

RESEARCH ARTICLE

Leveraging Statistical Shape Priors in GAN-Based ECG Synthesis

NOUR NEIFAR¹, ACHRAF BEN-HAMADOU², AFEF MDHAFFAR¹,
MOHAMED JMAIEL¹, AND BERND FREISLEBEN³

¹ReDCAD Lab, ENIS, University of Sfax, Sfax, Tunisia

²Laboratory of Signals, Systems, Artificial Intelligence and Networks (SM@RTS), Centre de Recherche en Numérique de Sfax, Technopôle de Sfax, Sfax 3021, Tunisia

³Department of Mathematics and Computer Science, Philipps-Universität Marburg, 35043 Marburg, Germany

Corresponding author: Nour Neifar (nour.neifar@redcad.org)

ABSTRACT Electrocardiogram (ECG) data collection during emergency situations is challenging, making ECG data generation an efficient solution for dealing with highly imbalanced ECG training datasets. In this paper, we propose a novel approach for ECG signal generation using Generative Adversarial Networks (GANs) and statistical ECG data modeling. Our approach leverages prior knowledge about ECG dynamics to synthesize realistic signals, addressing the complex dynamics of ECG signals. To validate our approach, we conducted experiments using ECG signals from the MIT-BIH arrhythmia database. Our results demonstrate that our approach, which models temporal and amplitude variations of ECG signals as 2-D shapes, generates more realistic signals compared to state-of-the-art GAN based generation baselines. Our proposed approach has significant implications for improving the quality of ECG training datasets, which can ultimately lead to better performance of ECG classification algorithms. This research contributes to the development of more efficient and accurate methods for ECG analysis, which can aid in the diagnosis and treatment of cardiac diseases.

INDEX TERMS GAN, deep learning, ECG, time series, physiological signals.

I. INTRODUCTION

Deep learning has been successfully applied in a wide range of fields, such as natural language processing, computer vision, and e-health systems [1], [2], [3]. Nevertheless, one of the persistent challenges in these fields, particularly in medicine, is the scarcity of training data [4], [5]. Data scarcity primarily arises from challenges associated with both data sharing and the complexities of data collection. Indeed, several ethical and legal regulations are imposed [6] such as the General Data Protection Regulation (GDPR) [7] in the European Union that set guidelines for ensuring confidentiality of patients and sharing of data. In addition, collecting medical data is a challenging process due to time-consuming and economic constraints. In particular, collecting pathological data during emergencies, such as heart attacks or epileptic seizures, is extremely difficult due to their unpredictable

nature, resulting in highly imbalanced datasets. Consequently, various techniques for ECG data synthesis have recently emerged in an attempt to address ECG data scarcity issues by providing additional and augmented data to support the training of downstream machine learning models, notably through the exploration of deep generative models [4], [5]. Among these models, Generative Adversarial Networks (GANs) have emerged as a widely adopted and popular choice known with its ability to seamlessly incorporate various forms of ECG shape prior knowledge [4], [5], [8], [9]. In particular, Generative Adversarial Networks (GANs) have shown their efficiency in generating high-quality facial images [10], medical imaging [11], [12], videos [13], and time series such as electrocardiograms (ECGs) [14], [15], [16]. In generating ECGs, one of the key challenges is capturing the complex dynamics of ECG signals. ECG signals are time series of physiological significance generated by recording electrical activities of the heart over time. They are characterized by a sequence of heartbeats, each of which corresponds to a cardiac

The associate editor coordinating the review of this manuscript and approving it for publication was Tao Huang¹.




FIGURE 1. Illustration of ECG waves including the P wave followed by the QRS complex and the T wave.

cycle represented by a series of electrical and mechanical events. A normal ECG signal exhibits distinct patterns for each event, including a P wave, a QRS complex, and a T wave, each with its own unique shape and timing (see Figure 1). The morphology and timing of ECG signals can vary significantly across individuals, making it challenging to generate realistic ECG signals that accurately capture the dynamics of physiological variations. Moreover, generating realistic ECG signals requires addressing several additional challenges, such as dealing with noise and artifacts in the ECG signal, handling variations in lead placement, and accounting for the effects of different body positions on the signal. These challenges are further complicated by the fact that pathological ECG signals can exhibit complex and irregular patterns, making it difficult to generate realistic synthetic data for training ECG classification models. These challenges underscore the importance of developing a more sophisticated approach to ECG signal generation that integrates prior knowledge of ECG dynamics and morphology, aiming to produce realistic and diverse ECG signals.

The main focus of this study is to leverage prior knowledge of ECG signal properties in order to generate accurate and realistic ECG waveforms. For this purpose, we integrate the capabilities of generative adversarial networks (GANs) with statistical shape modeling of ECG signals, enabling better control over the generation process.

Previous approaches based on deep learning models to generate synthetic ECG data [17], [18], [19] typically employ standard GAN architectures, which, however, fail to consider the dynamic properties inherent in complex physiological signals such as ECG. Several recent studies [8], [9], [20], [21] have made the efforts to address this challenge by integrating customized prior knowledge of ECG dynamics and patterns into the generation process, such as localization and order of ECG signal peaks. However, all of these features in [8], [9], and [20] are considered to be handcrafted. Indeed, they are based on domain-specific knowledge of ECG signals. Different from the approaches in [8], [9], and [20], our previous solution presented in [21] aimed to learn prior knowledge by utilizing ECG shape representatives and modeling the 1-D pattern dynamics of ECGs by disentangling the temporal and amplitude variations of the signal. However,

it should be noted that the dissociation between temporal and amplitude variations may not be as sufficient and effective in accurately modeling the complex patterns and variations.

In this paper, we present a novel approach that incorporates statistical prior knowledge of ECG signal dynamics into the generation process using GANs. This allows us to efficiently control the generation process and produce more realistic and diverse ECG signals. We designed a set of statistical 2-D shape models that incorporate prior knowledge about the general shape of main ECG signal clusters as well as the variability within signals belonging to the same cluster. Our contributions in this work can be summarized as follows:

- We present an efficient method for modeling 2-D shape variations in ECG data using statistical shape models.
- We introduce a novel GAN architecture that leverages statistical shape models as prior knowledge for a more realistic and stable ECG signal generation.
- We provide an extensive experimental evaluation, including both qualitative and quantitative assessments of our approach on the publicly available MIT-BIH arrhythmia benchmark, demonstrating the advantages of including statistical shape modeling in the generation process and the robustness of our method for generating ECG signals with realistic morphology.

The remaining of the paper is structured as follows: Section II discusses the related work, Section III presents our proposed approach, Section IV describes the experimental evaluation and the obtained results, and finally, Section V summarizes our findings and outlines potential future research directions.

II. RELATED WORK

In this section, we describe the principles of GANs and provide an overview of existing methods for generating ECG signals. GANs are powerful tools for data generation. They consist of two neural networks: a generator network and a discriminator network. Taking a noise vector as its input, the generator attempts to generate data that are similar to real data. The discriminator, on the other hand, attempts to distinguish between real and fake data (*i.e.*, the output of the generator). The training process is defined as a competition between these two networks. In other words, the discriminator's goal is to learn to distinguish between real and fake data, whereas the generator's goal is to fool the discriminator.

Several GAN-based methods for ECG signal synthesis have been proposed in the literature. They can be divided into two categories. The first category includes methods based on the adaptation of standard GANs. The second category of methods focuses on integrating prior knowledge of ECG signal dynamics into deep generative models in order to improve their ability to generate realistic ECG data.

A. ECG GENERATION USING STANDARD GANS

Wang et al. [22] proposed a simplified GAN architecture based on fully connected (FC) layers for generating ECG

heartbeat data. The authors qualitatively validated their models using a loss graph generated during the training process and by a visual comparison between the generated ECG heartbeat data and real ECG heartbeat data.

Antczak [23] presented an ECG generation method called ECG-GAN. The authors made few adjustments to the architecture of the deep convolutional generative adversarial network (DCGAN). To validate their approach, the generated data are added to real data to train an auto-encoder designed for filtering noise in the ECG signals.

Nankani and Baruah [24] proposed a conditional DCGAN model to synthesize different heartbeat classes. The synthesized heartbeats were quantitatively evaluated using various statistical metrics such as maximum mean discrepancy (MMD) and dynamic time warping (DTW).

Delaney et al. [17] designed different GAN models combining convolutional neural networks (CNN) and Long Short-Term Memory (LSTM) in order to generate time series data like ECG. The authors computed the MMD and DTW metrics after each training epoch to quantitatively evaluate the proposed models.

Zhu et al. [18] introduced a BiLSTM-CNN GAN framework for ECG generation. Three criteria were considered to evaluate the generated ECG signals: Fréchet distance (FD), percent root mean square difference (PRD) and root mean squared error (RMSE).

Hazra and Byun [25] developed SynSigGan, a framework for generating biomedical signals including ECG. It is based on a bidirectional grid LSTM (BiGridLSTM) for the generator network and a CNN for the discriminator network. The authors evaluated the quality of the generated synthetic signals using various metrics such as the Mean Absolute Error (MAE), RMSE, PRD, and FD scores. They assessed the correlation between real signals and generated synthetic signals by calculating the Pearson Correlation Coefficient (PCC). Their findings revealed that the synthetic data are highly correlated with the original data, and combining BiGridLSTM-CNN with a GAN produces better results.

Brophy [19] presented a Multivariate GAN to synthesize physiological, dependent and multivariate time series data. The generator and discriminator networks are based on LSTM and CNN layers, respectively. The multivariate dependent DTW (DTW_D) was proposed and computed at the end of each epoch to evaluate the generated data, in addition to the MMD metric. The (DTW_D) was used to compare the similarity between the dependent multi-channel real data and the generated data.

Later, Brophy et al. [26] improved their previous work by developing a GAN for the generation of dependent multivariate medical time series data. They looked into using the Loss Sensitive GAN (LS-GAN) objective function as a new loss function and adapted it to multivariate time series generation. To evaluate the proposed approach, the multivariate Dynamic Time Warping (MVDTW) and MMD metrics are used.

Golany and Radinsky [27] developed a set of generative models based on the DCGAN architecture in order to synthesize different classes of ECG heartbeats. To assess the quality of the synthetic data, the authors used it as additional training data to evaluate its impact on the performance of a neural sequence classifier.

Shaker et al. [28] developed a fully-connected layered GAN to generate ECG heartbeats. The synthetic data were combined with the training set to train two deep CNNs classification approaches to classify 15 arrhythmia types from the MIT-BIH dataset.

Yang et al. [29] proposed a gradually growing generative framework called ProEGAN-MS. Beginning with a reduced resolution of ECG signals, the two networks progressively grow throughout the training process by incrementing the convolutional layer appropriate to the features extraction. The authors assessed the fidelity and diversity of the synthetic data, as well as its impact on the performance of classification algorithms.

Yi et al. [30] introduced the TCGAN framework for ECG generation. The proposed architecture consists of a transformer generator and CNN discriminator. The generated ECG heartbeats were utilized as supplementary data alongside the real dataset to address the issue of data imbalance in ECG classification.

In [31], Jing et al. proposed the ECG-ADGAN model for ECG synthesis and cardiac abnormalities detection. A Bi-directional Long-Short Term Memory (Bi-LSTM) layer is used in the GAN generator and a mini-batch discrimination is used for the training of the proposed framework. The generated data were used with the real data for one-class ECG classification task.

B. ECG GENERATION USING CUSTOMIZED GANS

Golany et al. [20] introduced a Personalized GAN (PGAN) to generate patient-specific synthetic ECG signals. A customized loss was introduced in the training process of the generator to imitate the real heartbeats morphology. This loss is a combination of the standard cross-entropy loss and the mean squared error (MSE) which aims to produce natural waveforms similar to the real waveforms detected by the NeuroKit algorithm [32]. The generated ECG signals were combined with the real training set to train a patient-specific arrhythmia classifier to demonstrate the effect of adding synthetic data on the classification performance.

Later, Golany et al. [8] presented the SimGAN, a GAN framework that has been enhanced with supplementary knowledge from an ECG simulator. This ECG simulator was defined by three ordinary differential equations (ODE) and aims to understand the dynamic nature of ECG signal and to represent its heart dynamics. A specific loss was added to the standard cross-entropy loss of the generator optimization in order to synthesize fake heartbeats that are morphologically close to the heartbeats produced from the ECG simulator.

TABLE 1. Overview of the related work.

Papers	Year	Input Generator	Output Generator		Prior Knowledge modeling		Evaluation	Database
			end to end	decomposition	Handcrafted	Learned		
[17]	2019	noise	✓	✗	✗	✗	Qualitative eval., statistical metrics	MIT-BIH arrhythmia
[18]	2019	noise	✓	✗	✗	✗	statistical metrics	MIT-BIH arrhythmia
[22]	2020	noise	✓	✗	✗	✗	Qualitative eval.	MIT-BIH arrhythmia
[23]	2020	noise	✓	✗	✗	✗	Qualitative eval. Denoiser ECG	PTB-XL
[24]	2020	noise, class label	✓	✗	✗	✗	statistical metrics	MIT-BIH arrhythmia
[25]	2020	noise	✓	✗	✗	✗	Qualitative eval., statistical metrics	MIT-BIH arrhythmia
[19]	2020	noise	✓	✗	✗	✗	Qualitative eval., statistical metrics	MIT-BIH arrhythmia
[27]	2020	noise	✓	✗	✗	✗	Qualitative eval., ECG classification	MIT-BIH arrhythmia
[28]	2020	noise	✓	✗	✗	✗	Qualitative eval., ECG classification	MIT-BIH arrhythmia
[26]	2021	noise	✓	✗	✗	✗	Qualitative eval., statistical metrics	MIT-BIH arrhythmia, MIT-BIH NSR
[29]	2021	noise	✓	✗	✗	✗	Qualitative eval., ECG classification, statistical metrics	MIT-BIH arrhythmia
[30]	2023	noise	✓	✗	✗	✗	Qualitative eval., ECG classification	MIT-BIH arrhythmia
[31]	2023	noise	✓	✗	✗	✗	ECG classification	MIT-BIH arrhythmia
[20]	2019	noise	✓	✗	✓	✗	ECG classification	MIT-BIH arrhythmia
[8]	2020	noise	✓	✗	✓	✗	Qualitative eval., ECG classification	MIT-BIH arrhythmia
[9]	2021	Physical parameters of ECG beat, initial voltage value of ECG beat	✗	Ordinary Differential Equations	✓	✗	Qualitative eval., ECG classification	MIT-BIH arrhythmia
[21]	2022	noise, class label	✗	Temporal and amplitude variations	✗	✓	Qualitative eval., ECG classification	MIT-BIH arrhythmia
Our approach	2023	noise, class label	✗	2-D shape variations	✗	✓	Qualitative eval., Cardiologists eval., ECG classification, statistical metrics	MIT-BIH arrhythmia

More recently, Golany et al. [9] proposed the ECG-ODE-GAN framework where the generator is defined as an ODE to learn the dynamics of ECG signals. The authors also demonstrated how to incorporate physical considerations into the ECG-ODE-GAN by introducing physical parameters that characterize the ECG signals as supplementary input to the generator. To assess their model, the generated data were used to train heartbeats classifier.

In our previous work [21], we introduced a new GAN framework for generating ECG heartbeats. We demonstrated how to integrate prior knowledge about ECG waveforms by modeling the ECG signals as shape clusters (*i.e.*, signal averages) that serve as references to represent ECG patterns in the training distribution.

Separating between the amplitude and temporal variations in the data modeling was also introduced in the generation process to improve the modeling of the ECG complex characteristics. The proposed approach was evaluated by combining the synthetic heartbeats as additional data with the real training set to train three arrhythmia classification baselines.

C. OUR METHOD POSITIONING

Table 1 summarizes the previously discussed approaches. Although the solutions belonging to the first category can generate ECG heartbeats, they are still limited to synthesize “realistic” ECG signals “with morphological proprieties”. This is due to the use of the standard GAN architecture that does not include specific knowledge related to ECG signals.

Due to the complex dynamic nature of ECGs, recent solutions [8], [9], [20] are based on GAN models enriched with prior knowledge of ECG signal patterns to generate synthetic ECG heartbeats with real morphologies. They typically rely on integrating handcrafted physics knowledge.

As a result, methods for specifying data properties were required during the generation process such as using specific simulators, adding complementary algorithms to identify specific patterns in the waveforms. Also, the model in [8] is constrained by the capabilities of the ECG simulator and limited to synthesize heartbeats with morphology defined by a certain dynamical system. Moreover, generating ODEs using GAN in [9] remains a difficult task as it increases the model complexity. The training process is also highly unstable, and despite multiple efforts, attaining reproducibility proves elusive [33]. In our previous work [21], we suggested employing ECG anchors to leverage a learned prior knowledge into the generation process. However, in this previous solution, we consider 1-D ECG pattern dynamics modeling by dissociating the temporal and amplitude variations of ECG signals. We acknowledge that this decomposition of variations may not accurately capture the complexity of patterns and variations observed in ECG signals.

In contrast to all previous solutions, we believe that ECG signal generation requires an advanced prior knowledge modeling of the ECG shape and dynamics for better control

over the generation process. Statistical shape models have demonstrated their effectiveness in capturing shape variations in various applications, including facial analysis. Therefore, in this work, we propose the first ECG generation approach that integrates statistical shape modeling as prior knowledge of ECG signal 2-D dynamics and shape variations. The 2-D representation is introduced to enable statistical shape modeling of the ECG dynamics in both voltage and time axes. It allows for explicit modeling of the inherent temporal and voltage variation within the ECG signal dynamics. In addition, as each 2-D point is by construction referring to the same specific location on the ECG signal, our approach allows for the implicit encoding of the semantics of the generated ECG signals. In contrast, the naive approach to statistical modeling of ECG signals would be to apply it directly to a 1-D representation, which, however, only covers voltage variations.

The integration of statistical shape modeling offers several advantages.

Firstly, the statistical shape model provides a compact representation of the shape variations present in ECG signals. This enables the GAN architecture to generate a wide range of samples that align with the inherent shape characteristics and dynamics of real ECG signals. This integration enables the generation approach to effectively learn and produce realistic variations in shape while preserving the essential structural properties of the ECG waveform.

Moreover, the integration of a statistical shape model allows for the encoding of prior knowledge regarding shape variations. This prior knowledge serves as a valuable guidance mechanism for GANs during the generation process. By incorporating the statistical shape model, the GAN can leverage this prior knowledge to improve the fidelity and accuracy of the generated ECG signals. This ensures that the generated waveforms are not only diverse but also adhere to the expected shape characteristics of real ECG signals, enhancing the realism of the generated outputs.

III. PROPOSED APPROACH

The workflow of our proposed method is shown in Figure 2. We first start with a processing block designed for statistical modeling. Then, we move on to the training phase of our GAN model. The generator takes a noise vector and data labels as input to generate a specific linear combination that will be used to the output of the statistical model. Afterwards, the real ECG heartbeats, as well as the synthetic heartbeats, are passed to the discriminator, which attempts to distinguish between the real and fake heartbeats. The underlying principles are detailed below.

A. STATISTICAL MODEL CONSTRUCTION

The aim of our statistical model is to accurately model ECG signal shape variations. It consists of three consecutive steps: clustering, alignment, and statistical modeling for signal decomposition. The pseudo-code of this phase is outlined in Algorithm 1. First, the training dataset is composed of C

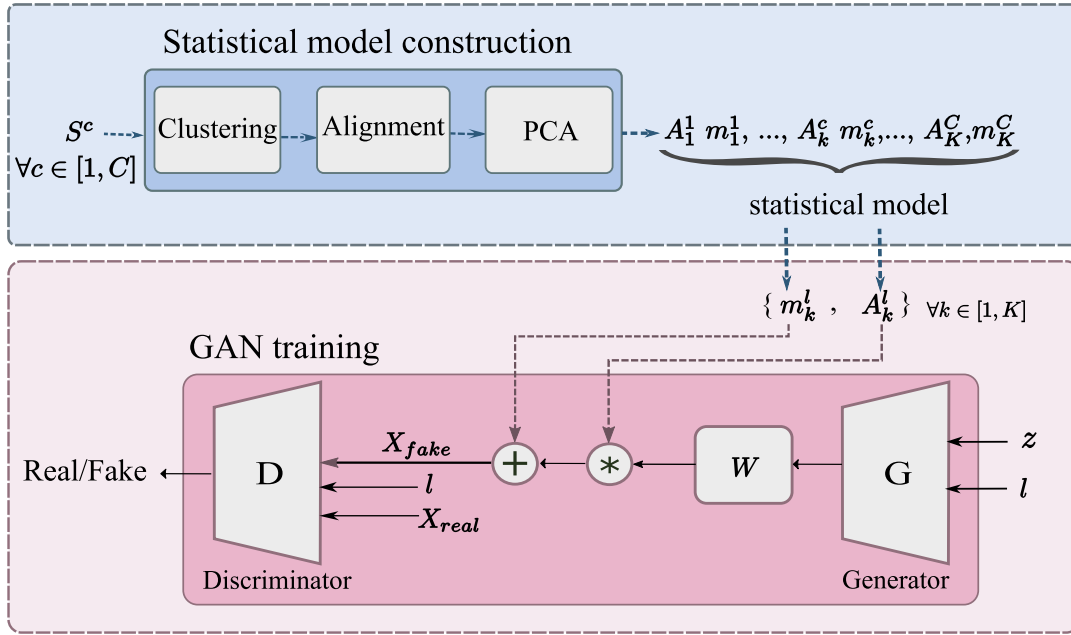


FIGURE 2. Workflow of the proposed approach.

subsets of signals $\{\mathcal{S}\}^c$ grouped by class, where $c \in [1, C]$ is the class index and \mathcal{S} refers to one signal. As discussed above, the ECG signal is represented by its 2-D point coordinates (*i.e.*, position in time and amplitude, respectively) where $\mathcal{S} \in \mathbb{R}^{2 \times T}$ and T is the length of the signal.

The signals of each subset are clustered into fixed K clusters of signals noted by $\{\mathcal{S}\}_k^c$, where k is the cluster index. For this step, we use the modified version of the original K-means algorithm, specifically designed for time-series data clustering [34]. In particular, we apply the DTW-based K-means for ECG clustering. The used implementation is provided in the Python package `tslearn`.¹

To model the time and amplitude variations of homologous points inside a cluster of signals, we must first localize the position of each homologous point in every signal of $\{\mathcal{S}\}_k^c$. As an example, Figure 3 shows a plot of homologous points corresponding to the peak R localized in different signals that belong to a same cluster.

Algorithm 1 Statistical Model Construction

Data: Number of classes C , Number of clusters K

for $c:=1$ to C **do**

$\{\mathcal{S}\}_k^c = \text{K-Means}(\{\mathcal{S}\}^c, K)$;

for $k:=1$ to K **do**

$\{\mathcal{p}\}_k^c = \text{DTAN}(\{\mathcal{S}\}_k^c)$;

$\{\hat{\mathcal{S}}\}_k^c = \text{interpolate}(\{\mathcal{S}\}_k^c, \{\mathcal{p}\}_k^c)$;

$\mathbf{X}_k^c = \begin{bmatrix} \{\hat{\mathcal{S}}[0]\}_k^c \\ \{\hat{\mathcal{S}}[1]\}_k^c \end{bmatrix}$;

$\mathbf{m}_k^c, \mathbf{A}_k^c = \text{PCA}(\mathbf{X}_k^c)$;

We use a Diffeomorphic Temporal Alignment Net (DTAN) [35] to align the signals $\{\mathcal{S}\}_k^c$. A DTAN is an efficient way to align time series via flexible temporal transformer layers. Compared to other existing solutions for data alignment, a DTAN shows better results on different types of time series data including ECG signals of the ECGFiveDays dataset [36]. Figure 4(a) and Figure 4(b) show examples of normal and pathological ECG signals before and after alignment.

In addition to the aligned signals, DTAN provides the positions $\{\mathcal{p}\}_k^c$ that will be used later to interpolate their corresponding time and amplitude values from $\{\mathcal{S}\}_k^c$ yielding to $\{\hat{\mathcal{S}}\}_k^c$. Then, the set of signals $\{\hat{\mathcal{S}}\}_k^c$ is arranged in a matrix $\mathbf{X}_k^c \in \mathbb{R}^{N_k^c \times 2T}$ by concatenating their time and amplitude values.

N_k^c is the total number of signals in $\{\mathcal{S}\}_k^c$.

The final step is to apply Principal Component Analysis (PCA) to decompose \mathbf{X}_k^c into a mean vector $\mathbf{m}_k^c \in \mathbb{R}^{2T}$ (*i.e.*, shape average of signals) and variation matrix $\mathbf{A}_k^c \in \mathbb{R}^{B_k^c \times 2T}$ regrouping B_k^c eigenvectors. This decomposition allows us to represent any signal belonging to a cluster k and a class c within a linear combination between \mathbf{m}_k^c and \mathbf{A}_k^c .

B. GENERATIVE ADVERSARIAL NETWORK ARCHITECTURE

The generator $G(z, l)$ receives as input a random noise vector z sampled from a normal distribution $\mathcal{N}(0,1)$ and a label $l \in [1, C]$ conditioning the class of the signals to generate, and outputs a matrix \mathbf{W} . \mathbf{W} gathers the weights used for a linear combination of the eigenvectors in \mathbf{A}_k^c where $c = l$. It is worth noticing that \mathbf{A}_k^c matrices have variable number of eigenvectors B_k^c . For the sake of readability, we define $\mathbf{W} \in \mathbb{R}^{K \times \max_k(B_k^c)}$ and we add row zero-padding on \mathbf{A}_k^c matrices so

¹`tslearn.clustering.TimeSeriesKMeans`

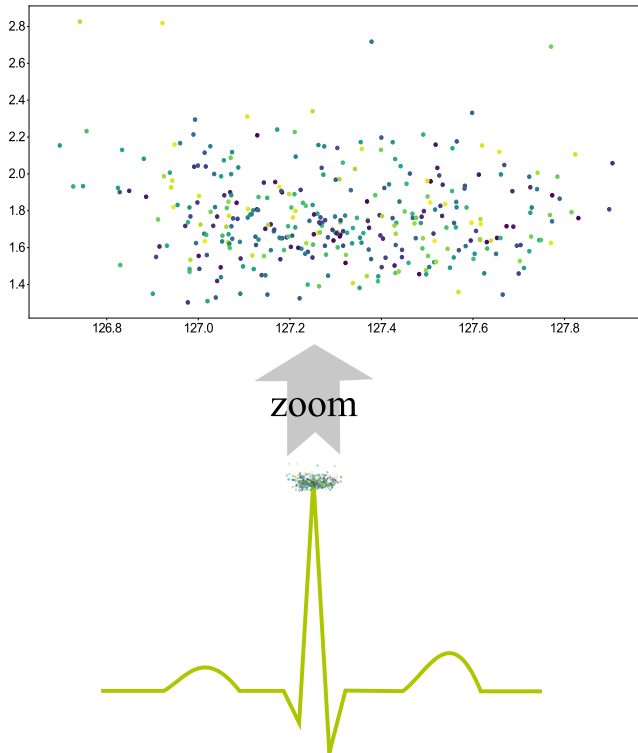


FIGURE 3. A set of homologous points corresponding to the peak R localized in different signals that belong to the same cluster.

that A_k^c becomes $\in \mathbb{R}_{c,k}^{\max(B_k^c) \times 2T} \forall c \in [1, C]$ and $k \in [1, K]$. In this way, the generation of fake signals X_{fake} is:

$$X_{fake} = M_l + W A_l \quad (1)$$

where: $W = G(z, l)$.

In Equation (1), the matrix $X_{fake} \in \mathbb{R}^{K \times 2T}$ arranges one fake signal by each cluster. $M_l \in \mathbb{R}^{K \times 2T}$ is a matrix that regroups the means $m_k^l, \forall k \in [1, K]$. $A_l \in \mathbb{R}_{c,k}^{K \times \max(B_k^c) \times 2T}$ is a tensor gathering all A_k^l matrices $\forall k \in [1, K]$.

The architectures of our generator and discriminator networks are described in Table 2. All used convolution layers apply one-dimension (1-D) convolution. The generator’s input noise vector and labels are first concatenated and then passed through three convolution layers, each followed by a batch normalization layer and a ReLU activation function. The output of the final convolution layer is then fed into a batch normalization layer, which is then processed by an FC layer to obtain the desired output shape of W . The discriminator consists of three convolution layers. A batch normalization followed by a ReLU activation function are used between these three layers. The output of the last convolutional layer is passed to a batch normalization layer and a FC layer. This is followed by an LSTM layer, three FC layers and finally a sigmoid activation function to produce the classification scores. Despite the simplicity of the proposed architecture, various architectures, including the transformer network, were investigated for both the generator and discriminator. However,

no notable difference in results was observed across these architectures.

The training of GAN is a min-max game between its components, which can be expressed as follows:

$$\min_G \max_D \mathbb{E}_{z \sim P_g} [\log(1 - D(G(z, l)))] + \mathbb{E}_{X_{real} \sim P_{real}} [\log(D(X_{real}, l))] \quad (2)$$

where X_{real} represents the input data sampled from the real distribution P_{real} , $D(X_{real}, l)$ and $D(G(z, l))$ are the probabilities estimated by the discriminator for classifying real instances as real and fake instances as real, respectively. The generator attempts to minimize this loss by synthesizing fake instances that are similar to real instances to deceive the discriminator. On the other hand, the discriminator aims to maximize this loss function by maximizing both terms of this function. Therefore, it attempts to estimate a probability equal to 1 when fed X_{real} and 0 when fed X_{fake} . To stabilize the training of our models, we used a Wasserstein GAN with gradient penalty loss [37] instead of the original GAN loss. It consists of adding an extra penalty term to Equation (2).

$$\mathcal{L} = \mathbb{E}_{X_{fake} \sim P_{fake}} [D(X_{fake}, l)] - \mathbb{E}_{X_{real} \sim P_{real}} [D(X_{real}, l)] + \lambda \mathbb{E}_{\hat{X} \sim P_{\hat{X}}} [(\|\nabla_{\hat{X}} D(\hat{X}, l)\|_2 - 1)^2] \quad (3)$$

where $\mathbb{E}_{\hat{X} \sim P_{\hat{X}}} [(\|\nabla_{\hat{X}} D(\hat{X}, l)\|_2 - 1)^2]$ is the gradient penalty, and λ is the penalty coefficient used to weight the gradient penalty term. $P_{\hat{X}}$ is the distribution obtained by randomly interpolating between samples from the distributions P_{real} and P_{fake} .

IV. EXPERIMENTAL RESULTS

We conducted two types of experiments to evaluate our proposed method. First, we performed a qualitative evaluation to assess the synthetic ECG heartbeats. Then we performed a quantitative evaluation that consists of assessing the effect of adding our synthetic ECG signals to the real training set on three arrhythmia classifiers performance. The synthetic ECG signals were then assessed using various metrics.

A. EVALUATION DATABASE

ECG signals taken from the MIT-BIH arrhythmia database² were considered for our models training [38]. The ECG represents the main diagnostic procedure for detecting cardiac anomalies. Various cardiac abnormalities such as arrhythmia can be the sources of changes in the normal ECG patterns. The arrhythmia is a disease specified by a disorder in the cardiac rhythm accompanied by modifications of the ECG. The MIT-BIH arrhythmia database is widely recognized as the benchmark dataset for arrhythmia analysis, and it has been extensively used in previous studies for evaluation purposes. The database contains 48 half-hour ECG recordings, obtained

²<https://physionet.org/content/mitdb/1.0.0/>




FIGURE 4. (A): Examples of signals before and after alignment taken from the training dataset of the normal class (N). (B): Examples of signals before and after alignment taken from the training dataset of pathological class.

TABLE 2. Details of the generator and discriminator networks.

Generator	Input: (z, l)
	Layer 1: conv1 (kernel_size=1), BatchNorm, ReLU output (batch_size, 16, 1)
	Layer 2: conv1 (kernel_size=1), BatchNorm, ReLU, output (batch_size, 32, 1)
	Layer 3: conv1 (kernel_size=1), BatchNorm, output (batch_size, 64, 1)
Layer 4: FC, output (batch_size, $K \times \max_{c,k}(B_k^c)$)	
Discriminator	Input: $(X_{real}/X_{fake}, l)$
	Layer 1: conv1 (kernel_size=1), BatchNorm, ReLU. output (batch_size, 16, 1)
	Layer 2: conv1 (kernel_size=1), BatchNorm, ReLU. output (batch_size, 32, 1)
	Layer 3: conv1(kernel_size=1), BatchNorm. output (batch_size, 64, 1)
	Layer 4: FC. output (batch_size,128)
	Layer 5: LSTM(hidden_size=256, number_layers=3). output (batch_size,256)
	Layer 6: FC, Tanh. output (batch_size,128)
	Layer 7: FC, Tanh. output (batch_size,64)
Layer 8: FC, Sigmoid. output (batch_size,1)	

from 47 patients examined between 1975 and 1979 by the BIH Arrhythmia Laboratory. Each record contains two 30-minute ECG lead signals that have been annotated by cardiologists and digitized at 360 samples per second. Our experiments focus on MLII lead as it was consistently employed for recording one channel for all patients and is the most widely used configuration within the research community. This choice simplifies comparisons with other methods. Cardiac cycles can be divided into different categories that can represent normal heartbeats or arrhythmias. The P, Q, R, S, and T wave patterns vary by the arrhythmia category. In this paper, three types of heartbeats were considered. The first type is the normal beats

(class N). The second is the premature ventricular contraction beats (class V), and the last one is the fusion beats (class F).

B. TRAINING SETTINGS

Our models were implemented using the PyTorch framework and trained on an Ubuntu server version 20.04, with a GeForce GTX 1080 Ti GPU having 11 GB memory. The optimization of our GAN is performed using the ADAM algorithm and a learning rate of 0.00001. The batch size used is equal to 64. The penalty coefficient λ used in loss equation (3) is set to 10 as in [23] and [37]. The input z of the generator is a noise

vector of length 100 from the normal distribution $\mathcal{N}(0,1)$ as in [9], [20], [24], and [27]. Each ECG signal is divided into heartbeats with 270 voltage values ($T = 270$). The T values correspond to 350 ms before the R-peak and 400 ms after the R-peak. In this study, we randomly selected 70% of the data as training dataset, while the remaining 30% of the data was allocated for testing purposes. Based on several related works [39], [40], [41], we adopted a signal-wise paradigm when splitting the database where each signal is assigned to either the training or testing set independently of other signals. A low-pass filter was applied to the ECG signals to remove the noise before being fed to our models.

C. RESULTS

1) QUALITATIVE EVALUATION

As an initial step for the qualitative assessment of our method, a set of synthetic heartbeats was randomly selected and compared to the distributions of real heartbeats taken from the training dataset. Figure 5 shows examples of synthetic heartbeats belonging to the normal class and the pathological classes (V and F), in addition to different clusters of the real distribution of these classes. We can observe that the synthetic heartbeats have realistic morphology and are closely similar to real distributions. In addition, the cardiac cycles demonstrate faithful dynamics according to their corresponding classes, visually proving the ability of our approach in generating realistic ECG heartbeats.

For a more representative evaluation, the quality of the generated ECG heartbeats was further evaluated by three cardiologists. The assessment consists of presenting shuffled examples of real and synthetic cardiac cycles of each class to the cardiologists. The objective is to assess the generated signals by having cardiologists visually differentiate between synthetic and real ECGs. Additionally, we aim to determine whether these signals can be correctly classified into their respective classes, thereby evaluating both visual appearance and class-specific characteristics of synthetic signals.

For this purpose, we randomly chose a set of 30 samples of real heartbeats from the used training dataset (N, V, and F classes) and a set of 30 heartbeats generated by our approach (N, V, and F classes).

The three cardiologists successfully identified the classes of the presented cardiac cycles. First, we calculated the recognition rate of real and fake heartbeats as real, disregarding their categories. The recognition rate of real heartbeats as real represents the rate achieved by a cardiologist in validating that real heartbeats belong to their corresponding classes. The recognition rate of fake heartbeats as real corresponds to the validation rate that fake heartbeats belong to their classes. The first cardiologist validated 90% of the real heartbeats as real ones and 100% of the fake heartbeats as real ones. The second cardiologist achieved 96.6% and 90% as the recognition rates for real and fake heartbeats as real ones, respectively. The third cardiologist identified 96.6% of both real heartbeats as real heartbeats and fake heartbeats as

TABLE 3. Recognition rate of real and fake heartbeats as real heartbeats for the classes (N, V, F) of cardiologists 1, 2, and 3.

		Class N	Class V	Class F
Cardiologist 1	real as real	94.1%	85.7%	83.3%
	fake as real	100%	100%	100%
Cardiologist 2	real as real	100%	100%	83.3%
	fake as real	100%	71.4%	83.3%
Cardiologist 3	real as real	100%	100%	83.3%
	fake as real	100%	100%	83.3%

real heartbeats. The obtained recognition rates of each class are detailed in Table 3. The three cardiologists achieved good rates for both real and fake cases with the different types of heartbeats. These results show the robustness of our proposed approach in generating ECG signals similar to real signals with realistic wave morphology. For example, the second cardiologist correctly classified all real beats from class N as real and all synthetic beats as real indicating that he could not reliably distinguish between real and synthetic heartbeats. Additionally, the cardiologist correctly assigned class N to the synthetic heartbeats, which can explain that the generated ECG closely resembles real ECG heartbeats in both visual appearance and class-specific characteristics.

2) QUANTITATIVE EVALUATION

a: ECG CLASSIFICATION

The quantitative evaluation is first carried out by training three state-of-the-art heartbeat classification baselines [42], [43], [44] on our data following three settings:

- Setting 1: The training is performed only with real training dataset without adding synthesized ECG signals.
- Setting 2: The MIT-BIH arrhythmia real dataset is augmented by synthetic ECG signals generated using a standard GAN [45] taking as its input noise vectors and class embeddings. The generator network is composed of a succession of FC and LeakyReLU layers, and finally outputs an entire ECG heartbeat signal.
- Setting 3: Same as setting 2, but the fake signals are generated using a state-of-the-art advanced GAN [21].
- Setting 4: Same as setting 2, but the MIT-BIH arrhythmia real dataset is augmented by synthetic ECG signals obtained by our generation approach.
- Setting 5: Same as setting 2, but the fake signals are generated using a state-of-the-art GAN called 2-CNN-GAN [17]. This GAN includes an LSTM-based generator, and CNN-based discriminator with two convolution ReLU-pooling layers.

These settings lead to 15 trained models that are then validated on the same test set consisting of real ECG signals randomly selected from the MIT-BIH arrhythmia dataset unseen during the training phases of the GAN and classifiers.

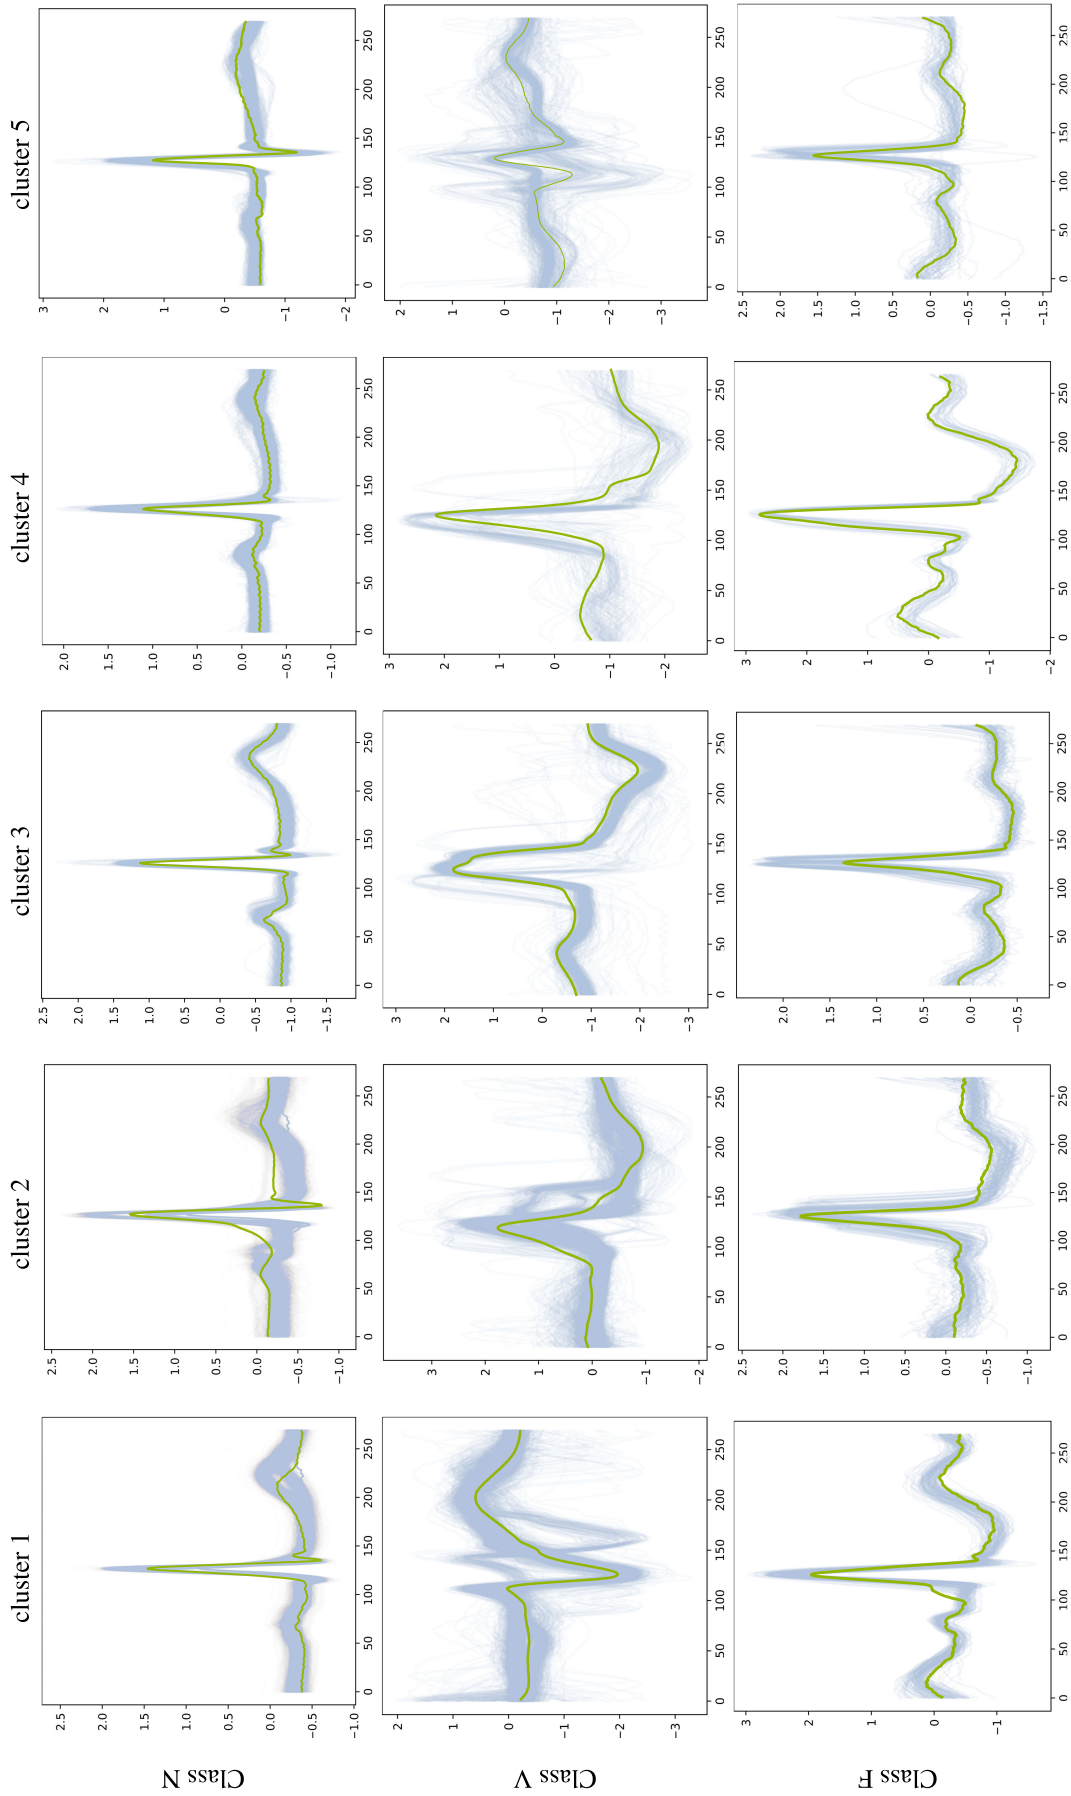


FIGURE 5. Examples of synthetic heartbeats for classes N,V and F. The blue background represents different clusters of the real dataset, while the green depicts synthetic heartbeats.

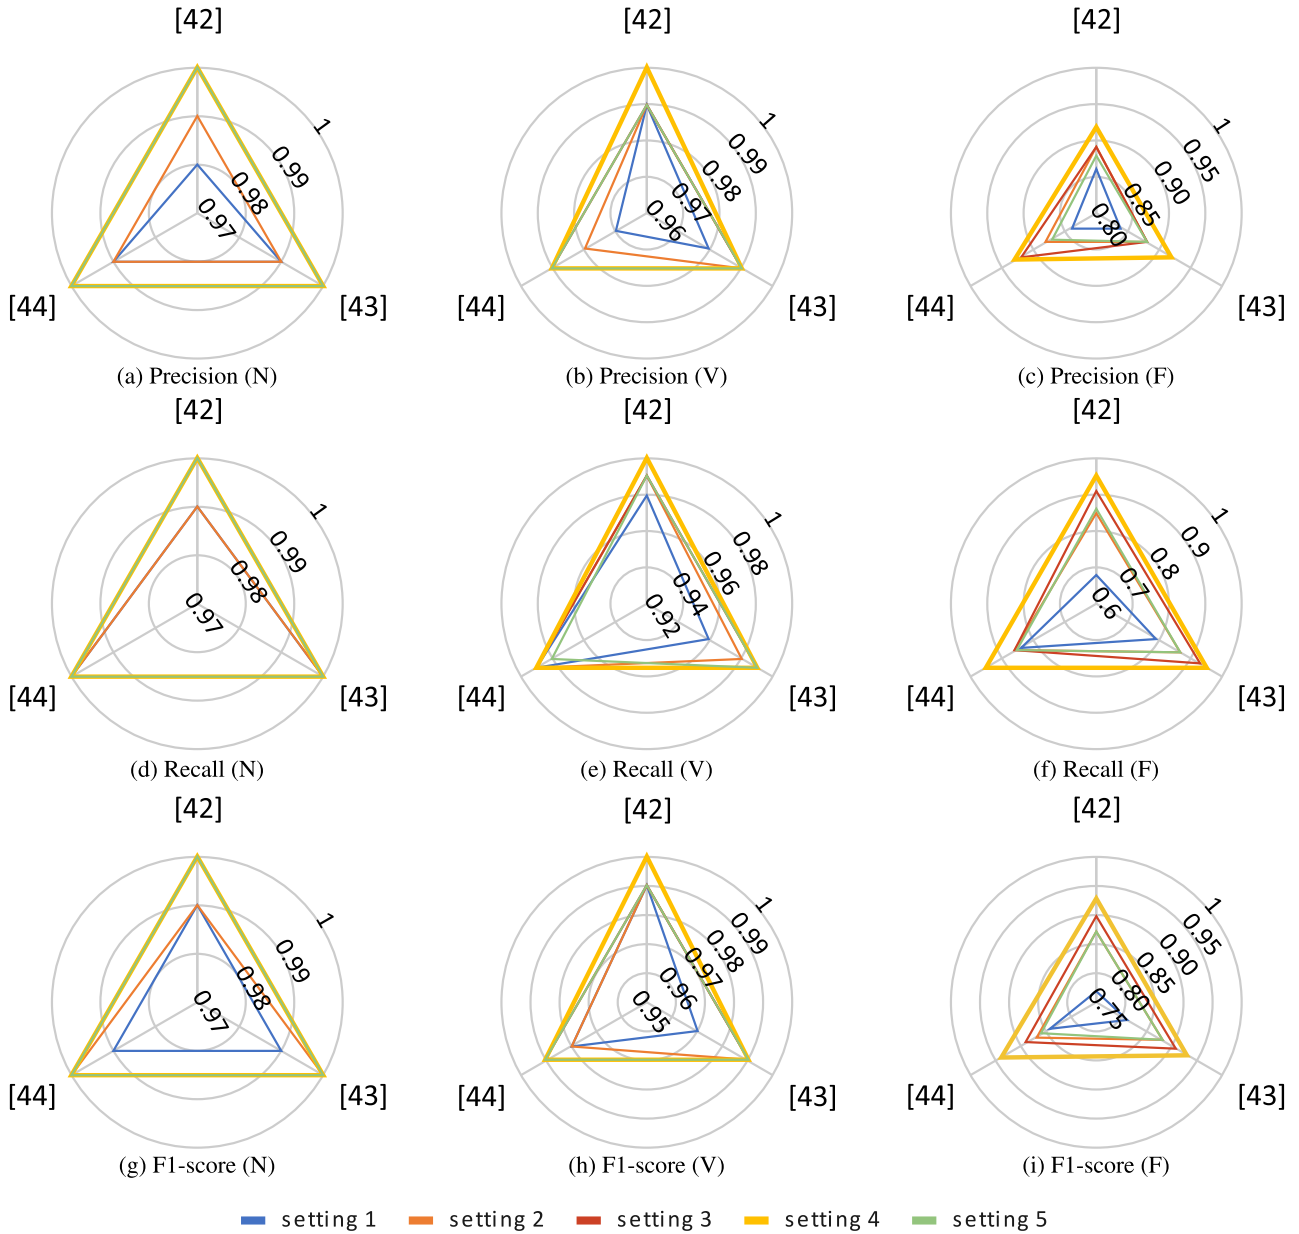


FIGURE 6. Performance results of the three classification baselines [42], [43], [44] on the classes N, V, and F. each axis of a triangle corresponds to a classification baseline.

The idea behind considering settings 1 and 4 is to highlight the contribution of data augmentation using our approach in improving ECG heartbeat classifiers, while the experiments related to settings 2, 3 and 5 demonstrate that our generation approach outperforms the standard GAN and the state-of-the-art standard and advanced GAN [21] for the purpose of ECG signal generation.

Before discussing the obtained results, we introduce the competing classification baselines. The classifier presented by Kachuee et al. [42] is composed of a 1-D convolutional layer fulfilled by five residual convolution sets, two FC layers and a softmax layer to generate the class probabilities. Every residual set is made up of two 1-D convolution layers,

two ReLU activation layers, a residual skip connection, and finally a pooling layer. The classifier architecture presented by Acharya et al. [43] consists of three 1-D convolution layers, each followed by a max-pooling layer. Then, three FC layers are applied to the output data with a softmax function at the end of the last layer. The classification model developed by Kumar et al. [44] is made up of four blocks, including a FC layer succeeded by a batch normalization layer and ReLU activation function. Finally, a FC layer and a softmax activation function are performed on the output of the last block.

In Figure 6, the radar charts in the shape of triangles provide a summary of the classification results achieved for the three

TABLE 4. Obtained evaluation metrics for our approach and state-of-the-art GAN approaches.

Methods	Classes	RMSE	MAE	MSE	EMD	DTW
[45]	Class N	1.86e-3	1.47e-3	3.45e-6	2.35e-2	29.05
	Class V	3.87e-3	2.82e-3	1.49e-5	4.96e-2	28.10
	Class F	5.37e-3	4.18e-3	2.88e-5	8.26e-2	8.63
[17]	Class N	1.85e-3	1.42e-3	3.45e-6	2.31e-2	28.91
	Class V	3.79e-3	2.78e-3	1.43e-5	4.89e-2	27.62
	Class F	5.49e-3	4.07e-3	3.02e-5	6.71e-2	8.35
[21]	Class N	1.83e-3	1.33e-3	3.38e-6	2.03e-2	25.67
	Class V	3.54e-3	2.64e-3	1.25e-5	3.52e-2	24.04
	Class F	5.16e-3	3.85e-3	2.66e-5	5.45e-2	7.47
Ours	Class N	1.66e-3	1.25e-3	2.76e-6	1.84e-2	24.42
	Class V	3.02e-3	2.22e-3	9.18e-6	2.52e-2	21.20
	Class F	4.84e-3	3.47e-3	2.34e-5	4.23e-2	6.61

heartbeat classes. The presented results demonstrate that using fake ECG heartbeats synthesized by our method as additional data with the training set of the real dataset improved the classification performance (*i.e.*, precision, recall, and F1 score). The triangles corresponding to setting 4 (in yellow) are not only larger than those of the other settings but setting 4 systematically improves performance regardless of the baseline and the class N, V and F. The three classifiers achieve significantly higher performance when adding synthetic ECG heartbeats of the minority class F to the training dataset. As an example, the classifier [42] achieved higher performance metrics in setting 4 compared to setting 1 for class F. The precision, recall, and F1 score were (0.92, 0.95, and 0.93) in setting 4, whereas they were (0.86, 0.68, and 0.77) in setting 1. The performance of the two other classes is faintly improved. For the same classifier, the performance has been improved by (2%, 1%, and 1%) and (1%, 1%, and 1%) from setting 1 to 4 in classes (N and V) respectively. Furthermore, the classification models trained with added fake ECG heartbeats, generated using our approach, significantly outperform classification models trained with added fake ECG heartbeats generated by the standard GAN approach and the state-of-the-art GAN in setting 5, with higher precision, recall, and F1-score values. It is also clear that the performance of the three classifiers in setting 4 is better than the classifiers performance in setting 3, where synthetic data from the state-of-the-art advanced GAN model are used. The classifier model [43] achieved in class F classification (0.93, 0.95, and 0.94) for (precision, recall, and F1 score) when using synthetic data from obtained by our approach in setting 4 while classifiers in settings 2, 3 and 5 obtained (0.88, 0.86, and 0.87), (0.92, 0.86, and 0.89) and (0.87, 0.85, and 0.86), respectively. These results show the robustness of our approach in generating higher quality and more realistic ECG signals compared to state-of-the-art GAN approaches, including both standard and advanced methods. This clearly indicates that incorporating statistical shape modeling and prior knowledge of ECG patterns significantly enhances the generation process, resulting in more realistic signals.

b: EVALUATION METRICS:

Several metrics were considered to assess further the quality of the generated signals based on [46], [47]: Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), Mean Squared Error (MSE), Earth Mover's Distance (EMD), Dynamic time warping (DTW). The RMSE, MAE, MSE and EMD metrics are expressed in millivolts (mV). The DTW is expressed in milliseconds (ms) in this study. The obtained values for our generation method and other GAN models are shown in Table 4. We can observe that our generation approach outperforms the other generation baselines across all heartbeats classes for the different metrics which demonstrates that our model is more effective in generating ECGs compared with other approaches. For instance, we achieved (1.66e-3, 1.84e-2) of (RMSE, EMD) in class N, while the standard GAN, the state-of-the-art GAN [17] and advanced GAN [21] achieved (1.86e-3, 2.45e-2), (1.85e-3, 2.31e-2) and (1.83e-3, 2.03e-2), respectively.

V. CONCLUSION

In this article, we proposed a novel GAN method for generating ECG signals. The proposed method leverages statistical shape modeling of the ECG signal dynamics to integrate prior knowledge of ECG patterns into the deep generation process. Furthermore, we proposed modeling ECG signals as a 2-D pattern representing its temporal and amplitude dynamics in order to generate realistic signals. The experimental results obtained on the MIT-BIH arrhythmia database showed the effectiveness and robustness of the proposed approach in synthesizing faithful ECG signals. Indeed, the synthetic ECG signals include a realistic ECG morphology and capture similar dynamics as real signals. Moreover, we demonstrated that our method outperforms other state-of-the-art ECG generation methods and can improve the performance of known arrhythmia classification algorithms. In future work, we plan to extend the validation of our approach to larger datasets covering different biological signals. In addition, we aim to investigate the transferability across different databases by exploring advanced techniques to bridge the gap between datasets, such as self-supervised learning and knowledge distillation. As another part of our future work, we are particularly interested in exploring diffusion models for physiological signals synthesis. During the clustering step, the number of clusters was optimized using a brute force approach, involving exhaustive iterations. However, future research should explore more efficient optimization methods to determine the optimal number of clusters for achieving the best variations modeling.

REFERENCES

- [1] D. W. Otter, J. R. Medina, and J. K. Kalita, "A survey of the usages of deep learning for natural language processing," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 2, pp. 604–624, Feb. 2021.
- [2] Q. Wu, Y. Liu, Q. Li, S. Jin, and F. Li, "The application of deep learning in computer vision," in *Proc. Chin. Autom. Congr. (CAC)*, Oct. 2017, pp. 6522–6527.

- [3] R. Zhao, R. Yan, Z. Chen, K. Mao, P. Wang, and R. X. Gao, "Deep learning and its applications to machine health monitoring," *Mech. Syst. Signal Process.*, vol. 115, pp. 213–237, Jan. 2019.
- [4] G. Monachino, B. Zanchi, L. Fiorillo, G. Conte, A. Auricchio, A. Tzovara, and F. D. Faraci, "Deep generative models: The winning key for large and easily accessible ECG datasets?" *Comput. Biol. Med.*, vol. 167, Dec. 2023, Art. no. 107655.
- [5] N. Neifar, A. Mdhaffar, A. Ben-Hamadou, and M. Jmaiel, "Deep generative models for physiological signals: A systematic literature review," 2023, *arXiv:2307.06162*.
- [6] S. Gerke, T. Minssen, and G. Cohen, "Ethical and legal challenges of artificial intelligence-driven healthcare," *Artif. Intell. Healthcare*, pp. 295–336, Jan. 2020.
- [7] *Regulation of the European Parliament and of the Council*, document EU 2016/679, E. Parliament of the Council of the European Union, 2016.
- [8] T. Golany, K. Radinsky, and D. Freedman, "SimGANs: Simulator-based generative adversarial networks for ECG synthesis to improve deep ECG classification," in *Proc. 37th Int. Conf. Mach. Learn.*, 2020, pp. 3597–3606.
- [9] T. Golany, D. Freedman, and K. Radinsky, "ECG ODE-GAN: Learning ordinary differential equations of ECG dynamics via generative adversarial learning," in *Proc. AAAI Conf. Artif. Intell.*, vol. 35, 2021, pp. 134–141.
- [10] Z. Zhang, X. Pan, S. Jiang, and P. Zhao, "High-quality face image generation based on generative adversarial networks," *J. Vis. Commun. Image Represent.*, vol. 71, Aug. 2020, Art. no. 102719.
- [11] Y. Skandarani, P.-M. Jodoin, and A. Lalonde, "GANs for medical image synthesis: An empirical study," *J. Imag.*, vol. 9, no. 3, p. 69, Mar. 2023.
- [12] H.-C. Shin, N. A. Tenenholtz, J. K. Rogers, C. G. Schwarz, M. L. Senjem, J. L. Gunter, K. P. Andriole, and M. Michalski, "Medical image synthesis for data augmentation and anonymization using generative adversarial networks," in *Simulation and Synthesis in Medical Imaging*. Cham, Switzerland: Springer, 2018, pp. 1–11.
- [13] S. Tulyakov, M.-Y. Liu, X. Yang, and J. Kautz, "MoCoGAN: Decomposing motion and content for video generation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 1526–1535.
- [14] J. Yoon, D. Jarrett, and M. van der Schaar, "Time-series generative adversarial networks," in *Proc. 33rd Int. Conf. Neural Inf. Process. Syst. (NIPS)*, Dec. 2019, pp. 5508–5518.
- [15] K. E. Smith and A. O. Smith, "Conditional GAN for timeseries generation," 2020, *arXiv:2006.16477*.
- [16] K. E. Smith and A. O. Smith, "A spectral enabled GAN for time series data generation," 2021, *arXiv:2103.01904*.
- [17] A. Marie Delaney, E. Brophy, and T. E. Ward, "Synthesis of realistic ECG using generative adversarial networks," 2019, *arXiv:1909.09150*.
- [18] F. Zhu, F. Ye, Y. Fu, Q. Liu, and B. Shen, "Electrocardiogram generation with a bidirectional LSTM-CNN generative adversarial network," *Sci. Rep.*, vol. 9, no. 1, pp. 1–11, May 2019.
- [19] E. Brophy, "Synthesis of dependent multichannel ECG using generative adversarial networks," in *Proc. 29th ACM Int. Conf. Inf. Knowl. Manage.*, Oct. 2020, pp. 3229–3232.
- [20] T. Golany and K. Radinsky, "PGANs: Personalized generative adversarial networks for ECG synthesis to improve patient-specific deep ECG classification," in *Proc. AAAI Conf. Artif. Intell. (AAAI)*, Jul. 2019, pp. 557–564.
- [21] N. Neifar, A. Mdhaffar, A. Ben-Hamadou, M. Jmaiel, and B. Freisleben, "Disentangling temporal and amplitude variations in ECG synthesis using anchored GANs," in *Proc. 37th ACM/SIGAPP Symp. Appl. Comput.*, Apr. 2022, p. 645.
- [22] H. Wang, Z. Ge, and Z. Wang, "Accurate ECG data generation with a simple generative adversarial network," *J. Phys., Conf. Ser.*, vol. 1631, no. 1, Sep. 2020, Art. no. 012073.
- [23] K. Antczak, "A generative adversarial approach to ECG synthesis and denoising," 2020, *arXiv:2009.02700*.
- [24] D. Nankani and R. D. Baruah, "Investigating deep convolution conditional GANs for electrocardiogram generation," in *Proc. Int. Joint Conf. Neural New. (IJCNN)*, Jul. 2020, pp. 1–8.
- [25] D. Hazra and Y.-C. Byun, "SynSigGAN: Generative adversarial networks for synthetic biomedical signal generation," *Biology*, vol. 9, no. 12, p. 441, Dec. 2020.
- [26] E. Brophy, M. De Vos, G. Boylan, and T. Ward, "Multivariate generative adversarial networks and their loss functions for synthesis of multichannel ECGs," *IEEE Access*, vol. 9, pp. 158936–158945, 2021.
- [27] T. Golany, G. Lavee, S. T. Yarden, and K. Radinsky, "Improving ECG classification using generative adversarial networks," in *Proc. AAAI Conf. Artif. Intell.*, vol. 34, no. 8, 2020, pp. 13280–13285.
- [28] A. M. Shaker, M. Tantawi, H. A. Shedeed, and M. F. Tolba, "Generalization of convolutional neural networks for ECG classification using generative adversarial networks," *IEEE Access*, vol. 8, pp. 35592–35605, 2020.
- [29] H. Yang, J. Liu, L. Zhang, Y. Li, and H. Zhang, "ProEGAN-MS: A progressive growing generative adversarial networks for electrocardiogram generation," *IEEE Access*, vol. 9, pp. 52089–52100, 2021.
- [30] Y. Xia, Y. Xu, P. Chen, J. Zhang, and Y. Zhang, "Generative adversarial network with transformer generator for boosting ECG classification," *Biomed. Signal Process. Control*, vol. 80, Feb. 2023, Art. no. 104276.
- [31] J. Qin, F. Gao, Z. Wang, D. C. Wong, Z. Zhao, S. D. Relton, and H. Fang, "A novel temporal generative adversarial network for electrocardiography anomaly detection," *Artif. Intell. Med.*, vol. 136, Feb. 2023, Art. no. 102489.
- [32] D. Makowski, T. Pham, Z. J. Lau, J. C. Brammer, F. Lespinaisse, H. Pham, C. Schölzel, and S. H. A. Chen, "NeuroKit2: A Python toolbox for neurophysiological signal processing," *Behav. Res. Methods*, vol. 53, no. 4, pp. 1689–1696, Aug. 2021.
- [33] T. Golany. (2020). *ODE_ECG*. Accessed: Sep. 12, 2023. [Online]. Available: https://github.com/tomerGolany/ode_ecg/commit/e1e678c7fe09e6ba35e33b0d3b0d71bcd8eeb06d
- [34] R. Tavenard, J. Faouzi, G. Vandewiele, F. Divo, G. Androz, C. Holtz, M. Payne, R. Yurchak, M. Rußwurm, K. Kolar, and E. Woods, "Tslearn, A machine learning toolkit for time series data," *J. Mach. Learn. Res.*, vol. 21, no. 1, pp. 4686–4691, 2020.
- [35] R. A. S. Weber, M. Eyal, N. Skafta, O. Shriki, and O. Freifeld, "Diffeomorphic temporal alignment nets," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 6570–6581.
- [36] Y. Chen, E. Keogh, B. Hu, N. Begum, A. Bagnall, A. Mueen, and G. Batista, "The UCR time series classification archive," Jul. 2015. [Online]. Available: www.cs.ucr.edu/~eamonn/time_series_data/
- [37] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville, "Improved training of Wasserstein GANs," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, Long Beach, CA, USA, 2017, pp. 5769–5779.
- [38] G. B. Moody, R. G. Mark, and A. L. Goldberger, "PhysioNet: A web-based resource for the study of physiologic signals," *IEEE Eng. Med. Biol. Mag.*, vol. 20, no. 3, pp. 70–75, May 2001.
- [39] A. Ullah, S. M. Anwar, M. Bilal, and R. M. Mehmood, "Classification of arrhythmia by using deep learning with 2-D ECG spectral image representation," *Remote Sens.*, vol. 12, no. 10, p. 1685, 2020.
- [40] M. Kang, X.-F. Wang, J. Xiao, H. Tian, and T.-L. Ren, "Beat-level interpretation of intra-patient paradigm based on object detection," *Frontiers Cardiovascular Med.*, vol. 9, Mar. 2022, Art. no. 857019.
- [41] L. Xiaolin, B. Cardiff, and D. John, "A 1D convolutional neural network for heartbeat classification from single lead ECG," in *Proc. 27th IEEE Int. Conf. Electron., Circuits Syst. (ICECS)*, Nov. 2020, pp. 1–2.
- [42] M. Kachuee, S. Fazeli, and M. Sarrafzadeh, "ECG heartbeat classification: A deep transferable representation," in *Proc. IEEE Int. Conf. Healthcare Informat. (ICHI)*, Jun. 2018, pp. 443–444.
- [43] U. R. Acharya, S. L. Oh, Y. Hagiwara, J. H. Tan, M. Adam, A. Gertych, and R. S. Tan, "A deep convolutional neural network model to classify heartbeats," *Comput. Biol. Med.*, vol. 89, pp. 389–396, Oct. 2017.
- [44] G. Kumar, U. Pawar, and R. O'Reilly, "Arrhythmia detection in ECG signals using a multilayer perceptron network," in *Proc. 27th Irish Conf. Artif. Intell. Cogn. Sci.*, 2019, pp. 353–364.
- [45] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, and S. Ozair, "Generative adversarial nets," in *Proc. 27th Int. Conf. Neural Inf. Process. Syst.* Cambridge, MA, USA: MIT Press, 2014, pp. 2672–2680.
- [46] E. Brophy, Z. Wang, Q. She, and T. Ward, "Generative adversarial networks in time series: A systematic literature review," *ACM Comput. Surv.*, vol. 55, no. 10, pp. 1–31, 2023.
- [47] M. R. Pavan Kumar and P. Jayagopal, "Generative adversarial networks: A survey on applications and challenges," *Int. J. Multimedia Inf. Retr.*, vol. 10, no. 1, pp. 1–24, Mar. 2021.



NOUR NEIFAR received the Engineering degree from the National Engineering School of Sfax (ENIS), University of Sfax, Tunisia, in 2017, where she is currently pursuing the Ph.D. degree. Her main research interests include deep learning and e-health.



ACHRAF BEN-HAMADOU received the M.Sc. degree in computer science from the Institut National des Sciences Appliquées de Rouen, in 2007, and the Ph.D. degree from the Institut National Polytechnique de Lorraine, in 2011. In 2012, he was a Postdoctoral Researcher with École des Ponts ParisTech. From 2013 and 2017, he was a Computer Vision Researcher with the Driving Assistance Research Center, Valeo, working on autonomous driving and artificial

intelligence projects. He is currently a Senior Researcher in computer science with the Centre de Recherche en Numérique de Sfax (Digital Research Center of Sfax). His current research interest includes machine learning methods applied to visual and numerical data.



MOHAMED JMIAEL received the Diploma degree in computer science engineering from Kiel University, Germany, in 1992, and the Ph.D. degree from the Technical University of Berlin, in 1996. Since 2009, he has been a Full Professor in computer science with the National School of Engineers of Sfax (ENIS), University of Sfax, Tunisia. He conducted many research projects and published more than 300 regular and invited papers in international conferences and journals.

His current research interests include software engineering of distributed systems, formal methods in model-driven architecture, self-adaptive and pervasive systems, and autonomic middleware.



AFEF MDHAFFAR received the M.Sc. degree in computer science from the National Engineering School of Sfax, University of Sfax, Tunisia, in 2009, the Ph.D. degree in computer science from the University of Marburg, Germany, in 2014, and the Habilitation degree in computer science from the University of Sfax, in 2023. She is currently an Assistant Professor with the University of Sfax and an Associate Researcher with the Centre de Recherche en Numérique de Sfax (Digital Research

Center of Sfax), Tunisia. She is involved in self-*architectures applied to cloud computing environments, data analytics methods, e-health, and precision agriculture.



BERND FREISLEBEN received the M.Sc. degree in computer science from The Pennsylvania State University, USA, in 1981, and the Ph.D. and Habilitation degrees in computer science from the Technical University of Darmstadt, Germany, in 1985 and 1993, respectively. He is currently a Full Professor of computer science with the Department of Mathematics and Computer Science, Philipps-Universität Marburg, Marburg, Germany.

His current research interests include data analytics, machine learning, and biomedical applications.

...