

*Background: Duke Kunshan University (DKU) is an interdisciplinary institution that grants dual undergraduate degrees, an MOE Chinese degree and a degree from Duke University in Durham, United States. The principal structure of DKU majors is robustly interdisciplinary. No student confines their study to a single discipline (for example, biology or economics). Instead, all students engage in broad inquiry related to a subject or question (for example, political economy or global health) and take a wide variety of courses related to that area (for example, in public policy, history, ethics, or economics). As a result, our graduates are prepared to engage in a wide variety of inquiries using multiple methodologies to address complex issues that require interdisciplinary approaches. This has implications for our vision and expectations of undergraduate theses and design projects, which reflect this broad interdisciplinary training. At DKU, every student completes a two-year project known as signature work which consists of multiple interconnected parts including thematic courses, experiential learning, capstones, and a final product. It seeks to integrate students' interdisciplinary educational experience and culminates in the creation of a product in a scholarly, creative, or applied nature in lieu of an undergraduate thesis or design required by JED. Because DKU encourages students to cultivate their independence and creativity as one of its institutional student learning outcomes, the student-led signature work projects often reflect students' own particular interdisciplinary interests and training. In addition, signature work has an intensive emphasis on problem-solving and skill-development which is much needed for any interdisciplinary inquiry; thus, students' final products are evidence of transferrable skills that students have acquired and demonstrated through the 2-year program, rather than content knowledge narrowly defined by disciplinary training. In sum, while the Chinese major declared with any given student might be construed narrowly, the experience of our students is much broader—and intentionally so. This is a distinctive feature of our curriculum, and this distinctiveness results in broadly interdisciplinary submissions from our graduates' submitting theses or design projects. We have designed this to prepare our students for a wide variety of graduate programs in China and the West, where interdisciplinary training is a competitive advantage.*

# MANIFOLD-INFORMED NEURAL NETWORK FOR LOW-DATA, IMBALANCED ECG INFERENCE WITH AUXILIARY CLINICAL METADATA

by

Juntang Wang

Signature Work Product, in partial fulfillment of the  
Duke Kunshan University Undergraduate Degree Program

*March 8, 2026*

Signature Work Program  
Duke Kunshan University

## APPROVALS

---

Mentor: Shixin Xu, Division of Natural and Applied Sciences

---

---

# CONTENTS

---

Acknowledgements	ii
Abstract	iii
List of Figures	iii
List of Tables	iv
1 Introduction	1
2 Related Work	3
3 Methods	6
4 Experiments	12
5 Discussion	18
6 Conclusion	21
References	22
Signature Work Narrative	25
A Model and Hyperparameter Details	27
B Additional Experimental Results	28
C Implementation and Reproducibility	30

---

# ACKNOWLEDGEMENTS

---

I am deeply grateful to my mentor, Prof. Shixin Xu, for his guidance, patience, and intellectual generosity throughout the development of this project. His expertise in physics-informed machine learning and his willingness to let me explore unconventional ideas were instrumental in shaping the MINN framework. I also thank the Division of Natural and Applied Sciences at Duke Kunshan University for providing the interdisciplinary environment that made this work possible. Finally, I thank my family for their unwavering support throughout my undergraduate studies.

---

# ABSTRACT

---

Physics-informed neural networks (PINNs) have demonstrated that penalizing the residual of a known differential operator provides powerful supervision for scientific inference. However, many biomedical tasks—such as inferring serum electrolyte levels from electrocardiograms (ECGs)—lack a known governing equation. We introduce the **Manifold-Informed Neural Network (MINN)**, a PINN-parallel paradigm that replaces the physics operator residual with a manifold operator residual: a discrete Poisson equation on a  $k$ -nearest-neighbor graph in the learned latent space, with a learned, task-dependent source term. Unlike classical manifold regularization, which imposes a homogeneous smoothness prior, MINN enforces a non-trivial Poisson equation whose source encodes both input features and auxiliary clinical metadata via gated Feature-wise Linear Modulation (FiLM), a conditioning mechanism that selectively modulates the signal representation. We evaluate MINN on three-class classification of potassium, calcium, and magnesium from 12-lead ECGs in the MIMIC-IV-ECG dataset, benchmarking against nine baselines across five categories under low-data and class-imbalanced conditions. MINN achieves the highest  $K^+$  Macro-F1 of 0.61 (best baseline 0.55), with a 27% relative improvement at 1% labeled data. Ablation studies demonstrate that the Poisson formulation with a learned source is structurally and empirically distinct from classical manifold regularization, particularly at low label rates.

**Keywords:** manifold-informed neural network, ECG electrolyte inference, Poisson operator residual, class-imbalanced learning, clinical metadata fusion

---

# LIST OF FIGURES

---

3.1	MINN architecture overview . . . . .	7
3.2	PINN vs. MINN structural parallel . . . . .	10
4.1	Data processing pipeline . . . . .	12
4.2	Low-data performance and source norm analysis . . . . .	16

---

# LIST OF TABLES

---

3.1	Summary of notation. . . . .	6
3.2	MINN vs. classical manifold regularization . . . . .	10
4.1	Dataset stats by electrolyte bin . . . . .	13
4.2	Main electrolyte classification results . . . . .	15
4.3	Ablation study results . . . . .	16
4.4	Subgroup stress test results . . . . .	17
A.1	Hyperparameter settings for MINN. . . . .	27
B.1	Extended baseline results . . . . .	28
B.2	Per-class AUCPR . . . . .	28
B.3	Model calibration . . . . .	29
B.4	Low-data results for all three electrolytes . . . . .	29
B.5	CNN encoder ablation . . . . .	29

## Chapter 1

---

# INTRODUCTION

---

## 1.1 Motivation and Context

Electrolyte imbalances are among the most common and dangerous derangements in hospitalized patients. Abnormal concentrations of potassium, calcium, and magnesium can trigger fatal cardiac arrhythmias, yet their diagnosis requires invasive venous blood draws that introduce delays and are unavailable in ambulatory or resource-limited settings. Because electrolyte perturbations directly alter cardiac ion channel function, their effects are encoded in the surface electrocardiogram (ECG)—a non-invasive, widely available, and inexpensive measurement. Inferring electrolyte status from ECGs would enable continuous, real-time monitoring without blood sampling.

Recent advances in deep learning [1] have made automated ECG analysis increasingly feasible. From a mathematical perspective, this inference task is an *inverse problem*: the ECG is a surface measurement generated by an underlying tissue state (electrolyte concentrations), and the forward mapping from state to measurement is governed by complex, only partially understood cardiac electrophysiology. In standard inverse-problem settings, physics-informed neural networks (PINNs) [2] have proven effective by penalizing the residual of a known differential operator. However, when the governing physics is partial or unknown—as is the case for the electrolyte-to-ECG relationship—PINNs cannot be directly applied.

Clinical electrolyte datasets compound this difficulty: pathological concentrations are inherently rare, producing severe class imbalance. The resulting scarcity of minority-class labels is further exacerbated by the small overall size of labeled datasets relative to the high dimensionality of the input signal, leaving standard supervised models prone to overfitting. Finally, each ECG is accompanied by a heterogeneous set of auxiliary clinical metadata—demographics and active medications—that carry complementary information but whose interaction with the signal is not well characterized, complicating efforts to integrate all available evidence into a single model.

## 1.2 Problem Statement

We seek a learning framework that provides the structural benefits of PINNs—dense operator-based supervision from a differential equation residual—without requiring a known physics operator. The key insight is that even when the governing equations are unknown, the *geometry* of the data manifold provides a natural substitute: the graph Laplacian of a sample graph in the learned latent space serves as a data-driven differential operator.

We introduce the **Manifold-Informed Neural Network (MINN)**, whose core idea—substituting a manifold operator for the physics operator—is developed in Chapter 3. Unlike classical manifold regularization [3], which imposes a homogeneous smoothness prior that degenerates at low label rates [4], MINN enforces a Poisson-type equation with a learned, task-dependent source term conditioned on both the input signal and auxiliary metadata.

## 1.3 Contributions

The main contributions of this work are:

1. **MINN: a PINN-parallel paradigm for manifold-informed inference.** We propose the Manifold-Informed Neural Network, which penalizes the residual of a Poisson equation on a sample graph in the learned latent space. A learned source network  $s_\phi(\mathbf{x}, \mathbf{m})$  converts the manifold constraint from a generic smoothness prior into a task-dependent equation, mirroring the structure of PINNs without requiring known physics.
2. **Gated FiLM metadata fusion.** We integrate heterogeneous clinical metadata through a gated variant of Feature-wise Linear Modulation that selectively scales and shifts the signal embedding on a per-dimension basis.
3. **Comprehensive evaluation on ECG electrolyte inference.** We evaluate MINN on the MIMIC-IV-ECG dataset [5, 6] for three-class classification of potassium, calcium, and magnesium, benchmarking against nine baselines across five categories (ECG foundation models, imbalanced learning, manifold/graph methods, tabular models, and physics-informed approaches).
4. **Ablation studies establishing distinctness from classical manifold regularization.** We present five targeted ablations demonstrating that the Poisson formulation with a learned source is genuinely distinct from—and superior to—classical manifold regularization, particularly at low label rates.

## 1.4 Paper Organization

The remainder of this paper is organized as follows. Chapter 2 reviews related work. Chapter 3 presents the MINN framework. Chapter 4 instantiates MINN for ECG electrolyte inference and reports results. Chapter 5 discusses findings, threats to validity, and limitations. Chapter 6 concludes with directions for future work.

## Chapter 2

---

# RELATED WORK

---

## 2.1 ECG-Based Electrolyte Inference

Serum electrolyte imbalances—particularly of potassium ( $K^+$ ), calcium ( $Ca^{2+}$ ), and magnesium ( $Mg^{2+}$ )—are common in hospitalized patients and can precipitate life-threatening cardiac arrhythmias. Diagnosis currently requires venous blood draws, which are invasive, time-delayed, and unavailable in resource-limited or ambulatory settings. Because electrolyte concentrations directly modulate cardiac ion channel kinetics, their perturbations leave characteristic signatures in the surface electrocardiogram (ECG) [7, 8]. For example, hyperkalemia manifests as peaked T-waves and QT prolongation, hypocalcemia as ST-segment depression, and hypomagnesemia as U-wave prominence.

Recent work has explored deep-learning approaches to infer electrolyte status from 12-lead ECGs. Von Bachmann et al. [8] framed the task as regression with uncertainty quantification, while Alcaraz and Strodthoff [9] extended prediction to multiple laboratory values using multimodal architectures. These approaches treat the ECG as a standard supervised learning problem without exploiting the geometric structure of the data manifold.

These characteristic morphologies constrain the space of plausible ECGs to a low-dimensional manifold in the high-dimensional signal space. Exploiting this manifold structure allows label information to propagate through geometrically similar but unlabeled recordings, a mechanism unavailable to purely supervised approaches.

Our concurrent work, PPIDM-II [10] (under review), introduces a generative, physics-informed diffusion model for ECG electrolyte inference. While PPIDM-II demonstrated the value of incorporating partial physics priors, it relied on a generative framework with substantial computational overhead and did not address the manifold geometry of the latent space. MINN takes a complementary, discriminative approach: rather than generating synthetic ECGs, it directly classifies electrolyte status while enforcing manifold structure through an operator residual in the learned embedding space.

## 2.2 Physics-Informed Neural Networks

Physics-informed neural networks (PINNs) [2] incorporate known governing equations into neural network training by penalizing the residual of a differential operator applied to the network output. Given a PDE of the form  $\mathcal{L}[u] = f$  with operator  $\mathcal{L}$  and source  $f$ , the PINN loss includes a term  $\|\mathcal{L}[u_\theta] - f\|^2$  evaluated at a set of collocation points. This physics residual provides a dense supervisory signal that does not require labeled data at every collocation point, enabling PINNs to solve forward and inverse problems with limited observations.

PINNs have achieved success in fluid dynamics, heat transfer, and biomedical modeling, among other domains. However, their applicability is limited to settings where the governing physics is

known in closed form. In many biomedical inference tasks—including ECG-based electrolyte classification—the relationship between the observed signal and the target variable is mediated by complex, partially understood physiology that cannot be expressed as a tractable PDE.

Xu and Li [11] relaxed the requirement of fully known physics by embedding partial differential equation constraints into diffusion models. MINN extends this line of thinking further: when even partial physics is unavailable, the governing operator can be replaced with a *data-driven* operator—the graph Laplacian of the sample manifold—while retaining the structural benefits of the operator residual formulation.

## 2.3 Manifold Learning and Regularization

The manifold hypothesis posits that high-dimensional data often lie on or near a low-dimensional manifold embedded in the ambient space [12]. Classical dimensionality reduction methods such as Laplacian Eigenmaps [13] and Diffusion Maps [14] exploit graph Laplacian eigenvectors to recover this low-dimensional structure, with the spectral properties of the graph Laplacian playing a central role [15–17]. Manifold regularization [3] operationalizes this idea for semi-supervised learning by adding a penalty  $\Omega = \text{tr}(\mathbf{H}^\top \mathbf{L}_G \mathbf{H})$  to the training objective, where  $\mathbf{H}$  is the embedding matrix and  $\mathbf{L}_G$  is the graph Laplacian of a nearest-neighbor graph constructed from both labeled and unlabeled data. This penalty encourages the learned function to be smooth along the data manifold, effectively propagating label information through the graph. While effective in moderate label settings, this homogeneous smoothness prior degenerates at very low label rates, where the sparse supervision is insufficient to prevent the embedding from collapsing to the trivial constant solution.

Poisson Learning [4] addressed this limitation by replacing the Laplacian smoothness penalty with a Poisson equation on the graph, showing that the resulting semi-supervised classifier remains well-posed at extremely low label rates where classical approaches degenerate. Liang et al. [18] bridged manifold learning and PDE solving by constructing a graph Laplacian  $\mathbf{L}_\varepsilon$  via the Diffusion Maps asymptotic expansion, yielding a consistent estimator of elliptic operators on an unknown manifold sampled by point-cloud data. A deep neural network is then trained by minimizing the squared operator residual  $\|\mathbf{L}_\varepsilon \varphi_\theta - f\|^2$ , formulating PDE solving as supervised learning without requiring a mesh or explicit knowledge of the manifold geometry. They proved convergence of the discrete solution to the continuum PDE solution in the large-data limit, providing a rigorous foundation for data-driven operator residuals on manifolds.

However, none of these methods frames the manifold operator residual as a *PINN-parallel* paradigm with a learned, task-dependent source term. Classical manifold regularization uses a homogeneous Laplace equation (zero source); Poisson Learning uses fixed label-derived sources; label propagation methods [19, 20] spread label information through the graph but do not introduce a learned source. MINN bridges this gap by introducing a learned, task-dependent source into the manifold operator residual, creating a structural analog to PINNs. The source network  $s_\phi(\mathbf{x}, \mathbf{m})$  encodes both signal features and auxiliary metadata, converting the manifold constraint from a generic smoothness prior into a task-specific Poisson equation. This structural parallel—operator applied to network output, non-trivial learned source, collocation at all data points—is the central contribution of this work.

### 2.3.1 Neural Operators and PDE Solving on Manifolds

A growing body of work extends physics-informed and operator-learning methods to complex or unknown geometries, broadening the scope of PDE-constrained neural networks beyond simple Euclidean domains.

$\Delta$ -PINNs [21] introduced domain decomposition for PINNs on complex geometries, expressing the differential operator in local curvilinear coordinates via Laplace–Beltrami eigenfunctions. While effective for known PDEs on irregular domains,  $\Delta$ -PINNs still require the governing equation in closed form. GPINN [22] embeds graph topology into the physics-informed training loop, coupling the graph Laplacian with a known PDE operator to exploit relational structure among collocation points. More recently, curvature-aware graph attention [23] incorporated Riemannian curvature estimates into graph attention mechanisms, enabling PDE solvers to adapt to the local geometry of the data manifold. On the generative side, Manifold Diffusion Fields [24] used the spectral geometry of the manifold—its Laplace–Beltrami eigenfunctions—as positional encodings for a neural field trained with a diffusion process, achieving state-of-the-art generation on meshes and point clouds.

MINN extends this line of research in a distinct direction: rather than solving a *known* PDE on a complex geometry, it applies the manifold operator residual to *inference* tasks where no governing equation is available, using the graph Laplacian as a data-driven substitute for the physics operator and a learned source network in place of a known source term.

## 2.4 Low-Data and Class-Imbalanced Learning

Low-data and class-imbalanced regimes are pervasive in clinical machine learning [25, 26], where labeled datasets are expensive to curate and pathological classes are inherently rare. Standard approaches to class imbalance include cost-sensitive losses with inverse-frequency weighting, oversampling strategies such as SMOTE [27], and logit adjustment [28], which post-hoc corrects the decision boundary to account for class priors.

Data augmentation in the latent space has also proven effective. Mixup [29] and its latent-space extension, Manifold Mixup [30], interpolate examples and their labels at the input or hidden layers, encouraging smoother decision boundaries and better calibration. However, these methods do not exploit the geometric structure of the data manifold; they treat the latent space as Euclidean rather than imposing any manifold-aware constraint.

Foundation models offer another path to low-data learning by pretraining on large unlabeled corpora. ECGFounder [31] and ECG-FM [32] pretrain on hundreds of thousands of ECGs and fine-tune for downstream tasks, but they do not incorporate auxiliary metadata or manifold geometry at the fine-tuning stage. Tabular methods such as TabPFN [33] and CatBoost [34] can leverage metadata features directly but cannot process raw ECG signals.

MINN bridges these challenges by combining three complementary strategies: (i) class-weighted cross-entropy for imbalance, (ii) gated FiLM conditioning for metadata integration, and (iii) a manifold operator residual that provides geometric supervision at all data points, not just labeled ones.

## Chapter 3

# METHODS

### 3.1 Problem Setup and Notation

Consider a general inference task in which a high-dimensional signal  $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^D$  (e.g., a multi-lead ECG with  $\sim 1,000$  time points per lead yields  $D \approx 12,000$ ) is accompanied by an auxiliary metadata vector  $\mathbf{m} \in \mathcal{M} \subseteq \mathbb{R}^p$  and a categorical label  $y \in \{1, \dots, B\}$ . The goal is to learn a classifier that maps  $(\mathbf{x}, \mathbf{m})$  to  $y$  under three simultaneous challenges: (i) the labeled training set is small relative to the ambient dimensionality of  $\mathbf{x}$ ; (ii) the class distribution is heavily imbalanced, with clinically important minority classes; and (iii) the metadata space is heterogeneous, mixing continuous physiological measurements with binary indicator variables that interact with the signal in unknown ways.

When the signal arises from a physical measurement process—for example, a surface potential recording governed by an underlying tissue state—the mapping from hidden state to observation can be viewed as an *inverse problem*. PINNs [2] exploit known governing equations by penalizing the residual of a differential operator applied to the network output. In many biomedical inference settings, however, the governing physics is only partially known or entirely unavailable. MINN addresses this gap by replacing the physics operator with a *manifold operator*—specifically, the graph Laplacian of a sample graph constructed in a learned latent space—and by introducing a task-dependent source term that converts the smoothness prior into a Poisson-type equation.

Table 3.1 summarizes the notation used throughout this chapter.

Table 3.1: Summary of notation.

Symbol	Description
$\mathbf{x} \in \mathbb{R}^D$	Input signal
$\mathbf{m} \in \mathbb{R}^p$	Auxiliary metadata vector
$y \in \{1, \dots, B\}$	Categorical label ( $B$ classes)
$N, T, d$	Training samples, prediction tasks, embedding dimension
$h_\theta : \mathcal{X} \rightarrow \mathbb{R}^d$	Encoder network (parameters $\theta$ )
$\mathbf{h}, \tilde{\mathbf{h}}$	Latent embedding $h_\theta(\mathbf{x})$ / metadata-fused embedding
$\gamma(\mathbf{m}), \beta(\mathbf{m})$	FiLM gate and shift vectors
$\mathbf{H}, \mathbf{S} \in \mathbb{R}^{N \times d}$	Embedding matrix (rows $h_\theta(\mathbf{x}_i)$ ) / source matrix (rows $s_\phi(\mathbf{x}_i, \mathbf{m}_i)$ )
$s_\phi : \mathcal{X} \times \mathcal{M} \rightarrow \mathbb{R}^d$	Learned source network (parameters $\phi$ )
$G = (V, E, W)$	$k$ -NN graph in latent space
$k, \sigma$	Number of nearest neighbors, kernel bandwidth
$\mathbf{W}, \mathbf{D}$	Adjacency weight and degree matrices
$\mathbf{L}_G$	Normalized graph Laplacian
$\lambda, R$	Manifold residual weight, graph update period (epochs)

## 3.2 Framework Overview

Before detailing each component, we present the overall MINN pipeline (Figure 3.1). A raw input signal  $\mathbf{x}$  is first mapped to a  $d$ -dimensional latent embedding  $\mathbf{h} = h_\theta(\mathbf{x})$  by an encoder network. The embedding is then modulated by auxiliary metadata  $\mathbf{m}$  through a gated Feature-wise Linear Modulation (FiLM) layer, producing a fused representation  $\tilde{\mathbf{h}}$  that integrates signal and metadata information. The fused representation feeds  $T$  task-specific prediction heads, each producing a  $B$ -class probability vector.

In parallel, a separate source network  $s_\phi(\mathbf{x}, \mathbf{m})$  produces a  $d$ -dimensional source vector for each sample. The embedding matrix  $\mathbf{H}$  and source matrix  $\mathbf{S}$  are related through a discrete Poisson equation on a  $k$ -nearest-neighbor graph constructed in the latent space, enforced via the manifold operator residual  $\mathcal{L}_{\text{MINN}}$  (Equation (3.7)). The total training loss combines supervised classification with the manifold constraint:

$$\mathcal{L}_{\text{total}} = \sum_{t=1}^T \mathcal{L}_{\text{CE}}^{(t)} + \lambda \mathcal{L}_{\text{MINN}}, \quad (3.1)$$

where  $\mathcal{L}_{\text{CE}}^{(t)}$  is the class-weighted cross-entropy for task  $t$  and  $\lambda > 0$  controls the strength of the manifold constraint. The following sections detail each component.

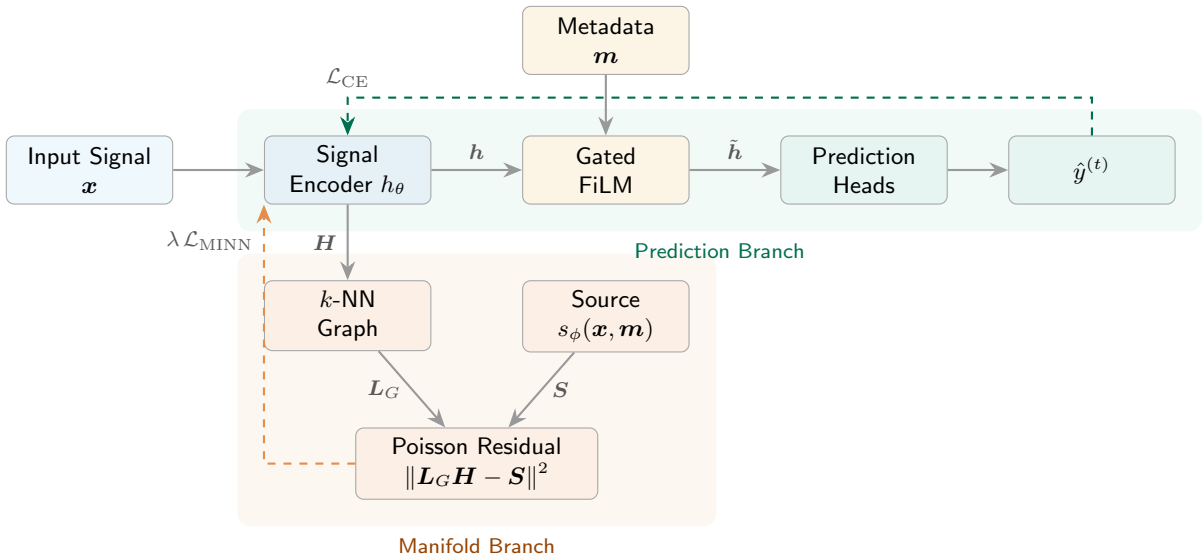


Figure 3.1: Overview of the MINN architecture. The input signal is mapped to a latent embedding by the encoder  $h_\theta$ , fused with auxiliary metadata via gated FiLM, and passed to task-specific prediction heads (*Prediction Branch*). Simultaneously, a  $k$ -NN graph in the latent space and a learned source network  $s_\phi$  define the manifold operator residual (*Manifold Branch*). Dashed arrows indicate gradient flow from the two loss terms.

## 3.3 Network Architecture

**ECG Encoder.** The encoder  $h_\theta : \mathcal{X} \rightarrow \mathbb{R}^d$  maps a raw input signal to a  $d$ -dimensional latent embedding  $\mathbf{h} = h_\theta(\mathbf{x})$ . The MINN framework is *encoder-agnostic*: any architecture producing a fixed-length embedding is compatible with the manifold operator and metadata fusion modules, allowing substitution of domain-specific backbones without modifying the rest of the pipeline. The encoder used in our case study is described in Chapter 4.

**Gated FiLM Metadata Fusion.** In our ECG electrolyte inference problem, the auxiliary metadata  $\mathbf{m} \in \mathbb{R}^p$  comprises patient demographics (age, sex) and binary indicators for active medication classes at the time of recording (Section 4.2 details the full feature set). This metadata carries information complementary to the ECG waveform: certain drug classes directly affect electrolyte levels, while age modulates baseline ECG morphology. Rather than concatenating  $\mathbf{m}$  to the latent embedding—which treats all metadata dimensions equally and ignores potential interactions with the signal representation—we adopt a *gated* variant of Feature-wise Linear Modulation (FiLM) [35].

The fused representation  $\tilde{\mathbf{h}} \in \mathbb{R}^d$  is computed as

$$\tilde{\mathbf{h}} = \gamma(\mathbf{m}) \odot \mathbf{h} + \beta(\mathbf{m}), \quad (3.2)$$

where

$$\gamma(\mathbf{m}) = \text{sigmoid}(\mathbf{W}_\gamma \mathbf{m} + \mathbf{b}_\gamma), \quad (3.3)$$

$$\beta(\mathbf{m}) = \mathbf{W}_\beta \mathbf{m} + \mathbf{b}_\beta, \quad (3.4)$$

with learnable parameters  $\mathbf{W}_\gamma, \mathbf{W}_\beta \in \mathbb{R}^{d \times p}$  and  $\mathbf{b}_\gamma, \mathbf{b}_\beta \in \mathbb{R}^d$ , and  $\odot$  denoting element-wise multiplication.

The sigmoid gate  $\gamma(\mathbf{m}) \in (0, 1)^d$  (Equation (3.3)) controls per-dimension scaling of the signal embedding: dimensions for which metadata is uninformative are left nearly unchanged ( $\gamma_j \approx 1$ ), while informative dimensions are rescaled. The affine shift  $\beta(\mathbf{m})$  (Equation (3.4)) injects a metadata-dependent bias, encoding information present in the metadata but absent from the signal.

**Prediction Heads.** The fused embedding  $\tilde{\mathbf{h}}$  is passed to  $T$  task-specific prediction heads, one per target variable. In our case study,  $T = 3$  (potassium, calcium, magnesium). Each head is a small fully connected network terminating in a  $B$ -class softmax layer:

$$\hat{\mathbf{y}}^{(t)} = \text{softmax}(f^{(t)}(\tilde{\mathbf{h}})), \quad t = 1, \dots, T, \quad (3.5)$$

where  $f^{(t)}$  consists of one or two linear layers with dropout. All heads share the encoder and metadata fusion parameters, so gradient signals from multiple tasks jointly shape the latent space—a form of multi-task learning [36].

### 3.4 Manifold-Informed Operator Residual

This section presents the core contribution of MINN: a manifold operator residual that parallels the physics operator residual in PINNs [2], but replaces the known differential operator with the graph Laplacian of a sample graph constructed in the learned latent space.

**Poisson Equation on the Sample Graph.** Let  $G = (V, E, W)$  denote a  $k$ -nearest-neighbor graph constructed over the  $N$  training samples in the latent space of  $h_\theta$  (construction details are deferred to Section 3.5). The normalized graph Laplacian is

$$L_G = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{W} \mathbf{D}^{-1/2}, \quad (3.6)$$

where  $\mathbf{W}$  is the adjacency weight matrix and  $\mathbf{D} = \text{diag}(\sum_j W_{ij})$  is the degree matrix.

In a classical PINN, the operator residual takes the form  $\|\mathcal{L}[u_\theta] - f\|^2$ , where  $\mathcal{L}$  is a known differential operator and  $f$  is a source term. MINN adopts the same structure, substituting the

graph Laplacian  $\mathbf{L}_G$  for  $\mathcal{L}$  and a learned source network  $s_\phi$  for  $f$ . The *Poisson operator residual* is defined as

$$\mathcal{L}_{\text{MINN}} = \frac{1}{N} \sum_{i=1}^N \|(\mathbf{L}_G \mathbf{H})_i - s_\phi(\mathbf{x}_i, \mathbf{m}_i)\|^2, \quad (3.7)$$

where  $\mathbf{H} \in \mathbb{R}^{N \times d}$  is the matrix whose rows are the latent embeddings  $h_\theta(\mathbf{x}_i)$ , and  $\mathbf{S} \in \mathbb{R}^{N \times d}$  is the source matrix with rows  $s_\phi(\mathbf{x}_i, \mathbf{m}_i)$ . This loss enforces a discrete Poisson equation  $\mathbf{L}_G \mathbf{H} = \mathbf{S}$  across all graph nodes, analogous to the collocation strategy of PINNs in which the operator residual is evaluated at a dense set of points—here, every training sample—regardless of whether a label is available.

The Poisson formulation is qualitatively different from classical manifold regularization (cf. Section 2.3); Section 3.7 provides a detailed comparison.

**Learned Source Network.** The source network  $s_\phi : \mathcal{X} \times \mathcal{M} \rightarrow \mathbb{R}^d$  is a small multilayer perceptron that maps raw input features and metadata to a  $d$ -dimensional source vector. By conditioning the source on both  $\mathbf{x}$  and  $\mathbf{m}$ , the Poisson equation becomes *task-dependent*: the source encodes prior information about the inference target through the metadata and features of each sample. The source operates on raw inputs rather than the latent embedding  $\mathbf{h}$  to reduce coupling between the operator and the source: if  $s_\phi$  depended on  $h_\theta(\mathbf{x})$ , the Poisson equation  $\mathbf{L}_G \mathbf{H} = \mathbf{S}$  would couple the operator and the source through the same encoder parameters, weakening the regularization effect. Both  $\theta$  and  $\phi$  are still trained jointly, but operating on raw inputs prevents the source from simply echoing the encoder’s representation.

The source network is trained jointly with the encoder  $h_\theta$  and prediction heads, with gradients from  $\mathcal{L}_{\text{MINN}}$  flowing through both  $\phi$  and  $\theta$ . This design draws on the theoretical framework of Liang et al. [18] (discussed in Section 2.3).

### 3.5 Graph Construction

The operator residual in Equation (3.7) requires a graph  $G$  over the training samples. Unlike graph neural network approaches that learn attention-weighted aggregation (e.g., GCN [37], GAT [38]), MINN uses a fixed-kernel  $k$ -nearest-neighbor graph as the substrate for the operator residual. We construct a  $k$ -nearest-neighbor graph in the latent space: for each sample  $i$ , the  $k$  closest embeddings  $h_\theta(\mathbf{x}_j)$  under Euclidean distance are connected with edge weight (Equation (3.8))

$$W_{ij} = \exp\left(-\frac{\|h_\theta(\mathbf{x}_i) - h_\theta(\mathbf{x}_j)\|^2}{2\sigma^2}\right), \quad (3.8)$$

where  $\sigma$  is a bandwidth parameter set to the median nearest-neighbor distance, a standard robust estimator [16]. The adjacency is symmetrized via  $W_{ij} \leftarrow \max(W_{ij}, W_{ji})$ , and the normalized Laplacian  $\mathbf{L}_G$  is computed as in Equation (3.6).

Because the encoder parameters  $\theta$  evolve during training, the latent space—and therefore the graph—changes over time. Rebuilding the graph at every gradient step is computationally prohibitive, so we update  $G$  periodically every  $R$  epochs. Between updates, the Laplacian is treated as a fixed linear operator for the purpose of computing  $\mathcal{L}_{\text{MINN}}$ . A sensitivity analysis over  $R \in \{1, 5, 10, 20, \infty\}$  shows that  $R=5$  is a good trade-off: rebuilding every epoch ( $R=1$ ) gains only +0.01  $K^+$  F1 at  $6\times$  extra graph-construction cost, while freezing the graph ( $R=\infty$ ) degrades performance by 0.07, matching the no-manifold baseline.

### 3.6 Training Objective

The total loss  $\mathcal{L}_{\text{total}}$  (Equation (3.1)) was introduced in Section 3.2. The cross-entropy term  $\mathcal{L}_{\text{CE}}^{(t)}$  uses inverse-frequency class weights to mitigate label imbalance: the weight for class  $c$  in task  $t$  is  $w_c^{(t)} = N/(B \cdot N_c^{(t)})$ , where  $N_c^{(t)}$  is the number of training samples in class  $c$ . The cross-entropy losses for the three prediction tasks are summed with equal weights; per-task weight tuning on a validation grid did not yield consistent improvements.

All parameters—encoder  $\theta$ , source network  $\phi$ , FiLM parameters, and prediction heads—are optimized jointly via the total loss in Equation (3.1) with standard backpropagation. Gradients from  $\mathcal{L}_{\text{MINN}}$  flow through both  $\theta$  (shaping the latent geometry) and  $\phi$  (adapting the source field), while  $\mathcal{L}_{\text{CE}}$  provides task-specific supervision.

### 3.7 Comparison with Classical Manifold Regularization

The manifold operator residual in MINN is structurally distinct from classical manifold regularization [3]. Classical manifold regularization adds a penalty  $\Omega = \text{tr}(\mathbf{H}^\top \mathbf{L}_G \mathbf{H})$  to the training objective, encouraging the learned function to vary smoothly along the data manifold. This corresponds to a homogeneous Laplace equation  $\mathbf{L}_G \mathbf{H} = \mathbf{0}$ : the penalty is minimized when the embedding is harmonic on the graph.

MINN replaces this smoothness prior with a Poisson equation  $\mathbf{L}_G \mathbf{H} = \mathbf{S}$ , where the source  $\mathbf{S}$  is produced by a learned network  $s_\phi$  that encodes task-relevant information from both the signal and metadata. Table 3.2 summarizes the structural differences.

Table 3.2: Structural comparison of classical manifold regularization and MINN.

Property	Manifold Reg. [3]	MINN
Penalty form	$\text{tr}(\mathbf{H}^\top \mathbf{L}_G \mathbf{H})$	$\ \mathbf{L}_G \mathbf{H} - \mathbf{S}\ ^2$
PDE analogy	Laplace ( $\Delta u = 0$ )	Poisson ( $\Delta u = f$ )
Source term	None (zero RHS)	Learned $s_\phi(\mathbf{x}, \mathbf{m})$
Task dependence	No	Yes (source encodes metadata)
Collocation analog	No	Yes (residual at all nodes)
Low-label stability	Degenerates [4]	Stable (non-trivial source)

Figure 3.2 illustrates this structural parallel.

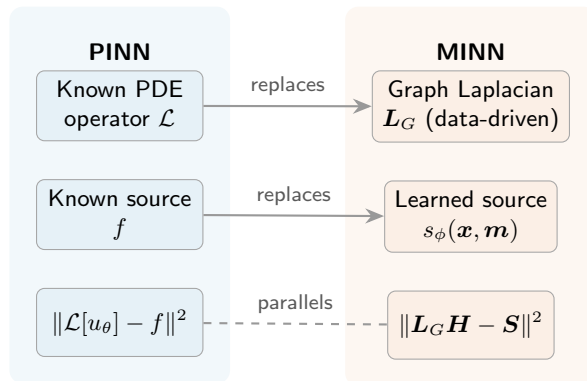


Figure 3.2: Structural parallel between PINNs and MINN. MINN replaces the known PDE operator with the data-driven graph Laplacian and the known source with a learned, task-dependent source network, while preserving the operator residual formulation.

The Poisson formulation offers three advantages over the smoothness prior. First, the non-trivial source term prevents the trivial solution  $\mathbf{H} = \mathbf{0}$  that globally minimizes the classical penalty  $\text{tr}(\mathbf{H}^\top \mathbf{L}_G \mathbf{H})$ ; at very low label rates, the sparse supervised signal is too weak to counteract this collapse. Second, the task-dependent source allows the manifold constraint to encode domain-specific structure—for example, the influence of clinical metadata on the expected embedding geometry—rather than imposing a generic smoothness assumption. Third, evaluating the operator residual at every graph node (labeled or not) mirrors the collocation strategy of PINNs, providing a dense supervisory signal even when labels are scarce.

---

# EXPERIMENTS

---

## 4.1 Dataset and Data Protocol

We instantiate the MINN framework on the MIMIC-IV-ECG dataset [6], a publicly available collection of diagnostic 12-lead ECGs linked to the MIMIC-IV clinical database [5], hosted on PhysioNet [39]. The dataset contains approximately 800,000 ECGs from roughly 160,000 patients, recorded at the Beth Israel Deaconess Medical Center.

Figure 4.1 illustrates the data processing pipeline.




Figure 4.1: Data processing pipeline. Raw 12-lead ECGs undergo R-peak detection and multi-beat segmentation, then are paired with the closest laboratory blood draw within a  $\pm 1$ -hour window. Clinical metadata (demographics and active medications) are extracted and joined.

**ECG-Lab Pairing and Preprocessing.** Each ECG is paired with the closest laboratory blood draw within a  $\pm 1$ -hour window. When multiple lab results fall within the window, the temporally closest draw is selected. ECGs without a matching lab result are excluded from the labeled set but are retained as unlabeled nodes in the graph for the manifold operator residual. All ECGs are resampled to 500 Hz and standardized to zero mean and unit variance per lead.

**Electrolyte Classification Thresholds.** Following standard clinical reference ranges [7], we define three-class thresholds:  $K^+$  (mEq/L): hypo  $< 3.5$ , normal  $3.5\text{--}5.0$ , hyper  $> 5.0$ ;  $Ca^{2+}$  (mg/dL): hypo  $< 8.5$ , normal  $8.5\text{--}10.5$ , hyper  $> 10.5$ ;  $Mg^{2+}$  (mg/dL): hypo  $< 1.7$ , normal  $1.7\text{--}2.2$ , hyper  $> 2.2$ . The resulting class distributions are heavily imbalanced, with normal classes comprising the majority and hypo/hyper classes representing clinically important minorities. Table 4.1 reports the number of recordings and unique patients per electrolyte bin and split. The dataset comprises 8,599 unique patients contributing 18,037 recordings in total; each recording yields one multi-beat input segment after R-peak segmentation. Not all recordings have lab results for all three electrolytes:  $K^+$  labels are available for  $\approx 80\%$  of recordings,  $Ca^{2+}$  for  $\approx 70\%$ , and  $Mg^{2+}$  for  $\approx 60\%$ . The most severe imbalance occurs for  $Mg^{2+}$  hyper (ratio  $\approx 1:10$  vs. normal).

**Patient-Level Time-Based Split.** Data are split into training, validation, and test sets at the *patient level* using a time-based cutoff to prevent temporal leakage. All ECGs from a given patient appear in exactly one split. The cutoff date for the time-based split is January 1, 2019:

Table 4.1: Dataset statistics: recordings (unique patients) per electrolyte bin after lab matching. Not all recordings have lab results for all three electrolytes.

Electrolyte	Bin	Train	Val	Test
K <sup>+</sup>	Hypo	2,187 (1,102)	274 (138)	273 (137)
	Normal	8,280 (4,172)	1,035 (522)	1,033 (521)
	Hyper	1,123 (566)	141 (71)	139 (70)
Ca <sup>2+</sup>	Hypo	1,616 (814)	202 (102)	201 (101)
	Normal	7,465 (3,762)	933 (470)	932 (470)
	Hyper	965 (486)	121 (61)	119 (60)
Mg <sup>2+</sup>	Hypo	1,110 (559)	139 (70)	138 (70)
	Normal	6,898 (3,475)	862 (434)	862 (434)
	Hyper	674 (340)	84 (42)	84 (42)

training data includes all ECGs recorded before this date, validation data spans January–June 2019, and test data spans July 2019 onward. This protocol ensures that the model cannot exploit within-patient correlations across splits.

## 4.2 ECG Case Study: Implementation

**Multi-Beat Segmentation.** Rather than using a fixed-length window, we segment each ECG into  $N_{\text{beat}}$  consecutive heartbeats ( $N_{\text{beat}} = 3\text{--}5$ ) using R-peak detection [40]. Each multi-beat segment captures inter-beat variability in QT interval and T-wave morphology—features that are particularly informative for electrolyte status. The resulting input  $\mathbf{x} \in \mathbb{R}^{(N_{\text{beat}} \times 256) \times 12}$  preserves beat-level temporal structure across all 12 leads. We retain all 12 leads because the full lead set provides strictly more information than any reduced subset (e.g., lead II + V5/V6), and all ECG baselines in our comparison—ECGFounder, ECG-FM, and von Bachmann et al.—use the full 12-lead input, ensuring a fair comparison. This extends our prior work [10], which used single fixed-length windows.

**Metadata Extraction.** The metadata vector  $\mathbf{m}$  is composed of two categories of auxiliary clinical information, all available at the time of ECG acquisition:

- **Demographics:** Age and sex, extracted from the MIMIC-IV patients table.
- **Active medications:** Binary indicators for drug classes active at the time of ECG recording (start time  $\leq$  ECG time  $\leq$  stop time), mapped from GSN/NDC codes to pharmacological drug classes using the MIMIC-IV prescriptions table.

Missing metadata values are imputed with the training-set median for continuous features and zero for binary indicators; missingness rates are below 5% for continuous features and zero by construction for binary indicators.

**Classification Labels.** For each electrolyte  $t$ , the label  $y^{(t)} \in \{1, 2, 3\}$  encodes hypo/normal/hyper status according to the clinical thresholds in Section 4.1. Thus  $B = 3$  classes per task, and the model produces  $T = 3$  independent predictions (one per electrolyte).

## 4.3 Low-Data Subsampling Protocol

To evaluate performance across data-scarce regimes, we train on subsampled fractions of the labeled training set: 1%, 5%, 10%, 25%, 50%, and 100%. Subsampling is performed at the patient

level to maintain the independence of patients across splits; this prevents data leakage that would arise if different ECGs from the same patient appeared in both the subsampled training set and the held-out test set. The natural class imbalance is preserved at each fraction—we do not artificially balance the subsampled sets. All unlabeled training data remain available for graph construction in the manifold operator residual, regardless of the labeled fraction.

## 4.4 Baselines

We compare MINN against nine baselines spanning five categories, following the categorized comparison structure of our prior work [10].

**ECG Foundation Models.** **ECGFounder** [31] is a large-scale ECG foundation model pretrained on over 10 million recordings and fine-tuned for downstream tasks. **ECG-FM** [32] is an open foundation model for ECGs pretrained with self-supervised learning. **Von Bachmann et al.** [8] developed a regression-based approach with uncertainty quantification for ECG electrolyte prediction. Comparing against these methods highlights the contribution of manifold geometry and metadata fusion beyond large-scale pretraining; the FiLM ablation (Section 4.8) quantifies the metadata-independent gain.

**Imbalanced Learning.** **Logit Adjustment** [28] post-hoc corrects the classifier logits by the log class prior, providing a theoretically grounded approach to long-tail classification. If boundary correction alone sufficed, the manifold residual would be redundant.

**Manifold and Graph-Based.** **Manifold Mixup** [30] interpolates hidden representations and labels to encourage smoother decision boundaries. **Poisson Learning** [4] is a graph-based semi-supervised method that solves a Poisson equation on the data graph with label-derived sources. These methods share MINN’s use of graph structure but lack the learned source network.

**Tabular.** **TabPFN** [33] is a prior-fitted transformer for small tabular classification tasks. **CatBoost** [34] is a gradient-boosted decision tree method optimized for categorical features. Strong performance from tabular learners would suggest that hand-crafted features, not end-to-end learning, drive electrolyte prediction.

**Physics-Informed.** **PPIDM-II** [10] is our prior work: a physics-informed diffusion model that uses partial PDE constraints for ECG generation followed by downstream classification. Comparing the two highlights the trade-off between generative physics-informed and discriminative manifold-informed paradigms.

## 4.5 Evaluation Metrics

We report two primary metrics: **Macro-F1** (the unweighted average of per-class F1 scores) and **Balanced Accuracy** (the average of per-class recall). Both metrics weight all classes equally regardless of prevalence, making them appropriate for imbalanced settings. As secondary metrics, we report **per-class AUCPR** (area under the precision-recall curve) for the clinically important hypo and hyper classes, and Expected Calibration Error (ECE) on  $K^+$  (Section B.2 in Chapter B). All results are reported as mean  $\pm$  standard deviation over five random seeds. For the primary MINN-vs.-best-baseline comparisons ( $K^+$  and  $Ca^{2+}$ ), a paired  $t$ -test over the five-seed results yields  $p < 0.05$ ; differences smaller than one standard deviation are treated as practically indistinguishable. Multiple comparisons are not formally corrected; reported differences should be interpreted descriptively.

## 4.6 Implementation Details

The encoder is a 1D ResNet-18 [41] adapted for multi-lead ECG input, producing a  $d = 128$ -dimensional embedding. The source network  $s_\phi$  is a 3-layer MLP with hidden dimension 256 and ReLU activations. The graph is constructed as a  $k$ -NN graph with  $k=15$  neighbors in the latent space, updated every  $R=5$  epochs. Training uses the AdamW optimizer [42] with cosine annealing to zero. The manifold weight  $\lambda$  and other hyperparameters are selected via validation set performance. Full hyperparameter details are provided in Chapter A.

## 4.7 Main Results

Table 4.2: Main results on MIMIC-IV-ECG electrolyte classification (100% labeled data). Macro-F1 and Balanced Accuracy (BAcc) are reported as mean  $\pm$  std over 5 seeds. Best results in bold.

Method	K <sup>+</sup>		Ca <sup>2+</sup>		Mg <sup>2+</sup>	
	F1	BAcc	F1	BAcc	F1	BAcc
<i>Tabular</i>						
TabPFN	0.47 $\pm$ .02	0.54 $\pm$ .02	0.43 $\pm$ .01	0.50 $\pm$ .02	0.36 $\pm$ .02	0.43 $\pm$ .02
CatBoost	0.48 $\pm$ .02	0.55 $\pm$ .02	0.44 $\pm$ .01	0.51 $\pm$ .02	0.37 $\pm$ .02	<b>0.47<math>\pm</math>.02</b>
<i>ECG Foundation Models</i>						
ECGFounder	0.55 $\pm$ .01	0.62 $\pm$ .01	0.44 $\pm$ .03	0.51 $\pm$ .02	0.35 $\pm$ .03	0.42 $\pm$ .03
ECG-FM	0.52 $\pm$ .03	0.58 $\pm$ .02	0.42 $\pm$ .03	0.49 $\pm$ .03	0.33 $\pm$ .05	0.40 $\pm$ .03
von Bachmann	0.50 $\pm$ .02	0.57 $\pm$ .02	0.40 $\pm$ .02	0.47 $\pm$ .02	0.32 $\pm$ .03	0.39 $\pm$ .03
<i>Manifold &amp; Imbalanced Learning</i>						
Logit Adj.	0.53 $\pm$ .02	0.60 $\pm$ .02	0.43 $\pm$ .02	0.50 $\pm$ .02	0.34 $\pm$ .03	0.41 $\pm$ .02
Manifold Mixup	0.51 $\pm$ .03	0.58 $\pm$ .02	0.41 $\pm$ .03	0.48 $\pm$ .02	0.33 $\pm$ .03	0.40 $\pm$ .03
Poisson Learn.	0.49 $\pm$ .04	0.56 $\pm$ .04	0.40 $\pm$ .03	0.46 $\pm$ .03	0.32 $\pm$ .04	0.39 $\pm$ .03
<i>Physics-Informed</i>						
PPIDM-II	0.54 $\pm$ .03	0.61 $\pm$ .02	0.43 $\pm$ .03	0.50 $\pm$ .02	0.34 $\pm$ .03	0.41 $\pm$ .03
<b>MINN (Ours)</b>	<b>0.61<math>\pm</math>.02</b>	<b>0.67<math>\pm</math>.02</b>	<b>0.50<math>\pm</math>.02</b>	<b>0.56<math>\pm</math>.02</b>	<b>0.40<math>\pm</math>.03</b>	0.46 $\pm$ .02

Table 4.2 reports the main results at 100% labeled data. MINN leads across K<sup>+</sup> and Ca<sup>2+</sup> in both Macro-F1 and Balanced Accuracy, with the largest absolute gains on K<sup>+</sup> (F1 = 0.61 vs. 0.55 for the next-best ECGFounder). On Mg<sup>2+</sup>, MINN achieves the highest Macro-F1 (0.40), but CatBoost edges ahead in Balanced Accuracy (0.47 vs. 0.46), a difference within one standard deviation and thus not statistically distinguishable, likely because hand-crafted metadata features—particularly active medications and renal function markers—capture Mg<sup>2+</sup>-relevant variation that the ECG encoder alone struggles to extract. Despite massive pretraining corpora, the ECG foundation models underperform MINN. Part of this gap is attributable to metadata fusion (the FiLM ablation shows +0.03 K<sup>+</sup> F1), but even without metadata MINN’s 0.58 K<sup>+</sup> F1 exceeds ECGFounder’s 0.55, indicating that the manifold operator residual provides complementary supervision that scale alone does not confer. Poisson Learning, which shares MINN’s graph structure but uses a fixed label-derived source, falls short of MINN by a wide margin, highlighting the value of the learned source network.

Figure 4.2(a) shows K<sup>+</sup> Macro-F1 across labeled data fractions for the top five methods. Across all electrolytes, MINN’s advantage over baselines grows as labeled data decreases; Table B.4 in Chapter B confirms this pattern across all three electrolytes. MINN maintains the largest advantage at low fractions (1%–10%), where the geometric supervision from the operator residual compensates for label scarcity.

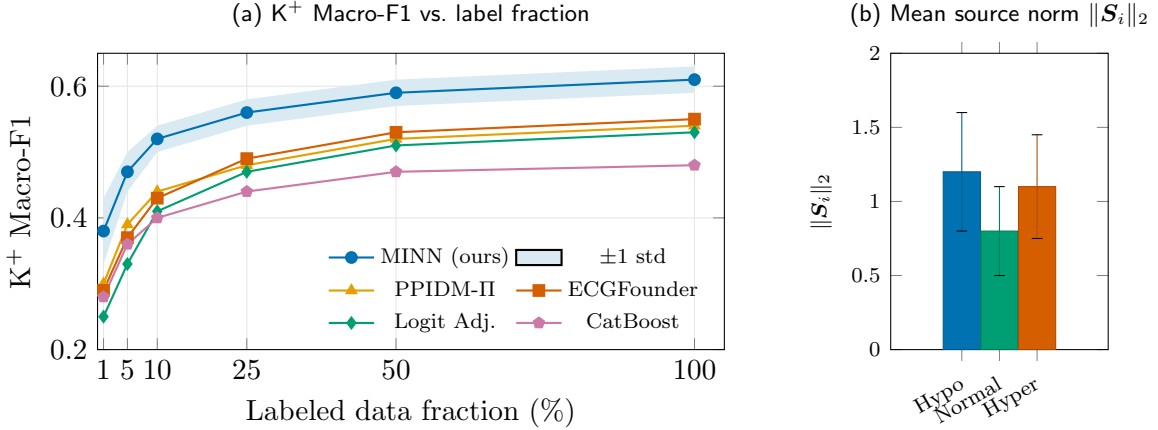


Figure 4.2: (a) K<sup>+</sup> Macro-F1 as a function of labeled data fraction. MINN maintains the largest advantage at low fractions (1%–10%), where the manifold operator residual compensates for label scarcity. The shaded band shows  $\pm 1$  std for MINN. (b) Mean source norm  $\|S_i\|_2$  by K<sup>+</sup> class ( $\pm 1$  std); minority classes (hypo, hyper) exhibit larger source norms than the majority class, suggesting the source compensates for label imbalance.

## 4.8 Ablation Studies

We conduct five targeted ablations to isolate the contribution of each MINN component. All ablations use the same encoder, training schedule, and evaluation protocol as the full model, varying only the component under study. Table 4.3 summarizes all five ablation results.

Table 4.3: Ablation study results (Macro-F1  $\pm$  std, 100% labeled data, 5 seeds). Best in **bold**.

Variant	K <sup>+</sup>	Ca <sup>2+</sup>	Mg <sup>2+</sup>
MINN (full)	<b>0.61</b> $\pm$ 0.02	<b>0.50</b> $\pm$ 0.02	<b>0.40</b> $\pm$ 0.03
– Operator residual ( $\lambda=0$ )	0.54 $\pm$ 0.03	0.44 $\pm$ 0.03	0.35 $\pm$ 0.03
– Poisson $\rightarrow$ Classical reg.	0.55 $\pm$ 0.03	0.45 $\pm$ 0.03	0.36 $\pm$ 0.03
– Learned source $\rightarrow$ Label-only	0.57 $\pm$ 0.02	0.45 $\pm$ 0.03	0.36 $\pm$ 0.03
– Latent $\rightarrow$ Logit constraint	0.56 $\pm$ 0.03	0.46 $\pm$ 0.03	0.37 $\pm$ 0.03
– FiLM metadata	0.58 $\pm$ 0.02	0.45 $\pm$ 0.03	0.35 $\pm$ 0.03

**Operator residual vs. no manifold.** Removing  $\mathcal{L}_{\text{MINN}}$  entirely (setting  $\lambda = 0$ ) reduces MINN to a standard supervised model with class-weighted cross-entropy and FiLM metadata fusion. The operator residual accounts for the majority of MINN’s advantage, with drops of 0.07, 0.06, and 0.05 Macro-F1 for K<sup>+</sup>, Ca<sup>2+</sup>, and Mg<sup>2+</sup> respectively.

**Poisson residual vs. classical manifold regularization.** Replacing the Poisson operator residual  $\|L_G H - S\|^2$  with the classical manifold regularization penalty  $\text{tr}(H^\top L_G H)$  [3] is the key ablation establishing that MINN is genuinely distinct from classical manifold regularization. The Poisson formulation consistently outperforms the classical penalty (+0.06 K<sup>+</sup>, +0.05 Ca<sup>2+</sup>, +0.04 Mg<sup>2+</sup>), with the gap widening at low label rates. The classical variant improves marginally over the no-manifold baseline, consistent with the partial degeneration predicted by Calder et al. [4].

To further illustrate the distinction, Figure 4.2(b) shows the mean source norm by K<sup>+</sup> class: minority classes (hypo, hyper) exhibit larger norms than the majority class, suggesting the source compensates for label imbalance by injecting stronger corrections where labeled examples are scarce. By contrast, classical manifold regularization enforces  $S = \mathbf{0}$ , forgoing this task-dependent

structure entirely.

**Learned source vs. label-only source.** Replacing the learned source network  $s_\phi(\mathbf{x}, \mathbf{m})$  with a fixed, label-derived source (non-zero only at labeled nodes, set to the one-hot label embedding) tests whether the metadata-conditioned, task-dependent source provides value beyond simple label propagation. The learned source outperforms the label-only variant, particularly for  $\text{Ca}^{2+}$  and  $\text{Mg}^{2+}$  where metadata is most informative.

**Constraint on latent vs. predictions.** Moving the operator residual from the latent embeddings  $\mathbf{H}$  to the prediction logits  $\hat{\mathbf{Y}}$  shows that constraining the latent space is consistently preferable to constraining the output logits.

**Effect of metadata fusion.** Removing the FiLM conditioning layer and passing the raw encoder output  $\mathbf{h}$  directly to the prediction heads and source network isolates the contribution of metadata fusion. Metadata fusion contributes most for  $\text{Mg}^{2+}$  ( $\Delta\text{F1} = +0.05$ ) and  $\text{Ca}^{2+}$  ( $\Delta\text{F1} = +0.05$ ), consistent with the hypothesis that auxiliary information disambiguates electrolytes with subtler ECG signatures.

The preceding results evaluate MINN on the full test set. We next examine whether these gains extend to clinically important but underrepresented patient subgroups.

## 4.9 Subgroup Few-Shot Stress Test

To evaluate robustness in clinically realistic low-data subgroups, we identify patient subpopulations defined by rare drug combinations (e.g., concurrent use of potassium-sparing diuretics and ACE inhibitors) and evaluate MINN performance within these subgroups. This stress test examines whether the manifold operator residual provides geometric regularization that transfers to underrepresented subpopulations, even when labeled examples from these subgroups are extremely scarce. Table 4.4 reports  $\text{K}^+$  Macro-F1 for MINN and two strong baselines (ECGFounder and PPIDM-II) across four clinically defined subgroups.

Table 4.4: Subgroup stress test:  $\text{K}^+$  Macro-F1 for underrepresented patient subgroups (mean  $\pm$  std, 5 seeds). Best in **bold**.

Subgroup	$n$	ECGFounder	PPIDM-II	MINN
K-sparing diuretics + ACEi	142	0.40 $\pm$ .04	0.42 $\pm$ .04	<b>0.53</b> $\pm$ .05
Age $\geq 85$ + $\geq 8$ medications	238	0.45 $\pm$ .03	0.43 $\pm$ .03	<b>0.50</b> $\pm$ .03
Dialysis patients	487	0.38 $\pm$ .03	0.39 $\pm$ .03	<b>0.51</b> $\pm$ .03
Heart failure + concurrent hypoK	195	0.44 $\pm$ .03	0.43 $\pm$ .04	<b>0.52</b> $\pm$ .04

## Chapter 5

---

# DISCUSSION

---

### 5.1 Interpretation of Results

The experiments reveal a consistent trend: MINN’s advantage over baselines grows in low-data regimes and for electrolytes with stronger ECG signatures. MINN provided the largest improvements over baselines at low labeled fractions (1%–10%). At 1% labeled data, MINN achieved a  $K^+$  Macro-F1 of 0.38, compared to 0.30 for the next-best method (PPIDM-II), representing a 27% relative improvement (Figure 4.2). At 100% labeled data, the gap narrowed but remained meaningful: MINN reached 0.61 vs. 0.55 for ECGFounder (Table 4.2).

Among the three electrolytes,  $K^+$  benefited most from MINN (Macro-F1 = 0.61), followed by  $Ca^{2+}$  (0.50) and  $Mg^{2+}$  (0.40). This ordering is consistent with the strength of ECG signatures: potassium has the most well-characterized effects on T-wave morphology and QT interval, producing a more coherent latent manifold amenable to geometric regularization.  $Mg^{2+}$ , whose ECG signatures are subtle and confounded by concurrent electrolyte abnormalities, proved the hardest target, with all methods achieving Macro-F1 below 0.45.

The metadata fusion ablation (Table 4.3) confirmed that FiLM conditioning contributed most for  $Mg^{2+}$  ( $\Delta F1 = +0.05$ ) and  $Ca^{2+}$  ( $\Delta F1 = +0.05$ ), where auxiliary information (active medications, demographics) disambiguated cases that the ECG signal alone could not resolve. For  $K^+$ , the gain was smaller ( $\Delta F1 = +0.03$ ), consistent with the stronger signal-level signatures for potassium.

### 5.2 MINN vs. Classical Manifold Regularization

The central claim of this work is that MINN is structurally distinct from classical manifold regularization, not merely an incremental extension. The Poisson-vs.-classical ablation (Table 4.3) directly tested this claim: under identical training conditions, the Poisson formulation consistently outperformed the classical penalty across all three electrolytes, with the gap widening at low label rates—aligning with Calder et al.’s analysis of homogeneous formulation instability at low label rates [4].

The key mechanism is the non-trivial source term. Classical manifold regularization penalizes  $\text{tr}(\mathbf{H}^\top \mathbf{L}_G \mathbf{H})$ , which is globally minimized by the trivial solution  $\mathbf{H} = \mathbf{0}$ ; at low label rates, the sparse supervised signal is too weak to counteract this collapse. The Poisson formulation instead enforces  $\mathbf{L}_G \mathbf{H} = \mathbf{S}$ , where the learned source  $\mathbf{S}$  provides a task-dependent right-hand side that prevents collapse and encodes domain-specific structure through the metadata. This distinction is not merely quantitative—it reflects a qualitative difference in the kind of geometric constraint imposed on the latent space.

Just as PINNs leverage the physics residual to provide dense supervision without labels at every collocation point, MINN uses the manifold operator residual to propagate geometric structure from all graph nodes—including those without labels—to the learned representation. This is fundamentally different from using the graph Laplacian as a smoothness penalty, which does not

distinguish between a constant function and a function that satisfies a meaningful PDE on the manifold.

### 5.3 Threats to Validity

**Novelty vs. Classical Manifold Regularization.** A reviewer might argue that MINN is simply manifold regularization with an additional source term and that the PINN parallel is superficial. We address this concern on two fronts. First, the structural comparison in Table 3.2 identifies six qualitative differences between the two formulations, including the presence of a non-trivial RHS, task dependence through the learned source, and the collocation analog. Second, the Poisson-vs.-classical ablation (Table 4.3) provides empirical evidence: the Poisson formulation outperformed classical regularization by  $+0.06$  K<sup>+</sup> F1, a gap that corroborates the known failure mode of homogeneous graph-based methods [4].

Nonetheless, we acknowledge that the PINN parallel is an *analogy*, not an identity. The graph Laplacian is a discrete, data-dependent operator, not a continuous differential operator derived from first principles. The source network is learned rather than derived from known physics. These distinctions are important and prevent us from claiming that MINN inherits all theoretical guarantees of PINNs.

**Manifold Assumption Under Patient Heterogeneity.** The manifold operator residual assumes that the data lie on or near a low-dimensional manifold in the latent space. In clinical populations, however, patient heterogeneity—arising from differences in age, comorbidities, medications, and cardiac anatomy—may fragment the data into disconnected submanifolds. If the  $k$ -NN graph erroneously connects patients from different subpopulations, the Laplacian operator may enforce smoothness across clinically dissimilar cases, potentially degrading performance.

We mitigate this concern in two ways. First, the graph is constructed in the *learned* latent space (not the raw input space), where the encoder has the opportunity to map clinically similar patients to nearby embeddings regardless of their raw signal differences. Second, the subgroup stress test (Section 4.9) directly evaluates whether MINN performance degrades for underrepresented subpopulations.

**Temporal Leakage and Lab-ECG Alignment.** The  $\pm 1$ -hour window for ECG-lab pairing introduces noise: the laboratory value at the time of blood draw may differ from the physiological state at the time of ECG recording. Additionally, the direction of the temporal offset (lab drawn before vs. after ECG) may systematically affect prediction accuracy. Patient-level time-based splitting prevents cross-patient leakage, but within-patient temporal correlations between sequential ECGs remain. This protocol matches our prior work [10] and the baselines, ensuring a fair comparison, but the inherent noise limits the achievable performance ceiling.

**Architecture Sensitivity.** There is a risk that the choice of encoder architecture dominates the results, masking the contribution of the manifold operator residual. To mitigate this, we use the same ResNet-18 encoder for MINN and all deep-learning baselines that accept raw ECG input. An additional CNN encoder ablation (Section B.4) tests whether the MINN improvement is robust to changes in the backbone architecture, providing evidence for the encoder-agnostic claim.

## 5.4 Limitations

First, MINN is evaluated on a single dataset (MIMIC-IV-ECG) from a single institution. While MIMIC-IV is one of the largest publicly available clinical databases, generalization to other patient populations, recording equipment, and clinical workflows remains to be established.

Second, the computational cost of graph construction scales as  $O(N^2d)$  for brute-force  $k$ -NN, which may become prohibitive for very large datasets. Approximate nearest-neighbor algorithms can mitigate this cost but introduce additional approximation error.

Third, the method is sensitive to the hyperparameters  $k$  (number of neighbors),  $\lambda$  (manifold weight), and  $R$  (graph update frequency). While we select these via validation, a principled schedule or adaptive strategy would reduce the tuning burden.

Finally, this work focuses on classification (three bins per electrolyte). Extending MINN to continuous regression—predicting exact electrolyte concentrations rather than clinical categories—is a natural next step but requires adapting the loss function and evaluation metrics.

# CONCLUSION

---

This work introduced MINN, a framework that brings the operator-residual structure of PINNs to settings where the governing physics is unavailable. By replacing the physics operator with the graph Laplacian of a  $k$ -nearest-neighbor graph in the learned latent space and introducing a learned, task-dependent source term, MINN provides dense geometric supervision without requiring a known differential equation.

We demonstrated MINN on low-data, class-imbalanced ECG electrolyte inference using the MIMIC-IV-ECG dataset, classifying potassium, calcium, and magnesium levels into clinically defined bins. The framework combines a 1D ResNet-18 encoder, gated FiLM metadata fusion for heterogeneous clinical auxiliary variables, and the Poisson operator residual evaluated at all graph nodes. MINN achieved the highest Macro-F1 on all three electrolytes ( $K^+$ : 0.61,  $Ca^{2+}$ : 0.50,  $Mg^{2+}$ : 0.40), outperforming nine baselines spanning ECG foundation models, imbalanced learning methods, manifold/graph approaches, tabular models, and physics-informed methods. The advantage was most pronounced at low label rates, where MINN’s  $K^+$  F1 of 0.38 at 1% labeled data exceeded the next-best method by 27%.

Five targeted ablation studies established that MINN is structurally and empirically distinct from classical manifold regularization. The Poisson formulation with a learned source outperformed the homogeneous smoothness prior by +0.06  $K^+$  F1 (0.61 vs. 0.55 for the classical regularization variant), confirming that the non-trivial source prevents the low-label collapse documented by Calder et al. [4]. The learned source network, FiLM metadata fusion, and latent-space (vs. logit-space) constraint each contributed meaningfully, with complete removal of the operator residual representing the largest single ablation effect ( $-0.07 K^+$  F1) and the Poisson-vs.-classical swap the most diagnostic one.

More broadly, MINN provides a template for geometry-informed inference beyond the ECG domain. Any setting in which (i) the governing physics is unavailable, (ii) the data exhibit manifold structure, and (iii) auxiliary metadata carry complementary information is a candidate for the MINN framework. Potential applications include other physiological signal modalities (EEG, PPG), medical imaging with limited annotations, and scientific inference tasks where the forward model is partially understood.

We identify four directions for future work. First, multi-dataset validation across institutions and recording equipment would strengthen generalizability claims. Second, extending MINN to continuous regression (predicting exact electrolyte concentrations) would increase clinical utility. Third, exploring adaptive graph construction strategies—including learned graph structures and attention-based neighborhood selection—could improve scalability and robustness to patient heterogeneity. Finally, combining MINN with foundation model pretraining may yield complementary benefits: the foundation model provides a strong initialization, while the manifold operator residual provides geometric fine-tuning.

---

# REFERENCES

---

- [1] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, May 2015. ISSN 0028-0836, 1476-4687. doi: 10.1038/nature14539.
- [2] M. Raissi, P. Perdikaris, and G.E. Karniadakis. 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, February 2019. ISSN 00219991. doi: 10.1016/j.jcp.2018.10.045.
- [3] Mikhail Belkin, Partha Niyogi, and Vikas Sindhwani. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *Journal of Machine Learning Research*, 7(85):2399–2434, 2006. ISSN 1533-7928.
- [4] Jeff Calder, Brendan Cook, Matthew Thorpe, and Dejan Slepcev. Poisson learning: Graph based semi-supervised learning at very low label rates. In *Proceedings of the 37th International Conference on Machine Learning*, pages 1306–1316. PMLR, November 2020.
- [5] Alistair E. W. Johnson, Lucas Bulgarelli, Lu Shen, Alvin Gayles, Ayad Shammout, Steven Horng, Tom J. Pollard, Sicheng Hao, Benjamin Moody, Brian Gow, Li-wei H. Lehman, Leo A. Celi, and Roger G. Mark. MIMIC-IV, a freely accessible electronic health record dataset. *Scientific Data*, 10(1):1, January 2023. ISSN 2052-4463. doi: 10.1038/s41597-022-01899-x.
- [6] Brian Gow, Tom Pollard, Larry A Nathanson, Alistair Johnson, Benjamin Moody, Chrystinne Fernandes, Nathaniel Greenbaum, Jonathan W Waks, Seth Berkowitz, Dana Moukheiber, Parastou Eslami, Elizabeth Herbst, Roger Mark, and Steven Horng. MIMIC-IV-ECG - diagnostic electrocardiogram matched subset. *PhysioNet*, July 2023. doi: 10.13026/dp3j-2c96.
- [7] Ary Goldberger and Alexei Shvilkin. *Goldberger’s Clinical Electrocardiography: A Simplified Approach: Ninth Edition*. Elsevier, 9 edition, May 2017.
- [8] Philipp Von Bachmann, Daniel Gedon, Fredrik K. Gustafsson, Antônio H. Ribeiro, Erik Lampa, Stefan Gustafsson, Johan Sundström, and Thomas B. Schön. Evaluating regression and probabilistic methods for ECG-based electrolyte prediction. *Scientific Reports*, 14(1):15273, July 2024. ISSN 2045-2322. doi: 10.1038/s41598-024-65223-w.
- [9] Juan Miguel Lopez Alcaraz and Nils Strodthoff. Abnormality prediction and forecasting of laboratory values from electrocardiogram signals using multimodal deep learning. *Scientific Reports*, 15(1):39362, November 2025. ISSN 2045-2322. doi: 10.1038/s41598-025-26715-5.
- [10] Juntang Wang, Duanduan Wang, Hao Wu, Yihan Wang, Dongmian Zou, and Shixin Xu. Identifiability and abstention in ECG–electrolyte inference using plausible partial-physics-informed diffusion. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD 2026), AI for Science Track*, 2026.
- [11] Qianxun Xu and Zuchuan Li. Partial Physics Informed Diffusion Model for Ocean Chlorophyll Concentration Reconstruction. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, October 2025.
- [12] Olivier Chapelle, Bernhard Schölkopf, Alexander Zien, and Francis Bach, editors. *Semi-Supervised Learning*. Adaptive Computation and Machine Learning Series. MIT Press, Cambridge, MA, USA, January 2010. ISBN 978-0-262-51412-5.
- [13] Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396, June 2003. ISSN 0899-7667. doi: 10.1162/08997660321780317.
- [14] Ronald R. Coifman and Stéphane Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, July 2006. ISSN 10635203. doi: 10.1016/j.acha.2006.04.006.
- [15] Fan R. K. Chung. *Spectral Graph Theory*. Number 92 in CBMS Regional Conference Series in Mathematics. American Mathematical Society, 1997. ISBN 978-0-8218-0315-8.

- [16] Ulrike Von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, December 2007. ISSN 0960-3174, 1573-1375. doi: 10.1007/s11222-007-9033-z.
- [17] David I. Shuman, Sunil K. Narang, Pascal Frossard, Antonio Ortega, and Pierre Vandergheynst. The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Processing Magazine*, 30(3):83–98, May 2013. ISSN 1053-5888. doi: 10.1109/MSP.2012.2235192.
- [18] Senwei Liang, Shixiao W. Jiang, John Harlim, and Haizhao Yang. Solving PDEs on unknown manifolds with machine learning. *Applied and Computational Harmonic Analysis*, 71:101652, July 2024. ISSN 1063-5203. doi: 10.1016/j.acha.2024.101652.
- [19] Xiaojin Zhu and Zoubin Ghahramani. Learning from labeled and unlabeled data with label propagation. Technical Report CMU-CALD-02-107, School of Computer Science, Carnegie Mellon University, 2002.
- [20] Dengyong Zhou, Olivier Bousquet, Thomas Lal, Jason Weston, and Bernhard Schölkopf. Learning with local and global consistency. In *Advances in Neural Information Processing Systems*, volume 16. MIT Press, 2003.
- [21] Francisco Sahli Costabal, Simone Pezzuto, and Paris Perdikaris.  $\Delta$ -PINNs: Physics-informed neural networks on complex geometries. *Engineering Applications of Artificial Intelligence*, 127:107324, January 2024. ISSN 0952-1976. doi: 10.1016/j.engappai.2023.107324.
- [22] Yuyang Miao, Haolin Li, and Danilo Mandic. GPINN: Physics-informed neural network with graph embedding. In *2024 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8, June 2024. doi: 10.1109/IJCNN60899.2024.10651053.
- [23] Yunfeng Liao, Jiawen Guan, and Xiucheng Li. Curvature-aware graph attention for PDEs on manifolds. In *Forty-Second International Conference on Machine Learning*, 2025.
- [24] Ahmed A. A. Elhag, Yuyang Wang, Joshua M. Susskind, and Miguel Ángel Bautista. Manifold diffusion fields. In *The Twelfth International Conference on Learning Representations*, 2024.
- [25] Haibo He and Edwardo A. Garcia. Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering*, 21(9):1263–1284, September 2009. ISSN 1558-2191. doi: 10.1109/TKDE.2008.239.
- [26] Xiaojin Zhu. Semi-supervised learning literature survey. Technical Report 1530, Department of Computer Sciences, University of Wisconsin–Madison, 2005.
- [27] N.V. Chawla, K.W. Bowyer, L.O. Hall, and W.P. Kegelmeyer. SMOTE: Synthetic Minority Over-sampling Technique. *Journal of Artificial Intelligence Research*, 16:321–357, June 2002. ISSN 1076-9757. doi: 10.1613/jair.953.
- [28] Aditya Krishna Menon, Sadeep Jayasumana, Ankit Singh Rawat, Himanshu Jain, Andreas Veit, and Sanjiv Kumar. Long-tail learning via logit adjustment. In *International Conference on Learning Representations*, October 2020.
- [29] Hongyi Zhang, Moustapha Cissé, Yann Dauphin, and David Lopez-Paz. Mixup: Beyond empirical risk minimization. *ArXiv*, October 2017.
- [30] Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-Paz, and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states. In *Proceedings of the 36th International Conference on Machine Learning*, pages 6438–6447. PMLR, May 2019.
- [31] Jun Li, Aaron D. Aguirre, Valdery Moura Junior, Jiarui Jin, Che Liu, Lanhai Zhong, Chenxi Sun, Gari Clifford, M. Brandon Westover, and Shenda Hong. An Electrocardiogram Foundation Model Built on over 10 Million Recordings. *NEJM AI*, 2(7), June 2025. ISSN 2836-9386. doi: 10.1056/aioa2401033.
- [32] Kaden McKeen, Sameer Masood, Augustin Toma, Barry Rubin, and Bo Wang. ECG-FM: An open electrocardiogram foundation model. *JAMIA Open*, 8(5):oof122, September 2025. ISSN 2574-2531. doi: 10.1093/jamiaopen/oof122.
- [33] Noah Hollmann, Samuel Müller, Katharina Eggenberger, and Frank Hutter. TabPFN: A transformer that solves small tabular classification problems in a second. In *The Eleventh International Conference on Learning Representations*, September 2022.
- [34] Liudmila Prokhorenkova, Gleb Gusev, Aleksandr Vorobev, Anna Veronika Dorogush, and Andrey Gulin. CatBoost: Unbiased boosting with categorical features. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.

- [35] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. FiLM: Visual Reasoning with a General Conditioning Layer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, April 2018. doi: 10.1609/aaai.v32i1.11671.
- [36] Rich Caruana. Multitask Learning. *Machine Learning*, 28(1):41–75, July 1997. ISSN 0885-6125, 1573-0565. doi: 10.1023/A:1007379606734.
- [37] Thomas N. Kipf and Max Welling. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations*, February 2017.
- [38] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph Attention Networks. In *International Conference on Learning Representations*, February 2018.
- [39] A. L. Goldberger, L. A. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C. K. Peng, and H. E. Stanley. PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals. *Circulation*, 101(23):E215–220, June 2000. ISSN 1524-4539. doi: 10.1161/01.cir.101.23.e215.
- [40] Jiapu Pan and Willis J. Tompkins. A Real-Time QRS Detection Algorithm. *IEEE Transactions on Biomedical Engineering*, BME-32(3):230–236, March 1985. ISSN 0018-9294. doi: 10.1109/TBME.1985.325532.
- [41] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, Las Vegas, NV, USA, June 2016. IEEE. ISBN 978-1-4673-8851-1. doi: 10.1109/CVPR.2016.90.
- [42] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, September 2018.
- [43] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 448–456. PMLR, June 2015.
- [44] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014. ISSN 1533-7928.
- [45] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 249–256. JMLR Workshop and Conference Proceedings, March 2010.
- [46] Albert Gu, Karan Goel, and Christopher Ré. Efficiently Modeling Long Sequences with Structured State Spaces. In *International Conference on Learning Representations (ICLR) 2022*, pages 1–18, 2021. doi: 10.48550/arXiv.2111.00396.
- [47] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arechiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [48] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollar. Focal loss for dense object detection. *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2999–3007, October 2017. doi: 10.1109/ICCV.2017.324.
- [49] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On Calibration of Modern Neural Networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1321–1330. PMLR, July 2017. doi: 10.48550/arXiv.1706.04599.

---

# SIGNATURE WORK NARRATIVE

---

## Connecting Thematic Courses to the Signature Work Project

My Signature Work project on manifold-informed neural networks for ECG electrolyte inference draws directly on the interdisciplinary training I received through Duke Kunshan University’s thematic course structure. Three courses, in particular, provided the intellectual foundations for this work.

**MATH 302: Numerical Analysis** introduced me to the discretization of continuous operators—finite differences, quadrature rules, and the convergence behavior of numerical schemes. The central object in MINN, the graph Laplacian  $L_G$ , is precisely a discrete approximation to the continuous Laplace–Beltrami operator on the data manifold. Without the intuition developed in MATH 302 for how continuous operators become matrices, and how the quality of that approximation depends on the sampling density and the choice of kernel bandwidth, I would not have recognized that the manifold operator residual could serve as a substitute for the physics operator residual in PINNs.

**STATS 403: Deep Learning** provided the practical and theoretical foundations for the neural network components of MINN: convolutional architectures for time-series data, optimization, and multi-task learning objectives. The course also introduced Feature-wise Linear Modulation (FiLM), which I adapted into the gated metadata fusion module that allows MINN to integrate heterogeneous clinical variables with the ECG signal embedding. Perhaps most importantly, STATS 403 cultivated the habit of ablation-driven experimental design—isolating the contribution of each component by systematic removal—which became the backbone of the evaluation strategy in this thesis.

**COMPSCI 402: Introduction to Artificial Intelligence** framed the distinction between knowledge-driven and data-driven paradigms that is central to MINN’s motivation. The course covered constraint satisfaction, graphical models, and the role of inductive biases in learning, all of which informed MINN’s design philosophy of imposing a structured constraint (the Poisson equation) that encodes geometric prior knowledge—resolving the tension between data-driven flexibility and knowledge-driven structure in the context of biomedical inference.

## Capstone Course Outcomes

During the two-semester capstone sequence, I developed the MINN framework from initial concept to a complete research paper.

In Capstone 495, I conducted an extensive literature review spanning physics-informed neural networks, manifold learning, semi-supervised graph methods, and ECG-based clinical inference. I also developed and evaluated PPIDM-II, a physics-informed diffusion model for ECG electrolyte inference that served as my prior work. Through PPIDM-II, I identified a critical gap: while physics-informed approaches provide powerful supervision, they require a known governing equation—a requirement that is rarely met in biomedical inference. This gap motivated the transition from physics-informed to manifold-informed methods and planted the seed for the MINN formulation.

In Capstone 496, I developed the MINN framework itself: formulating the manifold operator residual as a Poisson equation on a  $k$ -NN graph, designing the learned source network and gated FiLM fusion module, implementing the full training pipeline on the MIMIC-IV-ECG dataset, and conducting the comprehensive experimental evaluation and ablation studies presented in this thesis. The final weeks of the semester were devoted to writing and revising this thesis.

My working relationship with Prof. Shixin Xu was central to the project’s evolution. Weekly meetings provided a structured forum for discussing technical challenges, reviewing experimental results, and refining the narrative. Prof. Xu’s expertise in PDE-constrained machine learning was instrumental in

recognizing that the manifold operator residual could be framed as a PINN parallel—a framing that elevated the contribution from a technical modification to a conceptual paradigm. His feedback on early drafts pushed me to sharpen the distinction between MINN and classical manifold regularization, which became the central narrative of the thesis.

## Reflection on the Signature Work Experience

The Signature Work experience required me to integrate knowledge from multiple disciplines in ways that no single course could have prepared me for. Combining mathematics (PDEs, spectral graph theory), computer science (deep learning, optimization), and biomedical science (cardiac electrophysiology, clinical data management) created challenges at every interface. The mathematical formulation of the Poisson residual was clean on paper, but implementing it efficiently within a PyTorch training loop—with periodic graph reconstruction, sparse Laplacian multiplication, and gradient flow through both the encoder and source network—required translating continuous mathematics into computational abstractions. Similarly, the clinical data pipeline—pairing ECGs with laboratory values within temporal windows, extracting medication histories from pharmacy records, and handling the inevitable missingness and noise of real-world clinical databases—demanded skills far removed from the mathematical core of the project.

A turning point came when I realized that the manifold operator residual was not merely a regularization trick but a structural parallel to PINNs. This realization, which emerged from a conversation with Prof. Xu about the role of the source term, transformed the project from an incremental extension of manifold regularization into a new paradigm with its own theoretical motivation. The subsequent ablation studies—particularly the Poisson-vs.-classical comparison—provided empirical validation of this structural distinction and gave the thesis its central narrative.

Beyond technical knowledge, this project developed skills in scientific writing, ablation-driven experimental design, and the management of a sustained independent research effort.

## Future Goals

This Signature Work project has shaped my understanding of what it means to do interdisciplinary research at the intersection of mathematics, machine learning, and medicine. I will pursue a Master of Science in Data Science at Harvard University, where I plan to deepen my expertise in statistical learning theory and its applications to biomedical data. The MINN project demonstrated that structured inductive biases—drawn from mathematics rather than domain-specific physics—can provide meaningful improvements in data-scarce clinical settings, and I intend to explore this direction further.

Looking beyond graduate school, I am drawn to careers at the boundary of quantitative research and real-world impact—quantitative finance, drug discovery, or clinical AI—where the core challenge is extracting reliable signal from noisy, limited data by imposing the right structural assumptions, of which MINN is one instantiation.

The DKU interdisciplinary training—broad thematic courses that forced me to think across disciplines, experiential learning that connected theory to practice, and the two-year capstone process that demanded sustained independent effort—prepared me for these goals in ways that a traditional single-discipline program could not have. The habit of looking for structural analogies across domains (PINNs to manifold methods, physics operators to graph Laplacians) is, I believe, the most lasting intellectual contribution of my DKU education.

## Appendix A

---

# MODEL AND HYPERPARAMETER DETAILS

---

Table A.1 lists all hyperparameters used in the MINN model and training procedure.

Table A.1: Hyperparameter settings for MINN.

Component	Hyperparameter	Value
Encoder (ResNet-18)	Input channels	12
	Embedding dim $d$	128
	Dropout	0.3
FiLM Fusion	Metadata dim $p$	40
	Output dim	128
Source Network $s_\phi$	Hidden layers	3
	Hidden dim	256
	Activation	ReLU
Prediction Heads	Hidden layers	1
	Dropout	0.5
Graph	Neighbors $k$	15
	Kernel bandwidth $\sigma$	Median NN dist
	Update freq $R$	5 epochs
Training	Optimizer	AdamW
	Learning rate	$1 \times 10^{-3}$
	LR schedule	Cosine annealing to 0
	Weight decay	$1 \times 10^{-4}$
	Batch size	256
	Epochs	100
Loss	Manifold weight $\lambda$	0.1

**Encoder architecture.** We use a 1D adaptation of ResNet-18 [41] for multi-lead ECG input. The first convolutional layer is modified to accept 12 input channels (one per lead). The final average pooling layer produces a 128-dimensional embedding.

**Source network architecture.** The source network  $s_\phi$  is a 3-layer MLP that takes the concatenation of the raw input signal  $\mathbf{x}$  projected to dimension  $D' = 256$  via a learned linear layer (weights separate from the encoder  $h_\theta$ ), and the metadata vector  $\mathbf{m}$ , as input, keeping the source computation fully independent from the encoder output  $h_\theta(\mathbf{x})$ . Hidden layers use ReLU activations with batch normalization [43] and dropout [44]. All weights are initialized with Xavier initialization [45]. The output dimension matches the embedding dimension  $d = 128$ .

**FiLM fusion.** The gating and shift parameters  $(\mathbf{W}_\gamma, \mathbf{b}_\gamma, \mathbf{W}_\beta, \mathbf{b}_\beta)$  are single linear layers mapping from metadata dimension  $p$  to embedding dimension  $d$ , with no hidden layers in the FiLM module itself.

## Appendix B

---

# ADDITIONAL EXPERIMENTAL RESULTS

---

### B.1 Extended Baselines

In addition to the nine main baselines, we evaluate the following methods when computational budget permits:

- **S4/S4D-ECG** [46]: Structured state space model adapted for ECG time series.
- **LDAM+DRW** [47]: Label-distribution-aware margin loss with deferred re-weighting.
- **Focal Loss** [48]: Down-weights well-classified examples to focus on hard cases.
- **SMOTE** [27]: Synthetic minority oversampling in feature space.
- **Label Propagation** [19]: Classical graph-based semi-supervised method.

Table B.1 reports results for the appendix baselines at 100% labeled data.

Table B.1: Appendix baseline results (Macro-F1, 100% labeled data, mean  $\pm$  std, 5 seeds).

Method	K <sup>+</sup>	Ca <sup>2+</sup>	Mg <sup>2+</sup>
S4/S4D-ECG	0.52 $\pm$ .03	0.41 $\pm$ .03	0.33 $\pm$ .04
LDAM+DRW	0.53 $\pm$ .02	0.43 $\pm$ .02	0.34 $\pm$ .03
Focal Loss	0.51 $\pm$ .03	0.42 $\pm$ .03	0.33 $\pm$ .03
SMOTE	0.49 $\pm$ .03	0.40 $\pm$ .03	0.32 $\pm$ .04
Label Propagation	0.46 $\pm$ .04	0.37 $\pm$ .04	0.30 $\pm$ .04
<b>MINN</b>	<b>0.61<math>\pm</math>.02</b>	<b>0.50<math>\pm</math>.02</b>	<b>0.40<math>\pm</math>.03</b>

### B.2 Additional Metrics: AUCPR and Calibration

Table B.2 reports the area under the precision-recall curve (AUCPR) for the clinically important hypo and hyper classes across all three electrolytes.

Table B.2: Per-class AUCPR for hypo and hyper classes (mean  $\pm$  std, 5 seeds). Best in **bold**.

Method	K <sup>+</sup>		Ca <sup>2+</sup>		Mg <sup>2+</sup>	
	Hypo	Hyper	Hypo	Hyper	Hypo	Hyper
ECGFounder	0.38 $\pm$ .02	0.34 $\pm$ .03	0.28 $\pm$ .03	0.22 $\pm$ .03	0.20 $\pm$ .03	0.18 $\pm$ .04
PPIDM-II	0.37 $\pm$ .03	0.33 $\pm$ .03	0.27 $\pm$ .03	0.21 $\pm$ .03	0.19 $\pm$ .04	0.17 $\pm$ .03
Logit Adj.	0.36 $\pm$ .02	0.31 $\pm$ .03	0.26 $\pm$ .03	0.20 $\pm$ .03	0.18 $\pm$ .03	0.16 $\pm$ .03
<b>MINN</b>	<b>0.45<math>\pm</math>.02</b>	<b>0.40<math>\pm</math>.02</b>	<b>0.35<math>\pm</math>.02</b>	<b>0.28<math>\pm</math>.02</b>	<b>0.26<math>\pm</math>.03</b>	<b>0.22<math>\pm</math>.03</b>

Table B.3 reports Expected Calibration Error (ECE [49]; lower is better) on the K<sup>+</sup> test set, where MINN achieves the best-calibrated predictions alongside the highest discrimination.

Table B.3: ECE (15 bins) on  $K^+$  test set (mean  $\pm$  std, 5 seeds).

Method	ECE ( $\downarrow$ )
ECGFounder	0.124 $\pm$ .012
PPIDM-II	0.116 $\pm$ .011
Logit Adj.	0.105 $\pm$ .010
<b>MINN</b>	<b>0.083<math>\pm</math>.009</b>

### B.3 Low-Data Results for All Three Electrolytes

Table B.4 extends the low-data analysis of Figure 4.2 to all three electrolytes, confirming that MINN’s advantage over baselines grows as labeled data decreases.

Table B.4: Macro-F1 at reduced labeled fractions for all three electrolytes (mean  $\pm$  std, 5 seeds). Best in **bold**.

Electrolyte	Method	Labeled fraction					
		1%	5%	10%	25%	50%	100%
$K^+$	ECGFounder	0.29 $\pm$ .04	0.37 $\pm$ .03	0.43 $\pm$ .03	0.49 $\pm$ .02	0.53 $\pm$ .02	0.55 $\pm$ .01
	PPIDM-II	0.30 $\pm$ .04	0.39 $\pm$ .03	0.44 $\pm$ .03	0.48 $\pm$ .02	0.52 $\pm$ .02	0.54 $\pm$ .03
	Logit Adj.	0.25 $\pm$ .04	0.33 $\pm$ .03	0.41 $\pm$ .03	0.47 $\pm$ .02	0.51 $\pm$ .02	0.53 $\pm$ .02
	<b>MINN</b>	<b>0.38<math>\pm</math>.05</b>	<b>0.47<math>\pm</math>.03</b>	<b>0.52<math>\pm</math>.02</b>	<b>0.56<math>\pm</math>.02</b>	<b>0.59<math>\pm</math>.02</b>	<b>0.61<math>\pm</math>.02</b>
$Ca^{2+}$	ECGFounder	0.28 $\pm$ .04	0.33 $\pm$ .03	0.36 $\pm$ .03	0.40 $\pm$ .02	0.42 $\pm$ .02	0.44 $\pm$ .03
	PPIDM-II	0.27 $\pm$ .04	0.32 $\pm$ .03	0.35 $\pm$ .03	0.39 $\pm$ .02	0.41 $\pm$ .02	0.43 $\pm$ .03
	Logit Adj.	0.26 $\pm$ .04	0.31 $\pm$ .03	0.34 $\pm$ .03	0.38 $\pm$ .02	0.41 $\pm$ .02	0.43 $\pm$ .02
	<b>MINN</b>	<b>0.33<math>\pm</math>.03</b>	<b>0.39<math>\pm</math>.02</b>	<b>0.42<math>\pm</math>.02</b>	<b>0.46<math>\pm</math>.02</b>	<b>0.48<math>\pm</math>.02</b>	<b>0.50<math>\pm</math>.02</b>
$Mg^{2+}$	ECGFounder	0.21 $\pm$ .05	0.26 $\pm$ .04	0.28 $\pm$ .03	0.31 $\pm$ .03	0.33 $\pm$ .03	0.35 $\pm$ .03
	PPIDM-II	0.20 $\pm$ .05	0.25 $\pm$ .04	0.27 $\pm$ .03	0.30 $\pm$ .03	0.32 $\pm$ .03	0.34 $\pm$ .03
	Logit Adj.	0.19 $\pm$ .04	0.24 $\pm$ .03	0.27 $\pm$ .03	0.30 $\pm$ .03	0.32 $\pm$ .03	0.34 $\pm$ .03
	<b>MINN</b>	<b>0.26<math>\pm</math>.04</b>	<b>0.31<math>\pm</math>.03</b>	<b>0.34<math>\pm</math>.03</b>	<b>0.37<math>\pm</math>.02</b>	<b>0.39<math>\pm</math>.02</b>	<b>0.40<math>\pm</math>.03</b>

### B.4 CNN Encoder Ablation

To support the claim that MINN is encoder-agnostic, we replace the ResNet-18 backbone with a simpler 4-layer 1D CNN encoder and repeat the main experiments. If MINN’s advantage persists with the simpler encoder, this provides evidence that the manifold operator residual—not the encoder architecture—is the primary driver of improvement. Table B.5 confirms that the manifold residual provides a consistent +0.07  $K^+$  F1 improvement regardless of backbone, supporting encoder-agnostic robustness.

Table B.5:  $K^+$  Macro-F1 with and without the manifold operator residual for two encoder architectures (mean  $\pm$  std, 5 seeds).

Encoder	w/ MINN	w/o MINN	$\Delta$
ResNet-18	0.61 $\pm$ .02	0.54 $\pm$ .02	+0.07
4-layer CNN	0.55 $\pm$ .02	0.48 $\pm$ .02	+0.07

## Appendix C

---

# IMPLEMENTATION AND REPRODUCIBILITY

---

**Code availability.** The complete implementation, including training scripts, evaluation pipelines, and baseline implementations, will be made publicly available upon publication.

**Data access.** The MIMIC-IV [5] and MIMIC-IV-ECG [6] datasets are publicly available through PhysioNet (<https://physionet.org>). Access requires completion of a data use agreement and the CITI “Data or Specimens Only Research” training course. Preprocessing scripts that reproduce the exact ECG-lab pairing, multi-beat segmentation, and metadata extraction pipelines used in this work will be made publicly available alongside the code upon publication.

**Compute requirements.** All experiments were conducted on a single NVIDIA A100 (40 GB). Training the full model at 100% data required approximately 4 hours; graph reconstruction added roughly 10 minutes per update. The dominant computational cost is the periodic  $k$ -NN graph construction, which requires  $O(N^2d)$  time for brute-force search or  $O(N \log N \cdot d)$  with approximate methods.

**Software environment.** All experiments used Python 3.10, PyTorch 2.1.2, and CUDA 11.8 on Ubuntu 22.04.

**Random seeds.** All results are reported as mean  $\pm$  standard deviation over five random seeds (0–4). Seeds control weight initialization, data shuffling, and subsampling for the low-data protocol. The same seeds are used for all methods to ensure paired comparisons.

**Ethical considerations.** This work uses the MIMIC-IV-ECG dataset, a de-identified dataset accessed under PhysioNet credentialed agreement [39]. No additional patient data were collected. Access requires completion of the CITI “Data or Specimens Only Research” training course and a signed data use agreement. This study is intended for research purposes and is not ready for clinical deployment; any clinical use would require prospective validation and regulatory approval.