PDE-Driven Spatiotemporal Generative Modeling for Multilead ECG Synthesis

Yakir Yehuda, Kira Radinsky

Technion–Israel Institute of Technology
y.yakir@cs.technion.ac.il, kirar@cs.technion.ac.il

Abstract

Synthesizing realistic 12-lead electrocardiogram (ECG) data is a complex task due to the intricate spatial and temporal dynamics of cardiac electrophysiology. Traditional generative models often struggle to capture the nuanced interdependencies among ECG leads, which are essential for accurate medical analysis. In this paper, we propose Physics-Inspired Partial Differential Equation GAN for Multilead ECG Synthesis **(PhysioPDE-GAN)**, a generative framework designed to model the spatiotemporal structure of multilead ECG signals by incorporating physiological priors and spatial constraints directly into the generative process. By embedding PDE-based representations directly into the generative process, our approach effectively captures both the temporal evolution and spatial relationships between ECG leads. We conduct extensive experiments to evaluate the performance of various base classifiers trained on the synthetic 12-lead ECG data generated by PhysioPDE-GAN. These classifiers outperform those trained on data produced by other conventional methods, achieving statistically significant improvements in detecting cardiac abnormalities. Our work highlights the potential of combining PDE-driven cardiac models with advanced generative techniques to enhance the quality and utility of synthetic biomedical datasets.

1 Introduction

Electrocardiograms (ECGs) are indispensable in diagnosing heart conditions by recording electrical activity across multiple leads. Among these, the 12-lead ECG is particularly valuable, as it captures comprehensive spatial information on cardiac function from multiple angles. However, developing robust machine learning models for 12-lead ECG analysis requires access to large and diverse datasets. Privacy concerns, data scarcity (particularly for rare cardiac conditions), and logistical constraints make it difficult to assemble such extensive databases (Voigt and Bussche 2017). This limitation slows research advances and complicates the deployment of reliable diagnostic tools in clinical practice. To address this, researchers have turned to synthetic data generation as a promising solution (de Melo et al. 2022; Giuffrè and Shung 2023).

While generative models such as Generative Adversarial Networks (GAN) (Goodfellow et al. 2014) have achieved success in producing synthetic data, they often fail to capture the complex spatiotemporal dynamics of 12-lead ECG signals. Specifically, they neglect the intricate interdependencies and physical relationships between the 12 leads, which are vital for a synthetic signal to be physiologically plausible. Simultaneously, physics-informed approaches have shown that integrating domain knowledge such as partial differential equations (PDEs) into neural networks can significantly improve predictive accuracy (Raissi, Perdikaris, and Karniadakis 2019; Karniadakis et al. 2021). By encoding known physical laws, Physics-Informed Neural Networks (PINNs) ensure that the learned solutions respect fundamental constraints of the underlying domain. Nonetheless, these approaches typically rely on having a fully known or explicitly defined physical model. In many real-world applications, the exact physical model may be unknown or only partially specified.

In this work, we introduce a novel generative framework that addresses these limitations by embedding differential operators within the model to approximate the unknown nonlinear dynamics underlying 12-lead ECG signals. Instead of relying on a pre-defined physical model, we enable the generator to learn PDE-inspired constraints from the data itself. This allows the model to synthesize physiologically consistent ECG signals by capturing both the temporal evolution and spatial dependencies across leads.

To this end, we introduce PhysioPDE-GAN, a generative framework that integrates PDE-inspired constraints into a GAN architecture to synthesize realistic 12-lead ECG signals. This unique integration allows PhysioPDE-GAN to synthesize physiologically consistent ECG signals that capture both the temporal evolution (how a signal changes over time) and the spatial dependencies (how the leads relate to each other) across all 12 leads. Unlike prior methods that often focus on ordinary differential equations (ODEs), which capture only temporal dependencies, our PDE-based formulation explicitly models spatial interactions across the 12 leads. By learning from both observed ECG data and an enforced PDE constraint, PhysioPDE-GAN generates physiologically consistent signals that more accurately reflect real-world cardiac behavior.

The main contributions of this study are threefold. (1) We introduce a novel PhysioPDE-GAN framework that models the spatiotemporal dynamics of 12-lead ECG by learning

discrete derivative operators and nonlinear response functions. This framework, which embeds physiological constraints into the generative process, produces high-fidelity ECG waveforms. (2) Through extensive experiments on large-scale datasets, we show that classifiers trained on data generated by our framework achieve superior accuracy in detecting cardiac abnormalities compared to those trained on other synthetic datasets. This demonstrates the significant clinical value of integrating physics-informed constraints into data generation. (3) We bridge physiological modeling and machine learning by integrating domain knowledge to improve cardiac signal representation, and release our implementation to advance research in medical AI.

2 Related Work

The efficacy of deep learning models, particularly in medical applications, is heavily reliant on the availability of large, high-quality annotated datasets. However, data scarcity, privacy concerns, and stringent ethical constraints prevalent in the medical domain often hinder the collection of such extensive datasets (de Melo et al. 2022). To address these limitations, synthetic data generation has emerged as a critical solution to these challenges, providing a way to augment high-quality, diverse datasets essential for training robust deep neural networks (Cichy and Kaiser 2019).

Electrocardiogram (ECG) data are fundamental for diagnosing cardiovascular conditions (Yehuda, Freedman, and Radinsky 2023). Generative models, especially those based on Generative Adversarial Networks (GANs) (Goodfellow et al. 2014; Donahue, McAuley, and Puckette 2018), have been widely explored as a means to synthesize ECG data and enrich training sets. Early works (Zhu et al. 2019; Golany, Radinsky, and Freedman 2020; Golany, Freedman, and Radinsky 2021) demonstrated the utility of GANs for single-lead ECG generation. ODE-GAN (Golany, Freedman, and Radinsky 2021) introduced the use of ordinary differential equations (ODEs) to model the temporal dynamics of single-lead ECG. While this approach improved the physiological plausibility of the generated signals, it was inherently limited to a single lead and could not capture the crucial spatial relationships and inter-dependencies that exist across multiple leads.

Several subsequent works address some aspects of multi-lead or multi-view ECG generation. For instance, vector quantized VAEs (VQ-VAE) have been proposed for enhancing ECG classifiers (Liu et al. 2020), while 3KG (Gopal et al. 2021) uses 3D augmentations in the vector-cardiogram space. ME-GAN (Chen et al. 2022) extends GANs to multi-view ECG synthesis and (Huang, Wang, and Li 2023) adds unsupervised noise generation for ECGs. SSSD-ECG (Alcaraz and Strodthoff 2023) adopts diffusion models for 12-lead ECG synthesis, and MultiODE-GAN (Yehuda and Radinsky 2024) couples ODE frameworks for simulating multi-lead signals building on the works of (Golany, Radinsky, and Freedman 2020) and (McSharry et al. 2003). Although these methods improve realism, they still rely predominantly on ODE-like formulations or specialized architectures, which may not fully capture the spatial interplay across leads.

A parallel line of research, Physics-Informed Neural Networks (PINNs) (Raissi, Perdikaris, and Karniadakis 2019; Karniadakis et al. 2021), has demonstrated the power of integrating known physical laws into deep learning models. PINNs typically embed explicit differential equations as regularization constraints to solve forward or inverse problems. However, this approach is fundamentally limited by its reliance on a pre-existing, well-defined physical model—a prerequisite that is often unmet in complex biological systems like the heart.

Our work draws inspiration from a different, more data-driven branch of this field: data-driven PDE discovery. Frameworks like PDE-Net (Long et al. 2018; Long, Lu, and Dong 2019) and learned discretization methods (Bar-Sinai et al. 2019) have shown that it is possible to approximate unknown governing dynamics from observational data without a pre-specified PDE. These models learn discrete operators that mimic differential equations, enabling them to discover the underlying physical laws directly from the data.

Unlike prior work, our approach is not aimed at solving a known PDE or solely improving the realism of generated signals through specialized architectures. We introduce PhysioPDE-GAN, a novel generative framework that integrates Partial Differential Equation (PDE)-inspired constraints directly into a GAN's generator. Crucially, we do not assume a known physical model. Instead, our framework learns a physics-informed representation of the spatiotemporal dynamics of 12-lead ECGs by discovering differential operators from the data.

Applied to 12-lead ECG synthesis, this formulation enables the generator to capture both temporal evolution and inter-lead dependencies, overcoming limitations of ODE-based models like ODE-GAN (Golany, Freedman, and Radinsky 2021), which model only temporal dynamics. By embedding PDE structure into a generative process for multichannel biomedical time series, PhysioPDE-GAN produces more physiologically coherent signals and enhances downstream heart-disease classification, demonstrating the power of a data-driven physics-informed approach to generative modeling.

3 Method

In this section, we introduce PhysioPDE-GAN, a novel generative framework for synthesizing realistic 12-lead ECG heartbeats. Our approach integrates a system of learnable, coupled partial differential equations (PDEs) directly within a Generative Adversarial Network (GAN). Specifically, PhysioPDE-GAN learns to model the distinct electrophysiological dynamics of the heart's primary conduction components: the Sinoatrial (SA) node, the Atrioventricular (AV) node, and the His-Purkinje (HP) system. By learning a separate reaction-diffusion process for each component, our generator can produce ECG signals with high physiological fidelity.

Inspired by data-driven PDE discovery (Long et al. 2018; Long, Lu, and Dong 2019), our method learns differential operators directly from data. This encourages the synthesis of ECGs with physiologically coherent spatiotemporal

dynamics, moving beyond purely data-driven approaches to incorporate fundamental biophysical constraints.

3.1 Physiologically-Inspired PDE Formulation

The heart's electrical activity is inherently a spatiotemporal process. We model the 12-lead ECG signal $\mathbf{u}(t) \in \mathbb{R}^{12}$ as a spatially discrete field evolving over time. Classical bidomain and monodomain models (Potse et al. 2006; Bishop and Plank 2011) solve 3D reaction-diffusion PDEs before projecting to electrode positions. While the 12 leads do not form a regular spatial grid, they represent different projections of the same underlying 3D cardiac vector field. Instead of explicitly simulating this complex 3D field, our model learns a system of coupled 1D reaction-diffusion equations over the lead index $x \in \{1, 2, \ldots, 12\}$ where ordered leads represent different projections of the heart's 3D electrical vector. This *lead lattice* serves as an effective abstraction that captures essential inter-lead dependencies and temporal dynamics.

The dynamics are governed by a general PDE of the form:

$$\frac{\partial \mathbf{u}}{\partial t} = F(\mathbf{u}(t, \mathbf{x}), D_x\mathbf{u}, D_{xx}\mathbf{u}, \ldots), \quad \mathbf{x} \in \Omega, \quad t \in [0, T]$$

where $\mathbf{u}$ is the signal vector across all leads, F is a nonlinear function parameterized by a neural network, and D_x, D_{xx} are operators that approximate spatial derivatives across the lead dimension (see Fig. 1).

A Coupled System for Cardiac Conduction To capture the distinct physiological roles of different cardiac tissues, we model the ECG signal as a superposition of three distinct, interacting fields: the SA, AV, and His-Purkinje systems (Potse 2018; Quiroz-Juárez et al. 2019). Each field is governed by its own reaction-diffusion PDE, and they are coupled in a physiologically causal manner.

Let $u_{\text{SA}}(t, x)$, $u_{\text{AV}}(t, x)$, and $u_{\text{HP}}(t, x)$ represent the electrical potentials generated by the SA, AV, and His-Purkinje systems, respectively. Their dynamics are described by the following system of coupled PDEs:

$$\frac{\partial u_{\text{SA}}}{\partial t} = F_{\text{SA}}\big(u_{\text{SA}}, D_x u_{\text{SA}}, D_{xx} u_{\text{SA}}, \ldots\big) \tag{1}$$

$$\frac{\partial u_{\text{AV}}}{\partial t} = F_{\text{AV}}\big(u_{\text{AV}}, D_x u_{\text{AV}}, D_{xx} u_{\text{AV}}, \ldots ; S(u_{\text{SA}})\big) \tag{2}$$

$$\frac{\partial u_{\text{HP}}}{\partial t} = F_{\text{HP}}\big(u_{\text{HP}}, D_x u_{\text{HP}}, D_{xx} u_{\text{HP}}, \ldots ; S(u_{\text{AV}})\big) \tag{3}$$

where:

- D_x and D_{xx} denote discrete spatial derivative operators (approximating $\frac{\partial}{\partial x}$ and $\frac{\partial^2}{\partial x^2}$ respectively) with respect to the lead index x, see (Sec. 3.3).
- F_{SA}, F_{AV}, and F_{HP} are distinct nonlinear response functions, implemented as neural networks, that learn the unique reaction-diffusion dynamics.
- $S(\cdot)$ is a stimulus function (MLP) that models the causal coupling between systems. The AV node's activity is triggered by the SA node, and the HP system is triggered by the AV node.

The final synthesized 12-lead ECG signal, $\mathbf{u}(t, x)$, is the linear superposition of the three components:

$$\mathbf{u}(t, x) = u_{\text{SA}}(t, x) + u_{\text{AV}}(t, x) + u_{\text{HP}}(t, x) \tag{4}$$

3.2 Numerical Solution

The generator G acts as a numerical solver for the coupled PDE system described in Equations (1)–(3). It simulates the system's evolution over K discrete time steps to generate a full heartbeat sequence.

To numerically solve this system of PDEs, we use a forward-Euler-like scheme. This scheme is implemented as a stack of neural network blocks, where each block advances the solution one time step, Δt.

Initial Conditions The generative process is initiated from a random noise vector $\mathbf{z} \sim \mathcal{N}(0, I)$. This vector is mapped by a fully connected (FC) layer to define the initial state of the SA node, the heart's primary pacemaker:

$$\mathbf{u}_{\text{SA}}(0) = \text{FC}(\mathbf{z})$$

The AV and HP systems are initialized to a zero resting potential, as they are quiescent until activated:

$$\mathbf{u}_{\text{AV}}(0) = \mathbf{0}, \qquad \mathbf{u}_{\text{HP}}(0) = \mathbf{0}$$

Forward via Euler Steps We employ a forward Euler scheme to integrate the PDEs over time. The core of the generator is a PDE block that executes a single integration step of size Δt. For each subsystem $k \in \{\text{SA, AV, HP}\}$, the update rule is:

$$\mathbf{u}_k(t + \Delta t) = \mathbf{u}_k(t) + \Delta t \cdot F_k\big(\mathbf{u}_k(t),\, D_x\mathbf{u}_k(t),\, D_{xx}\mathbf{u}_k(t);\, S(\mathbf{u}_{\text{parent}}(t))\big) \tag{5}$$

where $\mathbf{u}_{\text{parent}}$ denotes the stimulating signal (e.g., $\mathbf{u}_{\text{SA}}$ for the AV system). The nonlinear function F_k is implemented as a small MLP (e.g., two hidden layers with ReLU activations).

The generator applies this update rule iteratively for K steps to produce the three component sequences $\{\mathbf{u}_k(t)\}_{t=0}^{T}$. These are then summed according to Equation (4) to produce the final 12-lead ECG $\mathbf{u}$. The learnable parameters of the generator, θ_G, include the weights of the FC layer, the MLPs (F_k), and the coefficients of the derivative filters (D_x, D_{xx}).

3.3 Learnable Finite-Difference Operators

We implement the spatial derivative operators, D_x and D_{xx}, as 1D convolutional filters applied across the lead dimension. To ensure these filters accurately approximate discrete derivatives, we enforce moment constraints on their kernels, a key idea from data-driven PDE discovery (Long, Lu, and Dong 2019).

In other words, a standard finite-difference can be viewed as convolution filters whose sum rules enforce a particular order of differentiation(Dong, Jiang, and Shen 2017).

To approximate spatial derivatives along this axis, we learn convolution filters $q[k]$, $k \in \left\{-\frac{N-1}{2}, \ldots, \frac{N-1}{2}\right\}$, of size N. When convolved with the ECG signal $u[x]$, each filter acts as a *finite-difference operator*:

$$(f \circledast q)[x] = \sum_k q[k] f_{\text{padded}}(x + k),$$

A 1D convolution with a filter q of size N can approximate a finite-difference operator. For a filter to approximate an m-th order derivative, its moments must satisfy specific sum rules. For a first-order derivative (D_x), the constraints are:

$$\sum_k q[k] = 0 \qquad \text{and} \qquad \sum_k k \cdot q[k] = 1$$

For a second-order derivative ($D_{xx}u$):

$$\sum_k q[k] = 0, \qquad \sum_k k \cdot q[k] = 0, \text{and} \qquad \sum_k k^2 \cdot q[k] = 2$$

These moment constraints ensure that the learned filters accurately approximate first and second order derivatives in the sense of finite difference schemes (Dong, Jiang, and Shen 2017; Long, Lu, and Dong 2019). Importantly, we learn separate filter kernels for each of the SA, AV, and HP modules, enabling the model to capture the distinct spatial coupling characteristics of each physiological subsystem (see A).

3.4 GAN Framework and Loss Functions

Our PDE-based solver serves as the generator (G) within a standard GAN framework. The discriminator (D) is a neural network trained to distinguish real from generated ECGs. The generator is optimized using a composite loss function that combines three key terms.

Adversarial Loss ($L_{\text{adv}}^{(G)}$) This is the standard non-saturating GAN loss, which encourages the generator to produce samples that are indistinguishable from real data by the discriminator:

$$L_{\text{adv}}^{(G)} = -\mathbb{E}_{z\sim p_z}\left[\log D(\widetilde{\mathbf{U}})\right]$$

where $\widetilde{\mathbf{U}}$ is the generated 12-lead ECG sequence.

Coupled PDE Loss ($\mathcal{L}_{\text{PDE}}$) This loss enforces that the generated signals adhere to the learned physiological dynamics. It measures the residual of each PDE component, the discrepancy between the numerical time derivative (from the forward-Euler step) and the output of the learned response function F_k.

For each subsystem $k \in \{\text{SA}, \text{AV}, \text{HP}\}$, the loss is defined as:

$$\mathcal{L}_{\text{PDE},k} = \sum_{t,x} \left\| \frac{\mathbf{u}_k(t+\Delta t, x) - \mathbf{u}_k(t,x)}{\Delta t} - F_k\Big(\mathbf{u}_k(t,x),\, D_x\mathbf{u}_k(t,x), \ldots; S(\mathbf{u}_k(t,x))\Big)\right\|^2 \tag{6}$$

where $S(\mathbf{u}_k(t,x))$ is the stimulus from the parent subsystem (e.g., $S(\mathbf{u}_{\text{SA}})$ for AV, and $S(\mathbf{u}_{\text{AV}})$ for HP).

The total PDE loss is a weighted sum:

$$\mathcal{L}_{\text{PDE}} = w_{\text{SA}}\mathcal{L}_{\text{PDE,SA}} + w_{\text{AV}}\mathcal{L}_{\text{PDE,AV}} + w_{\text{HP}}\mathcal{L}_{\text{PDE,HP}} \tag{7}$$

where w_{SA}, w_{AV}, and w_{HP} are hyperparameters balancing the contribution of each subsystem to the total loss.

Temporal Gating Loss ($\mathcal{L}_{\text{temporal}}$) To ensure each component is active only during its physiologically relevant time window, we introduce a temporal gating loss. This loss uses fixed binary masks $M_{\text{SA}}, M_{\text{AV}}, M_{\text{HP}}$ to penalize activity outside the expected P-wave, PR interval, and QRS-T segments, respectively,

$$\mathcal{L}_{\text{temporal}} = \sum_{k\in\{\text{SA,AV,HP}\}} w_k \sum_{t,x} \|\mathbf{u}_k(t,x) \odot (1 - M_k(t,x))\|_F^2 \tag{8}$$

where $\odot$ denotes element-wise multiplication, $\|\cdot\|_F$ is the Frobenius norm, and w_k are balancing hyperparameters, $M_k(t,x) = 1$ if component k is expected to be active at time t for lead x, and 0 otherwise.

Total Generator Loss The final generator loss is a weighted sum of the following components:

$$L_G = \lambda_{\text{adv}} L_{\text{adv}}^{(G)} + \lambda_{\text{PDE}}\mathcal{L}_{\text{PDE}} + \lambda_{\text{temporal}}\mathcal{L}_{\text{temporal}}$$

where λ_{adv}, λ_{PDE}, and $\lambda_{\text{temporal}}$ are hyperparameters controlling the relative importance of each loss component.

This integrated framework guides the generator to produce signals that satisfy both the empirical data distribution and the structured, interpretable model of cardiac electrophysiology (through PDE and temporal losses).

Discriminator Loss: The discriminator's adversarial loss aims to maximize the probability of correctly classifying real samples as real and fake samples as fake:

$$L_{\text{adv}}^{(D)} = -\mathbb{E}_{\mathbf{U}\sim p_{\text{data}}}\left[\log D(\mathbf{U})\right] - \mathbb{E}_{z\sim p_z}\left[\log(1 - D(\widetilde{\mathbf{U}}))\right]$$

Here, $\mathbf{U}$ represents real ECG signal sequences from the dataset, and $\widetilde{\mathbf{U}}$ denotes generated sequences.

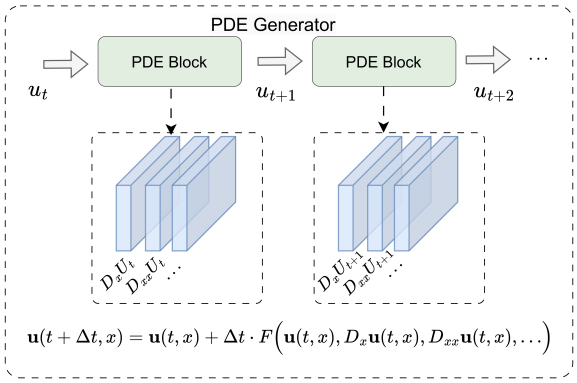

Figure 1: Architecture of the PhysioPDE-GAN generator. The network is composed of repeated PDE blocks, each applying learnable convolutional filters $D_x\mathbf{u}$ and $D_{xx}\mathbf{u}$ to approximate discrete spatial derivatives, followed by a forward–Euler update step used to evolve the ECG signal.

4 Experimental Framework

4.1 ECG Datasets

Georgia 12-Lead ECG Challenge (G12EC) Dataset The empirical evaluation of our proposed method relies on the Georgia 12-Lead ECG Challenge (G12EC) dataset, a key resource from the PhysioNet Challenge (Alday et al. 2020). This dataset is pivotal due to its comprehensive coverage and diverse representation of cardiac conditions. It consists of 10,344 12-lead ECG recordings from 7,871 patients, reflecting a broad demographic spectrum primarily from the southeastern United States. Each recording captures 10 seconds of ECG data, sampled at 500 Hz, resulting in 5,000 time points per recording. This high-frequency data collection ensures detailed ECG waveforms, facilitating accurate analysis and synthesis.

PTB-XL Dataset To further evaluate the effectiveness of our approach, we utilize the PTB-XL dataset (Wagner et al. 2020), a widely recognized benchmark dataset for ECG classification. Collected by the Physikalisch-Technische Bundesanstalt (PTB) in Germany, it contains 21,799 12-lead ECG recordings from 18,869 patients, offering a rich resource for testing model performance across a wide range of cardiac conditions. Each recording is 10 seconds long and includes metadata such as patient demographics, diagnostic annotations, and signal characteristics.

Dataset Splitting and Cross-Validation To ensure robust model evaluation, we employed 5-fold cross-validation, dividing the dataset into training and validation subsets. This approach provides reliable performance estimates by rotating the training and validation data across different splits. Folds are created inside the 80% training pool. Additionally, 20% of the dataset was set aside as an independent test set, reserved exclusively for evaluating the final model's performance, following the practices outlined in (Xu and Goodacre 2018).

To prevent data leakage, we enforce a patient-wise split, ensuring that ECG recordings from the same patient appear exclusively in either the training/validation set or the test set, but not both. This guarantees that model performance is evaluated on entirely unseen patients.

4.2 ECG Classifier for Synthetic Data Evaluation

To assess the utility of the synthetic 12-lead ECG data generated by PhysioPDE-GAN, we evaluate classifier performance using two state-of-the-art architectures widely adopted in 12-lead ECG analysis: a standard ResNet (Attia et al. 2019; Ribeiro et al. 2020) and a ResNet enhanced with attention mechanisms (Nejedly et al. 2021), which ranked among the top performers in the PhysioNet Challenge.

The first is a standard ResNet architecture, following the design in (Ribeiro et al. 2020). It begins with a convolutional layer followed by five residual blocks, each containing three convolutional layers with batch normalization, ReLU, dropout, and skip connections to ensure stable training. The number of filters increases progressively across blocks, with strided convolutions used for temporal downsampling. A global average pooling layer aggregates features before passing them to a dense output layer.

The second architecture builds on this backbone by incorporating a multi-head attention mechanism (Nejedly et al. 2021). This design allows the model to better capture long-range temporal dependencies and inter-lead relationships, which is particularly relevant for the complex spatiotemporal dynamics of 12-lead ECG signals.

4.3 Experimental Methodology

To evaluate the quality of synthetic data generated by PhysioPDE-GAN, we adopt a methodology proposed in (Alcaraz and Strodthoff 2023; Yehuda and Radinsky 2024) where classifiers are trained on different data mixtures and then evaluated on a real, held-out test set. If models trained on combined real and synthetic data perform comparably to those trained on real data alone, it indicates that the synthetic samples are of high quality and effectively support training. Conversely, a significant drop in test performance would suggest a distributional shift, reflecting lower-quality synthetic data. We consider three experimental settings:

1. Real Data Only: The classifier is trained exclusively on real data from the training set.

2. Real + Synthetic Data: The classifier is trained on real data augmented with synthetic ECG heartbeats generated by PhysioPDE-GAN.

3. Real + Alternative Synthetic Data: The classifier is trained on real data combined with synthetic ECGs generated by other baseline models.

We evaluate performance using sensitivity and specificity, which are standard and well-established metrics for ECG classification, as supported by prior studies (Wang et al. 2018; Bressman et al. 2020; Golany et al. 2022). All experiments fix classifier sensitivity at a set threshold and compare specificity on a held-out test set of unseen real samples, using clinically standard metrics.

4.4 Implementation Details

To prepare the raw ECG data for training and testing, we applied signal processing techniques using NeuroKit2, a Python library for neurophysiological signal processing (Makowski et al. 2021). Each signal was segmented into individual heartbeat cycles by detecting R-peaks, a process performed on Lead-II due to its high peak visibility and clinical reliability. The detected RR intervals were then used to synchronize cycle segmentation across all 12 leads.

For the reaction-diffusion modeling of the SA, AV, and HP systems, we additionally detected the P, Q, R, S, and T peaks using the same approach on Lead II.

The detected RR intervals were then used to synchronize cycle segmentation across all 12 leads, ensuring temporal alignment and preserving morphological coherence throughout the cardiac cycle. This alignment is critical for training a generative model that captures realistic inter-lead dynamics. Additionally, we select abnormalities that are consistently present across all cycles within a recording. This ensures stable training and avoids variability from intermittently expressed anomalies. Each segmented ECG cycle

Abnormality	Baseline CLS*		**DCGAN**	**WaveGAN**	**ME-GAN**	**MultiODE**	**SSSD-ECG**	**PhysioPDE-GAN**
	Sensitivity	Specificity	Specificity	Specificity	Specificity	Specificity	Specificity	Specificity
IAVB	0.94	0.82 ± 0.018	0.82 ± 0.012	0.83 ± 0.011	0.83 ± 0.013	0.85 ± 0.012	0.84 ± 0.014	$\mathbf{0.88 \pm 0.012}$
RBBB	0.94	0.89 ± 0.013	0.89 ± 0.012	0.89 ± 0.011	0.90 ± 0.010	0.92 ± 0.011	0.90 ± 0.012	$\mathbf{0.93 \pm 0.013}$
LBBB	0.97	0.96 ± 0.002	0.95 ± 0.003	0.96 ± 0.002	0.95 ± 0.002	0.96 ± 0.003	0.96 ± 0.002	0.96 ± 0.002
NSIVCB	0.78	0.72 ± 0.015	0.70 ± 0.013	0.73 ± 0.011	0.74 ± 0.010	0.76 ± 0.014	0.75 ± 0.012	$\mathbf{0.79 \pm 0.014}$
LAnFB	0.89	0.76 ± 0.006	0.76 ± 0.005	0.77 ± 0.006	0.77 ± 0.006	0.80 ± 0.008	0.77 ± 0.006	$\mathbf{0.83 \pm 0.008}$
QAb	0.81	0.70 ± 0.009	0.70 ± 0.008	0.70 ± 0.009	0.70 ± 0.008	0.72 ± 0.008	0.72 ± 0.009	$\mathbf{0.74 \pm 0.006}$
AFL	0.93	0.83 ± 0.011	0.82 ± 0.013	0.84 ± 0.012	0.85 ± 0.014	0.87 ± 0.014	0.85 ± 0.012	$\mathbf{0.88 \pm 0.019}$

Table 1: Classifier performance on real test data. Baseline uses only real ECGs, while other columns show performance when training is augmented with synthetic ECGs from each model. *Baseline classifier follows (Ribeiro et al. 2020)

is denoted as $\hat{x} \in \mathbb{R}^{12\times L}$, preserving the label from the full recording $x \in \mathbb{R}^{12\times 5000}$.

Experiments ran on a single NVIDIA A40 (48GB) GPU. Further implementation and hyperparameter details are provided in the Appendix C and the code repository.

5 Experimental Results

5.1 Main Result

We evaluate the effectiveness of PhysioPDE-GAN, a generative model that leverages partial differential equations (PDEs) to synthesize high-fidelity 12-lead ECG signals. To assess its utility, we compare the performance of classifiers trained on a combination of real and synthetic data against two benchmarks: 1) a classifier trained solely on real data and 2) classifiers trained on synthetic data from alternative state-of-the-art generative models.

We compare PhysioPDE-GAN against several generative baselines, including leading models for 12-lead ECG synthesis, to demonstrate the superior effectiveness of our proposed method in enhancing classifier performance. The generative models considered are:

- **DCGAN** (Radford, Metz, and Chintala 2016): A standard GAN architecture adapted for generating ECG waveforms.
- **ME-GAN** (Chen et al. 2022): A disease-aware GAN that synthesizes multi-lead ECGs using a panoptic representation.
- **WaveGAN** (Donahue, McAuley, and Puckette 2018): Originally designed for audio waveform synthesis, WaveGAN effectively models temporal dependencies and is adapted here for ECG time series.
- **MultiODE-GAN** (Yehuda and Radinsky 2024): An ODE-based model that captures the dynamics of 12-lead ECG signals, simulating via known ODEs.
- **SSSD-ECG** (Alcaraz and Strodthoff 2023): A recent diffusion-based generative model tailored for multi-lead ECG synthesis.
- **PhysioPDE-GAN**: Our proposed approach, which incorporates learnable partial differential equations to jointly capture the spatial and temporal dynamics of 12-lead ECG signals.

As shown in Table 1, classifiers trained with synthetic data from PhysioPDE-GAN consistently achieve superior performance. We observe statistically significant improvements in specificity across a range of cardiac abnormalities when augmenting the training data with our generated signals, while maintaining consistent sensitivity. These gains indicate that the physiologically-constrained nature of PhysioPDE-GAN leads to enhanced model robustness and better generalization on real-world test data.

Unlike black-box diffusion models like SSSD-ECG or purely temporal-focused methods like MultiODE-GAN, PhysioPDE-GAN explicitly embeds differential operators within the generator. This allows it to learn and enforce the spatiotemporal relationships that govern cardiac electrical activity, resulting in synthetic ECG signals that are not only realistic but also physiologically consistent.

The reported improvements were validated using a 5-fold cross-validated paired t-test, with statistically significant results confirmed by p-values less than 0.05. Statistically significant improvements are highlighted in bold in the tables.

5.2 Ablation Experiments

To comprehensively validate our approach, we performed several key experiments beyond the main results. These studies demonstrate the generalizability of PhysioPDE-GAN to a different dataset, provide qualitative evidence of its fidelity, and assess its compatibility with advanced classifier architectures.

Additional 12 lead ECG Dataset - PTB-XL To demonstrate the generalizability of PhysioPDE-GAN, we conducted a full comparative study using the PTB-XL dataset. The results, summarized in Table 2, reveal trends consistent with our primary findings. Classifiers trained on a combination of real and synthetic data from PhysioPDE-GAN exhibited significant improvements in specificity for a wide range of cardiac conditions. These results reinforce the model's adaptability across different datasets and its utility in improving the detection of both common and rare cardiac abnormalities. Moreover, conditions with similar prevalence in both datasets yielded closely aligned classification results, further demonstrating the consistency and robustness of PhysioPDE-GAN.

Qualitative Results Figure 2 provides a qualitative comparison of ECG heartbeats for Leads II and V1 generated by

Abnormality	Baseline CLS		PhysioPDE-GAN	
	Sensitivity	Specificity	Sensitivity	Specificity
IAVB	0.92	0.83	0.92	**0.88**
IRBBB	0.91	0.86	0.91	**0.90**
LBBB	0.97	0.97	0.97	0.97
NSIVCB	0.81	0.71	0.81	**0.76**
LAnFB	0.96	0.87	0.96	**0.92**
QAb	0.78	0.70	0.78	**0.76**
AFL	0.89	0.78	0.89	**0.83**

Table 2: Classifier performance on the PTB-XL dataset. The Baseline model is trained solely on real ECG data, while the augmented model is trained on real data combined with PhysioPDE-GAN synthetic ECGs. Statistically significant improvements are highlighted in bold.

PhysioPDE-GAN and MultiODE-GAN (Yehuda and Radinsky 2024). In each subplot, real signals (blue) are shown alongside generated signals (orange). The waveforms produced by PhysioPDE-GAN closely match the real ECG morphology and physiological characteristics, while MultiODE-GAN often exhibits artifacts and distortions. These examples visually demonstrate the superior ability of PhysioPDE-GAN to generate physiologically accurate and coherent ECG signals, a direct result of its PDE-based framework.

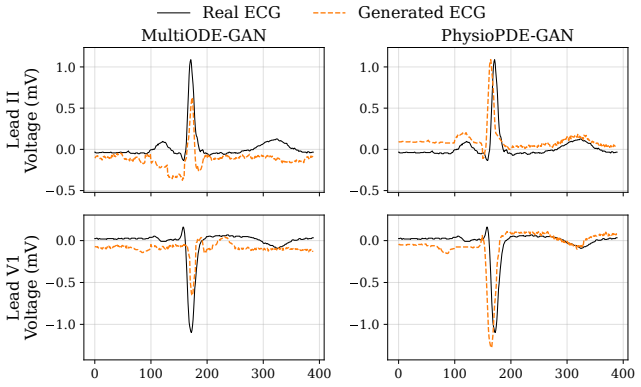

Figure 2: Comparison of generated ECG heartbeats for Lead II and Lead V1. Left: MultiODE-GAN (dashed orange) versus real ECG (solid black). Right: PhysioPDE-GAN (dashed orange) versus real ECG (solid black).

Impact of Classifier Architectures In this experiment, we investigate how different classifier architectures utilize synthetic data generated by PhysioPDE-GAN and whether advanced models can better leverage this data. We compare the standard ResNet to a ResNet with multi-head attention (Nejedly et al. 2021). This attention-based architecture ranked among the top performers in the PhysioNet Challenge. The classifier was trained on a combination of real and synthetic data from PhysioPDE-GAN and evaluated on a real test set. As shown in Table 3, the classifier exhibited a significant improvement in specificity when trained with synthetic data, compared to training on real data alone.

These findings highlight the compatibility of our synthetic data with advanced neural architectures. The physiological fidelity of PhysioPDE-GAN generated data allows sophisticated models to enhance their diagnostic performance, an essential advantage in medical applications where real-world data is often limited or imbalanced.

Abnormality	Alt-CLS *		PhysioPDE-GAN	
	Sensitivity	Specificity	Sensitivity	Specificity
IAVB	0.93	0.85	0.93	**0.89**
RBBB	0.93	0.90	0.93	**0.91**
LBBB	0.96	0.96	0.96	0.96
NSIVCB	0.80	0.73	0.80	**0.78**
LAnFB	0.90	0.76	0.90	**0.81**
QAb	0.83	0.72	0.83	**0.77**
AFL	0.92	0.82	0.92	**0.86**

Table 3: Evaluation results for an alternative classifier architecture, reported with and without data augmentation. *Alt-CLS architecture follows (Nejedly et al. 2021).

Additional ablation studies are provided in Appendix B.

6 Limitations

The benefits of our method are most pronounced for conditions where baseline classifier performance is low, as demonstrated by the significant gains observed for specific cardiac abnormalities. Conversely, for conditions like RBBB and LBBB, where classifiers already achieve over 96% sensitivity and specificity, the room for improvement is naturally limited. Additionally, while our model treats the 12-lead ECG index as an abstracted spatial dimension to capture inter-lead dynamics, this abstraction, while not anatomically literal, is justified by the Dower transform and is a standard approach for modeling inter-lead dynamics.

7 Conclusion

In this work, we introduced PhysioPDE-GAN, a novel generative framework that integrates learnable partial differential equations into a GAN architecture. By embedding domain-specific reaction-diffusion dynamics, our method generates synthetic 12-lead ECG signals with high physiological fidelity, capturing both the spatial and temporal coherence of cardiac electrical activity.

Our experiments demonstrate that classifiers trained on a combination of real and PhysioPDE-GAN generated data outperform those trained solely on real data, achieving significant improvements in specificity and generalization. This underscores the impact of incorporating domain-specific physiological modeling into generative frameworks for medical data synthesis.

By advancing the realism of synthetic ECG signals, PhysioPDE-GAN contributes to both machine learning research and practical applications in healthcare. This work may assist in training robust ECG classifiers in privacy-sensitive or data-scarce environments.

References

Alcaraz, J. M. L.; and Strodthoff, N. 2023. Diffusion-based conditional ECG generation with structured state space models. *Computers in Biology and Medicine*, 163: 107115.

Alday, E. A. P.; Gu, A.; Shah, A. J.; Robichaux, C.; Wong, A.-K. I.; Liu, C.; Liu, F.; Rad, A. B.; Elola, A.; Seyedi, S.; et al. 2020. Classification of 12-lead ecgs: the physionet/computing in cardiology challenge 2020. *Physiological measurement*, 41(12): 124003.

Attia, Z. I.; Kapa, S.; Lopez-Jimenez, F.; McKie, P. M.; Ladewig, D. J.; Satam, G.; Pellikka, P. A.; Enriquez-Sarano, M.; Noseworthy, P. A.; Munger, T. M.; et al. 2019. Screening for cardiac contractile dysfunction using an artificial intelligence–enabled electrocardiogram. *Nature medicine*, 25(1): 70–74.

Bar-Sinai, Y.; Hoyer, S.; Hickey, J.; and Brenner, M. P. 2019. Learning data-driven discretizations for partial differential equations. *Proceedings of the National Academy of Sciences*, 116(31): 15344–15349.

Bishop, M. J.; and Plank, G. 2011. Bidomain ECG simulations using an augmented monodomain model for the cardiac source. *IEEE Transactions on Biomedical Engineering*, 58(8): 2200–2210. Epub 2011 May 2.

Bressman, M.; Mazori, A. Y.; Shulman, E.; Chudow, J. J.; Goldberg, Y.; Fisher, J. D.; Ferrick, K. J.; Garcia, M.; Di Biase, L.; and Krumerman, A. 2020. Determination of Sensitivity and Specificity of Electrocardiography for Left Ventricular Hypertrophy in a Large, Diverse Patient Population. *The American Journal of Medicine*, 133(9): e495–e500.

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 Chaudhuri, K.; Jegelka, S.; Song, L.; Szepesvari, C.; Niu, G.; and Sabato, S., eds., *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, 3360–3370. PMLR.

Cichy, R. M.; and Kaiser, D. 2019. Deep Neural Networks as Scientific Models. *Trends in Cognitive Sciences*, 23(4): 305–317.

de Melo, C. M.; Torralba, A.; Guibas, L.; DiCarlo, J.; Chellappa, R.; and Hodgins, J. 2022. Next-generation deep learning based on simulators and synthetic data. *Trends in Cognitive Sciences*, 26(2): 174–187.

Donahue, C.; McAuley, J. J.; and Puckette, M. S. 2018. Synthesizing Audio with Generative Adversarial Networks. *CoRR*, abs/1802.04208.

Dong, B.; Jiang, Q.; and Shen, Z. 2017. Image Restoration: Wavelet Frame Shrinkage, Nonlinear Evolution PDEs, and Beyond. *Multiscale Modeling & Simulation*, 15(1): 606–660.

Giuffrè, M.; and Shung, D. L. 2023. Harnessing the power of synthetic data in healthcare: innovation, application, and privacy. *npj Digital Medicine*, 6(1): 186.

Golany, T.; Freedman, D.; and Radinsky, K. 2021. ECG ODE-GAN: Learning Ordinary Differential Equations of ECG Dynamics via Generative Adversarial Learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(1): 134–141.

Golany, T.; Radinsky, K.; and Freedman, D. 2020. SimGANs: Simulator-Based Generative Adversarial Networks for ECG Synthesis to Improve Deep ECG Classification. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 3597–3606. PMLR.

Golany, T.; Radinsky, K.; Kofman, N.; Litovchik, I.; Young, R.; Monayer, A.; Love, I.; Tziporin, F.; Minha, I.; Yehuda, Y.; Ziv-Baran, T.; Fuchs, S.; and Minha, S. 2022. Physicians and Machine-Learning Algorithm Performance in Predicting Left-Ventricular Systolic Dysfunction from a Standard 12-Lead Electrocardiogram. *Journal of Clinical Medicine*, 11(22): 6767.

Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative Adversarial Nets. In Ghahramani, Z.; Welling, M.; Cortes, C.; Lawrence, N.; and Weinberger, K., eds., *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc.

Gopal, B.; Han, R.; Raghupathi, G.; Ng, A.; Tison, G.; and Rajpurkar, P. 2021. 3KG: Contrastive Learning of 12-Lead Electrocardiograms using Physiologically-Inspired Augmentations. In Roy, S.; Pfohl, S.; Rocheteau, E.; Tadesse, G. A.; Oala, L.; Falck, F.; Zhou, Y.; Shen, L.; Zamzmi, G.; Mugambi, P.; Zirikly, A.; McDermott, M. B. A.; and Alsentzer, E., eds., *Proceedings of Machine Learning for Health*, volume 158 of *Proceedings of Machine Learning Research*, 156–167. PMLR.

Huang, S.; Wang, P.; and Li, R. 2023. Noise ECG generation method based on generative adversarial network. *Biomedical Signal Processing and Control*, 81: 104444.

Karniadakis, G. E.; Kevrekidis, I. G.; Lu, L.; Perdikaris, P.; Wang, S.; and Yang, L. 2021. Physics-informed machine learning. *Nature Reviews Physics*, 3(6): 422–440.

Liu, H.; Zhao, Z.; Chen, X.; Yu, R.; and She, Q. 2020. Using the VQ-VAE to improve the recognition of abnormalities in short-duration 12-lead electrocardiogram records. *Computer Methods and Programs in Biomedicine*, 196: 105639.

Long, Z.; Lu, Y.; and Dong, B. 2019. PDE-Net 2.0: Learning PDEs from data with a numeric-symbolic hybrid deep network. *Journal of Computational Physics*, 399: 108925.

Long, Z.; Lu, Y.; Ma, X.; and Dong, B. 2018. PDE-Net: Learning PDEs from Data. In Dy, J.; and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 3208–3216. PMLR.

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.

McSharry, P. E.; Clifford, G. D.; Tarassenko, L.; and Smith, L. A. 2003. A dynamical model for generating synthetic

electrocardiogram signals. *IEEE transactions on biomedical engineering*, 50(3): 289–294.

Nejedly, P.; Ivora, A.; Smisek, R.; Viscor, I.; Koscova, Z.; Jurak, P.; and Plesinger, F. 2021. Classification of ECG Using Ensemble of Residual CNNs with Attention Mechanism. In *2021 Computing in Cardiology (CinC)*, volume 48, 1–4.

Potse, M. 2018. Scalable and Accurate ECG Simulation for Reaction-Diffusion Models of the Human Heart. *Frontiers in Physiology*, 9: 370.

Potse, M.; Dubé, B.; Richer, J.; Vinet, A.; and Gulrajani, R. M. 2006. A comparison of monodomain and bidomain reaction-diffusion models for action potential propagation in the human heart. *IEEE Transactions on Biomedical Engineering*, 53(12): 2425–2435.

Quiroz-Juárez, M. A.; Jiménez-Ramírez, O.; Vázquez-Medina, R.; Breña-Medina, V.; Aragón, J. L.; and Barrio, R. A. 2019. Generation of ECG signals from a reaction-diffusion model spatially discretized. *Scientific Reports*, 9(1): 19000.

Radford, A.; Metz, L.; and Chintala, S. 2016. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. In Bengio, Y.; and LeCun, Y., eds., *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.

Raissi, M.; Perdikaris, P.; and Karniadakis, G. 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378: 686–707.

Ribeiro, A. H.; Ribeiro, M. H.; Paixão, G. M.; Oliveira, D. M.; Gomes, P. R.; Canazart, J. A.; Ferreira, M. P.; Andersson, C. R.; Macfarlane, P. W.; Meira Jr, W.; et al. 2020. Automatic diagnosis of the 12-lead ECG using a deep neural network. *Nature communications*, 11(1): 1–9.

Voigt, P.; and Bussche, A. v. d. 2017. *The EU General Data Protection Regulation (GDPR): A Practical Guide*. Springer Publishing Company, Incorporated, 1st edition. ISBN 3319579584.

Wagner, P.; Strodthoff, N.; Bousseljot, R.-D.; Kreiseler, D.; Lunze, F. I.; Samek, W.; and Schaeffter, T. 2020. PTB-XL, a large publicly available electrocardiography dataset. *Scientific Data*, 7(1): 154.

Wang, J. J.; Pahlm, O.; Warren, J. W.; Sapp, J. L.; and Horáček, B. M. 2018. Criteria for ECG detection of acute myocardial ischemia: Sensitivity versus specificity. *Journal of Electrocardiology*, 51(6, Supplement): S12–S17.

Xu, Y.; and Goodacre, R. 2018. On Splitting Training and Validation Set: A Comparative Study of Cross-Validation, Bootstrap and Systematic Sampling for Estimating the Generalization Performance of Supervised Learning. *Journal of Analysis and Testing*, 2(3): 249–262.

Yehuda, Y.; Freedman, D.; and Radinsky, K. 2023. Self-supervised Classification of Clinical Multivariate Time Series using Time Series Dynamics. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '23, 5416–5427. New York, NY, USA: Association for Computing Machinery. ISBN 9798400701030.

Yehuda, Y.; and Radinsky, K. 2024. Ordinary Differential Equations for Enhanced 12-Lead ECG Generation. arXiv:2409.17833.

Zhu, F.; Ye, F.; Fu, Y.; Liu, Q.; and Shen, B. 2019. Electrocardiogram generation with a bidirectional LSTM-CNN generative adversarial network. *Scientific Reports*, 9(1): 6734.