






A Tiny Transformer for Low-Power Arrhythmia Classification on Microcontrollers

Paola Busia , Matteo Antonio Scrugli , Victor Jean-Baptiste Jung , Luca Benini , *Fellow, IEEE*,
and Paolo Meloni 

Abstract—Wearable systems for the continuous and real-time monitoring of cardiovascular diseases are becoming widespread and valuable assets in diagnosis and therapy. A promising approach for real-time analysis of the electrocardiographic (ECG) signal and the detection of heart conditions, such as arrhythmia, is represented by the transformer machine learning model. Transformers are powerful models for the classification of time series, although efficient implementation in the wearable domain raises significant design challenges, to combine adequate accuracy and a suitable complexity. In this work, we present a tiny transformer model for the analysis of the ECG signal, requiring only 6k parameters and reaching 98.97% accuracy in the recognition of the 5 most common arrhythmia classes from the MIT-BIH Arrhythmia database, assessed considering 8-bit integer inference as required for efficient execution on low-power microcontroller-based devices. We explored an augmentation-based training approach for improving the robustness against electrode motion artifacts noise, resulting in a worst-case post-deployment performance assessment of 98.36% accuracy. Suitability for wearable monitoring solutions is finally demonstrated through efficient deployment on the parallel ultra-low-power GAP9 processor, where inference execution requires 4.28ms and 0.09mJ.

Index Terms—Arrhythmia, transformer, wearable monitoring.

I. INTRODUCTION

ACCORDING to the World Health Organization, cardiovascular diseases are the first cause of death worldwide [1]. The electrocardiographic (ECG) signal represents the commonly accepted non-invasive diagnostic and monitoring tool for the detection and recognition of anomalies in the functioning of the heart, whose early detection can help reduce their consequences and provide adequate treatment.

Manuscript received 12 February 2024; revised 23 April 2024; accepted 8 May 2024. Date of publication 16 May 2024; date of current version 5 February 2025. This work was supported by the Key Digital Technologies Joint Undertaking (KDT JU) in EdgeAI “Edge AI Technologies for Optimised Performance Embedded Processing” project under Grant 101097300. This paper was recommended by Associate Editor Jae-sun Seo. (*Corresponding author: Paola Busia.*)

Paola Busia, Matteo Antonio Scrugli, and Paolo Meloni are with the DIEE, University of Cagliari, 09123 Cagliari, Italy (e-mail: paola.busia@unica.it; matteoa.scrugli@unica.it; paolo.meloni@unica.it).

Victor Jean-Baptiste Jung is with the Integrated Systems Laboratory, ETH Zürich, 8000-8099Z Zürich, Switzerland (e-mail: jungvi@iis.ee.ethz.ch).

Luca Benini is with the Integrated Systems Laboratory, ETH Zürich, Zürich, Switzerland, and also with the DEI, University of Bologna, 40126 Bologna, Italy (e-mail: lbenini@iis.ee.ethz.ch).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TBCAS.2024.3401858>.

Digital Object Identifier 10.1109/TBCAS.2024.3401858

It is established that wearable solutions are of great importance for allowing continuous monitoring of chronic conditions while reducing the burden on clinicians and hospital structures [2]. In this context, where the privacy of the data is a relevant concern, near-sensor edge processing represents an increasingly attractive option. Particular benefits would result from continuous monitoring for the detection of arrhythmia, which represents an anomaly of the heartbeat, showing an irregular or abnormal rhythm.

Large-size transformers have recently demonstrated state-of-the-art accuracy in the recognition of different arrhythmia types [3], [4], surpassing the performance obtained with Convolutional Neural Networks (CNNs) [5]. As a consequence, the ability to capture non-local data dependencies provided by the attention mechanism in transformers appears very promising also in the ECG processing field. Nonetheless, the complexity of current models [3], [4] is not compatible with efficient execution on wearable monitoring devices: the neural network-based classification approaches presented in the literature for this processing domain are dominated by lightweight CNNs [6], [7], whose accuracy is still suboptimal compared to the state of the art, whereas the role of transformer models at this complexity scale still needs to be investigated.

In this work, we plan to combine the advantages deriving from the attention mechanism exploited in transformers with actual execution on wearable devices for out-of-ambulatory continuous and real-time monitoring. To this aim, we have designed a lightweight transformer architecture that can be executed on microcontrollers, resulting from a design exploration of the topological parameters having the highest impact on its storage requirements and complexity. Moreover, we have considered motion artifact noise resilience as a main objective within the overall design process, using signals with artifacts to augment the dataset during training.

The main contributions of this work can be summarized in the following points:

- we present a tiny transformer model for arrhythmia recognition, reaching accuracy aligned with the state of the art of transformer-based classifiers, with $60\times$ fewer parameters and $300\times$ fewer required operations;
- we evaluate the robustness of our proposed classifier considering common issues of real-time execution, resulting from the presence of artifacts and noise in the signal acquisition;
- we provide an efficient implementation on a low-power commercial MCU, resulting in less than 0.09mJ per

TABLE I

STATE OF THE ART OF NEURAL NETWORK-BASED ARRHYTHMIA CLASSIFICATION ON THE MIT-BIH ARRHYTHMIA DATABASE, INCLUDING CLASSIFICATION PERFORMANCE AND EFFICIENCY FIGURES EVALUATED ON THE TARGET HARDWARE. WE INDICATE WITH THE SYMBOL “-” THAT THE CORRESPONDING INFORMATION IS NOT REPORTED OR NOT ASSESSED IN THE CITED PAPER

Work	Year	Classes	Task	Model	Accuracy	Memory	MOPS	Inference Time	Energy Consumption	Hardware Target	Noise Addition
[4]	2019	4	intra-patient	ViT	99.62%	3 MB*	365*	-	-	-	×
[3]	2022	4	intra-patient	Transformer	99.49%	20 MB	70.53	-	-	-	×
[8]	2022	5	intra-patient	MLP	> 99%	-	30.41	-	-	-	×
[5]	2022	8	intra-patient	CNN	99.54%	16 MB	-	-	-	-	×
[9]	2023	4	intra-patient	CNN	99%	2.2 MB	151	-	-	-	×
[7]	2022	5	intra-patient	CNN	98.89%	96 kB	0.74	215 ms	0.66 mJ	ARM Cortex-M4	×
[6]	2023	3	inter-patient	CNN	98.18%	15 kB	0.05	1 ms	-	Cortex-ARMv8	✓
this work	2024	5	intra-patient	ViT 32-bit	99.05%						×
				ViT 8-bit	98.97%	49 kB	0.97	4.28 ms	0.09 mJ	GAP9	×
				ViT 8-bit	98.36%						✓

* Estimated from paper.

inference obtained with quantization up to 8-bit precision and parallel execution on 8 cores, thus showing the proposed model is suitable for continuous and real-time wearable monitoring.

We start with a review of the state of the art of classification approaches in Section II. In Section III, we provide a general overview of the considered ECG monitoring system and describe its components, focusing on the architecture of the proposed transformer-based classifier. The assessment of the classification performance is presented in Section IV. Finally, we report in Section V the evaluation of the performance of the proposed model deployed on the GAP9 low-power multicore microcontroller unit (MCU).

II. RELATED WORK

Table I summarizes the recent works presenting neural network-based models for arrhythmia recognition on the MIT-BIH dataset [10], [11]. For each of the listed works, we report in subsequent columns the publication year, the number of arrhythmia classes considered, the specific task addressed, the classification model, and the test accuracy obtained. The “Memory”, “MOPS”, “Inference Time”, and “Energy consumption” columns summarize the complexity and performance evaluated on the target platform, reported in the “Hardware Target” column. Finally, for each work, we indicate whether the degradation of the classification accuracy in the presence of growing levels of noise has been explored.

As can be noted, classification based on CNNs reaches remarkable results. In [5] the authors present a CNN adapted from the EfficientNet family, focusing on the 8-class arrhythmia recognition task and reaching 99.54% accuracy with around 5.3M parameters. A 1D-CNN exploiting 2M parameters is presented in [9], reporting 99% accuracy in the 4-class recognition problem. As an alternative, the work of [8] presents arrhythmia classification based on a Multi-Layer-Perceptron (MLP) model, scoring a classification accuracy higher than 99.47% for each of the 5 classes considered. These accurate models rely on a high number of parameters, not suitable for deployment on low-power microcontrollers. To better motivate this statement, we consider as a reference for wearable deployment a maximum memory footprint of 128 kB, which is the available storage

space in the L1 memory of our target platform, the GAP9 parallel ultra-low-power processor [12].

Tiny CNN models have also been exploited for real-time arrhythmia detection on resource-constrained devices. Ref. [7] presents two alternative CNNs for arrhythmia recognition on a commercial low-power microcontroller, reaching up to 98.69% and 98.89% accuracy in the 5-class recognition problem, with 18K and 93K parameters. The efficiency of inference execution was evaluated on the ST Sensortile microcontroller, where it requires 215ms and 0.66mJ. Hybrid models resulting from the combination of the CNN topology with the MLP, such as in [13], or the Long-Short-Term-Memory (LSTM) network, such as in [14], were also considered for binary classification, to distinguish between normal and ventricular, or generally abnormal, heartbeats. The proposed models reach up to 98.5% and 97% accuracy. The reduced complexity of these models results in a sub-optimal classification performance compared to the state of the art, encouraging further studies on the design of efficient models for arrhythmia recognition.

Similarly, the work of [6] presents a CNN working on the first derivative of the ECG signal and exploiting matched filters reproducing the different arrhythmia classes within the convolutional layers. The model also exploits the processing of the RR intervals (the distance between consecutive heartbeat peaks in the signal) through a stack of fully connected layers, obtaining 98.18% accuracy with a 15 kB memory footprint, addressing the more challenging inter-patient classification problem, where the classifier is trained and tested on the data from separate sets of patients. Edge deployment on a Raspberry Pi equipped with a Cortex-ARMv8 64-bit System on Chip results in less than 1 ms inference time. This very efficient model is thus suitable for low-power deployment, nonetheless the obtained accuracy, evaluated on only three arrhythmia classes, is still over 1% below the state of the art.

Another relevant contribution is presented in [15], introducing BioAIP, a specialized processor for biological signal processing, enabling the possibility to refine a CNN-based global classification model considering a small subset of patient-specific data, reporting up to 7.5% accuracy improvement on tests following the leave-one-patient-out approach, reaching up to 99.16% classification accuracy. However, this interesting work is mostly oriented to the description of the proposed

processing platform, rather than describing the specifics of the proposed ECG classification and testing approach.

As can be noticed from Table I, transformer-based classification represents the state-of-the-art approach, particularly based on the work of [4]. The authors presented a Vision Transformer (ViT) consisting of 3 encoding blocks and 4 parallel heads in the attention layer. They obtained 99.62% accuracy, exploiting a denoising and baseline-wandering removal step, and concatenating the information about the RR intervals within the last classification layer. An interesting approach is presented in the work of [3]. The proposed model is applied on 3s-long windows of signal and provides the positioning and classification of the heartbeats occurring during that period, reaching 99.86% heartbeat positioning accuracy, and 99.49% classification accuracy. The best model exploits 5M parameters, but competitive results were obtained with a similar architecture based on 0.75MB of parameters, reaching 99.05% accuracy.

Other relevant examples are represented by the transformer architectures presented in [16], [17], [18], which were designed and tested on different datasets, considering multiple-lead acquisition.

To the best of our knowledge, the transformer-based classifiers in the literature have storage and computational requirements non-compatible with efficient deployment on low-power MCUs. Considering our target platform, GAP9, the model in [4] cannot be allocated even in the larger L2 memory, which provides an additional 1.5 MB of storage space accessible with lower efficiency.

Finally, the robustness of the classification under different noise conditions is only explored in the work of [6], considering white noise addition up to 3db Signal to Noise Ratio (SNR). On the contrary, we assessed the classification performance under different noise conditions, considering real recordings of electrode motion artifacts noise.

In this work, we thus present an efficient transformer architecture, targeting accurate arrhythmia detection on low-power wearable devices, and extending the performance assessment to the typical challenges of real-time wearable operating conditions. As most of the referenced works, with the exception of [6], we refer to the intra-patient classification problem, where each heartbeat in the dataset is considered independently, and randomly assigned to the training, validation, or test set.

III. ECG CLASSIFIER DESIGN

A typical ECG signal is shown in Fig. 1, referring to acquisition from a single electrode. Each heartbeat is defined by a complex wave around a peak (R) in the amplitude, commonly known as the QRS complex. The shape of the wave around the R-peak, and the intervals between consecutive peaks (pre-RR and post-RR intervals), highlighted in the figure, are both meaningful diagnostic elements.

Fig. 2 describes a general overview of the real-time monitoring system that would be executed on a wearable device. The input to the system is represented by the continuous ECG signal. As reported in the plot, our proposed classification system exploits two main information elements: a window of




Fig. 1. Typical ECG signal. Highlighted is the considered window around a single heartbeat, the R peak position and the definition considered in the following for pre- and post- RR intervals.




Fig. 2. Overview of the ECG monitoring system.

the acquired signal selected around a single heartbeat, and a pair of values representing the pre-RR and post-RR intervals. The identification and processing of these two elements is the result of the peak preparation block, which involves a denoising step and a peak detection step. The ECG window and the pre-RR and post-RR intervals are thus provided to the transformer-based classifier, for the recognition of the considered classes of arrhythmia.

In the following, we describe in Section III-A the details of the data preparation block. The architectural description of our proposed transformer model for heartbeat classification is reported in Section III-B. Finally, Section III-C presents the main characteristics of the reference dataset considered for the design optimization and assessment.

A. Peak Preparation

The raw ECG signal goes through a denoising block performing baseline-wandering and noise removal. This block was structured according to the scheme suggested in [4], which includes two main filtering steps. The baseline wandering component is identified thanks to two median filters with respectively 200ms and 600ms window length, and subtracted from the signal. At this point, a low-pass filter with a cut-off frequency of 35 Hz is exploited for powerline noise removal.

The clean signal is then processed for the identification of the window of interest around the R peak. Among the common QRS detection algorithms in the literature, here we refer to the well-known Pan-Tompkins detection algorithm, which allows

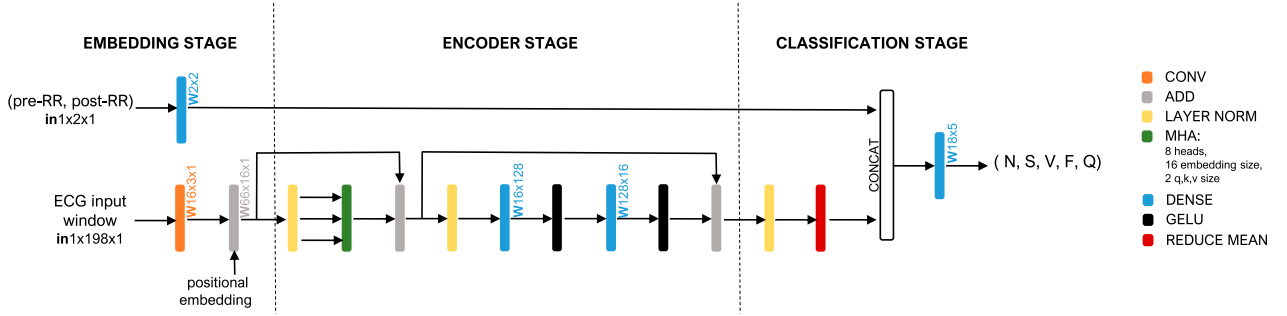


Fig. 3. Architecture of the proposed transformer model for arrhythmia classification. The shape of the weight tensor in each layer is indicated with the prefix W , whereas the prefix in refers to the shape of the input data tensors.

the detection of 99.67% of the heartbeats in the considered dataset [19].

B. Transformer for Arrhythmia Recognition

The proposed classifier is a Vision Transformer [20], [21], adapted for the processing of 1D medical signals, similarly to the approaches of [22], [23], which presented tiny vision transformer models for hand gesture recognition based on the electromyographic signal and electroencephalography processing for epilepsy monitoring. The model proposed in this work was developed starting from the architecture presented in [23] and further optimized for the arrhythmia recognition task.

The structure of the model is represented in Fig. 3. A typical ViT model consists of three main stages, an *embedding stage*, an *encoder stage*, and a *classification stage*. The embedding stage performs a data preparation task, producing as output a sequence of patches, inspired by the tokens of an input text, and including the positional information.

This sequence represents the input to the encoder stage, where most of the data processing is performed. The encoder is constituted by a stack of blocks including a Multi-Head-Attention (MHA) layer and a Feed Forward network, exploiting residual connections.

The MHA represents the most important mechanism within the transformer, allowing the evaluation of the mutual relevance between two points within the observed time series. As a first step for the attention evaluation, the input sequence of patches is projected into three different spaces, known as queries Q , keys K , and values V . The first two projections are used for the evaluation of the attention matrix, whose items allow to give weight to the elements within the third V projection, according to the Equation in 1, where d represents the size of the Q , K , and V projections.

$$\text{Attention}(Q, K, V) = \text{Softmax}\left(\frac{QK^T}{\sqrt{d}}\right)V \quad (1)$$

The attention mechanism is usually replicated considering multiple parallel executions, called heads, where independent projections of the input are evaluated. In MHA, a final linear projection is performed in order to restore the embedding data size. Finally, the classification stage is usually implemented as an MLP or a Dense layer.

In the following, we describe the design exploration performed for the definition of the proposed transformer topology. As a starting point, we fixed a few design choices, based on the literature and previous results in [23]. In detail, in alignment with the choice in [7], the input size was selected equal to 198 samples, equally distributed on the left and on the right of the considered R peak position within the heartbeat, and corresponding to roughly 0.5 s of data based on the sampling frequency of the considered dataset.

Additionally, we considered a general topology embedding a single encoder block, aligned with the one defined in [23], which featured 50k parameters. Furthermore, in our model, the embedding stage is obtained with a 1D-convolutional block, as proposed in [22]. This solution adapts the Vision Transformer embedding to 1D signals. In this approach, the input signal of shape $1 \times 198 \times 1$ (*channel, width, height*) is transformed into a tensor of size $E \times S \times 1$, where the number of output channels E is interpreted as the width of a single patch ($E \times 1$), whereas S defines the length of the sequence of patches to be processed and is obtained from the original 198×1 input based on the size k of the filtering kernels (applied considering no overlap, with stride $s = k$).

Finally, according to the original topology presented in [21], we considered a fixed relation between the size of the patches processed within the attention layer - the embedding size E -, the number of attention heads H , and the size P of the projected queries, keys, and values; the relation is set as $P = E/H$.

The design exploration was performed to select the optimal value of the parameters impacting the computational load of the transformer, and especially of the MHA layer, namely the embedding size E , the length of the sequence of patches S , the number of heads H , and the size h of the hidden layer in the feed-forward network executed after MHA. The exploration process aimed at maximizing the accuracy on the target dataset, under a defined storage constraint (128 kB). In this stage, the comparison of the candidate design points evaluated the classification accuracy on some left-out records, i.e. records “108”, “210”, “233”, “106”, and “203”.

The parameters were selected among a small set of alternatives, based on a sequential evaluation where all parameters were kept fixed except for the one being explored at a certain step. We started from the embedding size, and the choice $E = 32$ exploited in [23]: we evaluated values within the set $\{16, 32,$

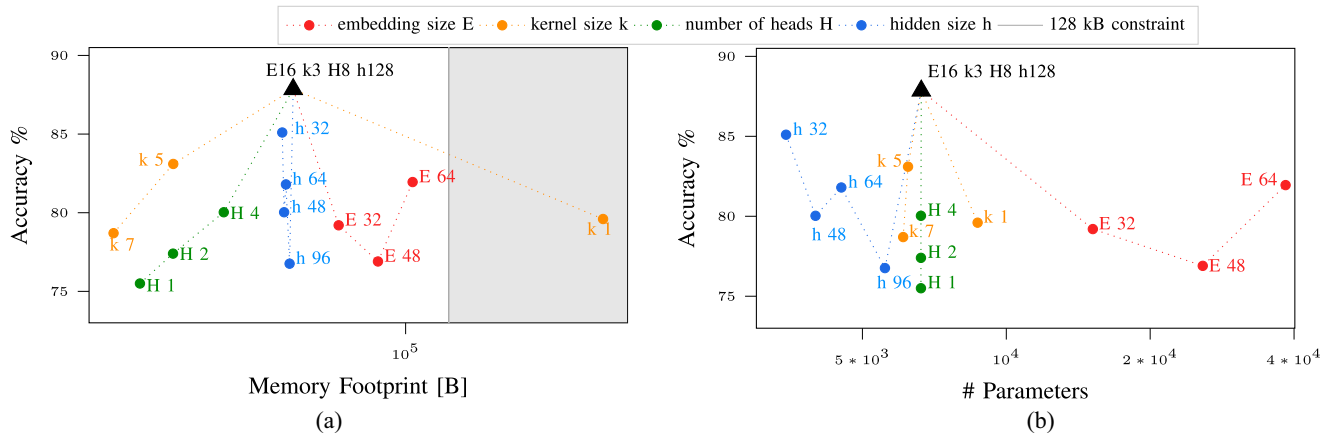


Fig. 4. Exploration of the architectural parameters based on left-out test accuracy on the MIT-BIH dataset. Each dot in the plot represents one of the evaluated design points, resulting from the selection of the parameters explored: embedding size, kernel size, number of heads, and hidden layer size. Dot placement in a) considers the memory footprint of the model, accounting for the storage of the partial results produced during inference computation, whereas placement in b) considers the number of parameters in the model. The selected point is highlighted in black. Shaded grey area: outside of the 128 kB feasibility region.

48, 64}. For this exploration step, the kernel size k was set to 3, H was set to 8, and h was set to 128. Based on the outcome of the exploration, $E = 16$ was selected and considered in the following. The second design choice was the kernel/stride size impacting the sequence length in MHA. We considered values within the set $\{1, 3, 5, 7\}$ and obtained the highest accuracy for $k = 3$, which was kept in the following. The third design choice was the number of heads H , selected equal to 8 among values within $\{1, 2, 4, 8\}$. Finally, the hidden layer size h was selected equal to 128 from a set including $\{32, 48, 96, 128\}$.

The outcome of the exploration is summarized in Fig. 4(a) and 4(b), where each evaluated point is placed according to its accuracy and memory footprint or parameters' count. The comparison between the two plots highlights how the memory footprint is highly impacted by the storage space allocated for the partially computed results, especially the attention matrix. As a consequence, design choices having a limited impact on the parameters' count, such as the kernel size in the embedding convolutional layer, have a very significant effect on the overall memory requirements of the model. As can be noticed, all the evaluated design points fit the memory constraint defined for the exploration, except for the solution corresponding to $k = 1$, $E = 16$, $H = 8$, and $h = 128$. The selected point, resulting in the highest accuracy, is highlighted in the plot.

The selected model architecture is summarized in Table II, where the first column highlights the composition of each stage, the second column lists the operands applied during the execution of the model, the third column reports the shape of the output data produced by the computation, the fourth column summarizes the values of the parameters selected based on the design exploration, and the fifth one details the parameters' count for each operand.

Considering the results reported in other works in the literature [3], [4], [6], [24], presenting neural network models for arrhythmia recognition and demonstrating the benefits of including the information about the RR intervals to differentiate the considered arrhythmia classes, we also included this

TABLE II
PARAMETERS OF THE PROPOSED TRANSFORMER MODEL FOR
ARRHYTHMIA CLASSIFICATION

Stage	Layer	Output	Parameters	
			Values	#
Embedding	Convolution	$16 \times 66 \times 1$	k: 3×1 s: 3	64
	Add Pos. Embedding	$16 \times 66 \times 1$		1056
	Dense	1×2		4
Encoder	Layer Norm	$16 \times 66 \times 1$		32
	Multi-Head-Attention	$16 \times 66 \times 1$	E: 16 S: 66 d: 2 H: 8	1088
	Layer Norm	$16 \times 66 \times 1$		32
	Dense	$128 \times 66 \times 1$		2176
	Gelu	$128 \times 66 \times 1$		
	Dense	$16 \times 66 \times 1$		2064
	Gelu	$16 \times 66 \times 1$		
Classification	Layer Norm	$16 \times 66 \times 1$		32
	Reduce Mean	1×16		
	Dense	1×5		95
Tot				6643

information as an additional input to the transformer model. We thus followed the example of [4] and introduced an additional input to the traditional transformer topology, representing the distance between the current R peak and the previous and following one. This second input is first linearly projected within the embedding stage and then concatenated to the output of the encoder which processes the ECG window around the heartbeat.

The tuple of values representing the RR intervals is normalized into the range $[-2, 2]$, to resemble the distribution of the encoder output at the concatenation level, thus simplifying the quantization process.

C. Reference Dataset

As a main reference for this study, we considered the MIT-BIH Arrhythmia Database [10], [11], which collects the ECG recordings of 47 subjects into 48 records, lasting 30 minutes. The data included in the dataset is a composition of randomly

selected ambulatory recordings from Boston’s Beth Israel Hospital, and recordings selected to include the clinically relevant arrhythmia examples. Each record reports the signal acquired by two channels, usually modified limb lead II (MLII) and V1, with a 360Hz sampling frequency. We refer in the following to processing based on the signal acquired by the MLII.

In alignment with most of the works in the literature [3], [4], [6], [24], we excluded from the training and test sets the records containing paced beats, namely “102”, “104”, “107”, “217”. Furthermore, we referred to the AAMI standard [25] to identify the 5 most relevant classification groups for the arrhythmia recognition problem: N (non-ectopic beats), S (supraventricular ectopic beats), V (ventricular ectopic beats), F (fusion beat), Q (unclassifiable beat).

Real-time heartbeat classification on wearable monitoring devices can be affected by different sources of noise. The MIT-BIH Noise Stress Test Database [11], [26] includes authentic ECG recordings accompanied by real noise samples, essential for simulating typical disturbances encountered in ambulatory ECG recordings. This dataset comprises twelve half-hour ECG recordings and three additional half-hour segments explicitly containing noise sourced from three primary disturbances: baseline wander, muscle artifacts, and electrode motion artifacts. The noise records were derived from recordings using standard ECG equipment and physically active volunteers, with electrodes positioned on the limbs to avoid interfering with the ECG signal visibility. This setup effectively replicates the challenging conditions often faced in real-world ECG data collection. The focus of this research is particularly on the noise caused by electrode motion artifacts, the most problematic type due to its potential to mimic ectopic beats and its resistance to simple filtering techniques. This real noise environment is critical for analyzing the robustness of ECG signal processing techniques under realistic conditions.

Here we describe the process leading to the creation of the noise-corrupted data referenced in the Experimental Section IV-C. The corruption of the signal is indicated as different Signal-to-Noise ratio (SNR) levels, expressed in dB. For brevity, we indicate as noiseless the original signal, with no additional noise deriving from electrode motion artifacts. Given a noiseless signal S from the MIT-BIH Arrhythmia Database and a noise component N from the MIT-BIH Noise Stress Test Database, the noise was scaled by a scaling factor α calculated based on the desired SNR, according to Equation 2.

$$\alpha = \sqrt{\frac{\text{mean}(|S|^2)}{\text{mean}(|N|^2) \times 10^{\frac{\text{desired SNR}}{10}}}} \quad (2)$$

The resulting noisy signal Y is generated in Equation 3, where the scaling factor α results from Equation 2.

$$Y = S + \alpha \times N \quad (3)$$

Each noisy ECG signal is annotated with the original arrhythmia labels from the MIT-BIH Arrhythmia Database.

The MIT-BIH dataset was chosen as a reference for the performance assessment due to the extensive literature targeting this resource. However, the ECG classification system

TABLE III
COMPOSITION OF THE MIT-BIH DATASET AND TYPICAL TRAINING - VALIDATION - TEST SPLIT CONSIDERED

Arrhythmia Class	# Samples	Train	Valid	Test
N	90098	63022	9057	18019
S	2781	1967	273	541
V	7007	4919	664	1424
F	802	573	75	154
Q	15	12	1	2

we are proposing can be exploited for arrhythmia recognition based on single-lead acquisition, provided that proper upsampling/downsampling of the data is performed to obtain a coherent input shape corresponding to 0.5 s observation interval around the detected R peak.

D. Evaluation Metrics

As will be further detailed, Table III reports the typical composition of the training and test sets from the MIT-BIH dataset. As can be noticed, the dataset is highly unbalanced, presenting a majority of normal heartbeats and a different number of instances for each of the represented classes. Due to this reason, for the performance assessment, we not only considered the overall classification accuracy but also the sensitivity and precision on each of the targeted arrhythmia classes. The definition we referenced for these metrics is reported in equations 4 and 5, where we apply them to class N:

$$Sens_N = \frac{TN_N}{TN_N + FS_N + FV_N + FF_N + FQ_N} \quad (4)$$

$$Prec_N = \frac{TN_N}{TN_N + FN_S + FN_V + FN_F + FN_Q} \quad (5)$$

where the notation TN_N indicates a normal sample, correctly classified as normal, FS_N indicates a normal sample which is classified as a supra-ventricular ectopic beat, and finally FN_S indicates a supra-ventricular ectopic beat classified as normal. A similar logic can be applied to the interpretation of the other symbols in the equations.

IV. EXPERIMENTAL RESULTS

In the following, we summarize the outcome of the tests performed to evaluate the performance of our proposed arrhythmia classifier.

A. Intra-Patient Arrhythmia Recognition

We assessed the performance of our proposed model in the intra-patient classification task, based on 5-fold cross-validation. The ECG records in the dataset were segmented around each heartbeat corresponding to one of the NSVFQ arrhythmia classes, considering the dataset annotations for the heartbeat positioning. All collected segments were shuffled and randomly split into a training set, a validation set, and a test set, with a 7:1:2 ratio, to enable a direct comparison with the literature [4]. The composition of the dataset and a typical train-valid-test split is reported in Table III.

TABLE IV
CONFUSION MATRIX SUMMARIZING PERFORMANCE ON THE CONSIDERED SPLITS FOR THE FULL-PRECISION MODEL, WHEN NO ADDITIONAL NOISE IS APPLIED TO THE SIGNAL

True Labels	N	89867	119	79	31	2
	S	289	2468	23	1	0
	V	139	31	6783	45	0
	F	120	4	62	616	0
	Q	10	0	4	0	1
		N	S	V	F	Q
		Predicted Labels				

TABLE V
TEST CLASSIFICATION PERFORMANCE, WHEN NO ADDITIONAL NOISE IS APPLIED TO THE SIGNAL. WE REPORT MEAN VALUES AND STANDARD DEVIATION IN 5-FOLD CROSS-VALIDATION WITH FULL-PRECISION INFERENCE

Class	Sensitivity	Precision
N	99.74% (0.08)	99.38% (0.04)
S	88.79% (2.02)	94.14% (1.72)
V	96.91% (0.58)	97.59% (0.24)
F	76.92% (5.48)	88.99% (1.93)
Q	10% (20)	10% (20)
Accuracy tot	99.05% (0.08)	

The training of the classifier was performed on one NVIDIA T4 Tensor Core GPU, using the TensorFlow framework inside the Google Colab environment. We exploited 200 epochs training, with Adam optimizer, adaptive reduce-on-plateau learning rate with starting value $2e - 3$, and batch size 128.

The classification performance for full-precision inference is reported in Table V, in terms of mean value and standard deviation across the 5 tests considered. The performance of the model is only slightly impacted by the extremely low sensitivity and precision registered for the Q class, which is very under-represented in the dataset, and excluded from most studies [3], [4]. Our proposed model reaches an average of 99.05% accuracy on the 5 splits considered. Table IV reports the summary confusion matrix referring to the cross-validation test. As can be derived from Table V, the main weakness is represented by a reduced sensitivity in the recognition of the F class.

B. Ablation Study

We report in the following the ablation study of the different data pre-processing choices explored, summarized in Table VI, limiting the analysis to a single train-test split.

The first line in the Table reports the accuracy of the selected model, depicted in Fig. 3 of Section III-B. We started from the removal of the concatenation of the second input, representing the RR interval information, which helps in the arrhythmia class discrimination, as reported in several previous studies [3], [4], [6], [24]. The result was a 0.36% accuracy drop, reported in the second line.

We assessed at this point the relevance of the denoising step. We report in the third line the performance obtained when considering the RR information, but the selected topology was trained, validated, and tested on unfiltered data, not pre-processed with the denoising step in Fig. 2. The resulting accuracy is lower than the best obtained by 0.51% points. To confirm

TABLE VI
COMPARISON OF TEST PERFORMANCE ON A SINGLE 7:1:2 SPLIT, BASED ON DIFFERENT INPUT PROCESSING CHOICES. ACCURACY NUMBERS REFER TO TESTS PERFORMED WITH NO ADDITIONAL NOISE APPLIED TO THE SIGNAL

Model	Test Accuracy
Selected Model	99.17%
no RR	98.81%
no denoising	98.66%
no RR and no denoising	98.69%




Fig. 5. Test classification performance considering increasing levels of noise deriving from electrode motion artifacts.

this result, we finally removed both the RR concatenation and the denoising step, thus obtaining the performance reported in the fourth line. As can be observed, there is no further accuracy degradation compared to the number reported in the third line, showing how the RR information only provides a relevant advantage when combined with an effective denoising approach.

C. Post-Deployment Evaluation

This work targets wearable devices for real-time heartbeat classification, we thus present in this section the assessment of the performance degradation deriving from two main challenges of wearable deployment, namely increasing levels of noise corrupting the ECG signal and accuracy loss due to data quantization.

1) *Noise Degradation*: As a first step for the noise degradation analysis, we used as a baseline the accuracy reported in Table VI, deriving from the evaluation on the noiseless signal, and the corresponding specific train-test split. We repeated the same test using noise-augmented versions of the dataset samples in training and test, to create Fig. 5. Different lines correspond to different levels of noise used for augmenting the training set, while different evaluation points in each line correspond to a different level of noise in the test. Considered levels are *noiseless*, *24dB SNR*, *10dB SNR*, *3dB SNR*, and *balanced mix*, where *noiseless* corresponds to the original non-augmented dataset samples and *balanced mix* corresponds to an equally partitioned test set, with different noise levels applied to each subset, to emulate real-life conditions where the noise level is expected to vary over an evaluation.

TABLE VII

TEST CLASSIFICATION PERFORMANCE, REPORTING MEAN VALUES AND STANDARD DEVIATION IN 5-FOLD CROSS-VALIDATION WITH FULL-PRECISION INFERENCE, WHEN TRAINING DATA AUGMENTATION WITH VARYING NOISE LEVELS IS EXPLOITED

Class	Noiseless Test		Balanced Mix Test	
	Sensitivity	Precision	Sensitivity	Precision
N	99.73% (0.08)	99.42% (0.08)	99.59% (0.12)	99.15% (0.08)
S	89.13% (1.1)	93.64% (1.34)	84.08% (1.2)	91.34% (1.62)
V	97% (0.42)	98.01% (0.33)	95.7% (0.42)	96.59% (0.47)
F	78.73% (4.97)	88.63% (3.42)	71.05% (4.34)	84.62% (5.15)
Q	4% (8)	10% (20)	2% (4)	2.5% (5)
Accuracy	99.07% (0.06)		98.65% (0.07)	

TABLE VIII

CONFUSION MATRIX SUMMARIZING *NOISELESS* TEST PERFORMANCE ON THE CONSIDERED SPLITS FOR THE FULL-PRECISION MODEL, WHEN TRAINING DATA AUGMENTATION WITH VARYING NOISE LEVELS IS EXPLOITED

True Labels	N	89858	136	50	37	17
	S		278	2479	22	2
V		130	33	6789	44	2
F		104	1	64	631	2
Q		11	0	3	0	1
		N	S	V	F	Q

The solid grey line shows the results obtained when performing classification on different test conditions with the model trained on the non-corrupted signal. As can be observed, in this case, the accuracy is dramatically affected by the noise, with up to almost 8% points drop registered with high noise levels and test SNR equal to 3dB.

In order to improve robustness during real-time inference, we considered the effect of including examples of corrupted signals during the training process, represented by the other lines in the plot. As can be observed, augmenting with a specific level of noise improves the resilience of the classification algorithm when a similar noise level is considered in the test, although in general, it determines a slight degradation for lower noise levels. The *balanced mix* strategy, on the other hand, performs better than specific levels in most evaluation points and especially in the leftmost one, representing real-life conditions. It reduces the accuracy drop on the most corrupted signal (SNR = 3dB) to less than 1.2% points. This gain in the robustness to high noise levels does not compromise the classification accuracy on the noiseless signal, which is only 0.02% lower than the highest one reported in Table VI. The results thus confirmed the advantages of a training approach considering a variable noise level for augmentation.

Based on these findings, we repeated 5-fold cross-validation using training data augmentation and reported the results in Table VII. The first two columns summarize the performance for the classification of the noiseless signal. The average accuracy is 99.07%, thus no degradation is observed compared to the un-augmented training approach. The confusion matrix is also included for completeness and reported in Table VIII.

On the other hand, a significant advantage was obtained on the *balanced mix* test. In this scenario, 98.65% was achieved on average on the 5 train-test splits considered. The corresponding cumulative confusion matrix is reported in Table IX.

TABLE IX

CONFUSION MATRIX SUMMARIZING *BALANCED MIX* TEST PERFORMANCE ON THE CONSIDERED SPLITS FOR THE FULL-PRECISION MODEL, WHEN TRAINING DATA AUGMENTATION WITH VARYING NOISE LEVELS IS EXPLOITED

True Labels	N	358901	710	483	232	66
	S		1630	9355	130	8
V		807	182	26802	192	9
F		596	4	325	2278	5
Q		47	0	11	0	2
		N	S	V	F	Q

TABLE X

TEST CLASSIFICATION PERFORMANCE, REPORTING MEAN VALUES AND STANDARD DEVIATION IN 5-FOLD CROSS-VALIDATION WITH 8-BIT INFERENCE, WHEN TRAINING DATA AUGMENTATION WITH VARYING NOISE LEVELS IS EXPLOITED

Class	Noiseless Test		Balanced Mix Test	
	Sensitivity	Precision	Sensitivity	Precision
N	99.74% (0.05)	99.3% (0.08)	99.41% (0.14)	99.01% (0.1)
S	84.73% (1.1)	94.65% (0.87)	78.3% (1.55)	92.12% (0.8)
V	97.52% (0.42)	97.24% (0.45)	96.34% (0.53)	94.11% (1.26)
F	76.58% (3.52)	89.56% (2.57)	68.94% (2.62)	81.33% (4.45)
Q	9% (11.14)	40% (48.99)	5.75% (7.57)	4.48% (6.84)
Accuracy	98.97% (0.04)		98.36% (0.08)	

As can be observed, each test sample was replicated in order to represent a different noise condition in real-time inference. The resulting evaluation metrics are summarized in the right columns of Table VII. The comparison with the performance on the *noiseless* test shows a generally reduced sensitivity and precision, especially in the recognition of the S and F classes.

2) *Quantization*: Finally, in order to optimize the storage requirements of our model and exploit at best the byte processing resources on edge-processing platforms, we made use of the Quantlab framework [27] to obtain 8-bit quantization. To allow for efficient integer-only execution, the quantization strategy applies the solutions proposed in [28] for the implementation of non-linear operators, especially Layer Norm, which is replaced with an iterative approximation, exploiting 32-bit integer parameters.

We considered for quantization the models trained with noise augmentation and whose full-precision performance is summarized in Table VII. We exploited 15 epochs of quantization-aware fine-tuning, obtaining an average test accuracy of 98.97% on the clean noiseless signal, with a limited 0.1% drop from the accuracy of the full-precision models. The performance metrics are summarized in the first two columns of Table X, whereas the confusion matrix is reported in Table XI. The accuracy drop resulted from restricting all computations to the integer type, and it mostly affects the sensitivity of the S class, as can be observed from the table.

We finally tested the performance on the classification of noise-corrupted signal, with the *balanced mix* test. The outcome is reported in the third and fourth columns of Table X, and a summary confusion matrix is reported in Table XII. The assessment resulted in an average 98.36% accuracy, which we can consider as a worst-case post-deployment real-time performance prediction.

TABLE XI

CONFUSION MATRIX SUMMARIZING *NOISELESS* TEST PERFORMANCE ON THE CONSIDERED SPLITS FOR THE 8-BIT MODEL, WHEN TRAINING DATA AUGMENTATION WITH VARYING NOISE LEVELS IS EXPLOITED

True Labels	N	89860	115	88	32	3
	S	388	2356	37	0	0
	V	115	17	6825	40	1
	F	121	1	66	614	0
	Q	10	0	3	0	2
		N	S	V	F	Q
		Predicted Labels				

TABLE XII

CONFUSION MATRIX SUMMARIZING *BALANCED Mix* TEST PERFORMANCE ON THE CONSIDERED SPLITS FOR THE 8-BIT MODEL, WHEN TRAINING DATA AUGMENTATION WITH VARYING NOISE LEVELS IS EXPLOITED

True Labels	N	358267	647	1061	344	73
	S	2126	8708	282	7	1
	V	743	93	26971	169	16
	F	645	3	345	2211	4
	Q	46	0	9	0	5
		N	S	V	F	Q
		Predicted Labels				

D. Comparison With the State of the Art

We discuss at this point how our proposed solution compares to the state-of-the-art context, considering as a main reference the work of [4], which represents the most accurate alternative. Table XIII summarizes the most relevant performance figures for the comparison. The first line refers to full-precision inference performed on noiseless data, as we indicate the data not corrupted with electrode motion noise. A limited accuracy drop is observed when evaluating the performance of the 8-bit model, in the second line. Our result presents only a slight degradation with respect to the performance of the state-of-the-art transformer model, corresponding to only a 0.65% accuracy drop, despite the reduced complexity in terms of the number of parameters and memory footprint, reported in the fifth and sixth lines.

Finally, although a direct comparison with the reference model is not possible, we provide an assessment of worst-case performance considering a deployment scenario corrupted by different levels of noise. The training data augmentation allowed us to limit the accuracy drop to 0.4% points for full-precision inference, whereas post-deployment assessment of the quantized model results in 98.36% average accuracy, which we report in Table I as the worst-case expected real-time deployment performance.

V. DEPLOYMENT

In this section, we assess the efficiency of the proposed transformer classifier, considering its deployment on the parallel ultra-low power (PULP) GAP9 processor [12]. This commercial platform embeds a computing cluster of 8 parallel processors, accessing a shared 128kB L1 memory to enable efficient parallel execution. The offload of computations to the cluster is handled by an additional core, working as a Fabric Controller. The memory system includes also a 1.5MB L2 memory and supports voltage and frequency tuning for power management.

TABLE XIII

COMPARISON WITH THE STATE-OF-THE-ART TRANSFORMER MODEL FOR ARRHYTHMIA RECOGNITION

Metric	This Work	[4]
32-bit Accuracy Noiseless	99.05%	99.62%
8-bit Accuracy Noiseless	98.97%	-
32-bit Accuracy Noisy	98.65%	-
8-bit Accuracy Noisy	98.36%	-
Parameters	6649	405711 ¹
Memory	49kB	3MB ¹

¹ Estimated from the paper.

TABLE XIV

ON-HARDWARE PERFORMANCE ON THE GAP9 PROCESSOR FOR PARALLEL EXECUTION ON THE 8 CORES OF THE COMPUTING CLUSTER

	GAP9	
	370 MHz	240 MHz
Frequency	370 MHz	240 MHz
Time/Inf	2.85 ms	4.28 ms
Power	42.60 mW	20.33 mW
Energy/Inf	0.12 mJ	0.09 mJ

Based on the recent assessment on the tiny-ML benchmarks [29], GAP9 shows exceptional energy efficiency, with as low as 0.33mW/GOP, thus representing a perfect fit for continuous and real-time wearable monitoring tasks. Furthermore, the computational workload of the transformer model is intrinsically parallel, thus the computing cluster can be efficiently exploited to reduce inference time during real-time execution.

For efficient deployment, we considered the 8-bit model reported in Table XIII. We exploited the Dory code generation tool [30] for assisted implementation on the target platform. As depicted in Fig. 6(a), over 67% of the computing time is occupied by the MHA layer, whose implementation exploits the parallel resources on the platform by balancing the computation among the cores of the computing cluster: each head is mapped into one of the eight cores, according to the strategy described in [31]. As can be noticed from the plots in Fig. 6(b) and 6(c), the MHA layer is still the most relevant workload when reducing the number of parallel heads to 4, and it is comparable to the dense layers when further reducing it to 1. To evaluate the impact of this design parameter on the overall efficiency, based on our parallelization strategy, we report in Fig. 6(d) the achievable speedup in terms of required execution time and required number of OPS when considering parallelization on the 8 cores of the processing platform. As can be noticed, our implementation choices result in limited performance gains when reducing the number of heads below the number of available processing cores, despite a linear reduction of the workload.

Table XIV summarizes the performance figures evaluated at the working frequencies resulting in the lowest required inference time or energy consumption. At the maximum working frequency supported by the platform, 370 MHz, inference requires as low as 2.85 ms. On the other hand, the most energy-efficient configuration, obtained at the maximum frequency supported with low voltage supply, i.e. 240 MHz, results in 4.28 ms/inference time and 0.09 mJ energy consumption, for parallel execution exploiting 8 cores.

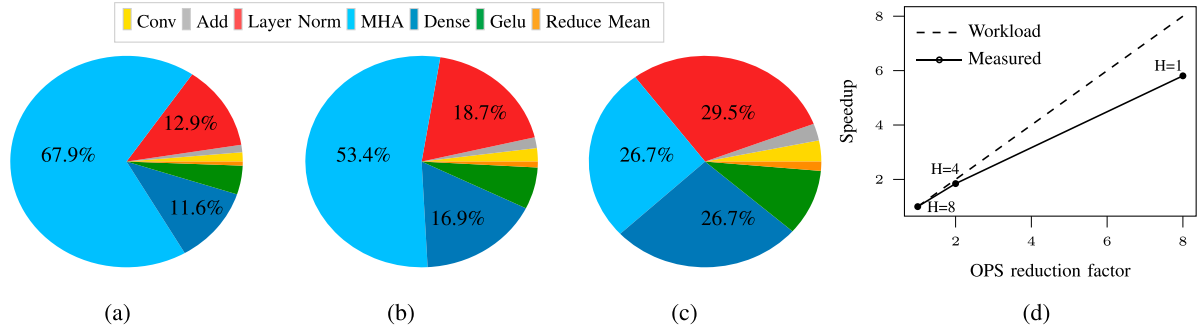


Fig. 6. Composition of the computational workload of the different operators expressed as percentage of the total inference time for parallel execution on 8 cores referring to (a) the proposed model, (b) the same model with $H = 4$, (c) the same model with $H = 1$. (d) Achievable speedup in the required execution time of MHA, reducing the number of heads from 8 to 4, to 1, resulting in OPS reduction by a factor of $2\times$ and $8\times$.

The execution of signal denoising and peak detection on the platform requires 0.77 ms for each newly acquired sample. System level latency accounts for 0.55 s, required for the acquisition of the window of 198 samples around the peak, plus a varying interval until the next peak detection, plus 0.77 ms to perform denoising on the last acquired sample, plus 4.28 ms inference time.

The power measurements were performed with a Power Profiler Kit II (PPK2) connected to the GAP9 Evaluation Kit, resulting in an average power consumption of 20.33mW. As a final consideration on the efficiency of the proposed transformer-based system, the comparison with the hardware performance metrics, reported in Table I for the referenced state-of-the-art works, demonstrates a reduced energy consumption per inference, despite a slightly increased computational complexity compared to the works of [6] and [7]. Nonetheless, this results in improved classification accuracy, obtained with a reduced number of trainable parameters compared to the model in [7].

VI. CONCLUSION

In this work, we presented an efficient transformer model for arrhythmia classification, reaching 99.05% accuracy in the classification of the 5 most common arrhythmia classes from the MIT-BIH Arrhythmia database. The classification performance was assessed considering different inference conditions, resembling real-time disturbance, by introducing different levels of noise corruption deriving from electrode motion artifacts, based on the noise samples from the MIT-BIH Noise Stress Test database.

Integer 8-bit inference resulted in 98.97% accuracy on the noiseless signal, which is only 0.65% lower than the state-of-the-art transformer model, although reducing by $60\times$ the number of parameters, and by $300\times$ the number of operations. The lean topology of the proposed model can be efficiently deployed on low-power devices for wearable monitoring, as demonstrated by the performance evaluated on the GAP9 parallel processor, where inference is executed in 4.28ms, with 0.09mJ energy consumption. The future directions for the development of this research include considering different implementation targets, aiming at improving the response time and energy efficiency of the proposed ECG processing system.

REFERENCES

- [1] "Cardiovascular diseases." World Health Organization. Accessed: Jun. 2023. [Online]. Available: https://www.who.int/health-topics/cardiovascular-diseases#tab=tab_1
- [2] S. M. A. Iqbal, I. Mahgoub, E. Du, M. A. Leavitt, and W. Asghar, "Advances in healthcare wearable devices," *npj Flexible Electron.*, vol. 5, no. 9, p. 9, 2021, doi: 10.1038/s41528-021-00107-x.
- [3] R. Hu, J. Chen, and L. Zhou, "A transformer-based deep neural network for arrhythmia detection using continuous ECG signals," *Comput. Biol. Med.*, vol. 144, 2022, Art. no. 105325. Accessed: 1 Aug. 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0010482522001172>
- [4] G. Yan, S. Liang, Y. Zhang, and F. Liu, "Fusing transformer model with temporal features for ECG heartbeat classification," in *Proc. IEEE Int. Conf. Bioinf. Biomed. (BIBM)*, 2019, pp. 898–905.
- [5] C.-F. Zhao, W.-Y. Yao, M.-J. Yi, C. Wan, and Y.-I. Tian, "Arrhythmia classification algorithm based on a two-dimensional image and modified EfficientNet," *Comput. Intell. Neurosci.*, vol. 2022, 2022.
- [6] M. M. Farag, "A tiny matched filter-based CNN for inter-patient ECG classification and arrhythmia detection at the edge," *Sensors*, vol. 23, no. 3, p. 1365, 2023. [Online]. Available: <https://www.mdpi.com/1424-8220/23/3/1365>
- [7] M. A. Scrugli, D. Loi, L. Raffo, and P. Meloni, "An adaptive cognitive sensor node for ECG monitoring in the internet of medical things," *IEEE Access*, vol. 10, pp. 1688–1705, 2022.
- [8] W. Wang, J. Guan, X. Che, and W. Wang, "MS-MLP: Multi-scale sampling MLP for ECG classification," in *Proc. 30th Eur. Signal Process. Conf. (EUSIPCO)*, 2022, pp. 1288–1292.
- [9] A. A. Ahmed, W. Ali, T. A. A. Abdullah, and S. J. Malebary, "Classifying cardiac arrhythmia from ECG signal using 1D-CNN deep learning model," *Math.*, vol. 11, no. 3, p. 562, 2023. [Online]. Available: <https://www.mdpi.com/2227-7390/11/3/562>
- [10] G. Moody and R. Mark, "The impact of the MIT-BIH arrhythmia database," *IEEE Eng. Med. Biol. Mag.*, vol. 20, no. 3, pp. 45–50, May/Jun. 2001.
- [11] A. Goldberger et al., "PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, pp. e215–e220, 2000.
- [12] "Ultra low power GAP processors." GreenWaves. Accessed: 1 Aug. 2023. [Online]. Available: <https://greenwaves-technologies.com/low-power-processor/>
- [13] D. L. T. Wong, Y. Li, D. John, W. K. Ho, and C.-H. Heng, "Low complexity binarized 2D-CNN classifier for wearable edge AI devices," *IEEE Trans. Biomed. Circuits Syst.*, vol. 16, no. 5, pp. 822–831, Oct. 2022.
- [14] G. Sivapalan, K. K. Nundy, S. Dev, B. Cardiff, and D. John, "ANNet: A lightweight neural network for ECG anomaly detection in IoT edge sensors," *IEEE Trans. Biomed. Circuits Syst.*, vol. 16, no. 1, pp. 24–35, Feb. 2022.
- [15] J. Liu et al., "An ultra-low power reconfigurable biomedical AI processor with adaptive learning for versatile wearable intelligent health monitoring," *IEEE Trans. Biomed. Circuits Syst.*, vol. 17, no. 5, pp. 952–967, Oct. 2023.
- [16] C. Che, P. Zhang, M. Zhu, Y. Qu, and B. Jin, "Constrained transformer network for ECG signal processing and arrhythmia classification," *BMC Med. Inform. Decis. Making*, 2021, Art. no. 184, doi: 10.1186/s12911-021-01546-2.

- [17] A. Natarajan et al., "A wide and deep transformer neural network for 12-lead ECG classification," in *Proc. Comput. Cardiol.*, 2020, pp. 1–4.
- [18] L. Meng, W. Tan, J. Ma, R. Wang, X. Yin, and Y. Zhang, "Enhancing dynamic ECG heartbeat classification with lightweight transformer model," *Artif. Intell. Med.*, vol. 124, 2022, Art. no. 102236. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0933366572200001X>
- [19] J. Pan and W. J. Tompkins, "A real-time QRS detection algorithm," *IEEE Trans. Biomed. Eng.*, vol. BME-32, no. 3, pp. 230–236, Mar. 1985.
- [20] D. Alexey et al., "An image is worth 16x16 words: Transformers for image recognition at scale," in *Proc. Int. Conf. Learn. Representations*, 2021.
- [21] A. Vaswani et al., "Attention is all you need," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 2017, 2017, p. 5999–6009.
- [22] A. Burrello et al., "Bioformers: Embedding transformers for ultra-low power sEMG-based gesture recognition," in *Proc. IEEE Design Automat. Test Eur. Conf. Exhibition (DATE)*, 2022, pp. 1443–1448.
- [23] P. Busia et al., "EEGformer: Transformer-based epilepsy detection on raw EEG traces for low-channel-count wearable continuous monitoring devices," in *Proc. IEEE Biomed. Circuits Syst. Conf. (BioCAS)*, 2022, pp. 640–644.
- [24] T. Wang, C. Lu, Y. Sun, M. Yang, C. Liu, and C. Ou, "Automatic ECG classification using continuous wavelet transform and convolutional neural network," *Entropy*, vol. 23, no. 1, p. 119, 2021. [Online]. Available: <https://www.mdpi.com/1099-4300/23/1/119>
- [25] Association for the Advancement of Medical Instrumentation American National Standards Institute. *Testing and Reporting Performance Results of Cardiac Rhythm and ST-segment Measurement Algorithms*(ser. ANSI/AAMI). The Association, 1999. [Online]. Available: <https://books.google.it/books?id=gzPdtgAACAAJ>
- [26] B. Moody, W. K. Muldrow, and R. G. Mark "A noise stress test for arrhythmia detectors," *Comput. Cardiol.*, vol. 11, pp. 381–384, 1984.
- [27] M. Spallanzani, G. Rutishauser, M. Scherer, A. Burrello, F. Conti, and L. Benini, "QuantLab: A modular framework for training and deploying mixed-precision NNs," Accessed: 1 Aug. 2023. [Online]. Available: <https://cms.tinyml.org/wp-content/uploads/talks2022/Spallanzani-Matteo-Hardware.pdf>
- [28] S. Kim, A. Gholami, Z. Yao, M. W. Mahoney, and K. Keutzer, "I-BERT: Integer-only BERT quantization," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 5506–5518.
- [29] "Inference: Tiny. v1.0 Results." ML Commons. Accessed: Aug. 1, 2023. [Online]. Available: <https://mlcommons.org/en/inference-tiny-10/>
- [30] A. Burrello, A. Garofalo, N. Bruschi, G. Tagliavini, D. Rossi, and F. Conti, "DORY: Automatic end-to-end deployment of real-world DNNs on low-cost IoT MCUs," *IEEE Trans. Comput.*, vol. 70, no. 8, pp. 1253–1268, Aug. 2021.
- [31] A. Burrello, M. Scherer, M. Zanghieri, F. Conti, and L. Benini, "A microcontroller is all you need: Enabling transformer execution on low-power IoT endnodes," in *Proc. IEEE Int. Conf. Omni-Layer Intell. Syst. (COINS)*, 2021, pp. 1–6.



Matteo Antonio Scrugli received the M.S. degree with honors from the University of Cagliari, in 2018 and the Ph.D. degree in electrical engineering from the University of Cagliari, in 2022, where he is currently a Research Fellow with the DIEE. His research interests include development of systems capable of managing at runtime the hardware and software configuration of low-power devices in order to adapt it to the required operating mode. His activity is currently focused on cognitive IoT devices, based on single-core or multi-core platforms.



Victor Jean-Baptiste Jung received the bachelor's degree in computer science and engineering physics from Juniata College, and the master's degree in computer science from the Institut Supérieur de l'Electronique et du Numérique of Lille (ISEN Lille), in 2022. After three months as a Research Intern with KU Leuven's MICAS Research group supervised by Prof. Marian Verhelst, he is currently working toward the Ph.D. degree with the Integrated Systems Laboratory with Prof. Dr. Luca Benini. His research interests include efficient deployment

of ML models on Microcontrollers, Tiny Transformers, Scheduling, and Quantization.



Luca Benini received the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, USA, in 1997. He is the Chair of digital circuits and systems with ETH Zürich, Zürich, Switzerland, and a Full Professor with the University of Bologna, Bologna, Italy. His research interests include energy-efficient computing, machine learning hardware, and smart microsystems. He is a Member of the Academia Europaea and a Fellow of the ACM.



Paola Busia received the Ph.D. degree in electronics and computer engineering from the Università degli studi di Cagliari, Italy, in 2023. She is currently an Assistant Professor with the Università degli studi di Cagliari. Her research interests include power efficient processing architectures for embedded applications and for at-the-edge AI.



Paolo Meloni is an Associate Professor with the University of Cagliari. His research interests include development of advanced digital systems, on the application-driven design and programming of multi-core on-chip architectures, and FPGAs.