Heart Disease Detection Architecture for Lead I Off-the-Person ECG Monitoring Devices

Pedro Sá
INESC-ID,
Instituto Superior Técnico,
Universidade de Lisboa
Lisbon, Portugal
pbritosa@tecnico.ulisboa.pt

Helena Aidos
Instituto de Telecomunicações &
LASIGE, Faculdade de Ciências,
Universidade de Lisboa
Lisbon, Portugal
haidos@fc.ul.pt

Nuno Roma
INESC-ID,
Instituto Superior Técnico,
Universidade de Lisboa
Lisbon, Portugal
nuno.roma@inesc-id.pt

Pedro Tomás
INESC-ID,
Instituto Superior Técnico,
Universidade de Lisboa
Lisbon, Portugal
pedro.tomas@inesc-id.pt

***Abstract*—With the rise of smart-watches and other wearables, off-the-person electrocardiography is gaining momentum as high-quality Lead 1 ECG signals can now be acquired from a persons hands or arms. Although several heart disease detection algorithms have been described in recent years, they are not designed considering Lead 1-only setups. This work bridges this gap with an architecture for a robust Lead 1 real-time heart disease detection system and an FPGA-based implementation. The proposed system is based on a signal processing pipeline composed of: ECG signal denoising; heartbeat detection and segmentation; extraction of dynamic morphological features; and heartbeat classification (standard and different abnormal heartbeats). Resorting to the only database from MITs Physiobank with Lead 1 annotated recordings, InCarTDb, the proposed pipeline resulted in a 4-class model with a classification accuracy of up to 96.5%. Moreover, when implemented in a Zynq-7 ZC702 Evaluation Board, the proposed architecture requires less than 30% of the FPGA resources and a total power consumption of 192 mW at a clock frequency of 35 MHz.**

***Index Terms*—ECG analysis, cardiac pathology identification, hardware architecture, real-time processing**

I. Introduction

Cardiovascular diseases account for the death of more than 17.3 million people per year around the world [1]. To diagnose multiple heart conditions associated with this pathology, the electrocardiogram (ECG) is still the most used tool, being most effective in the detection of arrhythmias and myocardial infarctions. Typically, the signal acquisition is made with a 12-lead device, which relies on ten adhesive sensors on the subject's torso and limbs. However, since this setup is impractical for daily life person monitoring, off-the-person ECG processing systems, supported on a smaller number of leads, are required. The most practical off-the-person setup consists of a single-lead view, which still provides relevant information on the subject's biometrics and allows the detection of many pathologies [2]. Aside from the non-intrusiveness and easiness of use, contact-based single-lead acquisitions also imply a smaller amount of data. Therefore, the required computational and processing power of the system decreases, making current embedded platforms a highly viable alternative for ECG processing in real-time environments. In this scenario, designing a system for Lead 1-only processing (acquired from both hands/arms) fits the off-the-person paradigm [3], as the development and adoption of wearable technologies including smart watches and bands is becoming gradually established in several application domains.

Over the years, the research on automatic ECG-based heartbeat classification methods have established (see [4], [5] and references therein). Most works make use of morphological features (e.g. obtained by applying a Discrete Wavelet Transform (DWT) [6]), eventually combined with dynamic features (the RR interval information) [7]. In [6], independent component analysis (ICA), or principal component analysis (PCA) were also applied for feature extraction, whereas in [8] a combination of PCA and linear discriminant analysis (LDA) was used. In [9], redundant features are eliminated through a genetic algorithm. For the classification of heartbeats, support vector machines (SVM) [6], [9], probabilistic neural networks (PNN) [6], [8], [9], bagging tree classifiers [10], or ensemble of classifiers [7] have been used. Furthermore, a semi-supervised approach based on consensus clustering can be used to identify abnormal heartbeats [11].

Although the majority of these works focus on multi-lead setups, single-lead ones are also emerging [9], [10]. However, they are not based on lead 1 setups, hence demanding more complex acquisition systems, with requirements for sensors on the wrists and legs (lead II) or on the torso (lead IV). In contrast, the setup that is now presented makes use of an alternative approach that facilitates off-the-person monitoring [3]. It gathers the information directly from the wrists or hands, which significantly simplifies the acquisition process and makes it easier for integration with wearable devices.

Hence, given the increasing interest in health monitoring devices, several systems have been proposed for real-time detection of heart conditions. This includes the detection of simple arrhythmia based on the heart rate or other straightforward features [12]–[14]; the implementation of more complex algorithms to detect P and T waves in real-time [15], or ST segment elevation/depression [16]. Approaches relying

This work was partially supported by Portuguese national funds through Fundação para a Ciência e a Tecnologia (FCT) under projects with reference UID/CEC/50021/2019 (INESC-ID), UID/EEA/50008/2019 (Instituto de Telecomunicações), UID/CEC/00408/2019 (LASIGE), and PTDC/EEI-HAC/30485/2017 (HAnDLE).

978-9-0827-9703-9/19/$31.00 ©2019 IEEE

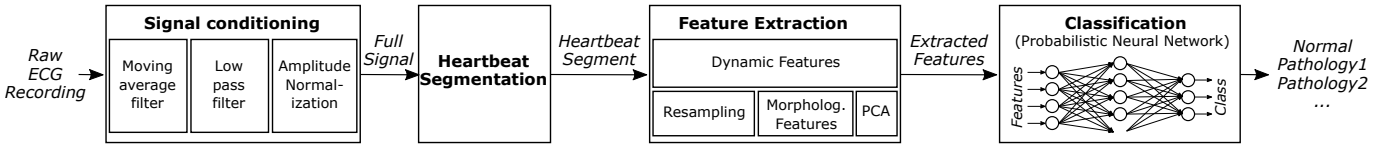

Fig. 1. Proposed lead I-only heart disease processing pipeline.

on neural networks [17], [18] or on fuzzy clustering based algorithms [19] have also been proposed. However, most works either do not fully implement a classification system or rely on software-based solutions (e.g., [16], [19]) or focus on the detection of a small subset of heart conditions [20]. Although [21] proposes an end-to-end architecture, from pre-processing to heart condition detection, they target modified lead II acquisition systems that send raw information to cloud servers, instead of locally processing the information.

In contrast, this manuscript proposes a novel system architecture that was specifically designed to be integrated on embedded devices and smart systems (where resource usage and power consumption are fundamental aspects). By following usual off-the-person acquisition setups, it considers the integration on daily-life objects (e.g., drivers wheel), where ECG acquisition must be performed using lead I. The architecture is scalable and programmable, which allows it to be tailored to: different acquisition setups; different parameters in most of the filtering, segmentation, feature extraction and classification stages (by updating the coefficients); and different abnormalities to be classified.

Hence, this paper introduces the following contributions:

1) a scalable architecture to be embedded in off-the-person setups;
2) the processing system parameters can be easily tuned to adapt to specific application needs (acquisition setup, number of leads, filtering parameters, number and type of abnormalities), by feeding the architecture with different configuration parameters;
3) it performs end-to-end ECG analysis, from raw signal pre-processing to classification.

The developed system was prototyped in a Zynq-7 ZC702 Evaluation board and evaluated using an inter-subject approach, showing accuracies of 96.5% for a 4 class setup.

II. Heart disease detection algorithm

The proposed processing methodology is composed of four stages, as depicted in Fig. 1: signal conditioning (pre-processing), segmentation, feature extraction and classification.

A. Signal conditioning

Signal conditioning comprehends three steps: baseline removal, high-frequency noise filtering, and amplitude normalization. First, the signal baseline is removed by subtracting the result of a moving average filter (with a 500 ms window). Then, a low-pass filter is applied to remove high-frequency noise. Since FIR filters have a low stopband attenuation for lower orders and they are computationally more complex than equivalent IIR filters, an IIR filter was designed with the lowest order possible, since the phase shifts do not significantly affect the signal. In particular, a 8-order Chebyshev Type II filter with a 45 Hz cut-off frequency and 80 dB stopband attenuation was used. Finally, since ECGs have a tremendous amplitude variability between patients (and sometimes even for the same patient), and also between acquisition setups, the third step normalizes the signal to keep the peaks around 1 mV. It does so by dividing the values of each 1000 ms window by a normalization parameter, ensuring proper signal conditioning across signal recordings.

B. Heartbeat segmentation

The heartbeat segmentation is performed using two techniques that showed to provide accurate results when compared with training database annotations: (i) detection of the signal peak (maximum value); followed by (ii) cross-correlation with a set of previously stored templates. Hence, the ECG signal is split into segments of 625 ms, corresponding to 250 ms before the detected peak and 375 ms after.

C. Feature extraction

Feature extraction comprehends dynamic and morphologic features. The former are obtained from the RR intervals of the heartbeat. From a given heartbeat, an average rhythm and a long-term average RR interval are computed by averaging all the RR intervals within a sliding window of the past 10 seconds and 5 minutes, respectively, from the heartbeat.

Afterwards, the segmented heartbeat is resampled (i.e., stretched or compressed, depending on the heartbeat rate) in order to have one, and only one, P-QRS-T wave inside the segment window. The resampling rates were experimentally assessed, giving rise to three intervals: (i) for a heart rate below 60 beats per minute (bpm), a resampling rate of 1.2 is employed, stretching the original signal; for a heart rate in the interval [60;133] bpm, no resampling is performed; for a heart rate above 133 bpm, a resampling rate of 0.8 is employed.

The morphological features are extracted from the 3rd and 4th order details and the 4th order approximation coefficients obtained from the application of the DWT with Daubechies 8th-order wavelet function (DB8), followed by a dimensionality reduction implemented with a PCA.

D. Classification

Several algorithms were evaluated to implement the pathology classification, including: k-Nearest Neighbors (kNN), Decision Trees, Support Vector Machines (SVMs) and Artificial

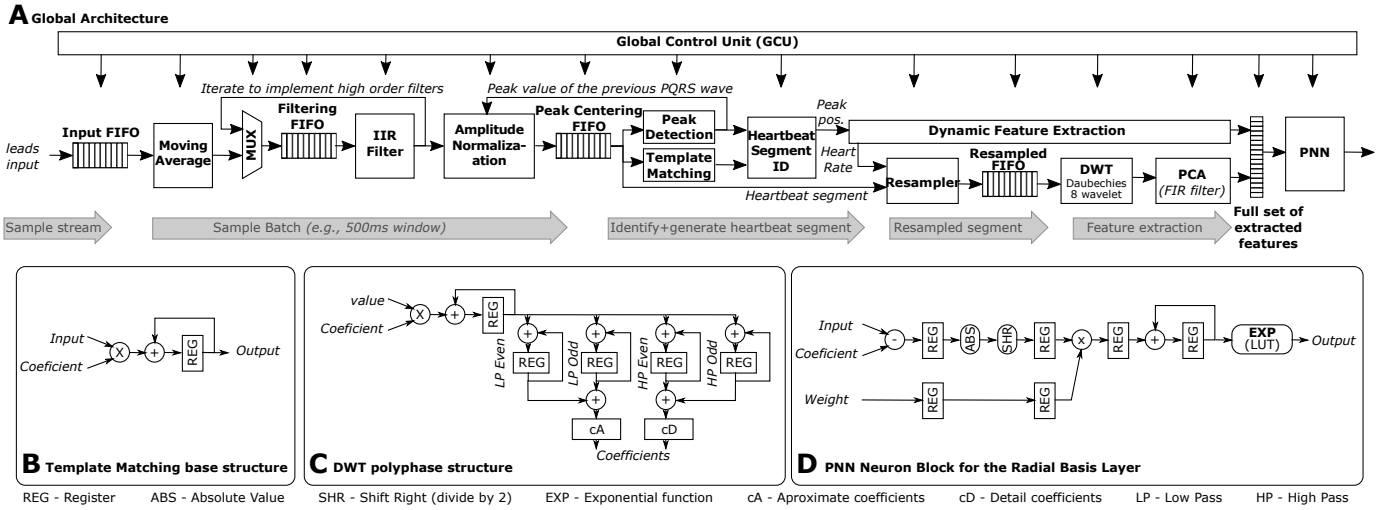

Fig. 2. End-to-end architecture of the implemented circuit.

Neural Networks (ANNs). Experimental evaluation allowed to conclude that ANNs provide the best classification scores, when using Probabilitic Neural Networks (PNNs). Hence, a PNN was designed that is composed of a set of three layers: input layer, radial basis function (rbf) and competitive layer.

III. Hardware implementation

The conceived architecture, illustrated in Fig. 2A, is globally commanded by a Global Control Unit (GCU), which manages the flow of data between each Processing Unit (PU) by means of a set of specifically designed flags. Hence, while flags flowing from the PUs to the GCU indicate the dataflow status (e.g., the PU has finished processing the input data), flags on the reverse direction command the PUs to start processing a new block. The GCU is implemented with a sequential state-machine using a hierarchical topology: each PU is controlled by a local state-machine, with the synchronization between machines being implemented at a global level.

Since the ECG sampling frequency is usually low, a multi-cycle structure was considered and the architecture of each PU was specifically designed to operate at the lowest frequency that still satisfies the aimed processing throughput. Contrasting to an equivalent pipelined architecture, this allows reducing the hardware resource utilization, while still allowing to process lead-I signals in real-time, even when sampled at high rates. All arithmetic operations (as well as input coefficients) use 16-bit fixed-point representations, as no significant difference in accuracy was observed when using a higher number of bits. The complete architecture was also specifically designed to support changes in the parameters and in the number of output classes, by feeding the design with different parameters and by changing the coefficient stored in memory.

The overall processing flow is as follows. The input sample stream (i.e., ECG lead signal, converted from the analog to the digital domain) is initially buffered by means of a FIFO memory block. Data processing is then performed in sample batches of 500 ms window (although the architecture supports windows of different sizes, this value was considered appropriate to the characteristics of the ECG signal). Each sample batch is then passed through the Moving Average filter and stored on another FIFO block.

To implement the pre-pocessing low-pass filtering stage, an IIR Filter Unit was specifically designed to implement a quadratic filtering section in direct form II. Hence, to support the implementation of high-order filters, the sample batch can flow through the IIR Filter Unit as many times as needed (4 times for the 8-order Chebyshev filter).

Amplitude Normalization would normally require looping through the sample batch twice: the first to compute the maximum value, and the second to implement the normalization. However, since the EEG signal is quasi-stationary, the normalization of each sample batch n is performed by considering the maximum value measured by the Peak Detection Unit when processing the previous batch, $n-1$.

The implementation of the Morphological Feature Extraction requires the ECG signal to be aligned according with the QRS complex. Two units collaborate to attain this goal: the Peak Detection Unit, which is implemented as a simple register (holds maximum value) and comparator; and the Template Matching Unit, which implements a cross-correlation between the sample batch and the disease templates. Since the templates are pre-computed when training the algorithm and remain constant thereafter, the Template Matching Unit is actually implemented as a Finite Impulse Response (FIR) filter. To minimize the resource usage, the operations are folded into a multiply-and-accumulate (MAC) unit that handles the actual computation (see Fig. 2B).

The Dynamic Feature Extraction Unit computes both the instantaneous and mean heartbeat rates (required for tachycardia identification). Computation of instantaneous heart rate is performed by counting the number of samples between successive hearbeat peaks. Average heart rates are computed by counting the total number of threshold crossings in fixed

time windows, followed by time window size normalization (multiplication with a constant term).

By taking the instantaneous heart rate as input, the Resampler Unit performs sample interpolation whenever the heart rate is below 60 bpm or above 133 bpm. This is achieved through a MAC-based interpolation unit, which consumes as many samples from the Peak Centering FIFO as required (depends on resampling factor), to compose a resampled segment.

To extract the morphologic features, a DWT Unit is used, which relies on a polyphase filter design (see Fig. 2C). Dimensionality reduction is attained through Principal Component Analysis. However, the employed eigenvectors are not re-estimated in real-time, but based on a previous training database. Hence, the PCA Unit can be simply implemented as another FIR filter.

The PNN classifier is fed with the final features and outputs the classification of the processed heartbeat. Its architecture is decomposed in a set of homogeneous sub-units, each featuring a multi-stage pipelined architecture (see Fig. 2D), and designed to identify a specific heart condition. All sub-units have access to a specific BRAM block to store input coefficients (including bias). The exponential function of the RBF layer is also implemented through a BRAM, configured as a look-up table (LUT).

IV. Experimental results

A. Dataset and experimental setup

To evaluate the proposed system (algorithm and architecture), the InCarTDb database (from PhysioNet) was selected, since it features Lead 1 recordings [22], the most typical setup in off-the-person ECG analysis. Therefore, in the context of this work, all of the seventy-five 30-minutes recordings at 257 Hz from 32 different patients were used. To train the classification models, four distinct conditions were considered: (i) normal beat; (ii) atrial premature contraction (APC); (iii) premature ventricular contraction (PVC); and (iv) right bundle branch block beat (RBBB). The remaining classes were discarded from this evaluation since the limited number of patients with such conditions invalidated the used cross-validation procedure. Fig. 3 presents average heartbeats segments for these classes.

B. Classification accuracy

To evaluate the heart disease detection algorithm, an inter-subject analysis is performed, with 47% of subjects used for training, the rest for validation. A cross-validation scheme is subsequently applied using 100 runs, each with 100 heartbeats per class. The Association for the Advancement in Medical Instrumentation (AAMI) [23] recommends the use of three metrics to assess the performance of the algorithm. Thus, for each class c, the following metrics are reported:

1) sensitivity, with sens. $= \mathrm{TP}_c/(\mathrm{TP}_c + \mathrm{FN}_c)$;
2) positive predictive value, with PPV $= \mathrm{TP}_c/(\mathrm{TP}_c + \mathrm{FP}_c)$;
3) false positive rate, with FPR $= \mathrm{FP}_c/(\mathrm{FP}_c + \mathrm{TN}_c)$;

TABLE I
HEART DISEASE DETECTION ALGORITHM RESULTS. $Ntemp$ MEANS NUMBER OF TEMPLATES, SENS., PPV, FPR AND ACC. CORRESPOND TO THE EVALUATION METRICS SENSITIVITY, POSITIVE PREDICTIVE VALUE, FALSE POSITIVE RATE AND ACCURACY, RESPECTIVELY.

Classes	$Ntemp$	Sens.(%)	PPV(%)	FPR(%)	Acc.(%)
Healthy	38 130	97.9	97.6	7.7	-
APC	1 543	42.3	50.9	1.3	-
PVC	7 626	100.0	99.4	0.1	-
RBBB	3 045	97.5	94.7	0.4	-
Overall	50 344	-	-	-	96.5

TABLE II
FPGA RESOURCE UTILIZATION, MAXIMUM OPERATING FREQUENCY AND POWER CONSUMPTION FOR A SYSTEM WITH 4 CLASSES, 8 AND 16 CLASSES.

		4 classes	**8 classes**	**16 classes**
Resource	**Avail.**	**Utilization**	**Utilization**	**Utilization**
LUTs	53200	20648 (39%)	24252 (46%)	28371 (53%)
LUTRAM	17400	2176 (13%)	3968 (23%)	7552 (43%)
FFs	106400	16542 (16%)	16905 (16%)	17678 (17%)
BRAMs	140	28 (20%)	28 (20%)	28 (20%)
DSPs	220	10 (5%)	14 (6%)	22 (10%)
Max. clock freq.		35 MHz	35 MHz	35 MHz
Static Power†		122 mW	122 mW	122 mW
Dynamic Power†		70 mW	78 mW	86 mW
Total Power†		192 mW	200 mW	208 mW

† Power consumption at the maximum operating frequency (35MHz).

where TP_c, TN_c, and FP_c represent the number of true positives, true negatives and false positives of class c, respectively.

The accuracy of the algorithm was also computed to allow studying the overall algorithm performance. Since no significant difference is observed between the software and hardware implementations, Table I reports the hardware accuracy for the considered 4 classes. As it can be observed, all the classes present highly accurate results (global accuracy of 96.5%), with the exception of the APC, whose lower accuracy is mostly related with the known difficulty in correctly identifying the P-wave changes using the information from lead I only (compare also the templates presented in Fig. 3).

Although such results are inferior to those obtained with a multi-lead configuration (e.g., on a 12 lead system accuracies of 99.4% can be obtained [10]), this still allows attaining high cardiac pathology detection levels with a much simpler acquisition system, that is particularly suited for portable, low-power and off-the-person approaches.

C. Resource utilization

To evaluate the hardware implementation of the devised architecture, a Zynq-7 ZC702 Evaluation Board was used for prototyping (the ARM cores were not used). Xilinx Vivado 2016.4 was used for architecture synthesis, map, and place&route. Table II presents the resource utilization, maximum clock frequency and power consumption values.

As it can be observed, the proposed architecture can easily scale for over 16 classes (heart conditions). Moreover, it

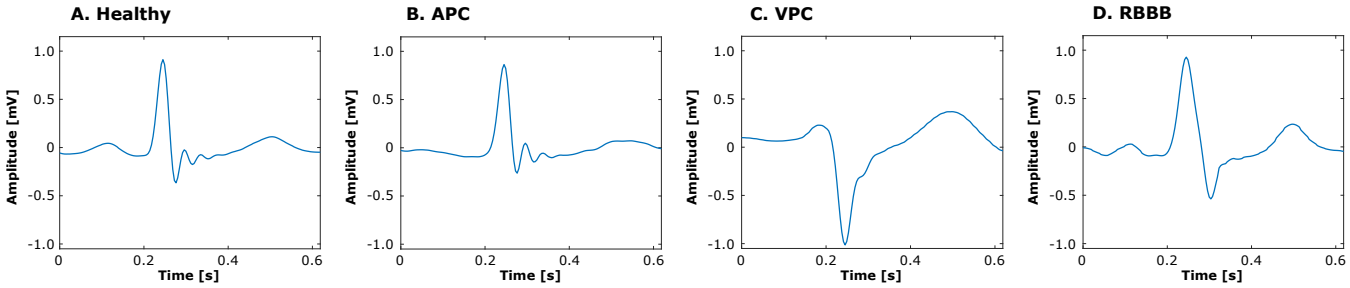

Fig. 3. Average heartbeats segments (after alignment and signal conditioning) for the selected classes.

is able to operate at up to 35 MHz, which would even allow supporting 12 leads with sampling frequencies of up to 20kHz. However, for the particular case of lead I-only analysis (the focus of this paper), with a data sampling rate of 257 Hz, and a maximum heartbeat rate of 180 bpm, the device operating frequency could be substantially reduced (as low as 72kHz), reducing dynamic power consumption to $<$ 1mW (maximum resolution of the Xilinx Power Estimator tool). Such results indicate that the proposed solution is a highly viable approach for integration in application-specific integrated circuits (ASICs).

V. Conclusions

To tackle the difficulties posed by current ECG monitoring and analysis devices in continued, ambulatory and daily life environments, this manuscript proposed a complete system (algorithm and architecture) for end-to-end real-time detection and processing of Lead 1 ECG signals in embedded devices, supporting a wide number of classes (heart conditions) and leads. According to the conducted experimental evaluation in a Zynq-7 ZC702 FPGA device, the developed prototype offers accuracy levels as high as 97% and requires a power consumption below 1mW@72kHz to process a single lead, but being able to scale up to a 12-lead scheme by operating at higher clock rates.

References

[1] A. S. Association *et al.*, "Heart disease, stroke and research statistics at-a-glance," 2016.

[2] H. Silva, C. Carreiras, A. Lourenço, and A. Fred, "Off-the-person electrocardiography," in *Int. Congress on Cardiovascular Technologies*, 2013.

[3] J. Ribeiro Pinto, J. S. Cardoso, and A. Loureno, "Evolution, Current Challenges, and Future Possibilities in ECG Biometrics," *IEEE Access*, vol. 6, pp. 34 746–34 776, 2018.

[4] E. Luz, W. Schwartz, G. Cámara-Chávez, and D. Menotti, "ECG-based heartbeat classification for arrhythmia detection: A survey," *Computer methods and programs in biomedicine*, vol. 127, pp. 144–164, 2016.

[5] S. Berkaya, A. Uysal, E. Gunal, S. Ergin, S. Gunal, and M. Gulmezoglu, "A survey on ECG analysis," *Biomedical Signal Processing and Control*, vol. 43, pp. 216–235, 2018.

[6] R. Martis, U. Acharya, and L. Min, "ECG beat classification using PCA, LDA, ICA and discrete wavelet transform," *Biomedical Signal Processing and Control*, vol. 8, no. 5, pp. 437–448, 2013.

[7] V. Mondéjar-Guerra, J. Novo, J. Rouco, M. Penedo, and M. Ortega, "Heartbeat classification fusing temporal and morphological information of ECGs via ensemble of classifiers," *Biomedical Signal Processing and Control*, vol. 47, pp. 41–48, 2019.

[8] J. Wang, W. Chiang, Y. Hsu, and Y. Yang, "ECG arrhythmia classification using a probabilistic neural network with a feature reduction method," *Neurocomputing*, vol. 116, pp. 38–45, 2013.

[9] P. Pławiak, "Novel methodology of cardiac health recognition based on ECG signals and evolutionary-neural system," *Expert Systems with Applications*, vol. 92, pp. 334–349, 2018.

[10] S. Khatun and B. Morshed, "Detection of myocardial infarction and arrhythmia from single-lead ECG data using bagging trees classifier," in *Int. Conf. on Electro Information Technology (EIT)*, 2017.

[11] H. Aidos, A. Lourenço, D. Batista, S. Bulò, and A. Fred, "Semi-supervised consensus clustering for ECG pathology classification," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2015, pp. 150–164.

[12] M. Cvikl and A. Zemva, "FPGA-oriented HW/SW implementation of ECG beat detection and classification algorithm," *Digital Signal Processing*, vol. 20, no. 1, pp. 238–248, 2010.

[13] H. Zairi, M. Kedir-Talha, S. Benouar, and A. Ait-Amer, "Intelligent system for detecting cardiac arrhythmia on FPGA," in *Int. Conf. on Information and Communication Systems (ICICS)*. IEEE, 2014, pp. 1–5.

[14] A. Patel, P. Gakare, and A. Cheeran, "Real time ECG feature extraction and arrhythmia detection on a mobile platform," *Int. Journal on Computer Applications*, vol. 44, no. 23, pp. 40–45, 2012.

[15] H. Chatterjee, R. Gupta, and M. Mitra, "Real time P and T wave detection from ECG using FPGA," *Procedia Technology*, vol. 4, pp. 840–844, 2012.

[16] M. Karim, M. El Kouache, M. Amarouch *et al.*, "An FPGA-based system for real-time electrocardiographic detection of STEMI," in *Int. Conf. on Advanced Technologies for Signal and Image Processing (ATSIP)*. IEEE, 2016, pp. 830–835.

[17] Y. Sun and A. Cheng, "Machine learning on-a-chip: A high-performance low-power reusable neuron architecture for artificial neural networks in ECG classifications," *Computers in biology and medicine*, vol. 42, no. 7, pp. 751–757, 2012.

[18] M. Wess, P. Manoj, and A. Jantsch, "Neural network based ECG anomaly detection on FPGA and trade-off analysis," in *IEEE Int. Symposium on Circuits and Systems (ISCAS)*. IEEE, 2017, pp. 1–4.

[19] H. de Carvalho Junior, R. Moreno, T. Pimenta, P. Crepaldi, and E. Cintra, "A heart disease recognition embedded system with fuzzy cluster algorithm," *Computer methods and programs in biomedicine*, vol. 110, no. 3, pp. 447–454, 2013.

[20] S. Pudukotai Dinakarrao and A. Jantsch, "ADDHard: Arrhythmia Detection with Digital Hardware by Learning ECG Signal," in *Proc. of the 2018 on Great Lakes Symposium on VLSI*. ACM, 2018, pp. 495–498.

[21] X. Wang, Y. Zhu, Y. Ha, M. Qiu, and T. Huang, "An FPGA-based cloud system for massive ECG data analysis," *IEEE Trans. on Circuits and Systems II: Express Briefs*, vol. 64, no. 3, pp. 309–313, 2017.

[22] A. Goldberger, L. Amaral, L. Glass, J. Hausdorff, P. Ivanov, R. Mark, J. Mietus, G. Moody, C. Peng, and H. Stanley, "Physiobank, physiotoolkit, and physionet: components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, no. 23, pp. e215–e220, 2000.

[23] Association for the Advancement of Medical Instrumentation and others, "Testing and reporting performance results of cardiac rhythm and ST segment measurement algorithms," *ANSI/AAMI EC57*, vol. 1998-(R)2008, 2008.