Arrhythmia Classification at the Edge Using a Weightless Neural Network Hardware Accelerator

Anshul Jha[1[0000-0002-5491-1884]], Shashank Nag[2[0009-0000-3619-6066]], Igor D. S. Miranda[3[0000-0002-8555-9609]], Felipe M. G. Franca[4[0000-0002-8980-6208]], Lizy K. John[2[0000-0002-8747-5214]], Priscila M. V. Lima[5[0000-0002-8515-9904]], and Eugene B. John[6[0000-0001-9494-4894]]

[1] University of The Pacific, California, USA
ajha@pacific.edu
[2] The University of Texas at Austin, Texas, USA
shashanknag@utexas.edu
ljohn@ece.utexas.edu
[3] Federal University of Reconcavo da Bahia (UFRB), Cruz das Almas, Brazil
igordantas@ufrb.edu.br
[4] Instituto de Telecomunicações, Portugal (now at Google LLC)
felipe@ieee.org
[5] Federal University of Rio de Janeiro, Brazil
priscilamvl@cos.ufrj.br
[6] The University of Texas at San Antonio, USA
Eugene.John@utsa.edu

Abstract. Recent advances in hardware architectures enable real-time, low-latency, and energy-efficient computation capabilities that are essential in medical applications where timely and reliable decisions can directly affect patient outcomes. While traditional Artificial Intelligence (AI) approaches such as Convolutional Neural Networks (CNNs) and Deep Neural Networks (DNNs) have shown strong performance in biomedical signal analysis, their reliance on computationally intensive operations often limits deployment in low-power, resource-constrained medical devices. This paper presents a lightweight hardware accelerator for real-time detection and classification of cardiac arrhythmias from electrocardiogram (ECG) signals using Weightless Neural Networks (WNNs). The proposed architecture leverages an ensemble of Wilkie–Stonham–Aleksander Recognition Device (WiSARD)-based classifiers, replacing conventional weight-based computation with RAM-based storage. Multi-threshold binarization is used to extract discriminative ECG features across multiple intensity levels, while a tie-breaking mechanism improves robustness during ensemble decision-making. The complete design is implemented in synthesizable SystemVerilog and validated through functional simulation. The proposed model achieves 92.57% accuracy on the MIT-BIH dataset across five Association for the Advancement of Medical Instrumentation (AAMI) standard heartbeat classes. ASIC synthesis in 45 nm and 90 nm technology nodes further demonstrates the accelerator's energy efficiency, with the 45 nm implementation consuming 30.18 mW at 100 MHz, underscoring its suitability for low-power, real-time point-of-care medical inference.

Keywords: Arrhythmia classification · Weightless Neural Network · Low-power hardware accelerator.

1 Introduction

Improvements in semiconductor technology and hardware accelerator architectures have substantially enhanced the ability to execute complex algorithms in real time while maintaining low power consumption. These developments are particularly critical in medical applications, where millisecond-level delays or inaccurate predictions can directly impact patient outcomes. With the growing adoption of point-of-care and wearable medical devices, there is an increasing demand for on-device intelligence capable of real-time signal processing, classification, and decision-making. Consequently, the design of energy-efficient hardware accelerators has emerged as a key enabler for next-generation medical systems.

AI techniques have demonstrated strong capability in biomedical signal and image analysis, supporting automated diagnosis and continuous patient monitoring. Deep learning (DL) models such as CNNs and DNNs have been widely deployed for tasks including medical image classification and physiological signal analysis. However, these models rely heavily on floating-point arithmetic and large-scale matrix multiplications, which result in significant computational complexity and energy consumption. Therefore, these limitations make the deployment of DL models on low-power, real-time medical hardware accelerators particularly challenging. To address this issue, an alternative AI model is explored, which is inherently more hardware-efficient.

WNNs offer an essentially different learning and inference approach as compared to conventional weight-based models. Instead of performing weighted summations, WNNs store learned patterns directly in Random Access Memory (RAM), and the inference is performed through simple memory access operations. Hence, the memory-centric architecture enables fast, deterministic inference with minimal computational resources. The WiSARD WNN architecture is especially attractive due to its single-pass training, deterministic classification, and suitability for lightweight, low-power hardware implementations.

To further enhance performance, the Ensemble WiSARD architecture employs multiple WiSARD classifiers, each trained on different subsets or distinct input mappings of the data [5]. The WiSARD outputs are combined using majority voting to yield the final classification. This approach reduces overfitting associated with single models and improves robustness and generalization.

The requirement for efficient and reliable AI-driven hardware accelerators is recognizable for cardiovascular health monitoring. Cardiac disorders are among the most prevalent health conditions globally. According to the World Health Organization (WHO) [15], cardiovascular diseases (CVDs) remain the leading cause of death worldwide, accounting for 17.9 million deaths annually. A major subset of these conditions is cardiac arrhythmia, which results from abnormalities in the heart's electrical activation sequence and leads to irregular heart rhythms. While preliminary indications of arrhythmia may be observed through pulse assessment or auscultation during routine examinations, accurate diagnosis requires signal-based and imaging techniques. Among available diagnostic modalities, including electrocardiography (ECG), cardiac magnetic resonance imaging (MRI), and computed tomography (CT), ECG remains the most widely used and clinically practical tool for arrhythmia detection. ECG is a non-invasive technique that records the heart's electrical activity during each contraction and relaxation cycle using surface electrodes placed on the chest and, in some cases, on the limbs. The resulting electrical signals are transmitted via leads to a recording device, which generates characteristic waveforms for clinical interpretation and treatment monitoring.

Building upon these observations, this paper focuses on the design and hardware realization of an ensemble of WiSARD-based WNN accelerator for real-time cardiac arrhythmia detection using ECG signals. Cardiac arrhythmias are characterized by irregular electrical patterns; if not detected promptly, can lead to severe or life-threatening conditions. Under normal conditions, ECG signals exhibit repetitive and structured waveforms, making them well-suited for pattern-based classification. The proposed model is evaluated at the algorithmic level using the MIT-BIH dataset and is subsequently translated into a fully synthesizable ASIC design and evaulated on ECG signals. Hardware implementations across multiple technology nodes are explored to assess power consumption and area efficiency.

The remainder of this paper is organized as follows. Section 2 provides background information and reviews related work on WNN-based approaches for arrhythmia classification and detection. Section 3 details the methodology and implementation of the proposed WNN-based hardware architecture. Section 4 presents the experimental results and hardware evaluation, followed by concluding remarks in Section 5.

2 Literature Review

The integration of AI into medical diagnostics has significantly reformed disease detection by enabling faster, more accurate, and scalable solutions that support clinical decision-making. In recent years, the global healthcare community has faced growing challenges in diagnosing and managing life-threatening conditions such as cardiac arrhythmias. These challenges highlight the need for intelligent and automated diagnostic systems that can deliver real-time, reliable results, particularly in resource-constrained and point-of-care settings.

Under normal conditions, the heart transmits electrical impulses through specialized conduction pathways that transmit signals from the atria (upper chambers) to the ventricles (lower chambers). This rhythmic electrical activity regulates cardiac contractions and is controlled by the sinoatrial and atrioventricular nodes, which ensure the timely excitation of myocardial tissues [7]. ECGs are a non-invasive test that captures these electrical signals using surface electrodes placed on the skin. The acquired ECG signal is typically filtered to remove noise and artifacts before being digitized using an analog-to-digital converter. A healthy ECG waveform exhibits periodic and consistent morphology, characterized by the P wave, PR interval, QRS complex, ST segment, T wave, and the subsequent resting phase.

To enable efficient detection of arrhythmias on edge and wearable platforms, WNN-based hardware accelerators have gained increasing attention. Unlike conventional DL models, WNN accelerators leverage lightweight operations and avoid computationally expensive multiply-accumulate (MAC) units. Large WNN arrays can be implemented as RAM-based neurons, which helps to achieve throughput with substantially lower energy consumption. Due to the lower hardware utilization, WNN accelerators are inherently well-suited for deployment in energy and area-constrained edge devices. Compared to weight-based neural networks such as CNNs and DNNs, which rely heavily on high-precision arithmetic operations, WNN-based hardware offers a powerful replacement for low-power medical inference.

Pillai et al. [12] proposed arrWNN, a two-class WNN architecture designed for normal and arrhythmia classification. The proposed architecture converts arriving serial ECG bit streams into 8-bit parallel words and instantiates 741 logic minterms derived using the Combinational Intelligent Networks (COIN) training methodology [10]. During inference, a multiplier selects the appropriate minterm group corresponding to each input tuple, while up-down counters accumulate class-wise votes. The final prediction is determined using an argmax operation. The arrWNN model was evaluated using a 50-iteration Monte Carlo cross-validation framework on the MIT-BIH arrhythmia database, achieving a mean classification accuracy of 88.27%. The trained model was synthesized in a 0.6-µm indium–gallium–zinc–oxide (IGZO)–based flexible integrated circuit (FlexIC) technology, resulting in a 24-mm^2 core with 5,706 NAND2-equivalent gates. Operating at 100 kHz and 3 V, the design consumed 9.4 mW, demonstrating its feasibility for low-power medical applications.

In subsequent work, Pillai et al. [13] proposed arrDWNN, an edge-optimized arrhythmia detection model inspired by LogicWiSARD [11] and COIN [10] architectures. ECG signals from the MIT-BIH database were preprocessed, unary-encoded, and mapped to 8-bit addresses as inputs to the WNN. The architecture employs 741 minterms, a serial-to-parallel converter, and up-down counters, with final classification performed via an argmax unit. The arrDWNN model achieved an accuracy of 88.27%, sensitivity of 68.62%, specificity of 99.83%, and an area under the curve (AUC) of 0.8422. Implemented in a 45-nm CMOS technology, the hardware consumed 4.24 mW of power, occupied 12.4 mm^2, and operated at 500 MHz, highlighting its suitability for low-power, real-time wearable ECG monitoring systems.

3 Methodology & Implementation

The paper proposes an algorithm-hardware co-design methodology. The proposed methodology utilizes a WiSARD-based WNN ensemble with an integrated tie-breaking mechanism to improve robustness under overlapping feature distributions. The WNN algorithm is translated to an application-specific hardware accelerator and synthesized on ASIC technology nodes to analyze area and power efficiency for edge medical applications.

3.1 Arrhythmia Dataset

The proposed implementation and evaluation of the ensembled WiSARD-based WNN are conducted using the MIT-BIH Arrhythmia Database, which is a widely accepted benchmark for cardiac signal analysis. The dataset consists of 48 ECG recordings obtained from 47 distinct subjects, with each record containing a 30-minute segment selected from continuous 24-hour ambulatory monitoring. The signals were acquired using a two-channel Holter monitor and digitized at a sampling frequency of 360 Hz, offering sufficient temporal resolution to capture the morphological characteristics of individual heartbeats. Fig. 1 represents the annotated beats of record 101.

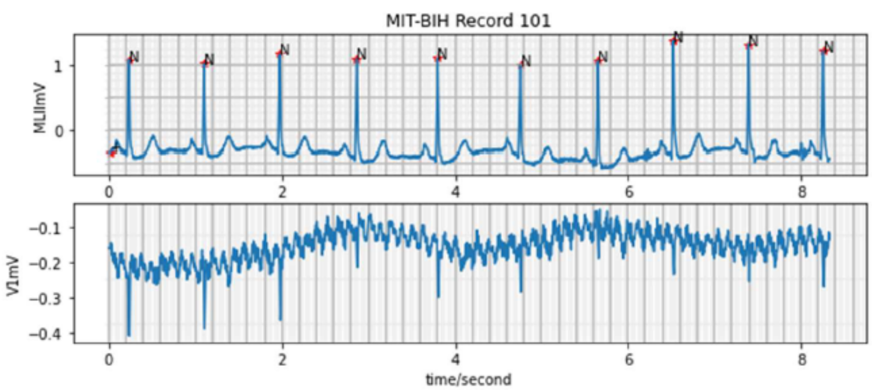

Fig. 1. Annotated beat of record 101 from the MIT-BIH Database. [6]

In this research, heartbeat annotations follow the Association for the Advancement of Medical Instrumentation (AAMI) standard [3], which consists of five primary beat types: normal (N), supraventricular (S), ventricular (V), fusion (F), and unknown (Q). Table 1 summarizes the subsequent MIT-BIH arrhythmia categories based on the AAMI standard. The data is first converted to grayscale images. Each heartbeat is segmented using a fixed window size of 179 $\times$ 179 images [2], ensuring preservation of local waveform morphology. The segmented beats are subsequently normalized and transformed into grayscale images to represent the temporal features of each heartbeat, as illustrated in Fig. 2.

To evaluate the model performance, the dataset is partitioned into 80%training samples and 20%testing samples. The distribution of heartbeats across different classes is reported in Table 2. This partitioning uncovers the class imbalance across different classes. Therefore, to mitigate the challenge of learning parameters from classes with an imbalanced dataset, robust modeling strategies are adopted, such as ensemble learning and multi-threshold processing.

Table 1. MIT-BIH arrhythmia dataset according to AAMI standards. [3]

AAMI Class	**Content Description**	**Subclass Description**
N	Non-Ectopic Beats	N: Left bundle branch block L: Normal ECG beat R: Right bundle branch block e: Atrial escape beat j: Borderline escape beat
S	Supraventricular Ectopic Beats	A: Atrial premature beats a: Abnormal atrial premature beats J: Borderline premature beats S: Supraventricular premature beats
V	Ventricular Ectopic Beats	V: Ventricular premature beats E: Ventricular escape beat
F	Fusion Beats	F: Ventricular fusion heartbeat
Q	Unknown Beats	/: Paced beat f: Pacing and normal fusion heartbeat Q: Unclassifiable beat ?: Beat not classified during learning

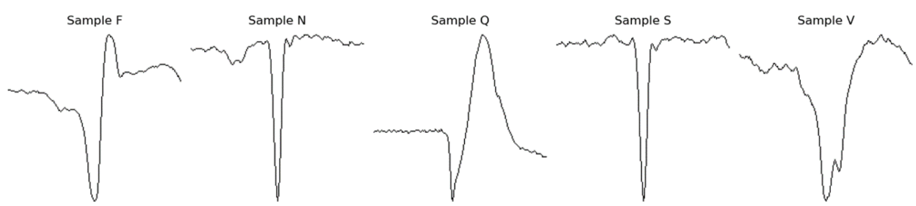

Fig. 2. Images of the segmented beats from each class of the MIT-BIH database.

3.2 Data Preprocessing and Multi-Threshold Binarization

Single-threshold binarization is often insufficient for capturing subtle variations in ECG signals, as clinical information can be impacted from small change in the waveform. In WiSARD-based WNNs, binarization is an essential requirement, as the model operates directly on binary inputs, with each bit corresponding to a unique address in the RAM-based architecture. Therefore, an effective binarization method is critical to preserve ECG signal characteristics. To address

Table 2. MIT-BIH arrhythmia dataset classification for training and testing of the proposed model

MIT-BIH Arrhythmia Classes	Total Heartbeats	Training Heartbeats	Testing Heartbeats
N	3195	2556	639
S	1382	1105	277
V	992	793	199
F	373	298	75
Q	2084	1667	417

the information loss associated with single-threshold binarization, this paper employs a multi-threshold binarization system that enhances feature representation without complex preprocessing. Three binary maps are generated for each input sample using different intensity thresholds of low (≥ 0.3), mid (≥ 0.5), and high (≥ 0.7). This approach enables the model to capture corresponding signal features across multiple intensity levels. Thus, improving sensitivity to subtle morphological variations associated with different arrhythmia classes.

The resulting binary maps are combined to form a three-channel binary representation of each input. This representation is flattened into a one-dimensional binary vector for compatibility with the tuple-mapping process in the WiSARD model. The threshold-encoded input provides a prolific and more discriminative binary representation.

3.3 Proposed Model Architecture

The proposed model adopts a WNN based on the WiSARD architecture [1], which replaces weight-based learning and backpropagation with a RAM-centric pattern storage mechanism. The output of each class is represented by a dedicated discriminator composed of multiple RAM nodes, where each RAM observes a fixed-length tuple of size n from the binary input vector. For an input length I, each discriminator contains $N \triangleq \frac{I}{n}$ RAMs. Pseudo-random but consistent input-to-RAM mappings are applied across all discriminators to ensure uniform feature interpretation among various classes.

Training of the model follows a single-pass learning paradigm in which only the discriminator corresponding to the true class is updated. Each RAM writes a binary value at the address generated by its input tuple, directly storing observed patterns. During inference, the input is simultaneously evaluated by all discriminators, and each RAM contributes a vote based on whether its addressed memory location is active. The discriminator response is processed as a pop count of activated RAMs, and the class with the highest response is the prediction. If the pattern matches precisely, it results in full activation, while partial matches results into generalization to unobserved but similar inputs. This characteristic makes WiSARD well-suited for ECG-based arrhythmia classification.

The tuple size n controls the trade-off between generalization and capacity to learn intricate input patterns. The larger tuple values capture more complex patterns at the risk of overfitting. This paper selected a tuple size n=12, resulting in 4096 memory locations per RAM, with each class discriminator consisting of 120 RAM nodes. To improve robustness and reduce bias, an ensemble of five independently trained WiSARD models is employed to leverage the variety in random mappings. A tie-breaker mechanism is integrated to resolve tied predictions by gradually increasing the activation threshold until a single class standout as a winner. The resulting ensemble WiSARD architecture is interpretable, memory-efficient, and hardware-friendly, making it well-suited for low-power edge deployment in ECG arrhythmia classification. Fig 3 illustrates the implemented ensemble WiSARD architecture with a tie-breaker, tailored for this classification task

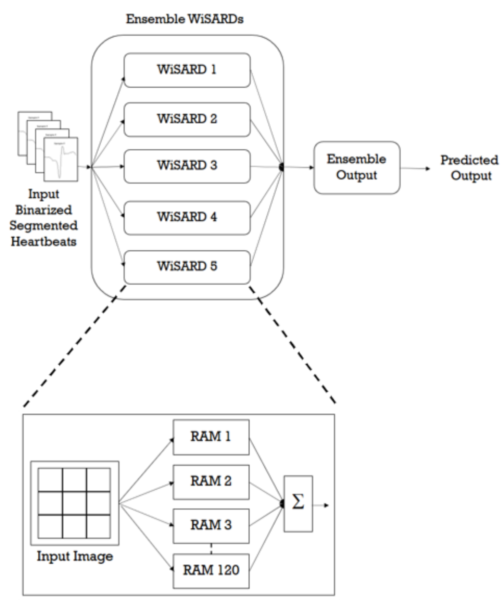

Fig. 3. Overview of the proposed Multi-threshold Ensemble WiSARD Architecture.

3.4 Proposed Model training

The WiSARD model employs a one-pass learning strategy that trades iterative weight updates with direct pattern storage in RAM. During training, each input is turned into fixed-size tuples, where each tuple forms an address to a RAM node in a class-specific discriminator. Only the discriminator corresponding to the correct class is updated. The addressed memory locations are written once with a binary value of '1'. Therefore, this non-iterative learning approach enables rapid training and deterministic behavior, making it well-suited for resource-constrained applications such as real-time ECG-based arrhythmia classification.

During inference, the model operates in read-only mode, evaluating all class discriminators concurrently. Each RAM node contributes a vote if its addressed memory location is active, and class-wise responses are computed by aggregating these votes. This process is independently performed across all ensemble components is combined k response vector element-wise, to produce an ensemble-level response vector.

$$R_{\text{ensemble}}(c) = \sum_{i=1}^{k} R_i(c) \tag{1}$$

where $R_i(c)$ is the count for model i for the class c.

The final prediction c^* is achieved by summing class responses and selecting the maximum.

$$c^* = \arg\max_c R_{\text{ensemble}}(c) \tag{2}$$

Further to improve reliability for overlapping or uncertain patterns, a tie-breaker mechanism is incorporated to resolve tied predictions by gradually increasing a confidence threshold, allowing only strongly activated RAMs to participate. The resulting ensemble WiSARD framework is interpretable, memory-efficient, and hardware-friendly, making it suitable for low-power edge deployment without reliance on gradient-based optimization.

3.5 HDL-Based Implementation of the Ensemble WiSARD

The HDL-based implementation of the ensemble WiSARD accelerator is generated using the Mako [9] template, which enables SystemVerilog code generation through Python-based templates. This approach allows complex hardware structures to be expressed compactly and flexibly ahead of the restrictions of SystemVerilog preprocessing. A Python script retrieves the trained model checkpoints and extracts multi-threshold binarization patterns aligned with the address mappings of individual WiSARD discriminators. These patterns are implanted directly into RAM structures via Mako templates, presenting synthesizable RTL that instantiates local pattern memory. Unlike conventional neural networks that rely on floating-point weights and external memory access, the proposed architecture employs logic-based memory, reducing memory access latency and energy consumption.

The hardware realization of the ensemble WiSARD focuses on efficiently mapping memory-based pattern recognition instead of arithmetic-intensive computation. Each trained RAM node is implemented as a synthesizable lookup structure indexed by input tuples. The proposed architecture supports parallel address generation, concurrent RAM access across class discriminators, and aggregation of ensemble responses using pipelined summation to enhance the throughput and reduce inference latency of the accelerator. The functional verification is accomplished by utilizing binarized test vectors, which are generated by using the same preprocessing pipeline as the software model. The RTL design predicted class outputs are compared against software outcomes. This step is used to verify the behavioral functionality between the trained model and hardware realization before synthesis.

3.6 Evaluation Criteria

Model performance is evaluated using standard classification metrics, including accuracy, precision, recall, F1-score, and the confusion matrix. Collectively, these metrics provide a thorough assessment of classification efficiency, whereas the confusion matrix offers complete insight into class-wise misclassifications. The model performance analysis is very crucial for arrhythmia detection, where errors within specific classes may result in different clinical consequences.

Precision, recall, F1-score, and accuracy are computed using the counts of true positives, true negatives, false positives, and false negatives. Precision reflects the correctness of positive predictions.

$$\text{Precision} = \frac{TP}{TP + FP} \tag{3}$$

Recall or sensitivity measures the model's ability to correctly identify positive cases.

$$\text{Recall} = \frac{TP}{TP + FN} \tag{4}$$

The F1-score captures the balance between precision and recall.

$$F1\text{-}score = 2 \cdot \frac{TP}{TP + FP + FN} \tag{5}$$

Accuracy represents the ratio of correctly classified samples to the total number of samples.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{6}$$

where
TP (True Positive): correctly identified positive samples.
TN (True Negative): correctly identified negative samples.
FP (False Positive): negative samples are incorrectly classified as positive.
FN (False Negative): positive samples incorrectly classified as negative.

These metrics are critical for medical diagnostic systems, where minimizing missed detections and false alarms is essential. The resulting performance comparison enables an objective evaluation of model reliability and suitability for clinical deployment.

To evaluate hardware accelerator efficiency, the SystemVerilog-HDL design is synthesized using Synopsys Design Compiler targeting 90 nm and 45 nm ASIC technology nodes under operating conditions: 1.2 V power supply, 100 MHz clock frequency, 12.5% switching activity, and 25 °C temperature. The synthesis process yields area and power reports, facilitating comparative analysis among technology nodes. Cadence Innovus is utilized for physical design exploration to synthesis, floorplanning, placement, routing, and the production of the final layout in GDS II format. Together, these results provide insight into the scalability, energy efficiency, and suitability of the ensemble WiSARD architecture for low-power ASIC-based edge deployment.

4 Results

An ensemble of five WiSARD-based WNNs was developed and evaluated for arrhythmia classification using the MIT-BIH database as per AAMI standards. The model classifies heartbeats into five clinically relevant classes: normal (N), supraventricular ectopic (S), ventricular ectopic (V), fusion (F), and unknown (Q). The dataset is split using an 80:20 training-testing ratio, and performance is evaluated using accuracy and cross-entropy loss, along with macro and weighted accuracy to account for class imbalance. The macro accuracy gives equal weight to all classes regardless of their sample sizes. In contrast, weighted accuracy considers the relative size of each class and is determined by dividing the total number of correctly identified samples by the total number of samples.

$$\text{Macro Accuracy} = \frac{1}{N}\sum_{i=1}^{N}\frac{TP_i}{\text{Total samples of class } i} \tag{7}$$

$$\text{Weighted (Overall) Accuracy} = \frac{\text{Total correct predictions}}{\text{Total number of samples}} \tag{8}$$

On the algorithmic-level the ensemble achieved a training accuracy of 92.7% and a validation accuracy of 93.3%, with training and validation losses of 0.5159 and 0.5425, respectively. The Table 3 presents the class-wise macro accuracy, which shows strong performance across most arrhythmia types, with particularly high accuracy for N, Q, and S classes, while, due to limited data availability, the F class shows expected prediction.

The Table 4 represents the precision, recall, and F1-score results, which further prove that the ensemble WiSARD model efficiently adapts to class imbalance, maintaining high recall for minority classes.

As shown in Fig. 4, the confusion matrix confirms the model's reliability across all classes, including the recognition of fusion beats. The results highlight the benefit of ensemble learning for clinically relevant minority patterns.

Algorithm-level evaluation confirms that the proposed WiSARD ensemble model can be efficiently translated from PyTorch to synthesize HDL without loss of classification accuracy. For HDL functional verification, 1,607 inference vectors are utilized, which shows the same behavior between software and hardware

Table 3. Macro Accuracy for each arrhythmia class

Class	Macro-Accuracy
F	47%
N	97%
Q	98%
S	95%
V	85%

Table 4. Classification report of Ensemble-WiSARD on the MIT-BIH arrhythmia dataset

Classes	Precision	Recall	F1-score
F	0.95	0.47	0.62
N	0.96	0.97	0.96
Q	0.87	0.98	0.93
S	0.96	0.95	0.96
V	0.96	0.85	0.90

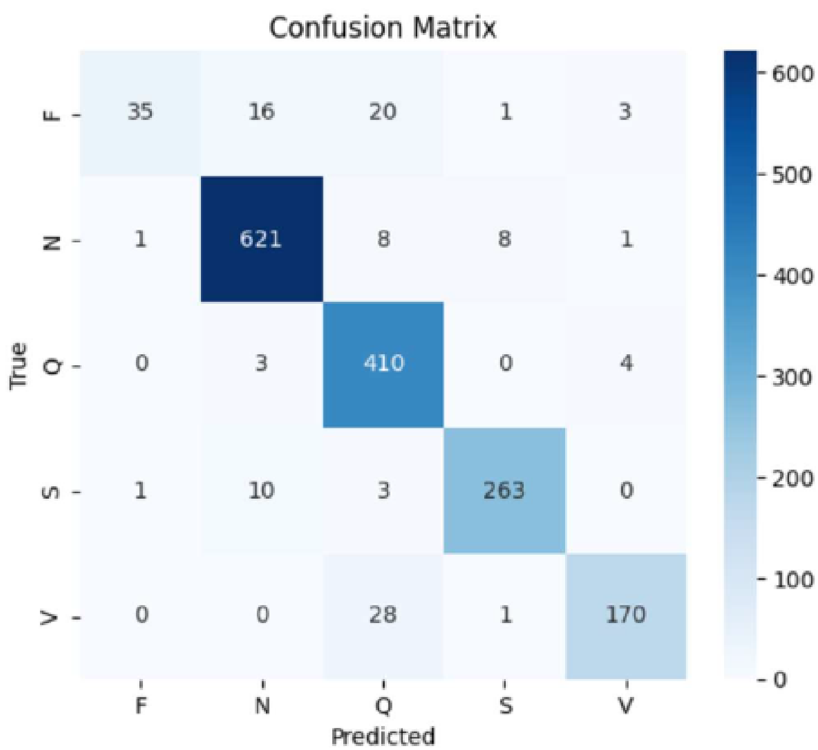

Fig. 4. Confusion matrix of Ensemble-WiSARD on the MIT-BIH dataset.

Table 5. Performance of Individual WiSARD Models and Ensemble

Model / Submodels	Test Acc. (%)
WiSARD1	86.02
WiSARD2	88.18
WiSARD3	79.06
WiSARD4	81.34
WiSARD5	82.18
Ensemble (5 models)	**92.57**

output, achieving a hardware accelerator accuracy of 92.57%. Table 5 presents a comparison of the test accuracy of each standalone model, along with the ensemble implementation. Each WiSARD model utilizes 1-bit precision per input, 12 inputs per filter, 4096 memory entries per filter, 120 RAMs, and a memory size of 500 KiB per WiSARD, leading to 2500 KiB for the entire ensemble model. The results clearly demonstrate the improvement in classification performance achieved through ensemble learning.

An ASIC synthesized using Synopsys Design Compiler demonstrates scaling across 90 nm and 45 nm technology nodes under identical operating conditions. Table 6 represents the total, dynamic, switching, and internal power consumption on the 45 nm and 90 nm technology nodes. Table 7represents the chip area consumption on 45 nm and 90 nm technology nodes.

Table 6. Power consumption of the ensemble WiSARD model on 45 nm and 90 nm technology nodes

Technology Node	Dynamic Power (mW)	Switching Power (mW)	Leakage Power (mW)	Total Power (mW)
45 nm	26.4433	0.528927	3.7429	30.181
90 nm	47.2034	1.77099	7.8228	55.026

Table 7. Chip area of the ensemble WiSARD model on 45 nm and 90 nm technology nodes

Technology Node	Area (μm^2)
45 nm	5,009,366.59
90 nm	13,982,588.92

The 45 nm implementation achieves significant reductions in both area and total power consumption. Fig. 5 presents the physical design placement results obtained using Cadence Innovus, validating the feasibility of a full place-and-route implementation, showing well-distributed logic, robust power delivery, and compliance with timing and design-rule constraints. Together these results demonstrate that the ensemble WiSARD architecture provides an effective balance between classification accuracy and hardware efficiency, making it suitable for low-power, real-time arrhythmia detection in edge medical devices.

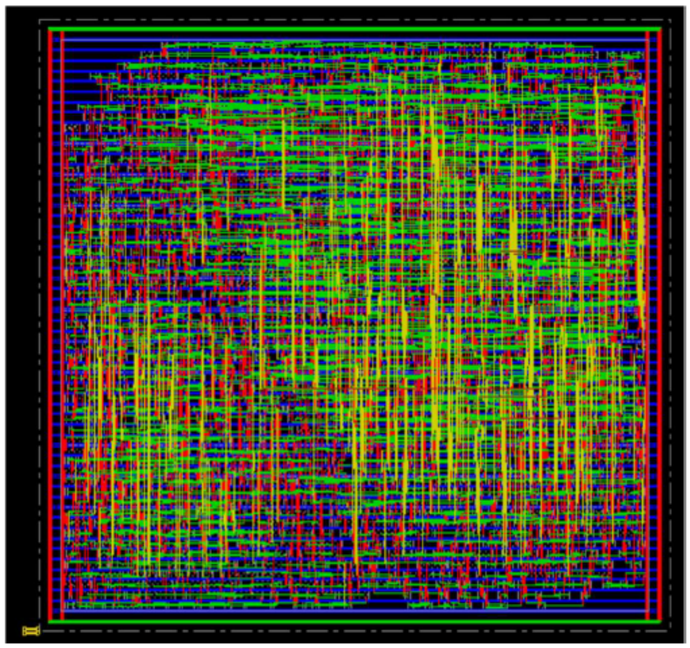

Fig. 5. Placement of the ensemble WiSARD model using Cadence Innovus.

At the system level, the proposed ensemble WiSARD model is synthesized and physically realized using Cadence Innovus with TSMC 0.18 μm CMOS technology, resulting in a complete GDS II layout. The design achieves an arrhythmia classification accuracy of approximately 92.57% on the MIT-BIH dataset. Power and area characteristics are evaluated using Synopsys Design Compiler across 45 nm and 90 nm technology nodes. At 45 nm, the implementation consumes 30.18 mW of total power and occupies 5009366.59 μm^2 of silicon area, demonstrating efficiency for resource-constrained biomedical edge applications. While Pillai et al. [12] and Pillai et al. [13] reported arrWNN and arrDWNN models based on LogicWiSARD and COIN-based WNN architectures with accuracies of 88.27% and 89%, respectively, the proposed ensemble of five WiSARD networks achieves a competitive accuracy of 92.57%.

5 Conclusion

This paper presents an ensemble-based WiSARD weightless neural network (WNN) for real-time arrhythmia classification. The hardware design is evaluated on the MIT-BIH dataset under AAMI standards. While WNNs have previously been explored for arrhythmia detection, to the best of our knowledge this is the first work to apply a WNN-based approach to multi-class arrhythmia classification. An ensemble of five WiSARD models delivers strong system-level performance, achieving 92.57% accuracy with consistently high precision, recall, and F1-scores across all arrhythmia classes. These results highlight the effectiveness of ensemble learning in addressing class imbalance and reducing classification ambiguity. The proposed model is implemented in synthesizable SystemVerilog HDL and validated through a complete ASIC design flow, with hardware results matching software accuracy. ASIC synthesis in 45 nm technology achieves 30.18 mW power consumption and a compact area of 5.01 mm^2 at 100 MHz. Overall, the results demonstrate the practicality of the ensemble WiSARD architecture as an energy-efficient, hardware-friendly solution for deployable, edge-level medical diagnostics.

Acknowledgments. This research was supported in part by NSF Grants 2326894 and 2425655. Any opinions, findings, conclusions, or recommendations are those of the authors and not of the funding agencies.

Disclosure of Interests. Authors have no competing interests.

References

1. Aleksander, I., Thomas, W., Bowden, P.: Wisard· a radical step forward in image recognition. Sensor review **4**(3), 120–124 (1984)
2. Aphale, S., Jha, A., John, E.: High accuracy arrhythmia classification using transfer learning with fine-tuning. In: 2022 IEEE 13th Annual Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON). pp. 0480–0487. IEEE (2022)
3. Aphale, S.S., John, E., Banerjee, T.: Arrhynet: a high accuracy arrhythmia classification convolutional neural network. In: 2021 IEEE international midwest symposium on circuits and systems (MWSCAS). pp. 453–457. IEEE (2021)
4. Carvalho, D.S., Carneiro, H.C., França, F.M., Lima, P.M.: B-bleaching: Agile overtraining avoidance in the wisard weightless neural classifier. In: ESANN (2013)
5. Filho, L.L., França, F.M., Lima, P.M.: Wisard-based ensemble learning. In: ESANN 2023 Proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. i6doc.com, Bruges, Belgium (oct 2023)
6. Goldberger, A.L., Amaral, L.A., Glass, L., Hausdorff, J.M., Ivanov, P.C., Mark, R.G., Mietus, J.E., Moody, G.B., Peng, C.K., Stanley, H.E.: Physiobank, physiotoolkit, and physionet: components of a new research resource for complex physiologic signals. circulation **101**(23), e215–e220 (2000)
7. Haddad, S.A., Serdijn, W.A.: The evolution of pacemakers: An electronics perspective. In: Ultra Low-Power Biomedical Signal Processing: An Analog Wavelet Filter Approach for Pacemakers, pp. 13–31. Springer (2009)

8. Jha, A., John, E., Banerjee, T.: Transfer learning for covid-19 and pneumonia detection using chest x-rays. In: 2022 IEEE 65th International Midwest Symposium on Circuits and Systems (MWSCAS). pp. 1–4. IEEE (2022)
9. Mako Templates: Welcome to mako! https://www.makotemplates.org/, accessed: Jan 2026
10. Miranda, I.D., Arora, A., Susskind, Z., Souza, J.S., Jadhao, M.P., Villon, L.A., Dutra, D.L., Lima, P.M., França, F.M., Breternitz, M., et al.: Coin: Combinational intelligent networks. In: 2023 IEEE 34th International Conference on Application-specific Systems, Architectures and Processors (ASAP). pp. 27–28. IEEE (2023)
11. Miranda, I.D., Arora, A., Susskind, Z., Villon, L.A., Katopodis, R.F., Dutra, D.L., De Araújo, L.S., Lima, P.M., França, F.M., John, L.K., et al.: Logicwisard: Memoryless synthesis of weightless neural networks. In: 2022 IEEE 33rd International conference on application-specific systems, architectures and processors (ASAP). pp. 19–26. IEEE (2022)
12. Pillai, V., Miranda, I.D., Musale, T., Jadhao, M., Neto, P.C.S., Susskind, Z., Bacellar, A.T., Lhostis, M., Lima, P.M., Dutra, D.L., et al.: arrwnn: Arrhythmia-detecting weightless neural network flexic. In: 2024 IEEE International Flexible Electronics Technology Conference (IFETC). pp. 1–4. IEEE (2024)
13. Pillai, V., Nag, S., Jadhao, M., Bacellar, A., Miranda, I.D., França, F.M., Dutra, D.L., Lima, P.M., John, L.K., John, E.B.: Edge-optimized weightless neural network for low-power wearable arrhythmia detection. In: 2025 IEEE 18th Dallas Circuits and Systems Conference (DCAS). pp. 1–6. IEEE (2025)
14. Weng, O., Andronic, M., Zuberi, D., Chen, J., Geniesse, C., Constantinides, G.A., Tran, N., Fraser, N.J., Duarte, J.M., Kastner, R.: Greater than the sum of its luts: Scaling up lut-based neural networks with amigolut. In: Proceedings of the 2025 ACM/SIGDA International Symposium on Field Programmable Gate Arrays. pp. 25–35 (2025)
15. World Health Organization: Cardiovascular diseases. World Health Organization: CVD (2025), https://www.who.int/health-topics/cardiovascular-diseases, accessed: Jan 2026