DETECTION OF EPILEPSY USING EEG SIGNALS AND MACHINE LEARNING

A Thesis Submitted

In Partial Fulfillment of the Requirements for the Degree of

MASTER OF TECHNOLOGY

in

Signal Processing and Digital Design

by

Bhargab Mahanta

24/SPD/21

Under the Supervision of

Dr. Rajesh Birok

Associate Professor, ECE Dept., DTU



Electronics and Communication Engineering

DELHI TECHNOLOGICAL UNIVERSITY

(Formerly Delhi College of Engineering)

Shahbad Daulatpur, Main Bawana Road, Delhi-110042,

India May, 2026



DELHI TECHNOLOGICAL UNIVERSITY

(Formerly Delhi College of Engineering)

Shahbad Daulatpur, Main Bawana Road, Delhi-42

ACKNOWLEDGEMENT

I wish to express my sincerest gratitude to Dr. Rajesh Birok for his continuous guidance and mentorship that he provided me during the project. He showed me the path to achieve my targets by explaining all the tasks to be done and explained to me the importance of this project as well as its industrial relevance. He was always ready to help me and clear my doubts regarding any hurdles in this project. Without his constant support and motivation, this project would not have been successful.

Bhargab Mahanta

Bhargab Mahanta

24/SPD/21

DELHI TECHNOLOGICAL UNIVERSITY

(Formerly Delhi College of Engineering)

Shahbad Daulatpur, Main Bawana Road, Delhi-42

CERTIFICATE BY THE SUPERVISOR

Certified that **Bhargab Mahanta**(24/SPD/21) has carried out their research work presented in this thesis entitled **"DETECTION OF EPILEPSY USING EEG SIGNALS AND MACHINE LEARNING"** for the award of **Master of Technology** from Department of Electronics and Communication Engineering, Delhi Technological University, Delhi, under my supervision. The thesis embodies results of original work, and studies are carried out by the student himself and the contents of the thesis do not form the basis for the award of any other degree to the candidate or to anybody else from this or any other University/Institution.

Signature

3-06-26

Dr. Rajesh Birok
Associate Professor,
ECE Dept., DTU

Date:

DELHI TECHNOLOGICAL UNIVERSITY

(Formerly Delhi College of Engineering) Shahbad Daulatpur

Main Bawana Road, Delhi-42

PLAGIARISM VERIFICATION

Title of the Thesis DETECTION OF EPILEPSY USING EEG SIGNALS AND MACHINE LEARNING

Total Pages Scholar 34 Name of the

Bhargab Mahanta

Supervisor(s)

(1) Dr. Rajesh Birtok

(2) _____

(3) _____

Department Electronics and Communications

This is to report that the above thesis was scanned for similarity detection.
Process and outcome is given below:

Software used: Turnitin Similarity Index: 6% Total Word Count: 8616

Date: 03/06/2026

Bhargab Mahanta
Candidate's Signature

03-06-26
Signature of Supervisor

DELHI TECHNOLOGICAL UNIVERSITY

(Formerly Delhi College of Engineering)

Shahbad Daulatpur, Main Bawana Road, Delhi-42

CANDIDATE'S DECLARATION

I hereby certify that the work being presented in the thesis entitled **"DETECTION OF EPILEPSY USING EEG SIGNALS AND MACHINE LEARNING"** in partial fulfillment of the requirements for the award of the degree of Master of Technology submitted in the Department of Electronics and Communication Engineering, Delhi Technological University, is an authentic record of my own work carried out during the period under the supervision of Dr. Rajesh Birok.

The matter presented in the thesis has not been submitted by me for the award of any other degree of this or any other Institute.

Bhargab Mahanta
24/SPD/21

This is to certify that the student has incorporated all the corrections suggested by the examiners in the thesis and the statement made by the candidate is correct to the best of our knowledge.

Signature of Supervisor **Signature of External Examiner**

Detection of Epilepsy Using EEG Signals and Machine Learning

ABSTRACT

Epilepsy is a debilitating chronic neurological condition characterised by recurrent, spontaneous seizures, affecting an estimated 50 million individuals globally across all demographic groups. Rapid and accurate seizure identification is indispensable for ensuring patient safety and guiding clinical decision-making. Electroencephalography (EEG) remains the principal neurophysiological tool for epilepsy diagnosis, offering a direct window into cerebral electrical dynamics; however, its manual review is labour-intensive, subjective, and heavily dependent on specialist expertise that is scarce in many healthcare settings. This thesis presents a systematic and comprehensive investigation of automated epilepsy seizure detection from EEG signals, employing a suite of machine learning (ML) and deep learning approaches.

The developed pipeline integrates four sequential stages: signal acquisition and preprocessing, multi-domain feature extraction (spanning time-domain, frequency-domain, and time-frequency representations), classification via multiple ML models, and rigorous performance evaluation. The publicly available University of Bonn EEG dataset is used throughout to ensure reproducibility and enable fair cross-method comparisons. Feature descriptors encompass time-domain statistical moments, Hjorth parameters, FFT-derived spectral power indices, Discrete Wavelet Transform (DWT) sub-band energies, and non-linear complexity measures including approximate entropy and sample entropy.

Five classifiers are systematically benchmarked: Support Vector Machine (SVM), Random Forest (RF), k-Nearest Neighbour (k-NN), Gradient Boosting, and a one-dimensional Convolutional Neural Network (CNN). Experimental evaluation shows that SVM with wavelet-based features achieves the best performance among conventional classifiers, reaching 98.7% accuracy, 97.9% sensitivity, and 99.2% specificity on the binary seizure-versus-non-seizure task. The end-to-end CNN operating directly on raw EEG windows attains 99.1% accuracy, demonstrating the substantial representational advantage of deep learning over handcrafted feature pipelines.

This work delivers a reproducible, end-to-end automated seizure detection framework with demonstrated potential for integration into clinical decision-support and real-time monitoring systems. The study further provides a thorough discussion of the practical limitations, translational challenges, future research priorities, and clinical implications of the proposed methodology.

Keywords: Epilepsy detection, EEG signals, machine learning, seizure classification, wavelet transform, support vector machine, deep learning, convolutional neural network.

TABLE OF CONTENTS

Title i
Acknowledgements ii
Candidate's Declarations iii
Certificate by the Supervisor iv
Abstract v
Table of Contents vi
List of Tables vii
List of Figures viii
List of Symbols, Abbreviations and Nomenclature ix
CHAPTER 1: INTRODUCTION 1
1.1 Background and Motivation 1
1.2 Problem Statement 1
1.3 Research Objectives 2
1.4 Scope and Limitations 2
1.5 Thesis Organization 3
CHAPTER 2: LITERATURE REVIEW 4
2.1 Epilepsy: Clinical Background 4
2.2 EEG in Epilepsy Diagnosis 4
2.3 Traditional Signal Processing Approaches 5
2.4 Machine Learning for EEG-based Epilepsy Detection 6
2.5 Deep Learning Approaches 7
2.6 Research Gaps and Contributions 8
CHAPTER 3: METHODOLOGY 9
3.1 Dataset Description 9
3.2 Preprocessing Pipeline 9
3.3 Feature Extraction 10
3.4 Classification Algorithms 11
3.5 Evaluation Metrics 12
CHAPTER 4: RESULTS AND DISCUSSION 14
4.1 Preprocessing Results 14
4.2 Feature Analysis 14
4.3 Classifier Performance Comparison 15
4.4 Deep Learning Results 17
4.5 Discussion 17
CHAPTER 5: CONCLUSION, FUTURE SCOPE AND SOCIAL IMPACT 19
5.1 Summary of Contributions 19

5.2 Future Directions 19
5.3 Social Impact 20
References 21
Curriculum Vitae/Brief Profile 24

LIST OF SYMBOLS, ABBREVIATIONS AND NOMENCLATURE

ANN	Artificial Neural Network
ApEn	Approximate Entropy
AUC	Area Under the ROC Curve
CNN	Convolutional Neural Network
DWT	Discrete Wavelet Transform
EEG	Electroencephalography
EMD	Empirical Mode Decomposition
FFT	Fast Fourier Transform
FN	False Negative
FP	False Positive
GNN	Graph Neural Network
HHT	Hilbert-Huang Transform
ICA	Independent Component Analysis
IED	Interictal Epileptiform Discharge
IMF	Intrinsic Mode Function
k-NN	k-Nearest Neighbor
LDA	Linear Discriminant Analysis
LSTM	Long Short-Term Memory
ML	Machine Learning
PCA	Principal Component Analysis
PSD	Power Spectral Density
RBF	Radial Basis Function
RF	Random Forest
RNN	Recurrent Neural Network
ROC	Receiver Operating Characteristic
SampEn	Sample Entropy
SNR	Signal-to-Noise Ratio
STFT	Short-Time Fourier Transform
SUDEP	Sudden Unexpected Death in Epilepsy
SVM	Support Vector Machine
TN	True Negative
TP	True Positive
WHO	World Health Organization

CHAPTER 1

INTRODUCTION

1.1 Background and Motivation

Epilepsy is among the most burdensome chronic brain conditions globally, characterized by spontaneously recurring seizures that arise from abnormal, hypersynchronous electrical discharges within neural networks. The World Health Organization estimates that over 50 million individuals across all age groups are living with epilepsy, placing it among the most significant neurological health challenges worldwide (WHO, 2022). The condition transcends geographic, ethnic, and socioeconomic boundaries; however, its impact is disproportionately concentrated in low- and middle-income nations, where nearly 80% of patients reside yet fewer than 10% receive timely and adequate diagnosis or treatment (WHO, 2022).

Seizure manifestations span a wide clinical spectrum, ranging from brief lapses in awareness (absence seizures) to violent convulsions involving the limbs and trunk (tonic-clonic seizures), each capable of causing physical harm, cognitive impairment, and a lasting decline in quality of life. Despite considerable advances in anti-epileptic drug (AED) pharmacotherapy, roughly one-third of patients remain refractory to medication, making drug-resistant epilepsy a persistent therapeutic challenge. For such individuals, precise seizure identification becomes indispensable in evaluating candidacy for interventions such as surgical resection of the seizure focus, Vagus Nerve Stimulation, or closed-loop Responsive Neurostimulation.

Electroencephalography (EEG) serves as the cornerstone diagnostic modality for epilepsy, capturing scalp-electrode recordings of the brain's electrical activity to identify pathological patterns such as interictal epileptiform discharges (IEDs), ictal rhythms, and post-ictal suppression. Continuous EEG monitoring may span several days, generating enormous volumes of data that must be thoroughly reviewed. The manual interpretation of such prolonged recordings is inherently labour-intensive and cognitively taxing, rendering it susceptible to inter-rater inconsistency and observer fatigue. A single night of EEG monitoring can produce gigabytes of data, and detecting epileptiform events within it demands extensive neurological expertise — a resource that is critically scarce in many developing-world healthcare settings.

The application of machine learning (ML) to large-scale EEG repositories offers a compelling solution to these limitations. Trained on the distinctive electrophysiological signatures of normal, interictal, and ictal brain states, ML models can perform automated classification with accuracy approaching that of expert neurologists. The last decade has witnessed rapid growth in research applying both conventional ML methods — including Support Vector Machines (SVM), Random Forests, and k-Nearest Neighbour classifiers — and advanced deep learning architectures such as Convolutional Neural Networks (CNN) and Recurrent Neural Networks (RNN) to the problem of automated seizure detection and diagnosis.

1.2 Problem Statement

Building a reliable automated system for epileptic seizure detection from EEG signals remains an open and formidable technical challenge. To begin with, EEG recordings are inherently non-stationary and non-linear, exhibiting pronounced variability both across subjects and within the same individual over time. Noise contamination from ocular movements, muscular activity, and electrode artefacts further complicates signal interpretation. Additionally, extracting informative features from high-dimensional, multi-channel EEG demands substantial domain expertise, and many existing pipelines rely on manually crafted features with limited generalisability. A further complication arises from the pronounced class imbalance between ictal and non-ictal segments in long-term recordings. Finally, the scarcity of large, rigorously annotated, publicly accessible EEG datasets constrains the development of robust and generalisable models. Many published systems exhibit methodological shortcomings — including data leakage, insufficient cross-validation, and reliance on unrepresentative subsets — making the translation of laboratory results into clinically deployable systems contingent on transparent, rigorously validated experimental design.

1.3 Research Objectives

The above issues are approached within this thesis, where the main objectives are:

- To design and develop an effective preprocessing pipeline for EEG signal denoising and artifact removal.
- To extract and compute a rich set of time-domain, frequency-domain, and time-frequency domain features from EEG signal.
- To build and compare different machine learning classifiers including SVM, Random Forest, k-NN, and Gradient Boosting for epilepsy detection.
- To construct a Convolutional Neural Network (CNN)-based deep learning model that achieves end-to-end epilepsy classification from raw EEG signals..
- To compare the developed models with the state-of-the-art by reporting results with standard evaluation metrics.

1.4 Scope and Limitations

The scope of this work is confined to offline (non-real-time) epilepsy detection using the widely available University of Bonn single-channel EEG database. Multi-channel analysis and source localisation are outside the present scope. While benchmark performance is encouraging, extending these methods to more heterogeneous patient populations, alternative EEG acquisition hardware, and varied recording settings constitutes important future work. Considerations regarding computational latency and hardware constraints for real-time deployment are discussed conceptually rather than implemented.

1.5 Thesis Organization

The remainder of this thesis is structured as follows. Chapter 2 presents an extensive review of the relevant literature, covering the clinical background of

epilepsy, the neurophysiological basis of EEG signals, classical signal processing approaches, and the state of the art in machine learning and deep learning for seizure detection. Chapter 3 details the proposed methodology, encompassing dataset description, the preprocessing pipeline, feature extraction strategies, and the classification algorithms employed. Chapter 4 presents and critically analyses the experimental results. Chapter 5 draws conclusions, outlines directions for future research, and reflects on the broader social implications of this work.

CHAPTER 2

LITERATURE REVIEW

2.1 Epilepsy: Clinical Background

The International League Against Epilepsy (ILAE) formally defines epilepsy as a brain disorder characterised by a sustained vulnerability to spontaneous seizure generation, together with the attendant neurobiological, cognitive, psychological, and social consequences (Fisher et al., 2014). Clinically, a diagnosis of epilepsy requires either a history of two or more unprovoked seizures separated by at least 24 hours, or a single unprovoked seizure accompanied by an elevated probability of recurrence.

At the cellular level, seizures result from pathological hypersynchronous neuronal discharges. Depending on whether the initial discharge is confined to one cortical region or engages both hemispheres simultaneously, seizures are categorised as focal (partial) or generalised, respectively; focal seizures may also undergo secondary generalisation. The clinical phenotype varies considerably: absence seizures present as transient interruptions of consciousness without motor involvement; tonic-clonic seizures feature sustained muscular rigidity followed by rhythmic convulsions with loss of consciousness; myoclonic seizures manifest as brief, shock-like involuntary jerks; and atonic seizures cause abrupt loss of postural muscle tone.

The aetiological landscape of epilepsy is heterogeneous. Structural abnormalities — including cortical dysplasias, hippocampal sclerosis (the predominant substrate in temporal lobe epilepsy), space-occupying lesions, and traumatic brain injury — represent major contributors. Genetic mutations affecting ion channel function (channelopathies) underlie a substantial proportion of idiopathic epilepsy syndromes. Metabolic disturbances, autoimmune encephalitis, and infectious conditions such as neurocysticercosis constitute additional aetiological categories. In a considerable subset of patients, no identifiable cause is established despite thorough investigation, a situation formerly designated as cryptogenic epilepsy.

Beyond its neurological effects, epilepsy exerts a substantial socioeconomic toll, restricting driving privileges, limiting occupational participation, and perpetuating social stigma that discourages care-seeking. Critically, epilepsy carries a non-trivial mortality risk: population-based studies report a rate of approximately 1.16 cases of sudden unexpected death in epilepsy (SUDEP) per 1000 patient-years in community cohorts (Thurman et al., 2014), underscoring that even seemingly benign uncontrolled seizures can prove fatal. There is therefore compelling motivation to develop reliable automated seizure monitoring systems capable of alerting patients or caregivers in real time, thereby potentially reducing SUDEP incidence.

2.2 EEG in Epilepsy Diagnosis

The electroencephalogram (EEG), first described by Hans Berger in 1929,

measures voltage fluctuations generated by the aggregate postsynaptic currents of cortical pyramidal neurons. In clinical practice, scalp EEG employs between 19 and 256 electrodes positioned according to the standardised international 10-20 system. The recorded signals decompose into characteristic frequency bands that correlate with distinct physiological and pathological brain states: delta (0.5–4 Hz), theta (4–8 Hz), alpha (8–13 Hz), beta (13–30 Hz), and gamma (>30 Hz).

Within epilepsy management, EEG serves multiple clinical roles: confirming the diagnosis through detection of epileptiform discharges, classifying the epilepsy syndrome, delineating the seizure onset zone for pre-surgical evaluation, and gauging therapeutic response. The hallmark EEG signatures of epilepsy encompass interictal epileptiform discharges (IEDs) such as spikes and spike-and-wave complexes, ictal patterns characterised by evolving rhythmic discharges, and post-ictal suppression with generalised slowing.

Prolonged video-EEG monitoring (LTM) is particularly valuable for pre-surgical evaluation, as it enables simultaneous behavioural and electrographic documentation of habitual seizures. Ambulatory EEG systems facilitate data acquisition outside the hospital environment. When scalp recordings are insufficient, invasive intracranial EEG (iEEG) — employing depth electrodes or subdural grid arrays — provides superior spatial precision for seizure onset localisation, though at the cost of procedural risk.risks.

Artefact identification and rejection constitutes one of the most demanding aspects of EEG interpretation. The principal sources of contamination include ocular artefacts — eye blinks and saccadic movements generating low-frequency noise (1–3 Hz) most prominent in frontal derivations — myogenic (EMG) artefacts above 30 Hz, electrode ‘pop’ artefacts presenting as high-amplitude transients, cardiac ECG contamination, and motion artefacts. Their morphological resemblance to genuine epileptiform discharges makes reliable differentiation challenging and represents a key bottleneck for automated detection systems.

2.3 Traditional Signal Processing Approaches

Early automated EEG seizure detection relied on deterministic, rule-based and threshold-based algorithms designed to flag patterns characteristic of epileptic activity. Among the pioneering contributions, Gotman (1982) introduced an algorithm exploiting half-wave decomposition to isolate seizure-specific rhythmic activity. While groundbreaking for their era, such rule-based approaches exhibited limited adaptability to the diversity of seizure semiology and inter-electrode variability encountered in clinical practice.

Classical spectral analysis via the Fourier Transform enables characterisation of EEG power distribution across frequency bands; however, its fundamental assumption of signal stationarity is violated by the inherently non-stationary nature of EEG. The Short-Time Fourier Transform (STFT) partially addresses this by computing localised spectra over sliding temporal windows, thereby providing joint time-frequency information. Its limitation, however, lies in the fixed window length, which imposes an irreconcilable trade-off between temporal and frequency resolution governed by the Heisenberg uncertainty principle.

Wavelet analysis overcomes the fixed-resolution constraint through multi-resolution decomposition, offering fine temporal resolution at high frequencies and fine frequency resolution at low frequencies simultaneously. The Discrete Wavelet Transform (DWT) achieves this by iteratively projecting the signal onto a set of scaled and translated basis functions, yielding approximation and detail coefficients at multiple scales. This multi-scale decomposition has become a standard preprocessing step for extracting band-specific energy, entropy, and statistical descriptors from EEG signals in the context of epilepsy detection (Adeli et al., 2003; Subasi, 2005).

Empirical Mode Decomposition (EMD) offers a fully data-driven alternative that decomposes non-linear, non-stationary signals into a set of Intrinsic Mode Functions (IMFs) without requiring a predetermined basis. When combined with Hilbert spectral analysis — forming the Hilbert-Huang Transform (HHT) — it has demonstrated utility in epileptic seizure detection (Bajaj and Pachori, 2011). Complementing these transform-based approaches, non-linear dynamical measures such as approximate entropy (ApEn), sample entropy (SampEn), Lyapunov exponents, fractal dimension, and Hjorth parameters quantify the irregularity and complexity of EEG dynamics, both of which undergo marked changes during seizure activity.

2.4 Machine Learning for EEG-based Epilepsy Detection

The emergence of machine learning marked a paradigm shift in EEG-based epilepsy detection, moving away from rigid rule-based heuristics toward flexible, data-driven classifiers. A comparative evaluation of SVM, ANN, and k-NN algorithms employing DWT-derived features by Subasi and Gursoy (2010) yielded classification accuracies reaching 95% for SVM on the Bonn dataset. Shoeb and Guttag (2010) independently demonstrated the utility of patient-individualised SVM models trained on channel selection and spectral power features, attaining 96% sensitivity for seizure onset detection.

Random Forest (RF), introduced by Breiman (2001), aggregates predictions from an ensemble of decorrelated decision trees trained on bootstrap samples of the data. Its inherent resistance to overfitting, ability to model high-dimensional feature spaces, and provision of intrinsic feature importance rankings have made it a widely adopted tool for EEG-based epilepsy research. Multiple studies have reported RF classification accuracies exceeding 95% (Alickovic et al., 2018; Yuan et al., 2012), while feature importance scores derived from RF models have provided neurophysiologically interpretable insights into the most discriminative EEG characteristics.

The Support Vector Machine (SVM) operates by locating the decision boundary that maximises the margin between classes in a transformed feature space. Non-linear kernel functions — most commonly the Radial Basis Function (RBF) and polynomial kernels — project the input features into higher-dimensional spaces where linear separability is achieved. This property makes SVM particularly well-suited to clinical EEG datasets, which are typically characterised by high dimensionality and limited sample sizes. Nicolaou and Georgiou (2012) exploited permutation entropy features within an SVM framework, reporting 97.77% accuracy

for seizure detection. Systematic meta-analyses consistently identify SVM as the most frequently applied and highest-performing conventional classifier in EEG-based epilepsy research.

The k-Nearest Neighbour (k-NN) algorithm assigns a query sample to the majority class among its k closest neighbours in the feature space. Although conceptually straightforward and interpretable, its performance is sensitive to the choice of k, the selected distance metric, and the presence of irrelevant features. Gradient boosting frameworks — notably XGBoost and LightGBM — construct an additive ensemble of shallow learners by sequentially minimising residual errors, and are regarded as state-of-the-art for tabular classification tasks. Their application to EEG-based seizure detection has yielded competitive results across multiple studies (Tzimourta et al., 2019).

When the extracted EEG feature vector is high-dimensional, dimensionality reduction and feature selection become essential preprocessing steps. Techniques including Principal Component Analysis (PCA), Linear Discriminant Analysis (LDA), filter-based methods (e.g., mutual information, correlation coefficients), and wrapper-based selection strategies are routinely applied to prune redundant features, improve classifier generalisation, and reduce computational cost.

2.5 Deep Learning Approaches

Deep learning architectures, particularly Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), have substantially advanced the state of the art in automated EEG analysis. Their key advantage is the capacity for end-to-end learning: rather than relying on manually engineered features, these models discover hierarchically structured representations directly from raw or minimally processed input signals.

CNNs apply learnable convolutional filters that capture localised temporal and spectral patterns across multiple scales. Acharya et al. (2018) applied a 13-layer CNN directly to EEG segments from the Bonn dataset, achieving 88.67% seizure detection accuracy. Subsequent work incorporating deeper architectures, data augmentation, and batch normalisation pushed accuracy beyond 99%. CNNs have additionally been applied to 2D time-frequency representations derived from wavelet scalograms and spectrograms, treating them as image classification problems.

Modelling the temporal dynamics of sequential EEG data is a natural application for recurrent architectures, especially Long Short-Term Memory (LSTM) networks. LSTMs incorporate gating mechanisms that selectively retain and update an internal memory state, enabling them to capture long-range temporal dependencies relevant to evolving seizure patterns. Zhao et al. (2021) proposed a hybrid CNN-LSTM framework in which a CNN first extracts spatial features before an LSTM models the temporal evolution, yielding improved detection performance relative to either architecture alone.

Transformer architectures, originally developed for natural language processing, have been adapted for EEG signal analysis. Their self-attention mechanism allows the model to dynamically weight different temporal positions, focusing on clinically salient events such as epileptiform transients embedded within

lengthy recordings. Graph Neural Networks (GNNs) offer a complementary approach by explicitly encoding the topology of the EEG electrode network, thereby capturing spatio-temporal dependencies across channels. Transfer learning strategies, which initialise deep models on large-scale EEG corpora and then fine-tune on smaller clinical datasets, address the chronic shortage of labelled epilepsy data and have shown promising results in cross-patient generalisation.

2.6 Research Gaps and Contributions

Notwithstanding the rapid progress in ML-based EEG analysis, several methodological gaps remain. A widespread weakness is reliance on evaluation protocols that partition data in ways that inflate reported performance metrics relative to what would be observed in genuine prospective clinical deployment. Furthermore, fair cross-study benchmarking is impeded by heterogeneous experimental setups that vary in dataset splits, feature sets, and evaluation metrics. A third gap concerns model transparency: the interpretability of learned representations — essential both for clinical credibility and regulatory scrutiny — is rarely addressed with rigour in current literature. Fourth, the pronounced imbalance between seizure and non-seizure segments and the substantial inter-individual variability in EEG morphology remain incompletely resolved challenges in most proposed systems.

The present thesis is designed to address these shortcomings through: (i) a reproducible, standardised preprocessing and feature extraction pipeline; (ii) a rigorous side-by-side evaluation of five classifiers under identical experimental conditions; (iii) a 1D CNN trained end-to-end on raw EEG; (iv) feature importance analysis to furnish neurophysiologically grounded explanations; and (v) a critical appraisal of the clinical translatability of the proposed system.

CHAPTER 3

METHODOLOGY

3.1 Dataset Description

The widely-used benchmark in EEG epilepsy research is the University of Bonn EEG dataset (Andrzejak et al., 2001), which provides a well-established basis for reproducible and comparable evaluation. It comprises five subsets (A through E), each containing 100 single-channel EEG epochs of 23.6 seconds duration (4097 samples per epoch) acquired at a sampling frequency of 173.61 Hz.

Set A: Surface (scalp) EEG acquired from neurologically healthy volunteers in an awake, eyes-open condition.

Set B: Scalp EEG recordings from the same healthy volunteer group, obtained during a relaxed, eyes-closed resting state.

Set C: Interictal intracranial EEG captured from epileptic patients, recorded from the cerebral hemisphere contralateral to the seizure focus (i.e., seizure-free hemisphere).

Set D: Intracranial EEG data from epilepsy patients recorded during clinically confirmed seizure-free intervals.

Set E: Intracranial EEG epochs containing active seizure (ictal) activity.

Three binary discrimination tasks of escalating clinical difficulty are investigated: (1) Set E against the combined non-seizure sets A+B+C+D, representing the broadest seizure-versus-non-seizure problem; (2) Set E versus Set A, distinguishing ictal from healthy normal activity; and (3) Set E versus Sets C+D, the most demanding scenario requiring differentiation of ictal from interictal states. A five-class problem encompassing all subsets is additionally explored.

All EEG segments are treated as temporally independent units with no overlap between training and test partitions. Binary ground-truth labels are assigned as 0 (non-seizure) and 1 (seizure). Identical preprocessing and feature extraction procedures are applied uniformly across all segments to ensure experimental consistency.

3.2 Preprocessing Pipeline

3.2.1 Artifact Removal

A fourth-order Butterworth band-pass filter with a 50 Hz upper cut-off is applied to the raw EEG signals, attenuating DC offsets, baseline drift below 0.5 Hz, high-frequency myogenic artefacts, and power-line interference above 40 Hz. For the Bonn dataset, which includes both intracranial (Sets C, D, E) and scalp (Sets A, B) recordings, the filter passband is carefully configured to preserve high-frequency gamma activity in the 30–40 Hz range, given its known relevance to ictal dynamics.

For multi-channel datasets containing oculographic or cardiac artefacts,

Independent Component Analysis (ICA) is available as an optional processing step; however, it was not applied to the Bonn data, which consists of single-channel recordings. Residual power-line interference is suppressed using 50/60 Hz notch filters. Segments exhibiting peak-to-peak amplitudes exceeding 500 μV are flagged as potentially corrupted and excluded from subsequent analysis.

3.2.2 Normalization

All EEG epochs are then standardised via z-score transformation to zero mean and unit variance. This step ensures that features derived from different segments and subjects occupy a comparable numerical range, which is a prerequisite for distance-sensitive classifiers (k-NN) and kernel-based methods (SVM). For neural network inputs, min-max rescaling to the [0, 1] interval may alternatively be employed.

3.2.3 Segmentation

Each 4097-sample Bonn dataset epoch serves as the fundamental analysis unit for conventional ML classifiers. For the deep learning model, each epoch is further subdivided into non-overlapping sub-windows of 256 samples (approximately 1.47 seconds) to increase effective training set size while preserving local temporal context. Overlapping windows at 50% stride are additionally employed as a data augmentation strategy.

3.3 Feature Extraction

Feature extraction proceeds across three complementary domains, yielding a consolidated 36-dimensional feature vector per EEG segment:

3.3.1 Time-Domain Features

Time-domain features characterise the amplitude statistics and morphological structure of the raw EEG waveform without frequency transformation. The following ten scalar descriptors are computed:

(1) Mean: central tendency of the amplitude distribution. (2) Variance: spread of signal amplitude around the mean, reflecting signal power. (3) Standard Deviation: square root of variance, in amplitude units. (4) Skewness: asymmetry of the amplitude distribution. (5) Kurtosis: peakedness or 'tailedness' of the distribution; elevated kurtosis is indicative of impulsive spike-like transients. (6) Maximum amplitude: largest positive sample value. (7) Minimum amplitude: most negative sample value. (8) Peak-to-peak amplitude: arithmetic difference of maximum and minimum values. (9) Zero-crossing rate: number of zero-crossings per unit time, reflecting dominant oscillation frequency. (10) Line length: cumulative sum of absolute sample-to-sample differences, quantifying signal complexity and activity.

Three Hjorth descriptors are additionally computed: Activity (signal variance, reflecting power), Mobility (square root of the ratio of first-derivative variance to signal variance, approximating the mean frequency), and Complexity (ratio of the Mobility of the first derivative to the Mobility of the signal, reflecting spectral bandwidth). Collectively, these nine time-domain and Hjorth measures summarise

the signal's power, frequency content, and waveform complexity.

3.3.2 Frequency-Domain Features

The power spectral density (PSD) of each EEG segment is estimated using Welch's periodogram method with Hann-windowed sub-segments of 256 samples at 50% overlap. Absolute and relative band power values are computed for the five canonical EEG frequency bands: delta (0.5–4 Hz), theta (4–8 Hz), alpha (8–13 Hz), beta (13–30 Hz), and gamma (30–40 Hz). A further twelve spectral features are extracted, including spectral entropy, peak frequency, median frequency, and the spectral edge frequency at 90% cumulative power.

3.3.3 Time-Frequency Features (DWT)

A five-level DWT decomposition using the Daubechies-4 (db4) mother wavelet is applied to each EEG epoch, producing five detail sub-bands (D1–D5) spanning progressively lower frequency ranges and one approximation sub-band (A5) encompassing frequencies below 5.42 Hz. Six statistical descriptors — mean, variance, standard deviation, maximum, minimum, and energy (sum of squared coefficients) — are computed from each sub-band, yielding 36 DWT-based features in total across all decomposition levels.

Irregularity and complexity of EEG dynamics are quantified using non-linear features such as Approximate Entropy (ApEn) and Sample Entropy (SampEn). The approximate entropy is calculated for m=2 and for r=0.2 times the standard deviation of the signal. Such entropy measures have been found to drop markedly when seizures occur, because the dynamics of the neurons become more synchronized, more regular.

3.4 Classification Algorithms

3.4.1 Support Vector Machine (SVM)

The SVM classifier seeks the decision boundary that maximises the geometric margin between the two classes in a transformed feature space. When the classes are not linearly separable in the original space, the RBF kernel function $K(x_i, x_j) = \exp(-\gamma\|x_i - x_j\|^2)$ implicitly maps samples into a higher-dimensional Hilbert space where linear separation becomes feasible. The two critical hyperparameters — C (regularisation strength) and γ (kernel width) — are jointly optimised through 5-fold cross-validated grid search over $C \in \{0.1, 1, 10, 100\}$ and $\gamma \in \{0.001, 0.01, 0.1, 1\}$.

3.4.2 Random Forest (RF)

The RF ensemble is constructed from 200 individual decision trees, each trained on an independently drawn bootstrap sample of the training data. At each internal node, a random subset of $\sqrt{(n_features)}$ candidate features is evaluated for the optimal split, promoting diversity among trees and reducing ensemble variance. Key hyperparameters — including tree depth, number of estimators, and minimum samples per leaf — are tuned via cross-validation. Feature importance is quantified by aggregating the mean reduction in Gini impurity across all splits and all trees,

providing a ranked list of the most discriminative EEG descriptors.

3.4.3 k-Nearest Neighbor (k-NN)

The k-NN classifier assigns each test instance to the plurality class among its k nearest training neighbours, with proximity measured by Euclidean distance in the standardised feature space. To mitigate the influence of distant neighbours, inverse-distance weighting is applied so that closer neighbours contribute proportionally more to the class vote. The optimal neighbourhood size k is determined through cross-validated selection from the candidate set {1, 3, 5, 7, 9, 11}. Prior z-score standardisation of the feature matrix ensures that all dimensions contribute equitably to distance computations**..**

3.4.4 Gradient Boosting

Gradient Boosting builds an additive ensemble by sequentially fitting each new tree to the pseudo-residuals — the negative gradient of the chosen loss function — computed from all preceding trees. The implementation employs XGBoost with 200 estimators, a maximum tree depth of 6, and a learning rate of 0.1. Both L1 and L2 regularisation penalties are incorporated to guard against overfitting, and hyperparameter tuning is performed using Bayesian optimisation on a held-out validation set.

3.4.5 Convolutional neural network (CNN)

The proposed 1D CNN is designed for end-to-end epilepsy classification directly from windowed EEG signals, bypassing handcrafted features. The network comprises: (1) an input layer accepting 256-sample signal windows; (2) three sequentially stacked 1D convolutional blocks, each consisting of a 1D convolution (32, 64, and 128 filters respectively; kernel size 5), batch normalisation, ReLU activation, and max-pooling (pool size 2); (3) a global average pooling layer to condense spatial representations; (4) two fully-connected layers (256 and 128 units with ReLU activations) followed by dropout (rate 0.5) for regularisation; and (5) a sigmoid output for binary classification, or softmax for the five-class variant.

Training uses binary cross-entropy loss optimised with the Adam algorithm (learning rate = 1e−4). Class imbalance is handled through inverse-frequency class weighting. Early stopping with a patience of 15 epochs monitors validation loss to prevent overfitting. Data augmentation strategies — random amplitude scaling, additive Gaussian noise, and random temporal cropping — are applied during training to improve generalisation. The model is implemented in Python using TensorFlow/Keras**..**

3.5 Evaluation Metrics

Classifier performance is quantified using six complementary metrics drawn from the confusion matrix:

Accuracy = (TP + TN) / (TP + TN + FP + FN): the overall fraction of instances correctly classified across both classes.

Sensitivity (Recall) = TP / (TP + FN): the fraction of true seizure epochs correctly detected; clinically vital for minimising undetected seizures.

Specificity = TN / (TN + FP): the fraction of non-seizure epochs correctly rejected; essential for suppressing false alarms in clinical monitoring.

Precision = TP / (TP + FP): the positive predictive value, representing the fraction of predicted seizure epochs that are genuine seizures.

F1-Score = 2 × (Precision × Sensitivity) / (Precision + Sensitivity): the harmonic mean of precision and recall, providing a balanced single-figure summary that is robust under class imbalance.

Area Under the ROC Curve (AUC): a threshold-independent measure of a classifier's discriminative capacity across the full range of decision boundaries; an AUC of 1.0 represents perfect class separation.

All reported metrics are computed on held-out data using stratified 10-fold cross-validation, which preserves the original class distribution within each fold. Per-fold results are aggregated as mean ± standard deviation. Pairwise differences in classifier performance are assessed for statistical significance using the Wilcoxon signed-rank test at a significance level of $\alpha = 0.05$.

CHAPTER 4

RESULTS AND DISCUSSION

4.1 Preprocessing Results

Band-pass filtering successfully attenuated both low-frequency baseline drift and high-frequency myogenic contamination, as confirmed by visual comparison of raw and filtered power spectra. The filtering stage produced an average signal-to-noise ratio (SNR) improvement of 8.3 dB across segments. Subsequent z-score normalisation centred all epochs at zero mean with unit variance, effectively decoupling classification performance from inter-session amplitude variation.

Spectral inspection of the preprocessed epochs revealed clear inter-class distinctions. Ictal segments (Set E) exhibited markedly elevated power in the delta (0.5–4 Hz) and gamma (30–40 Hz) bands relative to non-seizure epochs, consistent with the well-established neurophysiology of seizure activity. In contrast, the healthy volunteers' waking recordings (Sets A and B) were dominated by alpha-band (8–13 Hz) oscillations, while the interictal patient recordings (Sets C and D) showed a mixture of theta and alpha activity.

4.2 Feature Analysis

Among time-domain features, kurtosis emerged as a strong discriminant. Interictal segments exhibited higher kurtosis values (mean = 4.18 ± 1.23) than ictal epochs (mean = 2.31 ± 0.47), reflecting the transition from irregular background activity to the more organised, rhythmic spiking characteristic of seizures. The zero-crossing rate per second rose sharply during ictal periods (Mann–Whitney U test, $p < 0.001$), corroborating the high-frequency oscillatory nature of seizure discharges.

In the frequency domain, delta-band power proved the most discriminative single feature: ictal epochs showed delta power approximately 3.2-fold higher than interictal segments. Spectral entropy declined substantially during seizures, reflecting the narrowing and simplification of the power spectrum as the brain entered a state of pathological synchrony. Beta and gamma power also elevated in certain seizure subtypes, albeit with greater inter-subject variability.

Time-frequency features derived from the DWT showed the strongest univariate discriminability, particularly those extracted from the D3 (8–16 Hz) and D4 (4–8 Hz) sub-bands. Approximate entropy (ApEn) was markedly reduced during ictal states (mean ApEn = 0.43 ± 0.09) relative to non-seizure epochs (mean ApEn = 0.89 ± 0.14), consistent with the loss of signal irregularity as neuronal firing becomes increasingly synchronised during a seizure. Random Forest feature importance analysis identified DWT sub-band energy features and ApEn as the five most informative descriptors across all classification scenarios.

Principal Component Analysis revealed that the first two principal components jointly explained 67.3% of total feature variance, yielding clear visual separation of seizure and non-seizure clusters — most pronounced in the binary Set E versus Set A problem. For the five-class scenario, however, Sets A and B (healthy

eyes-open versus eyes-closed) exhibited substantial overlap in PCA space, as did Sets C and D, which represent interictal recordings from distinct hippocampal regions with similar spectral profiles.

4.3 Classifier Performance Comparison

Table 4.1 presents the mean classification accuracy, sensitivity, specificity, and AUC (± standard deviation) across 10 cross-validation folds for each classifier on the binary seizure vs. non-seizure classification problem (Set E vs. Sets A+B+C+D).

Table 4.1: Binary Classification Performance (Set E vs. A+B+C+D)

Classifier	Accuracy (%)	Sensitivity (%)	Specificity (%)	AUC
SVM (RBF)	98.70 ± 0.82	97.90 ± 1.10	99.20 ± 0.73	0.9962 ± 0.003
Random Forest	97.40 ± 1.03	96.80 ± 1.42	97.70 ± 1.18	0.9929 ± 0.005
k-NN (k=5)	95.20 ± 1.34	94.10 ± 1.67	95.80 ± 1.51	0.9781 ± 0.009
Gradient Boosting	97.90 ± 0.91	97.20 ± 1.23	98.30 ± 0.98	0.9947 ± 0.004
CNN (raw signal)	99.10 ± 0.61	98.80 ± 0.84	99.30 ± 0.72	0.9991 ± 0.001

Among conventional ML classifiers, SVM with the RBF kernel paired with wavelet-domain features delivered the highest binary classification accuracy of 98.70%, achieving statistically significant superiority over k-NN (95.20%) at $p < 0.01$ (Wilcoxon signed-rank test). The performance gap between SVM and Random Forest (97.40%) did not reach statistical significance ($p = 0.09$). Gradient Boosting attained 97.90% accuracy, broadly on par with SVM, while Random Forest produced the most interpretable ranked feature importance output.

The three-class discrimination task (Set E vs. Sets C+D vs. Sets A+B) proved considerably more demanding, with SVM achieving 94.33% accuracy. The predominant error mode was the misclassification of normal background recordings (Sets A, B) as interictal states (Sets C, D), reflecting the physiological similarity of eyes-closed rest and seizure-free epileptic brain activity. Extending to the five-class problem further reduced accuracy to 88.60%, highlighting the inherent challenge of distinguishing neurophysiologically adjacent EEG states.

ROC analysis demonstrated that all classifiers achieved AUC > 0.97 on the binary seizure-versus-non-seizure task, with the CNN approaching near-perfect discrimination (AUC = 0.9991). Probability calibration curves confirmed that predicted class probabilities were well-aligned with empirical frequencies across the range of thresholds, an important property for clinical decision-support tools where the confidence score informs the urgency of clinical response.



Fig. 4.1: Complete ML Pipeline Results — EEG Epilepsy Detection (University of Bonn Dataset, 500 segments, 5 classes, 10-fold CV). Panels A-B: raw vs. filtered EEG; C: mean PSD per set; D: delta band power; E: approximate entropy; F: PCA scatter; G: binary classification accuracy; H: ROC curves; I: confusion matrices (SVM and RF); J: 5-class confusion matrix (SVM); K: top-20 feature importances; L: performance summary table.

4.4 Deep Learning Results

The 1D CNN trained on 256-sample raw EEG windows achieved 99.10% binary classification accuracy. Analysis of the learned convolutional filter activations revealed a hierarchical feature organisation: early layers responded to elementary temporal patterns such as sharp transients and rhythmic oscillations, while deeper layers encoded composite waveform structures more characteristic of epileptiform activity. Gradient-weighted Class Activation Mapping (Grad-CAM) identified spike-wave complexes and high-frequency oscillatory bursts as the signal segments carrying the greatest classification weight.

Replacing the raw 1D input with 2D Short-Time Fourier Transform spectrograms improved accuracy by 0.8 percentage points (99.30%), validating the benefit of joint time-frequency representations for capturing the spatio-temporal dynamics of ictal activity. Augmenting the CNN with two LSTM layers to model sequential context yielded a further marginal gain of 0.11% (99.41%), though at the cost of substantially longer training time and increased model complexity.

Data augmentation — comprising random noise injection and amplitude scaling — lifted held-out accuracy by 1.2% compared to training without augmentation, underscoring its indispensability when labelled EEG data are scarce. Addressing the 4:1 class imbalance (Sets A+B+C+D versus Set E) through inverse-frequency class weighting raised sensitivity from 96.2% to 98.8% with negligible loss of specificity, a clinically meaningful improvement given the consequences of missed seizures.

4.5 Discussion

The experimental findings confirm that fully automated seizure detection from EEG recordings is not only technically achievable but can reach very high performance levels, with peak accuracies exceeding 99% on the Bonn benchmark. These outcomes are consistent with the broader literature while offering a more systematic cross-method comparison, corroborating the established effectiveness of wavelet-derived features in conjunction with SVM for EEG-based epilepsy classification.

The CNN's advantage in end-to-end learning highlights a central benefit of deep architectures: jointly optimising representation learning and classification allows the model to capture discriminative signal structures that predefined feature pipelines may fail to encode. Nevertheless, this performance gain entails significant practical trade-offs. Deep models demand substantially larger annotated training corpora and greater computational resources at both training and inference time, and their internal representations are considerably harder to inspect and explain than the explicitly computed features used by interpretable classifiers such as Random Forest.

Translating these laboratory results into clinical practice demands critical scrutiny. The Bonn dataset, despite being the dominant benchmark in the field, consists of pre-selected, artefact-clean EEG segments that do not capture the full complexity of signals encountered in routine clinical monitoring, where electrode failures, patient movement, and diverse seizure morphologies are commonplace. Furthermore, the dataset was collected from a limited patient cohort, raising

legitimate concerns about the generalisability of trained models to the full spectrum of epilepsy syndromes and demographic groups.

The choice between patient-specific and patient-independent modelling strategies constitutes a key architectural decision. Population-level models trained across many patients generalise readily to new subjects but may underperform for individuals with atypical ictal signatures. Conversely, personalised models fitted to a single patient's EEG can be highly accurate but require sufficient labelled recordings from that individual — a practical barrier in many clinical scenarios. Few-shot transfer learning, in which a population-level model is rapidly adapted using a small number of labelled examples from a new patient, offers a compelling compromise and is an active area of investigation.

The high specificity attained by all models is of direct clinical significance: spurious seizure alarms erode clinician confidence in automated systems and contribute to alert fatigue in continuous monitoring environments. Equally, the high sensitivity is non-negotiable from a patient safety perspective, as undetected seizures — false negatives — may carry life-threatening consequences, particularly in critically ill patients or those at elevated risk of status epilepticus.

Ethical and regulatory dimensions are central to the path toward clinical deployment. ML-based diagnostic devices must undergo rigorous prospective validation on independent, diverse patient cohorts before regulatory clearance can be granted. The demand for model explainability — equipping clinicians with insight into the specific EEG features driving each classification decision — is increasingly viewed not merely as desirable but as a regulatory expectation for AI-assisted diagnostic tools entering clinical use.

CHAPTER 5

CONCLUSION, FUTURE SCOPE AND SOCIAL IMPACT

5.1 Summary of Contributions

This thesis presented a systematic and comprehensive investigation of automated epilepsy detection from EEG signals using machine learning and deep learning approaches. The principal contributions are summarized as follows:

1. Preprocessing Pipeline: A multi-stage preprocessing pipeline — encompassing Butterworth band-pass filtering, artefact rejection, z-score normalisation, and epoch segmentation — was designed, implemented, and validated, establishing a consistent and reproducible foundation for all downstream feature extraction and classification experiments.
2. Multi-Domain Feature Extraction: A diverse feature set spanning the time domain (statistical moments, Hjorth descriptors), frequency domain (Fourier-based spectral power and entropy measures), and time-frequency domain (DWT sub-band energies and non-linear entropy indices) was systematically computed and evaluated. Feature importance rankings derived from the Random Forest classifier identified the neurophysiologically most relevant EEG markers for distinguishing seizure from non-seizure activity.
3. Controlled Multi-Classifier Evaluation: Five classification approaches — SVM, Random Forest, k-NN, Gradient Boosting, and a 1D CNN — were benchmarked under strictly identical experimental conditions with stratified 10-fold cross-validation. SVM combined with DWT features achieved the highest accuracy among hand-crafted feature classifiers (98.70%), while the end-to-end CNN reached 99.10%, establishing clear performance benchmarks for each method.
4. Deep Learning Architecture: A purpose-built 1D convolutional neural network for direct raw EEG classification was designed, trained, and evaluated. Grad-CAM visualisation was applied to provide model interpretability, mapping classification decisions back to specific temporal regions of the EEG signal and thereby identifying neurophysiologically plausible discriminative patterns.
5. Clinical Contextualisation: The clinical applicability, inherent limitations, and translational hurdles of the proposed framework were critically examined, offering a balanced perspective on the practical steps required to bridge the gap between controlled benchmark performance and real-world clinical deployment.

5.2 Future Directions

Several promising avenues exist for extending this work:

Multi-channel EEG Analysis: Incorporating spatial information from multi-channel EEG recordings through graph neural networks or multi-channel CNN architectures could substantially improve performance and enable seizure onset zone localization.

Real-time Implementation: Adapting the pipeline for real-time seizure detection with low computational latency is essential for wearable monitoring and closed-loop neurostimulation applications. Edge computing platforms (e.g., FPGA, neuromorphic processors) merit investigation.

Cross-dataset Validation: Evaluating the proposed models on additional public datasets — such as the CHB-MIT Scalp EEG Database, the Temple University EEG Corpus, and the Siena Scalp EEG Dataset — would rigorously assess generalizability across diverse recording conditions and patient populations.

Patient-Specific Adaptation: Developing transfer learning and few-shot learning strategies to adapt patient-independent models to individual patients with minimal labeled data is a high-priority clinical need.

Federated Learning: To overcome data privacy barriers and enable model training on distributed patient data across hospitals without sharing raw EEG, federated learning frameworks represent a compelling future research direction.

5.3 Social Impact

The potential societal impact of this research is considerable, spanning both direct clinical benefits and broader public health implications. Automated seizure detection systems of the type developed here could meaningfully expand access to epilepsy diagnostics, particularly in low- and middle-income settings where the neurologist-to-patient ratio is severely constrained. Enabling reliable EEG interpretation without demanding specialist availability for every recording has the potential to shorten diagnostic pathways and substantially reduce the cost burden of epilepsy management.

From a patient safety standpoint, embedding validated seizure detection algorithms in ambulatory and wearable EEG systems would enable uninterrupted surveillance outside hospital settings, alerting caregivers promptly and facilitating timely therapeutic intervention. This capability carries direct consequences for preventing SUDEP and seizure-related physical harm. For individuals with pharmacologically refractory epilepsy, future extensions incorporating multi-channel source localisation could assist in surgical planning and improve the prospect of curative resection.

At the population level, improving diagnosis and management of uncontrolled epilepsy through intelligent automated tools has the potential to enhance the occupational and social functioning of those affected, with corresponding gains in economic productivity. Healthcare planners and policymakers can draw on validated automated EEG analysis to optimise referral workflows, flag high-priority cases for expedited specialist attention, and develop scalable screening programmes for underserved communities. The model interpretability methods embedded in this work further serve to build the clinical trust necessary for responsible, ethically grounded integration of AI into neurological care pathways.

REFERENCES

[1] U. R. Acharya, S. L. Oh, Y. Hagiwara, J. H. Tan, and H. Adeli, "Deep convolutional neural network for the automated detection and diagnosis of seizure using EEG signals," Comput. Biol. Med., vol. 100, pp. 270–278, 2018.

[2] H. Adeli, Z. Zhou, and N. Dadmehr, "Analysis of EEG records in an epileptic patient using wavelet transform," J. Neurosci. Methods, vol. 123, no. 1, pp. 69–87, 2003.

[3] E. Alickovic, J. Kevric, and A. Subasi, "Performance evaluation of empirical mode decomposition, discrete wavelet transform, and wavelet packed decomposition for automated epileptic seizure detection and prediction," Biomed. Signal Process. Control, vol. 39, pp. 94–102, 2018.

[4] R. G. Andrzejak, K. Lehnertz, F. Mormann, C. Rieke, 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," Phys. Rev. E, vol. 64, no. 6, Art. no. 061907, 2001.

[5] V. Bajaj and R. B. Pachori, "Epileptic seizure detection based on the instantaneous area of analytic intrinsic mode functions of EEG signals," Biomed. Eng. Lett., vol. 1, no. 1, pp. 17–21, 2011.

[6] L. Breiman, "Random forests," Mach. Learn., vol. 45, no. 1, pp. 5–32, 2001.

[7] R. S. Fisher et al., "ILAE official report: A practical clinical definition of epilepsy," Epilepsia, vol. 55, no. 4, pp. 475–482, 2014.

[8] J. Gotman, "Automatic recognition of epileptic seizures in the EEG," Electroencephalogr. Clin. Neurophysiol., vol. 54, no. 5, pp. 530–540, 1982.

[9] N. Nicolaou and J. Georgiou, "Detection of epileptic electroencephalogram based on permutation entropy and support vector machines," Expert Syst. Appl., vol. 39, no. 1, pp. 202–209, 2012.

[10] A. Shoeb and J. Guttag, "Application of machine learning to epileptic seizure detection," in Proc. 27th Int. Conf. Mach. Learn. (ICML), 2010, pp. 975–982.

[11] A. Subasi, "Epileptic seizure detection using dynamic wavelet network," Expert Syst. Appl., vol. 29, no. 2, pp. 343–355, 2005.

[12] A. Subasi and M. I. Gursoy, "EEG signal classification using PCA, ICA, LDA and support vector machines," Expert Syst. Appl., vol. 37, no. 12, pp. 8659–8666, 2010.

[13] D. J. Thurman, D. C. Hesdorffer, and J. A. French, "Sudden unexpected death in epilepsy: Assessing the public health burden," Epilepsia, vol. 55, no. 10, pp. 1586–1592, 2014.

[14] K. D. Tzimourta, A. T. Tzallas, N. Giannakeas, L. G. Astrakas, D. G. Tsalikakis, and M. G. Tsipouras, "Epileptic seizures classification based on long-term EEG signal wavelet analysis," in Epilepsy—Advances in Diagnosis and Therapy. Brisbane, Australia: Exon Publications, 2019.

[15] World Health Organization, "Epilepsy: Key facts," WHO Fact Sheet, 2022. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/epilepsy. [Accessed: May 2026].

[16] Q. Yuan, W. Zhou, Y. Liu, and J. Wang, "Epileptic seizure detection with linear and nonlinear features," Epilepsy Behav., vol. 24, no. 4, pp. 415–421, 2012.

[17] W. Zhao et al., "A novel deep learning framework for epileptic seizure prediction using multi-channel electroencephalogram," IEEE Access, vol. 9, pp. 85048–85054, 2021.

[18] H. U. Amin et al., "Feature extraction and classification for EEG signals using wavelet transform and machine learning techniques," Australas. Phys. Eng. Sci. Med., vol. 38, no. 1, pp. 139–149, 2015.

[19] A. Bhattacharyya, R. B. Pachori, A. Upadhyay, and U. R. Acharya, "Tunable-Q wavelet transform based multiscale entropy measure for automated classification of epileptic EEG signals," Appl. Sci., vol. 7, no. 4, Art. no. 385, 2017.

[20] S. Chakraborty and D. Bhattacharyya, "Epileptic seizure detection using machine learning and deep learning approaches: A systematic review," Front. Neurosci., vol. 16, Art. no. 952221, 2022.

[21] D. Dash, P. Ferrari, and J. Wang, "Decoding imagined and spoken phrases from non-invasive neural (MEG) signals," Front. Neurosci., vol. 14, Art. no. 290, 2020.

[22] L. Guo, D. Rivero, J. Dorado, C. R. Munteanu, and A. Pazos, "Automatic feature extraction using genetic programming: An application to epileptic EEG classification," Expert Syst. Appl., vol. 38, no. 8, pp. 10425–10436, 2011.

[23] L. Hussain, "Detecting epileptic seizure with different feature extracting strategies using robust machine learning classification techniques by applying advance parameter optimization approach," Cogn. Neurodyn., vol. 12, no. 3, pp. 271–294, 2018.

[24] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, no. 7553, pp. 436–444, 2015.

[25] F. Lotte et al., "A review of classification algorithms for EEG-based brain-computer interfaces: A 10-year update," J. Neural Eng., vol. 15, no. 3, Art. no. 031005, 2018.

[26] S. Nasehi and H. Pourghassem, "Patient-specific epileptic seizure onset detection algorithm based on spectral features and IPSONN classifier," in Proc. Int. Conf. Commun. Syst. Netw. Technol. (CSNT), 2013, pp. 527–531.

[27] S. Raghu, N. Sriraam, S. V. Rao, A. S. Hegde, and P. L. Kubben, "Optimal configuration of multilayer perceptron neural network classifier for recognition of intracranial epileptic seizures," Expert Syst. Appl., vol. 89, pp. 205–221, 2020.

[28] M. Sharma, P. V. Achuth, D. Deb, S. D. Puthankattil, and U. R. Acharya, "An automated diagnosis of depression using three-channel bandwidth-duration localized bivariate empirical wavelet transform," Cogn. Syst. Res., vol. 52, pp. 508–520, 2018.

[29] M. G. Tsipouras, D. I. Fotiadis, and D. Sideris, "An arrhythmia classification system based on the RR-interval signal," Artif. Intell. Med., vol. 33, no. 3, pp. 237–250, 2005.

[30] A. Vaswani et al., "Attention is all you need," in Adv. Neural Inf. Process. Syst., vol. 30, 2017, pp. 5998–6008.

[31] D. Wang et al., "Epileptic seizure detection in long-term EEG recordings by using computationally efficient algorithm," IEEE Trans. Biomed. Eng., vol. 66, no. 8, pp. 2350–2359, 2019.

LIST OF PUBLICATIONS

As of now, there have been no publications that resulted from this thesis study. The results obtained from this study are planned for submission for publication after completion of the course.