

DiffECG: A Versatile Probabilistic Diffusion Model for ECG Signals Synthesis

Nour Neifar

ReDCAD Lab, ENIS, University of Sfax, Tunisia
nour.neifar@redcad.org

Achraf Ben-Hamadou

SMARTS Lab, Digital Research Center of Sfax, Tunisia
achraf.benhamadou@crns.rnrt.tn

Afef Mdhaffar

ReDCAD Lab, ENIS, University of Sfax, Tunisia
afef.mdhaffar@enis.tn

Mohamed Jmaiel

ReDCAD Lab, ENIS, University of Sfax, Tunisia
mohamed.jmaiel@redcad.org

Abstract—Within cardiovascular diseases detection using deep learning applied to ECG signals, the complexities of handling physiological signals have sparked growing interest in leveraging deep generative models for effective data augmentation. In this paper, we introduce a novel versatile approach based on denoising diffusion probabilistic models for ECG synthesis, addressing three scenarios: (i) heartbeat generation, (ii) partial signal imputation, and (iii) full heartbeat forecasting. Our approach presents the first generalized conditional approach for ECG synthesis, and our experimental results demonstrate its effectiveness for various ECG-related tasks. Moreover, we show that our approach outperforms other state-of-the-art ECG generative models and can enhance the performance of state-of-the-art classifiers.

Index Terms—Deep generative models, Diffusion model, ECG synthesis, Data augmentation, ECG forecasting

I. INTRODUCTION

Cardiovascular diseases (CVDs) are the leading global cause of death, emphasizing the need for heart health diagnostic tools [1]. Electrocardiograms (ECG) represent the most significant non-invasive method for identifying cardiovascular problems [2]. ECG recordings capture the heart’s electrical activity, with each heartbeat characterized by distinct waves—the P wave, QRS complex, and T wave [2]. These waves, with their unique shapes representing specific electrical activities in the heart, provide valuable insights into the heart’s rhythm, playing a crucial role in the detection of various cardiac problems. However, ECG signals present several challenges. The recording process is particularly challenging due to the imposed regulations for personal data protection and sharing [3]. Additionally, collecting ECG data is complex due to economic constraints and time consumption [3]. The unpredictable nature of sudden cardiac issues further complicates ECG recording, resulting in imbalanced datasets. Moreover, technical issues, including equipment failures or data transmission problems, introduce additional hurdles in ECG recording, resulting in missing data. These challenges collectively impose limitations on the effectiveness of deep learning techniques proposed for preventing CVDs. Data synthesis, imputation and forecasting using deep generative models are well-known and effective solutions for addressing these challenges. However, the synthesis of

ECG signals presents a challenging task due to the complex dynamics of ECG signals [4]–[6], which significantly vary across individual conditions and among different individuals. These complexities make it challenging to generate realistic ECG signals.

Recently, diffusion models have emerged as a highly effective class of deep generative models for these tasks [7], [8]. These models offer several advantages over Generative Adversarial Networks (GANs) [9], such as training stability and the ability to generate diverse synthetic samples. Diffusion models have been shown to be effective in a wide range of applications, including time series modeling [10].

In this context, we propose a versatile framework based on Diffusion Denoising Probabilistic Models (DDPMs) for one ECG signal generation, imputation (*i.e.*, completion), and forecasting. In contrast to other related work, our proposed method is designed to be versatile and generalized, allowing seamless adaptation across various tasks. In addition, our approach introduces a simple yet efficient conditioning encoding, allowing for an explicit transition between different synthesis tasks. Moreover, by using the spectrogram representation of ECG signals for conditioning the reverse diffusion, our method leverages insights into the frequency components of the signal. This differs from standard diffusion models for 1D signals, as it incorporates information about the frequency patterns present in ECG signals. Our contributions are as follows:

- We introduce the first versatile DDPM model for ECG signal generation, imputation, and forecasting, incorporating an efficient conditioning encoding for flexible task transitions.
- We effectively condition the reverse diffusion based on the spectrogram representation of ECG signals to guide the ECG signal synthesis for all three tasks.
- We provide an extensive evaluation on the MIT-BIH arrhythmia database including a comparison with the state of the art for the three different tasks, demonstrating the effectiveness of our approach.

The remainder of this paper is organized as follows: section II provides an overview of related work, section III details

our proposed approach, section IV presents the obtained experimental results, and finally, section V summarizes our contributions and outlines some future research directions.

II. RELATED WORK

This section provides an overview of current research on deep generative models applied to ECG signals, as well as an introduction to the basic principles of diffusion models.

A. Deep generative models applied to ECG

Several previous studies investigated the use of deep learning techniques for time series generation and imputation, with deep generative models being a popular choice [11], [12]. GANs have been widely employed for the related ECG tasks [4]–[6], [13]. For instance, the authors of [5], [6] proposed leveraging shape prior knowledge on ECG into the generation process by using a set of anchors and 2D statistical modeling. For imputation task, GAN-based approaches were proposed. In [13], the authors employed a GAN framework based on a modified Gate Recurrent Unit (GRU) in the generated and discriminator networks to learn the original distribution as well as to capture the characteristics of incomplete time series. Similarly, the authors in [11] adopted a GRU in the denoising auto-encoder generator of the proposed GAN with the goal of an end-to-end imputation. In [12], the GP-VAE framework was introduced for time series imputation, where a Gaussian process (GP) prior is used in the latent space to transform the data into a smoother and more comprehensive representation. One common challenge in GANs is mode collapse, leading to limited generation diversity. VAE models, on the other hand, struggle with learning latent variables for imputation and sampling. As stated in [14], interpreting the imputation process is challenging, since these variables could not accurately capture the specific characteristics of time series data.

Diffusion models have emerged as a successful alternative to GANs, offering improved training stability and superior generation quality, as demonstrated in [10], [15]–[17]. They have found effective applications in various areas, including time series generation and imputation. Adib *et al.* proposed an unconditional generation method for one-channel ECG signals [16]. They transformed the ECG data from 1D to 2D by employing Cartesian-to-polar coordinate mapping and using techniques such as Gramian Angular Summation Fields, Gramian Angular Difference Fields, and Markov Transition Fields to create three distinct 2D embedding matrices, serving as input for a diffusion model. Unlike spectrograms, these transformations lack insight into ECG signal frequency components. Spectrograms, on the other hand, offer a comprehensive view of frequency components over time, enabling the identification of critical frequency patterns and abnormalities for accurate analysis. Furthermore, statistical metrics showed the DDPM model did not outperform a GAN-based method [10]. Alcaraz *et al.* [10] introduced the SSSD-ECG framework, a short 12-lead ECG generation approach based on DDPM. They employed structured state-space models (S4) as the primary component to capture long-term dependencies in

time series data, in contrast to transformer layers or dilated convolutions. However, the experimental results indicated a restricted impact of synthetic ECG data on enhancing cardiac anomaly classification. The authors also introduced a method for ECG signal imputation and forecasting [15], which is based on a conditional diffusion model combining DiffWave [18] and S4 models. The two distinct conditioning strategies in [10], [15] yield a non-generalized synthesis approach, with model performance highly reliant on input type and conditional data design. Zama *et al.* [17] introduced DSAT-ECG, a novel framework for generating 12-lead ECG signals that combines a diffusion model with a State Space Augmented Transformer (SPADE). The DSAT architecture draws inspiration from SSSD-ECG, replacing S4 layers with SPADE layers. Similar to [10], the authors observed limited improvement in ECG signal classification tasks when using DSAT-generated synthetic ECG data.

In this paper, unlike the discussed approaches, we propose a versatile and generalized DDPM-based approach designed for easy adaptation across different tasks including ECG generation, imputation, and forecasting. In particular, we present a simple encoding of the proposed conditioning, enabling flexible and explicit transitions between distinct tasks, making it a valuable tool for scenarios where different synthesis objectives are required.

B. Principle of diffusion models

Diffusion models involve two Markovian processes: forward and reverse diffusion. In the forward process, Gaussian noise is incrementally added to the input data \mathbf{x}_0 over T steps, converting it to a standard Gaussian distribution $q(\mathbf{x}_T) \sim \mathcal{N}(\mathbf{x}_T; 0, I)$ using a variance schedule $\beta \in [\beta_1, \beta_T]$. In the reverse diffusion process, a neural network parameterized by θ is trained to remove the noise. The forward process is defined as :

$$q(\mathbf{x}_1, \dots, \mathbf{x}_t, \dots, \mathbf{x}_T | \mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t | \mathbf{x}_{t-1}) \quad (1)$$

where $q(\mathbf{x}_t | \mathbf{x}_{t-1}) := \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t I)$. The closed-form expression for sampling \mathbf{x}_t is $\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{(1 - \bar{\alpha}_t)} \epsilon$, where $\epsilon \sim \mathcal{N}(0, I)$, $\alpha_t = 1 - \beta_t$, and $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$. The reverse diffusion process learns to recursively denoise \mathbf{x}_t to retrieve \mathbf{x}_0 . It starts with pure Gaussian noise sampled from $p(\mathbf{x}_T) = \mathcal{N}(\mathbf{x}_T, 0, I)$, and the reverse diffusion process is described by a Markov chain as follows:

$$p_\theta(\mathbf{x}_{0:T}) = p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t) \quad (2)$$

Ho *et al.* [7] showed that the reverse process can be trained with the following objective:

$$L = \min_{\theta} E_{\substack{\mathbf{x}_0 \sim D \\ \epsilon \sim \mathcal{N}(0, I) \\ t \sim U(0, T)}} \|\epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{(1 - \bar{\alpha}_t)} \epsilon, t)\|_2^2 \quad (3)$$

where t follows the discrete uniform distribution, D is the original data distribution, and the denoising function ϵ_θ estimates the noise ϵ added to get the noisy input \mathbf{x}_t .

III. PROPOSED APPROACH

The general principle of our approach is depicted in Fig. 1, which involves three main blocks: (i) ECG signal transformation, (ii) diffusion model training, and (iii) conditioning. Let us first denote x_0^h as the input ECG signal, where h is its heartbeat index. We use a spectrogram transformation \mathcal{S} and its inverse \mathcal{S}^{-1} as bidirectional transformations between 1D and 2D representations of ECG signals. Before going on the diffusion process, the input ECG signal x_0^h is transformed via $\mathcal{S}(x_0^h)$ to obtain S_0^h . The diffusion process is then applied recursively to S_0^h by gradually adding noise to it through T iterations, resulting in the noisy spectrogram S_T^h . In our approach, we adapted the U-Net architecture [19] as the reverse diffusion model, requiring conditioning not only to specify the tasks to accomplish but also to encode the signal or part of it to be considered, depending on the selected task. Also, conditioning is crucial for task switching while maintaining a unified reverse diffusion model for all tasks; generation, imputation, or forecasting. To simplify notation, we define $\Phi(h, s)$ as the function that selects the appropriate heartbeat signal depending on x_0^h , and the chosen task s . $\Phi(h, s) = x_0^h$ if s is either generation or imputation. However, if s is forecasting, $\Phi(h, s) = x_0^{h-1}$ as forecasting requires the previous heartbeat. Additionally, we define the binary mask m_s , a vector that explicitly indicates the values from $\Phi(h, s)$ to retain in the input ECG based on the selected task s . For the generation task, $m_s(i) = 0, \forall i \in [0, l]$ where l is the length of the signal to generate. For signal forecasting, $m_s(i) = 1, \forall i \in [0, l]$. For the imputation task, we set $m_s(i) = 0, \forall i \in [randstart, randend]$ where $randstart$ and $randend$ represent the start and the end of the interval where the signal values are missing. During the reverse diffusion process, S_T^h is first concatenated to the class embedding label e_l of x_0^h , the time step embedding e_T , in addition to the task-aware static embedding e_s . Moreover, we introduce two additional conditions C_1 and C_2 . C_1 represents the spectrogram of the mask m_s corresponding to the chosen task s . C_2 is the spectrogram of $(x_0^h \odot m_s)$ if s is the generation or imputation tasks, and the spectrogram of $(x_0^{h-1} \odot m_s)$ if s is the forecasting task. The resulting concatenated tensor is then fed to our diffusion model to remove the noise added in step T , and generate a less noisy spectrogram S_{T-1}^h . The model consists of seven blocks, in addition to max-pooling and deconvolution layers. Each block is comprised of three CNN sub-blocks, where each sub-block includes two convolution layers and uses the Silu activation function. This iterative process is repeated over the T steps to obtain the new spectrogram \tilde{S}_0^h .

This iterative reverse diffusion process can be formulated as follows:

$$p_\theta(S_{0:T}^h | l, s, C_1, C_2) = p(S_T^h) \prod_{t=1}^T p_\theta(S_{t-1}^h | S_t^h, l, s, C_1, C_2) \quad (4)$$

Finally, we apply the \mathcal{S}^{-1} to \tilde{S}_0^h to produce the synthetic ECG signal \tilde{x}_0^h . To enhance the convergence of the reverse

diffusion in imputation and forecasting tasks, we consider an MSE regression loss computed between generated \tilde{x}_0^h and ground truth x_0^h signals. This is in addition to the reverse diffusion loss previously detailed in equation 3.

IV. EXPERIMENTAL EVALUATION

This section presents the used dataset, details our model training settings, presents the conducted experiments, and discusses the obtained results.

A. Dataset

To train our model, we use the MIT-BIH arrhythmia database [20], a widely recognized standard dataset for arrhythmia detection and classification [21]. This dataset contains 48 half-hour ECG recordings of different patients. Each recording consists of two annotated 30-minute ECG leads, digitally recorded at 360 samples per second. This dataset includes more than 100,000 ECG heartbeats, with the majority classified as normal ECG. Three classes of heartbeats are taken into consideration in this study: the normal beats, the premature ventricular contraction beats, and fusion beats (classes N, V, and F respectively). We focused our experiments on using MLII lead since it was the most commonly employed lead to record a single ECG channel for all patients in the database and the most frequently used configuration in the previous research works.

B. Training settings

We implemented our model using the PyTorch library and trained it on an Ubuntu server equipped with a GeForce GTX 1080 Ti GPU with 11 GB of memory. The ADAM algorithm with a learning rate of 0.001 was used for stochastic gradient optimization. The number of steps in the diffusion process was set to 1000, while the minimum of schedule noise is $\beta_0 = 0.0001$ and the maximum is $\beta_T = 0.02$. We used the TorchAudio library for transferring modalities between ECG signals and their spectrograms, in both forward and reverse directions¹. An ECG signal is divided into heartbeats, also known as cardiac cycles, with each heartbeat consisting of 270 voltage values. A cardiac cycle is therefore a vector with a length of 270 values ($l = 270$), corresponding to 350 and 400 milliseconds before and after the R-peak. We randomly chose 70% of the data for the training steps, while the remaining 30% of the data was used for model testing. A signal-wise paradigm is adopted when splitting the database.

C. Experiments and results

Two steps of evaluation were considered to evaluate our approach: quantitative and qualitative evaluations. The quantitative evaluation involves assessing the impact of augmenting the real training dataset with our synthetic ECG on the performance of state-of-the-art arrhythmia classification baselines. This assessment includes a comparison of the performance of our generation method with other competing generation approaches. Additionally, the evaluation involves the use of

¹<https://pytorch.org/audio/stable/transforms.html>

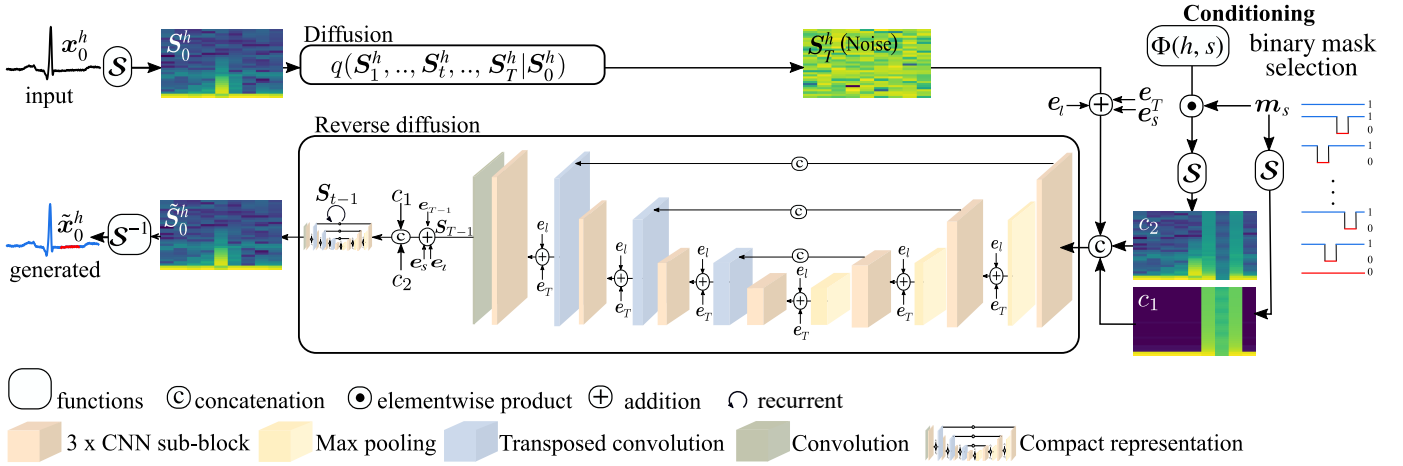


Fig. 1: The general principle of the proposed approach.

various standard metrics to quantitatively assess our approach in the three considered scenarios: generation, imputation, and forecasting.

As a qualitative evaluation, we visually checked the generated ECG signals for the different scenarios to identify any inconsistencies and visual incoherence and compare them to other generated signals obtained by competing approaches.

1) Quantitative evaluation:

a) *Classification evaluation:* In this evaluation, we used three state-of-the-art arrhythmia classification baselines [22]–[24] to assess the impact of adding synthetic data obtained by different generation methods. As a first generation baseline, we adapted the seminal GAN work [9] to our context. We considered also a state-of-the-art standard GAN approach [25] and two other recent GAN-based ECG generation methods [5], [6] that successfully incorporate shape prior to guide the generation process leading to prominent results.

The arrhythmia classification baselines [22]–[24] were trained following these five settings:

- Setting 1: only the real training dataset is used.
- Setting 2: + synthetic ECG signals generated by [9].
- Setting 3: + synthetic ECG signals generated by [25].
- Setting 4: + synthetic ECG signals generated by [5].
- Setting 5: + synthetic ECG signals generated by [6].
- Setting 6: + synthetic ECG signals generated using our approach.

Table I show the obtained classification results for the different settings. We can observe that adding synthetic ECG in the training phase systematically improves the arrhythmia classification performances of the three baselines. Moreover, our generation method outperforms the standard GANs [9], [25] and [5] and achieves comparable results to the most advanced GAN-based approach [6]. It is worth noticing that these approaches are used only for data generation. However, our method is a generalized approach adapted for generation, imputation, and forecasting without any fine-tuning and reconfiguration.

b) *Evaluation metrics:* Following previous studies such as [26]–[28], we consider the Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), Fréchet Inception Distance (FID), Dynamic Time Warping (DTW), Earth Mover’s Distance (EMD) and Maximum Mean Discrepancy (MMD) as performance metrics.

We first present the obtained metrics for the generation task, followed by a comparison with the metrics obtained for the imputation and forecasting tasks.

Table II reports the values of the used metrics for our ECG generation approach and GAN-based approaches [5], [6], [9], [25] on a set of real ECG from the test set and synthetic ECG. For all metrics, lower scores mean good results. Overall, our method produces results that are competitive with the four other competing generation methods. Indeed, for all heartbeat classes and metrics, our generation method outperforms the standard GANs generation baseline. In addition, our approach outperforms the advanced GAN [5] for the three classes on all metrics. On the other hand, we achieved better results on some metrics and obtained comparable results on other metrics with the most advanced GAN [6], across the 3 classes. This could be explained by the fact of incorporating advanced prior knowledge about ECG complex dynamics in the generation process in [6]. For instance, for class N, we obtained (22.10, 1.54e-2) for (DTW, EMD); while [6] obtained (24.42, 1.84e-2). For example, we obtained for RMSE (1.68e-3, 3.32e-3, 4.97e-3) for classes (N, V, and F) which is comparable to (1.66e-3, 3.02e-3, 4.84e-3) obtained by [6].

For the evaluation of our approach in the imputation and forecasting tasks, we compared our method to Long short-term memory (LSTM) and Variational auto-encoder (VAE) baselines in addition to a state-of-the-art method [29] based on bottleneck dilated convolutional (BDC) self-Attention architecture. Tables III and IV show the obtained results of the used metrics for the imputation and forecasting tasks, respectively. We can observe that our method achieves better results globally for all metrics in both imputation and forecasting tasks. Indeed, our approach outperformed the three selected

TABLE I: Classification results for the three baselines for the different settings. Acc., Pr., Re., F1.sc. stand for Accuracy, Precision, Recall, and F1 score.

	[22]				[23]				[24]			
	Acc.	Pr.	Re.	F1.sc.	Acc.	Pr.	Re.	F1.sc.	Acc.	Pr.	Re.	F1.sc.
Setting 1	0.97	0.93	0.89	0.91	0.98	0.87	0.82	0.84	0.96	0.87	0.74	0.77
Setting 2	0.98	0.94	0.91	0.92	0.98	0.93	0.91	0.92	0.97	0.87	0.79	0.82
Setting 3	0.98	0.93	0.93	0.92	0.98	0.93	0.93	0.93	0.98	0.90	0.92	0.91
Setting 4	0.98	0.95	0.93	0.94	0.98	0.96	0.94	0.95	0.99	0.96	0.95	0.95
Setting 5	0.99	0.97	0.95	0.94	0.99	0.97	0.95	0.96	0.99	0.96	0.96	0.96
Setting 6	0.99	0.95	0.93	0.94	0.99	0.96	0.95	0.95	0.99	0.95	0.96	0.96




Fig. 2: Examples of synthetic heartbeats for classes N, V, and F classes obtained from our generation approach (a), [6] (b), [5] (c), [9] (d) and [25] (e). The gray background represents the distribution of the real dataset, while the red signals depict the generated heartbeats.

baselines in the imputation and forecasting tasks across all

TABLE II: Obtained results of the used metrics across classes (N, V, and F) for generation task.

	RMSE	MAE	FID	DTW	EMD	MMD
N	[9]	1.86e-3	1.47e-3	2.14e-2	29.05	2.35e-2
	[25]	1.85e-3	1.42e-3	2.20e-2	28.91	2.31e-2
	[5]	1.83e-3	1.33e-3	2.10e-2	25.67	2.03e-2
	[6]	1.66e-3	1.25e-3	1.74e-2	24.42	1.84e-2
	Ours	1.68e-3	1.20e-3	1.29e-2	22.10	1.54e-2
V	[9]	3.87e-3	2.82e-3	5.03e-2	28.10	4.96e-2
	[25]	3.79e-3	2.78e-3	4.98e-2	27.62	4.89e-2
	[5]	3.54e-3	2.64e-3	3.18e-2	24.04	3.52e-2
	[6]	3.02e-3	2.22e-3	1.99e-2	21.20	2.52e-2
	Ours	3.32e-3	2.35e-3	2.61e-2	22.03	2.87e-2
F	[9]	5.37e-3	4.18e-3	8.22e-2	8.63	8.26e-2
	[25]	5.49e-3	4.07e-3	6.18e-2	8.35	6.71e-2
	[5]	5.16e-3	3.85e-3	4.48e-2	7.74	5.45e-2
	[6]	4.84e-3	3.47e-3	2.96e-2	6.61	4.23e-2
	Ours	4.97e-3	3.35e-3	3.97e-2	6.50	3.93e-2

TABLE III: Obtained results of the used metrics across classes (N, V, and F) for the imputation task.

		RMSE	MAE	FID	EMD	MMD
N	LSTM	6.90e-4	5.59e-4	9.88e-3	1.00e-2	0.61
	VAE	8.58e-4	5.76e-4	1.40e-2	1.35e-2	0.81
	[29]	6.51e-4	4.62e-4	5.24e-3	5.90e-3	0.32
	Ours	3.33e-4	2.21e-4	2.28e-3	2.35e-3	0.24
V	LSTM	3.62e-3	2.09e-3	4.00e-2	4.47e-2	0.68
	VAE	3.51e-3	2.06e-3	5.78e-2	5.64e-2	1.02
	[29]	3.74e-3	2.28e-3	4.04e-2	4.25e-2	0.56
	Ours	1.18e-3	7.71e-4	7.13e-3	8.45e-3	0.14
F	LSTM	4.93e-3	2.80e-3	7.23e-2	6.75e-2	0.53
	VAE	5.09e-2	2.54e-3	6.20e-2	5.11e-2	0.28
	[29]	4.86e-3	3.24e-3	3.26e-2	3.45e-2	0.12
	Ours	4.09e-3	3.53e-3	3.57e-2	3.46e-2	0.27

the metrics in classes N and V. For example, we obtained lower FID in the imputation task (2.28e-3, 7.13e-3) in classes N and V, while ((9.88e-3, 4.00e-2), (1.40e-2, 5.78e-2), (5.24e-3, 4.04e-2)) in LSTM, VAE and [29], respectively. On the other hand, we obtained superior results on RMSE and FID metrics for class F (*i.e.*, minority class) in the imputation task, while obtaining comparable results for the other metrics. Additionally, our method outperformed the other baselines on all metrics in the forecasting task except the MMD metric. For instance, we achieved 0.27 as MMD; whereas 0.14 by [29]. These results demonstrate the effectiveness of our approach in handling missing data and making accurate forecasts for ECG signals.

TABLE IV: Obtained results of the used metrics across classes (N, V, and F) for the forecasting task.

		RMSE	MAE	FID	EMD	MMD
N	LSTM	8.57e-4	5.52e-4	7.87e-3	6.96e-3	0.30
	VAE	8.87e-4	5.93e-4	1.42e-2	1.38e-2	0.81
	[29]	9.35e-4	6.42e-4	1.47e-2	1.44e-2	0.86
	Ours	3.49e-4	2.57e-4	2.22e-3	2.65e-3	0.26
V	LSTM	4.03e-3	2.44e-3	5.43e-2	5.58e-2	0.87
	VAE	4.17e-3	2.36e-3	6.29e-2	6.30e-2	0.99
	[29]	3.63e-3	2.03e-3	3.38e-2	3.75e-2	0.48
	Ours	1.07e-3	6.99e-4	6.22e-3	6.92e-3	0.11
F	LSTM	5.39e-2	4.55e-3	7.33e-2	7.21e-2	0.57
	VAE	5.29e-2	3.37e-3	7.64e-2	7.06e-2	0.53
	[29]	4.78e-3	3.18e-3	3.66e-2	3.69e-2	0.14
	Ours	3.83e-3	2.77e-3	2.74e-2	3.17e-2	0.27




Fig. 3: Examples of synthetic heartbeats for classes N, V, and F classes obtained from our generation approach (a), LSTM (b), VAE (c), and [29] (d). The gray background represents the distribution of the real dataset, while the red portions depict the completed portions of the heartbeats.

2) *Qualitative evaluation:* For the qualitative assessment, we start by visually comparing the synthesized signals in the three considered scenarios. We select random heartbeats obtained by our approach and the selected baselines and compare them with the distribution of the training dataset. Fig.2 shows examples of synthetic heartbeats from classes (N, V, and F) obtained from our generation approach, [6], [5], [9] and [25] alongside the real distribution of these classes. Generated heartbeats from our approach and the advanced GAN methods exhibit realistic morphology and closely follow the real distribution. The generated heartbeats from standard GANs do not exhibit complete ECG morphology and are more noisy than other synthetic beats. However, the beats generated by [5] are quite noisier than those obtained from our method and [6].

Fig. 3 presents examples of heartbeats with missing values

TABLE V: Recognition rate of both real and fake heartbeats as real obtained by the two cardiologists.

		Overall	Class N	Class V	Class F
Cardio. 1	RR_{real}	90%	94.1%	85.7%	83.3%
	RR_{fake}	100%	100%	100%	100%
Cardio. 2	RR_{real}	96.6%	100%	100%	83.3%
	RR_{fake}	96.6%	100%	100%	83.3%

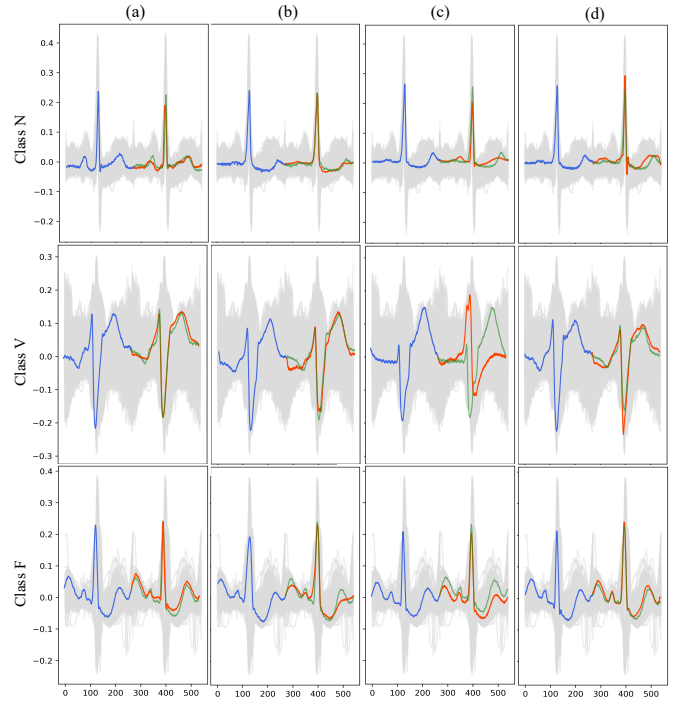


Fig. 4: Examples of synthetic heartbeats for classes N, V, and F classes obtained from our generation approach (a), LSTM (b), VAE (c), and [29] (d). The gray background represents the distribution of the real dataset, while the red and green portions depict the completed and ground-truth the heartbeats.

in the three classes. The completed part of heartbeats by the LSTM (b) VAE (c) and [29] (d) are overly smoothed compared to completed beats by our approach (a) which is not representative of real ECG signals. Our method accurately completes the missing values in these heartbeats, demonstrating a high imputation performance. Additionally, Fig. 4 displays examples of heartbeats forecasting for classes (N, V, and F) using our approach and the selected baselines. The generated beats from LSTM (b), VAE (c), and [29] (d) do not always follow the ground-truth heartbeats. For example, the generated heartbeat from class V (c) does not contain the same waves as the ground-truth beat. On the other hand, Fig. 4 (a) shows our approach's ability to accurately forecast ECG signals with realistic morphology.

The ECG signals were additionally assessed by cardiologists for a qualitative evaluation of our method. Shuffled sets containing 30 real heartbeats and 30 synthetic heartbeats from the different classes (N, V, and F) were presented to two cardiologists. The main goal is to visually differentiate between synthetic and real ECGs to assess the visual appearance of the generated ECG signals. In addition, this evaluation aimed to determine if these heartbeats could be accurately categorized into their appropriate classes, assessing the signals' class-specific properties. For this purpose, we computed the recognition rates for both real heartbeats classified as real RR_{real} and fake heartbeats classified as

real RR_{fake} . These two metrics indicate how the generated signals are realistic and how accurately they belong to their classes. Table V presents the obtained recognition rates by the two cardiologists. The cardiologists achieved high recognition rates for both real and synthetic heartbeats across all types. Notably, the first cardiologist classified all synthetic ECGs from the three classes as real ECGs. These results demonstrate the effectiveness of our proposed approach in generating ECG signals with realistic wave morphology.

V. CONCLUSION

In this paper, we presented the first versatile conditional diffusion framework for ECG synthesis that can perform three different tasks: heartbeats generation, imputation, and forecasting. The obtained results demonstrated the effectiveness of our approach, as well as its ability to enhance state-of-the-art classifiers' performance. For future work, we plan to investigate the combination of diffusion models with adversarial training to further enhance ECG synthesis. Additionally, we aim to extend our approach to generate other classes of arrhythmia and multi-lead ECG signals, while also considering the synthesis of other physiological signals.

REFERENCES

- [1] G. A. Mensah, G. A. Roth, and V. Fuster, "The Global Burden of Cardiovascular Diseases and Risk Factors: 2020 and Beyond," *Journal of the American College of Cardiology*, vol. 74, no. 20, p. 2529–2532, November 2019.
- [2] N. H. Kamaruddin, M. Murugappan, and M. I. Omar, "Early prediction of Cardiovascular Diseases using ECG signal: Review," in *2012 IEEE Student Conference on Research and Development (SCORED)*, 2012, pp. 48–53.
- [3] 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?" *Computers in biology and medicine*, vol. 167, p. 107655, 2023.
- [4] T. Golany, D. Freedman, and K. Radinsky, "ECG ODE-GAN: Learning Ordinary Differential Equations of ECG Dynamics via Generative Adversarial Learning," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, pp. 134–141, 2021.
- [5] N. Neifar, A. Mdhaffar, A. Ben-Hamadou, M. Jmaiel, and B. Freisleben, "Disentangling Temporal and Amplitude Variations in ECG Synthesis Using Anchored GANs," in *The 37th ACM/SIGAPP Symposium on Applied Computing*. New York, NY, USA: ACM, 2022, pp. 645–652.
- [6] N. Neifar, A. Ben-Hamadou, A. Mdhaffar, M. Jmaiel, and B. Freisleben, "Leveraging Statistical Shape Priors in GAN-based ECG Synthesis," *IEEE Access*, 2024.
- [7] J. Ho, A. Jain, and P. Abbeel, "Denoising Diffusion Probabilistic Models," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Virtual: Curran Associates, Inc., 2020, pp. 6840–6851.
- [8] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-Based Generative Modeling through Stochastic Differential Equations," in *International Conference on Learning Representations*, 2021.
- [9] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative Adversarial Nets," *Advances in neural information processing systems*, vol. 27, pp. 2672–2680, 2014.
- [10] J. M. L. Alcaraz and N. Strodthoff, "Diffusion-based conditional ECG generation with structured state space models," *Computers in Biology and Medicine*, p. 107115, 2023.
- [11] Y. Luo, Y. Zhang, X. Cai, and X. Yuan, "E²GAN: End-to-End Generative Adversarial Network for Multivariate Time Series Imputation," in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*. Macao, China: International Joint Conferences on Artificial Intelligence Organization, 2019, pp. 3094–3100.
- [12] V. Fortuin, D. Baranchuk, G. Rätsch, and S. Mandt, "GP-VAE: Deep Probabilistic Time Series Imputation," in *International conference on artificial intelligence and statistics*. PMLR, 2020, pp. 1651–1661.
- [13] Y. Luo, X. Cai, Y. ZHANG, J. Xu, and Y. xiaojie, "Multivariate Time Series Imputation with Generative Adversarial Networks," in *Advances in Neural Information Processing Systems*. Montréal, Canada: Curran Associates, Inc., 2018, pp. 1596–1607.
- [14] W. Du, D. Côté, and Y. Liu, "Saits: Self-attention-based imputation for time series," *Expert Systems with Applications*, vol. 219, p. 119619, 2023.
- [15] J. L. Alcaraz and N. Strodthoff, "Diffusion-based Time Series Imputation and Forecasting with Structured State Space Models," *Transactions on Machine Learning Research*, 2023.
- [16] E. Adib, A. S. Fernandez, F. Afghah, and J. J. Prevost, "Synthetic ECG Signal Generation using Probabilistic Diffusion Models," *IEEE Access*, 2023.
- [17] M. H. Zama and F. Schwenker, "ECG Synthesis via Diffusion-Based State Space Augmented Transformer," *Sensors*, vol. 23, no. 19, 2023.
- [18] Z. Kong, W. Ping, J. Huang, K. Zhao, and B. Catanzaro, "DiffWave: A Versatile Diffusion Model for Audio Synthesis," in *International Conference on Learning Representations*, 2021.
- [19] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, Eds. Munich, Germany: Springer International Publishing, 2015, pp. 234–241.
- [20] G. Moody, R. Mark, and A. Goldberger, "PhysioNet: a Web-based resource for the study of physiologic signals," *IEEE Engineering in Medicine and Biology Magazine*, pp. 707–75, 2001.
- [21] N. Neifar, A. Mdhaffar, A. Ben-Hamadou, and M. Jmaiel, "Deep Generative Models for Physiological Signals: A Systematic Literature Review," *arXiv preprint arXiv:2307.06162*, 2023.
- [22] U. R. Acharya, S. L. Oh, Y. Hagiwara, J. H. Tan, M. Adam, A. Gertych, and R. San Tan, "A deep convolutional neural network model to classify heartbeats," *Computers in biology and medicine*, vol. 89, pp. 389–396, 2017.
- [23] G. Kumar, U. Pawar, and R. O'Reilly, "Arrhythmia Detection in ECG Signals Using a Multilayer Perceptron Network," in *The 27th Irish Conference on Artificial Intelligence and Cognitive Science*. Galway, Ireland: AICS, 2019, pp. 353–364.
- [24] M. Kachuee, S. Fazeli, and M. Sarrafzadeh, "ECG Heartbeat Classification: A Deep Transferable Representation," in *2018 IEEE International Conference on Healthcare Informatics (ICHI)*. New York, NY, USA: IEEE, 2018, pp. 443–444.
- [25] A. M. Delaney, E. Brophy, and T. E. Ward, "Synthesis of realistic ECG using generative adversarial networks," *arXiv preprint arXiv:1909.09150*, 2019.
- [26] D. Hazra and Y.-C. Byun, "SynSigGAN: Generative adversarial networks for synthetic biomedical signal generation," *Biology*, vol. 9, no. 12, p. 441, 2020.
- [27] F. Zhu, F. Ye, Y. Fu, Q. Liu, and B. Shen, "Electrocardiogram generation with a bidirectional LSTM-CNN generative adversarial network," *Scientific Reports*, vol. 9, no. 6734, pp. 1–11, 2019.
- [28] Q. Xu, G. Huang, Y. Yuan, C. Guo, Y. Sun, F. Wu, and K. Weinberger, "An empirical study on evaluation metrics of generative adversarial networks," *arXiv preprint arXiv:1806.07755*, 2018.
- [29] M. Xu, A. Moreno, S. Nagesh, V. B. Aydemir, D. W. Wetter, S. Kumar, and J. M. Rehg, "PulseImpute: A Novel Benchmark Task for Pulsative Physiological Signal Imputation," in *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022.