirectional traveling waves echoing ictal discharges," *eLife*, vol. 11, p. e73541, 2022.

[33] A. Leal, J. Curty, F. Lopes, M. F. Pinto, A. Oliveira, F. Sales, A. M. B. Bianchi, M. G. Ruano, A. Dourado, J. Henriques, and C. A. Teixeira, "Unsupervised eeg preictal interval identification in patients with drug-resistant epilepsy," *Scientific Reports*, vol. 13, no. 784, 2023.

[34] G. Segal, N. Keidar, M. Herskovitz, and Y. Yaniv, "Personalized preictal eeg pattern characterization: do timing and localization matter?" *Frontiers in Neuroscience*, vol. 19, p. 1526963, 2025.

[35] Y. G. Chung, A. Cho, H. Kim, and K. J. Kim, "Single-channel seizure detection with clinical confirmation of seizure locations using chb-mit dataset," *Frontiers in Neurology*, vol. 15, p. 1389731, 2024.

[36] J. C. M. Pottkämper, J. Hofmeijer, J. A. van Waarde, and M. J. A. M. van Putten, "The postictal state – what do we know?" *Epilepsia*, vol. 61, no. 6, pp. 1045–1061, 2020.

[37] F. Lopes, A. Leal, M. F. Pinto, and A. Dourado, "Removing artefacts and periodically retraining improve performance of neural network-based seizure prediction models," *Scientific Reports*, vol. 13, p. 5918, 2023.

[38] M. Koutroumanidis, A. Arzimanoglou, R. Caraballo, S. Goyal, A. Kaminska, P. Laoprasert, H. Oguni, G. Rubboli, W. Tatum, P. Thomas *et al.*, "The role of eeg in the diagnosis and classification of the epilepsy syndromes: a tool for clinical practice by the ilae neurophysiology task force (part 1)," *Epileptic Disorders*, vol. 19, no. 3, pp. 233–298, 2017.

[39] P. P. Nair, R. Aghoram, and M. L. Khilari, "Applications of artificial intelligence in epilepsy," *International Journal of Advanced Medical and Health Research*, vol. 8, no. 2, pp. 41–48, 2021.

[40] E. Trinka and M. Leitinger, "Which EEG patterns in coma are nonconvulsive status epilepticus?" *Epilepsy and Behavior*, vol. 49, pp. 203–222, 2015.

[41] M. S. Nafea and Z. H. Ismail, "Supervised machine learning and deep learning techniques for epileptic seizure recognition using eeg signals-a systematic literature review," *Bioengineering*, vol. 9, no. 12, 2022.

[42] V. Barone, M. van Putten, and G. Visser, "Absence epilepsy: Characteristics, pathophysiology, attention impairments, and the related risk of accidents. a narrative review," *Epilepsy and Behavior*, vol. 112, pp. 107 342–107 342, 2020.

[43] C. E. Elger and C. Hoppe, "Diagnostic challenges in epilepsy: seizure under-reporting and seizure detection," *The Lancet Neurology*, vol. 17, no. 3, pp. 279–288, 2018.

[44] J. Guo, T. Feng, P. Wei, J. Huang, Y. Yang, Y. Wang, G. Cao, Y. Huang, G. Kang, and G. Zhao, "Adaptive graph learning with seeg data for improved seizure localization: Considerations of generalization and simplicity," *Biomedical Signal Processing and Control*, vol. 101, p. 107148, 2025.

[45] X. Zhou, C. Liu, R. Yang, L. Zhang, L. Zhai, Z. Jia, and Y. Liu, "Learning robust global-local representation from eeg for neural epilepsy detection," *IEEE Transactions on Artificial Intelligence*, vol. 5, no. 11, pp. 5720–5732, 2024.

[46] D. Klepl, F. He, M. Wu, D. J. Blackburn, and P. Sarrigiannis, "Eeg-based graph neural network classification of alzheimer's disease: An empirical evaluation of functional connectivity methods," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 30, pp. 2651–2660, 2022.

[47] T. L. Richards, T. J. Grabowski, P. Boord, K. Yagle, M. Askren, Z. Mestre, P. Robinson, O. Welker, D. Gulliford, W. Nagy, and V. Berninger, "Contrasting brain patterns of writing-related dti parameters, fmri connectivity, and dti–fmri connectivity correlations in children with and without dysgraphia or dyslexia," *NeuroImage: Clinical*, vol. 8, pp. 408–421, 2015.

[48] K. J. Friston, "Functional and effective connectivity: a review," *Brain connectivity*, vol. 1, no. 1, pp. 13–36, 2011.

[49] E. J. Pegg, J. R. Taylor, S. S. Keller, and R. Mohanraj, "Interictal structural and functional connectivity in idiopathic generalized epilepsy: A systematic review of graph theoretical studies," *Epilepsy and Behavior*, vol. 106, p. 107013, 2020.

[50] H. Xiong *et al.*, "Graph convolution network-based eeg signal analysis: a review," *Medical & Biological Engineering & Computing*, 2023.

[51] H. Mohammadi and W. Karwowski, "Graph neural networks in brain connectivity studies: Methods, challenges, and future directions," *Brain Sciences*, vol. 15, no. 1, p. 17, 2024.

[52] N. Khaleghi, S. Hashemi, M. Peivandi, S. Z. Ardabili, M. Behjati, S. Sheykhivand, and S. Danishvar, "Eeg-based functional

connectivity analysis of brain abnormalities: A systematic review study," *Informatics in Medicine Unlocked*, vol. 47, p. 101476, 2024.

[53] S. A. Rahman, M. I. Khalil, H. Zhou, Y. Guo, Z. Ding, X. Gao, and D. Zhang, "Advancement in graph neural networks for eeg signal analysis and application: A review," *IEEE Access*, 2025.

[54] C. J. Stam, G. Nolte, and A. Daffertshofer, "Phase lag index: assessment of functional connectivity from multi channel eeg and meg with diminished bias from common sources," *Human brain mapping*, vol. 28, no. 11, pp. 1178–1193, 2007.

[55] D. Andac, K.-A. Toshiaki, W. Ye, H. Masaki, and E. Deniz, "Eeg-gnn: Graph neural networks for classification of electroencephalogram (eeg) signals," in *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, Nov. 2021.

[56] D. Klepl, M. Wu, and F. He, "Graph neural network-based eeg classification: A survey," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 32, pp. 493–503, 2024.

[57] M. Jia, W. Liu, J. Duan, L. Chen, C. P. Chen, Q. Wang, and Z. Zhou, "Efficient graph convolutional networks for seizure prediction using scalp eeg," *Frontiers in Neuroscience*, vol. 16, p. 967116, 2022.

[58] L. Yang, C. Chatelain, and S. Adam, "Dynamic graph representation learning with neural networks: A survey," *IEEE Access*, vol. 12, pp. 43 460–43 484, 2024.

[59] L. Bai, G. Litscher, and X. Li, "Epileptic seizure detection using machine learning: A systematic review and meta-analysis," *Brain Sciences*, vol. 15, no. 6, 2025.

[60] J. Xu, K. Yan, Z. Deng, Y. Yang, J.-X. Liu, J. Wang, and S. Yuan, "Eeg-based epileptic seizure detection using deep learning techniques: A survey," *Neurocomputing*, vol. 610, p. 128644, 2024.

[61] Y. Roy, H. Banville, I. Albuquerque, A. Gramfort, T. Falk, and J. Faubert, "Deep learning-based electroencephalography analysis: a systematic review." *Journal of Neural Engineering*, vol. 16, no. 5, p. 051001, 2019.

[62] M. K. Siddiqui, R. Morales-Menendez, X. Huang, and N. Hussain, "A review of epileptic seizure detection using machine learning classifiers," *Brain informatics*, vol. 7, no. 1, p. 5, 2020.

[63] S. Raghu, N. Sriraam, Y. Temel, S. V. Rao, A. S. Hegde, and P. L. Kubben, "Performance evaluation of dwt based sigmoid entropy in time and frequency domains for automated detection of epileptic seizures using svm classifier," *Computers in biology and medicine*, vol. 110, pp. 127–143, 2019.

[64] Y. Xiong, J. Li, D. Wu, F. Dong, J. Liu, L. Jiang, J. Cao, and Y. Xu, "Seizure detection algorithm based on fusion of spatio-temporal network constructed with dispersion index," *Biomedical Signal Processing and Control*, vol. 79, p. 104155, 2023.

[65] A. K. Singh and S. Krishnan, "Trends in eeg signal feature extraction applications," *Frontiers in artificial intelligence*, vol. 5, p. 1072801, 2023.

[66] E. Vafaei and M. Hosseini, "Transformers in eeg analysis: A review of architectures and applications in motor imagery, seizure, and emotion classification," *Sensors*, vol. 25, no. 5, p. 1293, 2025.

[67] Y. Wu, L. Lu, A. Xu, Y. Wang, Z. Li, Z. Yang, L. Zeng, and Q. Li, "Neural networks for epilepsy detection and prediction with eeg signals: a systematic review," *Artificial Intelligence Review*, vol. 59, no. 1, p. 30, 2026.

[68] N. Moutonnet, S. White, B. P. Campbell, S. Sanei, T. Tanaka, H. Ji, D. Mandic, and G. Scott, "Clinical translation of machine learning algorithms for seizure detection in scalp electroencephalography: systematic review," *arXiv preprint arXiv:2404.15332*, 2024.

[69] V. J. Lawhern, A. J. Solon, N. R. Waytowich, S. M. Gordon, C. P. Hung, and B. J. Lance, "Eegnet: a compact convolutional neural network for eeg-based brain–computer inter-

faces," *Journal of neural engineering*, vol. 15, no. 5, p. 056013, 2018.

[70] A. Pandey, S. K. Singh, S. S. Udmale, and K. K. Shukla, "Epileptic seizure classification using battle royale search and rescue optimization-based deep lstm," *IEEE journal of biomedical and health informatics*, vol. 26, no. 11, pp. 5494–5505, 2022.

[71] Y. Li, Z. Yu, Y. Chen, C. Yang, Y. Li, X. Allen Li, and B. Li, "Automatic seizure detection using fully convolutional nested lstm," *International journal of neural systems*, vol. 30, no. 04, p. 2050019, 2020.

[72] W. Cui, W. Jeong, P. Thölke, T. Medani, K. Jerbi, A. A. Joshi, and R. M. Leahy, "Neuro-gpt: Towards a foundation model for eeg," in *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*. IEEE, 2024, pp. 1–5.

[73] M. Saeidi, W. Karwowski, F. V. Farahani, K. Fiok, R. Taiar, P. A. Hancock, and A. Al-Juaid, "Neural decoding of eeg signals with machine learning: a systematic review," *Brain sciences*, vol. 11, no. 11, p. 1525, 2021.

[74] S. Zhao, J. Yang, and M. Sawan, "Energy-efficient neural network for epileptic seizure prediction," *IEEE Biomedical Engineering Journal*, 2021.

[75] M. N. Anastasiadou, M. Christodoulakis, E. S. Papathanasiou, S. S. Papacostas, A. Hadjipapas, and G. D. Mitsis, "Graph theoretical characteristics of eeg-based functional brain networks in patients with epilepsy: The effect of reference choice and volume conduction," *Frontiers in Neuroscience*, vol. 13, p. 221, 2019.

[76] A. Kamrud, B. Borghetti, and C. Schubert Kabban, "The effects of individual differences, non-stationarity, and the importance of data partitioning decisions for training and testing of eeg cross-participant models," *Sensors*, vol. 21, no. 9, p. 3225, 2021.

[77] A. Van Esbroeck, L. Smith, Z. Syed *et al.*, "Multi-task seizure detection: addressing intra-patient variation in seizure morphologies," *Machine Learning*, vol. 102, pp. 309–321, 2016.

[78] M. I. Maturana, C. Meisel, K. Dell, P. J. Karoly, W. D'Souza, D. B. Grayden, A. N. Burkitt, P. Jiruska, J. Kudlacek, J. Hlinka *et al.*, "Critical slowing down as a biomarker for seizure susceptibility," *Nature communications*, vol. 11, no. 1, p. 2172, 2020.

[79] A. Shoeb, "Application of machine learning to epileptic seizure onset detection and treatment," Ph.D. dissertation, Massachusetts Institute of Technology, September 2009.

[80] P. Detti, "Siena Scalp EEG Database (version 1.0.0)," 2020. [Online]. Available: https://doi.org/10.13026/5d4a-j060

[81] P. Detti, G. Vatti, and G. Z. M. de Lara, "Eeg synchronization analysis for seizure prediction: A study on data of noninvasive recordings," *Processes*, vol. 8, no. 7, p. 846, 2020.

[82] R. G. Andrzejak, K. Lehnertz, C. Rieke, F. Mormann, P. David, and C. E. Elger, "Indications of nonlinear deterministic and finite-dimensional structures in time series of brain electrical activity: Dependence on recording region and brain state," *Physical Review E*, vol. 64, no. 6, p. 061907, 2001.

[83] V. Shah, E. Von Weltin, S. Lopez, J. R. McHugh, L. Veloso, M. Golmohammadi, I. Obeid, and J. Picone, "The temple university hospital seizure detection corpus," *Frontiers in neuroinformatics*, vol. 12, p. 83, 2018.

[84] P. Resque, A. Barros, D. Rosário, and E. Cerqueira, "An investigation of different machine learning approaches for epileptic seizure detection," in *2019 15th International Wireless Communications and Mobile Computing Conference (IWCMC)*, 2019, pp. 301–306.

[85] T. Ingolfsson, A. Cossettini, X. Wang, E. Tabanelli, G. Tagliavini, P. Ryvlin, L. Benini, and S. Benatti, "Towards long-term non-invasive monitoring for epilepsy via wearable eeg devices." in *IEEE Biomedical Circuits and Systems Conference*, 2021.

[86] D. Sopic, A. Aminifar, and D. Atienza, "e-glass: A wearable system for real-time detection of epileptic seizures," in *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, 2018, pp. 1–5.

[87] R. Zanetti, A. Aminifar, and D. Atienza, "Robust epileptic seizure detection on wearable systems with reduced false-alarm rate." in *42nd Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2020, pp. 4248–4251.

[88] B. M. A. Daoud Hisham, "Efficient epileptic seizure prediction based on deep learning." *IEEE transactions on biomedical circuits and systems*, vol. 13, pp. 804–813, 2019.

[89] I. Ullah, M. Hussain, H. Aboalsamh *et al.*, "An automated system for epilepsy detection using eeg brain signals based on deep learning approach," *Expert Systems with Applications*, vol. 107, pp. 61–71, 2018.

[90] P. Thuwajit, P. Rangpong, P. Sawangjai, P. Autthasan, R. Chaisaen, N. Banluesombatkul, P. Boonchit, N. Tatsaringkansakul, T. Sudhawiyangkul, and T. Wilaiprasitporn, "Eegwavenet: Multiscale cnn-based spatiotemporal feature extraction for eeg seizure detection," *IEEE transactions on industrial informatics*, vol. 18, no. 8, pp. 5547–5557, 2021.

[91] W. Huang, Y. Xue, L. Hu, and H. Liuli, "S-eegnet: Electroencephalogram signal classification based on a separable convolution neural network with bilinear interpolation," *IEEE Access*, vol. 8, pp. 131 636–131 646, 2020.

[92] T. M. Ingolfsson, U. Chakraborty, X. Wang, S. Beniczky, P. Ducouret, S. Benatti, P. Ryvlin, A. Cossettini, and L. Benini, "EpiDeNet: An energy-efficient approach to seizure detection for embedded systems," in *Proc. 2023 IEEE Biomedical Circuits and Systems Conference (BioCAS)*. IEEE, 2023, pp. 1–5.

[93] T. Alotaiby, F. E. A. El-Samie, S. A. Alshebeili, and I. Ahmad, "A review of channel selection algorithms for eeg signal processing," *EURASIP Journal on Advances in Signal Processing*, vol. 2015, pp. 1–21, 2015.

[94] J. Duun-Henriksen, T. W. Kjaer, R. E. Madsen, L. S. Remvig, C. E. Thomsen, and H. B. D. Sorensen, "Channel selection for automatic seizure detection," *Clinical Neurophysiology*, vol. 123, no. 1, pp. 84–92, 2012.

[95] M. Yildiz and E. Bergil, "The investigation of channel selection effects on epileptic analysis of eeg signals," *Journal of Electrical and Computer Engineering*, vol. 3, pp. 236–241, 2015.

[96] J. Birjandtalab, B. P. Maziyar, D. Cogan, N. Mehrdad, and J. Harvey, "Automated seizure detection using limited-channel eeg and non-linear dimension reduction." *Computers in biology and medicine*, vol. 82, pp. 49–58, 2017.

[97] A. A. E. Shoka, M. H. Alkinani, A. El-Sherbeny, A. El-Sayed, and D. M. M., "Automated seizure diagnosis system based on feature extraction and channel selection using eeg signals," *Brain Informatics*, vol. 8, no. 1, 2021.

[98] R. Jana and I. Mukherjee, "Deep learning based efficient epileptic seizure prediction with eeg channel optimization." *Biomedical Signal Processing and Control*, vol. 68, p. 102767, 2021.

[99] E. L. Glassman and J. V. Guttag, "Reducing the number of channels for an ambulatory patient-specific eeg-based epileptic seizure detector by applying recursive feature elimination." in *IEEE Engineering in Medicine and Biology Society.* IEEE, 2006, pp. 2175–2178.

[100] G. Peng, M. Nourani, J. Harvey, and H. Dave, "Personalized feature selection for wearable eeg monitoring platform," in *2020 IEEE 20th International Conference on Bioinformatics and Bioengineering (BIBE).* IEEE, 2020, pp. 380–386.

[101] V. Mihajlovic, B. Grundlehner, R. Vullers, and J. Penders, "Wearable, wireless eeg solutions in daily life applications: what are we missing?" *IEEE journal of biomedical and health informatics*, vol. 19, no. 1, pp. 6–21, 2014.

[102] M. E. Menshawy, A. Benharref, and M. Serhani, "An automatic mobile-health based approach for eeg epileptic seizures detection," *Expert systems with applications*, vol. 42, no. 20, pp. 7157–7174, 2015.

[103] L. Swinnen, C. Chatzichristos, K. Jansen, L. Lagae, C. Depondt, L. Seynaeve, E. Vancaester, A. Van Dycke, J. Macea, K. Vandecasteele *et al.*, "Accurate detection of typical absence

seizures in adults and children using a two-channel electroencephalographic wearable behind the ears," *Epilepsia*, vol. 62, no. 11, pp. 2741–2752, 2021.

[104] M. G. Bleichner and S. Debener, "Concealed, unobtrusive ear-centered eeg acquisition: ceegrids for transparent eeg," *Frontiers in human neuroscience*, vol. 11, p. 163, 2017.

[105] M. A. Frankel, M. J. Lehmkuhle, M. C. Spitz, B. J. Newman, S. V. Richards, and A. M. Arain, "Wearable reduced-channel eeg system for remote seizure monitoring," *Frontiers in Neurology*, vol. 12, p. 728484, 2021.

[106] X. Chen, Y. Zheng, Y. Niu, and C. Li, "Epilepsy classification for mining deeper relationships between eeg channels based on gcn," in *2020 international conference on computer vision, image and deep learning (CVIDL)*. IEEE, 2020, pp. 701–706.

[107] Q. Lian, Y. Qi, G. Pan, and Y. Wang, "Learning graph in graph convolutional neural networks for robust seizure prediction," *Journal of neural engineering*, vol. 17, no. 3, p. 035004, 2020.

[108] Y. Zhao, M. Xue, C. Dong, J. He, D. Chu, G. Zhang, F. Xu, X. Ge, and Y. Zheng, "Automatic seizure identification from eeg signals based on brain connectivity learning," *International Journal of Neural Systems*, vol. 32, no. 11, p. 2250050, 2022.

[109] K. Yan, X. Luo, L. Ye, W. Geng, J. He, J. Mu, X. Hou, X. Zan, J. Ma, F. Li *et al.*, "Automated seizure detection in epilepsy using a novel dynamic temporal-spatial graph attention network," *Scientific Reports*, vol. 15, no. 1, p. 16392, 2025.

[110] Y. Liu, H. Zhou, M. Guan, F. Feng, and J. Duan, "Scalp eeg-based automatic detection of epileptiform events via graph convolutional network and bi-directional lstm co-embedded broad learning system," *IEEE Signal Processing Letters*, vol. 30, pp. 448–452, 2023.

[111] J. Xiang, Y. Li, X. Wu, Y. Dong, X. Wen, and Y. Niu, "Synchronization-based graph spatio-temporal attention network for seizure prediction," *Scientific Reports*, vol. 15, no. 1, p. 4080, 2025.

[112] J. Xu, S. Yuan, J. Shang, J. Wang, K. Yan, and Y. Yang, "Spatiotemporal network based on gcn and bigru for seizure detection," *IEEE Journal of Biomedical and Health Informatics*, vol. 28, no. 4, pp. 2037–2046, 2024.

[113] A. Tegon, T. M. Ingolfsson, X. Wang, L. Benini, and Y. Li, "Femba: Efficient and scalable eeg analysis with a bidirectional mamba foundation model," *arXiv preprint arXiv:2502.06438*, 2025.

[114] S. Qiu, W. Wang, and H. Jiao, "Lightseizurenet: A lightweight deep learning model for real-time epileptic seizure detection," *IEEE Journal of Biomedical and Health Informatics*, vol. 27, no. 4, pp. 1845–1856, 2022.

[115] "A wearable system for real-time detection of epileptic seizures," patent WO2019162850A1.

[116] H. Ahmadi, E. Fatemizadeh, and A. Motie-Nasrabadi, "Deep sparse graph functional connectivity analysis in ad patients using fmri data," *Computer Methods and Programs in Biomedicine*, vol. 201, p. 105954, 2021.

[117] W. A. Rios-Herrera, P. V. Olguin-Rodriguez, J. D. Arzate-Mena, M. Corsi-Cabrera, J. Escalona, A. Marin-Garcia, J. Ramos-Loyo, A. L. Rivera, D. Rivera-Lopez, J. F. Zapata-Berruecos, and M. F. Muller, "The influence of eeg references on the analysis of spatio-temporal interrelation patterns," *Frontiers in neuroscience*, vol. 13, p. 941, 2019.

[118] G. Nolte, O. Bai, L. Wheaton, Z. Mari, S. Vorbach, and M. Hallett, "Identifying true brain interaction from eeg data using the imaginary part of coherency," *Clinical neurophysiology*, vol. 115, no. 10, pp. 2292–2307, 2004.

[119] G. Chen, Y. Liu, and X. Zhang, "Eeg–fnirs-based emotion recognition using graph convolution and capsule attention network," *Brain Sciences*, vol. 14, no. 8, p. 820, 2024.

[120] I. Siviero, G. Menegaz, and S. F. Storti, "Functional connectivity and feature fusion enhance multiclass motor-imagery brain–computer interface performance," *Sensors*, vol. 23, no. 17, p. 7520, 2023.

[121] M. Wu, R. Ouyang, C. Zhou, Z. Sun, F. Li, and P. Li, "A study on the combination of functional connection features and riemannian manifold in eeg emotion recognition," *Frontiers in Neuroscience*, vol. 17, p. 1345770, 2024.

[122] Z. Li, K. Hwang, K. Li, J. Wu, and T. Ji, "Graph-generative neural network for eeg-based epileptic seizure detection via discovery of dynamic brain functional connectivity," *Scientific reports*, vol. 12, no. 1, p. 18998, 2022.

[123] T. Barbera, S. Zini, S. Bianco, and P. Napoletano, "Lightweight graph neural network for dementia assessment from eeg recordings," in *2024 IEEE 8th Forum on Research and Technologies for Society and Industry Innovation (RTSI)*. IEEE, 2024, pp. 190–195.

[124] S. Mazurek, R. Blanco, J. Falco-Roget, and A. Crimi, "Explainable graph neural networks for eeg classification and seizure detection in epileptic patients," in *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*, 2024, pp. 1–5.

[125] Z. Ying, D. Bourgeois, J. You, M. Zitnik, and J. Leskovec, "Gnnexplainer: Generating explanations for graph neural networks," *Advances in neural information processing systems*, vol. 32, 2019.

[126] L. S. Vidyaratne and K. M. Iftekharuddin, "Real-time epileptic seizure detection using eeg," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 25, no. 11, pp. 2146–2156, 2017.

[127] A. B. Justesen, M. T. Foged, M. Fabricius, C. Skaarup, N. Hamrouni, T. Martens, O. B. Paulson, L. H. Pinborg, and S. Beniczky, "Diagnostic yield of high-density versus low-density eeg: The effect of spatial sampling, timing and duration of recording," *Clinical Neurophysiology*, vol. 130, no. 11, pp. 2060–2064, 2019.

[128] L. Hadady, T. Robinson, E. Bruno, M. P. Richardson, and S. Beniczky, "Users' perspectives and preferences on using wearables in epilepsy: A critical review," *Epilepsia*, Jan. 2025.

[129] M. A. Frankel, M. J. Lehmkuhle, M. Watson, K. Fetrow, L. Frey, C. Drees, and M. C. Spitz, "Electrographic seizure monitoring with a novel, wireless, single-channel eeg sensor," *Clinical Neurophysiology Practice*, vol. 6, pp. 172–178, 2021.

[130] S. You, B. H. Cho, Y.-M. Shon, D.-W. Seo, and I. Y. Kim, "Semi-supervised automatic seizure detection using personalized anomaly detecting variational autoencoder with behind-the-ear EEG," *Computer Methods and Programs in Biomedicine*, vol. 213, p. 106542, 2022.

[131] P. Busia, A. Cossettini, T. M. Ingolfsson, S. Benatti, A. Burrello, V. J. Jung, M. Scherer, M. A. Scrugli, A. Bernini, P. Ducouret *et al.*, "Reducing false alarms in wearable seizure detection with EEGformer: A compact transformer model for mcus," *IEEE Transactions on Biomedical Circuits and Systems*, 2024.

[132] R. S. Fisher, J. H. Cross, C. D'Souza, J. A. French, S. R. Haut, N. Higurashi, E. Hirsch, F. E. Jansen, L. Lagae, S. L. Moshe, J. Peltola, E. Roulet Perez, I. E. Scheffer, A. Schulze-Bonhage, E. Somerville, M. Sperling, E. M. Yacubian, and S. M. Zuberi, "Instruction manual for the ilae 2017 operational classification of seizure types," *Epilepsia*, vol. 58, no. 4, pp. 531–542, 2017.

[133] J. A. Ramirez-Bautista, S. L. Chaparro-Cárdenas, C. Esmer, and J. A. Huerta-Ruelas, "Artificial intelligence approaches to physiological parameter analysis in the monitoring and treatment of non-communicable diseases: A review," *Biomedical Signal Processing and Control*, vol. 87, p. 105463, 2024.

[134] N. S. Amer and S. B. Belhaouari, "Exploring new horizons in neuroscience disease detection through innovative visual signal analysis," *Scientific Reports*, vol. 14, no. 1, p. 4217, 2024.

[135] R. Jana and I. Mukherjee, "Efficient seizure prediction and eeg channel selection based on multi-objective optimization," *IEEE Access*, vol. 11, pp. 54 112–54 121, 2023.

[136] R. Gifford, C. Sollars, J. Singh, and S. Krening, "Locally-adaptive feature selection for nonconvulsive seizure detection," *Biomedical Signal Processing and Control*, vol. 105, p. 107535, 2025.

[137] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, vol. 58, no. 1, pp. 267–288, 1996.

[138] C. M. Bishop and N. M. Nasrabadi, *Pattern recognition and machine learning*. Springer, 2006, vol. 4, no. 4.

[139] X. Li, Y. Qiao, L. Duan, and J. Miao, "Epilepsy eeg signals classification based on sparse principal component logistic regression model," *Computer Methods in Biomechanics and Biomedical Engineering*, vol. 28, no. 8, pp. 1295–1303, 2025.

[140] J. Mateo, A. Torres, C. Soria, and J. Santos, "A method for removing noise from continuous brain signal recordings," *Computers & Electrical Engineering*, vol. 39, no. 5, pp. 1561–1570, 2013.

[141] H. Abdi and L. J. Williams, "Principal component analysis," *Wiley interdisciplinary reviews: computational statistics*, vol. 2, no. 4, pp. 433–459, 2010.

[142] F. Rahmat, Z. Zulkafli, A. J. Ishak, R. Z. Abdul Rahman, S. D. Stercke, W. Buytaert, W. Tahir, J. Ab Rahman, S. Ibrahim, and M. Ismail, "Supervised feature selection using principal component analysis," *Knowledge and Information Systems*, vol. 66, no. 3, pp. 1955–1995, 2024.

[143] Y. Sato, S. M. Wong, Y. Iimura, A. Ochi, S. M. Doesburg, and H. Otsubo, "Spatiotemporal changes in regularity of gamma oscillations contribute to focal ictogenesis," *Scientific reports*, vol. 7, no. 1, p. 9362, 2017.

[144] K. Pearson, "Vii. note on regression and inheritance in the case of two parents," *proceedings of the royal society of London*, vol. 58, no. 347-352, pp. 240–242, 1895.

[145] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *International conference on machine learning*. pmlr, 2015, pp. 448–456.

[146] "Eeg-based seizure detection using linear graph convolution network with focal loss," *Computer Methods and Programs in Biomedicine*, vol. 208, p. 106277, 2021.

[147] S. Nie, "Epilepsy detection based on graph convolutional neural network and transformer," in *Proceedings of BIO Web of Conferences*, vol. 111. EDP Sciences, Jan. 2024, p. 3017.

[148] F. Jibon, A. Tasbir, M. Miraz *et al.*, "Graph attention network and radial basis function neural network-based hybrid framework for epileptic seizure detection from eeg signal," *Journal of Autonomous Intelligence*, vol. 7, pp. 1–21, 2024.

[149] "Gcn-lstm: A hybrid graph convolutional network model for schizophrenia classification," *Biomedical Signal Processing and Control*, vol. 105, p. 107657, 2025.

[150] J. He, J. Cui, G. Zhang, M. Xue, D. Chu, and Y. Zhao, "Spatial-temporal seizure detection with graph attention network and bi-directional lstm architecture," *Biomedical Signal Processing and Control*, vol. 78, p. 103908, 2022.

[151] F. A. Jibon, M. R. Amin, M. A. Islam, K. M. Ahmed, and M. S. Hossain, "Sequential graph convolutional network and deeprnn based hybrid framework for epileptic seizure detection from eeg signal," *Digital Health*, vol. 10, p. 20552076241249874, May 2024.

[152] M. Awais, S. B. Belhaouari, and K. Kassoul, "Graphical insight: Revolutionizing seizure detection with eeg representation," *Biomedicines*, vol. 12, no. 6, 2024.

[153] D. Grattarola, L. Livi, C. Alippi, R. Wennberg, and T. A. Valiante, "Seizure localisation with attention-based graph neural networks," *Expert Systems with Applications*, vol. 203, p. 117330, 2022.

[154] A. T. Adebisi, H.-W. Lee, and K. C. Veluvolu, "Eeg-based brain functional network analysis for differential identification of dementia-related disorders and their onset," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 32, pp. 1198–1209, 2024.

[155] D. Büchel, T. Lehmann, Sandbakk, and J. Baumeister, "Eeg-derived brain graphs are reliable measures for exploring exercise-induced changes in brain networks," *Scientific Reports*, vol. 11, no. 1, p. 20803, 2021.

[156] L. di Biase, L. Ricci, M. L. Caminiti, P. M. Pecoraro, S. P. Carbone, and V. Di Lazzaro, "Quantitative high density eeg brain connectivity evaluation in parkinson's disease: the phase locking value (plv)," *Journal of Clinical Medicine*, vol. 12, no. 4, p. 1450, 2023.

[157] Z. Xu, S. Tang, C. Liu, Q. Zhang, H. Gu, X. Li, Z. Di, and Z. Li, "Temporal segmentation of eeg based on functional connectivity network structure," *Scientific Reports*, vol. 13, no. 1, p. 22566, 2023.

[158] D. Chicco, A. Sichenze, and G. Jurman, "A simple guide to the use of Student's t-test, Mann-Whitney U test, Chi-squared test, and Kruskal-Wallis test in biostatistics," *BioData Mining*, vol. 18, no. 1, p. 56, Aug. 2025.

[159] Y. Benjamini and Y. Hochberg, "Controlling the false discovery rate: a practical and powerful approach to multiple testing," *Journal of the Royal statistical society: series B (Methodological)*, vol. 57, no. 1, pp. 289–300, 1995.

[160] N. Cliff, "Dominance statistics: Ordinal analyses to answer ordinal questions." *Psychological bulletin*, vol. 114, no. 3, p. 494, 1993.

[161] Y. Wang, Y. Shi, Y. Cheng, Z. He, X. Wei, Z. Chen, and Y. Zhou, "A spatiotemporal graph attention network based on synchronization for epileptic seizure prediction," *IEEE Journal of Biomedical and Health Informatics*, vol. 27, no. 2, pp. 900–911, 2022.

[162] S. Abadal, P. Galvàn, A. Màrmol, N. Mammone, C. Ieracitano, M. L. Giudice, A. Salvini, and F. C. Morabito, "Graph neural networks for electroencephalogram analysis: Alzheimer's disease and epilepsy use cases," *Neural Networks*, vol. 181, p. 106792, 2025.

[163] V. Gabeff, T. Teijeiro, M. Zapater, L. Cammoun, S. Rheims, P. Ryvlin, and D. Atienza, "Interpreting deep learning models for epileptic seizure detection on eeg signals," *Artificial intelligence in medicine*, vol. 117, p. 102084, 2021.

[164] National Institute of Neurological Disorders and Stroke, "Epilepsy and seizures," https://www.ninds.nih.gov/

health-information/disorders/epilepsy-and-seizures, national Institutes of Health, accessed 2026-03-09.

[165] T. M. Ingolfsson, S. Benatti, X. Wang, A. Bernini, P. Ducouret, P. Ryvlin, S. Beniczky, L. Benini, and A. Cossettini, "Minimizing artifact-induced false-alarms for seizure detection in wearable eeg devices with gradient-boosted tree classifiers," *Scientific reports*, vol. 14, no. 1, p. 2980, 2024.

[166] S. Chakrabarti, A. Swetapadma, and P. K. Pattnaik, "A channel selection method for epileptic eeg signals," in *Emerging Technologies in Data Mining and Information Security: Proceedings of IEMIS 2018, Volume 1.* Springer, 2018, pp. 565–573.

[167] A. Bahr, M. Schneider, M. A. Francis, H. M. Lehmann, I. Barg, A.-S. Buschhoff, P. Wulff, T. Strunskus, and F. Faupel, "Epileptic seizure detection on an ultra-low-power embedded risc-v processor using a convolutional neural network," *Biosensors*, vol. 11, no. 7, p. 203, 2021.

[168] H. Ke, D. Chen, X. Li, Y. Tang, T. Shah, and R. Ranjan, "Towards brain big data classification: Epileptic eeg identification with a lightweight vggnet on global mic," *IEEE Access*, vol. 6, pp. 14 722–14 733, 2018.

[169] C. Wang, L. Liu, W. Zhuo, and Y. Xie, "An epileptic eeg detection method based on data augmentation and lightweight neural network," *IEEE Journal of Translational Engineering in Health and Medicine*, vol. 12, pp. 22–31, 2023.

[170] I. Dokare and S. Gupta, "Optimized seizure detection leveraging band-specific insights from limited eeg channels," *Health Information Science and Systems*, vol. 13, no. 1, p. 30, 2025.

[171] A. Affes, A. Mdhaffar, C. Triki, M. Jmaiel, and B. Freisleben, "Personalized attention-based eeg channel selection for epileptic seizure prediction," *Expert Systems with Applications*, vol. 206, p. 117733, 2022.

[172] M. Vendrame, M. Tracy, R. Das, F. Duffy, T. Loddenkemper, and S. V. Kothare, "Clinical correlations of midline spikes in children," *Epilepsy & Behavior*, vol. 18, no. 4, pp. 460–465, 2010.

[173] E. Ali, M. Angelova, and C. Karmakar, "Epileptic seizure detection using chb-mit dataset: The overlooked perspectives," *Royal Society open science*, vol. 11, no. 5, p. 230601, 2024.

[174] S. Varma and R. Simon, "Bias in error estimation when using cross-validation for model selection," *BMC Bioinformatics*, vol. 7, p. 91, 2006.

[175] J. Pilet, S. A. Beardsley, C. Carlson, C. T. Anderson, C. Ustine, S. Lew, W. Mueller, and M. Raghavan, "Predicting seizure onset zones from interictal intracranial eeg using functional connectivity and machine learning," *Scientific Reports*, vol. 15, no. 1, p. 17801, 2025.

[176] I. Dallmer-Zerbe, J. Kopal, A. Pidnebesna, J. Curot, M. Denuelle, A. De Barros, J.-C. Sol, L. Valton, E. Barbeau, and J. Hlinka, "Pro-ictal, rather than pre-ictal, brain state marked by global critical slowing and local gamma power increase," *Clinical Neurophysiology*, p. 2110742, 2025.

[177] S. J. Ruiz-Gómez, R. Hornero, J. Poza, A. Maturana-Candelas, N. Pinto, and C. Gomez, "Computational modeling of the effects of eeg volume conduction on functional connectivity metrics. application to alzheimer's disease continuum," *Journal of neural engineering*, vol. 16, no. 6, p. 066019, 2019.

[178] J. Pastor and L. Vega-Zelaya, "Normative structure of resting-state eeg in bipolar derivations for daily clinical practice: A pilot study," *Brain Sciences*, vol. 13, no. 2, p. 167, 2023.

[179] T. Mitsuhashi, Y. Iimura, H. Suzuki, T. Ueda, K. Nishioka, K. Nomura, M. Nakajima, H. Sugano, and A. Kondo, "Bipolar and laplacian montages are suitable for high-gamma modulation language mapping with stereoelectroencephalography," *Frontiers in Neurology*, vol. 15, p. 1380644, 2024.

[180] V. Harpale and V. Bairagi, "An adaptive method for feature selection and extraction for classification of epileptic eeg signal in significant states," *Journal of King Saud University-Computer and Information Sciences*, vol. 33, no. 6, pp. 668–676, 2021.

[181] C. Yang, M. Westover, and J. Sun, "Biot: Biosignal transformer for cross-data learning in the wild," *Advances in Neural Information Processing Systems*, vol. 36, pp. 78 240–78 260, 2023.

[182] W.-B. Jiang, L.-M. Zhao, and B.-L. Lu, "Large brain model for learning generic representations with tremendous eeg data in bci," *arXiv preprint arXiv:2405.18765*, 2024.

[183] G. Wang, W. Liu, Y. He, C. Xu, L. Ma, and H. Li, "Eegpt: Pretrained transformer for universal and reliable representation of eeg signals," *Advances in Neural Information Processing Systems*, vol. 37, pp. 39 249–39 280, 2024.

List of Figures

1.1 The six most typical EEG applications [8]. 7
1.2 An overview of the targets and contributions of the thesis, including the objectives, research areas, proposed methods, interpretability aspects, and conclusion. 8
1.3 Thesis structure. 11

2.1 The 10-20 electrode placement system based on electrode distances. (Created with app *BioRender* [21]) . 19
2.2 (a) Bipolar and (b) monopolar EEG montage. (Created with app *BioRender* [21]) 20
2.3 Distinction between focal and generalized seizures. . 22
2.4 Different classification tasks for diagnosis of epilepsy: (a) seizure detection; (b) seizure prediction [41]. . . . 25
2.5 Flowchart of a common seizure detection method. . . 27
2.6 Example of graph computation from EEG recordings. (a) Calculation of FC between EEG channels. (b) Graph representation: raw EEG data represent feature nodes, FC values represent edges between nodes. (Created with app *BioRender* [21]) 32

3.1 Nested donut chart of neural networks applied in epilepsy detection and prediction [67]. 44
3.2 Deep Pyramidal 1D-CNN architecture (P-1D-CNN) proposed by Ullah *et al.* [89]. 54
3.3 EEGNet architecture proposed by Lawhern *et al.* [69]. 55
3.4 Target architecture on which the NAS algorithm has been applied [74]. 56
3.5 Comparison between ML and DL in terms of energy-efficiency and performance [74]. 57

3.6 (A) Traditional ambulatory acquisition system. (B) Home-setting video-EEG monitoring. (C) *Emotiv Epoc* devices. (D) *Byteflies* Sensor Dot with four electrodes placed on the mastoid bone behind the ears of the patient and wired connected to the device attached to the upper back. (E) *cEEGrid* sensor characterized by 10 electrodes arranged in a C-shape and attached around the ear. (F) *Epilog* sensor attached through "stickers" that serve as the interface between Epilog and the scalp when used below hairline 61

3.7 Recent trends in EEG-GNN applications. 64

4.1 Block diagram of the proposed reduced-channel framework. Training phase: EEG channels are selected and used to train a lightweight CNN. Operating phase: the test records from the selected channels are processed, and the model output is post-processed to detect seizure events. 79

4.2 Channel weight coefficient calculation steps. (a) PCA loadings matrix computed in a 1-second time window. (b) PCA loadings matrix multiplied by the explained variance associated with each PC. (c) Weighted mean of the PCA loadings, obtained by summing up the values in the corresponding row computed in (b). (d) Topographic map representation of (c). 83

4.3 Temporal evolution of CWC decomposed into five clinical frequency ranges (a-e). x-axes represent time in seconds over 1 hour, while the y-axes show the 21 bipolar channels (defined in the text). The intensity of the CWCs is indicated by the colorbar in (e). Dotted t_1 and t_2 lines denote the start and end of the ictal event, respectively, according to the ground truth. The t_0 dotted line marks the beginning of the pre-ictal phase, characterized by trend discontinuities. 85

4.4 Channel selection method. EEG topographic maps showing the $\overline{\text{CWC}}$ values (normalized to the range [0,1]) for each channel, represented in an easily interpretable format. The maps were obtained from all the records of a single patient. **(a)** $\overline{\text{CWC}}$ values during the no-seizure phase. **(b)** $\overline{\text{CWC}}$ values during the ictal phase. In both maps, areas in red represent higher CWC values according to the colormap. **(c)** XOR map of the digitalized values (with a threshold at 0.9) from (a) and (b) highlighting the two most significant channels. 87

4.5 Schematic view of the proposed GNN-based pipeline for the detection of epileptic seizures. 97

4.6 FC matrices between ictal (left) and non-ictal (right) for (a) PLV measure and (b) PCC measure. The x and y axes represent bipolar channels, while each pixel in the matrices indicates FC value between each pair of channels. 100

4.7 FC matrix linearization procedure [157]. Linearization is useful for both statistical analysis and comparison between distributions. 102

4.8 Proposed model for graph-based EEG signal analysis. (Top) ChebNet feature extractor, trained to encode the spatial properties of EEG graphs. (Bottom) Full ST-GCN architecture, in which the pre-trained ChebNet extractor is combined with an LSTM to model temporal dependencies between sequences of three consecutive graphs. 105

5.1 Schematic representation of the leave-one-record-out (LORO) cross-validation procedure. At each iteration, one record is excluded from the dataset and used as test sample, while all the remaining records are used to train the model. The training set is further split into training (80%) and validation (20%) subsets. The process is repeated until each record has been used once as test. 113

5.2 Post-processing procedure for detecting epileptic seizures based on consecutive four-second segments with a two-second overlap, using a three-window buffer rule to define ictal and no-seizure phases. 114

5.3 Maps of the digitalized $\overline{\text{CWC}}$ values highlighting the most significant channels across all patients of the CHB-MIT dataset. 117

5.4 Maps of the digitalized $\overline{\text{CWC}}$ values highlighting the most significant channels across patients of the Siena dataset. 133

5.5 Channel selection framework for extracting a single channel in three configurations: (i) ictal, (ii) no-seizure (inter-ictal), and (iii) bad-channel. 135

5.6 Average CWCs maps derived from all the records of patient chb01 highlighting how the informative area shifts across different regions of the scalp during a seizure. 140

5.7 Schematic view of the four analyzed configurations, including Fully Generalized System (FGS), Hybrid System-Specific Model (HSSM), Hybrid System-Specific Channels (HSSC), Fully Personalized System (FPS). 146

5.8 Average CWCs maps highlighting the neural activity localization during the inter-ictal (a,c) and ictal phases (b,d). In particular, (a,b) are the CWC maps derived from all the records of patient chb01, taken as example, while (c,d) are the CWC maps derived from all the records of the dataset of the CHB-MIT dataset, i.e. across all the patient 148

5.9 Average CWCs maps highlighting the neural activity localization during the inter-ictal (a) and ictal phases (b). Both CWC maps derive from all the records of the Siena Scalp dataset, i.e. across all the patient. . 151

6.1 Boxplots of PLV and PCC values during seizure and non-seizure phases for patient chb01 from the CHB-MIT dataset. Dots represent the single FC values computed between each electrode pair. Symbol * denotes statistical significance. 165

6.2 Significant FC connections for PLV (a) and PCC (b) according to Cliff's effect size (δ) between ictal and no-seizure phases. Only channel pairs with statistically significant differences are shown. The color red ($\delta > 0$) indicates stronger connectivity in ictal; the color blue ($\delta < 0$) indicates stronger connectivity in no-seizure. 168

6.3 Relative probability distributions of connection strength for the two FC measures, i.e. PLV (a) and PCC (b). The histograms show the percentage of connections falling within each bin interval under ictal (light color) and non-ictal (dark color) conditions. 169
6.4 Balanced accuracy of the proposed ST-GCN model for patient chb01 as a function of the threshold τ, evaluated for values ranging from 0.1 to 0.6. 171
6.5 Graph representation showing the channels that appeared most frequently among the top-3 nodes identified by GNNExplainer in the analyzed ictal windows. The color intensity reflects the normalized occurrence of each channel, with darker nodes indicating higher frequency. 175

7.1 Visual synthesis of the main contributions of the thesis. 189

List of Tables

2.1 A detailed analysis of conventional EEG frequency bands and their associated cognitive and physiological states. The bands typically fall within a range of 0.5 Hz to 80 Hz [19]. 17

3.1 Main feature extraction domains adopted in EEG analysis for seizure detection, together with representative features commonly used in the literature. . 43

3.2 Overview of the two datasets used in this study for seizure detection. 49

3.3 Accuracy results and computational complexity of ML algorithms [84]. 51

3.4 Accuracy results and computational complexity (in terms of number of learnable parameters) of DL algorithms [88, 89] on the CHB-MIT [79] and Bonn [82] datasets, respectively. 53

4.1 Architecture of the proposed deep learning model. . 90

5.1 Selected channel pairs with their corresponding average correlation $\overline{\rho}$ for the CHB-MIT dataset. Channel A corresponds to the inter-ictal phase and Channel B to the ictal phase, respectively. 116

5.2 Performance comparison for patients in whom the two-selected-channel configuration outperformed the four-temporal configuration. 119

5.3 Performance comparison between the four-temp. and two-selected channels configurations for each patient. The final row shows the results (mean ± standard deviation) across all CHB patients. 120

5.4 Performance comparison for patients in whom the two-channel configuration performed slightly worse than the four-channel setup. 121
5.5 Performance comparison between the four-temporal and two-selected channels configurations across all CHB patients reported as mean ± standard deviation. 121
5.6 Comparison against existing literature approaches in terms of characteristics and achieved performance. . 126
5.7 Selected channel pairs with their corresponding average correlation $\overline{\rho}$ for Siena Scalp dataset. Channel A refers to the inter-ictal phase, while Channel B refers to the ictal phase. 132
5.8 Performance comparison between the four-temporal and two-selected channels configurations across all Siena scalp patients reported as mean ± standard deviation. 133
5.9 EEG channels selected for each patients of the CHB-MIT dataset. 137
5.10 EEG channels selected for each patient of the Siena Scalp dataset. 137
5.11 System performance across all configurations on the CHB-MIT dataset. 138
5.12 System performance across all configurations on the Siena Scalp dataset. 139
5.13 Selected channel pairs for individual patients and for the dataset including all patients (CHB-MIT dataset). 149
5.14 Selected channel pairs for individual patients and for the dataset including all patients (Siena Scalp dataset). 150
5.15 System performance for the two-optimized channel configuration on the CHB-MIT dataset. 152
5.16 System performance for the two-optimized channel configuration on the Siena Scalp dataset. 152
5.17 System performance for the four-fixed channel configuration on the CHB-MIT dataset. 155
5.18 System performance for the two HSSM configurations, one temporal pair and one central-temporal pair, on the CHB-MIT dataset. 156

6.1 Descriptive statistics of average PLV and PCC connectivities over time in ictal and non-ictal conditions. 165

6.2 Statistical comparison of PLV and PCC between ictal and non-ictal conditions across patients from the CHB-MIT dataset. The table reports the p-values obtained from the distributional comparisons and the number of significant connections detected for each FC measure using the FDR-corrected Mann-Whitney U test to control for multiple comparisons. 166

6.3 Statistical comparison of PLV and PCC between ictal and non-ictal conditions across patients from the Siena Scalp dataset. The table reports the p-values obtained from the distributional comparisons and the number of significant connections detected for each FC measure using the FDR-corrected Mann-Whitney U test to control for multiple comparisons. 167

6.4 Threshold τ value selected for each patient of the CHB-MIT dataset. 170

6.5 Threshold τ value selected for each patient of the Siena Scalp dataset. 170

6.6 Comparison of GNN-based models for EEG seizure detection on the CHB-MIT dataset, including the current state-of-the-art. 173