

Region-Disentangled Diffusion Model for High-Fidelity PPG-to-ECG Translation

Debaditya Shome , Pritam Sarkar , Ali Etemad

Queen’s University, Canada

{debaditya.shome, pritam.sarkar, ali.etemad}@queensu.ca

Abstract

The high prevalence of cardiovascular diseases (CVDs) calls for accessible and cost-effective continuous cardiac monitoring tools. Despite Electrocardiography (ECG) being the gold standard, continuous monitoring remains a challenge, leading to the exploration of Photoplethysmography (PPG), a promising but more basic alternative available in consumer wearables. This notion has recently spurred interest in translating PPG to ECG signals. In this work, we introduce Region-Disentangled Diffusion Model (RDDM), a novel diffusion model designed to capture the complex temporal dynamics of ECG. Traditional Diffusion models like Denoising Diffusion Probabilistic Models (DDPM) face challenges in capturing such nuances due to the indiscriminate noise addition process across the entire signal. Our proposed RDDM overcomes such limitations by incorporating a novel forward process that selectively adds noise to specific regions of interest (ROI) such as QRS complex in ECG signals, and a reverse process that disentangles the denoising of ROI and non-ROI regions. Quantitative experiments demonstrate that RDDM can generate high-fidelity ECG from PPG in as few as 10 diffusion steps, making it highly effective and computationally efficient. Additionally, to rigorously validate the usefulness of the generated ECG signals, we introduce CardioBench, a comprehensive evaluation benchmark for a variety of cardiac-related tasks including heart rate and blood pressure estimation, stress classification, and the detection of atrial fibrillation and diabetes. Our thorough experiments show that RDDM achieves state-of-the-art performance on CardioBench. To the best of our knowledge, RDDM is the first diffusion model for cross-modal signal-to-signal translation in the bio-signal domain.

1 Introduction

The persistent rise in cardiovascular diseases (CVDs) is a leading cause of mortality worldwide, with an estimated 17.9 million fatalities attributed to CVDs in a year (Kauffman 2020). Early detection, timely intervention, and widely-accessible health-monitoring devices are crucial factors in preventing CVDs and reducing related healthcare costs. Democratizing accurate cardiac monitoring is critical, as it empowers individuals and communities to take control of their cardiovascular health regardless of socioeconomic status or geographic location (Bayoumy et al. 2021).

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Electrocardiography (ECG) has long been the gold standard for cardiac monitoring, offering crucial insights to diagnose numerous CVDs (Hannun et al. 2019). However, their requirement for specialized equipment and skilled medical personnel limits their accessibility and affordability for daily life cardiac activity monitoring. Continuous ECG monitoring is always preferred over conventional clinical settings as it offers a comprehensive assessment of cardiac activity by collecting and analyzing data from patients uninterruptedly in their day-to-day life, and is not limited to a single point in time. For example, the detection of Atrial Fibrillation (AFib) remains challenging due to its asymptomatic and intermittent nature (Camm, Corbucci, and Padeletti 2012), which often goes undetected in clinical settings that usually take place over the short duration of a patient’s visit to the physician (Steinhubl et al. 2018).

To overcome these limitations, Photoplethysmography (PPG) has emerged as a promising non-invasive and cost-effective alternative for continuous cardiac monitoring in wrist-based wearable devices. PPG is a non-invasive optical signal that measures blood volume fluctuations to capture cardiac activity. However, PPG lacks detailed cardiac information in comparison to ECG, which limits its use in diagnosing CVDs. Due to the non-invasive nature and ability of PPGs to be integrated into wearable devices, it provides an unmatched level of ubiquity.

To leverage the convenience and continuous monitoring ability of PPG and the diagnostic utility of ECG, PPG-to-ECG translation has emerged as a promising, yet challenging alternative (Sarkar and Etemad 2021; Vo et al. 2021; Lan 2023). The goal of PPG-to-ECG translation is to use deep learning techniques to generate ECG signals from PPG, with the aim that intricate CVD-related information in PPG signals become more prominent and detectable in the output ECG.

We identify two key areas that require significant improvement in the context of PPG-to-ECG translation. **First**, recent advances in Diffusion models have shown great promise in producing hyper-realistic content with high fidelity (Saharia et al. 2022; Blattmann et al. 2023; Ruan et al. 2023). Yet, to our knowledge, they have not been successfully explored for PPG-to-ECG translation. More specifically, designing Diffusion models capable of effectively learning the nuances of both spatial and temporal dynamics of bio-signals like ECG has not been explored. **Second**, despite promising results in

PPG-to-ECG translation, to our knowledge, experiments to comprehensively evaluate the predictive power of the generated ECG in detecting CVDs in comparison to the original PPG signals has not been conducted.

In this paper, we leverage diffusion models (Sohl-Dickstein et al. 2015) and develop a framework for PPG-conditioned high-fidelity ECG generation. We observe that Diffusion models like Denoising Diffusion Probabilistic Models (DDPMs) (Ho, Jain, and Abbeel 2020) struggle to capture the nuances of bio-signal translation due to their inability to identify regions of interest (ROIs), and therefore fail to learn both the global structures and fine-grained local details. The underlying reason for this limitation is rooted in the forward process, which uniformly adds noise to the entire signal without considering the morphological differences of different bio-signals (e.g., QRS complex in ECG or systolic peak in PPG). To overcome this challenge, we introduce a novel Diffusion model called Region-Disentangled Diffusion Model (RDDM) which is able to capture the complex temporal dynamics and region-specific intricate details of ECG signals. In particular, we introduce a novel forward process that selectively adds noise to specific ROIs in ECG and a reverse process that disentangles the denoising of QRS and Non-QRS regions. Our experiments show that our method outperforms state-of-the-art methods in PPG-to-ECG translation in terms of standard error metrics (e.g., RMSE). Moreover, to tackle the second problem mentioned earlier, we introduce CardioBench, a comprehensive evaluation benchmark for PPG-to-ECG translation. CardioBench consists of a series of 5 cardiac-related tasks selected to evaluate the utility of the generated ECG signals. The tasks include (1) Heart rate (HR) estimation, (2) AFib detection, (3) Stress and Affect classification, (4) Diabetes detection, and (5) Blood Pressure (BP) estimation. Evaluating our proposed method on CardioBench demonstrates that the generated ECG signals can be effectively used for cardiac-related tasks, and can outperform the use of the original PPG for the same tasks by large margins.

Following is a summary of our contributions:

- We introduce RDDM, a novel diffusion model for high-fidelity PPG-to-ECG translation. To the best of our knowledge, this is the first diffusion model for cross-modal signal-to-signal translation in the bio-signal domain.
- Our proposed RDDM introduces a novel approach that disentangles the diffusion process into two distinct components. The first component captures the global temporal structure of ECG signals, while the second component is dedicated to capture the fine-grained local details. This enables RDDM to generate high-fidelity ECG and do so in just 10 sampling steps making RDDM highly-efficient.
- To evaluate the quality and utility of the generated ECG, we introduce CardioBench, a comprehensive evaluation benchmark comprised of 5 challenging cardiac-related tasks.
- Our extensive evaluations demonstrate that the ECG signals generated by RDDM are significantly better in quality than DDPMs and prior works in both standard quantitative measures and their performance on CardioBench tasks.
- We make the code public to the research community¹.

¹<https://github.com/DebadityaQU/RDDM>

2 Related Work

ECG synthesis with Generative models. Generative models, particularly Generative Adversarial Networks (GANs) have been used as a promising approach in ECG synthesis (Chen et al. 2022; Golany and Radinsky 2019). Recently, Diffusion models have also been used for ECG generation (Alcaraz and Strodthoff 2023; Adib et al. 2023; Neifar et al. 2023). Unconditional ECG generation using 2D Gramian Angular Field representations has been explored in (Adib et al. 2023). DDPMs have been used for ECG generation from clinical statements (Alcaraz and Strodthoff 2023). DiffECG (Neifar et al. 2023) explored the application of DDPM for class-conditional ECG synthesis.

PPG-to-ECG translation with Generative models. The research landscape of PPG-to-ECG translation remains relatively under-explored. Earlier attempts to solve this problem primarily focused on estimating various ECG parameters from PPG signals rather than direct reconstruction using methods like Feature selection (Banerjee et al. 2014), Linear regression (Zhu et al. 2021), and Dictionary learning (Tian et al. 2022). Such methods suffer a degradation in performance when evaluating in a subject-independent manner. CardioGAN (Sarkar and Etemad 2021) made the first attempt at using Generative models for direct PPG-to-ECG translation. They used unpaired training using CycleGAN (Zhu et al. 2017) with an improved architecture by incorporating attention and dual time-frequency discriminators. Empirically, they demonstrated significantly better HR estimation from the generated ECG compared to the original PPG. Similarly, (Vo et al. 2021) utilized conditional Wasserstein GANs for PPG-to-ECG translation. (Lan 2023) developed a Transformer-based framework for reconstructing ECG from PPG, and further fine-tuned the model for CVD detection using original PPG and generated ECG together.

3 Region-Disentangled Diffusion Model

Diffusion models (Sohl-Dickstein et al. 2015) consist of a forward and a reverse process. The forward diffusion process $q(\mathbf{x}_t|\mathbf{x}_{t-1})$ can be characterized as a Markov chain that progressively adds Gaussian noise at every timestep t :

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

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

Here, $\mathbf{x}_0 \sim q(\mathbf{x}_0)$ is a clean signal and β_t is a small positive constant obtained from a fixed variance schedule β_1, \dots, β_T . Let $\alpha_t := 1 - \beta_t$ and $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$. Then, the forward diffusion process yields a sample at timestep t , denoted by \mathbf{x}_t , which can be obtained in a single step as follows:

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, \quad \text{where} \quad \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (2)$$

Since the reverse step of the forward process, $q(\mathbf{x}_{t-1}|\mathbf{x}_t)$ is computationally infeasible, DDPMs (Ho, Jain, and Abbeel 2020) maximize the variational lower bound (ELBO) using a parameterized Gaussian transition $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$ with parameter θ . Consequently, the reverse process is approximated as

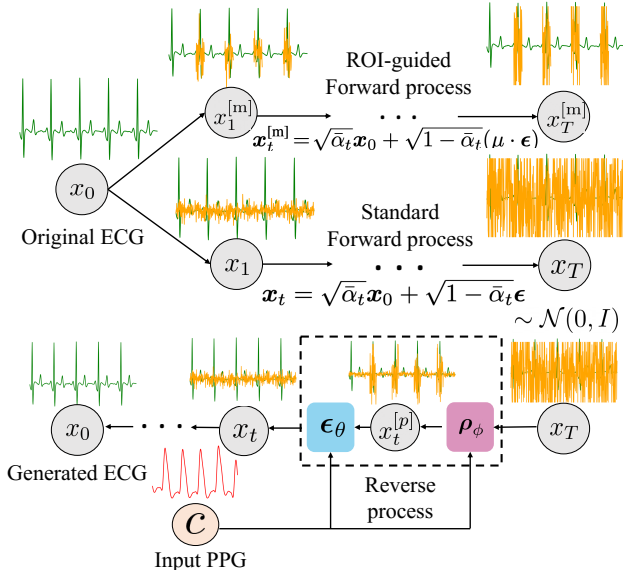


Figure 1: PPG-to-ECG translation with RDDM. Top: ROI-guided Forward process selectively adds Gaussian noise to the QRS regions, while standard forward process adds noise uniformly across the ECG. Bottom: Our reverse process involves disentangled denoising of the QRS and Non-QRS regions by ϵ_θ and ρ_ϕ respectively.

a Markov chain with a learned mean and a fixed variance, initiated from $p(x_T) = \mathcal{N}(x_T; \mathbf{0}, \mathbf{I})$:

$$p_\theta(x_{0:T}) := p_\theta(x_T) \prod_{t=1}^T p_\theta(x_{t-1}|x_t), \quad (3)$$

$$\text{where, } p_\theta(x_{t-1}|x_t) := \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \sigma_t^2 \mathbf{I}), \quad (4)$$

$$\text{and, } \mu_\theta(x_t, t) := \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t) \right), \quad (5)$$

The diffusion model ϵ_θ can be conditioned upon the input signal c , which makes the conditional objective:

$$L(\theta) := \mathbb{E}_{t, x_0, \epsilon} \left[\left| \epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, c, t) \right|^2 \right]. \quad (6)$$

PPG-to-ECG translation can be performed by adopting Equation 6 using the PPG signal as input condition c to generate the corresponding ECG signal x_0 . To enable the Diffusion model to learn both global structures and fine-grained local details of ECG for nuanced bio-signal translation, we propose RDDM which modifies the forward and reverse process of DDPMs as discussed below.

Forward Process. The forward process of RDDM consists of 2 stages, (i) we obtain a noise-added sample x_t following the standard forward process mentioned in Equation 2; (ii) in parallel, we also obtain a selectively noise-added sample $x_t^{[m]}$, in a process which we refer to as the ROI-guided forward process. Given the salience of cardiac-related information in the QRS complex of ECG signals, we define the ROI to enable the Diffusion model to learn these regions effectively

Algorithm 1: RDDM Training

- 1: **repeat**
- 2: $x_0 \sim q(x_0)$;
- 3: $t \sim \text{Uniform}(\{1, \dots, T\})$
- 4: $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
- 5: Apply ROI mask: $\epsilon^{[m]} = (\mu \cdot \epsilon)$
- 6: Compute $x_t = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon$
- 7: Estimate $x_t^{[p]} = \rho_\phi(x_t, c, t)$
- 8: Compute $x_t^{[m]} = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon^{[m]}$
- 9: Calculate gradient as
$$\nabla_{\theta, \phi} \left(\lambda_1 \|\epsilon^{[m]} - \epsilon_\theta(x_t^{[m]}, c, t)\|^2 + \lambda_2 \|x_t^{[m]} - x_t^{[p]}\|^2 \right)$$
- 10: **until** converged

alongside the global structure of the ECG which is facilitated through the standard forward process. Accordingly, we define a binary mask μ with the same length of x as:

$$\mu[i] = \begin{cases} 1, & \text{if } i_r - \frac{\Gamma}{2} \leq i \leq i_r + \frac{\Gamma}{2} \\ 0, & \text{otherwise} \end{cases} \quad (7)$$

where, i_r denotes the indices of R-peaks and Γ is a hyper-parameter defined as the window size to capture the ROI centered at i_r . Finally, the ROI-guided forward process is defined as:

$$x_t^{[m]} = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}(\mu \cdot \epsilon) \text{ where } \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (8)$$

Reverse Process. Since μ is not available during inference, we approximate the parameterized transition $\rho_\phi(x_t^{[m]}|x_t)$ by our training objective defined as:

$$L(\theta, \phi) := \mathbb{E} \left[\lambda_1 |(\mu \cdot \epsilon) - \epsilon_\theta(x_t^{[m]}, c, t)|^2 + \lambda_2 |x_t^{[m]} - \rho_\phi(x_t, c, t)|^2 \right]. \quad (9)$$

Here, λ_1 and λ_2 are the loss coefficients. In our proposed objective, the first part is responsible for learning to denoise the ROIs (QRS), while the second part learns to denoise the rest of the regions (Non-QRS), effectively disentangling the denoising of the important regions of ECG signals.

More specifically, the objective function first guides the learning of ϵ_θ , where ϵ_θ is responsible for refining the ROIs. The loss term $|(\mu \cdot \epsilon) - \epsilon_\theta(x_t^{[m]}, c, t)|^2$ encourages ϵ_θ to estimate noise that, when subtracted from the signal, refines and enhances the morphology of the ROIs. This is achieved by minimizing the difference between the noise generated by ϵ_θ and the ROI-specific noise $(\rho \cdot \epsilon)$ in the original data distribution. Therefore, **the role of ϵ_θ is to learn the fine-grained details of the ROI.**

Second, the loss term $|x_t^{[m]} - \rho_\phi(x_t, c, t)|^2$ encourages ρ_ϕ to minimize the discrepancy between $x_t^{[m]}$ and the signal recovered by $\rho_\phi(x_t, c, t)$. Consequently, the function ρ_ϕ learns to denoise the non-ROI parts of the ECG signal, which implicitly involves learning ‘where’ to place the QRS complexes. **This results in ρ_ϕ capturing the temporal dynamics of the ECG signal.** This disentanglement makes RDDM particularly suitable for PPG-to-ECG translation, which requires capturing both the spatial and temporal characteristics of

Algorithm 2: RDDM Sampling

```

1:  $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
2: for  $t = T, \dots, 1$  do
3:    $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  if  $t > 1$ , else  $\mathbf{z} = \mathbf{0}$ 
4:    $\mathbf{x}_t^{[p]} = \rho_\phi(\mathbf{x}_t, c, t)$ 
5:    $\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left( \mathbf{x}_t^{[p]} - \frac{1-\alpha_t}{\sqrt{1-\alpha_t}} \epsilon_\theta(\mathbf{x}_t^{[p]}, c, t) \right) + \sigma_t \mathbf{z}$ 
6: end for
7: return  $\mathbf{x}_0$ 

```

ECG signals. The overall training algorithm of RDDM has been outlined in Algorithm 1.

Sampling algorithm. As outlined in Algorithm 2, RDDM generates samples by initializing a random signal \mathbf{x}_T from a Gaussian distribution and iteratively applying the reverse process. For each time step t from T to 1, we sample a latent variable \mathbf{z} from a Gaussian distribution if $t > 1$, otherwise, $\mathbf{z} = \mathbf{0}$. We then obtain the ROI masked signal $\mathbf{x}_t^{[p]}$ as $\rho_\phi(\mathbf{x}_t, c, t)$. Finally, we update the signal \mathbf{x}_{t-1} using ϵ_θ . The resulting signal \mathbf{x}_0 is the final output of the sampling algorithm. The forward and reverse processes of our proposed method are depicted in Figure 1.

4 Experiment Setup

4.1 Implementation Details

Datasets. Following, we present the details of the datasets that are used in this study.

- WESAD (Schmidt et al. 2018) comprises approximately 24 hours of synchronized Lead-II ECG data (sampled at 700Hz) and PPG data (sampled at 64 Hz) from 15 subjects with annotations of stress and affect states.
- MIMIC AFib (Bashar et al. 2019) is a subset of the MIMIC-III waveform database (Johnson et al. 2016) with binary labels for AFib. It consists of 35 critically ill adults, out of which 19 suffer from AFib. Both ECG and PPG signals are sampled at 125 Hz.
- PPG-BP (Liang et al. 2018) consists of 657 pre-windowed PPG signals from 219 subjects, along with rich metadata describing the health conditions of each subject. The PPG was recorded at a sampling rate of 1 kHz. Please note that unlike the other datasets used in this study, this dataset only consists of PPG signals and does not contain ECG.
- Cuffless BP (Kachuee et al. 2015) is a dataset derived from the MIMIC-II waveform database (Saeed et al. 2011). It consists of thousands of PPG, ECG, and ABP signals recorded from various hospitals between 2001 and 2008, sampled at a frequency of 125 Hz.
- CAPNO (Karlen 2021) comprises approximately 5.6 hours of paired Lead-II ECG and PPG signals sampled at 300 Hz. The data was recorded from 42 subjects under medical supervision.
- BIDMC (Pimentel et al. 2016) is a dataset collected from 53 ICU patients, consisting of around 7 hours of concurrent Multi-lead ECG (Lead II, V, and AVR) and PPG signals, both recorded at a frequency of 125 Hz. We utilize only the Lead-II ECG in our work.

- DALIA (Reiss et al. 2019) is a dataset with approximately 35 hours of synchronized PPG and Lead-II ECG signals, recorded from 15 subjects performing daily-life activities like walking and driving. The ECG and PPG were recorded at sampling rates of 700 and 64 Hz respectively.

To train RDDM, we follow a similar setup to CardioGAN (Sarkar and Etemad 2021) and create a diverse training set by combining the paired ECG and PPG signals from WESAD, DALIA, CAPNO, BIDMC, and MIMIC-AFib. In particular, we use the data from 80% of the subjects for training and the remaining 20% of subjects are used for cross-subject evaluation. Additionally, for CardioBench, we use WESAD, DALIA, MIMIC AFib, PPG-BP, and Cuffless BP.

Data pre-processing. We follow standard pre-processing steps for both ECG (Pan and Tompkins 1985; Makowski et al. 2021) and PPG (Makowski et al. 2021). First, we resample both signals at 128 Hz. Subsequently, we apply a high-pass Butterworth filter with a cut-off frequency of 0.5 Hz on ECG signals. Similarly, a band-pass Butterworth filter with a cut-off frequency between 0.5 to 8 Hz is applied on PPG signals. Additionally, subject-specific z-score normalization is applied to compensate for individual variances. Further, min-max scaling is applied to re-scale the ECG and PPG signals between $[-1, 1]$. Finally, we segment both signals into 4-second windows which are used for RDDM training.

Network architecture. We adapt the UNet architecture from (Ho, Jain, and Abbeel 2020; Dhariwal and Nichol 2021) and modify for 1D signals, which is used as the backbone for ϵ_θ and ρ_ϕ modules. Additionally, we replace the default self-attention blocks with cross-attention modules to better condition the generated ECG for a given PPG input (Saharia et al. 2022). We apply cross-attention between *ECG at the previous timestep* and *input PPG*. Diffusion timestep t is added to each UNet block using Sinusoidal Position Embedding, similar to (Ho, Jain, and Abbeel 2020; Dhariwal and Nichol 2021; Saharia et al. 2022). The UNet architecture comprises 6 down-sampling and 6 up-sampling stages. During down-sampling, the number of filters is increased from 64 to 1024, doubling at each stage. Conversely, in up-sampling stages, the number of filters is decreased from 1024 to 64, halving at each stage. The final output is obtained from a single-channel convolutional layer. For a fair comparison, we also use the same architecture for the baseline DDPMs.

Training setup and hyper-parameters. We train RDDM on $4 \times$ NVIDIA A100 GPUs with a batch size of 512. We choose AdamW (Loshchilov and Hutter 2018) as the optimizer with a Cosine learning rate scheduler, having a base learning rate of 10^{-4} for a total of 500 epochs. We use a fixed linear variance scheduler with $\beta \in (0.0001, 0.2)$. We empirically set the sampling steps to 10 and the ROI window size to $\Gamma = 32$ (see related study in Sec. 5.5). The loss coefficients in Equation 9 are set to $\lambda_1 = 100$ and $\lambda_2 = 1$.

4.2 Evaluation

We perform two types of evaluation to ensure a robust assessment of generated ECGs. First, we perform ‘standard’ evaluation to report the error/distance metrics between generated ECG and ground-truth ECG. Additionally, to understand

Task	Dataset(s)	Arch.	Input	Epochs	Optim	LR	Batch	Window	Metrics	GPUs
HR	DALIA & WESAD	-	ECG (raw)	-	-	-	-	8 sec	MAE	-
AFib	MIMIC-AFib	VGG-13	ECG (STFT)	25	AdamW	$1e^{-4}$	64	4 sec	F1, Acc.	$2 \times A100$
Stress & Affect	WESAD	Transformer	ECG (raw)	10	AdamW	$1e^{-4}$	64	60 sec	F1, Acc.	$2 \times A100$
Diabetes	PPG-BP	VGG-11	ECG (STFT)	50	AdamW	$1e^{-4}$	64	2 sec	F1, Acc.	$2 \times A100$
BP	Cuffless BP	1D UNet	ECG & PPG (raw)	50	Adam	$5e^{-4}$	512	10 sec	MAE (SBP), MAE (DBP)	$4 \times A100$

Table 1: Implementation details for evaluation on CardioBench.

the utility and predictive power of the generated ECG in analyzing cardiac activity, we introduce a comprehensive evaluation benchmark, CardioBench. Following, we provide the details for both setups.

Standard Evaluation. The generated ECG signals are evaluated on standard quantitative metrics in a subject-independent manner. Following CardioGAN (Sarkar and Etemad 2021), we generate 4-second ECG signals from RDDM and measure Root Mean Squared Error (RMSE) and Fréchet Distance (FD) with respect to the original ECG.

CardioBench. To test the utility of the generated ECG in analyzing cardiac activity, e.g., detecting CVDs, we introduce CardioBench. CardioBench is designed as a comprehensive and challenging benchmark aimed at evaluating the generated ECG signals, with a focus on their ability to facilitate the detection of a range of heart conditions. Our proposed benchmark consists of 5 cardiac-related tasks across 6 datasets:

(1) **Heart rate estimation** on *WESAD* (Schmidt et al. 2018) and *DALIA* (Reiss et al. 2019) datasets: Accurate HR estimation from synthetic ECG depends on the presence of R-peaks at the correct temporal locations with consistent RR intervals. Therefore, HR estimation is included to evaluate the temporal fidelity of generative models. In line with previous literature (Sarkar and Etemad 2021), we perform HR estimation by calculating Mean Absolute Error (MAE) in beats per minute (bpm) which measures the difference between the ground truth HR and the HR estimated from the generated ECG signals. We use (Hamilton 2002) for HR estimation from the generated ECG using 8-second windows.

(2) **Atrial Fibrillation detection** on *MIMIC-AFib* (Bashar et al. 2019) dataset: AFib detection in synthetic ECG demands accurate generation of non-periodic irregular rhythms, absent P waves, variable ventricular rates, possible fibrillatory waves, and temporally compressed QRS complexes (Burns and Buttner 2023). This ability serves as a vital measure of a generative model’s proficiency in capturing both detailed morphological features and non-periodic temporal dynamics, crucial for AFib detection. We evaluate the binary AFib classification using 4-second ECG windows. We apply Short-time Fourier Transform (STFT) on the generated ECG and use it as input to a standard VGG-13 (Sallem et al. 2020).

(3) **Stress and Affect classification** on *WESAD* (Schmidt et al. 2018) dataset: Stress and affect classification using

generated ECG requires an accurate representation of nuanced changes in heart rate variability (Kim et al. 2018), sympathetic-parasympathetic balance (Tewari et al. 2006), and emotional state-induced ECG patterns (Sarkar and Etemad 2020). This capability acts as an indicator of a generative model’s adeptness at capturing subtle morphological features and dynamic temporal changes. We perform three-class classification of neutral, stress, and amusement. We adopt a Transformer-based classifier (Behinaein et al. 2021) using a window size of 60 seconds for a fair comparison with previous methods.

(4) **Diabetes detection** on *PPG-BP* (Liang et al. 2018) dataset: Diabetes detection from synthetic ECG involves accurate capture of the anomalies in the ECG signals induced by the disease, such as prolonged QT intervals, alterations in the QRS complex, and changes in the ST-T segment (Veglio, Chinaglia, and Cavallo-Perin 2004). The ability to generate these distinguishing spatiotemporal variations serves as a crucial metric for ECG generation. We evaluate the generated ECG on the binary classification task of Diabetes detection. We adopt a VGG-11 classifier (Reiss et al. 2019) and apply STFT on the input. We use 2-second windows as per the pre-windowed dataset.

(5) **Blood Pressure estimation** on *Cuffless-BP* (Kachuee et al. 2015) dataset: Estimation of BP from synthetic ECG require an accurate representation of ECG-derived features like heart rate, P wave, QRS duration, and QT interval, all of which correlate with blood pressure levels (Peng et al. 2006; Bird et al. 2020). This serves as a vital test of a generative model’s capacity to reproduce fine-grained morphological and dynamic temporal features. We estimate the Systolic Blood Pressure (SBP) and Diastolic Blood Pressure (DBP), and evaluate using MAE. We use a simple 1D-UNet (Long and Wang 2023) as our estimator, and use 10 second windows following prior methods.

Comparison setup. HR estimation using (Hamilton 2002) does not require specific training, and generated ECG is used directly for estimation. The other 4 tasks, however, require training neural nets. As a result, for AFib detection and Stress and Affect classification, we train the models using real ECG and evaluate on both real and generated ECG signals. We treat performance on the real ECG of the test splits as the upper bound and compare the performance of the generated ECG in each specific task. A similar approach is taken for

Dataset	Method	RMSE(\downarrow)	FD(\downarrow)
WESAD	Sarkar and Etemad	0.37	29.15
	DDPM (T = 10)	0.36	30.58
	DDPM (T = 25)	0.29	14.50
	DDPM (T = 50)	0.26	9.17
	RDDM (Ours)	0.21	3.93
CAPNO	Sarkar and Etemad	0.38	31.10
	DDPM (T = 10)	0.33	21.56
	DDPM (T = 25)	0.28	8.98
	DDPM (T = 50)	0.24	5.15
	RDDM (Ours)	0.19	2.94
DALIA	Sarkar and Etemad	0.42	27.51
	DDPM (T = 10)	0.38	19.58
	DDPM (T = 25)	0.32	7.62
	DDPM (T = 50)	0.30	5.95
	RDDM (Ours)	0.25	5.84
BIDMC	Sarkar and Etemad	0.63	154.64
	DDPM (T = 10)	0.37	22.00
	DDPM (T = 25)	0.33	10.38
	DDPM (T = 50)	0.29	7.35
	RDDM (Ours)	0.24	6.72
MIMIC	Sarkar and Etemad	0.54	99.03
	DDPM (T = 10)	0.36	20.97
	DDPM (T = 25)	0.29	8.75
	DDPM (T = 50)	0.27	6.72
	RDDM (Ours)	0.22	6.71

Table 2: Quantitative comparison with the state-of-the-art methods for PPG-to-ECG translation. RDDM consistently outperforms prior works by a large margin in all the setups. Moreover, RDDM also outperforms the best DDPM which uses 50 sampling steps compared to ours using 10.

BP estimation, with the distinction that it is treated as a multimodal task (Elgendi et al. 2019), for which we use the PPG and real ECG for training. Finally, to conduct Diabetes detection, we directly train the classifier on the generated ECG due to the unavailability of original ECG signals and use the ECG generated from the test set for testing. We present the key implementation details for evaluation on CardioBench in Table 1.

5 Results

5.1 Standard Evaluation

The results presented in Table 2 clearly show that RDDM surpasses the current state-of-the-art PPG-to-ECG translation model (Sarkar and Etemad 2021) by a significant margin on all the datasets. For instance, RDDM improves RMSE from 0.63 to 0.24 on BIDMC. Additionally, RDDM outperforms the best version of the baseline DDPM by a large margin even when RDDM uses just 10 sampling steps compared to the best DDPM that uses 50 sampling steps.

5.2 Evaluation on CardioBench

In this section, we evaluate the ECG generated by our method (RDDM) on the proposed CardioBench benchmark. For comparison, we use the latest baselines from the literature, as well as DDPM with 50 sampling steps.

HR estimation. Table 3 demonstrates that RDDM outperforms prior works and DDPM by achieving a substantially lower MAE on both datasets. On DALIA, RDDM achieves state-of-the-art MAE of 4.49 bpm, outperforming the best PPG-based method (Song, Nam, and Kim 2021) by 1.53 bpm. RDDM demonstrates an improvement of 1.17 bpm over DDPM. Next, on WESAD, RDDM achieves state-of-the-art MAE of 1.40 bpm compared to the best PPG-based method (Reiss et al. 2019), CardioGAN, and DDPM, which achieve 9.50, 8.60, and 3.22 respectively.

AFib detection. As reported in Table 4, RDDM outperforms both existing PPG-based methods and DDPM. Specifically, RDDM achieves F1=0.71 and Acc. = 0.65, demonstrating improvements of 11% and 22% over DDPM, and approximately 16% and 4% when compared to the top-performing PPG-based method (Aliamiri and Shen 2018). Moreover, our method shows the closest performance to the upper bound, achieving an F1 of 0.71 vs. the upper bound 0.74.

Stress and Affect classification. As highlighted in Table 5, our method achieves superior performance compared to prior methods, with F1 = 0.64 and Acc. = 0.71. Compared to the best PPG-based method (Schmidt et al. 2018), RDDM exhibits a boost of 7% in F1 and 6% in Acc. Furthermore, RDDM also outperforms DDPM which achieves similar performance to (Schmidt et al. 2018). Moreover, our method shows the closest performance to the upper bound, achieving an F1 of 0.64, compared to the upper bound F1 of 0.68.

Diabetes detection. As illustrated in Table 6, RDDM significantly improves upon existing PPG-based methods and DDPM in diabetes detection. Notably, our method achieves an Acc. of 0.80 and an F1 of 0.52, indicating substantial improvements of 15% in Acc. and 8% in F1, compared to the leading PPG-based method (Avram et al. 2020), which registers an Acc. of 0.65 and an F1 of 0.44. Furthermore, RDDM outperforms DDPM by an impressive margin of 33% and 9% in terms of Acc. and F1 respectively. It is notable that despite the multimodal classifier introduced by Performer (Lan 2023) being pre-trained on a large corpus, a classifier trained from scratch only on ECG signals generated by RDDM still outperforms it, marking an Acc. improvement of 4%. This result shows the high-fidelity of the ECG generated by RDDM.

BP estimation. As demonstrated in Table 7, RDDM attains a state-of-the-art performance in BP estimation, achieving an MAE-SBP of 2.48 and an MAE-DBP of 1.20. Remarkably, RDDM makes significant improvements compared to the best PPG-based method (Vardhan et al. 2021) with a margin of 2.68 in MAE-SBP and 2.25 in MAE-DBP. This performance gain supports the notion that the ECG generated by RDDM provides complementary information learned through training on a large dataset of PPG-ECG pairs, which could enhance the performance beyond what is possible with PPG-only BP estimation methods. Furthermore, RDDM even outperforms the upper bound by achieving MAE-SBP = 2.48 vs. the upper bound MAE-SBP = 3.31, and performs closely

Dataset	Test modality	Estimator	MAE(↓)
DALIA	PPG	Schäck et al.	20.5
	PPG	Reiss et al.	11.1
	PPG	Wójcikowski and Pankiewicz	6.77
	PPG	Song, Nam, and Kim	6.02
	Gen. ECG (Sarkar and Etemad)	Hamilton	8.30
	Gen. ECG (DDPM)	Hamilton	5.66
	Gen. ECG (Ours)	Hamilton	4.49
WESAD	PPG	Schäck et al.	19.90
	PPG	Reiss et al.	9.50
	Gen. ECG (Sarkar and Etemad)	Hamilton	8.60
	Gen. ECG (DDPM)	Hamilton	4.67
	Gen. ECG (Ours)	Hamilton	1.40

Table 3: Evaluation of HR estimation. RDDM outperforms the prior works by a significant margin.

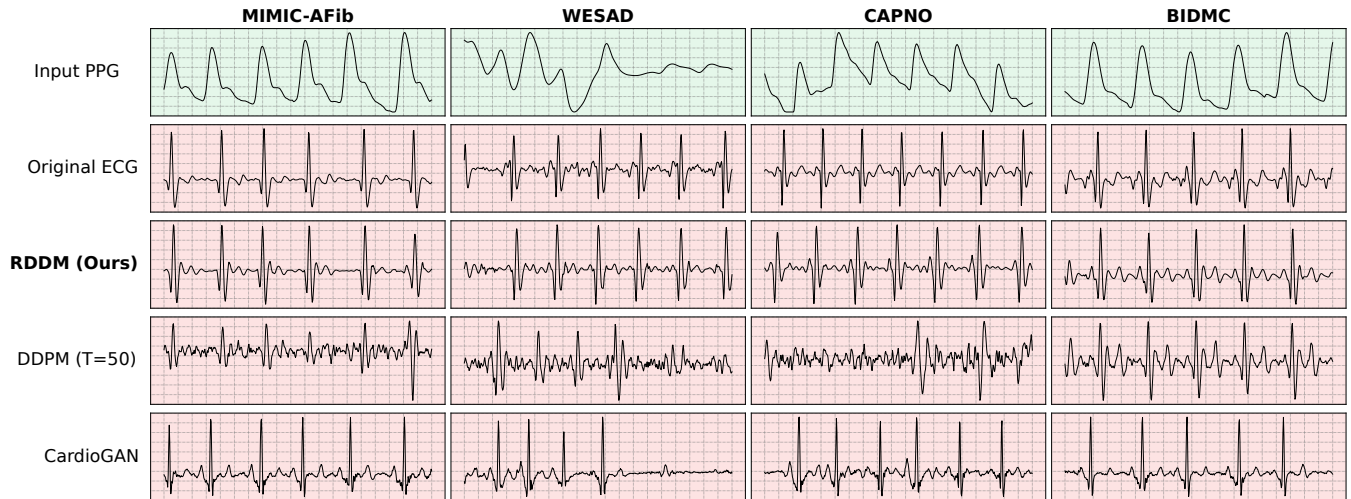


Figure 2: Qualitative comparison on generated ECG using varying nature of input PPG. RDDM consistently generates ECG that closely resembles the original ECG, even when the input PPG is noisy (e.g., WESAD), unlike CardioGAN and DDPM.

Test modality	Classifier	Acc.(↑)	F1(↑)
ECG (Upper bound)	VGG-13	0.73	0.74
PPG	Aliamiri and Shen*	0.61	0.55
PPG	Shen et al.*	0.51	0.47
Gen. ECG (DDPM)	VGG-13	0.43	0.60
Gen. ECG (Ours)	VGG-13	0.65	0.71

Table 4: Evaluation of AFib detection on MIMIC subset. RDDM significantly outperforms prior works and reports an F1 score close to the upper bound. *reimplemented by us.

to the upper bound in terms of MAE-DBP. We believe this improvement over upper-bound is observed as the original ECG signals were quite noisy. Although DDPM achieves the lowest MAE-DBP = 0.63, it performs poorly in terms of estimating SBP, with an MAE-SBP = 5.95. In contrast, RDDM provides a balanced performance for estimating both SBP and DBP, demonstrating real-world utility.

Test modality	Classifier	Acc.(↑)	F1(↑)
ECG (Upper bound)	Behinaein et al.	0.75	0.68
PPG	Schmidt et al.	0.65	0.57
PPG	Lisowska et al.	0.58	-
Gen. ECG (DDPM)	Behinaein et al.	0.65	0.58
Gen. ECG (Ours)	Behinaein et al.	0.71	0.64

Table 5: Evaluation of Stress and Affect detection on WESAD dataset. RDDM not only outperforms prior works, also achieves a very close performance to the upper bound.

5.3 Qualitative Analysis

In Figure 2, we perform qualitative comparisons between RDDM with the best DDPM (50 sampling steps) and CardioGAN. Our comparison includes samples from MIMIC-AFib, WESAD, CAPNO, and BIDMC, covering a diverse set of PPG inputs of varying natures. The ECG signals generated by RDDM demonstrate high-fidelity compared to the ground-

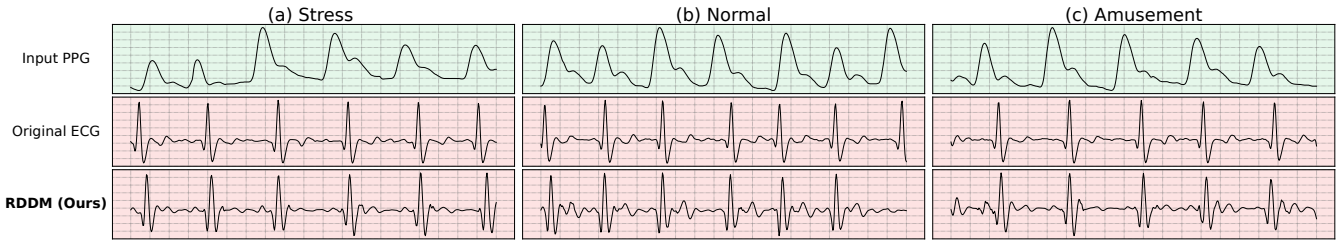


Figure 3: RDDM-generated ECG signals of subjects from the WESAD dataset experiencing different affective states.



Figure 4: RDDM-generated ECG signals of 3 AFib patients from the test set of MIMIC-AFib dataset.

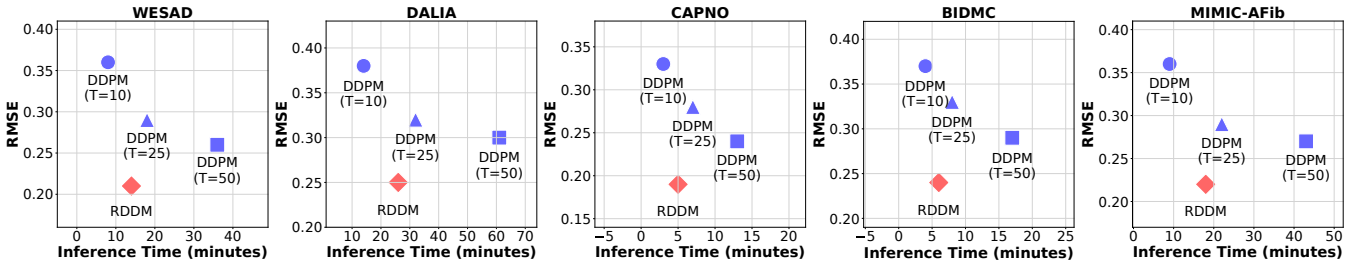


Figure 5: The plot illustrates the tradeoff between RMSE and inference time. RDDM demonstrates the best balance across all datasets, outperforming DDPM with (T= 50) even though the latter is approx. $2.5\times$ slower, highlighting the efficiency and effectiveness of the proposed method.

Test modality	Classifier	Acc.(\uparrow)	F1(\uparrow)
PPG	Avram et al.*	0.65	0.44
PPG + Gen. ECG (Lan)	Lan**	0.76	-
Gen. ECG (DDPM)	VGG-11	0.47	0.43
Gen. ECG (Ours)	VGG-11	0.80	0.52

Table 6: Evaluation of Diabetes detection on PPG-BP dataset. RDDM outperforms both uni-modal (PPG) and multimodal (PPG + Gen. ECG) prior works. *reimplemented by us. **uses pre-trained weights.

truth ECG signals, across all datasets. While comparing samples from MIMIC and BIDMC, generated ECG signals from CardioGAN and DDPM fail to accurately reproduce the QRS complex structure and significantly deviate from the original ECG morphology. Additionally, the generated ECG samples from DDPM exhibit higher noise as shown in MIMIC. In stark contrast, RDDM generates ECG signals that are almost identical to the ground truth ECG. Furthermore, in case of noisy PPG input, RDDM maintains temporal fidelity (as shown in WESAD) and also approximates the fine-grained

structure of the QRS complex with remarkable precision. In contrast, DDPM and CardioGAN struggle to accurately generate QRS complex, leading to distorted temporal fidelity which results in irregular RR intervals. We provide additional generated ECG samples corresponding to different affective states in Figure 3 and samples corresponding to AFib patients in Figure 4.

5.4 Inference Time Comparison

Figure 5 presents an inference time comparison between RDDM and DDPM variants. We measure the total inference time required to translate the PPG test set of each dataset into ECG signals for standard evaluation (i.e., RMSE). We use a single Nvidia 2080Ti GPU for this experiment. Across all datasets, RDDM consistently exhibits the best balance between RMSE and inference time. In contrast, DDPM with 50 sampling steps not only performs worse in RMSE but is also slower by a factor of approx. $2.5\times$, highlighting the superiority of RDDM in both speed and fidelity.

Test modality	Estimator	MAE	
		SBP(↓)	DBP(↓)
PPG + ECG (Upper bound)	UNet	3.31	0.75
PPG	Ibtehaz et al.	5.73	3.45
PPG	Vardhan et al.	5.16	2.89
PPG + Gen. ECG (DDPM)	UNet	5.95	0.63
PPG + Gen. ECG (Ours)	UNet	2.48	1.20

Table 7: Evaluation of BP estimation on Cuffless-BP dataset. RDDM clearly outperforms prior methods, and even shows a lower MAE-SBP than the upper bound.

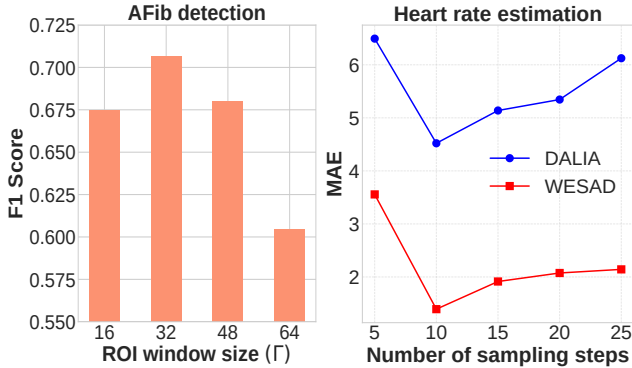


Figure 6: Left: RDDM shows an improvement in performance when Γ is increased from 16 to 32, but beyond that, the performance degrades. Right: On both datasets, RDDM achieves the best performance in just 10 sampling steps.

5.5 Analysis of RDDM

ROI window size. To find the optimal ROI window size, we vary the value of Γ between 16 and 64 with a step size of 16 and track their performance on AFib detection. Figure 6 (left) shows that both extremely small ($\Gamma = 16$) and large window ($\Gamma = 64$) sizes degrade performance as they fail to capture the ROI accurately. An ROI window size of 32 works best in our setup, which is used in all the experiments.

Sampling steps. To find the optimal value, we vary the number of sampling steps from 5 to 25 with a step size of 5 and evaluate on HR estimation. The results presented in Figure 6 (right) exhibit a similar trend for both WESAD and DALIA. The performance significantly improves when the number of sampling steps is increased from 5 to 10. However, the performance declines when the sampling steps are increased further. Thus, we set the sampling steps as 10 for RDDM.

6 Conclusion

In this paper, we introduce a novel diffusion model RDDM, specifically designed for high-fidelity PPG-to-ECG translation. To the best of our knowledge, RDDM is the first diffusion model for cross-modal signal-to-signal translation within the domain of bio-signals. The issue of indiscriminate noise addition during the forward process limits diffusion models like DDPMs from capturing both temporal structures

and morphological details. RDDM improves upon them by introducing a novel ROI-guided forward process and by disentangling the reverse process into 2 distinct components: (i) capturing the global temporal structure of ECG signals and (ii) capturing fine-grained local details. This allows RDDM to generate high-fidelity ECG signals in just 10 sampling steps. To assess the utility of generated ECG, we introduce a comprehensive evaluation benchmark CardioBench. CardioBench includes 5 popular cardiac-related tasks: HR and BP estimation, stress classification, and detection of AFib and diabetes. RDDM achieves superior performance compared to DDPM and prior works, excelling in both standard quantitative measures and performance on CardioBench.

Acknowledgments

This work was supported by Mitacs, Vector Institute, and Ingenuity Labs Research Institute.

References

- Adib, E.; Fernandez, A.; Afghah, F.; and Prevost, J. J. 2023. Synthetic ECG signal generation using probabilistic diffusion models. *arXiv preprint arXiv:2303.02475*.
- Alcaraz, J. M. L.; and Strodthoff, N. 2023. Diffusion-based conditional ECG generation with structured state space models. *Computers in Biology and Medicine*, 107115.
- Aliamiri, A.; and Shen, Y. 2018. Deep learning based atrial fibrillation detection using wearable photoplethysmography sensor. In *IEEE EMBS International Conference on Biomedical & Health Informatics*, 442–445.
- Avram, R.; Olgin, J. E.; Kuhar, P.; Hughes, J. W.; Marcus, G. M.; Pletcher, M. J.; Aschbacher, K.; and Tison, G. H. 2020. A digital biomarker of diabetes from smartphone-based vascular signals. *Nature Medicine*, 26(10): 1576–1582.
- Banerjee, R.; Sinha, A.; Choudhury, A. D.; and Visvanathan, A. 2014. PhotoECG: Photoplethysmography to estimate ECG parameters. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, 4404–4408.
- Bashar, S. K.; Ding, E.; Walkey, A. J.; McManus, D. D.; and Chon, K. H. 2019. Noise detection in electrocardiogram signals for intensive care unit patients. *IEEE Access*, 7: 88357–88368.
- Bayoumy, K.; Gaber, M.; Elshafeey, A.; Mhaimed, O.; Dineen, E. H.; Marvel, F. A.; Martin, S. S.; Muse, E. D.; Turakhia, M. P.; Tarakji, K. G.; et al. 2021. Smart wearable

- devices in cardiovascular care: where we are and how to move forward. *Nature Reviews Cardiology*, 18(8): 581–599.
- Behinaein, B.; Bhatti, A.; Rodenburg, D.; Hungler, P.; and Etemad, A. 2021. A transformer architecture for stress detection from ecg. In *International Symposium on Wearable Computers*, 132–134.
- Bird, K.; Chan, G.; Lu, H.; Greeff, H.; Allen, J.; Abbott, D.; Menon, C.; Lovell, N. H.; Howard, N.; Chan, W.-S.; et al. 2020. Assessment of hypertension using clinical electrocardiogram features: a first-ever review. *Frontiers in Medicine*, 7: 583331.
- Blattmann, A.; Rombach, R.; Ling, H.; Dockhorn, T.; Kim, S. W.; Fidler, S.; and Kreis, K. 2023. Align your latents: High-resolution video synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22563–22575.
- Burns, E.; and Buttner, R. 2023. Guidelines for the management of atrial fibrillation.
- Camm, A. J.; Corbucci, G.; and Padeletti, L. 2012. Usefulness of continuous electrocardiographic monitoring for atrial fibrillation. *The American Journal of Cardiology*, 110(2): 270–276.
- Chen, J.; Liao, K.; Wei, K.; Ying, H.; Chen, D. Z.; and Wu, J. 2022. ME-GAN: Learning panoptic electrocardio representations for multi-view ECG synthesis conditioned on heart diseases. In *International Conference on Machine Learning*, 3360–3370. PMLR.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34: 8780–8794.
- Elgendi, M.; Fletcher, R.; Liang, Y.; Howard, N.; Lovell, N. H.; Abbott, D.; Lim, K.; and Ward, R. 2019. The use of photoplethysmography for assessing hypertension. *NPJ Digital Medicine*, 2(1): 60.
- Golany, T.; and Radinsky, K. 2019. Pgens: Personalized generative adversarial networks for ecg synthesis to improve patient-specific deep ecg classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 557–564.
- Hamilton, P. 2002. Open source ECG analysis. In *Computers in Cardiology*, 101–104. IEEE.
- Hannun, A. Y.; Rajpurkar, P.; Haghpanahi, M.; Tison, G. H.; Bourn, C.; Turakhia, M. P.; and Ng, A. Y. 2019. Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network. *Nature Medicine*, 25(1): 65–69.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Ibtehaz, N.; Mahmud, S.; Chowdhury, M. E.; Khandakar, A.; Salman Khan, M.; Ayari, M. A.; Tahir, A. M.; and Rahman, M. S. 2022. PPG2ABP: Translating photoplethysmogram (PPG) signals to arterial blood pressure (ABP) waveforms. *Bioengineering*, 9(11): 692.
- Johnson, A. E.; Pollard, T. J.; Shen, L.; Lehman, L.-w. H.; Feng, M.; Ghassemi, M.; Moody, B.; Szolovits, P.; Anthony Celi, L.; and Mark, R. G. 2016. MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3(1): 1–9.
- Kachuee, M.; Kiani, M. M.; Mohammadzade, H.; and Shabany, M. 2015. Cuff-less high-accuracy calibration-free blood pressure estimation using pulse transit time. In *IEEE International Symposium on Circuits and Systems*, 1006–1009.
- Karlen, W. 2021. CapnoBase IEEE TBME respiratory rate benchmark.
- Kauffman, D. 2020. Cardiovascular disease burden, deaths are rising around the world. *American College of Cardiology*, 9.
- Kim, H.-G.; Cheon, E.-J.; Bai, D.-S.; Lee, Y. H.; and Koo, B.-H. 2018. Stress and heart rate variability: a meta-analysis and review of the literature. *Psychiatry Investigation*, 15(3): 235.
- Lan, E. 2023. Performer: A novel PPG-to-ECG reconstruction transformer for a digital biomarker of cardiovascular disease detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 1991–1999.
- Liang, Y.; Chen, Z.; Liu, G.; and Elgendi, M. 2018. A new, short-recorded photoplethysmogram dataset for blood pressure monitoring in China. *Scientific Data*, 5(1): 1–7.
- Lisowska, A.; et al. 2021. Catching patient’s attention at the right time to help them undergo behavioural change: stress classification experiment from blood volume pulse. In *19th International Conference on Artificial Intelligence in Medicine*, 72–82. Springer.
- Long, W.; and Wang, X. 2023. BPNet: A multi-modal fusion neural network for blood pressure estimation using ECG and PPG. *Biomedical Signal Processing and Control*, 86: 105287.
- Loshchilov, I.; and Hutter, F. 2018. Decoupled weight decay regularization. In *International Conference on Learning Representations*.
- Makowski, D.; Pham, T.; Lau, Z. J.; Brammer, J. C.; Lespinasse, F.; Pham, H.; Schölzel, C.; and Chen, S. H. A. 2021. NeuroKit2: A Python toolbox for neurophysiological signal processing. *Behavior Research Methods*, 53(4): 1689–1696.
- Neifar, N.; Ben-Hamadou, A.; Mdhaffar, A.; and Jmaiel, M. 2023. DiffECG: A generalized probabilistic diffusion model for ECG signals synthesis. *arXiv preprint arXiv:2306.01875*.
- Pan, J.; and Tompkins, W. J. 1985. A real-time QRS detection algorithm. *IEEE Transactions on Biomedical Engineering*, (3): 230–236.
- Peng, S.; Yu, Y.; Hao, K.; Xing, H.; Li, D.; Chen, C.; Huang, A.; Hong, X.; Feng, Y.; Zhang, Y.; et al. 2006. Heart rate-corrected QT interval duration is significantly associated with blood pressure in Chinese hypertensives. *Journal of Electrocardiology*, 39(2): 206–210.
- Pimentel, M. A.; Johnson, A. E.; Charlton, P. H.; Birrenkott, D.; Watkinson, P. J.; Tarassenko, L.; and Clifton, D. A. 2016. Toward a robust estimation of respiratory rate from pulse oximeters. *IEEE Transactions on Biomedical Engineering*, 64(8): 1914–1923.

- Reiss, A.; Indlekofer, I.; Schmidt, P.; and Van Laerhoven, K. 2019. Deep PPG: large-scale heart rate estimation with convolutional neural networks. *Sensors*, 19(14): 3079.
- Ruan, L.; Ma, Y.; Yang, H.; He, H.; Liu, B.; Fu, J.; Yuan, N. J.; Jin, Q.; and Guo, B. 2023. Mm-diffusion: Learning multi-modal diffusion models for joint audio and video generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10219–10228.
- Saeed, M.; Villarroel, M.; Reisner, A. T.; Clifford, G.; Lehman, L.-W.; Moody, G.; Heldt, T.; Kyaw, T. H.; Moody, B.; and Mark, R. G. 2011. Multiparameter intelligent monitoring in intensive care II (MIMIC-II): a public-access intensive care unit database. *Critical Care Medicine*, 39(5): 952.
- Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E. L.; Ghasemipour, K.; Gontijo Lopes, R.; Karagol Ayan, B.; Salimans, T.; et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35: 36479–36494.
- Sallem, M.; Ghrissi, A.; Saadaoui, A.; and Zarzoso, V. 2020. Detection of cardiac arrhythmias from varied length multi-channel electrocardiogram recordings using deep convolutional neural networks. In *Computing in Cardiology*, 1–4.
- Sarkar, P.; and Etemad, A. 2020. Self-supervised ECG representation learning for emotion recognition. *IEEE Transactions on Affective Computing*, 13(3): 1541–1554.
- Sarkar, P.; and Etemad, A. 2021. Cardiogan: Attentive generative adversarial network with dual discriminators for synthesis of ecg from ppg. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 488–496.
- Schäck, T.; et al. 2017. Computationally efficient heart rate estimation during physical exercise using photoplethysmographic signals. In *European Signal Processing Conference*, 2478–2481.
- Schmidt, P.; Reiss, A.; Duerichen, R.; Marberger, C.; and Van Laerhoven, K. 2018. Introducing wesad, a multimodal dataset for wearable stress and affect detection. In *Proceedings of the 20th ACM international conference on multimodal interaction*, 400–408.
- Shen, Y.; Voisin, M.; Aliamiri, A.; Avati, A.; Hannun, A.; and Ng, A. 2019. Ambulatory atrial fibrillation monitoring using wearable photoplethysmography with deep learning. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1909–1916.
- Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, 2256–2265. PMLR.
- Song, S. B.; Nam, J. W.; and Kim, J. H. 2021. NAS-PPG: PPG-based heart rate estimation using neural architecture search. *IEEE Sensors Journal*, 21(13): 14941–14949.
- Steinhuibl, S. R.; Waalen, J.; Edwards, A. M.; Ariniello, L. M.; Mehta, R. R.; Ebner, G. S.; Carter, C.; Baca-Motes, K.; Felicione, E.; Sarich, T.; et al. 2018. Effect of a home-based wearable continuous ECG monitoring patch on detection of undiagnosed atrial fibrillation: the mSToPS randomized clinical trial. *Jama*, 320(2): 146–155.
- Tewari, H. K.; Gadia, R.; Kumar, D.; Venkatesh, P.; and Garg, S. P. 2006. Sympathetic–parasympathetic activity and reactivity in central serous chorioretinopathy: a case-control study. *Investigative Ophthalmology & Visual Science*, 47(8): 3474–3478.
- Tian, X.; Zhu, Q.; Li, Y.; and Wu, M. 2022. Cross-domain joint dictionary learning for ECG inference from PPG. *IEEE Internet of Things Journal*.
- Vardhan, K. R.; Vedanth, S.; Poojah, G.; Abhishek, K.; Kumar, M. N.; and Vijayaraghavan, V. 2021. BP-Net: Efficient deep learning for continuous arterial blood pressure estimation using photoplethysmogram. In *IEEE International Conference on Machine Learning and Applications*, 1495–1500.
- Veglio, M.; Chinaglia, A.; and Cavallo-Perin, P. 2004. QT interval, cardiovascular risk factors and risk of death in diabetes. *Journal of Endocrinological Investigation*, 27: 175–181.
- Vo, K.; Naeini, E. K.; Naderi, A.; Jilani, D.; Rahmani, A. M.; Dutt, N.; and Cao, H. 2021. P2E-WGAN: ECG waveform synthesis from PPG with conditional wasserstein generative adversarial networks. In *Proceedings of the 36th Annual ACM Symposium on Applied Computing*, 1030–1036.
- Wójcikowski, M.; and Pankiewicz, B. 2020. Photoplethysmographic time-domain heart rate measurement algorithm for resource-constrained wearable devices and its implementation. *Sensors*, 20(6): 1783.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on computer vision*, 2223–2232.
- Zhu, Q.; Tian, X.; Wong, C.-W.; and Wu, M. 2021. Learning your heart actions from pulse: ECG waveform reconstruction from PPG. *IEEE Internet of Things Journal*, 8(23): 16734–16748.