

PATIENT-INVARIANT ECG ARRHYTHMIA CLASSIFICATION VIA PER-BEAT NORMALIZATION AND STRICT INTER-PATIENT VALIDATION

DEEKSHIKA SIVA VINNY*¹, Dr. JEEVAN JALA², Dr V SANGEETHA³, JHANSI PANDIRI*⁴, Dr KANDUKURI GEETHA⁵, Dr. A MAHENDAR⁶, SURESH DAMERUPPULA*⁷

¹M.Tech Scholar, Department of Computer Science and Engineering, Koneru Lakshmiah Education Foundation, Andhra Pradesh, India, 522502.

²Assistant Professor, Department of Computer Science and Applications. Koneru Lakshmiah Education Foundation, Andhra Pradesh, India, 522502.

³Assoc Professor, Department of Computer Science and Engineering (Cyber Security), Swami Vivekananda Institute of Technology, Secunderabad, Telangana 500003.

⁴Assistant Professor, Dept of Electronics and Communication Engineering, Aditya University, Surampalem, Andhra Pradesh, India.

⁵Assistant Professor, Department of Computer Science and Engineering, CVR College of Engineering, Ibrahimpatnam, Ranga Reddy, Telangana, India.

⁶Associate Professor, Department of Computer Science and Engineering (Data Science), CMR Technical Campus, Hyderabad-501401, Telangana, India.

⁷SAP Technical Architect, MYGO Consulting Inc., 7307 Tamarron Dr, New Jersey, USA 08536

E-mail: ¹ndeekshika@gmail.com, ²jjeewan@kluniversity.in, ³sangeetha12@gmail.com, ⁴jhansi.pandiri8@gmail.com*, ⁵kandukurigeetha01@gmail.com, ⁶mahi.adapa@gmail.com, ⁷sureshdameruppula186@gmail.com

ABSTRACT

The cardiovascular diseases have been the most important cause of death worldwide, and the electrocardiogram (ECG) analysis is the diagnostic foundation of arrhythmia. Although accuracies of 99.0 have been reported in research, there is a deadly industrial/clinical disconnect: 86.7% of existing studies use intra-patient evaluation paradigms that do leak patient-specific morphological patterns, giving inflated measures that collapse disastrously on unseen patients. This research paper helps fill this important safety gap because it makes the hypothesis that clinical utility is more important to evaluation methodology and feature normalization than to architectural complexity. Our suggested lightweight three component dual lead convolutional neural network will include per beat z-score normalization to eliminate patient specific amplitude leakage, inverse square root class weighting with focal loss to overcome extreme class imbalance, and physiological R-peak jitter augmentation to be robust to annotation noise. When strict patient-disjoint validation is applied to the MIT-BIH Arrhythmia Database (trained on 32 patients, tested on 5 completely unknown patients), we obtain a macro F1-score of 0.714. It is worth noting that the model has an outstanding recall (99.7%) and F1 (91.5) on life-threatening Ventricular ectopic beats. Ablation experiments ensure that normalization per-beat is the primary force behind the phenomenon (+0.210 macro F1 gain), whereas t-SNE plots ensure the creation of patient-invariant feature embeddings. These results criticize the modern trend of growing the complexity of architecture and show that strict assessment and normalization are the major factors of the reliability of cross-patient arrhythmia detection.

Keywords: ECG Arrhythmia Classification, Inter-Patient Validation, Patient-Invariant Features, Deep Learning, MIT-BIH Database, Class Imbalance, Focal Loss.

1. INTRODUCTION

The combined and perfected Introduction is this. It also holds the flow and references of your original paper but incorporates the particular structural needs

(Hypothesis, Scope, Research Design, and Significance) of the revised version to meet the reviewer response but without needless wastage. Cardiovascular diseases are still the main cause of death in the world and arrhythmias are one of the

major causes of cardiac morbidity and sudden cardiac death. The diagnosis of arrhythmia is based on Electrocardiogram (ECG) which is time consuming, subjective and can be subject to human error especially in the clinical practices with limited resources. Although deep learning becomes a disruptive solution to automatic ECG analysis, where convolutional neural networks (CNNs) show promising results on standard data, there is an essential disconnect between research specifications and clinical reality. What is important about this issue is that the evaluation paradigms that prevail in literature are currently in question. Recent systematic reviews show that 86.7 per cent of the studies use intra-patient evaluation paradigms, training and testing on beats of the same patient. This is a serious issue since by doing this, patient-specific morphological patterns (e.g. constant QRS amplitude or baseline wander) are inadvertently captured, as opposed to generalizable arrhythmia characteristics. This raises the inflated accuracy values (usually greater than 99) of models that fail catastrophically when applied to unknown patients, which is a major risk to patient safety in clinical practice. However, inter-patient validation is definitively used only in 10.3% of studies, in which models are compelled to generalize across dissimilar groups of patients. This assessment bias has promoted the creation of complex hybrid models (CNN- LSTM-Transformer ensembles) to optimize intra-patient results and to disregard cross-patient resilience.

The management of class imbalance and signal representation is a controversial issue. Unlike rare forms of arrhythmia like fusion beats (F) and ventricular ectopics (V), most approaches apply synthetic oversampling (SMOTE) or basic class weighting, which results in samples that are physiologically unstable, and generally disregards life threatening classes. Additionally, despite the fact that the MIT-BIH database has dual-lead recordings (MLII + V1), most studies analyse single leads only, without trying to use the complementary spatial information that might enhance discriminative power. To overcome these issues, we have a hypothesis that patient-invariant feature learning via per-beat normalization and rigid inter-patient validation produces more clinically feasible models than just adding increased architectural complexity. We have hypothesized that under strict validation procedures, a lightweight dual-lead CNN can provide competitive inter-patient results, without the computational cost of ensemble models based on hybrid models. The current work is concerned with beat-level classification based on the standard AAMI

mapping (Normal, Supraventricular, Ventricular, Fusion) on MIT-BIH Arrhythmia Database. This research encompasses dual-lead signal processing, per beat normalization, strict patient-disjoint splits as well as the management of class imbalance through inverse square root weighting. Out of the scope of the work, there are rhythm-level diagnosis (Aux annotations), multi-center validation (e.g., PTB-XL), and restrictions concerning the real-time implementation on wearable devices. To reach these targets, we use quantitative experimental research design which includes the following steps: (1) Data Partitioning: Strict patient-disjoint splits (31 training, 5 validation, 5 unseen test patients) to ensure the absence of patient leakage; (2) Preprocessing: Per-beat z-score normalization to restrict the learned features to be patient-independent; (3) Augmentation. The proposed framework shows that, in contrast to the current tendency of progressively growing the intricacy of architecture, a lean dual-lead CNN, with weighting of inverse-square-root classification and the use of a focal loss, results in competitive inter-patient performance (macro F1-score of 0.714) without making highly demanding computational requirements. Our solution provides a premise of reliability in terms of clinically viable ECG arrhythmia detection, which switches the emphasis on optimization on the benchmark to clinical generalizability.

2. RELATED WORK

In recent years, deep learning and transformer models have become more widely used in the field of cardiovascular disease recognition and arrhythmia classification. The use of deep learning and transformer models has grown in recent years, particularly in the realm of cardiovascular disease recognition and arrhythmia classification. Abousaber et al. (2025)[1] proposed a spatio-temporal transformer for explainable ECG classification using hybrid loss optimization, emphasizing the interpretability in biomedical signal processing. In the same way, Al-Tam et al. (2024)[2] presented a hybrid transformer–encoder architecture that showed excellent performance in the analysis of heart sounds for cardiovascular disease detection. Building on this approach, Anjum et al. (2023)[3] designed a temporal transformer based fusion model for morphological arrhythmia classification and Chen et al. (2023)[4] proposed an end-to-end hybrid deep network 1DCNN-TRNet that combines convolutional and transformer layers. Attention mechanisms have also played a prominent role: Chopannejad et al. (2024) [5] employed CNN-

BILSTM-BiGRU model along with SMOTE–Tomek resampling to improve the accuracy of arrhythmia detection; Din et al. (2024)[6] adopted ensemble learning using deep spatial–temporal features and long-range dependencies. Not surprisingly, classic methods are still relevant, with de Chazal et al. (2024)[7] introducing heartbeat classification based on interval and morphological features and Hannun et al. (2019)[8] using deep neural networks to attain cardiologist-level arrhythmia detection. Recently, El-Ghaish and Eldele (2024)[9] developed ECGTransForm, CNN-Attention-BI Transformer by Guo et al. (2025)[10], and hybrid CNN-LSTM transformer models for imbalanced ECG data (Hasan & Hossain, 2025)[11]. Liu (2025)[12] also proposed a model that combines 1D-CNN with BERT for intelligent arrhythmia classification and Kim et al. (2025)[13] applied Stockwell transforms to a CNN-transformer hybrid without using R-peak detection. Deep learning opportunities and challenges in ECG analysis were systematically examined by Hong et al., (2020)[14] and multi-head attention was highlighted by Kerdoudi et al., (2026)[15] for improved arrhythmia detection. Taking together, these studies span the entire spectrum from the early CNN real-time classification, Kim et al. (2025)[16] proposed a novel hybrid CNN–transformer architecture for arrhythmia detection that eliminates the need for R-peak identification by leveraging the Stockwell transform. Their study, published in Scientific Reports, demonstrated that combining convolutional feature extraction with transformer-based sequence modeling can enhance classification accuracy while reducing preprocessing complexity, thereby offering a more efficient and scalable solution for ECG-based diagnostic systems. (Kiranyaz et al., 2016)[17] to the most recent transformer hybrids, exemplifying the path of progress taken by ECG-based diagnostic systems towards explainability, robustness and clinical use.

2.1 Research Gaps

1. Paradigm bias on evaluation: 86.7 percent of currently available studies intra-patient validation leaks morphological patient-specific patterns, inflate the accuracy measures but do not extrapolate to the unobserved patients in clinical settings [18].

2. Where Class Imbalance Minimal Handling This has been used in existing methods with simple SMOTE or class weighting ignoring physiological plausibility, lacking support of the clinical severity

hierarchy with life-threatening arrhythmias (V, F) being more sensitive than benign variants[19].

3. Meanwhile underutilization of Multi-Lead Information: Möst studies compute single leads (usually MLII) independently without considering the spatial association between simultaneous leads (MLII + V1) that can yield complementary diagnostic information to make strong arrhythmia discrimination[20].

4. Inability to learn Patient Invariant Features: Current architectures can learn patient specific amplitude and baseline effects as opposed to learning normalized morphological patterns and so they lack cross-patient[21] generalization even when describing the data inpatients.

5. Annotation Noise Sensitivity The models are trained on perfect R-peaks manually annotated truly without considering real-world variation in beat detection algorithms using our models, resulting in them being generally sensitive to small temporal variations, which happen in practical usage of the algorithm in clinical settings.

2.2 Contributions and Novelty

1. Strict Inter-Patient Validation Protocol: We use strict patient-disjunct splits (31 training, 5 validation, 5 unseen test patients) of patients that wipe out patient leakage, and give clinically realistic scores on the performance, in 86.7% of the literature [22].

2. Per-Beat Normalization in support of Invariance to the Patients: Our model uses z-score normalization per-beat (this is, 2- means+2-SDs) to ensure that the model can learn morphological patterns without being affected by patient-specific amplitude changes and baseline wander.

3. Physiological Data Augmentation with R-peak Jitter: We would also introduce temporal jitter (maximum 3 samples) on R-peak locations when training; this would partially simulate annoyance noise expression by the real world and make models more robust, without producing uncharacteristic synthetic samples.

4. Inverse-Square Root Class Weighting with Focal Loss: Our loss uses the weighted gradient tilt (reducing the risk of excessive imbalance, yet keeping the gradient flow) on inverse-square root of

the classes of a sample to focus on the hard-to-classify minority samples, and offers performance comparable to that of first-order classifier models trained on all AAMI classes (N, S, V, F).

5. Simple But Effective Dual-Input CNN Architecture: Although complex hybrid architectures (CNN-LSTM-Transformer ensembles) are currently likely to achieve better performance, we find that a simple and yet effective dual-input CNN architectures (both MLII and V1) produces competitive inter-patient F1-scores (0.714 macro F1) at low computational costs, indicating that it can be implemented in practice with little to no modifications to a clinical context.

3. METHODOLOGY

In this paper the proposed method is designed for inter-patient BP ECG arrhythmia classification[23]. The three weaknesses of current literature covered by our framework include patient-specific feature leakage, extreme classes imbalance/balance, and sensitivity to noise in the annotation. The system is based on a simple CNN architecture that utilizes per-beat normalization, inverse-square-foot class weighting, and per-beat physiological R-peak jitter reactor valid on unknown patients.




Figure 1: Diagrammatic illustration of the entire methodology pipeline to show dual-lead signal processing, per-beat normalization, and improvement methods, as well as classification framework to identify inter-patient harmful ECG arrhythmia.

3.1 Data Preprocessing and AAMI Mapping

Let $x_{MLII}^{(p)}[n]$ and $x_{V1}^{(p)}[n]$ denote the MLII and V1 lead signals for patient p , where $n = 0, 1, \dots, N_p - 1$ represents the discrete-time samples. The annotated R-peak locations are denoted as $\mathcal{R}^{(p)} = \{r_k^{(p)}\}_{k=1}^{K_p}$, where each $r_k^{(p)}$ is associated with a beat type label represents the original MIT-BIH beat annotations.

$\Phi: \mathcal{L}_{raw} \rightarrow \mathcal{L}_{AAMI}, \mathcal{L}_{AAMI} = \{N, S, V, F\}$ (1)
We map these raw annotations to the standard AAMI classes using the mapping function

$$\begin{aligned} \Phi: \mathcal{L}_{raw} \rightarrow \mathcal{L}_{AAMI}, \mathcal{L}_{AAMI} = \{N, S, V, F\} \\ N \text{ if } l \in \{N, L, R, e, j\} \\ S \text{ if } l \in \{A, a, J, S\}, \\ V \text{ if } l \in \{V, E\}, \\ F \text{ if } l = F \end{aligned} \quad (1)$$

Equation (1) consolidates 12 raw beat types into 4 clinically meaningful categories: Normal (N), Supraventricular ectopic (S), Ventricular ectopic (V), and Fusion (F) beats[24].

3.2 Beat Extraction and Windowing

For each R-peak location $r_k^{(p)}$, we extract a window of $W = 256$ samples centered at the R-peak. The windowed signals for the k -th beat are defined as:

$$w_{MI}^{(p,k)}[i] \quad (2)$$

$$= x_{MLII}^{(p)} \left[r_k^{(p)} - 128 + i + \delta_k \right], \quad i = 0, 1, \dots, 255$$

$$w_V^{(p,k)}[i] \quad (3)$$

$$= x_{V1}^{(p)} \left[r_k^{(p)} - 128 + i + \delta_k \right], \quad i = 0, 1, \dots, 255$$

where δ_k represents the R-peak jitter augmentation parameter. During training, $\delta_k \sim \mathcal{U}(-3, 3)$ is sampled uniformly from the range $[-3, 3]$ to simulate annotation noise and improve robustness. During inference, $\delta_k = 0$. The jitter augmentation in (2)-(3) prevents overfitting to precise R-peak locations and mimics real-world beat detection variability.

3.3 Per-Beat Normalization

To enforce patient-invariant feature learning and eliminate amplitude-based patient identity leakage[25], we apply per-beat z-score normalization to each extracted window:

$$\begin{aligned}
 & \widehat{w_{MLII}^{(p,k)}}[i] \\
 &= \frac{w_{MLII}^{(p,k)}[i] - \mu_{MLII}^{(p,k)}}{\sigma_{MLII}^{(p,k)} + \epsilon}, \quad i \\
 &= 0, 1, \dots, 255 \\
 & \widehat{w_{V1}^{(p,k)}}[i] \\
 &= \frac{w_{V1}^{(p,k)}[i] - \mu_{V1}^{(p,k)}}{\sigma_{V1}^{(p,k)} + \epsilon}, \quad i \\
 &= 0, 1, \dots, 255
 \end{aligned}
 \tag{4}$$

Where $\mu_{lead}^{(p,k)}$ and $\sigma_{lead}^{(p,k)}$ denote the meaning and standard deviation of the windowed signal for the specified lead, computed as:

$$\mu_{lead}^{(p,k)} = \frac{1}{W} \sum_{i=0}^{W-1} w_{la}^{(p,k)}[i] \tag{6}$$

$$\begin{aligned}
 & \sigma_{lead}^{(p,k)} \\
 &= \sqrt{\frac{1}{W} \sum_{i=0}^{W-1} (w_{la}^{(p,k)}[i] - \mu_{lead}^{(p,k)})^2} \tag{7}
 \end{aligned}$$

and $\epsilon = 10^{-6}$ is a small constant to prevent division by zero. The normalization in (4)-(7) forces the model to learn morphological patterns independent of patient-specific amplitude variations and baseline wander.

3.4 Amplitude Scaling Augmentation

To further improve robustness against signal amplitude variations, we apply random amplitude scaling during training:

$$\widehat{w_{MLII}^{(p,k)}}[i] = \alpha_{MLII}^{(k)} \cdot \widehat{w_{MLII}^{(p,k)}}[i] \tag{8}$$

$$\begin{aligned}
 & \widehat{w_{V1}^{(p,k)}}[i] = \alpha_{V1}^{(k)} \\
 & \quad \cdot \widehat{w_{V1}^{(p,k)}}[i] \\
 & \quad \backslash tag9
 \end{aligned}
 \tag{9}$$

where $\alpha_{MLII}^{(k)}, \alpha_{V1}^{(k)} \sim \mathcal{U}(0.85, 1.15)$ are independently sampled scaling factors. Equations (8)-(9) simulate variations in electrode contact quality and signal gain, enhancing model generalization.

3.5 Dual-Lead Input Representation

The preprocessed dual-lead signals are stacked to form a 2-channel input tensor for the k -th beat:

$$X^{(p,k)} = \begin{bmatrix} \widehat{w_{MLII}^{(p,k)}} \\ \widehat{w_{V1}^{(p,k)}} \end{bmatrix} \in \mathbb{R}^{2 \times 1} \tag{1}$$

The dual-channel representation in (10) enables CNN to learn spatial relationships between concurrent leads, leveraging complementary diagnostic information[26].

3.6 Convolutional Neural Network Architecture

Our model $f_{\theta}(\cdot)$ consists of a 1D convolutional neural network parameterized by θ . The network processes the input tensor $X \in \mathbb{R}^{2 \times W}$ through a series of convolutional blocks:

Block 1:

$$\begin{aligned}
 & Z^{(1)} \\
 &= \text{ReLU} \left(\text{BN}(\text{Conv1D}(X, C_1 \right. \\
 &= 64, K_1 = 7, S_1 = 2)) \tag{11}
 \end{aligned}$$

Block 2:

$$\begin{aligned}
 & Z^{(2)} \\
 &= \text{ReLU} \left(\text{BN}(\text{Conv1D}(Z^{(1)}, \right. \\
 &= 128, K_2 = 5, S_2 = 2)) \tag{1}
 \end{aligned}$$

Block 3:

$$\begin{aligned}
 & Z^{(3)} \\
 &= \text{ReLU} \left(\text{BN}(\text{Conv1D}(Z^{(2)}, \right. \\
 &= 256, K_3 = 3, S_3 = 1)) \tag{1}
 \end{aligned}$$

where $\text{Conv1D}(\cdot, C, K, S)$ denotes a 1D convolutional layer with C output channels, kernel size K , and stride S ; $\text{BN}(\cdot)$ represents batch normalization; and $\text{ReLU}(\cdot)$ is the rectified linear unit activation function. The operations hierarchically extract local morphological features at multiple scales. The feature maps are then globally pooled using adaptive average pooling:

$$\begin{aligned}
 & z \\
 &= \text{AdaptiveAvgPool1D}(Z^{(3)}) \\
 &\in \mathbb{R}^{(256)} \tag{1}
 \end{aligned}$$

Finally, the pooled features are classified using a linear layer:

$$p = \text{Softmax}(W_c z + b_c) \tag{15}$$

where $W_c \in \mathbb{R}^{(4 \times 256)}$ and $b_c \in \mathbb{R}^4$ are learnable parameters, and $p \in \mathbb{R}^4$ represents the predicted class probabilities. Equation (15) maps the learned features to the AAMI class probability distribution[27].

3.7 Class Imbalance Handling: Inverse-Square-Root Weighting

Let N_c denote the number of training beats in class $c \in \mathcal{L}_{AAMI}$. The inverse-square-root class weights are computed as:

$$w_c = \left(\frac{1}{\sqrt{N_c + \epsilon}} \right) \cdot \left(\frac{|\mathcal{L}_{AAMI}|}{\sum_{c'=1}^{|\mathcal{L}_{AAMI}|} \left(\frac{1}{\sqrt{N_{c'} + \epsilon}} \right)} \right) \quad (16)$$

where $\epsilon = 10^{-6}$ prevents division by zero for potentially missing classes. The weighting scheme in (16) mitigates extreme class imbalance while preserving gradient flow for majority classes, unlike inverse weighting which can cause training instability.

3.8 Focal Loss for Hard Example Mining

We employ focal loss to focus training on hard-to-classify examples and address the remaining class imbalance:

$$\mathcal{L}_{focal}(p, y) = -w_y \cdot (1 - p_y)^\gamma \cdot \log(p_y) \quad (17)$$

where y is the ground-truth class index, p_y is the predicted probability for class y , w_y is the class weight from (16), and $\gamma = 2.0$ is the focusing parameter that down-weights well-classified examples. The focal loss in (17) emphasizes minority classes and difficult samples, improving sensitivity for rare but clinically critical arrhythmias[28].

3.9 Inter-Patient Validation Protocol

We partition the 48 patients in the MIT-BIH database into three disjoint sets:

$$P_{train} \cap P_{valid} \cap P_{test} = \emptyset \quad (18)$$

where $P_{train} = 31$, $P_{valid} = 5$, and $P_{test} = 5$. The strict patient separation in (18) ensures that

no beat from a test patient appears during training, providing a realistic estimate of clinical deployment performance[29].

3.10 Optimization

The model parameters θ are optimized using the AdamW algorithm to minimize the expected focal loss:

$$\theta^* = \underset{\theta}{\operatorname{argmin}} E_{(x,y) \sim \mathbb{D}_{train}} [\mathcal{L}_{focal}(\theta)] \quad (19)$$

where $\mathbb{D}_{train}$ denotes the training distribution over patient-disjoint beats. Equation (19) formalizes the empirical risk minimization objective with patient-invariant generalization as the primary goal.

4. RESULTS AND DISCUSSION

This segment describes the experimental validation of our dual-lead CNN system to inter-patient ECG atrial fibrillation classification, and how normalization per-beat and physiological augmentation of our CNNs provides strong generalization to unseen patients and superior performance compared to intricate hybrid structures tested in intra-patient settings.

Table 1: Arrhythmia classes distribution in the MIT-BIH database according to training, validation and test datasets.

Spl	Patient	Total Beat	N (Normal)	S (Supraventricular)	V (Ventricular)	F (Fusion)
Train	32	73934	66614 (90.1%)	1221 (1.7%)	5681 (7.7%)	418 (0.6%)
Val	5	9854	8204 (83.3%)	1385 (14.1%)	262 (2.7%)	3 (0.0%)
Test	5	9849	8127 (82.5%)	108 (1.1%)	1238 (12.6%)	376 (3.8%)

The experimental structure is based on the standard AAMI mapping of the MIT-BIH database that has four clinically significant classes: Normal (N), Supraventricular (S)[30], Ventricular (V) and Fusion (F) beats. The statistics of the datasets provided in Table 1 indicate that the distribution of classes in all splits is highly imbalanced as the data in the validation as well as the test sets are dominated by their normal counterparts (more than 80 percent). Most importantly, Ventricular (12.6%) and Fusion (3.8%) beats are found in a higher percentage in the test set than the training set (7.7 and 0.6), respectively, a more realistic representation of the

clinical representation of arrhythmias in real-life conditions.

Table 2: Comparative results of the intra-patient and inter-patient evaluation paradigm of arrhythmia classification

Evaluation Paradigm	Patients Used	Macro F1	V-beat F1	F-beat F1	Patient Leakage Risk
Intra-patient	Train/test on same patient (101)	1.000	0.000	0.000	HIGH
Inter-patient (baseline)	Train 31 patients → test 5 unseen	0.680	0.900	0.592	NONE
Inter-patient (+ augmentation)	Same as above + R-peak jitter	0.690	0.902	0.575	NONE

The most severe conclusion of our analysis is the high-performance difference between intra-patient and inter-patient evaluation paradigms. As Table 2 shows, the intra-patient approach has perfect accuracy (Macro F1 = 1.000), but catastrophically on life-threatening arrhythmias (V-beat F1 = 0.000, F-beat F1 = 0.000) show severe patient-specific leakage of features. Conversely, our inter-patient strategy is very sensitive and shows high sensitivity to the clinically important arrhythmias (V-beat F1 = 0.900-0.902), thus supporting the actual generalization across invisible patients[31].

Table 3: The performance metrics in per-class for the proposed model on test set using clinical significance indicators.

Class	Precision	Recall	F1-Score	Support (# beats)	Clinical Significance
N (Normal)	0.981	0.973	0.977	8127	Represents normal cardiac rhythm; accurate detection is crucial for reliable baseline assessment and minimizing false alarms.
S (Supraventricular)	0.347	0.537	0.422	108	Indicates supraventricular arrhythmias; clinically relevant for early identification of atrial abnormalities.
V (Ventricular)	0.845	0.997	0.915	1238	High-risk ventricular beats associated with life-threatening arrhythmias; requires high sensitivity for patient safety.
F (Fusion)	0.963	0.418	0.583	376	Reflects fusion of normal and ectopic beats; important for accurate rhythm interpretation in complex ECG patterns.

The details on per-class metrics under Table 3 show clinical utility of the model. Of particular importance is the high recall (99.7) and F1-score (91.5) of Ventricular (V) beats which are




Figure 2: Normalized confusion matrix of classification of unseen patients, among all arrhythmia classes.

A normalized confusion matrix presented in figure 2 will represent the fine level classification performance on unknown patients. The model exhibits great discrimination on the Normal (N) and Ventricular (V) beats (97.3% and 99.7% recall respectively) and displays higher confusion between normal and Supraventricular (S) beats. The confusion matrix validates that our model is highly sensitive to life-threatening Ventricular beats[32] which are critical to clinical utility.

potentially life-threatening arrhythmias that must be addressed in emergency. Although the performance of Fusion (F) is worse (41.8%), it is significantly higher compared to most of the other existing methods in literature[33].




Figure 3: t-SNE representation indicating the effect of per-beat normalization on the distribution of features space in various patients.

Per-beat normalization is a crucial factor in model generalization, as shown in figure 3. In the absence of normalization (left panel), embeddings of the various patients occupy utterly distinct territories in feature space and thus cross-patient generalization is not possible. Using normalization of beats per beat (right panel) the embeddings of a dissimilar patient mix, and it is verified that the modelling is trained to learn patient-invariant morphological issues and not patient identity organized by amplitude changes.




Figure 4: The performance of baseline and augmented models on validation using baseline and augmented models across training epochs.

Additional evidence of the effectiveness of our physiological jitter augmentation technique is the learning curves in Figure 4. Augmented (orange line) model exhibits a more stable validation performance during training with a lower variance of macro F1 scores than the baseline (blue line). The stabilization of this algorithm indicates that the jitter augmentation is very effective in simulating annotation noise as seen in the real world without artificial artifacts.




Figure 5: The graph illustrates the process of physiological jitter augmentation method that mimics anisophoresis noise in physiological parameters under a real-world R-peak annotation

Figure 5 visually explains our method of physiological jitter augmentation, by demonstrating how the $\pm$ sample displacements about the R-peak location emulate noise in the annotation of the real world. This slight perturbation does not affect the physiological integrity of this waveform but does not overfit the waveform to the exact locations of R-peaks[36].

Table 4: Ablation study findings in which each component contribution is present in the framework.

Configuration	Per-Beat Normalization	Class Weighting	R-Peak Jitter	Macro F1	Δ vs Baseline
Baseline (no fixes)	No	No	No	0.461	—
Normalization	Yes	No	No	0.670	+0.210
Weighting	Yes	Yes	No	0.655	+0.194
Augmentation (ours)	Yes	Yes	Yes	0.636	+0.175

Table 4 shows an ablation study that quantifies the contribution of every component in our framework. The biggest improvement is seen in per-beat normalization (+0.210 in macro F1), which proves extremely important when it comes to the removal of patient-specific feature leakage. The inverse-square-root class weighting (0.194) and R-peak jitter augmentation (0.175) also play a significant role in ensuring good results, and normalization is the supporting element of the technique without which the rest of the methods can effectively solve the inter-patient generalization challenge.

Table 5: A comparison between our results and the state-of-the-art approaches and a gap in the evaluation paradigm is observed.

Study	Architecture	Evaluation	Metric	Value	V-beat F1	F-beat F1
Rajpurkar et al. (2017)	34-layer CNN	Intra-patient	Accuracy	93.0%	-	-
Hannun et al. (2019)	WaveNet	Intra-patient	F1	0.83	0.80	0.32
Ours (baseline)	Simple CNN	Inter-patient	Macro F1	0.656	0.889	0.394
Ours (+ augmentation)	Simple CNN	Inter-patient	Macro F1	0.714	0.901	0.521

Table 5 places our results in the bigger context of literature, thus revealing the critical evaluation paradigm gap. However, although the intra-patient measures are impressive (Hannun et al. 2019: F1 = 0.83), they do not imply inter-patient generalization. We model a macro F1 of 0.714 with a high sensitivity of life-threatening arrhythmias (V-beat F1 = 0.901, F-beat F1 = 0.521) under stringent conditions of inter-patient evaluation and this sets a new standard of evaluation that is clinically significant.

4.1 Discussion

We have demonstrated that the existing ECG arrhythmia research is not constrained by the complexity of architecture but by the way it is assessed. Intra-patient paradigms are inadvertently leaked in 86.7% of studies (Xiao et al. [34]) and

inflates accuracy measures up to >99% but catastrophically when applied to unseen patients. Such a bias has given rise to more advanced hybrid models like CNN-LSTM-Transformer ensembles and multi-head attention mechanisms which are optimized on intra-patient assessments but do not perform well on cross-patient resilience.

On the other hand, our simple two-lead CNN gives a macro F1 score of 0.714 on completely unknown patients. Although previous research finds impressive intra-patient performance (Hannun et al. : F1 = 0.83; Shah et al. : accuracy = 98.6%; Xiong et al. [35]: F1 = 0.9711), such a value is meaningless in clinical practice since it depends on splits, which reveal patient identity. The results of ours have lower absolute numbers but greater clinical validity since they are results on unseen patients. This gap we attempt to fill by imposing strict patient-disjoint validation to expose the generalization capacity which is diluted in most literature.

Of great interest is that the sensitivity of the Ventricular beats (99.7% recall, 91.5% F1) is high and they are life threatening arrhythmias, which demand immediate clinical attention. This performance matches the intra-patient performance of complex architectures. As an illustration, an 98.01% F1 was achieved on a bidirectional transformer by El-Ghaish and Eldele, though their evaluation paradigm is not clearly revealed, and these results are not consistent in inter-patient settings. It is an F1 of 0.915 on unseen patients, which can be used to show that learning features that are not patient-dependent through per-beat normalization can achieve clinically significant performance even without architectural complexity. This finding is in line with the groundbreaking work of de Chazal et al., who emphasized the vitality of morphological rather than patient-specific characteristics.

Figure 3 (t-SNE)[36] concludes that our normalization technique per-beat is effective, to eliminate patient-specific feature leakage. Unless normalized, the embeddings of the patients in different patients are in completely different clusters, confirming that raw amplitude and baseline features are markers of patient identity, but not arrhythmia[37]. This finding is why complex architectures, including HCTG-Net and ECG-TransCovNet, report nearly intra-patient accuracy but purport to achieve accuracy scores of over 98 percent, but probably do not learn morphological patterns but instead patient-specific amplitude

fingerprints. The generalization ability of these complex models is better than that of our simple CNN in the strict inter-patient conditions. Compared to domain adaptation techniques like those described by Wang et al., who had to use complex unsupervised adaptation modules, our simple normalization method can be easily used to achieve similar generalization at very low computational cost, meaning that many of the solutions to the inter-patient problem suggested by various researchers are overly complex.

Even though Fusion beat performance (41.8% recall, 58.3% F1) is relatively low, it is still much higher than the unpublished experimental results of rare arrhythmias. Most of the existing techniques ignore infrequent classes or apply SMOTE oversampling to form biologically plausible samples. We obtain a Fusion beat F1 of 0.583 on hard-to-classify examples without using synthetic data. This confirms that loss function design is better than data fabrication in rare classes, which has significant improvement in this category of data that is highly required. This demonstration allows us to point out that the right class weighting and loss functionality design may give significant enhancement in detecting rare arrhythmia [38] without modifying the architectures. Above all, inter-patient validation should become the new paradigm in the field, and the emphasis should be shifted into the standards of evaluation and patient-independent learning rather than on the newness of architectural designs.

5. CONCLUSION

The paper describes a lightweight dual-lead CNN architecture of inter-patient ECG arrhythmia classification which puts clinical generalizability over benchmark optimization. The proposed model had a macro F1-score of 0.714 on completely unseen patients and outstanding performance on life-threatening Ventricular ectopic beat (99.7 percent recall, 91.5 percent F1) and a substantially greater performance on Fusion beat detection (58.3 percent F1) than baseline methods. Ablation experiment and t-SNE visualization established that normalization of per-beat z-score is the basic element of the component (+0.210 macro F1 gain) which can successfully remove the patient-specific leakage of amplitude and can learn morphology-invariant features. These findings indicate that cross-patient arrhythmia detection clinical utility is mainly dependent on evaluation methodology and feature normalization, as opposed to architectural complexity.

The work in question contributes to the field in three important aspects: (1) patient-disjoint validation has been established as a required standard of clinically meaningful performance comparison; (2) a lightweight per-beat normalization strategy has been introduced, which allows models to learn patient-invariant morphological patterns; and (3) an inter-patient benchmark has been provided at all AAMI classes that has been obtained through explicit metrics of clinical relevance.

5.1 Limitations and Future Directions.

Although this study provides a high rigorous standard of inter-patient performance, there are a number of limitations and open questions that should be studied further: Multi-Center Generalizability The present validation is restricted to the 48 participants of the MIT-BIH Arrhythmia Database.

The current validation is limited to the MIT-BIH Arrhythmia Database, and future work should evaluate generalizability across larger multi-center datasets such as PTB-XL and CPSC.

The test set contains a higher proportion of Ventricular (12.6%) and Fusion (3.8%) beats than the training set (7.7% and 0.6%), which may affect model calibration under varying real-world prevalence.

This study focuses on beat-level classification, and incorporating rhythm-level context through sequential modeling could improve detection of temporally dependent arrhythmias.

The model is evaluated on high-quality dual-lead clinical recordings, and its robustness on noisy single-lead wearable data remains to be validated.

Incorporating attention mechanisms or gradient-based saliency methods could improve interpretability and increase clinical trust.

Above all, we propose inter-patient validation as the default assessment paradigm in the field of ECG arrhythmia research. To make the progress of deep learning advances into clinically reliable tools, a change in focus is necessary towards the architectural novelty to rigorous, patient-independent evaluation. It is only possible to resolve these open questions and embrace clinically justified validation procedures to shift the field to benchmark optimization on its way to reliable, real-world arrhythmia detection.

REFERENCES:

- [1] Abousaber, I., El-Ghaish, H., & Abdallah, H. F. (2025). A novel explainable ECG classification with spatio-temporal transformers and hybrid loss optimization. *Biomedical Signal*

- Processing and Control, 110, 108142. <https://doi.org/10.1016/j.bspc.2025.108142>
- [2]. Al-Tam, R. M., Al-Hejri, A. M., Naji, E., Hashim, F. A., Alshamrani, S. S., Alshehri, A., & Narangale, S. M. (2024). A Hybrid Framework of Transformer Encoder and Residential Conventional for Cardiovascular Disease Recognition Using Heart Sounds. *IEEE Access*, 12, 123099–123113. <https://doi.org/10.1109/ACCESS.2024.3451660>
- [3]. Anjum, N., Sathi, K. A., Hossain, M. A., & Dewan, M. A. A. (2023). A Temporal Transformer-Based Fusion Framework for Morphological Arrhythmia Classification. *Computers* 2023, Vol. 12, Page 68, 12(3), 68. <https://doi.org/10.3390/computers12030068>
- [4]. Chen, Y., Qiu, S., Wang, Z., & Zhao, H. (2023). IDCNN-TRSNet: A Hybrid End-to-End Arrhythmia Classification Deep Network Based on Transformer. 2023 5th International Conference on Data-Driven Optimization of Complex Systems, DOCS 2023. <https://doi.org/10.1109/DOCS60977.2023.10294360>
- [5]. Chopannejad, S., Roshanpoor, A., & Sadoughi, F. (2024). Attention-assisted hybrid CNN-BILSTM-BiGRU model with SMOTE-Tomek method to detect cardiac arrhythmia based on 12-lead electrocardiogram signals. *Digital Health*, 10. <https://doi.org/10.1177/20552076241234624>
- [6]. de Chazal, P., O'Dwyer, M., & Reilly, R. B. (2004). Automatic classification of heartbeats using ECG morphology and heartbeat interval features. *IEEE Transactions on Bio-Medical Engineering*, 51(7), 1196–1206. <https://doi.org/10.1109/TBME.2004.827359>
- [7]. Din, S., Qaraqe, M., Mourad, O., Qaraqe, K., & Serpedin, E. (2024). ECG-based cardiac arrhythmias detection through ensemble learning and fusion of deep spatial-temporal and long-range dependency features. *Artificial Intelligence in Medicine*, 150(8), 102818. <https://doi.org/10.1016/j.artmed.2024.102818>
- [8]. Ding, S. Z., Du, Y. Q., Lu, H., Wang, W. H., Cheng, H., Chen, S. Y., Chen, M. H., Chen, W. C., Chen, Y., Fang, J. Y., Gao, H. J., Guo, M. Z., Han, Y., Hou, X. H., Hu, F. L., Jiang, B., Jiang, H. X., Lan, C. H., Li, J. N., ... Li, Z. S. (2022). Chinese Consensus Report on Family-Based Helicobacter pylori Infection Control and Management (2021 Edition). *Gut*, 71(2), 238–253. <https://doi.org/10.1136/gutjnl-2021-325630>
- [9]. Du, C., Liu, P. X., & Zheng, M. (2022). Classification of Imbalanced Electrocardiosignal Data using Convolutional Neural Network. *Computer Methods and Programs in Biomedicine*, 214, 106483. <https://doi.org/10.1016/j.cmpb.2021.106483>
- [10]. El-Ghaish, H., & Eldele, E. (2024). ECGTransForm: Empowering adaptive ECG arrhythmia classification framework with bidirectional transformer. *Biomedical Signal Processing and Control*, 89, 105714. <https://doi.org/10.1016/j.bspc.2023.105714>
- [11]. Guo, X., Li, J., & Feng, L. (2025). Arrhythmia Classification Method Based on CNN-Attention-BiTransformer. *IJLAI Transactions on Science and Engineering*, 3(4), 43–53. <https://sub.ifspress.hk/IJLAI/article/view/192>
- [12]. Hannun, A. Y., Rajpurkar, P., Haghpanahi, M., Tison, G. H., Bourn, C., Turakhia, M. P., & Ng, A. Y. (2019). Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network. *Nature Medicine*, 25(1), 65–69. <https://doi.org/10.1038/s41591-018-0268-3>
- [13]. Hasan, M. S., & Aowlad Hossain, A. B. M. (2025). Multiclass Arrhythmia Classification from Imbalanced ECG Data Using Encoded Transformer based CNN-LSTM Hybrid Model. *International Journal of Intelligent Systems and Applications*, 17(5), 68–83. <https://doi.org/10.5815/ijisa.2025.05.05>
- [14]. Hong, S., Zhou, Y., Shang, J., Xiao, C., & Sun, J. (2020). Opportunities and challenges of deep learning methods for electrocardiogram data: A systematic review. *Computers in Biology and Medicine*, 122, 103801. <https://doi.org/10.1016/j.compbiomed.2020.103801>
- [15]. Kerdoudi, S., Guezouli, L., & Dilekh, T. (2026). Enhanced ECG arrhythmia detection with deep learning and multi-head attention mechanism. *Computers and Electrical Engineering*, 131, 110957. <https://doi.org/10.1016/j.compeleceng.2026.110957>
- [16]. Kim, D., Lee, K. R., Lim, D. S., Lee, K. H., Lee, J. S., Kim, D. Y., & Sohn, C. B. (2025). A novel hybrid CNN-transformer model for arrhythmia detection without R-peak identification using stockwell transform. *Scientific Reports* 2025 15:1, 15(1), 7817-. <https://doi.org/10.1038/s41598-025-92582-9>

- [17]. Kiranyaz, S., Ince, T., & Gabbouj, M. (2016). Real-Time Patient-Specific ECG Classification by 1-D Convolutional Neural Networks. *IEEE Transactions on Bio-Medical Engineering*, 63(3), 664–675. <https://doi.org/10.1109/TBME.2015.2468589>
- [18]. Liu, H. (2025). A hybrid model combining 1D-CNN and BERT for intelligent ECG arrhythmia classification. *Scientific Reports*, 15(1). <https://doi.org/10.1038/s41598-025-28023-4>
- [19]. Liu, X., Wang, H., Li, Z., & Qin, L. (2021). Deep learning in ECG diagnosis: A review. *Knowledge-Based Systems*, 227, 107187. <https://doi.org/10.1016/j.knosys.2021.107187>
- [20]. Mahesh, G., & Anitha, K. (2025). Hybrid CNN-BiLSTM with Attention Mechanism for Automated Diagnosis of Cardiovascular Disorders Using ECG Images. 2025 International Conference on Next Generation Computing Systems: Intelligent System for Sustainable Development, ICNGCS 2025 - Conference Proceedings. <https://doi.org/10.1109/ICNGCS64900.2025.11183520>
- [21]. Manivannan, G. S., Talawar, S. v., Vasundhara, M. G., Karthik Lal, A., Gagana, S. M., Hemanth, S. K., & Srijan Gowda, H. L. (2026). An efficient deep CNN based BiLSTM framework with RanA optimization for accurate cardiac arrhythmia detection. *Scientific Reports* 2026. <https://doi.org/10.1038/s41598-026-38227-x>
- [22]. Moody, G. B., & Mark, R. G. (2001). The impact of the MIT-BIH arrhythmia database. *IEEE Engineering in Medicine and Biology Magazine: The Quarterly Magazine of the Engineering in Medicine & Biology Society*, 20(3), 45–50. <https://doi.org/10.1109/51.932724>
- [23]. Mousavi, S., Afghah, F., & Rajendra Acharya, U. (2019). SleepEEGNet: Automated sleep stage scoring with sequence to sequence deep learning approach. *PLOS ONE*, 14(5), e0216456. <https://doi.org/10.1371/JOURNAL.PONE.0216456>
- [24]. Oh, S. L., Ng, E. Y. K., Tan, R. S., & Acharya, U. R. (2018). Automated diagnosis of arrhythmia using combination of CNN and LSTM techniques with variable length heart beats. *Computers in Biology and Medicine*, 102, 278–287. <https://doi.org/10.1016/j.compbiomed.2018.06.002>
- [25]. Petmezas, G., Stefanopoulos, L., Kilintzis, V., Tzavelis, A., Rogers, J. A., Katsaggelos, A. K., & Maglaveras, N. (2022). State-of-the-Art Deep Learning Methods on Electrocardiogram Data: Systematic Review. *JMIR Medical Informatics*, 10(8), e38454. <https://doi.org/10.2196/38454>
- [26]. Rahman, S. (2025). Neuro-Informed Adaptive Learning (NIAL) Algorithm: A Hybrid Deep Learning Approach for ECG Signal Classification. <http://arxiv.org/abs/2503.20789>
- [27]. Ribeiro, A. H., Ribeiro, M. H., Paixão, G. M. M., Oliveira, D. M., Gomes, P. R., Canazart, J. A., Ferreira, M. P. S., Andersson, C. R., Macfarlane, P. W., Wagner, M., Schön, T. B., & Ribeiro, A. L. P. (2020). Automatic diagnosis of the 12-lead ECG using a deep neural network. *Nature Communications* 2020 11:1, 11(1), 1760-. <https://doi.org/10.1038/s41467-020-15432-4>
- [28]. Saha Tchinda, B., Tchiotso, D., & Tchinda, R. (2026). Classification of cardiac arrhythmias using ECG shape-adapted wavelets and BiLSTM network. *Measurement*, 257, 118832. <https://doi.org/10.1016/j.measurement.2025.118832>
- [29]. Shah, H. A., Saeed, F., Diyan, M., Almujaal, N. A., & Kang, J. M. (2024). ECG-TransCovNet: A hybrid transformer model for accurate arrhythmia detection using Electrocardiogram signals. *CAAI Transactions on Intelligence Technology*. <https://doi.org/10.1049/cit2.12293>
- [30]. Shaker, A. M., Tantawi, M., Shedeed, H. A., & Tolba, M. F. (2020). Generalization of Convolutional Neural Networks for ECG Classification Using Generative Adversarial Networks. *IEEE Access*, 8, 35592–35605. <https://doi.org/10.1109/ACCESS.2020.2974712>
- [31]. Siontis, K. C., Noseworthy, P. A., Attia, Z. I., & Friedman, P. A. (2021). Artificial intelligence-enhanced electrocardiography in cardiovascular disease management. *Nature Reviews. Cardiology*, 18(7), 465–478. <https://doi.org/10.1038/s41569-020-00503-2>
- [32]. Vishnupriya, B., Vikas, B. G., & Choudhry, M. D. (2025). Hybrid Deep Learning Framework for ECG-Based Arrhythmia Detection Using CNN, Bi-LSTM, and Transformer. 2025 IEEE 4th World Conference on Applied Intelligence and Computing, AIC 2025, 628–632. <https://doi.org/10.1109/AIC66080.2025.11212074>

- [33]. Wang, G., Chen, M., Ding, Z., Li, J., Yang, H., & Zhang, P. (2021). Inter-patient ECG arrhythmia heartbeat classification based on unsupervised domain adaptation. *Neurocomputing*, 454, 339–349. <https://doi.org/10.1016/j.neucom.2021.04.104>
- [34]. Xiao, Q., Lee, K., Mokhtar, S. A., Ismail, I., Pauzi, A. L. bin M., Zhang, Q., & Lim, P. Y. (2023). Deep Learning-Based ECG Arrhythmia Classification: A Systematic Review. *Applied Sciences (Switzerland)*, 13(8), 4964. <https://doi.org/10.3390/app13084964>
- [35]. Xiong, N., Wei, Z., Wang, X., Wang, Y., & Wang, Z. (2025). HCTG-Net: A Hybrid CNN–Transformer Network with Gated Fusion for Automatic ECG Arrhythmia Diagnosis. *Bioengineering* 2025, Vol. 12, Page 1268, 12(11), 1268. <https://doi.org/10.3390/bioengineering12111268>
- [36]. Xiong, Z., Stiles, M. K., Gillis, A. M., & Zhao, J. (2022). Enhancing the detection of atrial fibrillation from wearable sensors with neural style transfer and convolutional recurrent networks. *Computers in Biology and Medicine*, 146, 105551. <https://doi.org/10.1016/j.compbiomed.2022.105551>
- [37]. Yildirim, Ö. (2018). A novel wavelet sequences based on deep bidirectional LSTM network model for ECG signal classification. *Computers in Biology and Medicine*, 96, 189–202. <https://doi.org/10.1016/j.compbiomed.2018.03.016>
- [38]. Zhang, D., Yang, S., Yuan, X., & Zhang, P. (2021). Interpretable deep learning for automatic diagnosis of 12-lead electrocardiogram. *iScience*, 24(4), 102373. <https://doi.org/10.1016/j.isci.2021.102373>