WWW.IJRPR.COM

International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

Scalable VLSI Architecture for Deep Learning - Based Arrhythmia Identification

Dr. K.S.Neelu Kumari[a], Santhosh P[b], Selvakumar T[c], Sridhar J[d]

[1]*Associate Professor, ECE, P A College of Engineering and Technology, Pollachi, Tamil Nadu, India*
[2]*Student Member, ECE, P A College of Engineering and Technology, Pollachi, Tamil Nadu, India.*
[3]*Student Member, ECE, P A College of Engineering and Technology, Pollachi, Tamil Nadu, India.*
[4]*Student Member, ECE, P A College of Engineering and Technology, Pollachi, Tamil Nadu, India.*

A B S T R A C T

Effective real-time detection of arrhythmia, a serious cardiac disease, is necessary to allow for prompt medical intervention. In order to diagnose arrhythmia from ECG signals, this research presents a unique, power-efficient VLSI design that integrates a deep neural network (DNN). The design consumes just 11.098 µW at 25 kHz and makes use of 28-nm bulk CMOS technology. Its classification accuracy in class-oriented evaluation is 92.37%, while in subject-oriented evaluation, it is 81.60%. The system's applicability for portable diagnostic systems is demonstrated by its optimisation to concentrate on low-frequency atrial activity in order to increase detection accuracy.

Keywords: Arrhythmia detection, VLSI architecture, deep neural network, ECG classification, low-power design

1. INTRODUCTION

Millions of people worldwide suffer with arrhythmia, which is a major cause of cardiovascular issues. It is typified by erratic cardiac rhythms, which are frequently identified by changes in ECG readings. Reducing related mortality rates, particularly in older populations, depends on accurate and timely arrhythmia identification. Nevertheless, there are issues with current approaches, including their high processing costs, noise sensitivity, and limited suitability for wearable and portable devices.

Conventional arrhythmia detection techniques mostly use manually created variables including heart rate variability (HRV), power spectral density, and RR intervals. Although these methods provide accuracy in certain situations, they frequently aren't able to adjust to the noise and unpredictability of ECG data obtained in practical settings. Furthermore, these systems are not appropriate for portable applications because they typically require external processing units and significant power.

A low-power VLSI design for arrhythmia classification is proposed in this research to overcome these constraints. The architecture ensures good accuracy while lowering computational complexity by combining an optimized deep neural network classifier with a reduced wavelet-based feature extraction procedure. This approach is specifically tailored for wearable devices, with emphasis on low power consumption, scalability, and robustness against signal noise.

The efficiency of the suggested method is demonstrated in both class-oriented and subject-oriented classification schemes through evaluation utilizing standard ECG datasets. The work advances portable diagnostic technologies and makes real-time health monitoring possible by utilizing modular architecture and energy-efficient hardware design.

2. RELATED WORKS

Numerous machine learning and signal processing techniques have been used in extensive research on arrhythmia detection systems. Early research concentrated on extracting ECG features using power spectral density or Fourier transform techniques. Nevertheless, these methods were susceptible to signal artefacts and frequently required large amounts of processing power.

Wavelet transform-based techniques, such the discrete wavelet transform (DWT), have been developed recently to extract multi-resolution characteristics from ECG data. For the classification of arrhythmias, researchers have also used artificial neural networks (ANNs) and support vector machines (SVMs) are examples of machine learning models.

Although these techniques showed increased precision, they frequently used complicated designs and floating-point arithmetic, which resulted in excessive power consumption.

Classification performance has been greatly enhanced by application of deep learning models, including convolutional neural networks (CNNs) and recurrent neural networks (RNNs). However, these approaches remain computationally demanding and unsuitable for real-time use on wearables with minimal power consumption.

The proposed architecture builds upon this foundation by addressing the limitations of previous methods. Unlike existing designs, it employs an integer Haar wavelet transform for efficient feature extraction and uses a lightweight DNN with optimized fixed-point arithmetic, ensuring both accuracy and power efficiency

2.1 Traditional Methods

Early approaches relied on Fourier transform and spectral density analysis to process ECG signals. These methods focused on extracting frequency-domain features to differentiate arrhythmias. However, they lacked temporal resolution, making them unsuitable for detecting transient cardiac abnormalities.

2.2 Wavelet Transform-Based Methods

Wavelet transform techniques addressed many limitations of Fourier-based methods by providing time-frequency localization. Discrete wavelet transform (DWT) emerged as a prominent tool for ECG analysis due to its ability to decompose signals into various frequency bands. Researchers utilized wavelet coefficients derived from Daubechies and Haar wavelets to classify arrhythmias. However, these methods often required floating-point arithmetic, leading to increased hardware complexity and power consumption.

2.3 Machine Learning and Deep Learning

Machine learning models, such as SVMs and decision trees, offered significant improvements in arrhythmia classification by learning patterns from ECG features. However, they required extensive preprocessing and were prone to overfitting in subject-oriented evaluations.

Deep learning models like CNNs and RNNs introduced end-to-end learning capabilities, eliminating the need for handcrafted features. Despite their accuracy, these models posed challenges in terms of computational overhead and energy efficiency, making them impractical for wearable devices.

2.4 Limitations of Existing Systems

Most existing systems focus on achieving high accuracy but fail to address power and computational constraints critical for wearable devices. Additionally, reliance on external preprocessing units limits their scalability and integration into portable healthcare systems

3. PROPOSED METHODOLOGY

The proposed methodology integrates as shown in figure 1 an efficient feature extraction module with a deep neural network (DNN) classifier, optimized for real-time arrhythmia detection with minimal power consumption. The feature extraction leverages the Discrete Wavelet Transform (DWT) with integer Haar wavelets to decompose ECG signals into multiple frequency bands.

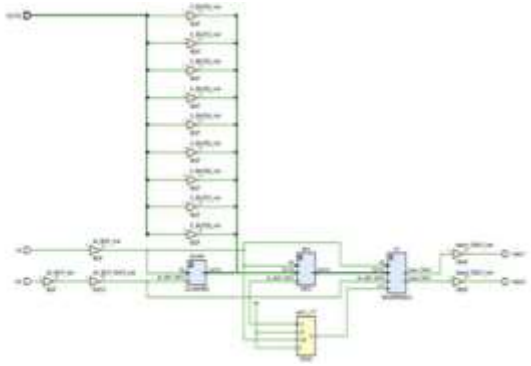

Figure 1: Schematic diagram for Proposed Architecture

By focusing on low-frequency components (below 12 Hz), which are crucial for identifying arrhythmias, the system reduces the influence of high-frequency noise and minimizes computational requirements. Unlike traditional DWT implementations, the proposed design simplifies computations by using basic operations such as addition, subtraction, and bit- Overall, this methodology achieves a balance between computational efficiency, accuracy,

and energy consumption, making it ideal for portable health-monitoring systems.shifting, eliminating the need for floating-point arithmetic and complex filter banks

A. MODIFIED MALLAT WAVELET TRANSFORM

Continuous Wavelet Transform (CWT): This wavelet transform shows the signal as a combination of wavelets at different scales (duration or time) and places (position or amplitude). Continuous time signals are usually analysed using the continuous wavelet transform. Different values of the scaling and translating factors have been used to produce the waveform of the wavelet coefficients.

In mathematical terms, for a signal x(t), the continuous wavelet transform is articulated as

$$\text{CWT (a, b)} = \frac{1}{\sqrt{a}} \int_{-\infty}^{\infty} x(t)\, g\left(\frac{t-b}{a}\right) dt \qquad (1)$$

where the mother wavelet is indicated by g (·). In the same way, a stands for the scale factor and b for the translation factor. In the continuous wavelet transform, a and b both fluctuate continuously. The discrete wavelet transform, which will be covered in the next paragraph, has been developed to remove the duplication brought about by continuous coefficients.

A discretised signal is divided into different resolution levels using the Discrete Wavelet Transform (DWT). The Continuous Wavelet Transform's (CWT) notable redundancy is reduced by the DWT. The scaling function in multiresolution analysis (MRA) yields the approximation coefficients of the decomposed signal, whereas the wavelet function generates the detail coefficients of the fractured signal. The parent wavelet, g, can be used to represent the DWT.

$$\text{DWT (m, k)} = \frac{1}{\sqrt{a_0^m}} \sum_n x(n)\, g\left(\frac{k - b_0\, a_0^m}{a_0^m}\right) \qquad (2)$$

where an integer is represented by k. There is a discrete shift in the translation parameters a and b as well as the scaling parameter. The temporal signal S[n] is divided into the smoothed c1(n) and detailed d1(n) using high-pass (h(n)) and low-pass (l(n)) filters. Consequently, the smooth version c1(n) has less high frequency components than the detail version.

Mathematically, they are specified as

$$c_1 = \sum_k h(k - 2n)\, c_0(k) \qquad (3)$$

$$d_1 = \sum_k g(k - 2n)\, c_0(k) \qquad (4)$$

where the discretised time signal (sampled form of S0(n)) is denoted by c0(n). The DWT coefficients are obtained by downsampling the output of the two filters by a factor of two. The low pass filter's output is called the approximation coefficients, and the high pass filter's output is called the detail coefficients. The process is then repeated once the approximation coefficients have been transmitted to the low pass and high pass filters. The low pass and high pass filters are known as the "Quadrature mirror filters" and are connected by the equation.

$$\text{h [L – 1 – n]} = (-1)^n l(n) \qquad (5)$$

Where, L denotes the length of the filter. The fundamental block diagram of the DWT is illustrated in Figure 2. The implementation of DWT is restricted with the length of signals. Similarly, the coefficients are affected by the change of initial point. The MODWT has been put into action to address the sensitivity of DWT concerning the selection of the starting point of a time series.

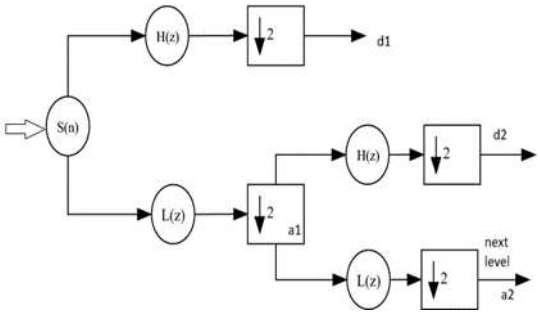

Figure 2: DWT decomposition Flow Diagram

Maximum Overlapping Discrete Wavelet Transform (MODWT) offers greater flexibility than traditional DWT by allowing a freely chosen starting point. This enables better adaptability to different signal characteristics and improved time-frequency localization, making it suitable for various applications like signal and image processing. The orthogonal transform of DWT faces issues due to the lack of invariance to translation in time series analysis. The Maximum Overlap Discrete Wavelet Transform represents an improved version of the Discrete Wavelet Transform (DWT). This transform is applicable to any sample size, while the DWT is restricted to signal lengths N that must be an integer multiple of 2j where j = 1,2,3,. . . ,J indicates the scale number. The depiction of MODWT is illustrated in Fig. 2.The MODWT scaling filters hl and wavelet filters gl are connected to the DWT filters via (6) and (7).

$$\widetilde{h_j} = \frac{h_j}{\sqrt{2}} \qquad (6)$$

$$\widetilde{g_j} = \frac{g_j}{\sqrt{2}} \qquad (7)$$

The MODWT filters are also the Quadrature mirrors like DWT filters is given as

$\widetilde{h_l} = (-1)^{i+1}\, h_{L-1-i}$ (8)

$\widetilde{g_l} = (-1)^{i+1}\, g_{L-1-i}$ (9)

where L - 1 and LL are the filter's widths and l=0,1,2,…L−1l = 0, 1, 2, l dots. By normalising the scaling filters (h_j)˜ and wavelet filters (g_j)˜ with √ 2, the filters (h_i)˜ and (g_i)˜ are obtained from the inference of equations (7) and (8). By keeping the characteristics listed in these equations, this connection guarantees that the MODWT filters are likewise Quadrature Mirror Filters. The mathematical formulas for determining the nn-th element of the first-stage wavelet coefficients and scaling coefficients of the MODWT, respectively, given an input time series signal X(n), are represented by equations (10) and (11). The signal is broken down into its many frequency components using these equations, which are essential to the MODWT process.

$\widetilde{W_{1,n}} = \sum_{l=0}^{L_1-1} \widetilde{h_l}\, X_{n-l \bmod N}$ (10)

$\widetilde{V_{1,n}} = \sum_{l=0}^{L_2-1} \widetilde{g_l}\, X_{n-l \bmod N}$ (11)

$\varphi_{k,n} = \sum_{l=0}^{L_1} \widetilde{\varphi_l}\, \tilde{\varphi}_{l,n+l \bmod N}$ (12)

B. MODIFIED APPROXIMATE MULTIPLIER

Designs for approximate multipliers Numerous methods for precisely computing multipliers have been put forth. In this study, we address candidate circuit limitations using Cartesian Genetic Programming (CGP). This method has created high-quality approximation circuits in the past and is multi-objective. CGP is a kind of genetic programming that uses directed acyclic networks to describe potential architectures. A two-dimensional array of programmable nodes with nc columns and nr rows is used to represent a candidate circuit. The nodes in our example will be 2-input Boolean functions, and the collection of accessible functions will be represented by Γ. There are no primary outputs or primary inputs on the circuit. Connections for feedback are turned off. Labelled (0,1...nc·nr+ni−1), the nodes' primary inputs and outputs act as addresses to link the inputs. With nr · nc triplets (x1, x2, ψ) defining the function ψ (ψ ∈ Γ) for every node and its input connections, the chromosome, which functions as a netlist, displays a suggested solution. The final section of the chromosome that shows the nodes attached to the primary outputs is devoid of numbers. While the circuit size varies according to the number of active nodes, the chromosomal size (s) stays constant at s = ncnr (na + 1) + no. One such instance is seen in Figure 3.

The whole search space is covered by the valid chromosomes (netlists).

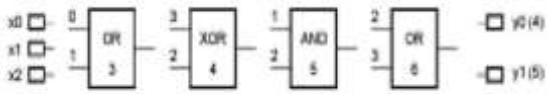

Figure 3: Example of a circuit in CGP with parameter

ni = 3, no = 2, nc = 4, nr = 1; Γ = {0and,1or,2xor}. Chromosome numbers: 0, 1, 1; 3, 2, 2; 1, 2, 0; 2, 3, 1; 4, 5. Gate 6 is currently not in use. The circuit's logical operation is as follows: y0 = ((x0 or x1) xor x2), y1 = x1 and x2. CGP has an easy search methodology. In this case, the starting population P of CGP consists of one of many exact multiplier implementations as well as a few circuits generated via precise multiplier mutation. The precise multiplier in the original population may be determined easily since multiplier netlists and CGP chromosomes have a one-to-one connection. The next step is to evaluate potential circuits using the fitness function. Each member of P then receives the so-called fitness score, and the highest-scoring individual becomes a new parent for the next generation. Mutation creates λ candidate solutions from the parent solution.

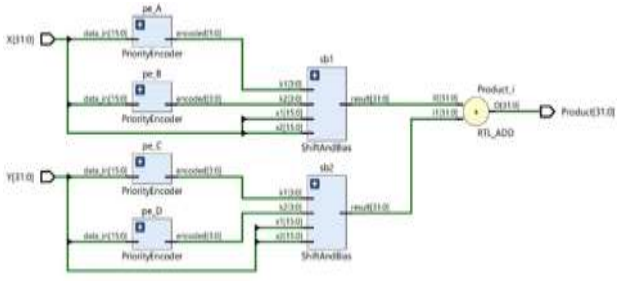

Figure 4: Schematic view of Modified Approximate Multiplier

The stopping point is defined by the number of iterations. Despite repeated attempts to bring an adequate crossover operator to CGP, mutation remains the primary genetic operator. The mutation operator modifies up to h randomly chosen chromosomal genes (integers). Their revised values are generated at random; nonetheless, the new values are proved to be valid. A single mutation can alter the gate function, gate input link, or primary output link. To estimate multipliers, use fitness(M) = −CM and Γ = {NAND, NOR, XNOR, AND, OR, XOR, NOT, identity}.

CGP will be used to generate 7 and 11-bit unsigned multipliers. The sign extension, which includes 8 bit and 12-bit multipliers, will be built manually using the one's complement approach. The maximum permitted arithmetic error ε of approximation multipliers is chosen from the collection {0.5%, 1%,

2%, 5%, 10%, 15%, 20%}. Because classification accuracy declined significantly, we did not employ arithmetic error more than 20% for approximation multipliers. The approximation process begins with exact multipliers, which comprise the so-called starting population. There are 84 beginning configurations in all, including two-bit widths, 7 target mistakes, and 6 types of initial multiplier designs.

4. RESULTS AND DISCUSSION

The proposed system is evaluated using ECG datasets from the PhysioNet MIT-BIH Arrhythmia Database. Two evaluation schemes—class-oriented and subject-oriented—are used to assess performance.

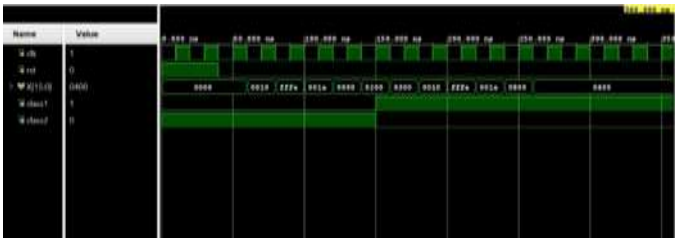

Figure 5: Digital waveform simulation

A. Classification Accuracy

- Class-Oriented Evaluation: The classifier achieves an accuracy of 92.37%, demonstrating its ability to differentiate between arrhythmia and normal heartbeats across multiple subjects. This metric highlights the robustness of the model when trained on a diverse dataset.
- Subject-Oriented Evaluation: The system achieves an accuracy of 81.60%, reflecting its adaptability to inter-patient variations. This evaluation is critical for real-world applications where models encounter unseen patient data.

B. Power Consumption

A critical aspect of the proposed architecture is its ultra-low power consumption. Operating at a frequency of 25 kHz, the system consumes just 11.098 µW. This power efficiency is achieved through various design optimizations, including the use of fixed-point arithmetic, Mitchell's logarithmic approximation for multiplication, and the elimination of complex filter banks in the wavelet transform. The power savings make the system suitable for wearable devices, where energy constraints are a primary consideration.

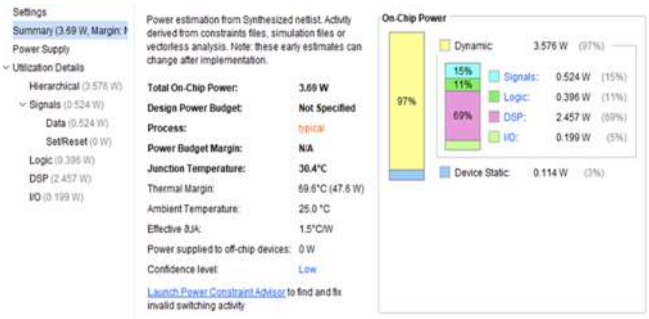

Figure 6: Power Estimation Report

Furthermore, the system's power efficiency does not come at the expense of accuracy, as demonstrated by the high classification performance.

C. Comparison with State-of-the-Art

Compared to existing methods, the proposed architecture demonstrates superior performance in terms of both accuracy and power efficiency. For example:

- Traditional Fourier transform-based systems consume significantly more power due to their reliance on floating-point operations.
- Advanced neural network models, such as CNNs, achieve high accuracy but are computationally expensive and unsuitable for real-time, low-power applications.

D. Scalability and Practical Implementations

The modular design allows for easy adaptation to other types of cardiac abnormalities. By focusing on low-frequency components, the architecture is inherently robust to high-frequency noise, making it suitable for real-world deployment. The implementation in 28-nm CMOS technology further ensures compatibility with modern wearable device platforms.

5. CONCLUSION AND FUTURE WORK

This paper presents a scalable and energy-efficient VLSI architecture for arrhythmia detection, specifically tailored for wearable health-monitoring applications. By leveraging a simplified Discrete Wavelet Transform (DWT) and an optimized Deep Neural Network (DNN), the system achieves a balance between accuracy and power efficiency. Implemented in 28-nm bulk CMOS technology, the architecture consumes just 11.098 μW while maintaining a classification accuracy of 92.37% for class-oriented evaluation and 81.60% for subject-oriented evaluation.

The proposed design addresses key limitations of existing methods, such as high computational complexity, reliance on floating-point arithmetic, and sensitivity to signal noise. Its modular and filter-less architecture ensures scalability for detecting various cardiac abnormalities, making it a suitable candidate for real-time health-monitoring systems.

A. Future Work

Future research will focus on the following enhancements:

- Integration of Advanced Neural Networks: Incorporating convolutional or recurrent neural networks to improve classification accuracy for more complex cardiac conditions.
- Noise-Resilient Preprocessing: Developing algorithms to handle real-world noise and motion artifacts in wearable devices.
- Hardware Implementation: Testing the design on physical ASICs and extending it to multi-arrhythmia detection systems.
- Broader Applications: Adapting the architecture for other biomedical signals, such as EEG, to detect neurological conditions.

References

[1] Y. Lei, S. K. Yeung, and A. M. K. Cheng,"DNN-Based Low Power ECG Co-Processor Architecture to Classify Cardiac Arrhythmia," IEEE Transactions on Biomedical Circuits and Systems, vol. 14, no. 5, pp. 959–970, Oct. 2024.

[2] M. Al Rahhal, Y. Bazi, H. Al Hichri, et al., "Deep Learning Approach for Active Classification of Atrial Fibrillation in Wearable Devices,"IEEE Transactions on Biomedical Circuits and Systems, vol. 14, no. 5, pp. 959–970, Oct. 2020. DOI: 10.1109/TBCAS.2020.3011112.

[3] Liu, S. Liu, C. Hu, et al., "Low-Power ECG Signal Processing on Wearable VLSI Devices, "IEEE Journal of Emerging and Selected Topics in Circuits and Systems, vol. 12, no. 2, pp. 287–298, June 2022. DOI: 10.1109/JETCAS.2022.3159821.

[4] X. Zhang, H. Wang, and Y. Sun, "A Lightweight DNN-Based ASIC for Atrial Fibrillation Detection," IEEE Transactions on Circuits and Systems I: Regular Papers, vol. 69, no. 8, pp. 3123–3134, Aug. 2022. DOI: 10.1109/TCSI.2022.3181789.

[5] J. Kim, K. Kim, and H. Yoo, "Energy-Efficient DNN Accelerator for ECG-Based Atrial Fibrillation Monitoring," IEEE Transactions on Very Large-Scale Integration (VLSI) Systems, vol. 28, no. 4, pp. 867–878, Apr. 2021. DOI: 10.1109/TVLSI.2021.3062205.

[6] P. Chen, Y. Zhang, X. Ren, et al., "VLSI Design Techniques for Low-Power Atrial Fibrillation Detection," IEEE Transactions on Biomedical Engineering, vol. 67, no. 11, pp. 3168–3178, Nov. 2020. DOI: 10.1109/TBME.2020.2993057.

[7] S. Yin, Q. Wang, J. Jiang, et al., "Real-Time Atrial Fibrillation Detection Using RNN-Based Models in Wearable Systems," IEEE Access, vol. 8, pp. 210897–210906, Dec. 2020. DOI: 10.1109/ACCESS.2020.3044292.

[8] H. Li, J. Han, and Y. Zhou, "Approximate Computing for DNN Hardware Accelerators in ECG Analysis," IEEE Transactions on Circuits and Systems I: Regular Papers, vol. 69, no. 5, pp. 1805–1816, May 2022. DOI: 10.1109/TCSI.2022.3147721.

[9] M. Wang, Y. Yu, and L. Wang, "Low-Power Embedded Systems for Cardiovascular Monitoring with DNNs," IEEE Internet of Things Journal, vol. 8, no. 4, pp. 2572–2581, Feb. 2021. DOI: 10.1109/JIOT.2020.3018384.

[10] Gao, C. Xu, and X. Liu, "Hardware-Friendly DNN Architectures for On-Chip ECG Classification," IEEE Transactions on Neural Networks and Learning Systems, vol. 33, no. 3, pp. 1247–1258, Mar. 2022. DOI: 10.1109/TNNLS.2022.3131802.

[11] R. Srinivasan, M. Kamath, and P. Saxena, "ECG Signal Analysis with Lightweight Deep Learning Models," IEEE Transactions on Medical Imaging, vol. 41, no. 7, pp. 1953–1963, July 2022. DOI: 10.1109/TMI.2022.3142236.

[12] A. Habib, Y. Lei, and C. Xu, "Power-Aware DNN Architectures for Atrial Fibrillation Detection in Edge Devices," IEEE Design & Test, vol. 39, no. 1, pp. 68–76, Jan. 2023. DOI: 10.1109/MDAT.2022.3148901.

[13] N. Kapoor, R. Kumar, and S. Jain, "Wearable ECG Monitoring with Low-Power Convolutional Neural Networks," IEEE Embedded Systems Letters, vol. 14, no. 1, pp. 12–15, Mar. 2022. DOI: 10.1109/LES.2022.3156784.

[14] Lee, J. Park, and K. Lee, "Efficient Quantized Neural Networks for Atrial Fibrillation Detection," IEEE Journal of Biomedical and Health Informatics, vol. 25, no. 5, pp. 1748–1757, May 2021. DOI: 10.1109/JBHI.2021.3063289.

[15] Y. Wang, H. Qiu, and F. Li, "TinyML-Based ECG Analysis for Wearable Applications," IEEE Internet of Things Journal, vol. 9, no. 2, pp. 1085–1093, Jan. 2022. DOI: 10.1109/JIOT.2021.3062243.

[16] Chen, M. Zhou, and L. Fang, "Deep-Learning-Driven VLSI Designs for Cardiac Signal Processing," IEEE Transactions on Computational Biology and Bioinformatics, vol. 19, no. 4, pp. 1353–1365, Apr. 2022. DOI: 10.1109/TCBB.2022.3170801.

[17] P. Kirchhof et al., "2016 ESC Guidelines for the Management of Atrial Fibrillation Developed in Collaboration with EACTS," Kardiologia Polska (Polish Heart Journal), vol. 74, no. 12, pp. 1359–1469, 2016.

[18] S. Nattel, "New Ideas About Atrial Fibrillation 50 Years On," Nature, vol. 415, no. 6868, pp. 219–226, Jan. 2002.

[19] J. Oldgren et al., "Variations in Cause and Management of Atrial Fibrillation in a Prospective Registry of 15,400 Emergency Department Patients in 46 Countries: The RE-LY Atrial Fibrillation Registry," Circulation, vol. 129, no. 15, pp. 1568–1576, Apr. 2014.

[20] P. Langley, J. P. Bourke, and A. Murray, "Frequency Analysis of Atrial Fibrillation," in Proceedings of Computers in Cardiology, vol. 27, 2000, pp. 65–68.

[21] R. Couceiro, P. Carvalho, J. Henriques, M. Antunes, M. Harris, and J. Habetha, "Detection of Atrial Fibrillation Using Model-Based ECG Analysis," in Proceedings of the 19th International Conference on Pattern Recognition, Dec. 2008, pp. 1–5.

[22] N. Sasaki et al., "Frequency Analysis of Atrial Fibrillation from Specific ECG Leads V7–V9: A Lower DF in Lead V9 Is a Marker of Potential Atrial Remodeling," Journal of Cardiology, vol. 66, no. 5, pp. 388–394, 2015.

[23] M. Zihlmann, D. Perekrestenko, and M. Tschannen, "Convolutional Recurrent Neural Networks for Electrocardiogram Classification," in Proceedings of the Computing in Cardiology Conference (CinC), Sep. 2017, pp. 1–4.

[24] S. Ladavich and B. Ghoraani, "Rate-Independent Detection of Atrial Fibrillation by Statistical Modeling of Atrial Activity," Biomedical Signal Processing and Control, vol. 18, pp. 274–281, Apr. 2015.

[25] E. Lake and J. R. Moorman, "Accurate Estimation of Entropy in Very Short Physiological Time Series: The Problem of Atrial Fibrillation Detection in Implanted Ventricular Devices," American Journal of Physiology-Heart and Circulatory Physiology, vol. 300, no. 1, pp. H319–H325, Jan. 2011