

# DeScoD-ECG: Deep Score-Based Diffusion Model for ECG Baseline Wander and Noise Removal

Huayu Li, Gregory Ditzler, Janet Roveda and Ao Li

**Abstract—Objective:** Electrocardiogram (ECG) signals commonly suffer noise interference, such as baseline wander. High-quality and high-fidelity reconstruction of the ECG signals is of great significance to diagnosing cardiovascular diseases. Therefore, this paper proposes a novel ECG baseline wander and noise removal technology. **Methods:** We extended the diffusion model in a conditional manner that was specific to the ECG signals, namely the Deep Score-Based Diffusion model for Electrocardiogram baseline wander and noise removal (DeScoD-ECG). Moreover, we deployed a multi-shots averaging strategy that improved signal reconstructions. We conducted the experiments on the QT Database and the MIT-BIH Noise Stress Test Database to verify the feasibility of the proposed method. Baseline methods are adopted for comparison, including traditional digital filter-based and deep learning-based methods. **Results:** The quantities evaluation results show that the proposed method obtained outstanding performance on four distance-based similarity metrics with at least 20% overall improvement compared with the best baseline method. **Conclusion:** This paper demonstrates the state-of-the-art performance of the DeScoD-ECG for ECG baseline wander and noise removal, which has better approximations of the true data distribution and higher stability under extreme noise corruptions. **Significance:** This study is one of the first to extend the conditional diffusion-based generative model for ECG noise removal, and the DeScoD-ECG has the potential to be widely used in biomedical applications.

**Index Terms**—ECG signal processing, Baseline wander, diffusion models

## I. INTRODUCTION

CARDIOVASCULAR diseases are the leading cause of sudden cardiac death in the world [1]. Thus, specialized medical services such as diagnostic tools for the study and treatment are in great demand. The electrocardiogram (ECG) is the most straightforward, non-invasive technique and monitoring approach for diagnosing heart diseases. An ECG records the electrical signal from the heart to check for different heart conditions. ECG signals are typically collected during exercise stress tests or in an ambulatory way to increase diagnostic sensitivity. However, the ECG signals acquired in these conditions are strongly influenced by different types of noise, such as baseline wander. The noisy ECG signals are not

H. Li is with the Department of Electrical & Computer Engineering at the University of Arizona, Tucson, AZ 85719 USA. E-mail: hl459@arizona.edu

G. Ditzler is with the Department of Electrical & Computer Engineering at Rowan University, Glassboro, NJ 08028 USA. E-mail: ditzler@rowan.edu

J. Roveda is with the Department of Electrical & Computer Engineering, Department of Biomedical Engineering, and Bio5 Institute at the University of Arizona, Tucson, AZ 85719 USA. E-mail: meilingw@arizona.edu

A. Li is with the Department of Electrical & Computer Engineering and Bio5 Institute at the University of Arizona, Tucson, AZ 85719 USA. E-mail: aoli1@arizona.edu

conducive to both manual and automatic signal analysis, which significantly reduces the usability of the collected signals and increases the cost of clinical experiments and diagnosis.

High-quality and high-fidelity noise removal in the ECG signals is of great significance to the auxiliary diagnosis of heart diseases via ECG. Several solutions have been proposed to restore the ECG signals through a signal processing perspective. Classical digital filter-based methods use finite impulse response (FIR) and infinite impulse response (IIR) filters to remove the noise presented in certain frequencies [2]. Wavelet transforms decompose the ECG signals into different frequency sub-bands and take advantage of the energy of the signals in different scales to isolate baseline wander from ECG signals. In [3], independent component analysis (ICA) is used for removing baseline wander from ECG under the assumptions of different statics between the noise and original ECG signals. Unfortunately, the traditional methods typically can only succeed in situations where the ECG is not corrupted severely and fail when the noise has high amplitude.

Recently, data-driven methods which take advantage of deep neural networks have been used to remove the baseline wander. Work has shown that the neural networks trained on clean and noisy ECG records surpass the performance of the traditional methods by a significant difference. In [4], a novel approach to denoise ECG signals with deep recurrent denoising neural networks (DRNN) is proposed. By modeling the ECG time series with an LSTM module [5], DRNN achieved outstanding results compared to the traditional digital filters. Further, the fully convolutional denoising autoencoder (FCN-DAE) [6] was proposed to denoise the ECG signals via denoising autoencoder with hourglass architecture. The denoising autoencoders [7] first learn architectures to extract the most relevant features automatically and then produce the reconstructions in the presence of noise at the input. Recently, a new deep neural network architecture named DeepFilter [8] was presented and achieved state-of-the-art performance that beat the traditional digital filters and other deep learning-based methods. It is inspired by the designation of the Inception architecture [9] and further proposed a Multi-Kernel Linear And Non-Linear (MKLANL) filter module for denoising the ECG signals with multi-scale features. Also, the generative model has been applied to ECG denoising. An ECG signal denoising method Using Conditional Generative Adversarial Net (CGAN) is proposed in [10]. It trained a neural network with skip connection (well known as UNet [11] architecture) under an adversarial framework for better approximation of

the true data distribution.

Unfortunately, the existing deep learning-based methods perform poorly with extremely high noise corruption, which limits the applicability of these methods for physical exercise stress tests. Thus, there is a significant need in the research community to develop more accurate ECG reconstruction techniques and find more precise approximations of the conditional distribution of the clean ECG records, given the noisy observations. In this work, we propose DeScoD-ECG, a novel probabilistic baseline wander removal method that learns the conditional distribution with conditional score-based diffusion models [12], [13]. We extend the score-based diffusion models into a novel conditional generative model specific for processing ECG signals. Instead of directly denoising/filtering the noisy observations, the DeScoD-ECG iteratively generates signal reconstructions from random Gaussian distributions conditioned on the noisy observations. The experimental results show that the DeScoD-ECG achieves state-of-the-art performance and beats the baseline methods with a large gap. Also, compared with the baseline methods, the DeScoD-ECG is more stable and consistent with different noise levels. As one of the first evaluations of the conditional score-based diffusion model for ECG denoising, our main contributions are: (1) a conditional score-based diffusion model for high-quality and high-fidelity baseline wander removal is proposed. To the best of our knowledge, this is the first study to deploy the score-based diffusion model for ECG signal restoration; (2) state-of-the-art performance was achieved in baseline wander removal on ECG records from the QT Database with baseline wander noise from MIT-BIH Noise Stress Test Database; (3) a self-ensemble strategy was adopted to further improve the reconstruction quality and the algorithm stability under extreme noise conditions.

## II. RELATED WORKS

In this section, we introduce the basic concept of electrocardiogram and the cause of baseline wander. We discussed the significance of baseline wander removal from a clinical point of view. Further, we introduce the Score-based generative models and their advantages compared with other generative models.

### A. Electrocardiogram and Baseline wander

An ECG signal records the electrical signals in the heart. It is a common and painless test used to detect heart problems and quickly monitor heart health [14]. Electrodes placed at certain spots on the chest and legs are used to detect the potential difference in the human body. An ECG sensor records these impulses to show how fast the heart is beating, the rhythm of the heartbeats (steady or irregular), and the strength and timing of the electrical impulses as they move through the different parts of the heart. Changes in an ECG can be a sign of cardiovascular diseases. An exercise stress test helps determine how well the cardiac responds when working its hardest. It typically involves walking on a treadmill or pedaling on a stationary bike while hooked up to an ECG to monitor the

heart activity [15], [16]. Baseline wander usually appears in the ECG signals during exercise stress tests which could be caused by factors such as respiration, variations in electrode impedance, and excessive body movements. Baseline wander significantly affects the accuracy of analysis and diagnosis from the ECG. Thus, high-quality and high-fidelity removal of baseline wander can improve the efficiency of clinical diagnosis and enable more flexible exercise stress tests so patients can move more freely.

### B. Score-based Diffusion model

Score-based diffusion models are a class of deep generative models that generate samples by gradually converting noise into a plausible data sample through denoising. Score-based diffusion models have recently been used for different tasks such as image generation [12], [17] and text-to-audio synthesis [18], [19] with state-of-the-art sample quality that outperform counterparts including generative adversarial nets (GANs) [20] and variational autoencoders (VAEs) [21]. Score-based diffusion models address the drawbacks of GANs and VAEs. More specifically, GANs are known for potentially unstable training and less diversity in generation due to their adversarial training nature. VAEs rely on a surrogate loss and have relatively poor generation quality. Unfortunately, tractability and flexibility are two conflicting objectives in generative modeling. Tractable models can be analytically evaluated and cheaply fit data but are unable to describe the structure in rich datasets aptly. In contrast, flexible models can fit arbitrary structures in data, but evaluating and training these models is more computationally expensive. Diffusion models mitigate the existing drawbacks, which are analytically tractable and flexible. Furthermore, the diffusion models do not need the adversarial training of GAN while having comparable or even better generation qualities than GAN. Therefore, in this work, we introduce the score-based diffusion model to the ECG baseline wander and noise removal tasks and demonstrate the efficiency of this type of generative model.

## III. METHODS

In this section, we first introduce the background of diffusion models and then introduce the proposed DeScoD-ECG for ECG baseline wander and noise removal. We also present the implementation details of the neural network and optimization procedure used in this work.

### A. Score-based Diffusion Model

Diffusion models are generative models, meaning they are used to generate data similar to the data on which they are trained. Fundamentally, diffusion models work by corrupting training data through the successive addition of Gaussian noise and then learning to recover the data by reversing this *diffusion process*. Thus, a well-trained diffusion model is able to generate data from randomly sampled noise through the learned denoising process.

More specifically, given samples from a data distribution  $q(x_0)$ , we are interested in learning a model distribution

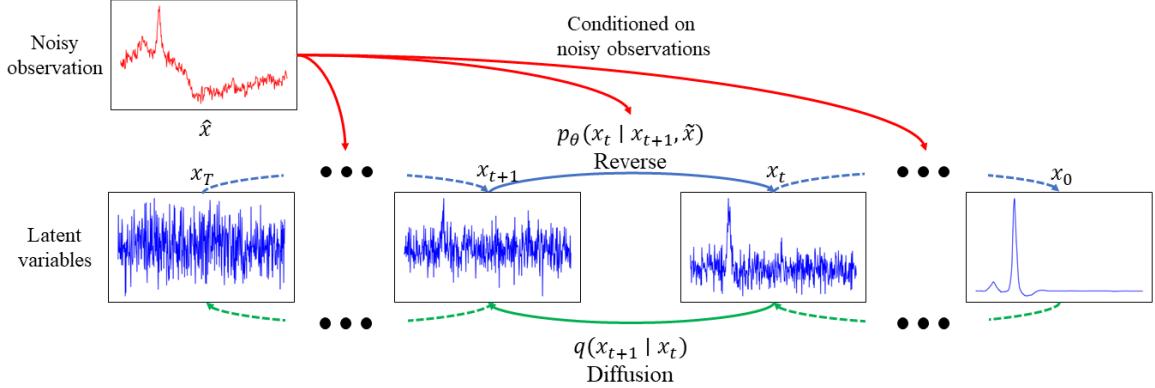


Fig. 1. The *diffusion* ( $q$ ) and *reverse* ( $p_\theta$ ) procedure of DeScoD-ECG. Intuitively, the *diffusion* process gradually corrupts the ground truth ECGs by adding noise, while the *reverse* process aims to learn a denoising function to remove this noise. It is noticeable that *reverse* process is conditioned by the corresponding observed ECGs. By training the denoising function (or the model  $\theta$ ) in this way, we can generate the clean ECGs from random Gaussian noise conditioned on the noisy observations.

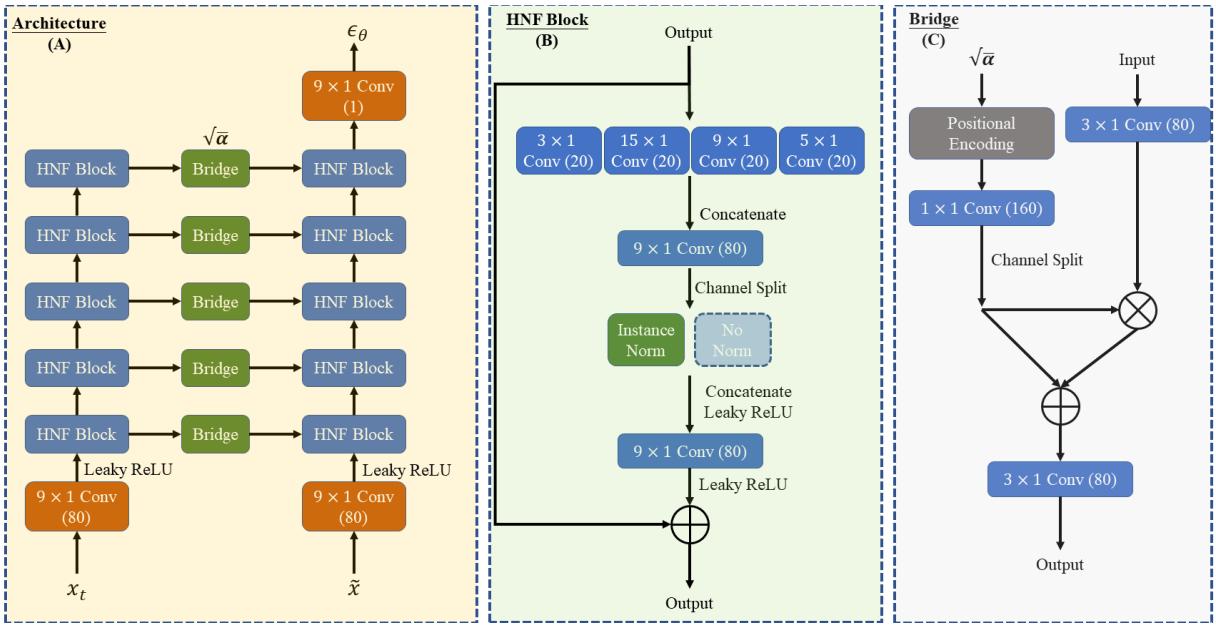


Fig. 2. The architecture of the neural network and each block used in it. We demonstrate the network architecture (A) used in DeScoD-ECG. The whole network is made up by two streams equipped with the HNF block (B) and communicated via the Bridge block (C).

$p_\theta(x_0)$  that approximates  $q(x_0)$ . Denoising diffusion probabilistic models (DDPMs) [12], [22] are latent variable models that consist of a *forward process* or *diffusion process*, in which the data are diffused to a well-behaved distribution by progressively adding noise, and a *reverse process*, in which noise is transformed back into a sample from the target distribution.

The *diffusion process* gradually adds noise to the data to approximate the posterior  $q(x_{1:T} | x_0)$ , where  $x_1, \dots, x_T$  are the latent variables with the same dimensionality as  $x_0$ . In [12], the *diffusion process* is set to a simple parameterization of a fixed Markov Chain with conditional Gaussian translation as

each step:

$$q(x_{1:T} | x_0) := \prod_{t=1}^T q(x_t | x_{t-1}), \quad (1)$$

$$q(x_t | x_{t-1}) := \mathcal{N}(x_t; \sqrt{1 - \beta_t} x_{t-1}, \beta_t \mathbf{I}), \quad (2)$$

where  $\beta_1, \dots, \beta_T$  is a variance schedule (either learned or fixed) which, if well-behaved, ensures that  $x_T$  is nearly an isotropic Gaussian for sufficiently large  $T$ , and  $\mathcal{N}(x; \mu, \Sigma)$  is a Gaussian pdf with parameters  $\mu$  and  $\Sigma$ . The sampling of  $x_t$  at an arbitrary timestep  $t$  has the closed-form of  $q(x_t | x_0) = \mathcal{N}(x_t; \sqrt{\alpha_t} x_0, (1 - \bar{\alpha}_t) \mathbf{I})$ , where  $\alpha_t := 1 - \beta_t$  and  $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$ . Thus,  $x_t$  can be expressed directly as:  $x_t = \sqrt{\alpha_t} x_0 + (1 - \alpha_t) \epsilon$ , where  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ .

---

**Algorithm 1** Training.

**Input:** distribution of clean and noisy ECG pairs  $q(x_o, \tilde{x})$ , the number of training iterations  $N$ , maximal diffusion steps  $T$ , a predefined noise schedule  $S = \{1, \sqrt{\bar{\alpha}_0}, \dots, \sqrt{\bar{\alpha}_T}\}$ ,

**Output:** Trained denoising function  $\epsilon_\theta$

```

1: Initiate:  $i = 0$ 
2: while  $i \leq N$  do
3:    $t \sim \text{Uniform}(\{1, \dots, T\})$ 
4:    $x_o, \tilde{x} \sim q(x_o, \tilde{x})$ 
5:    $\bar{\alpha} \sim \text{Uniform}(S_{t-1}, S_t)$ 
6:    $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
7:   Take gradient descent step on
      
$$\nabla_\theta \left\| \epsilon - \epsilon_\theta \left( \sqrt{\bar{\alpha}} x_0 + \sqrt{1 - \bar{\alpha}} \epsilon, \tilde{x}, \bar{\alpha} \right) \right\|^2$$

8:    $i = i + 1$ 
9: end while

```

---

**Algorithm 2** Inference.

**Input:** Trained denoising function  $\epsilon_\theta$ , a noisy observation  $\tilde{x}$ , maximal diffusion steps  $T$ ,  $M$ -shots reconstruction

**Output:** Reconstructed ECG  $x_0$

```

1: Initiate:  $x_0^m = 0$ ,  $m = M$ ,
2: while  $m > 0$  do
3:    $t = T$ ,  $x_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
4:   while  $t > 0$  do
5:     Let
        
$$x_{t-1} = \frac{1}{\alpha_t} \left( x_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, \tilde{x}, \sqrt{\bar{\alpha}_t}) \right)$$

6:     if  $t > 1$  then
7:        $x_{t-1} = x_{t-1} + \sigma_t$ 
8:     end if
9:      $t = t - 1$ 
10:   end while
11:    $x_0^m = x_0^{m+1} + x_0$ 
12:    $m = m - 1$ 
13: end while
return  $x_0^0/M$ 

```

---

The *reverse process* learns to reverse this *diffusion process* to recover  $x_0$  from  $x_t$ . Starting with the pure Gaussian noise sampled from  $p(x_T) := \mathcal{N}(x_T, \mathbf{0}, \mathbf{I})$ , the *reverse process* is defined by the following Markov chain:

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

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

where the time-dependent parameters of the Gaussian transitions are learned. In the DDPMs setting, the following specific

parameterization of  $p_\theta(x_{t-1} | x_t)$  is proposed:

$$\mu_\theta(x_t, t) = \frac{1}{\alpha_t} \left( x_t - \frac{\beta_t}{\sqrt{1 - \alpha_t}} \epsilon_\theta(x_t, t) \right), \quad (5)$$

$$\sigma_\theta(x_t, t) = \sqrt{\tilde{\beta}_t}, \quad \text{where } \tilde{\beta}_t = \begin{cases} \frac{1 - \alpha_{t-1}}{1 - \alpha_t} \beta_t & t > 1 \\ \beta_1 & t = 1 \end{cases} \quad (6)$$

where  $\epsilon_\theta(\cdot, \cdot)$  is a learnable denoising function which estimates the noise vector  $\epsilon$  that was added to a noisy input  $x_t$ . This parameterization leads to the following alternative loss function:

$$L_{\text{simple}}(\theta) := \mathbb{E}_{t, x_0, \epsilon} \left[ \left\| \epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t} x_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, t) \right\|^2 \right] \quad (7)$$

This objective can be viewed as a weighted combination of denoising score matching which enables more stable training and better results than the original score matching losses [23], [24].

### B. Conditional Diffusion Model for ECG Baseline Wander and Noise Removal

In the task of removing the baseline wander and noise in ECG signals, we consider recovery of the clean signal  $x_0$  given a noisy ECG signal  $\tilde{x}$ . The noisy ECG signal contains baseline wander and noise of unknown type. Thus, we estimate the true conditional data distribution  $q(x_0 | \tilde{x})$  via modeling the conditional distribution  $p_\theta(x_0 | \tilde{x})$  with a diffusion model. Then, we extend the Eq. (3) to a conditional fashion:

$$p_\theta(x_{0:T} | \tilde{x}) := p(x_T) \prod_{t=1}^T p_\theta(x_{t-1} | x_t, \tilde{x}), \quad x_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (8)$$

$$p_\theta(x_{t-1} | x_t, \tilde{x}) := \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t | \tilde{x}), \sigma_\theta(x_t, t | \tilde{x}) \mathbf{I}). \quad (9)$$

The conditional *reverse process*  $p_\theta(x_{t-1} | x_t, \tilde{x})$  iteratively reconstructs the clean ECG  $x_0$  from a Gaussian distribution with the noisy observations as conditions. We visualized the diffusion and reverse mechanism of DeScoD-ECG in Figure 1. The conditional *reverse process*  $p_\theta(x_{t-1} | x_t, \tilde{x})$  gradually denoises the Gaussian noise by a small step conditioned on the noisy observations and then produce the clean signals. Algorithm 1 details the training process of DeScoD-ECG. We first initiate the noise schedule  $\beta$ s for  $T$  steps diffusion in a quadratic form, that the  $\{\beta_1, \dots, \beta_n\}$  are defined as:

$$\beta_1 = 0.0001, \beta_T = 0.5, \quad (10)$$

$$\beta_t = \left( \frac{T-t}{T-1} \sqrt{\beta_1} + \frac{t-1}{T-1} \sqrt{\beta_t} \right)^2. \quad (11)$$

This noise schedule can improve the synthetic quality with fewer diffusion/reverse steps as prior works have shown [13], [25]. Recall that  $\alpha_t := 1 - \beta_t$  and  $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$ , we define a new noise schedule specific for training process that can avoid the denoising function  $\epsilon_\theta$  by directly conditioning on the discrete diffusion step index  $t$  and will improve the model performance. More specifically, we use the predefined noise

schedule  $S = \{1, \sqrt{\bar{\alpha}_0}, \dots, \sqrt{\bar{\alpha}_T}\}$  in the training process then let the denoising function condition on the continuous noise level  $\bar{\alpha} \sim \text{Uniform}(S_{t-1}, S_t)$ . The denoising function  $\epsilon_\theta$  now takes three inputs with conditioned on the noisy observations  $\tilde{x}$ , and the loss function is formally given by:

$$\mathbb{E}_{x_0, \bar{\alpha}, \epsilon} \left[ \left\| \epsilon - \epsilon_\theta \left( \sqrt{\bar{\alpha}} x_0 + \sqrt{1 - \bar{\alpha}} \epsilon, \tilde{x}, \bar{\alpha} \right) \right\|^2 \right] \quad (12)$$

The pseudo code for the inference phase is shown in Algorithm 2. The inference phase in our proposed approach is straightforward to implement, which is simply applying the Eq. (8) to sample the latent variables  $x_t$ . To further improve the reconstruction accuracy, we take advantage of the stochasticity of the diffusion models by averaging multiple runs. This averaging process produces more precise results on the reconstructed signals (we test the 1, 3, 5, and 10-shots averages and report the results in Section IV-D). Recall that the DeScoD-ECG produces reconstruction of the ECG signals from a random Gaussian distribution, therefore, we could assume each reconstruction of multiple individual runs are independent to each others. Let  $y$  be the clean signal and  $y_m$  be the reconstruction at  $m$ th run, we have:

$$y_{\text{comb}} = \frac{1}{M} \sum_{m=1}^M y_m, \quad (13)$$

where  $y_{\text{comb}}$  is the ensemble of multiple runs. And the output of each reconstruction can be written as the true value plus an error in the form:

$$y_m = y + \epsilon_m, \quad (14)$$

and the average error made by the models acting individually (which is the 1-shot reconstruction) is therefore:

$$\epsilon_{\text{avg}} = \frac{1}{M} \sum_{m=1}^M (y - y_m)^2 = \sum_{m=1}^M \frac{1}{M} \epsilon_m^2. \quad (15)$$

Similarly, the expected error from the multi-shot reconstruction is given by:

$$\epsilon_{\text{comb}} = (y - \frac{1}{M} \sum_{m=1}^M y_m)^2 = \left( \sum_{m=1}^M \frac{1}{M} \epsilon_m \right)^2. \quad (16)$$

We can see obviously the multi-shots reconstruction will achieve lower error than single shot reconstruction since:

$$\left( \sum_{m=1}^M \frac{1}{M} \epsilon_m \right)^2 \leq \sum_{m=1}^M \frac{1}{M} \epsilon_m^2. \quad (17)$$

### C. Network Architecture

The architecture of the neural network used in the DeScoD-ECG is shown in Figure 2. The network is comprised of two stream backbones that extract features of the noisy observation  $\tilde{x}$  and the latent variable  $x_t$ , then the Bridges to combine information of the features. The main computational module is named *Half Normalized Filters* (HNF) block, which is motivated by the Multi-Kernel filter design in DeepFilter [8] and the Half-instance-norm in HINet [26]. The input is first fed

TABLE I  
ECG RECORDS USED AS TEST SET. WE USE THE EXACT SAME SELECTIONS AS DEEPFILTER [8] FOR FAIR COMPARISON.

Database Name	Select ID
MIT-BIH Arrhythmia Database	<i>sel123, sel233</i>
MIT-BIH ST Change Database	<i>sel302, sel307</i>
MIT-BIH Supraventricular Arrhythmia Database	<i>sel820, sel853</i>
MIT-BIH Normal Sinus Rhythm Database	<i>sel16420, sel16795</i>
European ST-T Database	<i>sel0106, sel0121</i>
Sudden death patients from BIH	<i>sel32, sel49</i>
MIT-BIH LongTerm ECG Database	<i>sel14046, sel15815</i>

into the multi-scale filters to extract the features under different receptive fields. Then the multi-scale features are aggregated by a convolutional layer after channel-wise concatenation. The half-instance-norm is used to normalize half of the features, which has been shown to improve the training stability and not degrade the natural statistical characters of the features [26]. After the half-instance-norm, the normalized and unnormalized features are fed into another convolutional layer. Then the processed features are connected with the input features with residual shortcut [27] for producing the outputs. The *Bridge* block is modified based on the feature-wise linear modulation (FiLM) [28]. The bridge block directly conditions on the noise level  $\sqrt{\bar{\alpha}}$  that the  $\sqrt{\bar{\alpha}}$  is first encoded by sinusoidal positional embeddings [29] then further encoded by a  $1 \times 1$  convolutional layer that create the learnable scaling and shift vectors. Further, the scaling and shift vectors interact with the input features via multiplications and additions. Notably, there is no downsampling and upsampling in the whole network. Thus, the network is compatible with input signals of any length.

## IV. EXPERIMENTS

In this section, we give the details of the experiments. We first introduce the definitions and formulations of the four similarity metrics that quantify the quality of the reconstructed ECG signals. Then we summarize the datasets and introduce our experimental procedure. Finally, we introduce the baseline comparison methods and discuss the quantitative evaluations of the DeScoD-ECGs and the baselines. We also visualize reconstruction results and how the signals change during the reverse process.

### A. Evaluation Metrics

This work evaluates the distortion of the ECG signals in a quantitative way with four distance-based metrics as suggested by multiple previous studies [30]–[32]. According to those studies, distance-based metrics are considered more suitable than statistics-based metrics (such as Root Mean square error (RMSE), Cross-correlation (CC), and Signal noise ratio

TABLE II

EVALUATION METRICS RESULTS FOR DIFFERENT METHODS FOR RANDOM NOISE AMPLITUDE BETWEEN 0.2 TO 2. COMPARISON RESULTS ARE OBTAINED FROM [8].

Method/Model	SSD (au)	MAD (au)	PRD (%)	Cosine Sim
FIR filter	47.464 $\pm$ 88.496	0.676 $\pm$ 0.572	66.158 $\pm$ 22.104	0.692 $\pm$ 0.216
IIR filter	35.645 $\pm$ 70.189	0.594 $\pm$ 0.532	61.125 $\pm$ 23.220	0.736 $\pm$ 0.208
FCN-DAE	10.633 $\pm$ 14.417	0.598 $\pm$ 0.381	94.504 $\pm$ 70.717	0.770 $\pm$ 0.171
CGAN	7.834 $\pm$ 13.224	0.442 $\pm$ 0.323	57.133 $\pm$ 41.816	0.852 $\pm$ 0.141
DRNN	6.702 $\pm$ 11.316	0.480 $\pm$ 0.340	50.035 $\pm$ 24.802	0.878 $\pm$ 0.126
DeepFilter	5.203 $\pm$ 7.915	0.386 $\pm$ 0.283	50.445 $\pm$ 29.598	0.897 $\pm$ 0.097
Ours (1-shot)	4.912 $\pm$ 7.822	0.369 $\pm$ 0.289	44.218 $\pm$ 27.980	0.904 $\pm$ 0.113
Ours (3-shot)	4.057 $\pm$ 6.202	0.340 $\pm$ 0.266	41.523 $\pm$ 26.169	0.920 $\pm$ 0.095
Ours (5-shot)	3.887 $\pm$ 5.911	0.334 $\pm$ 0.261	40.987 $\pm$ 26.562	0.923 $\pm$ 0.091
Ours (10-shot)	<b>3.771 <math>\pm</math> 5.713</b>	<b>0.329 <math>\pm</math> 0.258</b>	<b>40.527 <math>\pm</math> 26.258</b>	<b>0.926 <math>\pm</math> 0.087</b>

TABLE III  
SSD IN DIFFERENT NOISE SEGMENTS.

Method/Model	0.2 to 0.6	0.6 to 1.0	1.0 to 1.5	1.5 to 2.0
FIR filter	7.162 $\pm$ 9.485	22.542 $\pm$ 30.795	51.381 $\pm$ 72.620	95.983 $\pm$ 134.621
IIR filter	5.615 $\pm$ 7.508	17.262 $\pm$ 25.049	38.334 $\pm$ 58.050	71.980 $\pm$ 107.880
FCN-DAE	7.396 $\pm$ 8.032	8.638 $\pm$ 8.247	10.929 $\pm$ 12.883	14.586 $\pm$ 21.429
DRNN	4.000 $\pm$ 7.437	4.932 $\pm$ 6.652	7.016 $\pm$ 10.971	10.031 $\pm$ 15.665
CGAN	3.461 $\pm$ 6.054	5.518 $\pm$ 7.994	8.579 $\pm$ 12.565	12.382 $\pm$ 18.598
DeepFilter	2.401 $\pm$ 3.419	3.816 $\pm$ 4.771	5.717 $\pm$ 7.781	8.078 $\pm$ 11.104
Ours (1-shot)	2.324 $\pm$ 3.798	3.762 $\pm$ 5.869	5.328 $\pm$ 7.195	7.466 $\pm$ 10.724
Ours (3-shot)	2.036 $\pm$ 3.225	3.252 $\pm$ 5.261	4.421 $\pm$ 5.739	5.935 $\pm$ 8.164
Ours (5-shot)	1.977 $\pm$ 3.201	3.147 $\pm$ 4.963	4.193 $\pm$ 5.474	5.686 $\pm$ 7.772
Ours (10-shot)	<b>1.953 <math>\pm</math> 3.138</b>	<b>3.046 <math>\pm</math> 4.902</b>	<b>4.081 <math>\pm</math> 5.315</b>	<b>5.483 <math>\pm</math> 7.433</b>

TABLE IV  
MAD IN DIFFERENT NOISE SEGMENTS.

Method/Model	0.2 to 0.6	0.6 to 1.0	1.0 to 1.5	1.5 to 2.0
FIR filter	0.280 $\pm$ 0.183	0.510 $\pm$ 0.337	0.759 $\pm$ 0.513	1.045 $\pm$ 0.713
IIR filter	0.248 $\pm$ 0.171	0.449 $\pm$ 0.322	0.666 $\pm$ 0.487	0.918 $\pm$ 0.676
FCN-DAE	0.480 $\pm$ 0.278	0.532 $\pm$ 0.301	0.617 $\pm$ 0.369	0.727 $\pm$ 0.474
DRNN	0.390 $\pm$ 0.286	0.432 $\pm$ 0.273	0.496 $\pm$ 0.330	0.576 $\pm$ 0.409
CGAN	0.297 $\pm$ 0.175	0.370 $\pm$ 0.215	0.468 $\pm$ 0.306	0.587 $\pm$ 0.419
DeepFilter	0.293 $\pm$ 0.214	0.341 $\pm$ 0.222	0.401 $\pm$ 0.270	0.483 $\pm$ 0.351
Ours (1-shot)	0.230 $\pm$ 0.188	0.321 $\pm$ 0.225	0.403 $\pm$ 0.275	0.484 $\pm$ 0.351
Ours (3-shot)	0.216 $\pm$ 0.180	0.299 $\pm$ 0.213	0.369 $\pm$ 0.251	0.442 $\pm$ 0.323
Ours (5-shot)	0.213 $\pm$ 0.178	0.295 $\pm$ 0.211	0.362 $\pm$ 0.247	0.432 $\pm$ 0.314
Ours (10-shot)	<b>0.211 <math>\pm</math> 0.178</b>	<b>0.291 <math>\pm</math> 0.210</b>	<b>0.356 <math>\pm</math> 0.243</b>	<b>0.426 <math>\pm</math> 0.313</b>

TABLE V  
PRD IN DIFFERENT NOISE SEGMENTS.

Method/Model	0.2 to 0.6	0.6 to 1.0	1.0 to 1.5	1.5 to 2.0
FIR filter	42.673 $\pm$ 15.504	61.370 $\pm$ 18.213	73.491 $\pm$ 17.577	81.627 $\pm$ 15.250
IIR filter	38.653 $\pm$ 15.668	55.867 $\pm$ 19.482	67.972 $\pm$ 19.714	76.644 $\pm$ 18.290
FCN-DAE	76.445 $\pm$ 41.696	86.350 $\pm$ 56.163	98.331 $\pm$ 68.010	112.012 $\pm$ 94.656
DRNN	40.298 $\pm$ 19.983	45.582 $\pm$ 21.936	52.317 $\pm$ 24.336	59.312 $\pm$ 27.389
CGAN	39.090 $\pm$ 19.669	48.972 $\pm$ 24.694	60.514 $\pm$ 36.254	74.830 $\pm$ 59.722
DeepFilter	34.026 $\pm$ 14.234	43.346 $\pm$ 19.925	53.777 $\pm$ 27.189	66.167 $\pm$ 38.040
Ours (1-shot)	29.762 $\pm$ 15.583	38.945 $\pm$ 20.183	47.678 $\pm$ 25.366	56.559 $\pm$ 35.969
Ours (3-shot)	28.066 $\pm$ 14.923	36.681 $\pm$ 19.165	44.506 $\pm$ 23.090	53.225 $\pm$ 33.843
Ours (5-shot)	27.697 $\pm$ 14.706	36.187 $\pm$ 18.974	43.774 $\pm$ 22.834	52.715 $\pm$ 35.341
Ours (10-shot)	<b>27.475 <math>\pm</math> 14.631</b>	<b>35.710 <math>\pm</math> 18.718</b>	<b>43.321 <math>\pm</math> 22.959</b>	<b>52.081 <math>\pm</math> 34.681</b>

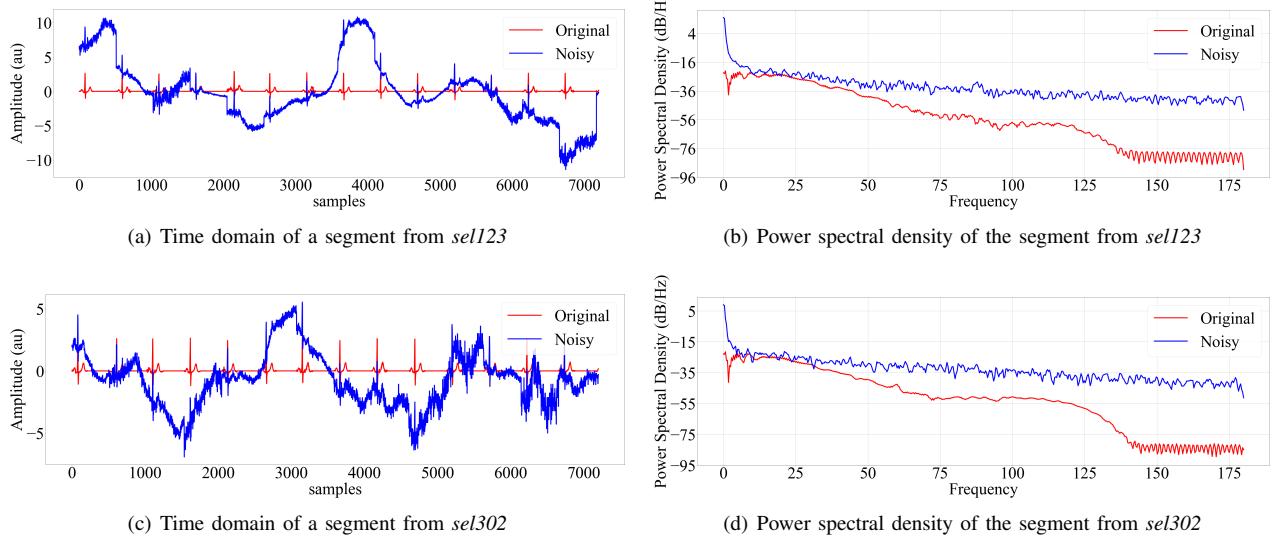


Fig. 3. Visualization of the clean (red lines) and noisy (blue lines) ECG records in the time domain and the corresponding power spectral density plot.

TABLE VI  
COSINE SIMILARITIES IN DIFFERENT NOISE SEGMENTS.

Method/Model	0.2 to 0.6	0.6 to 1.0	1.0 to 1.5	1.5 to 2.0
FIR filter	$0.891 \pm 0.082$	$0.760 \pm 0.149$	$0.637 \pm 0.192$	$0.532 \pm 0.209$
IIR filter	$0.910 \pm 0.074$	$0.800 \pm 0.143$	$0.690 \pm 0.191$	$0.589 \pm 0.217$
FCN-DAE	$0.844 \pm 0.110$	$0.807 \pm 0.136$	$0.757 \pm 0.167$	$0.692 \pm 0.205$
DRNN	$0.926 \pm 0.070$	$0.906 \pm 0.084$	$0.871 \pm 0.122$	$0.824 \pm 0.167$
CGAN	$0.928 \pm 0.060$	$0.891 \pm 0.085$	$0.838 \pm 0.131$	$0.773 \pm 0.184$
DeepFilter	$0.948 \pm 0.042$	$0.921 \pm 0.064$	$0.888 \pm 0.093$	$0.844 \pm 0.124$
Ours (1-shot)	$0.955 \pm 0.049$	$0.929 \pm 0.068$	$0.895 \pm 0.104$	$0.853 \pm 0.156$
Ours (3-shot)	$0.961 \pm 0.042$	$0.939 \pm 0.058$	$0.913 \pm 0.083$	$0.879 \pm 0.132$
Ours (5-shot)	$0.961 \pm 0.041$	$0.941 \pm 0.056$	$0.918 \pm 0.078$	$0.883 \pm 0.129$
Ours (10-shot)	<b><math>0.962 \pm 0.040</math></b>	<b><math>0.943 \pm 0.055</math></b>	<b><math>0.920 \pm 0.077</math></b>	<b><math>0.889 \pm 0.124</math></b>

(SNR)) for calculating the dissimilarity/similarity between the clean and restored signals. For convenience, we use the notations  $x$ ,  $\hat{x}$ , and  $\hat{x}$  to refer to the clean, the noisy, and denoised ECGs, respectively. The definitions and formulations of the four metrics are as follows:

- **Sum of the square of the distances (SSD) [30]:** SSD is a variant of RMSE which is defined as:

$$\text{SSD} = \sum_{n=0}^{N-1} [x(n) - \hat{x}(n)]^2. \quad (18)$$

Unlike RMSE, the SSD of a reconstructed signal provides an idea of how similar the signals are along their entire duration instead of average difference.

- **Absolute maximum distance (MAD) [30]:** MAD is a well-known similarity metric to quantify the ECG quality. Specifically, MAD measures the maximum absolute distance between the original and filtered signals. MAD is given by:

$$\text{MAD} = \max |x(n) - \hat{x}(n)|, \quad \text{for } 0 \leq n \leq N. \quad (19)$$

- **Percentage root-mean-square difference (PRD) [30]:** PRD computes the total distortion present in the denoised

signal. The calculation of PRD is given by:

$$\text{PRD} = \sqrt{\frac{\sum_{n=0}^{N-1} [x(n) - \hat{x}(n)]^2}{\sum_{n=0}^{N-1} [x(n) - \frac{1}{N} \sum_{n=0}^{N-1} x(n)]^2}} \times 100\%. \quad (20)$$

- **Cosine similarity (Cosine Sim) [30]:** The cosine similarity is a measure of similarity between two vectors through the normalized dot product by the Euclidean L2 normalization. The cosine similarity is given by:

$$\text{Cosine Sim} = \frac{\langle x, \hat{x} \rangle}{\|x\| \|\hat{x}\|}. \quad (21)$$

## B. Dataset

We benchmark the DeScoD-ECG on the QT Database [33] and the Massachusetts Institute of Technology – Beth Israel Hospital Noise Stress Test Database (MIT-BIH NST) [34]. The QT Database and MIT-BIH Noise Stress Test Database are archived and available at [www.physionet.org](http://www.physionet.org) [35]. We use the ECG records in the QT database, and corrupt the ECG records with the noise profiles taken from the MIT-BIH NST. QT Database is formed of 105 15-min two-lead ECG recordings containing a wide range of QRS and ST-T

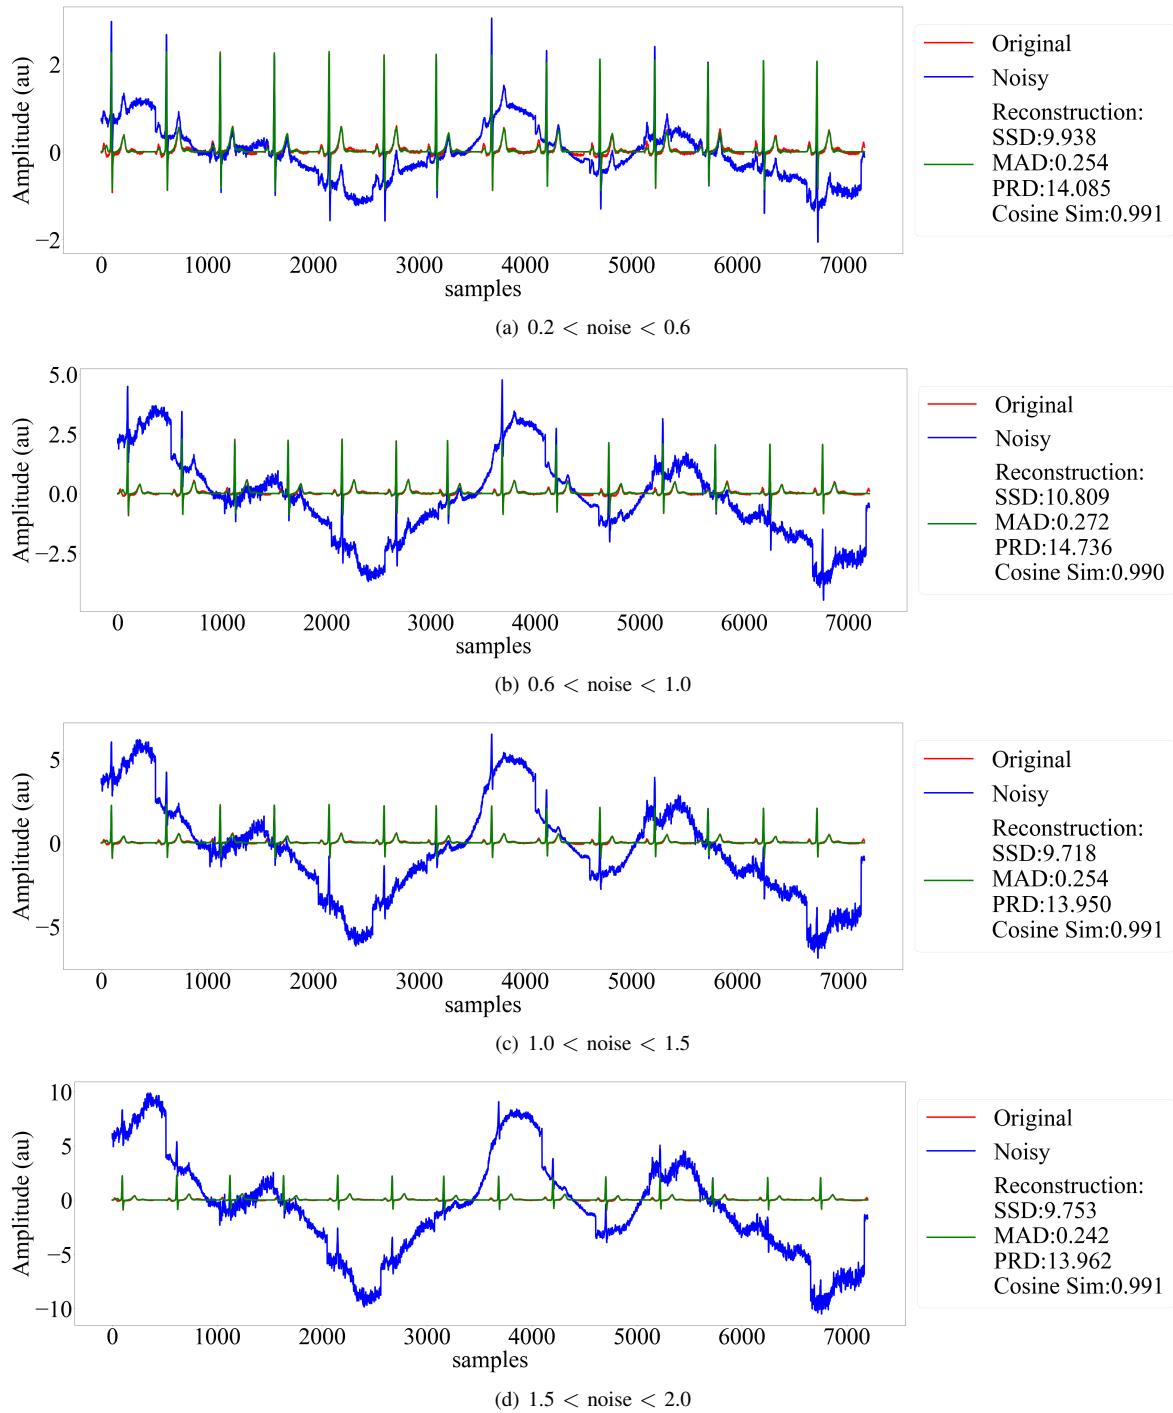


Fig. 4. The reconstructions of a segment intercepted from the *sele123* records under different noise amplitudes.

morphologies in all the possible variability. The MIT-BIH NST contains three noise recordings of typical noise that appears in stress tests which are caused by motion-related interference. The electrodes are placed on the patients' limbs to collect the noises without the presence of ECG signals.

When adding the noise to the ECG records, we first normalized the noise records to the range of min and max value of the corresponding ECG signals and rescaled the noise amplitudes by multiplying a random factor between 0.2 to 2.0. By doing that, we have a better chance to observe the performance of the denoising methods under different noisy conditions. Figure. 3 visualized some examples of the ECG signals with noise in both the time and frequency domains. We can observe that the baseline wander noises used in this work are more than just low-frequency noise that warps the signals. The noises contain both low and high-frequency components that can be used to prove the generalization capability of the denoising algorithms.

### C. Training and implementation details

We use the 14 signals listed in Table I as a test set, which are the same as the experiments in DeepFilter [8] for a fair comparison. Also, selecting test signals from different sources than the training set can ensure the generalization of the algorithms. The remaining ECG records are used for training and validation. We randomly shuffle the remaining records and split them by 70% for training and 30% for validation. For the convenience of training and testing, we split the signals into small patches with 512 points. Therefore, we obtained 50723 samples as training data, 21738 samples as the validation set, and 13337 as the test set. The random seed was set to the one used in the DeepFilter experiments to add noise to the original ECG signal. We ran an experiment to verify this was the case by re-running the FIR and IIR experiments and obtaining the same results.

We implemented the proposed using the PyTorch framework [36], and the source code will be freely available if this work is accepted<sup>1</sup>. During training, a batch size of 96 is used. We trained the model for 400 epochs by Adam optimizer [37] with an initial learning rate of 0.001, and the learning rate decays by multiplying 0.1 every 150 epochs. We save the model weights with the lowest loss on the validation set during training and evaluate the performance only on the validation set for model selection. That means, before we determine the final model for testing, the test set is invisible to the model. We do that because we do not want to adjust the network architecture and the hyper-parameters to overfit the test set.

### D. Results

We compare the performance of the DeScoD-ECG with several representative baseline methods, namely: FIR and IIR filters, DRNN [4], FCN-DAE [6], CGAN [10], and DeepFilter [8]. We present the comparisons in Table II which are shown as the mean values and standard deviations of

TABLE VII  
NUMBERS OF PARAMETERS OF DIFFERENT MODELS.

Method	FCN-DAE	DRNN	CGAN	DeepFilter	Ours
Parameters	0.08 M	0.03 M	4.89 M	0.17 M	1.22 M

the metrics. We can observe that the DeScoD-ECG achieves better performance than the baselines even with 1-shot reconstruction. With 1-shot reconstruction, the DeScoD-ECG outperforms the comparative methods in all the metrics. That indicates the score-based diffusion model achieved better results than naïvely end-to-end training neural networks with better distribution approximation capability. When we use multi-shot reconstruction, the performance is further improved. For example, the 10-shot reconstruction obtained the best results on all four similarity metrics with  $3.771 \pm 5.713$  au,  $0.329 \pm 0.258$  au,  $40.527 \pm 26.258$  %, and  $0.926 \pm 0.087$  compared to the existing state-of-the-art method, DeepFilter, which achieved  $5.203 \pm 7.915$  au,  $0.386 \pm 0.283$  au,  $50.455 \pm 29.598$  %, and  $0.897 \pm 0.097$ . The quantitative evaluations on multi-shot reconstruction support our hypothesis that the DeScoD-ECG can produce more accurate results with a self-ensemble strategy, which averages the multiple predictions.

Further, we evaluate the model performance on the test set under different noise amplitudes. Table III to VI present the four metrics of different noise segments. The results on each noise segment and different shots are consistent with the overall performance in Table II. We notice several observations that could further prove the efficiency of the DeScoD-ECGs. First, the DeScoD-ECG with 1-shot reconstruction still outperforms the baselines and achieves better results than the existing methods in 16 out of 16 evaluations under different noise segments. Second, the gain brought by multi-shots reconstruction still holds under different noise strengths. Third, the DeScoD-ECG boosted by the multi-shots reconstruction is more stable under extreme noisy scenarios than the baselines. For example, the results of 10-shot reconstruction under noise amplitude between 1.5 and 2.0 outperforms the results of DeepFilter under noise amplitude between 1.0 and 1.5. Figure 4 shows the reconstruction results of an ECG beat extracted from the signal *sel123* under different noise amplitude. We observed that the DeScoD-ECG could produce precise and consistent reconstruction under different noise levels and keep the morphology of the signal with a significant baseline wander.

Another interesting observation is how the DeScoD-ECG reconstructs the signal starting with random Gaussian noise. A sequence of results sampled from the reverse process is visualized in Figure 5. We can see that the diffusion model denoises the Gaussian noise iteratively by a small step. After 50 iterations, the Gaussian noise is denoised and transformed into a clean ECG signal.

## V. DISCUSSION AND FUTURE WORKS

This work proposed DeScoD-ECG, a deep score-based diffusion model for ECG baseline wander and noise removal.

<sup>1</sup><https://github.com/HuayuLiArizona/Score-based-ECG-Denoising.git>

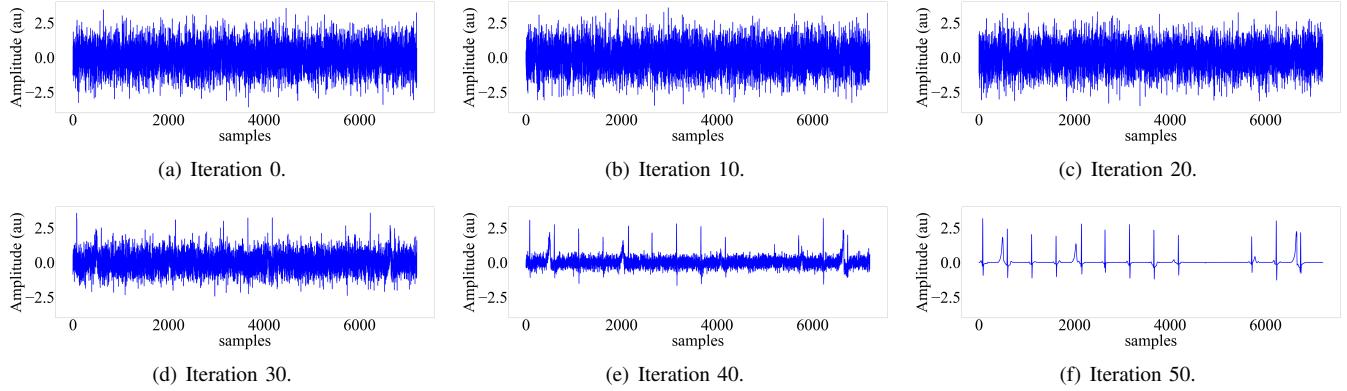


Fig. 5. Visualization of the inference process which the diffusion model denoise the Gaussian noise iteratively to the approximation of true data.

Experiments conducted on the QT database and MIT-BIH NST demonstrate the facts in two aspects. First, we show that the DeScoD-ECG achieved the best performance compared to the state-of-the-art denoising solutions. Taking advantage of the excellence of score-based diffusion models in approximating the true data distribution, DeScoD-ECG is proven able to beat the baseline methods with high accuracy and the preservation of fine details. The results reported in Table II show that DeScoD-ECG has the best overall denoising performance in four distance-based metrics. The results reported in Table III to VI show the denoising performance under different noise levels is consistent with overall performance and also surpasses the baseline methods. Second, we demonstrate that the self-ensembles strategies can further boost the precision of the reconstructions. The DeScoD-ECG also shows upward trends in the general performance and the performance by noise level segments for the four metrics when increasing ensemble size. In summary, the first fact proves the outstanding performance of the score-based diffusion model, and the second shows the existence of strategies for further improving the score-based diffusion model.

The DeScoD-ECG is of great significance for healthcare research. The motion-related noises, including electrode motion artifacts and baseline wander, exist widely in wearable ECG monitoring devices [38], [39]. Therefore, the noises used in this work are rational to prove the feasibility of applying the DeScoD-ECG in filtering the ECG collected in exercise stress tests. The excellent performance of the DeScoD-ECG under extreme noisy environments enables the exercise stress tests to take place in more flexible scenarios. This work can be regarded as a preliminary study and a critical data processing tool for our future analysis of exercise stress tests beyond the traditional test items, such as walking on a treadmill or riding a stationary bike. The routine examination items limit the patients' body movement, activity space, and exercise intensity to prevent the influence of the noises. However, the applications of stable and accurate reconstruction of the noisy ECG allow us conducted stress tests in more flexible ways. For instance, the impacts of COVID-19 on elite athletes are studied in [40]. Athletes risk persistent cardiac or incapacitating

tating symptoms that stop them from resuming competitive sports. ECG monitoring in the exercise stress test is critical for the cardiovascular evaluation of athletes. In particular, these stress tests are vital to evaluate whether an athlete's health condition allows them to return to competition. However, their hidden cardiovascular dangers may only be exposed when they go through their own routine training [41]. Therefore, we need ECG denoising techniques that can perform well in extreme noise conditions to fulfill the flexible exercise stress test requirements.

In addition to the performance achieved by DeScoD-ECG, we find some interesting future work directions. One direction is to improve the computation efficiency. We evaluated the running speed of DeScoD-ECG on a regular configuration laptop with i7-10750H 2.60GHz CPU and GTX 1660 Ti GPU. We chose this configuration due to the consideration of verifying DeScoD-ECG could be run by affordable personal computers for researchers who want to use DeScoD-ECG for ECG denoising. Also, we provided the numbers of parameter of DeScoD-ECG and the baseline methods in Tabel VII. We test the running time of the DeScoD-ECG on a 20-second ECG record (7200 points with a sample rate of 360 Hz) that the 1, 3, 5, and 10-shot reconstructions take 0.682, 1.970, 3.284, and 7.136 second. Although the running time is shorter than the record duration, and the running speed can be further improved by batch computation, we should think about a fast version of DeScoD-ECG in our future works. The major computational overhead of diffusion models is brought by the sampling/inference process. Thus, speeding up the sampling efficiency of the diffusion model is of great significance in improving DeScoD-ECG. Several recent studies [17], [18] were conducted to accelerate the reversing process of diffusion models by leveraging ODE solvers. Thus, our next step will focus on combining the DeScoD-ECG with ODE solvers to mitigate the computational overhead. Another direction is to enhance the self-ensemble strategy. The DeScoD-ECG used a simple averaging over the multi-shots reconstruction, while we believe that developing a weighted ensemble strategy will further improve the quality of the restored signals. Therefore, in our future works, we will explore the possibility of using

quality assignment techniques [42] to give each reconstruction a weight factor to achieve the weighted self-ensemble strategy. Also, we will evaluate the DeScoD-ECG on more types of noise and deploy the denoised results to more downstream tasks.

## VI. CONCLUSION

In this paper, we presented DeScoD-ECG, a novel ECG baseline wander and noise removal approach based on a conditional score-based diffusion model, which has achieved state-of-the-art performance compared with existing baseline methods. Unlike the traditional end-to-end training frameworks of deep neural networks, which directly learn the mapping from noisy observations to noise-free reconstructions, the proposed DeScoD-ECG learns the conditional distribution with stochasticity. The DeScoD-ECG starts from Gaussian white noise and iteratively reconstructs the signal through a fixed Markov Chain with Gaussian translation conditioned on the previous step reconstructions and the noisy ECG observations. The experimental results demonstrate the potential of the DeScoD-ECG for biomedical signal processing. Compared to the naively training deep neural networks, the diffusion models have better approximations of the true data distribution and higher stability under extreme noise corruptions. The high-quality and high-fidelity ECG reconstruction methodologies could serve the cardiac activity monitoring in various clinical settings, thereby increasing the flexibility of Cardiovascular disease diagnosis.

## ACKNOWLEDGEMENT

This work was supported by grants from the National Heart, Lung, and Blood Institute (#R21HL159661), the National Science Foundation (#2052528 and CAREER #1943552), and the Department of Energy (#DE-NA0003946). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsors.

## REFERENCES

- [1] W. H. Organization *et al.*, *Global status report on noncommunicable diseases 2014*. World Health Organization, 2014, no. WHO/NMH/NVI/15.1.
- [2] K. S. Kumar, B. Yazdanpanah, and P. R. Kumar, “Removal of noise from electrocardiograms using digital fir and iir filters with various methods,” in *2015 International conference on communications and signal processing (ICCP)*. IEEE, 2015, pp. 0157–0162.
- [3] Z. Barati and A. Ayatollahi, “Baseline wandering removal by using independent component analysis to single-channel ecg data,” in *2006 International Conference on Biomedical and Pharmaceutical Engineering*. IEEE, 2006, pp. 152–156.
- [4] K. Antczak, “Deep recurrent neural networks for ecg signal denoising,” *arXiv preprint arXiv:1807.11551*, 2018.
- [5] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [6] H.-T. Chiang, Y.-Y. Hsieh, S.-W. Fu, K.-H. Hung, Y. Tsao, and S.-Y. Chien, “Noise reduction in ecg signals using fully convolutional denoising autoencoders,” *Ieee Access*, vol. 7, pp. 60 806–60 813, 2019.
- [7] P. Vincent, H. Larochelle, Y. Bengio, and P.-A. Manzagol, “Extracting and composing robust features with denoising autoencoders,” in *Proceedings of the 25th international conference on Machine learning*, 2008, pp. 1096–1103.
- [8] F. P. Romero, D. C. Piñol, and C. R. Vázquez-Seisdedos, “Deepfilter: an ecg baseline wander removal filter using deep learning techniques,” *Biomedical Signal Processing and Control*, vol. 70, p. 102992, 2021.
- [9] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, “Going deeper with convolutions,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 1–9.
- [10] X. Wang, B. Chen, M. Zeng, Y. Wang, H. Liu, R. Liu, L. Tian, and X. Lu, “An ecg signal denoising method using conditional generative adversarial net,” *IEEE Journal of Biomedical and Health Informatics*, 2022.
- [11] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [12] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 6840–6851, 2020.
- [13] J. Song, C. Meng, and S. Ermon, “Denoising diffusion implicit models,” *arXiv preprint arXiv:2010.02502*, 2020.
- [14] P. J. Zimetbaum and M. E. Josephson, “Use of the electrocardiogram in acute myocardial infarction,” *New England Journal of Medicine*, vol. 348, no. 10, pp. 933–940, 2003.
- [15] T. Y. Goraya, S. J. Jacobsen, P. A. Pellikka, T. D. Miller, A. Khan, S. A. Weston, B. J. Gersh, and V. L. Roger, “Prognostic value of treadmill exercise testing in elderly persons,” *Annals of internal medicine*, vol. 132, no. 11, pp. 862–870, 2000.
- [16] R. Gosselink, T. Troosters, and M. Decramer, “Exercise testing: why, which and how to interpret,” *Breathe*, vol. 1, no. 2, pp. 120–129, 2004.
- [17] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, “Score-based generative modeling through stochastic differential equations,” *arXiv preprint arXiv:2011.13456*, 2020.
- [18] Z. Kong, W. Ping, J. Huang, K. Zhao, and B. Catanzaro, “Diffwave: A versatile diffusion model for audio synthesis,” *arXiv preprint arXiv:2009.09761*, 2020.
- [19] N. Chen, Y. Zhang, H. Zen, R. J. Weiss, M. Norouzi, and W. Chan, “Wavegrad: Estimating gradients for waveform generation,” *arXiv preprint arXiv:2009.00713*, 2020.
- [20] 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, 2014.
- [21] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” *arXiv preprint arXiv:1312.6114*, 2013.
- [22] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, “Deep unsupervised learning using nonequilibrium thermodynamics,” in *International Conference on Machine Learning*. PMLR, 2015, pp. 2256–2265.
- [23] Y. Song and S. Ermon, “Generative modeling by estimating gradients of the data distribution,” *Advances in Neural Information Processing Systems*, vol. 33, 2019.
- [24] P. Vincent, “A connection between score matching and denoising autoencoders,” *Neural computation*, vol. 23, no. 7, pp. 1661–1674, 2011.
- [25] A. Q. Nichol and P. Dhariwal, “Improved denoising diffusion probabilistic models,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 8162–8171.
- [26] L. Chen, X. Lu, J. Zhang, X. Chu, and C. Chen, “Hinet: Half instance normalization network for image restoration,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 182–192.
- [27] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [28] V. Dumoulin, E. Perez, N. Schucher, F. Strub, H. d. Vries, A. Courville, and Y. Bengio, “Feature-wise transformations,” *Distill*, vol. 3, no. 7, p. e11, 2018.
- [29] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [30] R. Nygaard, G. Melnikov, and A. K. Katsaggelos, “A rate distortion optimal ecg coding algorithm,” *IEEE Transactions on biomedical engineering*, vol. 48, no. 1, pp. 28–40, 2001.
- [31] M. S. Manikandan and S. Dandapat, “Ecg distortion measures and their effectiveness,” in *2008 First International Conference on Emerging Trends in Engineering and Technology*. IEEE, 2008, pp. 705–710.

- [32] F. P. Romero, L. V. Romaguera, C. R. Vázquez-Seisdedos, M. G. F. Costa, J. E. Neto *et al.*, “Baseline wander removal methods for ecg signals: A comparative study,” *arXiv preprint arXiv:1807.11359*, 2018.
- [33] P. Laguna, R. G. Mark, A. Goldberg, and G. B. Moody, “A database for evaluation of algorithms for measurement of qt and other waveform intervals in the ecg,” in *Computers in cardiology 1997*. IEEE, 1997, pp. 673–676.
- [34] G. Moody, W. Muldrow, and R. Mark, “A noise stress test for arrhythmia detectors,” in *Computers in Cardiology*, 1984, pp. 381–384.
- [35] P. PhysioBank, “Physionet: components of a new research resource for complex physiologic signals,” *Circulation*, vol. 101, no. 23, pp. e215–e220, 2000.
- [36] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, “Pytorch: An imperative style, high-performance deep learning library,” *Advances in neural information processing systems*, vol. 32, 2019.
- [37] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [38] J. G. Webster, *Medical instrumentation: application and design*. John Wiley & Sons, 2009.
- [39] X. An and G. K. Stylios, “Comparison of motion artefact reduction methods and the implementation of adaptive motion artefact reduction in wearable electrocardiogram monitoring,” *Sensors*, vol. 20, no. 5, p. 1468, 2020.
- [40] T. J. Yeo, “Sport and exercise during and beyond the covid-19 pandemic,” pp. 1239–1241, 2020.
- [41] S. Sridi-Cheniti, S. Benhenda, S. Doutreleau, S. Cade, S. Guerard10, J.-M. Guy11, P. Trimoulet12, S. Picard13, B. Dusfour, A. Pouzet14 *et al.*, “Resuming training in high-level athletes after mild covid-19 infection: A multicenter prospective study (asccovid-19),” 2022.
- [42] U. Satija, B. Ramkumar, and M. S. Manikandan, “A review of signal processing techniques for electrocardiogram signal quality assessment,” *IEEE reviews in biomedical engineering*, vol. 11, pp. 36–52, 2018.