Wearable PPG-to-Multi-Lead ECG Conversion for Cardiac Monitoring

Chongxin Zhong*, Zhishan Guo*, Anil Kishin Gehi†, Chenhan Xu*, Huining Li*
*Department of Computer Science, North Carolina State University, USA
Email: {czhong4, zguo32, cxu34, hli83}@ncsu.edu
†School of Medicine, University of North Carolina (UNC), USA
Email: {anil_gehi}@med.unc.edu

***Abstract*—The electrocardiogram (ECG) has been the gold standard for heart disease evaluation due to the rich information about the electrical activity of the heart contained in it. However, existing ECG monitoring devices either lack the capability for continuous monitoring or are unable to support multi-lead ECG recordings. To address the issues, we propose an approach for generating multi-lead ECG from photoplethysmogram (PPG), which can be passively monitored by wearable devices such as smartwatches. The PPG collected from wearable devices is first passed to a trained conditional diffusion model to generate the single-lead ECG, and then through a long short-term memory (LSTM) model to construct and predict the multi-lead ECG. The final outputs can be used to monitor and detect abnormal cardiac patterns in daily life. We evaluate the performance of our proposed approach with the dataset collected from daily-life scenarios involving 32 subjects. The results show that our approach can generate multi-lead ECGs accurately. In addition, a case study is conducted using data collected from the hospital, which demonstrates the effectiveness of our approach in detecting ST elevation.**[1]

***Index Terms*—PPG, ECG, Cardiac Monitoring, Wearable**

I. Introduction

Cardiovascular disorders are the leading cause of death in many countries with continually rising rates due to the change to modern lifestyle. As a result, the early detection of abnormal cardiac activity and timely intervention are essential for preventing the progression of cardiovascular disease and reducing associated economic and social burden. The electrocardiogram (ECG) is widely regarded as the gold standard for heart disease evaluation, owing to its ability to provide detailed insights into the heart's electrical activity. It is important to have real-time ECG monitoring, especially as part of daily out-of-clinic care, to capture transient abnormalities in cardiac patterns that might otherwise go unnoticed during routine clinical assessments.

Portable devices, like Holter monitors [1], can enable multi-lead ECG monitoring outside clinical settings. Despite their functionality, such devices are often large and require consistent skin contact via multiple electrodes, making them unsuitable for prolonged use. As an alternative, wearable devices with ECG capabilities, like the Zio Patch [2] and Apple Watch, have been introduced. These devices offer enhanced convenience but are typically limited to fixed body locations and only support single-lead ECG recordings. This limitation may compromise their reliability for comprehensive clinical evaluations. Moreover, wrist-worn devices often require users to manually maintain finger contact with the device for approximately 30 seconds to complete a single ECG recording.

In light of this situation, we study the feasibility of generating multi-lead ECG signals from the existing biosensors on off-the-shelf wearables, particularly the photoplethysmogram (PPG) sensor. This approach leverages the intrinsic physiological relationship between ECG and PPG signals: ECG reflects the heart's electrical activity, while PPG captures its mechanical response. Specifically, the QRS complex in an ECG represents ventricular depolarization, which triggers cardiac contraction and blood ejection, leading to volumetric changes that are readily detectable by PPG [3]. Recent studies have demonstrated the potential of generative models [4], such as GANs and diffusion models, to synthesize one-lead ECG from PPG measurements. Separately, other works have shown success in reconstructing a full 12-lead ECG from limited or single-lead ECG signals using signal completion or domain translation techniques [5]. However, none of these works are able to directly generate the multi-lead ECG from the PPG. One significant and often overlooked reason for this is the accumulation of modeling errors in a multi-stage pipeline, especially when starting from a non-electrical signal like PPG, which is the focus of this paper.

In this paper, we present a wearable approach for converting PPG signals into multi-lead ECG data to support continuous and wearable cardiac monitoring. At the high level, the proposed approach contains two phases: (i) generating the single-lead ECG from the PPG and (ii) predicting the multi-lead ECG with the generated single-lead ECG. Specifically, we combine the conditional diffusion model and the long short-term memory model, leveraging their abilities to generate outputs based on specific input conditions and handle time series inputs. Moreover, we incorporate the physiological constraints between ECG leads into the loss function to minimize the accumulated errors incurred when combining the two phases. The final results can be used to monitor transient abnormalities in cardiac patterns. Our approach offers continuous monitoring without requiring additional user cooperation, while providing

[1]ST elevation refers to an upward deviation of the ST segment on an electrocardiogram (ECG) from the baseline, indicating a potential heart attack or other cardiac issues. It is a crucial diagnostic finding in acute myocardial infarction (heart attack) and requires prompt medical attention. It is a key indicator of myocardial ischemia in practice.

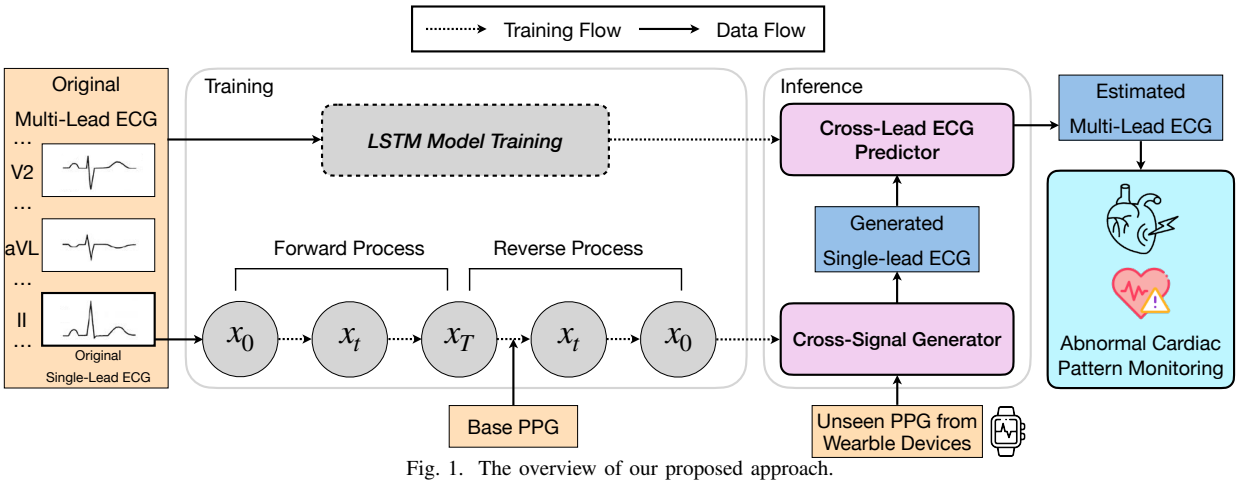

Fig. 1. The overview of our proposed approach.

rich diagnostic information through multi-lead ECG.

We evaluate our approach on a public dataset with both PPG and multi-lead ECG from daily-life scenarios involving 32 subjects. Our proposed approach demonstrates satisfactory accuracy in generating multi-lead ECGs from PPG, with an average root mean square error of 0.24 and an average Fréchet distance of 3.108.

II. Methodology

A. Data preprocessing

Due to the nature of ECG signal acquisition, various electrical signals from muscles other than the heart are also collected, which can introduce muscular artifacts into the recordings. Moreover, poor electrical contact between the ECG electrodes and the skin can further contribute to signal noise. Since the frequencies of the ECG typically range from 0.5 to 150 Hz, we apply a low-pass filter with a 150 Hz cutoff and a zero-phase filter to remove the noise. The signals are segmented into windows of 4 seconds for the following modules.

B. Cross-signal generator for single-lead ECG

The cross-signal generator takes PPG signals as inputs and generates ECG signals accordingly. Prior studies [4] have demonstrated that the diffusion model outperforms other generative models, such as GAN. In this work, we adapt the denoising diffusion probabilistic model (DDPM) [4] to generate the single-lead ECG from PPG. It progressively adds noise to the training data in a forward process and then learns to reverse this process to synthesize new data samples. As shown in Figure 1, given a segment of the original ECG signal x_0 at timestep 0, the forward process adds Gaussian noise gradually to the data at each timestep t until the T-th step as a Markov chain:

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

where

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1-\beta_t}\mathbf{x}_{t-1}, \beta_t \boldsymbol{I}). \tag{2}$$

β_t is the noise variance obtained from a pre-defined variance schedule $\beta_1, \ldots, \beta_T$ that guides the addition of noise at timestep t, and $\boldsymbol{I}$ is the identity matrix. The forward process of each step can be formulated as:

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1-\bar{\alpha}_t}\epsilon, \tag{3}$$

where $\alpha_t = 1 - \beta_t$, $\bar{\alpha}_t = \prod_{i=1}^{t} \alpha_i$, and $\epsilon \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$. At the end of the forward process, x_T is considered as an isotropic Gaussian.

The reverse process can also be defined as a Markov chain to obtain the original signal x_0 from x_t using the transition kernel $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$ and the posterior distribution $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$. According to DDPM [6], the diffusion model can be trained on the following loss function, maximizing the variational lower bound (ELBO) to learn the probability distribution of the original single-lead ECG dataset:

$$L(\theta) = \mathbb{E}_{t,\mathbf{x}_0,\epsilon}[\|\epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1-\bar{\alpha}_t}\epsilon, c, t)\|^2], \tag{4}$$

where $\epsilon = (\mathbf{x}_t - \sqrt{\bar{\alpha}_t}\mathbf{x}_0)/\sqrt{1-\bar{\alpha}_t}$ is the added noise in each step, $\epsilon_\theta(\mathbf{x}_t, t)$ is a noise-prediction network to estimate ϵ, and c is the input control signal that the diffusion model is conditioned on. In our work, c is the corresponding PPG of the original ECG. We adopt a 1D U-Net model as the backbone of the conditional diffusion model. After training, the cross-signal generator uses the trained diffusion model to generate the single-lead ECG signals with unseen PPG signals. We specifically generate Lead II in the signal generation phase and use it in subsequent predictions, given its clinical importance and clarity in visualizing ECG waveforms.

C. Error-resilient cross-lead ECG predictor

Model overview and data representation: The core of the cross-lead ECG predictor is a Long Short-Term Memory (LSTM) neural network, chosen for its strong capability in modeling temporal dynamics commonly seen in physiological signals such as ECGs. The training signals are processed with a sliding window of length m. The training data at time t is represented as $d_t = \{x_t, y_{it}\}$, $y_{it} \in \{y_{1t}, y_{2t}, \ldots, y_{nt}, \}$, where n is the number of leads to be predicted. The output of model f, denoted by $\hat{y_{it}}$, has the same dimensionality as y_{it}. The training objective minimizes the mean squared

error (MSE) across all predicted leads. The loss function is formulated as:

$$L_p = \frac{1}{n}\sum_{i=1}^{n} MSE(y_{it}, \hat{y}_{it}). \quad (5)$$

Reducing accumulated error via physiological constraints: A key limitation of prior approaches is that they typically treat the lead generation and prediction stages independently, ignoring the accumulated errors that arise when early prediction errors propagate into later stages. This accumulation can lead to clinically misleading reconstructions. To address this, our method integrates mathematical relationships among ECG leads, which reflect underlying cardiac vector projections, directly into the learning objective. For example, the well-known Einthoven's law: $Lead_{II} = Lead_I + Lead_{III}$ allows us to derive one limb lead from the other two. Similarly, augmented leads such as aVF or aVL can be expressed using limb lead combinations. By incorporating these constraints during training, we reduce accumulated error by ensuring that the model's predictions remain internally consistent across leads. That is, not only are the leads predicted to match ground truth values, but they are also constrained to satisfy known inter-lead dependencies. The relationship-based loss quantifies deviations from these expected inter-lead constraints:

$$L_m = \frac{1}{n}\sum_{i=1}^{n} MSE(y_{it}, y_{mit}), \quad (6)$$

where y_{mit} represents the value calculated using the mathematical relationships with the already generated leads. Therefore, the total loss is calculated as: $L = L_p + L_m$. This dual-objective formulation ensures that the model not only produces accurate lead estimates but also maintains fidelity to the physiological laws governing ECG signal morphology.

D. Cardiac function monitoring

Our system can facilitate cardiac function monitoring outside clinics. Specifically, wearable devices passively collect users' PPG data, process it, and forward it to the cross-signal generator. The generator produces single-lead ECG signals, which are then input to the cross-lead ECG predictor to synthesize the desired multi-lead ECG signals. These outputs can subsequently be fed into pre-trained diagnostic models [7] for monitoring and detecting abnormal cardiac patterns.

III. Evaluation

A. Experimental Setup

1) Data preparation: We evaluate the quality of generated multi-lead ECG signals using the SensSmartTech database of cardiovascular signals [8] that contains synchronized ECG and PPG collected from 32 healthy subjects. The heart rates of the subjects range from 58 bpm - 173 bpm. The dataset contains 4 leads of ECG signals (Lead I, II, V3, and V4). Lead I and Lead II are measured by limb electrodes, and Lead V3 and Lead V4 are measured by precordial electrodes. We chose the PPG collected from the brachial artery instead of one from the left carotid because it is closer to the position of wearing smartwatches. The ECG signals are collected with a sampling rate of 500 Hz, and the PPG signals are collected with a sampling rate of 100 Hz. Both ECG and PPG are resampled to 128 Hz and normalized to the range (-1, 1) in the experiments. We choose to generate Lead II of ECG from PPG and then generate the other three leads of ECG because Lead II is the most commonly monitored ECG lead. We process the signals with a 4-second window and a window shift of 0.5 seconds. In total, we have 17689 samples of synchronized multi-lead ECG signals and PPG signals. 80% of the samples are used as training data and the rest 20% are used for testing.

2) Software implementation: The experiments are performed on two NVIDIA RTX A600 GPUs. The diffusion model is trained for 1000 epochs using the Adam optimizer with a 1e-4 learning rate. For the LSTM model predicting the multi-lead ECG signals, we use the Adam optimizer with a 1e-3 learning rate. The LSTM model is trained for 200 epochs.

B. Performance Metrics

1) Data fidelity: We evaluate the performance of generating multi-lead ECG from PPG using Root Mean Square Error (RMSE) and Fréchet distance (FD) [9]. The RMSE calculates the quadratic mean of the differences between the original ECG signals and the generated ones. A small RMSE indicates the high accuracy of the generated ECG compared with the original ECG. The FD measures the similarity between the original ECG and the generated ECG, taking into account the location and ordering of the points along the signal curves. A small FD means the high similarity of the original ECG and the generated ECG.

2) Feature authenticity: We compute normalized errors of Q-T interval and R-R interval of the generated multi-lead ECG signals against the original ones to evaluate the outputs' ability for heart electrical activity computing. The Q-T interval and R-R interval are crucial for the detection of several diseases, such as ventricular proarrhythmia, diabetes, and insomnia. The Q-T interval represents the time elapsed between the start of the QRS complex and the end of the T wave. The R-R interval represents the time elapsed between R waves of two successive QRS complexes. The normalized error is computed as $\frac{|I_g - I_o|}{Heartbeat\ duration}$, where I_g and I_o represent the intervals of the generated ECG and the original ECG, respectively.

C. Overall Performance

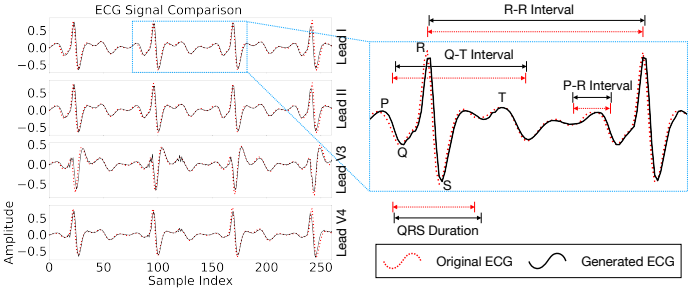

Fig. 2. The comparison between the original ECG and the generated ECG.

1) Data evaluation: Figure 2 depicts the visualization of the original and generated multi-lead ECG. The generated multi-lead ECG shows nearly identical cycle length and amplitude. This demonstrates the high fidelity of the generated multi-lead ECG. Figure 2 also zooms in two cycles of the lead I ECG to show the comparison of the fiducial points between the original ECG and the generated ECG. The generated ECG preserves consistency with the original ECG in most key measurements, including the R-R interval and the Q-T interval.

ECG Lead	RMSE	FD
Lead I	0.239	3.074
Lead II	0.260	3.017
Lead V3	0.227	3.425
Lead V4	0.234	2.916

TABLE I
THE PERFORMANCE OF GENERATED MULTI-LEAD ECG SIGNALS.

Table I shows the RMSE and FD of each lead of the generated ECG signals. The baseline performance from previous work [4] is as follows: The average RMSE is 0.22, and the average FD is 5.23. The results indicate that our method can generate high-quality multi-lead ECG signals with similar performance compared with state-of-the-art approaches.

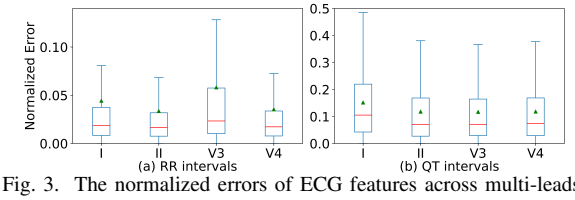

Fig. 3. The normalized errors of ECG features across multi-leads.

2) Feature evaluation: Figure 3 shows the ECG feature evaluation results. We calculate the normalized errors of the features, such as R-R interval and Q-T interval, between the original ECG and the generated ECG. The normalized errors of the R-R interval range from 0.035 to 0.058. The normalized errors of the Q-T interval range from 0.117 to 0.151. The generated multi-lead ECG signals achieve better performance on R-R interval extraction compared with that of the Q-T interval. This is primarily because existing algorithms can detect R peaks more reliably, as they correspond to the points of maximum amplitude in the ECG waveform.

D. Case Study: ST Elevation Detection

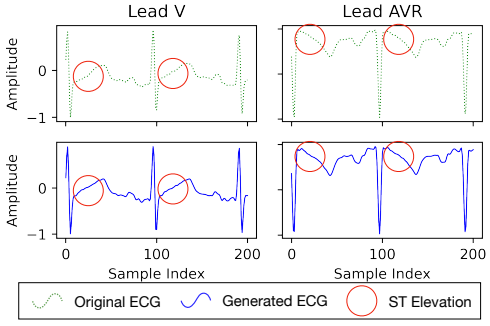

Fig. 4. The visualization of the generated ECG showing ST Elevation.

We present a case study using the MIMIC-III waveform database [10] to demonstrate that our method can generate ECG signals exhibiting clinically relevant features, such as ST elevation. ST elevation refers to an upward deviation of the ST segment and is commonly associated with myocardial ischemia. Its presence across multiple leads is a critical diagnostic indicator. As shown in Figure 4, the generated ECG signals for lead V and lead aVR—based on data collected from intensive care unit (ICU) patients—clearly display ST elevation in both leads. This result highlights the potential clinical applicability of our approach in generating diagnostically meaningful ECG waveforms.

IV. CONCLUSION

In this paper, we present an approach for generating a multi-lead ECG from PPG, combining the advantages of continuous PPG monitoring and the rich information provided by the multi-lead ECG. It consists of data preprocessing, a cross-signal generator for single-lead ECG, and an error-resilient cross-lead ECG predictor. The results show that our work can generate accurate multi-lead ECG signals, which can be used to detect real-world heart diseases. In future work, we will include a more diverse patients with various cardiac conditions in our study. We aim to evaluate the generalizability of our method across different pathological scenarios.

V. ACKNOWLEDGMENT

This work is supported under NSF CMMI 2246672.

REFERENCES

[1] J. P. DiMarco and J. T. Philbrick, "Use of ambulatory electrocardiographic (holter) monitoring," *Ann. Intern. Med.*, vol. 113, no. 1, pp. 53–68, Jul. 1990.

[2] M. A. Rosenberg, M. Samuel, A. Thosani, and P. J. Zimetbaum, "Use of a noninvasive continuous monitoring device in the management of atrial fibrillation: a pilot study," *Pacing Clin. Electrophysiol.*, vol. 36, no. 3, pp. 328–333, Mar. 2013.

[3] C. Orphanidou, "Quality assessment for the photoplethysmogram (PPG)," in *SpringerBriefs in Bioengineering*. Cham: Springer International Publishing, 2018, pp. 41–63.

[4] D. Shome, P. Sarkar, and A. Etemad, "Region-disentangled diffusion model for high-fidelity ppg-to-ecg translation," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 38, no. 13, 2024, pp. 15 009–15 019.

[5] L. Bedin, G. Cardoso, J. Duchateau, R. Dubois, and E. Moulines, "Leveraging an ecg beat diffusion model for morphological reconstruction from indirect signals," *Advances in Neural Information Processing Systems*, vol. 37, pp. 84 409–84 446, 2024.

[6] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.

[7] X. Zhang, K. Gu, S. Miao, X. Zhang, Y. Yin, C. Wan, Y. Yu, J. Hu, Z. Wang, T. Shan, S. Jing, W. Wang, Y. Ge, Y. Chen, J. Guo, and Y. Liu, "Automated detection of cardiovascular disease by electrocardiogram signal analysis: a deep learning system," *Cardiovasc. Diagn. Ther.*, vol. 10, no. 2, pp. 227–235, Apr. 2020.

[8] A. Lazović, P. Tadić, N. Đorđević, V. Atanasoski, M. Tiosavljevic, M. Ivanovic, L. Hadzievski, A. Ristic, V. Vukcevic, and J. Petrovic, "Senssmarttech database of cardiovascular signals synchronously recorded by an electrocardiograph, phonocardiograph, photoplethysmograph and accelerometer," 2024. [Online]. Available: https://physionet.org/content/senssmarttech/1.0.0/

[9] Efrat, Guibas, S. Har-Peled, and Murali, "New similarity measures between polylines with applications to morphing and polygon sweeping," *Discrete & Computational Geometry*, vol. 28, pp. 535–569, 2002.

[10] A. E. W. Johnson, T. J. Pollard, L. Shen, L.-W. H. Lehman, M. Feng, M. Ghassemi, B. Moody, P. Szolovits, L. A. Celi, and R. G. Mark, "MIMIC-III, a freely accessible critical care database," *Sci. Data*, vol. 3, no. 1, p. 160035, May 2016.