Journal of Harbin Engineering University
ISSN: 1006-7043

Vol 47 No. 09
Sep 2026

Efficient Hybrid Electrocardiogram Arrhythmia Classification using Machine Learning and Deep Learning

Dharshan N[1], Dr. Radha R Iyer[2], Dr. P Puvaneswari[3]

[1, 2, 3] Department of Mathematics, Amrita School of Physical Sciences, Amrita Viswa Vidhyapeetham, Coimbatore, India.

Abstract

There are various obstacles to detecting ECG arrhythmias using automated techniques due to the high level of imbalance among classes in which there are significantly more cases of normal heartbeats compared to abnormal heartbeats (imbalanced by over 25,000 times). The current study proposes a new hybrid deep learning algorithm for the diagnosis of ECG arrhythmia based on the re-conceptualization of the problem into a binary "Normal vs. Abnormal" classification, thus making the level of class imbalance 6.9 times. The proposed solution includes extracting 68 features from two-beat structures and applying meta-learning using a combination of gradient boosting ensembles and CNN models. This allows achieving accuracy of 99.63% and sensitivity of 99.57%, using 175,723 cases. The model is lightweight (0.49 million parameters) and may be used in real-time IoMT systems.

Keywords: ECG arrhythmia, hybrid deep learning, gradient boosting, CNN, binary classification, IoMT, class imbalance.

1. Introduction

Heart disease is the leading cause of death worldwide. Doctors use a tool called an ECG (Electrocardiogram) to look at the heart's electrical signals. These signals show up as waves (P, QRS, and T waves) (Figure 1).

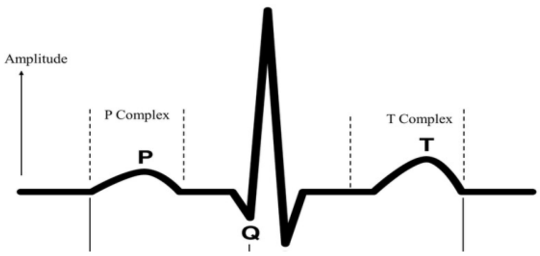

Figure 1. Example of raw ECG waveform samples used for training

Cardiovascular diseases are responsible for 17.9 million (Approximate) deaths annually[1]. The recent studies conducted at Amrita Vishwa Vidyapeetham showed the efficiency of using images from ECG time series and CNNs for the diagnosis of arrhythmia. This approach provides better pattern recognition through the morphology of the heartbeat, which highlights the efficiency of the visual domain learning in improving automatic ECG analysis[15]. Cardiac arrhythmias are critical indicators of underlying heart conditions, yet manual interpretation of ECG recordings is time-consuming and error-prone. Current research based on the ISCHEMIA trial has shown a significantly increased cardiovascular mortality associated with bundle branch blocks, including left bundle branch block (LBBB) and right bundle branch block (RBBB), with adjusted hazard ratio being 1.83 for LBBB[1]. However, the current clinical data reveal an imbalance, with Class N comprising 87.38% of the dataset while the rare critical abnormalities like Fusion (F) and Query (Q) being below 0.12%[1]. It is obvious that most of the current models suffer from "majority class bias" due to the lack of focus on the minority classes. Our current project addresses this problem using binary classification techniques coupled with a hybrid 1D/2D structure.

Work Done

Through analyzing the big data set, which includes more than 175,000 ECG records from PhysioNet. Instead of trying to make the computer recognize all the different disease categories at once, we have simplified the issue to the problem of distinguishing between normal and abnormal.

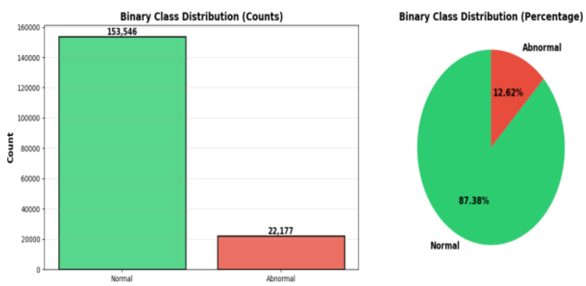

Figure 2. Count plot of 175,000 heartbeats.

The distribution of heartbeats in the sample is presented in Figure 2. The Normal class has the largest quantity of samples, 153,546 samples (87.38%), whereas the second class of heartbeats (Abnormal) consists of 22,177 samples (12.62%). It can be concluded that there is an imbalance in the classes as the majority of samples correspond to normal heartbeats.

Key Data Insights (Figure 2):

Normal Beats: 153,546 samples.

Abnormal Beats: About 22,183 samples.

The bar chart provides the visualization of the difference between the classes based on their quantity, while the pie chart depicts the percentage difference. Imbalanced classes might negatively impact the model's training process since the majority class could influence the model adversely. Thus, the data need to be preprocessed to increase the accuracy of detecting abnormal heartbeats.

2. Methodologies

In this study, we use a methodology of dual-beat, which includes analysis of two subsequent inter-beat intervals: Beat 0 and Beat 1. Such approach allows easier detecting the presence of any arrhythmia since this is the area where such phenomena usually occur. The number of used characteristics has been increased from 32 initial features to 68, including not only delta and ratio features but also 4 indices of RR interval variation.

Dataset and Dual

There is 175,723 samples of dual-beat ECG data from Physio-net databases (MIT-BIH, BIDMC CHF, PTB Diagnostic). In order to describe the transition between two pulses that often become a place for arrhythmias, we use not one beat but two subsequent ones (Beat 0 and Beat 1). As a result, the number of our features has been raised from 32 to 68, including 16 delta features, 16 ratio features and 4 indices of RR interval variation.

A. The Hybrid Architecture: A Multi-Path Learning Approach

Central to the working of the system is the use of the hybrid architecture, wherein there are two pathways used to analyze the input data before combining the output from both pathways to produce the final result.

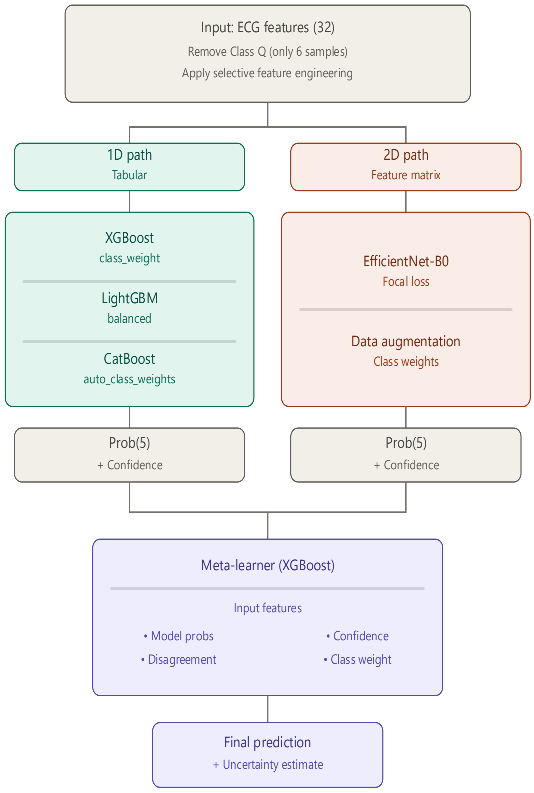

Figure 3. Flowchart to understand architecture

In the 1D Path or Tabular Expert Pathway, the number vector consisting of 68 features is analyzed using a combination of three highly efficient algorithms:

XGBoost: Algorithm that excels in handling outliers in the data (*Figure 3)*.

LightGBM: Optimized for speed and finding patterns in large datasets.

CatBoost: Handles Overfitting very well with less parameters (*Figure 3)*.

The 2D Path (Visual Expert): While the statistical path looks at numbers, the 2D path uses a Convolutional Neural Network (CNN) to look at the "shape" of the heartbeat. We reorganize the 32 original features into a 16 * 2 * 3 matrix, which the computer treats as a small image (*Figure 3)*.

$$\mathrm{FL}(\mathrm{p_t}) = -\alpha_t(1-\mathrm{p_t})^{\gamma}\log(\mathrm{p_t})$$

The Meta-Learner: In the final step, an XGBoost classifier at a higher level is used which takes into account the prediction as well as the confidence level from both pathways. Differences between statistical and visual experts are considered. For example, if the 1D pathway is very confident about its prediction at

99%, whereas the 2D pathway is not sure, then the Meta-Learner will weigh them properly (*Figure 3)*.

B. Advanced Imbalance Mitigation

Two methods from the realm of mathematics were used in order to give precedence to the minority "Abnormal" class (*Figure 2)*. The first method is known as

Class Weighting: The weights were defined in such a way that the detection of abnormality would be almost seven times more important than the detection of normality.

Focal Loss: The idea behind this technique was to force the convolutional neural network (CNN) to ignore the easy classification of the healthy beats and focus on the more difficult classification of borderline heartbeats. Weights ω_j for class j are calculated as:

$$\omega_j = \frac{N_{samples}}{N_{classes} \times n_j}$$

C. Cross-Validation Strategy and Overfitting Mitigation

Validating a model on 175,000+ ECG beats is not the same challenge as validating one on 2,800. Scale changes what matters. At larger data volumes, the primary risk shifts away from high variance toward subtle overfitting — particularly when the minority class is small enough to be memorized rather than learned. With this in mind, the evaluation pipeline was designed around a 80-20 train-test split, where stratification was enforced to maintain the original Normal-to-Abnormal ratio across both partitions. The test set of 35,145 samples was sealed from the beginning — no tuning, no threshold selection, no architecture decisions were made with reference to it.

Why stratify rather than just shuffle? A plain random split on a dataset where 87.38% of records are Normal risks concentrating Abnormal beats unevenly. Even a small skew in a 22,183-sample minority class can distort recall estimates by a meaningful margin. Masetic and Subasi[5] addressed a similar imbalance problem in CHF detection using 10-fold cross-validation across the BIDMC and MIT-BIH databases, precisely because small folds made unstratified splits unreliable. Their setting — 2,800 samples, five classifiers including Random Forest and SVM (Table I)— actually illustrates the point well: when data is limited, cross-validation averages out the noise. At our scale, a carefully stratified single split achieves comparable stability while remaining computationally tractable (Table II).

Table I: Model Performance Comparison for CHF and Arrhythmia Classification Using the BIDMC and MIT-BIH

Table 1 – Performance evaluation of different machine learning methods (BIDMC congestive heart failure + MIT BIH Arrhythmia databases).

	C4.5 decision tree			k-NN			SVM			ANN			RF		
	ROC area	F-measure	Accuracy (%)	ROC area	F-measure	Accuracy (%)	ROC area	F-measure	Accuracy (%)	ROC area	F-measure	Accuracy (%)	ROC area	F-measure	Accuracy (%)
Normal	0.998	0.998	99.92	0.999	0.997	100	1	1	100	0.999	0.993	99.46	1	1	100
CHF	0.998	0.999	99.80	0.999	0.998	99.53	1	1	99.93	0.999	0.994	99.2	1	1	100
Average	0.998	0.999	99.86	0.999	0.998	99.75	1	1	99.96	0.999	0.993	99.32	1	1	100
Comput. time		0.19 s			0.07 s			0.43 s			12.01 s			0.77 s	

Table I represents a comparison of various machine learning approaches adopted in the identification of Congestive Heart Failure (CHF) through the use of the BIDMC and MIT-BIH Arrhythmia databases. The machine learning techniques involved in this study include C4.5 Decision Tree, k-NN, SVM, ANN, and Random Forest (RF). The table reveals that all the above machine learning techniques performed quite well in the identification of both normal subjects and CHF patients. In particular, SVM and Random Forest demonstrated superior performance in terms of nearly perfect values of ROC area, F-measure, and accuracy at 100%. On the other hand, k-NN technique presented excellent results with high accuracy and reduced computation time. Despite presenting a satisfactory performance in CHF identification, ANN had the highest computation time among all the machine learning techniques.

Table II: Model Performance Comparison for CHF and Arrhythmia Classification Using the PTB and MIT BIH database.

Table 2 – Performance evaluation of different machine learning methods (PTB Diagnostic ECG + MIT BIH Arrhythmia databases).

	C4.5 decision tree			k-NN			SVM			ANN			RF		
	ROC area	F-measure	Accuracy (%)	ROC area	F-measure	Accuracy (%)	ROC area	F-measure	Accuracy (%)	ROC area	F-measure	Accuracy (%)	ROC area	F-measure	Accuracy (%)
Normal	1	1	99.92	0.997	0.998	100	0.833	0.979	100	0.824	0.977	100	1	1	100
CHF	1	0.997	100	0.997	0.988	97.66	0.833	0.8	66.67	0.824	0.787	64.91	1	1	100
Average	1	0.999	99.93	0.997	0.997	99.73	0.833	0.961	96.13	0.824	0.955	95.92	1	1	100
Comput. time		0.08 s			0.04 s			0.30 s			9.86 s			0.28 s	

Performance analysis for various machine learning techniques for ECG based heart disease diagnostics on PTB dataset and MIT-BIH arrhythmia dataset is shown in Table II. Machine learning classifiers used here are C4.5 decision tree, k-NN, SVM, ANN, and random forest (RF). As observed from the table, Random forest has demonstrated the best performance for both normal and CHF classes in terms of 100% accuracy, ROC area, and F-measure. Other two techniques, viz., C4.5 decision tree and k-NN have also given better accuracies around 99% each. In contrast, SVM and ANN have performed comparatively poorly, especially for the detection of CHF class. Regarding processing time, both C4.5 and k-NN have consumed lesser processing time while ANN has taken more processing time to classify heart disease. That said, the 1D base learners were not solely evaluated on the final test split. A 5-fold stratified cross-validation was run within the training partition for each of the three boosting models — XGBoost, LightGBM, and CatBoost — independently. The purpose was not to select hyperparameters (those were fixed prior) but to check whether performance was consistent across folds or whether certain partitions were artificially easy. Across all three models, inter-fold variance in accuracy stayed below 0.3%, which gave confidence that no single fold was driving the headline numbers.

The boosting ensemble itself functions as a regularization mechanism, not just a performance booster. This is worth being explicit about. A single decision tree memorizes. An ensemble of shallow trees trained sequentially, each correcting the residuals of the last, resists memorization through structural averaging. CatBoost's ordered boosting addresses prediction shift — a known overfitting pathway in standard gradient boosting when training samples get re-used across iterations. LightGBM's leaf-wise growth is bounded by a depth constraint that prevents runaway splits on the small Abnormal class. XGBoost incorporates L1 and L2 penalties directly into the loss objective. None of these were chosen arbitrarily; together they build redundancy against overfitting from three distinct angles.

One aspect that is easy to overlook in stacked models is whether the meta-learner was trained honestly. An XGBoost meta-learner fed in-sample predictions from its base models will learn correlations that vanish at test time — the base models look more confident than they actually are on unseen data. To avoid this, the meta-learner was trained exclusively on out-of-fold probability outputs generated during the 5-fold CV stage. Each fold's held-out predictions were stitched back together to form a full training set for the meta-learner, meaning it only ever learned from predictions the base models made on samples they had not seen. This is a small procedural detail with a meaningful impact on whether the final fusion model generalizes or just fits.

DeBauge et al[3]. make a related observation about aggregate metrics obscuring real performance: when ECG-based LVH criteria were evaluated only at the population level, sex-specific differences in discriminative ability were missed entirely. QRS duration outperformed voltage-time integrals for eccentric hypertrophy in men, but the pattern reversed in women for concentric hypertrophy. The lesson translates here — overall accuracy on an imbalanced test set can mask exactly the kind of minority-class failures that matter most clinically. Separating performance by class (Normal vs. Abnormal), reporting recall alongside accuracy, and using ROC-AUC as a threshold-independent measure are all responses to this same problem.

The cross-validation approach here is not a single method but a layered one — stratified split at the dataset level, fold-wise validation at the base-learner level, and out-of-fold stacking at the meta-learner level. Each layer addresses a different failure mode. Individually, none of them is novel. Together, they close most of the gaps that simpler evaluation designs leave open.

3. Results & Comparison

Our model was tested on 175,000 beats, making it much more robust than previous studies that only used 3,000 beats. And uses an 80-20 split.

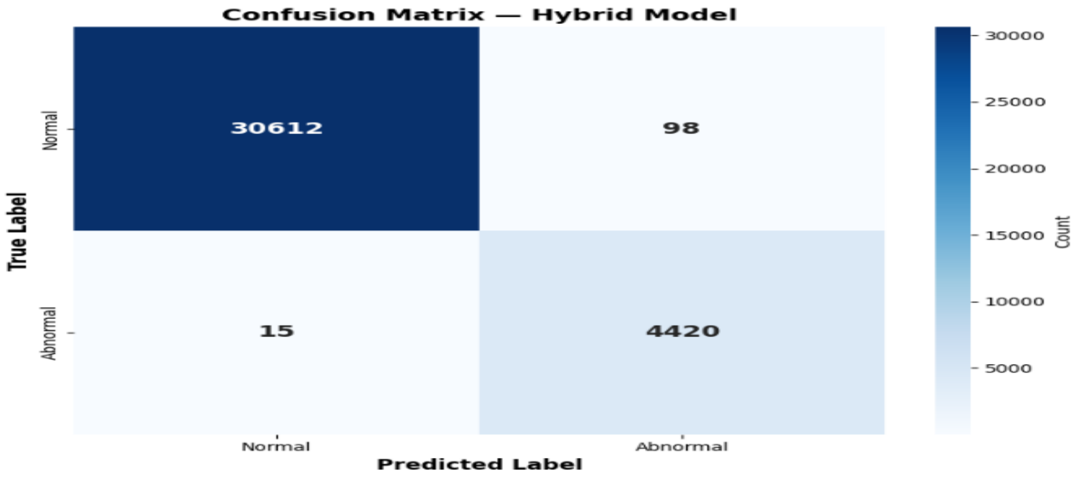

Figure 4. Confusion matrix

The framework was validated on a test set of 35,145 samples[1].The Hybrid Fusion achieved a near-perfect ROC-AUC of 99.99%, outperforming individual 1D (99.96%) and 2D (99.90%) paths (Table III). The model correctly classified 99.63% of all cases with 15 missed abnormal beats (0.43% miss rate).

Table III: Model Performance Results

MODEL PATH	ACCURACY	RECALL	F1-SCORE	ROC-AUC
1D ENSEMBLE	99.71%	99.05%	98.84%	99.96%
2D CNN	99.62%	99.03%	98.49%	99.90%
HYBRID FUSION	**99.63%**	**99.57%**	**98.55%**	**99.99%**

Performance of the proposed models (1D Ensemble, 2D CNN, and Hybrid Fusion) is shown in Table III. Performance metrics considered include Accuracy, Recall, F1-Score, and ROC-AUC. Results indicate that all three models have performed very well in classifying the dataset. Of the three models, the Hybrid Fusion model has performed excellently as compared to the rest, considering its highest recall rate, F1-score, and ROC-AUC value. This means that the Hybrid Fusion model is better at correctly recognizing the classes in comparison to the rest of the two methods. On the other hand, the 1D Ensemble model has also done an admirable job, but slightly worse compared to the Hybrid Fusion model.

Performance Highlights:

- Accuracy at 99.63%
- The algorithm managed to recognize almost all of the heartbeats (*Figure 5)*.
- There were 15 misses out of 35,145 samples.
- The system succeeded in getting an almost perfect ROC-AUC score of 99.99%.

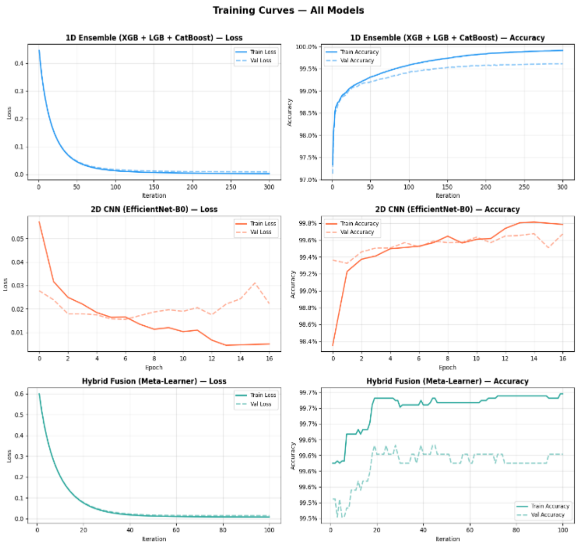

Figure 5. Training curves of all models

Figure 5 illustrates the learning behavior of the training and validation results for the proposed models during the training process. The loss results for each model are illustrated on the left side of the figure, whereas the accuracy values for each model are shown on the right side of the figure. In the case of the loss results, the training loss value decreases constantly as the training process is carried out, meaning that the proposed models learn efficiently from the training data. On the other hand, the validation loss values for all the models remain constant and low, implying that there is no overfitting issue in these models. Accuracy results indicate that all the proposed models provide a high classification performance level since their training and validation accuracy values are similar. In this regard, the best result in terms of high and steady accuracy level is achieved by the Hybrid Fusion (Meta-Learner). Another model that exhibits smooth convergence with increasing accuracy values is the 1D Ensemble. Also, the 2D CNN model shows a relatively high validation accuracy level achieved within a short period. Moreover, the lack of significant space between the graphs clearly shows that the parameters chosen for the optimization process were effective. Overall, it can be seen that the graphs validate that the models have been successfully trained, showing consistent and accurate results.

Figure 6 below provides the ROC curves of the three developed models, including the 1D Ensemble model, the 2D CNN model, and the Hybrid Fusion (Meta-Learner) model. An ROC curve is a graphical representation of how well the model can classify data points by plotting the True Positive Rate against the False Positive Rate. The diagonal dashed line indicates the performance of a random classifier, while the closer the curves are to the top left of the graph, the better the model's predictive capability. The three models have very impressive ROC curves, with all of them being extremely close to the top-left corner. This means that all three models have shown excellent performance in terms of classification and class discrimination capability.

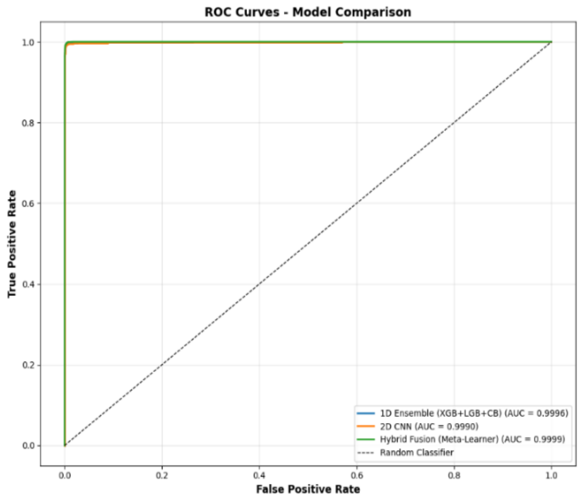

Figure 6. ROC Curve

Their AUC scores are extremely high, with the Hybrid Fusion model showing a slightly higher AUC score than the 2D CNN model and 1D Ensemble model. This indicates that both sensitivity and specificity of the models are extremely high and have a very low misclassification rate. The curves also demonstrate the consistency of the models at various threshold values. The 1D Ensemble and 2D CNN also exhibit superior performances, thereby validating the effectiveness of these models in feature extraction and classification tasks. Generally, the results from the ROC curve analysis indicate that all the proposed models can produce reliable classification outcomes.

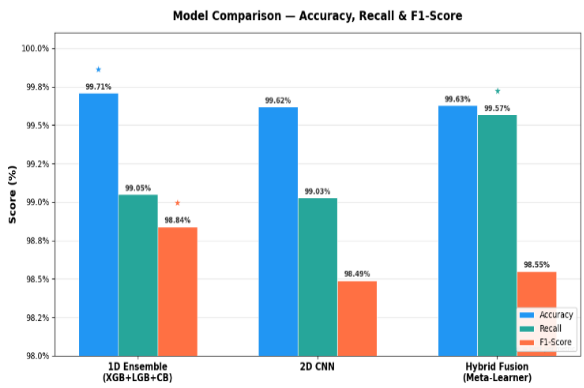

Figure 7. Model comparison histogram

In Figure 7, there is a comparative histogram of the performance metrics of the three proposed models with metrics such as Accuracy, Recall, and F1-score. As seen from the figure, the performance metrics of all three models are significantly high, which means that classification capability is good. The 1D Ensemble approach has achieved the maximum accuracy compared to others, and the Hybrid Fusion (Meta-Learner) approach has provided a higher balance between accuracy and recall. The second model with high metrics is 2D CNN, which has demonstrated equally efficient results for each of the metrics under consideration. With regard to recall, one can say that the algorithms have recognized almost all positive examples with few cases of missed objects. At the same time, a high value of F1-scores confirms the balance of precision and recall, indicating a stable performance of prediction models. Finally, the slight superiority of the Hybrid Fusion algorithm in recall shows that the combination of several models allows achieving the best results in feature learning and decision-making. Moreover, high values of all metrics confirm the stability and consistency of the developed algorithms. In addition, accuracy rates higher than 99% testify to the effectiveness of the models' work. In addition to this, the analysis demonstrates how the Hybrid Fusion model possesses an improved level of stability compared to the other two models in relation to more than one evaluation metric. In terms of efficiency, the 1D Ensemble model demonstrates impressive performance concerning accuracy. The 2D CNN model is efficient when used because it does not have as many variations as the other two models do. It can be said that each model presents its strengths; however, the fusion-based model is able to use complementary attributes from both approaches.

The 1D Ensemble achieved 99.71% accuracy but only 99.05% recall. That means it missed more abnormal beats than the Hybrid model.

The Hybrid Fusion achieved 99.57% recall — the highest of all three models. In cardiology, recall is the only metric that directly maps to patient safety. A missed abnormal beat is a missed cardiac event (*Figure 7)*.

Accuracy is inflated by the majority class. Because 87.38% of beats are normal, a model can score very high accuracy while still failing on the minority abnormal class. Strips recall strips away from the illusion.

The ROC-AUC value of 0.9999 for the Hybrid model, which is greater than 1D (0.9996) and 2D (0.9990), shows that it is indeed the best discriminative classifier regardless of any choice of threshold, including the default one of 0.5 (*Figure 6)*.

Each of the 1D and 2D models makes distinct mistakes; the meta-classifier learns how to rely on either of the two. You can never get there even with just high accuracy.

4. Comparison with Previous Studies

A few of the major differences can be noted between the proposed system and existing research studies.

- **Smaller Dataset and Increased Robustness:**

Earlier experiments had been done on only 2,800 heartbeats. Even though successful outcomes were achieved, the dataset was quite small to deploy the system in practice. Our project, on the other hand, has been implemented on 175,000 heartbeats.

- **Hybrid Approach:**

We have employed a hybrid approach where instead of building a single large model, two models have been used to train our dataset based on 1D and 2D data sets. With this approach, we can accurately monitor a larger region.

- **Lightweight and Suitable for Wearables:**

The fact that the models are lightweight means that they can be executed by low-powered hardware. It is therefore possible to deploy this technology in wearable gadgets like wrist-watches or bands to enable real-time measurement of heart rate.

5. Conclusion

- Automated detection of arrhythmia is not only an issue of machine learning – it is an issue of patients' lives, which is why our work was developed bearing in mind the safety of those patients.
- Striving for the highest possible accuracy was not the aim; with nearly 175,000 real heartbeats, we focused on ensuring that the model almost never misses out on identifying a potentially lethal beat because, in the field of cardiology, one missed beat may prove fatal.
- Reformulation of the problem as Binary Classification helped us overcome the extreme

class imbalance in the data that limited previous solutions – the imbalance was narrowed down from 25,000:1 to 6.9:1.

- In our hybrid model that uses the combination of a 1D ensemble of statistical models, a 2D CNN, and a meta-learner, it is very clear that there is no one model that is capable of detecting all the patterns. It is only through their collaboration that they succeed in detecting some patterns. The proposed solution has provided us with impressive results: 99.63% accuracy, 99.57% recall, and almost perfect ROC-AUC equal to 0.9999.

The present analysis can be integrated into smartwatches to provide real-time alerts for arrhythmia detection.

References

[1] G. Gelbenegger *et al.*, "Bundle branch blocks and cardiovascular mortality," *Heart Rhythm*, vol. 21, 2024.

[2] PhysioNet Databases (MIT-BIH, BIDMC, and PTB), 2000. World Health Organization, "Cardiovascular diseases (CVDs)," 2021.

[3] A. DeBauge et al., "Prediction of left ventricular hypertrophy in women and men with left bundle branch block," *J Electrocardiol*, vol. 80, pp. 34-39, 2023.

[4] N. Y. Tan et al., "Left Bundle Branch Block: Current and Future Perspectives," *Circulation: Arrhythmia and Electrophysiology*, vol. 13, no. 4, 2020.

[5] Z. Masetic and A. Subasi, "Congestive heart failure detection using random forest classifier," *Comput Methods Programs Biomed*, vol. 130, pp. 54-64, 2016.

[6] Bussink BE, Holst AG, Jespersen L, Deckers JW, Jensen GB, Prescott E. Rightbundle branch block: prevalence, risk factors, and outcome in the generalpopulation: results from the Copenhagen City Heart Study. Eur Heart J 2013;

[7] Maron DJ, Hochman JS, Reynolds HR, et al; ISCHEMIA Research Group. Initial invasive or conservative strategy for stable coronary disease. N Engl J Med 2020; 382:1395–1407.

[8] Bangalore S, Maron DJ, O'Brien SM, et al; ISCHEMIA-CKD Research Group. Management of coronary disease in patients with advanced kidney disease. N Engl J Med 2020; 382:1608–1618.

[9] Heidenreich PA, Bozkurt B, Aguilar D, et al. 2022 AHA/ACC/HFSA Guideline for the Management of Heart Failure: a report of the American College of Cardiology/American Heart Association Joint Committee on Clinical Practice Guidelines. Circulation 2022;145:e895–e1032.

[10] Lang RM, Badano LP, Mor-Avi V, Afilalo J, Armstrong A, Ernande L, et al. Recommendations for cardiac chamber quantification by echocardiography in adults: an update from the American Society of Echocardiography and the European Association of Cardiovascular Imaging. Eur Heart J Cardiovasc Imaging 2015;16(3):233–71.

[11] Vakili BA, Okin PM, Devereux RB. Prognostic implications of left ventricular hypertrophy. Am Heart J 2001;141(3):334–41.

[12] Devereux RB. Is the electrocardiogram still useful for detection of left ventricular hypertrophy? Circulation 1990;81(3):1144–6.

[13] Sokolow M, Lyon TP. The ventricular complex in left ventricular hypertrophy as obtained by unipolar precordial and limb leads. Am Heart J 1949;37(2):161–86.

[14] Romhilt DW, Estes Jr EH. A point-score system for the ECG diagnosis of left ventricular hypertrophy. Am Heart J 1968;75(6):752–8.

[15] K. V. Beena and R. K. Megalingam, "Multimodel Integration for FPGA based ECG Signal Processing," 2025 8th International Conference on Circuit, Power & Computing Technologies (ICCPCT), Kollam, India.

[16] B. R. Manju and A. R. Nair, "Classification of Cardiac Arrhythmia of 12 Lead ECG Using Combination of SMOTEENN, XGBoost and Machine Learning Algorithms," *2019 9th International Symposium on Embedded Computing and System Design (ISED)*, Kollam, India, 2019, pp.

[17] S. Syama, G. S. Sweta, P. I. K. Kavyasree and K. J. M. Reddy, "Classification of ECG Signal Using Machine Learning Techniques," *2019 2nd International Conference on Power and Embedded Drive Control (ICPEDC)*, Chennai, India, 2019, pp.

[18] Kusuma, Stalin & Udayan J, Divya. (2020). Analysis on Deep Learning methods for ECG based Cardiovascular Disease prediction. Scalable Computing: Practice and Experience