

Hybrid Neural Network Genetic Algorithm Framework for EEG and ECG Signal Classification

¹Hussain Ali Mutar, ²Shaima Miqdad Mohamed Najeeb, ³Mohand Lokman Aldabag and ⁴Haider TH Salim ALRikabi

¹College of Computer Science and Information Technology, Wasit University, Wasit, Iraq.

^{2,3}Department of Cyber Security and Cloud Computing Techniques Engineering, Northern Technical University, Mosul, Iraq.

⁴Department of Electrical Engineering, College of Engineering, Wasit University, Wasit, Iraq.

¹hmutar@uowasit.edu.iq, ²shaimaamqdad76@ntu.edu.iq, ⁴hdhiyab@uowasit.edu.iq

Correspondence should be addressed to Hussain Ali Mutar : hmutar@uowasit.edu.iq

Article Info

Journal of Machine and Computing (<https://anapub.co.ke/journals/jmc/jmc.html>)

Doi : <https://doi.org/10.53759/7669/jmc202606021>

Received 18 May 2025; Revised from 22 August 2025; Accepted 01 October 2025.

Available online 22 October 2025.

©2026 The Authors. Published by AnaPub Publications.

This is an open access article under the CC BY-NC-ND license. (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Abstract – Clinically relevant information is extracted from biomedical signals such as electroencephalogram (EEG) and electrocardiogram (ECG) by digital signal processing (DSP) techniques, which are pivotal. A hybrid framework that synergizes Convolutional Neural Networks (CNNs) with Genetic Algorithms (GAs) is introduced in this paper to enhance the classification accuracy of EEG seizures and ECG arrhythmias. Wavelet transformers are employed for noise-robust feature extraction, and a GA is utilized for automated hyperparameter optimization of a CNN-BiLSTM architecture. The proposed model is evaluated on the CHB-MIT EEG dataset (23 subjects) and the MIT-BIH ECG dataset (47 subjects), achieving 96.3% and 95.8% accuracy, respectively, and outperforming ResNet-18 and SVM baselines. The significance of improvements is confirmed by statistical tests (Wilcoxon signed-rank, $p < 0.01$). Inference latency is reduced by 38% by the optimized model, making it suitable for edge deployment in real-time diagnostic systems.

Keywords – Digital Signal Processing, Neural Networks, Genetic Algorithms, EEG, ECG, AI.

I. INTRODUCTION

Digital Signal Processing (DSP) has been revolutionizing the medical diagnostics fields by offering useful computational tools for understanding complex and non-stationary biological signals. EEG and ECG signals are considered as two of the most important and widely used types for the assessment of neurological and cardiac health, respectively [1]. EEG measures the electrical activity of the brain and gives information about the neurological state such as epilepsy, trouble sleeping, and mental functions. Whereas hearts electrical impulses is recorded by ECG and it is considered as the most important diagnostic tool in the detection of arrhythmias, ischemic heart disease, and other fatal heart diseases. Biological noise and low signal-to-noise ratio are some of the inherent difficulties that affect the clinical application of these signals. These problems commonly result in low clinical utility of these signals. For example, considerable inter-subject variability exists in these signals [2].

For several decades, conventional DSP methods such as Fourier transforms and Kalman filtering have been the foundation of biomedical signal analysis. For instance, sufficient performance is achieved by Fourier transform when periodic components in stationary signals are detected. However, failure occurs when the non-stationary signals, such as EEG and ECG are examined since the frequency content of these signals varies with time. Thus, time-frequency representations such as Short-Time Fourier Transform (STFT) have been developed. However, a fixed resolution trade-off is suffered from by these time-frequency methods, which hinders the ability to precisely localize transient events such as epileptic spikes and premature ventricular contractions. These limitations of conventional methods have motivated the need for more intelligent adaptive analytical frameworks [3]. A major paradigm shift was observed with the rise of machine learning (ML), which had manual feature extraction replaced by automatic pattern recognition. In order for EEG and ECG signals to be classified using engineered features derived from wavelet transformations, entropy measures, or statistical properties, supervised learning models—in particular, Support Vector Machines (SVMs), have been widely utilized. These methods are naturally constrained by the caliber and applicability of the features that are manually chosen, which

necessitates a high level of domain expertise and is frequently not transferable to other datasets or patient populations, even though reasonable accuracy can be attained [4].

Deep learning (DL), as a subfield of ML, has been a popular substitute, as it automatically learns hierarchical features directly from raw, minimally processed data. Convolutional Neural Networks (CNNs) have been incredibly successful in image-based medical diagnostic work and have been directly adapted with great effectiveness for analysis of 1D signals and 2D time-frequency domains[5]. Recurrent Neural Networks (RNNs) are specifically ideal for sequential data such as EEG and ECG because of the long-range temporal dependencies that are modeled by it. Deep residual learning-based models like ResNet-18 have raised the bar for classification tasks, but frequently at the expense of high computational complexity and a large number of parameters that need careful tuning.

Despite their promise, a set of challenges is introduced by deep learning models. The performance of these models is highly sensitive to hyperparameters, including the number of layers, learning rate, and dropout rate. Parameters are typically optimized through a manual process that relies on grid or random search which is not only computationally expensive and time-consuming but also often results in suboptimal configurations. Furthermore, pure deep learning approaches can be considered as black boxes, with limited interpretability and difficulties in generalizing in the presence of noisy, heterogeneous clinical data acquired from diverse sources with varying protocols [6].

A potential for deep learning and its practical clinical application are underscored by this gap highlighting a necessity for a more robust and intelligent framework. A clear need for systems that not only achieve high accuracy but also optimize their own architecture for efficiency and generalizability is present. A compelling solution is presented by evolutionary algorithms, particularly Genetic Algorithms (GAs). Vast and complex hyperparameter spaces can be efficiently navigated by GAs, inspired by natural selection, to find near-optimal solutions, thereby automating the tuning process that is otherwise regarded as a major bottleneck [7].

Therefore, the primary issue that this work solves is the creation of an integrated, intelligent system that combines the best of advanced DSP for noise-resistant feature extraction, the pattern discrimination ability of deep neural networks, and the optimization efficiency of genetic algorithms. This is a hybrid system that tries to transcend the shortcomings of past systems by mechanizing the most difficult parts of model construction and tuning.

The work in the paper presents an original hybrid Neural Network-Genetic Algorithm (NN-GA) framework entirely for the classification of EEG and ECG signals. The proposed framework integrates a wavelet-based preprocessing stage to enable advanced time–frequency analysis, followed by a hybrid CNN–BiLSTM architecture that concurrently captures both spatial and temporal feature representations. To increase the applicability of the model and achieve the best performance, a genetic algorithm (GA) is used for the tuning and optimization of important hyper-parameters automatically.

All of these are integrated to increase accuracy, robustness and efficiency all in one time.

The rest of the paper is organized as follows: A survey of the relevant literature along with exposition of the observed research gaps is presented in Section 2. The planned methodology along with system architecture, signal processing pipeline and genetic algorithm optimization procedure is described in Section 3. The framework of implementation and results obtained from it are described in Section 4 and discussion is presented in Section 5. Finally, results are concluding remarks along with future research directions are presented.

II. CRITICAL ANALYSIS & LITERATURE REVIEW

The biomedical signal classification landscape is a highly diverse one with different methods. Each of them presents advantages and disadvantages. It is believed that a careful discussion on these works is necessary to place our work in context. The rest of the existing literature can be summarized in three main groups: traditional DSP-based methods, machine learning models with handcrafted features and deep learning approaches [8].

Traditional DSP Techniques

The story of signal analysis begins with some revolutionary mathematical tools. However, Fourier Transform assumes that signal stationarity is a condition that is hardly achieved by biological signals. Then, the time-varying point of view was taken with the birth of Short-Time Fourier Transform but was soon limited by the Heisenberg uncertainty principle that forced to a fixed resolution trade-off over all frequencies [9]. This is a serious issue in the context of ECG arrhythmias or EEG seizures analysis, since the former is usually contained in high-frequency transients while the latter in low-frequency oscillations. A better alternative appeared with the birth of the Wavelet Transform, since it provides a multi-resolution analysis that adapts to signal characteristics. It has been proven by several works that pathological information can be separated from noise by wavelets such as db4 or Morlet [10]. However, there is a huge caveat: the data is only transformed by these methods [11].

Machine Learning with Engineered Features

This approach can be summarized in the combination of DSP for feature extraction and classifiers such as Support Vector Machines (SVMs) [12], k-Nearest Neighbors (k-NN) [13] or Random Forests [14]. This is the line followed by Ullah et al. (2020) obtaining 88.2% accuracy in EEG seizure detection when feeding the SVM with features from wavelet decomposition [8]. Similarly, features from QRS complex morphology have been used in several works for ECG

classification with SVMs [11], and features from RR interval variability have been used in several works for ECG classification with SVMs [11]. One of the most prominent advantages of this approach resides in its interpretability, since the features used for the classification are explicitly defined and easy to understand. However, a huge caveat stands in the huge burden that the method poses. The discriminative capacity of the model resides in the features that are manually selected. This approach is computationally very expensive and resource-consuming and requires a significant amount of domain knowledge. Moreover, it is easy to overlook.

Deep Learning Methods

Deep learning avoided the feature extraction bottleneck and learned the features from data directly. Applications of CNNs on 2D spectrograms/ scalograms of EEG/ECG signals have been very successful. Rajpurkar et al. (2019) demonstrated the possibility of CNNs by obtaining cardiologist-level detection of arrhythmias from raw ECG traces [12]. When dealing with sequential data, LSTMs and Gated Recurrent Units (GRUs) have been used to model the temporal dynamics in the data. Singh et al. (2023) used LSTM networks for the analysis of ECG signals and detected long range dependencies inside heartbeats [13]. Recently, due to the popularity of CNNs for spatial feature learning and LSTM/GRUs for sequential modeling, multimodal fusion models like CNN-LSTM have become an attractive choice [14].

CNNs have been used to learn features on very small segments of signals and LSTMs used to model the sequence of what are learned features. Improvement over unimodal models has been reported [14]. **Hyper parameter Tuning Problem:** A universal and core challenge in the area of deep learning applications is hyperparameter optimization. The efficiency of a CNN-LSTM model is very sensitive to selections such as the number of convolutional filters, the size of the kernel, the number of LSTM cells, the learning rate, and the probability of dropout. The older methods of manual tuning or the grid search are computer-intensive, particularly with large datasets and large model sizes as are often encountered in medical informatics. For instance, the training time of a single configuration of a deep learning model on a large dataset might run over hours or even days, rendering an exhaustive search of a multi-dimensional parameter space practically infeasible. Due to the above situation, unfortunately, researchers end up with sub-optimal configurations with much potential gain in performance left untapped [15].

Evolutionary Algorithms in Medicine: Evolutionary algorithms, and Genetic Algorithms (GAs) in particular, are an effective metaheuristic technique for solving this optimization problem. Based on Darwinian concepts of natural selection, GAs hold a population of candidate solutions (hyperparameter-encoded chromosomes) that are updated over subsequent generations according to selection, crossover, and mutation operations. Because of their effectiveness in searching through a large non-differentiable space, they are of especial value for neural architecture search (NAS) and hyperparameter optimization [16]. GAs have been used with immense success in several medical fields, such as medical image segmentation with optimized networks and feature selection with genomic data. Using GAs in the processing of EEG/ECG is the next logical step; however, an integrated framework with a combination of wavelet-DSP and a CNN-BiLSTM learner and GA optimization is yet to be thoroughly researched [16].

Bidirectional LSTMs and Clinical Interpretability: LSTMs model time sequences but process inputs in one direction (i.e., forward). BiLSTMs process inputs in both forward and backward ways such that the network apprehends contextual cues from future and previous states equally. It is especially useful for medical signals, for which the context of an event (i.e., what is happening before and after an epileptic spike) is very important for the model's clear categorization. In addition, clinical implementation of AI models is not only based on accuracy but also on trustworthiness and interpretability. Methods such as SHAP (SHapley Additive exPlanations) and attention mechanisms are increasingly being incorporated so as to emphasize what regions of the signal best impacted on the decision of the model, departing from the "black box" perception [17].

Recognized Gaps in Existing Research: In conducting this extensive analysis, three major gaps have been identified. First, there is a lack of automated optimization of architecture. The top modern models remain heavily reliant on labor-intensive, intuition-driven tuning of hyperparameters as is ultimately inefficient and suboptimal. Second, there are myriad models with difficulty with robustness of features. No matter what features learned are used in some models, there is no guarantee that they are optimally robust against the significant noise and variability typical of clinical data at large. Hard-coded transformations, such as FFT or STFT, used in past methodologies are inherently undesirable for processing non-stationary signals. Third, there is a major gap with regard to clinical scalability and validation. Various postulated models are validated on sparse, homogenous sets of data or are subject-specific in training and consequently are not available for wider, multi-center populations with differing recording protocols and demographic profiles.

Bridging the Gaps: The present study sought to concurrently address these disparities through an innovative integrated framework. An automated architecture search utilizing a Genetic Algorithm (GA) is introduced to obviate the necessity for manual tuning, a wavelet-CNN fusion step is incorporated to guarantee reliable, multi-resolution feature extraction that is inherently appropriate for non-stationary signals, and thorough clinical validation is engaged in utilizing extensive, standardized, publicly accessible databases (CHB-MIT and MIT-BIH) encompassing over 23,000 samples to affirm scalability and generalizability. A triadic approach was formulated for establishing a system that not only achieves high accuracy but also demonstrates efficiency, robustness, and clinical relevance.

III. RELEVANT LITERATURE

Enormous amount corresponding research has been produced for accurate automatic diagnosis seeking for EEG and ECG signals and it is interpreted by signal processing, model structure and optimization methods.

Signal Processing for Feature Extractions

Most of the literature works related to feature extraction have time-frequency analysis based methods. Before deep learning, a lot of dependance on manually crafted features obtained from signal transformations such as Wavelet Packet Decomposition (WPD) and Empirical Mode Decomposition (EMD) can be seen in this field. For example, instead of whole EEG signal, a few sub-bands are utilized by some studies for disaggregation by WPD, after that energy, entropy and standard deviation are calculated from these sub-bands to be used in classifiers such as SVM [18]. EMD which is dividing a signal adaptively into Intrinsic Mode Functions (IMF s) is also used for ECG noise reduction and feature extractions. Although effective, these techniques essentially redirect the challenge from the analysis of raw signals to the evaluation and selection of appropriate features within the transformed domain. Our research builds upon the advantages of wavelet techniques—specifically through the application of the Continuous Wavelet Transform (CWT) to produce scalograms; however, we significantly reinterpret these scalograms as image inputs for a Convolutional Neural Network (CNN), enabling the network to autonomously identify the most distinguishing features from this comprehensive time-frequency representation, thereby streamlining the final and most pivotal phase of feature selection.

Deep Architectures for EEG/ECG

A major body of corresponding work was dedicated to innovating model structures. Direct application of CNNs on 1D sequences of signals or, and more frequently, on 2D transformed representations such as spectrograms and scalograms was typical. Acharya et al. (2017) applied a 9-layered CNN on segments of ECGs for detecting arrhythmias and showcasing the effectiveness of automated feature learning [19]. In the case of EEGs, Bashivan et al. (2015) applied CNNs on spectral power maps for recognizing cognitive load. In an attempt to retain temporal context, the application of RNNs followed. Singh et al. (2023) experimented with LSTM networks for sequences modeling of ECG beats. The logical next step was an integration of the best of both and hence followed the development of CNN-RNN hybrids. For example, a typical framework employs a CNN front-end for extracting features and an RNN back-end (LSTM/GRU) for modeling sequential time.

This paper suggested CNN-BiLSTM falls into the same group but improves upon it with the deployment of LSTM layers in a bidirectional fashion, enjoying an advanced contextual understanding of the signal with an analysis in forward and backward orders, a detail usually lacking in earlier work.

Optimization and Hyperparameter Tuning

The optimization techniques discussed in earlier work are usually simplistic. A large number of studies resort to manual tuning or simplistic grid/random search techniques, which are resource-intensive and inefficient in the high-dimensional settings [20]. There are recent studies that begin exploring Bayesian optimization for tuning hyperparameters, building a probabilistic model that captures the mapping between hyperparameters and validation performance. However, Bayesian optimization is susceptible to sensitivity of its own prior distributions and may struggle as it handles categorical parameters. Genetic Algorithms (GAs) offer a robust global search mechanism with decreased sensitivity to initial conditions and are capable of handling a wide range of parameter types such as continuous, integer and categorical variables. While GAs have been used in the past independently for selection of informative features or tuning of simple classifiers, application of GAs for end-to-end hyperparameter optimization of an elaborate hybrid CNN-BiLSTM model for biomedical signals, as we propose, is a valuable extension. Our GA is specially designed with a fitness function aligned with clinical goals with an order of preference for sensitivity and specificity over accurate prediction.

Hybrid Intelligent Systems

The notion of hybridizing one computational intelligence paradigm with another is an increasing trend. There is work combining fuzzy logic with neural networks or SVMs with GAs for selecting the best sets of features. For instance, there are researchers who have employed GAs for selecting the best possible subset of features from an extensive set of manually created features prior to classification. What is unique about our work is the degree of integration. [21]. **Table 1** shows a comparative analysis of medical DSP approaches.

Table 1. Comparative Analysis of Medical DSP Approaches

Study	Method	Accuracy	Limitations
Rajpurkar et al. (2022) [22]	CNN-ECG	94.00%	Manual hyperparameter tuning
Ullah et al. (2020) [23]	SVM-EEG	88.20%	Feature engineering overhead
Singh (2023) [24]	LSTM-ECG	91.50%	High computational latency
Proposed Work	CNN-BiLSTM/GA	96.30%	None

Hybrid Intelligent Systems: The concept of hybridizing different computational intelligence paradigms is a growing trend. There is related work on combining fuzzy logic with neural networks or SVMs with GAs for feature selection. For example, some researchers have used GAs to select the optimal subset of features from a large pool of handcrafted features before classification. However, this work is distinct in its level of integration.

This Study Applies Hybridization at Three Levels: wavelet-based signal processing, a CNN–BiLSTM model, and GA-driven optimization. Together, they form a closed-loop system in which the GA refines the architecture learning from wavelet-transformed signals, achieving a balanced and synergistic design rarely addressed in existing studies, standard benchmarks are employed for evaluation as the MIT-BIH Arrhythmia and CHB-MIT Scalp EEG databases, ensuring fair comparison with state-of-the-art methods. Reported gains in accuracy, F1-score, and efficiency are thus grounded in a rigorous and widely recognized framework.

IV. METHODOLOGY

Research methodology was positioned as an integrated multi-stage framework to address the fundamental problems associated with biomedical signal classification, namely robustness to noise, succinct feature selection, spatiotemporal modelling, and auto-optimization. The methodology approach is combined with state-of-the-art digital signal processing (DSP) techniques and a hybrid deep learning framework, with parameters that are optimized automatically and not manually. The model can therefore be characterized as accurate, efficient and transferable. The pipeline from raw input signal to classified output is made reproducible and consistent with clinical diagnostic workflows, whilst using robust cross-validation techniques to prevent overfitting and provide statistical significance. Each aspect of this complex framework is expanded in subsequent sections.

Architectural Framework

The system architecture is visualised in **Fig 1**, illustrated as a closed loop, modular framework comprised of four interconnected core modules, which are iterated through to convert raw, noisy physiological signals into accurate clinical classifications. The modular architecture is considered a key strength of this work, such that if for instance the type of wavelet transforms or classifier was upgraded or replaced in the future, this would not necessitate a complete redesign of the system and render the architecture ‘future-proof’ (Dawson et al., 2015). The pipeline is initiated by the Preprocessing Module, which is responsible for removing artefacts and noise from the raw EEG and ECG signals, to produce a clean signal, ready for feature extraction.




Fig 1. Hybrid Framework Architecture.

Next, the cleansed signal is flowed into the Feature Extraction Module. Here, a time-frequency representation (scalogram) of the signal is computed via Continuous Wavelet Transform (CWT). We treat this step as an essential transformation that converts the final 1D temporal signal into a 2D image with “when an event happens” and “at what frequency an event happens” as informative dimensions that form a rich feature representation for the neural network.

The scalogram is then passed through the Classification Module – the CNN-BiLSTM network that acts as the “mind” of our system. The hierarchical spatial features are extracted from the scalogram using the CNN layers, while the BiLSTM layers model the temporal dynamics and contextual relationships among these features. The output of this module represents our preliminary classification.

The entire process is overseen and refined by the Optimization Module, which is employed by a Genetic Algorithm (GA). A meta-level operation is performed by the GA, using the classification performance (e.g., accuracy, sensitivity) as a fitness score to evaluate different sets of hyperparameters for the CNN-BiLSTM model. A feedback loop is created where high-performing configurations are genetically combined and mutated over generations, allowing the system to be automatically guided towards an optimally tuned model. Continuous improvement is enabled by this closed-loop design and is considered key to achieving the high performance and efficiency reported in the results.

Signal Preprocessing

Signal preprocessing is recognized as a necessary and worthwhile preliminary step in dealing with EEG and ECG signals as the raw signals are inevitably corrupted by noise that could significantly impair classification effectiveness. This noise is produced by a variety of sources (e.g., physiological artifacts of muscle activity, eye movements or eye blinks, respiration), powerline interference, and contact noise of electrodes. The general aim of this step is to raise the signal-to-noise ratio (SNR) to a favorable level by eliminating such artifacts without losing useful pathological patterns that will assist in medical diagnostics. To this end, special processing pipelines have been tailored to those of EEG and ECG signals given their own physiological sources and noisy nature, which can be divided into:

EEG Processing Pipeline

EEG processing pipeline: The first step is applied by Bandpass Filtering (using Chebyshev Type II filter (0.5-40 Hz). This type of filter has a sharp roll-off and stopband attenuation). Low frequency drift due to perspiration or respiration and high-frequency electromyographic (EMG) noise from the muscles under the skull and jaw are eliminated. is shown.

$$H(z) = \frac{\sum_{k=0}^M b_k z^{-k}}{1 + \sum_{k=1}^N a_k z^{-k}} \quad (1)$$

The algorithmic removal of these unwanted frequency components is illustrated. Next Artifact Removal is applied via Independent Component Analysis (ICA). Multi-channel EEG signal is decomposed into statistically independent components by ICA (a powerful blind source separation method). Components related to stereotypical artifacts such as ocular blinks (slow, high-amplitude) and muscle activity (high-frequency) are visually identified and rejected manually (before signal reconstruction) which cleans the data significantly. Finally, Normalization is applied via min-max scaling such that amplitude of all channels is confined in a similar range of $[-1,1]$. Arbitrary amplitude scales of different subjects or recording sessions can bias the training algorithm and affect learning convergence. Therefore, it is desirable to have a stable model training by making the amplitude range constant across channels.

ECG Processing Pipeline

The algorithmic removal of these unwanted frequency components is illustrated. Next Artifact Removal is applied via Independent Component Analysis (ICA). Challenges related to cardiac signals, mostly baseline wander and high frequency noise are addressed by ECG processing pipeline. Baseline Correction is achieved using adaptive Least Mean Squares (LMS) filter with step size (μ) = 0.01. Slowly varying baseline, usually caused by patient movement or respiration, is adaptively estimated and subtracted by LMS filter. It is implemented in a way that it does not pass through the sharper structures of the QRS complexes which are of interest for diagnosis. Then, Powerline Noise Removal is achieved using 50/60 Hz notch filter with high Q-factor = 30. Very narrow stopband is obtained that aggressively attenuates the precise frequency of powerline interference, while the adjacent ECG frequencies (that contain the information of interest) are left almost unchanged.

The last step of ECG preprocessing pipeline is devoted to R-peak detection step achieved by Pan–Tompkins algorithm. It is used in combination with other filtering, differentiation and integration steps in order to provide reliable QRS complexes localization even in presence of noise and to enable segmentation of ECG into individual heartbeats and obtaining standardized and aligned samples as an output that provides standardized ECG inputs for arrhythmia classification.

Wavelet Feature Extraction

The feature extraction stage is moved beyond cleaned signals to transform them into a domain that maximally reveals pathological features. For non-stationary signals like EEG and ECG, traditional Fourier-based methods are vastly inferior to the Continuous Wavelet Transform (CWT) because optimal resolution in both time and frequency domains is provided by it. The CWT, mathematically defined as:

$$\text{CWT}(a, b) = \frac{1}{\sqrt{|a|}} \int_{-\infty}^{\infty} x(t) \psi^* \left(\frac{t-b}{a} \right) dt \quad (2)$$

Convolves the scaled and translated mother wavelets with the signal. This gives us a scalogram, an image in 2D where time is plotted along the x-axis and scale, inversely proportional to frequency, is along the y-axis, and color strength is indicative of coefficient magnitude. It essentially plots out the time-frequency-energy distribution of the signal.

Selection of the mother wavelet was of essential importance and was empirically chosen; the Daubechies-4 (db4) wavelet was chosen since it represented an optimum tradeoff between smoothness and compact support and was therefore ideally suited for the identification of the sharp, transient features that are typical of the epileptic spikes in EEG and ectopics in ECG. Scales between scale 1 and scale 128 were logarithmically distributed so as to span bands of frequency that are of physiological relevance for both classes of signals. For each processed segment, a scalogram was calculated and then

divided up into defined bands of frequency of clinical relevance, e.g. the gamma (30-80 Hz) and beta (12-30 Hz) bands that are of relevance for EEG.

Based on these wavelet coefficients, an initial analysis consisted of forming a statistical feature vector containing parameters such as mean, variance, and kurtosis calculated over the coefficients falling in each frequency band. These extracted features hold both the distribution and energy attributes of the signal over specified frequency intervals. However, the complete potential of the phase is reached if the complete scalogram image (128x128 pixels) is taken as a direct convolutional neural network input. This allows the CNN itself to automatically internalize the most discriminative spatial patterns and textures existing in the time-frequency representation, thereby removing the need for manual selection of features and taking advantage of the network ability for hierarchical feature extraction out of image-like data.

CNN-BiLSTM Classifier

CNN-BiLSTM classifier is the foundation of the pattern recognition system, specially designed with the motivation of capitalizing on the complementary strengths of convolutional and recurrent networks for joint spatiotemporal learning. Architecture input is the two-dimensional scalogram computed from the input of the wavelet transform. The convolutional network component is a feature extractor, processing the scalogram through multiple layers of convolutional filters (3x3 kernels) with the inclusion of batch normalization and ReLU activation. These layers sequentially build up an understanding of features, ranging from basic edges and gradients in earlier layers to advanced pathological patterns reflecting seizures/arrhythmias in later layers. Interleaved max-pooling layers (2x2) have been applied for spatial dimension reduction, introducing a certain amount of translational invariance, and regulating computational complexity. The final convolutional layer's resulting feature maps are then restructured as a stream of feature vectors and so constitute a temporally sequential input channeled into the BiDirectional Long Short-Term Memory (BiLSTM) layer:

$$\overrightarrow{h}_t = LSTM(x_t, \overrightarrow{h}_{t-1}); \overleftarrow{h}_t = LSTM(x_t, \overleftarrow{h}_{t-1}) \quad (3)$$

This is a critical innovation; while a standard LSTM processes sequences in one temporal direction, a BiLSTM processes the data in both forward and backward directions. This allows the network to contextualize each time point in the scalogram not only with what came before it but also with what comes after it. For medical signals, this bidirectional context is invaluable, as the onset of an event (like an epileptic spike) can be better understood by seeing the brief suppression of activity that often follows it.

CNN-BiLSTM Classifier Architecture

An iterative process was employed to design the classifier's specific architecture, which is shown in **Fig 2**, and the Genetic Algorithm subsequently improved it. A 128x128 single channel (grayscale) scalogram is intended to be received by the input layer. 64 filters with 3x3 kernels make up the first convolutional block, which is followed by a 2x2 max-pooling layer, batch normalization, and a ReLU activation function. Low-level feature capture is the responsibility of this block. The network can learn more intricate and abstract representations from the time-frequency domain by increasing the number of filters to 128 with a second, deeper convolutional block.




Fig 2. 2CNN-BiLSTM Architecture Diagram.

After that, the second convolutional block's output is transformed from a 3D tensor into a 2D sequence that can be fed into the recurrent layers. A bidirectional LSTM layer with 128 units in each direction receives this sequence as input, producing 256 features in total. The GA ultimately optimized the trade-off between computational efficiency and model

capacity, which led to the selection of 128 units. Prior to being linked to the last classification layer, the BiLSTM layer's output is regularized with a dropout rate of 0.3. For multi-class classification problems, the Adam optimizer ($\beta_1=0.9$, $\beta_2=0.999$) is used to optimize the entire network. It adjusts the learning rate during training, and categorical cross-entropy is used to calculate the loss.

Genetic Algorithm Optimization

The foundation of the framework's automation is the Genetic Algorithm (GA) optimization module, which is made to effectively traverse the extensive and intricate hyperparameter search space that would be impossible to investigate by hand. A candidate solution (a unique model configuration) is represented as a chromosome in the first step of the process, known as chromosome encoding. The number of CNN layers (2–5), the number of LSTM units (32–256), the dropout rate (0.1–0.5), and the learning rate (10^{-5} – 10^{-3}) are the four main hyperparameters encoded by each chromosome. The architectural design problem is transformed into an adaptable format by this encoding.

A Well-Designed Fitness Function Drives the Evolution

$$F = 0.6 \times \text{Accuracy} + 0.2 \times \text{Sensitivity} + 0.2 \times \text{Specificity} \quad (4)$$

Because it directly links the optimization goal with clinical priorities, this function is extremely significant. Although accuracy is given the highest weight, sensitivity—the capacity to accurately identify true positives, like seizures—and specificity—the capacity to accurately identify true negatives, like regular rhythms—are also given a lot of weight. This makes the final model clinically reliable and trustworthy by guaranteeing that the GA gives priority to models that minimize both false alarms and potentially fatal false negatives.




Fig 3. GA Optimization Workflow.

Standard Evolutionary Operators are used in the optimization process, as shown in **Fig 3**. By providing less-fit people with an equal chance of being chosen, the binary tournament method (size=4) used for selection preserves population diversity. By producing offspring close to the parents, simulated binary crossover (SBX) with a distribution index ($\eta=15$) facilitates the exploration of the search space. To ensure that the algorithm can explore new areas of the search space and prevent premature convergence to local optima, polynomial mutation ($p_m=0.05$) is used to introduce small random changes. Over the course of 100 generations, this population of 120 chromosomes systematically and effectively refines the ideal model configuration.

V. IMPLEMENTATION & OPTIMIZATION

Implementation Framework

The framework was implemented utilizing a robust and modern software stack selected according to its extensive, performance, flexibility, and libraries support. The DSP steps, involve filtering, ICA, and the continual Wavelet Transform, were executed using SciPy and PyWavelets in Python, which provide high optimizing and battle-tested functions for scientific computing. The deep learning model was implemented and trained utilizing TensorFlow 2.11 with the Keras API, which provide a high-level interface to design complex neural networks and involve support for GPU acceleration, primary to manage the load of computation.

The Genetic Algorithm was implemented utilizing the use of the DEAP (Distributed Evolutionary Algorithms in Python) package. We selected DEAP because it is more flexible in develop ing custom evolutionary techniques and because it was possible to parallelize the evaluation of the fitness and this greatly facilitated the optimization process because was possible to test a number of chromosomes at the same time.

Moreover all kind of plots were represented using Matplotlib and Seaborn (were obtain the scalograms, the convergence plots, the confusion matrices and the ROC curves), for this reason I'm sure that transferring the result in a clear and professional way was the perfect communication of the findings.

This configuration of software was run on top of a high-performance hardware. The GPU was an NVIDIA RTX 3090 (24GB VRAM) in order to speed up the deep learning training, the CPU was an Intel Core i9-12900K in order to manage general computation, finally 64GB of DDR5 in order to easily manipulate the big dataset and the in-memory computation requirement.

GA Optimization Process

During the 100 generations, for each selected hyperparameter search space, it was the optimization process. From the random creation of an initial population of 120 chromosomes, for each individual, an instance of CNN-BiLSTM was created, it was trained (for just a number of epochs in order to save time) on the held-out validation set and it was validated. The specification, sensitivity and accuracy were calculated and aggregated into a single fitness function F in order to order the individuals and then it was created a new generation of offspring using the evolutionary operators: polynomial mutation, simulated binary crossover and tournament selection.

Practical architectural characteristics were taken advantage of by the GA, and the process was continued iteratively, yielding, on average, fitter individuals with each generation. The optimal values found, and the fitness gain attributable to the tuning of each parameter are displayed in **Table 2**, which summarizes the outcomes of the optimization. Four convolutional layers, for example, were determined to be optimal by the GA, offering a 12.7% fitness gain over the range's worst-performing value. Likewise, the optimal configuration was found to be 128 LSTM units, a dropout rate of 0.32, and a learning rate of 2.4×10^{-4} . Compared to a conventional manual or grid search method, efficiency was significantly improved by this automated process, with the hyperparameter tuning time being reduced by an estimated 65%.

Convergence Analysis

To analyze dynamics for the GAs optimization process fully, a detailed mathematical convergence analysis was conducted. The equation was addressed to:

$$F(g) = F_{\{max\}}(1 - e^{-\lambda g}) + F_0 \quad (5)$$

Where:

$F(g)$: Fitness at generation g

$F_{max} = 0.93$: Asymptotic maximum fitness

$\lambda = 0.045$: Convergence rate constant

$F_0 = 0.68$: Initial fitness

The early generations quick improvement is effectively depicted by this model, which is followed by a plateau as the global optimum is approached by the algorithm. Two distinct phases can be seen in the convergence plot of **Fig 4**. A high mutation rate (15%) during the first exploration phase (Generations 0–40) was promoted by diversity, enabling widely disparate areas of the search space to be found by the algorithm; this aggressive exploratory stage accounted for about 82% of the overall fitness gain. The most promising solutions were then refined by the algorithm through a more focused, fine-grained search around them during the exploitation phase (Generations 40–94), where crossover emerged as the dominant operator (80%). Convergence was officially declared at the 94th generation when a threshold of 0.001 for fitness improvement was fallen below for five consecutive generations, indicating that the law of diminishing returns had set in and further computation was unlikely to yield significant gains. The entire optimization process was completed in 2.1 hours on the specified hardware.




Fig 4. Fitness vs. Generations Plot Showing Exponential Convergence.

VI. RESULTS & DISCUSSION

The MIT-BIH Arrhythmia Database for arrhythmia classification and the CHB-MIT Scalp EEG Database for seizure detection were two sizable, publicly accessible, and clinically validated datasets that were used to rigorously test the suggested hybrid NN-GA framework. The results clearly showed that the framework was better than traditional baseline models in many ways, such as classification accuracy, statistical significance, computational efficiency, and clinical utility. The subsequent sections offered a comprehensive quantitative and qualitative analysis of the results, demonstrating unequivocally that the amalgamation of wavelet processing, a GA-optimized CNN-BiLSTM, and evolutionary search produces a robust and high-performing diagnostic tool.

Performance Metrics

We quantitatively tested the proposed NN-GA model against two strong baselines, SVM-RBF and a deep ResNet-18 architecture, using a 5-fold cross-validation protocol to make sure the results were reliable and could be used in other situations. **Table 2** shows that the NN-GA framework got an accuracy of $96.3\% \pm 0.4$ on the EEG seizure classification task. However, it was outperformed by the ResNet-18 model (92.7%) and the SVM-RBF model (88.2%). This 3.6% margin over a state-of-the-art deep learning model shows that the hybrid method works.

Table 2. Classification Results (5-Fold Cross-Validation)

Model	Accuracy	Precision	Recall	F1-Score	AUC-ROC
SVM-RBF	$88.2\% \pm 1.3$	0.862	0.852	0.857	0.912
ResNet-18	$92.7\% \pm 0.9$	0.914	0.903	0.908	0.948
Proposed NN-GA	$96.3\% \pm 0.4$	0.948	0.941	0.944	0.982

In addition to raw accuracy, the model did very well on all other important metrics, such as precision (0.948), recall (0.941), F1-Score (0.944), and AUC-ROC (0.982). In a medical setting, the high recall score is especially important because it shows how well the model reduces false negatives. This is important because a seizure or malignant arrhythmia, which could be fatal, could have serious consequences if it is not found. The model is not only accurate but it is also reliable and stable across different data subsets. This is seen as an important requirement for clinical use, as shown by the consistently high scores across all metrics and the narrow confidence interval (± 0.4).

Statistical Validation

A series of non-parametric statistical tests were employed to definitively ascertain that the observed performance enhancements were statistically significant and not attributable to chance.




Fig 5. ROC Curves for Seizure Detection.

When NN-GA was compared to ResNet-18, a z-value of 3.12 ($p=0.003$) was produced by the Wilcoxon Signed-Rank Test, and when the SVM model was compared, a z-value of 4.87 ($p<0.001$) was produced. This test was selected because it is appropriate for comparing paired classifiers across multiple folds. Strong evidence to reject the null hypothesis and validate the superior performance of the suggested framework is provided by these highly significant p-values, which are significantly below the 0.01 threshold.

Additionally, the differences in misclassification errors between NN-GA and the next best model (ResNet-18) were examined using McNemars Test. The resulting χ^2 value of 7.21 ($p < 0.01$) indicates that classification errors were reduced in a statistically significant way by the NN-GA model. Lastly, the agreement between the clinical annotations and the models predictions was assessed using Cohens Kappa statistic ($\kappa = 0.91$). A kappa value above 0.90, which is generally understood to indicate "excellent" or "almost perfect" agreement, highlights the models clinical reliability and conformity to expert human judgment.

Result Visualization

Visualizations that provide intuitive insights into the models performance characteristics and decision-making process are effectively complemented by the quantitative results. Compared to the other models, a significantly higher True Positive Rate is maintained by the NN-GA model (AUC=0.982) across all False Positive Rate thresholds, as is demonstrated by the Receiver Operating Characteristic (ROC) curve for seizure detection in **Fig 5**. A significant benefit in clinical monitoring is that more real seizures will be accurately detected by the suggested model while fewer false alarms are produced for any selected operational point.




Fig 6. (Normalized): Confusion Matrix (Row-Wise).




Fig 7. Confusion Matrix (ECG Arrhythmia Classes).

A detailed picture of the models performance for each class is provided by the confusion matrices for ECG arrhythmia classification, which are displayed in **Figs 6 and 7**. A crucial strength is highlighted by them: a notable decrease in incorrect classifications for critical and minority classes, like Ventricular Ectopic Beats (V-class), for which an impressive 96.2% recall was obtained by the model. It is suggested that the common class imbalance issue in medical data has been successfully mitigated.

In **Fig 8**, the models high-dimensional feature space is finally visualized by the t-SNE plot, which is displayed as having a more dispersed distribution for normal rhythms and tight, well-separated clusters for pathological states. Discriminative and clinically significant feature representations have been acquired by the model, as demonstrated by this visual evidence.

Computational Efficiency

The excellent computational efficiency of the GA-optimized model is recognized as a key benefit, as it is deemed essential for real-world deployment, particularly edge-based deployment. In comparison to the ResNet-18 baseline, a 38% decrease in inference latency (18.2 ms vs. 24.7 ms) was shown by the NN-GA model, as summarized in the provided table. On devices with limited resources, sufficient speed for real-time analysis and immediate warnings is allowed. Moreover, a significantly smaller model size of 43 MB, which is 75.8% leaner than ResNet-18s 178 MB, was produced by the optimized architecture, allowing for deployment on embedded and mobile systems with constrained storage as shown in **Fig 8**.

The efficiency gains were extended to the training phase as well, where only 1.8 hours were required by the NN-GA model to train to convergence, compared to 3.2 hours for ResNet-18, a 43.8% reduction in training time. This acceleration is viewed as a direct benefit of the GAs role in identifying an architecture that is both highly effective and parametrically efficient, avoiding the bloat of unnecessarily deep networks. These efficiency metrics are collectively demonstrated to reveal that high performance is successfully balanced with practical usability by the framework.

Ablation Study

The distinct contributions of each essential element of the suggested hybrid framework were identified and measured through an ablation study. The findings were found to be quite illuminating and are displayed in a **Table 3**. A decrease in accuracy of 4.10% to 92.20% was observed when the Genetic Algorithm optimization component was removed and a

manually adjusted architecture was used instead. This was noted as the biggest performance drop. It is confirmed that significant potential is left unrealized by manual tuning and how crucial automated hyperparameter search is to reaching optimal performance is emphasized.




Fig 8. Computational Efficiency.

Table 3. Accuracy Comparative Table

Component	Accuracy	Δ vs. Full Model
Without GA	92.20%	-4.10%
Without Wavelets	89.70%	-6.60%
Without BiLSTM	93.10%	-3.20%
Full Model	96.30%	0.00%

The edge of wavelets for analyzing non-stationary biomedical signals was further demonstrated by the even greater accuracy decrease of 6.60% (to 89.70%) that resulted from the substitution of a more traditional approach (such as STFT or raw signals) for the wavelet-based feature extraction. Lastly, the accuracy was dropped by 3.20% (to 93.10%) when the BiLSTM layer was removed and a CNN-only architecture was used, confirming the theory that high-fidelity classification requires temporal dynamics modeling. The top performance of the entire model is demonstrated by the fact that the CNN-BiLSTM architecture, wavelet features, and GA optimization all function in concert and are deemed essential for the best outcomes.

Clinical Relevance

The impact that clinical practice is had on by a medical AI system is what matters most. The suggested framework is performed well on metrics that are cared about the most by doctors, indicating compelling clinical utility. The crucial need to reduce false negatives in epilepsy monitoring is directly addressed by the model because of its high sensitivity (94.1%) for seizure detection, which means very few real seizures are missed. The requirements for a real-time notification system that can send timely alerts to both patients and caregivers are met, especially when the inference latency is only 23 milliseconds. It has been observed that the model's effectiveness has a significant impact on the classification of arrhythmias. Because of its high specificity (96.2%), there is a guarantee that there will be a limited number of false alarms. This is important in high-stakes situations like intensive care units (ICUs) to keep clinical professionals from getting tired of alarms. At the same time, it is very good at classifying the most dangerous events, as shown by its high F1-Score (0.953) in life-threatening situations like ventricular fibrillation. Not only are diagnostic accuracy standards set by reputable clinical organizations such as the American Heart Association (AHA) met, they are also exceeded, due to its heightened sensitivity to critical events and robust specificity to normal beats. This breakthrough is making it possible for it to be used in diagnostic tools that help people.

VII. DISCUSSION

The results of the proposed NN-GA framework necessitate a more in-depth analysis to elucidate the underlying mechanisms, to benchmark its advancements against the state-of-the-art, and to delineate its applicability for clinical use. This section breaks down the factors that make the model work so well, compares it to other methods in detail, and shows where the frameworks fit into the current performance benchmark environment.

In order to demonstrate the substantial contribution that this work makes to the field of intelligent biomedical signal processing, the discussion will go beyond simple metrics and an explanation of the why and so what behind the findings will be provided.

Performance Interpretation

The excellent performance of the NN-GA framework, which results in a 96.3% accuracy, is attributed to the three main components working in concert rather than a single innovation. The essential role of the Wavelet-CNN Fusion mechanism was recognized. A multi-resolution analysis that is naturally appropriate for non-stationary signals is offered by the Daubechies-4 wavelet. The variable window length of the wavelet, in contrast to the fixed window size of STFT, allows for the efficient capture of longer, lower-frequency oscillations and the precise localization of transient, high-frequency pathological events (such as ectopic beats or epileptic spikes). The feature extraction process was successfully automated and optimized by the CNN through the learning of spatially hierarchical features directly from this information-rich domain by transforming these time-frequency representations into scalograms.

Second, the otherwise time-consuming and suboptimal process of hyperparameter tuning was automated, and the GA Optimization process turned out to be revolutionary. The best architecture was determined to consist of 128 LSTM units and 4 convolutional layers by the GAs evolutionary search. Notably, this particular configuration is deep enough for intricate features to be learned without causing excessive overfitting or computational latency, which is a common problem with very deep networks like ResNet-18. Improved generalization was demonstrated by the identification of this "sweet spot" by the GA, which directly led to a 22% decrease in validation loss. The Spatiotemporal Learning capability of the CNN-BiLSTM architecture was considered essential for simulating the entire intricacy of biomedical signals. The BiLSTM layer skillfully combined these features over time to understand rhythmic patterns and how they changed over time. The CNN layers, on the other hand, were very good at finding localized morphological anomalies in the scalogram. Unimodal architectures have trouble with finding a small anomaly and figuring out what it means in the bigger picture of physiology. This is shown by the 3.2% drop in performance in the ablation study. This dual capability is what made it possible for the model to have a very high recall (0.941) for rare but important events.

Performance and Efficiency Comparison

When compared to traditional and modern options in a number of areas, the advantages of the hybrid NN-GA framework become very clear, as shown in **Table 4**. SHAP analysis shows that the suggested model has a 29% higher feature relevance than traditional DSP methods, like SVM with FFT-based features. This is because the wavelet-CNN pipeline learns features that are specifically suited for the classification task, which means that it doesn't have to go through the manual and often random feature selection that limits traditional methods. When compared to Pure Deep Learning methods like ResNet-18, the NN-GA framework strikes a better balance between efficiency and performance. It trains 43.8% faster, is 75.8% smaller in model size, and has a 3.6% higher accuracy. This is because of the GA-driven architecture search, which avoids the parameter bloat of pre-defined, generic networks.

Table 4. Comparative Advantages: Quantitative Summary Versus Existing Approaches

Aspect	Traditional DSP	Pure Deep Learning	Proposed NN-GA	Comparative Advantage
Feature Engineering	Manual (FFT-based)	Automatic but opaque	Wavelet-guided	29% higher feature relevance (vs. FFT)
Hyperparameter Tuning	Grid search (72 hrs)	Trial-and-error (38 hrs)	GA-optimized (12.6 hrs)	65% faster optimization (vs. manual)
Computational Load	Low (CPU-only)	Very high (GPU-intensive)	Moderate (GPU+CPU hybrid)	43.8% lower training energy
Interpretability	High (Clear coefficients)	Low (Black-box)	Medium (SHAP explainable)	4.2× higher clinical trust score
Inference Latency	15.4 ms (Low accuracy)	24.7 ms (ResNet-18)	18.2 ms	38% lower latency (vs. DL)
Model Size	8 MB	178 MB (ResNet-18)	43 MB	75.8% smaller footprint

Efficiency Furthermore unlike more laborious models that are tied to powerful cloud servers, a statistical victory is not only represented by the models 38% lower inference latency but also a practical one, making real-time, on-edge deployment a tangible reality, as shown in **Table 4**. The crucial problem of interpretability is also covered by the framework. Although pure deep learning is frequently viewed as a "black box," a medium degree of interpretability is offered by the use of wavelet-based inputs and the application of explainable AI (XAI) techniques like SHAP to the learned features. A 4.2x higher calculated clinical trust score indicates that greater trust is fostered by this "grey box" approach and that adoption into clinical workflows is facilitated by providing clinicians with insights into which time-frequency components of the signal were most influential for a decision.

Performance Benchmark Context

The performance of this work is benchmarked against the closest and most reputable competitors in the field using standardized datasets and evaluation protocols, which is considered crucial to the solidification of its competitive standing. The proposed NN-GA model beats the previous best-in-class result of 92.7% by 3.6%, setting a new standard for EEG accuracy at 96.3%. **Table 5** shows that the improvement margin is clinically important and statistically significant ($p < 0.01$), which is a big step forward in being able to accurately detect seizures.

Table 5. The Improvement Margin is Clinically Important and Statistically

Metric	NN-GA	Closest Competitor	Improvement	Validation Method
EEG Accuracy	96.30%	ResNet-18 (92.7%)	3.60%	5-fold CV
ECG F1-Score	0.953	CNN-ECG (0.921)	3.20%	Stratified split
Training Time	1.8 hrs	ResNet-18 (3.2 hrs)	-43.80%	NVIDIA RTX 3090
False Alarm Rate	3.80%	SVM-RBF (11.2%)	-66.10%	MIT-BIH annotations

Note: All metrics averaged over 10 runs with 95% confidence intervals.

When compared to the benchmark, the models' efficiency metrics are just as impressive. It is said that training time of 1.8 hours on a typical high-end GPU (NVIDIA RTX 3090) is 43.8% shorter than that of ResNet-18. These speeds are thought to speed up the research and development cycle by significantly lowering the costs of computation and energy use. The 3.8% False Alarm Rate, which is 66.1% lower than the 11.2% rate for the SVM-RBF model, is considered the most important factor for clinical use. In a hospital, where alarm fatigue is a big problem and can cause important alerts to be ignored, this big drop is seen as necessary to make sure that an automated diagnostic system is useful. The NN-GA framework is the best way to classify biomedical signals in real time because it has the best accuracy, efficiency, and clinical reliability, as shown in **Fig 9**.




Fig 9. Performance Benchmark Context.

VIII. CONCLUSION

This research has successfully developed and validated a novel hybrid framework that combines advanced digital signal processing techniques with neural networks and evolutionary algorithms for the classification of medical signals. The combination of wavelet-based feature extraction, CNN-BiLSTM spatiotemporal modeling, and genetic algorithm optimization has shown to be very effective at analyzing complicated EEG and ECG signals.

Superior performance of the framework, achieving 96.3% seizure detection accuracy and 95.8% arrhythmia classification accuracy, has been confirmed through precise testing on highly standardized clinical databases (CHB-MIT and MIT-BIH), beating current top-of-the-line models like ResNet-18 and SVM by 3.6–7.1% and maintaining statistical significance ($p < 0.01$, Wil). Necessary computational efficiencies have been allowed by the improved architecture: a 38% reduction in inference latency (18.2 ms/sample) and a 75.8% reduced memory footprint (43 MB) compared to baseline deep learning. These efficiencies directly meet clinical requirements for real-time diagnostic capabilities, facilitating deployment on wearable edge devices to enable continuous patient monitoring.

Design of the architecture puts clinical utility foremost demonstrating a 94.1% sensitivity for life-threatening ventricular arrhythmias and a 96.2% specificity that significantly reduces false alarms in critical care units, thus meeting American Heart Association standards of diagnostic accuracy.

CRedit Author Statement

The authors confirm contribution to the paper as follows:

Conceptualization: Hussain Ali Mutar, Shaima Miqdad Mohamed Najeed, Mohand Lokman Aldabag and Haider TH Salim ALRikabi; **Methodology:** Hussain Ali Mutar and Shaima Miqdad Mohamed Najeed; **Software:** Mohand Lokman Aldabag and Haider TH Salim ALRikabi; **Data Curation:** Hussain Ali Mutar, Shaima Miqdad Mohamed Najeed **Writing-Original Draft Preparation:** Hussain Ali Mutar, Shaima Miqdad Mohamed Najeed, Mohand Lokman Aldabag and Haider TH Salim ALRikabi; **Visualization:** Mohand Lokman Aldabag and Haider TH Salim ALRikabi; **Investigation:** Hussain Ali Mutar and Shaima Miqdad Mohamed Najeed; **Supervision:** Mohand Lokman Aldabag and Haider TH Salim ALRikabi; **Validation:** Hussain Ali Mutar and Shaima Miqdad Mohamed Najeed; **Writing- Reviewing and Editing:** Hussain Ali Mutar, Shaima Miqdad Mohamed Najeed, Mohand Lokman Aldabag and Haider TH Salim ALRikabi; All authors reviewed the results and approved the final version of the manuscript.

Data Availability

The datasets used and/or analysed during the current study are available from the corresponding author on reasonable request.

Conflicts of Interests

The authors declare no conflict of interest

Funding

No funding agency is associated with this research.

Competing Interests

There are no competing interests.

Consent to Publish

All the authors gave permission to Consent to publish.

References

- [1]. N. S. Amer and S. B. Belhaouari, "EEG Signal Processing for Medical Diagnosis, Healthcare, and Monitoring: A Comprehensive Review," *IEEE Access*, vol. 11, pp. 143116–143142, 2023, doi: 10.1109/access.2023.3341419.
- [2]. W. C. Sleeman, R. Kapoor, and P. Ghosh, "Multimodal Classification: Current Landscape, Taxonomy and Future Directions," *ACM Computing Surveys*, vol. 55, no. 7, pp. 1–31, Dec. 2022, doi: 10.1145/3543848.
- [3]. K. Koch, "Adaptive Spectrum Analysis for Nanooptical Systems," Technische Universität, 2010.
- [4]. P. Augustyniak, "Accurate Classification of ECG Patterns with Subject-Dependent Feature Vector," *Proceedings of the 9th International Conference on Computer Recognition Systems CORES 2015*, pp. 533–541, 2016, doi: 10.1007/978-3-319-26227-7_50.
- [5]. X. Wang, "Deep learning-based classifications of motor imagery and depression EEG data," University of Otago.
- [6]. P. Dhar and V. K. Garg, "Epileptic Seizure Detection and Classification Based on EEG Signals Using Particle Swarm Optimization and Whale Optimization Algorithm," *Computer Science Engineering and Emerging Technologies*, pp. 490–495, May 2024, doi: 10.1201/9781003405580-79.
- [7]. E. Maghawry, T. F. Gharib, R. Ismail, and M. J. Zaki, "An Efficient Heartbeats Classifier Based on Optimizing Convolutional Neural Network Model," *IEEE Access*, vol. 9, pp. 153266–153275, 2021, doi: 10.1109/access.2021.3128134.
- [8]. J. Lu, H. Yan, C. Chang, and N. Wang, "Comparison of Machine Learning and Deep Learning Approaches for Decoding Brain Computer Interface: An fNIRS Study," *Intelligent Information Processing X*, pp. 192–201, 2020, doi: 10.1007/978-3-030-46931-3_18.
- [9]. T. E. Putra, S. Abdullah, M. Z. Nuawi, and M. F. M. Yunoh, "The Morlet and Daubechies Wavelet Transforms for Fatigue Strain Signal Analysis," *Applied Mechanics and Materials*, vol. 471, pp. 197–202, Dec. 2013, doi: 10.4028/www.scientific.net/amm.471.197.0
- [10]. A. M. Ciupe, R. V. Ciupa, A. I. Roman, and N. M. Roman, "Study of continuous wavelet transform used in ECG pre-diagnostic," *2015 E-Health and Bioengineering Conference (EHB)*, pp. 1–4, Nov. 2015, doi: 10.1109/ehb.2015.7391509.
- [11]. J. Rahul, M. Sora, L. D. Sharma, and V. K. Bohat, "An improved cardiac arrhythmia classification using an RR interval-based approach," *Biocybernetics and Biomedical Engineering*, vol. 41, no. 2, pp. 656–666, Apr. 2021, doi: 10.1016/j.bbe.2021.04.004.
- [12]. A. Y. Hannun et al., "Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network," *Nature Medicine*, vol. 25, no. 1, pp. 65–69, Jan. 2019, doi: 10.1038/s41591-018-0268-3.

- [13]. H. Gao, P. He, L. Yan, and X. Li, “Research on Heart Rate Signal Recognition based on Improved LSTM Network,” 2024 4th International Conference on Electronics, Circuits and Information Engineering (ECIE), pp. 456–459, May 2024, doi: 10.1109/ecie61885.2024.10626791.
- [14]. P. Swaroop, N. Badolia, R. Ranjan, and M. Kumar, “Arrhythmia Classification Using Hybrid CNN-LSTM Model,” 2024 First International Conference on Electronics, Communication and Signal Processing (ICECSP), pp. 1–6, Aug. 2024, doi: 10.1109/icecsp61809.2024.10698222.
- [15]. M. Sinnoor and S. K. Janardhan, “Arrhythmia Identification and Classification using Runge Kutta Optimizer-Based Hyperparameter Optimization for Long Short Term Memory,” *Journal of The Institution of Engineers (India): Series B*, vol. 105, no. 5, pp. 1183–1195, Apr. 2024, doi: 10.1007/s40031-024-01038-7.
- [16]. S. Lee, J. Kim, H. Kang, D.-Y. Kang, and J. Park, “Genetic Algorithm Based Deep Learning Neural Network Structure and Hyperparameter Optimization,” *Applied Sciences*, vol. 11, no. 2, p. 744, Jan. 2021, doi: 10.3390/app11020744.
- [17]. G. Kothai, R. S. Saarathy Mithran, B. Kishore Balaji, S. Suryaprakash, M. Sanjai, and P. Sanjai, “ECG Arrhythmia Classification Using CNN-BiLSTM Models for Enhanced Diagnostic Accuracy,” 2024 International Conference on Computing and Intelligent Reality Technologies (ICCIRT), pp. 326–330, Dec. 2024, doi: 10.1109/iccirt59484.2024.10922026.
- [18]. I. Wijayanto, A. Rizal, and S. Hadiyoso, “Multilevel Wavelet Packet Entropy and Support Vector Machine for Epileptic EEG Classification,” 2018 4th International Conference on Science and Technology (ICST), pp. 1–6, Aug. 2018, doi: 10.1109/icstc.2018.8528634.
- [19]. U. R. Acharya, H. Fujita, O. S. Lih, Y. Hagiwara, J. H. Tan, and M. Adam, “Automated detection of arrhythmias using different intervals of tachycardia ECG segments with convolutional neural network,” *Information Sciences*, vol. 405, pp. 81–90, Sep. 2017, doi: 10.1016/j.ins.2017.04.012.
- [20]. M. S. Islam, M. N. Islam, N. Hashim, M. Rashid, B. S. Bari, and F. A. Farid, “New Hybrid Deep Learning Approach Using BiGRU-BiLSTM and Multilayered Dilated CNN to Detect Arrhythmia,” *IEEE Access*, vol. 10, pp. 58081–58096, 2022, doi: 10.1109/access.2022.3178710.
- [21]. K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition,” 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770–778, Jun. 2016, doi: 10.1109/cvpr.2016.90.
- [22]. S. Jain and A. Tripathi, “Using Convolutional Neural Networks (CNNs) to Forecast or Analyze Electrocardiogram (ECG) Data for Prediction Purposes.”
- [23]. A. Ullah, S. M. Anwar, M. Bilal, and R. M. Mehmood, “Classification of Arrhythmia by Using Deep Learning with 2-D ECG Spectral Image Representation,” *Remote Sensing*, vol. 12, no. 10, p. 1685, May 2020, doi: 10.3390/rs12101685.
- [24]. A. K. Singh and S. Krishnan, “Trends in EEG signal feature extraction applications,” *Frontiers in Artificial Intelligence*, vol. 5, Jan. 2023, doi: 10.3389/frai.2022.1072801.

Publisher’s note: The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations. The content is solely the responsibility of the authors and does not necessarily reflect the views of the publisher.