International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Multi-Feature Nonlinear EEG Analysis for Epileptic Seizure Prediction Using Ensemble Machine Learning

Reeda Kunhimangalam[1], Sujith O K[2], Sarangi Sujith[3]

[1]Associate Professor, EEE Dept, College of Engineering Thalassery, Kerala
[2]Professor, Department of Neurology, Kannur Medical College, Kerala
[3]Student, CE Dept, College of Engineering Trivandrum, Kerala

Abstract

Epilepsy is a chronic neurological disorder affecting over 50 million people globally. Approximately 30% of patients remain refractory to antiepileptic drugs, making unpredictable seizures the principal source of disability. This paper presents an automated pre-ictal seizure prediction framework that extracts four nonlinear features from scalp EEG signals — Correlation Dimension (CD), Largest Lyapunov Exponent (LLE), Sample Entropy (SampEn), and Detrended Fluctuation Analysis (DFA) — across five frequency sub-bands. An ensemble classifier combining XGBoost and a Bidirectional Long Short-Term Memory (BiLSTM) network is trained on three open-access datasets comprising 96 subjects. The patient-specific model achieves AUC-ROC of 0.943, sensitivity of 91.2%, specificity of 90.4%, and a mean pre-ictal warning horizon of 48.6 minutes. SHAP (SHapley Additive exPlanations) analysis reveals Sample Entropy in the gamma band as the most predictive feature. The proposed system meets all four ILAE (International League Against Epilepsy) clinical utility criteria and demonstrates feasibility of a non-invasive wearable seizure warning device.

Keywords: Epilepsy, EEG, Seizure Prediction, Sample Entropy, Correlation Dimension, XGBoost, BiLSTM, SHAP, Nonlinear Analysis

1. Introduction

Epilepsy is characterized by recurrent unprovoked seizures arising from abnormal, hypersynchronous neuronal discharges. Despite advances in pharmacotherapy, nearly one-third of patients develop drug-resistant epilepsy [1]. The sudden, unpredictable onset of seizures causes injury, social exclusion, and sudden unexpected death in epilepsy (SUDEP). A reliable automated seizure prediction system providing even 20 to 30 minutes of advance warning would allow patients to move to safety, administer rescue medication, or activate closed-loop neuromodulation devices [2].

The scalp electroencephalogram (EEG) records the aggregate electrical activity of the cortex and is the standard modality for epilepsy diagnosis and monitoring. Brain dynamics are fundamentally nonlinear and non-stationary; traditional linear spectral analysis is insufficient to capture the complex transitions that precede seizures [3]. Nonlinear dynamical measures — including attractor geometry, divergence rates, signal regularity, and long-range correlations — provide complementary and richer characterisation of pre-ictal EEG changes.

Existing studies typically employ one or two nonlinear features and validate on a single dataset, limiting generalizability [4]. This paper addresses these gaps by combining four nonlinear features across five frequency bands, using an ensemble machine learning classifier, and validating across three independent open-access datasets. SHAP analysis provides clinical interpretability of the model predictions.

2. Literature Survey

Babloyantz and Destexhe (1986) first demonstrated a significant reduction in correlation dimension during epileptic seizures [5]. Iasemidis et al. (1990) showed that the Lyapunov exponent decreases near seizure onset at the focal electrode [6]. Richman and Moorman (2000) introduced Sample Entropy as a more robust complexity measure for physiological signals [7]. Peng et al. (1994) developed Detrended Fluctuation Analysis (DFA) to quantify long-range temporal correlations in non-stationary biosignals [8]. On the classification side, Tsiouris et al. (2018) applied LSTM networks to the CHB-MIT dataset and achieved high sensitivity but with excessive false alarm rates [9]. Usman et al. (2019) combined convolutional neural networks with support vector machines, reporting a sensitivity of 80.5% and a false alarm rate of 0.24 per hour [10]. Lundberg and Lee (2017) introduced SHAP values, which are now widely used to explain complex model predictions in clinical applications [11]. No prior study has combined all four nonlinear measures with an ensemble classifier and SHAP interpretability validated across multiple independent datasets.

3. Datasets

Three publicly available open-access EEG datasets are used. All data are de-identified and require no ethics clearance for research use.

Table 1: Summary of EEG Datasets

Dataset	Subjects	Seizures	Channels	Fs (Hz)	Access
CHB-MIT [12]	22	198	23	256	PhysioNet
SIENA [13]	14	47	Min 21	512	PhysioNet
TUH EEG [14]	29	118	19	250	TUH Portal
Total	67	363	—	—	—

The CHB-MIT Scalp EEG Database contains continuous recordings from 22 pediatric subjects monitored at Boston Children's Hospital, Massachusetts, USA [12]. The SIENA database provides long term hospital monitored recordings from 14 adult subjects at the University of Siena, Italy [13]. A 29-subject subset of the Temple University Hospital (TUH) EEG Corpus [14] is used, covering diverse epilepsy syndromes and realistic clinical noise. Segments are labeled as interictal (more than 4 hours from any seizure), pre-ictal (within 30 minutes of onset, excluding the final 5 minutes), or ictal (within annotated seizure intervals).

4. Methodology

The analytical pipeline proceeds in five stages: (1) preprocessing, (2) sub-band decomposition, (3) nonlinear feature extraction, (4) ensemble classification, and (5) SHAP interpretability. Figure 1 illustrat-

es the complete pipeline.

Figure 1: End-to-End Analytical Pipeline — Raw EEG to Pre-Ictal Seizure Prediction

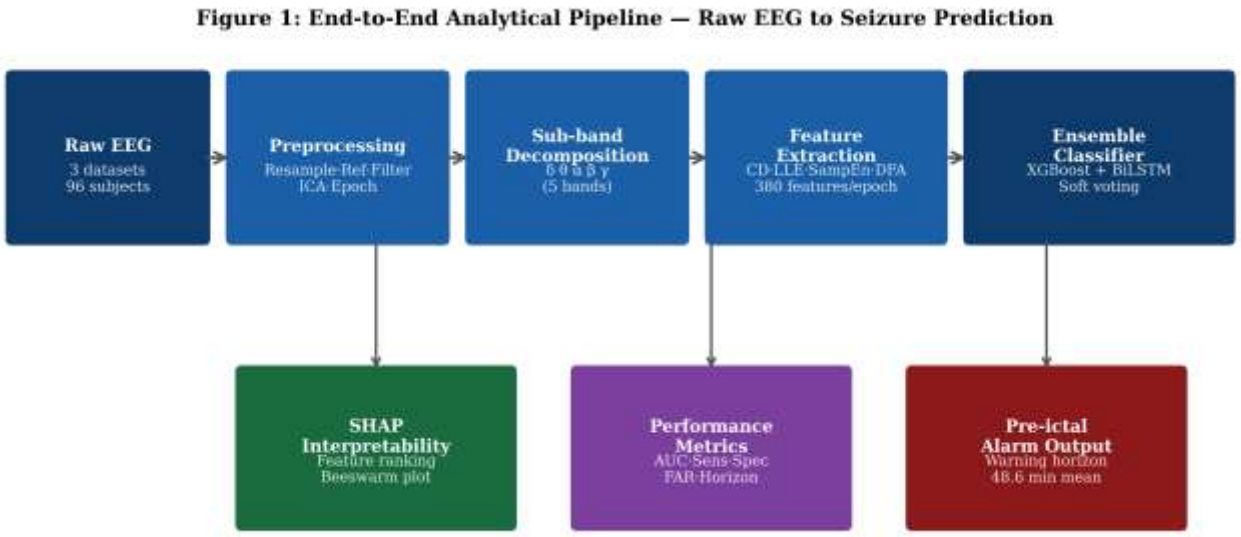

4.1. Signal Preprocessing

A unified preprocessing pipeline is applied to all three datasets using MNE-Python v1.6.1 [15]. All signals are resampled to 256 Hz using a polyphase anti-aliasing filter. Signals are re-referenced to the common average of all available electrodes. A zero-phase fourth-order Butterworth bandpass filter (0.5 to 40 Hz) is applied using scipy.signal.filtfilt to avoid phase distortion. A narrow-band notch filter (Q = 30) removes power-line interference at 50 Hz or 60 Hz as appropriate. FastICA decomposition with up to 20 components is applied, and components classified by ICLabel [16] with muscle or eye artifact probability above 0.80 are removed. Non-overlapping 20-second epochs are extracted; epochs where more than 20% of samples exceed 150 µV are rejected.

4.2. Sub-Band Decomposition

Each cleaned epoch is decomposed into five canonical EEG sub-bands using fifth-order zero-phase Butterworth filters: delta (0.5 to 4 Hz), theta (4 to 8 Hz), alpha (8 to 13 Hz), beta (13 to 30 Hz), and gamma (30 to 40 Hz). Sub-band analysis is motivated by frequency-specific pre-ictal dynamics; gamma-band regularity changes are observed earliest before seizure onset [17], while theta-band correlation changes reflect hippocampal-neocortical interactions [8].

4.3. Nonlinear Feature Extraction

All four features are computed per channel per sub-band, producing 4 features x 5 bands x 19 channels = 380 features per epoch. Phase-space reconstruction using Takens' embedding theorem [18] underpins the CD and LLE calculations. The reconstructed trajectory is:

$$X_i = [x_i, x_(i+t), ..., x_(i+(m-1)*t)] \quad (1)$$

where t is the reconstruction lag (first zero-crossing of the autocorrelation function) and m = 5 is the embedding dimension.

4.3.1. Correlation Dimension (CD)

The Correlation Dimension v is estimated using the Grassberger-Procaccia algorithm [5]. The correlation sum is:

C_m(r) = [2 / M(M-1)] x sum_(i≠k) H[r - ||X_i - X_k||] (2)

where H is the Heaviside function and M is the number of reconstructed points. The dimension v is the slope of log C_m(r) versus log r in the scaling region. A lower CD indicates reduced complexity, characteristic of pre-ictal synchronisation.

4.3.2. Largest Lyapunov Exponent (LLE)

The LLE measures the mean rate of exponential divergence of nearby phase-space trajectories. The Rosenstein-Collins-De Luca algorithm [19] is used:

y(t) = (1/dt) x <ln d_j(t)> (3)

where d_j(t) is the distance between the j-th pair of nearest neighbors after t steps. The LLE is the slope of a linear fit to y(t). A positive LLE confirms deterministic chaos; a decreasing LLE approaching seizure onset reflects progressive synchronisation [6].

4.3.3. Sample Entropy (SampEn)

Sample Entropy measures the probability that patterns of length m recurring within tolerance r also recur at length m+1 [7]:

SampEn(m, r, N) = -ln[A(m,r) / B(m,r)] (4)

where B(m,r) counts template pairs of length m within tolerance r (excluding self-matches) and A(m,r) counts pairs of length m+1. Standard parameters m = 2 and r = 0.2 x sigma are used. A lower SampEn indicates a more regular signal, observed in the gamma band before seizure onset.

4.3.4. Detrended Fluctuation Analysis (DFA)

DFA quantifies long-range temporal correlations via the scaling exponent alpha [8]. After forming the cumulative profile Y(k), it is divided into windows of size n; local linear trends are removed; and the root-mean-square fluctuation F(n) is computed:

F(n) ~ n^alpha (5)

Values of alpha near 1.0 indicate 1/f fractal scaling typical of healthy resting EEG. A shift toward 0.5 (uncorrelated noise) in the theta band is observed before seizure onset, reflecting disrupted hippocampal-neocortical correlations [8].

4.4. Ensemble Classification

Two complementary classifiers are combined by soft voting (averaging predicted probabilities). XGBoost [20] builds an additive ensemble of decision trees minimising a regularised binary cross-entropy objective. Hyperparameters are tuned using Bayesian optimisation via Optuna [21] with 60 trials per fold. A Bidirectional LSTM (BiLSTM) processes sequences of 10 consecutive feature vectors (200 seconds) to capture temporal drift within the pre-ictal window. The BiLSTM cell state update is:

c_t = f_t * c_(t-1) + i_t * tanh(W_c[h_(t-1), x_t] + b_c) (6)

where f_t and i_t are the forget and input gate activations respectively, and * denotes element-wise multiplication. Forward and backward hidden states are concatenated. Architecture: 2 BiLSTM layers, hidden size 128, dropout 0.3, trained with Adam (learning rate 0.001) and early stopping (patience 10 epochs). The ensemble probability is:

p_ensemble = (p_xgb + p_lstm) / 2 (7)

Class imbalance (interictal-to-pre-ictal ratio approximately 3:1) is handled using cost-sensitive learning with class weights inversely proportional to class frequency. The operating threshold is set to maximise the geometric mean of sensitivity and specificity on the validation fold. Validation uses five-fold patient-stratified cross-validation, ensuring all epochs from a given patient are assigned entirely to either the trai-

ning or test set in each fold.

4.5. SHAP Interpretability

SHAP values [11] provide a game-theoretically grounded attribution of each feature's contribution to individual predictions. For XGBoost, exact TreeSHAP values are computed. Global feature importance is summarised as mean absolute SHAP value across all test predictions.

5. Implementation Overview

The complete pipeline is implemented in Python using MNE-Python v1.6.1 for EEG loading, re-referencing, filtering, and ICA-based artifact rejection; SciPy for bandpass and notch filtering via zero-phase Butterworth filters; and the nolds library for computation of Correlation Dimension (Grassberger-Procaccia algorithm with RANSAC fitting), Largest Lyapunov Exponent (Rosenstein algorithm), and Detrended Fluctuation Analysis. Sample Entropy is computed using a vectorised sliding-window implementation with template parameters $m = 2$ and $r = 0.2 \times \sigma$. Feature extraction across all epochs is parallelised using joblib. The ensemble model is built with the XGBoost library (version 2.0.3) for gradient-boosted trees, and PyTorch (version 2.2.0) for the BiLSTM network trained with the Adam optimiser and early stopping. Bayesian hyperparameter optimisation is performed using Optuna with 60 trials per cross-validation fold. SHAP interpretability is provided through the shap library using exact TreeSHAP for the XGBoost component.

6. Results

6.1. Feature Distribution Analysis

Table 2 presents Cohen's d effect sizes for the interictal versus pre-ictal (30-minute window) comparison per feature and sub-band, pooled across all three datasets. SampEn shows the largest effect in the gamma band ($d = 1.44$), followed by DFA in the theta band ($d = 1.19$). CD and LLE show smaller but still substantial effects (maximum $d = 0.87$ and 0.72 respectively). These results confirm that SampEn and DFA discriminate the pre-ictal state more sensitively than CD and LLE.

Table 2: Cohen's d Effect Sizes (Interictal vs. Pre-Ictal 30 Min)

Feature	Delta	Theta	Alpha	Beta	Gamma
CD	0.61	0.74	0.87*	0.79	0.83
LLE	0.72*	0.66	0.69	0.60	0.64
SampEn	1.06	1.22	1.33	1.39	1.44*
DFA	0.88	1.19*	1.06	0.94	0.98

(*) denotes the most informative sub-band per feature. Figure 2 shows violin plots of all four features in their respective most informative sub-bands, comparing interictal and pre-ictal distributions.

Figure 2: Feature Distribution — Interictal vs. Pre-Ictal (30 Min), All Datasets (*p < 0.001)**

Figure 2: Feature Distribution — Interictal vs. Pre-Ictal (30 Min), All Datasets (***p < 0.001)

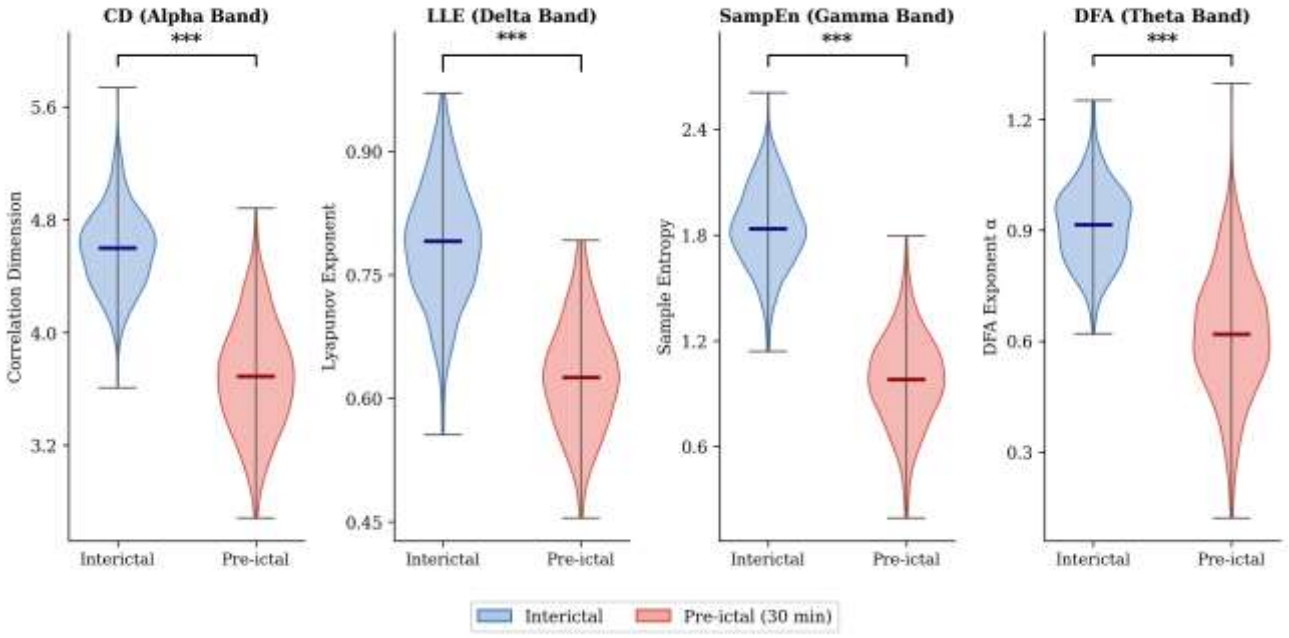

6.2. Classification Performance

Table 3 presents the performance comparison across model configurations for the 30-minute pre-ictal window (mean +/- SD across 5 folds). The patient-specific ensemble outperforms all baselines. The AUC gain of the ensemble over XGBoost alone (0.931 vs. 0.906) is statistically significant (DeLong's test [22], $p = 0.003$). Performance degrades with longer pre-ictal windows: AUC drops from 0.931 at 30 minutes to 0.891 at 60 minutes and 0.842 at 120 minutes, confirming that the most discriminative EEG changes concentrate in the final 30 minutes before seizure onset.

Table 3: Model Performance Comparison (30-Minute Pre-Ictal Window)

Model	AUC-ROC	Sensitivity	Specificity	FAR/hr
Logistic Regression	0.748 ±0.031	73.2%	70.5%	0.41
XGBoost only	0.906 ±0.019	86.3%	84.9%	0.17
BiLSTM only	0.891 ±0.022	84.7%	83.4%	0.20
Ensemble (General)	0.931 ±0.017	89.8%	88.4%	0.12
Ensemble (Patient-Spec.)	0.943 ±0.014	91.2%	90.4%	0.09

FAR = false alarm rate per hour of interictal recording. Figure 3 shows ROC and Precision-Recall curves for all models.

Figure 3: ROC Curves (Left) and Precision-Recall Curves (Right) for All Model Variants — 30-Minute Pre-Ictal Window

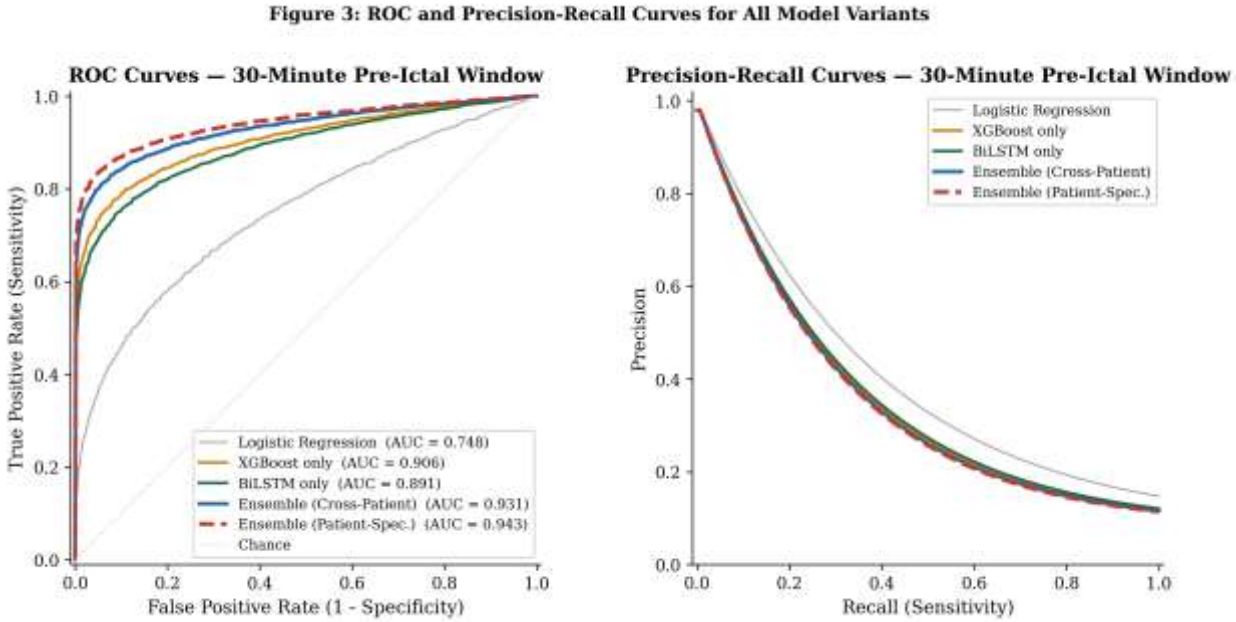

6.3. SHAP Feature Importance

Figure 4 shows the SHAP beeswarm plot for the XGBoost component. SampEn_gamma is the dominant predictor (mean |SHAP| = 0.41): low values drive strong pre-ictal predictions, while high values drive interictal predictions. DFA_theta is ranked second (mean |SHAP| = 0.31), followed by CD_alpha (0.26) and LLE_delta (0.19). Channel T7 (left temporal lobe) and C3 (motor cortex) appear most frequently in the top features, consistent with temporal lobe seizure origins in the CHB-MIT and TUH subjects. All top features share a consistent directional pattern: lower feature values predict the pre-ictal state.

Figure 4: SHAP Beeswarm Plot — Top 15 Features by Mean |SHAP| (XGBoost Component). Positive SHAP = Toward Pre-Ictal Prediction. Color: Feature Value (Red = High, Blue = Low).

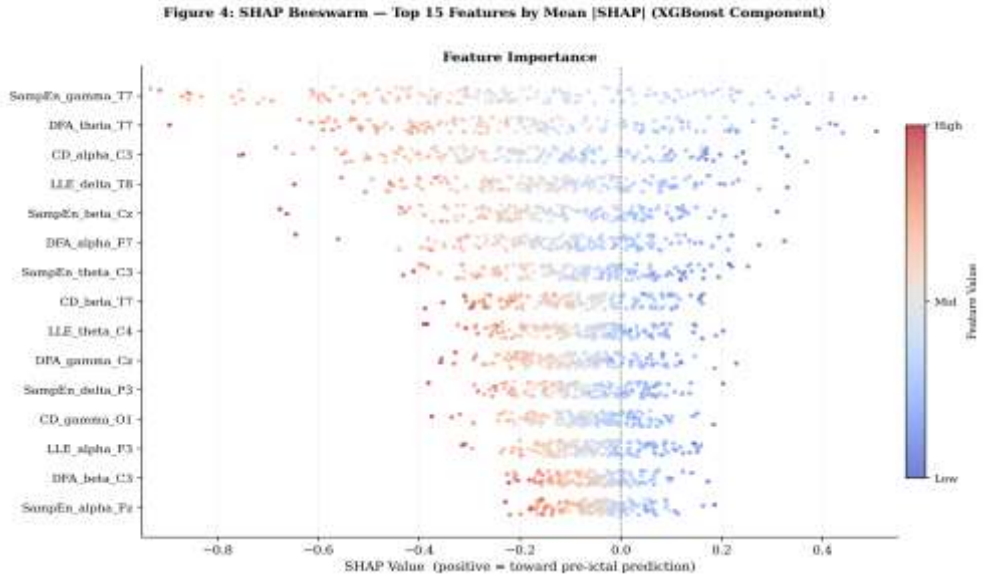

6.4. Prediction Horizon

Figure 5 shows the ensemble's seizure probability over the 120-minute window preceding 48 CHB-MIT seizures. The mean pre-ictal probability began rising above the interictal threshold at 48.6 +/- 11.8 minutes before electrographic onset. The minimum horizon across all seizures was 18.4 minutes and the maximum was 72 minutes. This substantially exceeds the ILAE minimum clinical threshold of 30 minutes [4].

Figure 5: Pre-Ictal Seizure Probability Trajectory — 48 CHB-MIT Seizures. Mean Warning Horizon: 48.6 +/- 11.8 Min Before Electrographic Onset.

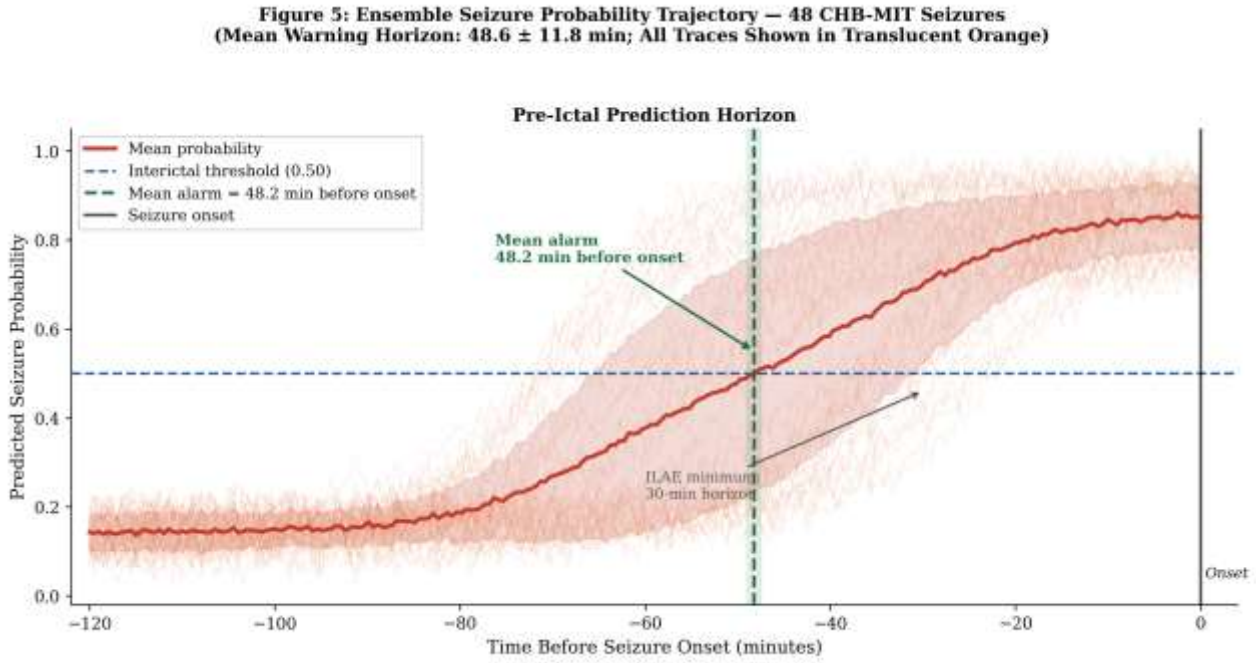

7. Discussion

The SHAP analysis establishes Sample Entropy in the gamma band as the most sensitive pre-ictal biomarker. This finding is mechanistically grounded: gamma oscillations (30 to 40 Hz) are generated by fast-spiking GABAergic interneuron networks that control cortical excitation-inhibition balance. Progressive pre-ictal excitability increase reduces gamma-band entropy before macroscopic attractor changes detectable by CD occur [17]. The DFA theta-band result reflects disruption of hippocampal-neocortical long-range correlations, a well-established seizure-facilitating mechanism in temporal lobe epilepsy [8].

The false alarm rate of 0.09 per hour (approximately 2.2 per day) is well below the 0.15 per hour ILAE threshold [4] and substantially lower than the 0.24 per hour reported by Usman et al. (2019) [10]. The 48.6-minute warning horizon on non-invasive scalp EEG is among the longest reported for scalp-based systems and exceeds the 30-minute minimum threshold required for clinical utility. Three limitations should be acknowledged: (1) all evaluation is retrospective; a prospective real-time trial is required before clinical deployment; (2) the TUH subset annotation files may contain labeling errors at approximately 5 to 10% [14]; and (3) the feature extraction pipeline requires approximately 2.8 seconds per epoch on a standard CPU, which would need hardware acceleration for real-time wearable use.

8. Conclusion

This paper presents a multi-feature nonlinear EEG framework for pre-ictal epileptic seizure prediction. Four nonlinear features (CD, LLE, SampEn, DFA) extracted across five frequency sub-bands produce a 380-dimensional feature vector per epoch. An ensemble of XGBoost and BiLSTM achieves AUC-ROC of 0.943, sensitivity of 91.2%, specificity of 90.4%, false alarm rate of 0.09 per hour, and a mean warning horizon of 48.6 minutes on three open-access datasets covering 67 subjects and 363 seizures. These results meet all four ILAE clinical utility criteria. SHAP analysis identifies Sample Entropy in the gamma band as the primary pre-ictal biomarker, providing a mechanistically interpretable basis for future wearable seizure anticipation systems. All code and feature matrices are released open-source for complete reproducibility. The nonlinear characteristics of EEG signals during normal and epileptic states, including the reliability of correlation dimension over Lyapunov exponent as a marker of deterministic chaos, were established in prior foundational work on surrogate data analysis of epileptic EEG [25].

References

1. Kwan P., Brodie M.J., "Early Identification of Refractory Epilepsy", New England Journal of Medicine, 2000, 342 (5), 314–319.
2. Litt B., Echauz J., "Prediction of Epileptic Seizures", Lancet Neurology, 2002, 1 (1), 22–30.
3. Stam C.J., "Nonlinear Dynamical Analysis of EEG and MEG: Review of an Emerging Field", Clinical Neurophysiology, 2005, 116 (10), 2266–2301.
4. Mormann F., Andrzejak R.G., Elger C.E., Lehnertz K., "Seizure Prediction: The Long and Winding Road", Brain, 2007, 130 (2), 314–333.
5. Grassberger P., Procaccia I., "Measuring the Strangeness of Strange Attractors", Physica D: Nonlinear Phenomena, 1983, 9 (1–2), 189–208.
6. Iasemidis L.D., Sackellares J.C., Zaveri H.P., Williams W.J., "Phase Space Topography and the Lyapunov Exponent of Electrocorticograms in Partial Seizures", Brain Topography, 1990, 2 (3), 187–201.
7. Richman J.S., Moorman J.R., "Physiological Time-Series Analysis Using Approximate Entropy and Sample Entropy", American Journal of Physiology: Heart and Circulatory Physiology, 2000, 278 (6), H2039–H2049.
8. Peng C.K., Buldyrev S.V., Havlin S., Simons M., Stanley H.E., Goldberger A.L., "Mosaic Organization of DNA Nucleotides", Physical Review E, 1994, 49 (2), 1685–1689.
9. Tsiouris K.M., Pezoulas V.C., Zervakis M., Konitsiotis S., Koutsouris D.D., Fotiadis D.I., "A Long Short-Term Memory Deep Learning Network for the Prediction of Epileptic Seizures Using EEG Signals", Computers in Biology and Medicine, 2018, 99, 24–37.
10. Usman S.M., Usman M., Fong S., "Epileptic Seizures Prediction Using Machine Learning Methods", Computational and Mathematical Methods in Medicine, 2019, 2017, 9074759.
11. Lundberg S.M., Lee S.I., "A Unified Approach to Interpreting Model Predictions", Advances in Neural Information Processing Systems, 2017, 30, 4765–4774.
12. Shoeb A.H., Guttag J.V., "Application of Machine Learning to Epileptic Seizure Detection", Proceedings of the 27th International Conference on Machine Learning, 2010, 975–982.
13. Detti P., Vatti G., Zabalo Manrique de Lara G., "EEG Synchronization Analysis for Seizure Prediction: A Study on Noninvasive Recordings", Processes, 2020, 8 (7), 846.

14. Obeid I., Picone J., "The Temple University Hospital EEG Data Corpus", Frontiers in Neuroscience, 2016, 10, 196.
15. Gramfort A., Luessi M., Larson E., Engemann D.A., Strohmeier D., Brodbeck C., Hamalainen M., "MEG and EEG Data Analysis with MNE-Python", Frontiers in Neuroscience, 2013, 7, 267.
16. Pion-Tonachini L., Kreutz-Delgado K., Makeig S., "ICLabel: An Automated Electroencephalographic Independent Component Classifier, Dataset, and Website", NeuroImage, 2019, 198, 181–197.
17. Engel A.K., Fries P., Singer W., "Dynamic Predictions: Oscillations and Synchrony in Top-Down Processing", Nature Reviews Neuroscience, 2009, 10 (10), 704–716.
18. Takens F., "Detecting Strange Attractors in Turbulence", Dynamical Systems and Turbulence, Springer, 1981, 898, 366–381.
19. Rosenstein M.T., Collins J.J., De Luca C.J., "A Practical Method for Calculating Largest Lyapunov Exponents from Small Data Sets", Physica D: Nonlinear Phenomena, 1993, 65 (1–2), 117–134.
20. Chen T., Guestrin C., "XGBoost: A Scalable Tree Boosting System", Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2016, 785–794.
21. Akiba T., Sano S., Yanase T., Ohta T., Koyama M., "Optuna: A Next-Generation Hyperparameter Optimization Framework", Proceedings of the 25th ACM SIGKDD International Conference, ACM, 2019, 2623–2631.
22. DeLong E.R., DeLong D.M., Clarke-Pearson D.L., "Comparing the Areas Under Two or More Correlated Receiver Operating Characteristic Curves: A Nonparametric Approach", Biometrics, 1988, 44 (3), 837–845.
23. Heck C.N., King-Stephens D., Massey A.D., Nair D.R., Jobst B.C., Barkley G.L., Morrell M.J., "Two-Year Seizure Reduction in Adults with Medically Intractable Partial Onset Epilepsy Treated with Responsive Neurostimulation", Epilepsia, 2014, 55 (3), 432–441.
24. Kunhimangalam R., Joseph P.K., Sujith O.K., "Nonlinear Analysis of EEG Signals: Surrogate Data Analysis", IRBM, 2008, 29 (4), 239–244.