DPIKAN: DECOMPOSED PHYSICS-INFORMED KANS FOR PHYSIOLOGICAL SIGNAL REPRESENTATION LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Learning representations from physiological signals is challenging because heterogeneous sensing modalities provide indirect and modality dependent observations of underlying cardiovascular dynamics. In this work, we propose the Decomposed Physics-Informed Kolmogorov–Arnold Network (DPIKAN), a compact framework that incorporates physiological knowledge through two mechanisms: multi-scale waveform modeling and cardiovascular descriptor guidance. For multi-scale waveform modeling, DPIKAN decomposes physiological signals into fast- and slow-varying components and assigns them component-specific functional capacities through a dual-order KAN. For cardiovascular descriptor guidance, it integrates cardiovascular descriptors through a physics-informed KAN and regularizes how predictions respond to variations in these descriptors. During experiments, we evaluate DPIKAN for blood pressure estimation on four datasets spanning photoplethysmography (PPG), bioimpedance (BioZ), and millimeter-wave (mmWave) sensing. We further introduce Ring-HGCPT, a smart ring PPG dataset collected during controlled cardiovascular perturbations, and evaluate DPIKAN using leave-one-subject-out cross-validation to assess its generalization to unseen participants. Across these heterogeneous sensing settings, DPIKAN achieves competitive performance for BP estimation. On Ring-HGCPT, it achieves mean errors of 1.08 and 0.09 mmHg for systolic blood pressure (SBP) and diastolic blood pressure (DBP), respectively, with corresponding error standard deviations of 5.41 and 4.81 mmHg. Overall, DPIKAN offers a compact, physiology-informed approach to representation learning from heterogeneous physiological signals. We will release our source code and data.

1 INTRODUCTION

Physiological signals acquired through wearable and contactless sensors offer a non-invasive window into cardiovascular activity. These signals contain information related to key cardiovascular indicators, including heart rate, heart rate variability, vascular pulse characteristics, and blood pressure (BP). Among these, BP is particularly important for assessing cardiovascular health, yet conventional cuff-based measurements are inconvenient for frequent monitoring (Flint et al., 2019; Kireev et al., 2022). This has motivated cuffless BP estimation from photoplethysmography (PPG), bioimpedance (BioZ), and millimeter-wave (mmWave) signals (Sel et al., 2023b; Liang et al., 2023; Zhu et al., 2025). These modalities, however, do not measure BP directly: each captures a different manifestation of cardiovascular activity, shaped by its sensing mechanism and the individual being measured. Consequently, BP estimation depends on learning representations that retain cardiovascular information despite variations in observed waveforms (Zheng et al., 2026).

Learning such representations is difficult for two main reasons. First, physiological waveforms contain patterns at different temporal scales: beat-level pulsations coexist with slower variations that may reflect both cardiovascular state and sensing conditions (Lee et al., 2026). These components have different approximation requirements, making uniform representation modeling across temporal scales suboptimal. Second, reliable BP estimation requires model predictions to remain consistent with established physiological relationships. Without explicit physiological constraints,

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

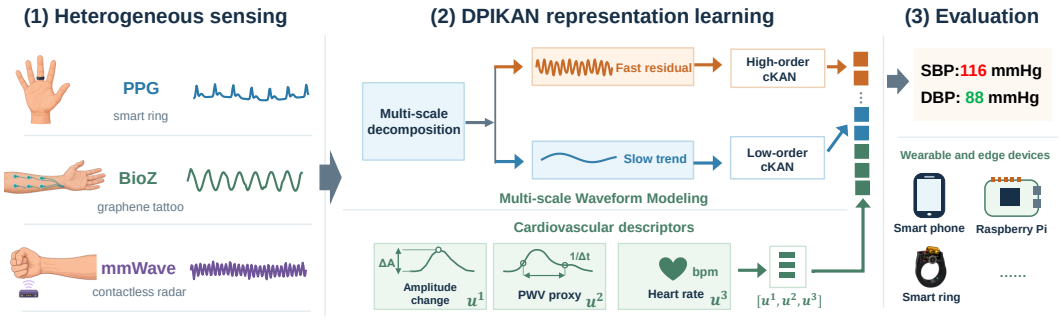

Figure 1: Overview of DPIKAN. (1) PPG, BioZ, and mmWave provide heterogeneous physiological signals. (2) DPIKAN combines multi-scale waveform modeling with cardiovascular descriptor guidance. (3) Its compact design supports deployment on wearable and edge devices.

a data-driven model may achieve low prediction error while responding implausibly to changes in cardiovascular descriptors, thereby reducing the credibility of its predictions (Wang et al., 2026).

Recent physiological representation learning shows the value of incorporating waveform morphology and temporal resolution into learned features (Pillai et al., 2025; Lee et al., 2026; Yoon et al., 2026). Yet recognizing structure does not by itself determine how much modeling capacity different temporal components require. KAN-based time-series methods have explored flexible modeling of frequency-specific and offset-dependent patterns (Huang et al., 2025; Zhang et al., 2026), but their forecasting-oriented designs do not use cardiovascular knowledge to guide the learned representations. Conversely, physics-informed BP models introduce cardiovascular descriptors and constraints (Sel et al., 2023a; Wang et al., 2026), without explicitly matching functional capacity to decomposed fast- and slow-varying components. The remaining challenge is therefore to couple component-specific capacity allocation with explicit cardiovascular priors in a compact model for physiological signal representation learning.

To address this gap, we propose DPIKAN, a compact framework that combines multi-scale waveform modeling with cardiovascular descriptor guidance, as illustrated in Figure 1. Multi-scale decomposition separates each input into a fast-varying residual and a slow-varying trend, while a dual-order KAN assigns them component-specific functional capacities rather than uniformly increasing the capacity of the entire network (Liu et al., 2025). A physics-informed KAN further integrates cardiovascular descriptors, with sensitivity regularization constraining how BP predictions respond to variations in these descriptors.

We evaluate DPIKAN on three public datasets spanning BioZ, PPG, and mmWave sensing: Graphene-HGCPT and Ring-CPT with BioZ signals, and Blumio with PPG and mmWave signals (Gomes et al., 2021; Kireev et al., 2022; Sel et al., 2023b). We further introduce Ring-HGCPT, a smart-ring PPG dataset collected during controlled Handgrip and Cold Pressor Tests to elicit BP variations. Subject-specific evaluation on the public datasets and leave-one-subject-out evaluation on Ring-HGCPT assess performance across sensing modalities and generalization to unseen subjects.

Our contributions are summarized as follows:

- We propose DPIKAN, a physiology-informed framework that combines multi-scale decomposition, dual-order KAN modeling, and explicit cardiovascular priors to learn compact representations from heterogeneous physiological signals.
- We introduce Ring-HGCPT, a dataset collected using a wearable smart ring for non-invasive physiological sensing under controlled cardiovascular perturbations, supporting evaluation of physiological representations across changing cardiovascular states and unseen participants.
- Extensive experiments are conducted across four datasets spanning PPG, BioZ, and mmWave sensing. Quantitative comparisons, ablation studies, and representation analyses demonstrate the effectiveness of DPIKAN for BP estimation with substantially reduced model complexity.

2 Related Work

Physiological Signal Representation Learning. Learning representations from physiological signals has progressed from task-specific supervised modeling toward transferable pretraining and physiology-informed learning. Supervised approaches use recurrent and attention-based architectures to capture temporal patterns (Ma et al., 2022), while FewShotBP (Fan et al., 2023) explores data-efficient personalized learning in cardiovascular sensing. PaPaGei (Pillai et al., 2025) leverages large-scale unlabeled PPG data and signal morphology knowledge to learn transferable representations for diverse health-related tasks. HiMAE (Lee et al., 2026) learns multi-resolution representations through hierarchical masked autoencoding, revealing task-relevant temporal scales. Beyond Hearing (Yoon et al., 2026) incorporates physiological knowledge into multi-band tokenization for task-agnostic ExG representation learning. Complementing these pretraining approaches, PINN (Sel et al., 2023a) and PITN (Wang et al., 2026) incorporate cardiovascular priors into physiological signal modeling for BP estimation. DPIKAN combines component-specific KAN capacity allocation with explicit cardiovascular descriptors and physiology-informed sensitivity regularization to learn compact representations for cuffless BP estimation.

Kolmogorov–Arnold Networks. Kolmogorov–Arnold Networks (KANs) (Liu et al., 2025) parameterize nonlinear mappings through learnable univariate functions on edges, offering an alternative to conventional MLPs that we investigate for modeling the nonlinear associations between cardiovascular descriptors and blood pressure. KANs have been explored in diverse applications, including vision transformers (Yang & Wang, 2025), time-series anomaly detection (Zhou et al., 2025), and graph attention (Fang et al., 2025). For time-series forecasting, TimeKAN (Huang et al., 2025) combines frequency decomposition with multi-order KANs to represent frequency-specific temporal patterns, while Time-TK (Zhang et al., 2026) integrates KANs with Transformer-based temporal interaction to capture dependencies among multiple offset subsequences. In physics-informed learning, PIKAN (Shukla et al., 2024) and KINN (Wang et al., 2025) employ KAN-based representations for solving forward and inverse problems constrained by differential equations. Building on these temporal and physics-informed modeling perspectives, DPIKAN combines fast- and slow-component modeling with explicit cardiovascular descriptors and physiology-informed sensitivity regularization for cuffless BP estimation.

Physiological Sensing and Datasets. Physiological sensing provides indirect observations of biological processes through diverse modalities, with data resources varying in acquisition conditions, population coverage, and reference labels. Within cardiovascular sensing, MIMIC-BP (Sanches et al., 2024) provides clinical waveforms for model development and evaluation. Wearable studies explore wrist PPG (Cao et al., 2021), smart-ring sensing (Zhu et al., 2025), and data-efficient personalized modeling (Fan et al., 2023). Non-contact approaches employ RF (Wang et al., 2024b) and mmWave sensing (Liang et al., 2023; Hu et al., 2024), while BP3 (Zheng et al., 2024) integrates physiological factors in a cardiovascular disease cohort. Complementing these settings, Ring-HGCPT uses a wearable smart ring for non-invasive physiological signal acquisition during controlled Handgrip and Cold Pressor Tests. It supports evaluation of physiological representations under changing cardiovascular states and across unseen participants, using BP estimation as the downstream task.

3 Methodology

Let $\mathcal{D} = \{(\mathbf{X}_i, y_i)\}_{i=1}^{N}$ denote a dataset of N paired physiological signal segments and reference BP measurements. Here, $\mathbf{X}_i \in \mathbb{R}^{T\times C}$ contains T time steps and C sensing channels, and $y_i \in \mathbb{R}$ denotes the corresponding systolic or diastolic BP value. The goal is to learn a predictor f_θ that estimates BP from the observed signal, such that $\hat{y}_i = f_\theta(\mathbf{X}_i)$.

As shown in Fig. 2, DPIKAN combines multi-scale waveform modeling with cardiovascular descriptor guidance. Multi-Scale Decomposition (Sec. 3.1) separates the input into fast- and slow-varying components, while the Dual-Order KAN (Sec. 3.2) assigns them component-specific functional capacities. The Physics-Informed KAN (Sec. 3.3) then integrates cardiovascular descriptors into BP prediction, while the associated sensitivity regularization constrains the model's response to variations in these descriptors.

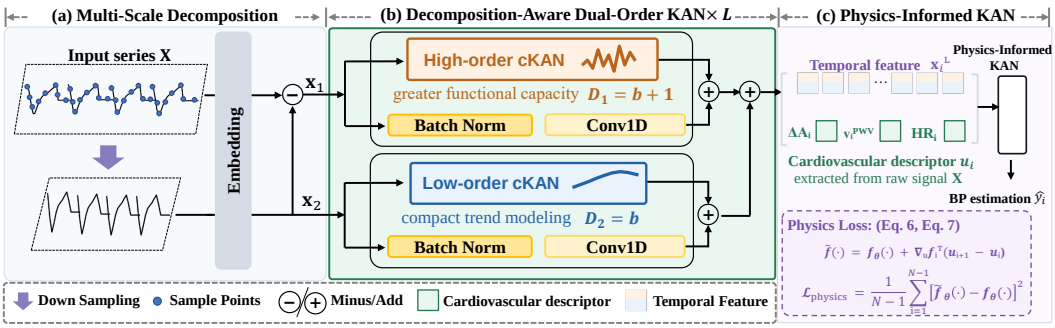

Figure 2: Model architecture. (a) Multi-scale decomposition separates the input into high- and low-frequency components. (b) Dual-order KANs model the two streams in parallel using Chebyshev KAN (cKAN) blocks. (c) Physics-informed KAN fuses the learned representations with physiological features for BP estimation.

3.1 Multi-Scale Decomposition

As shown in Fig. 2(a), we represent the input signal $\mathbf{X} \in \mathbb{R}^{T\times C}$ at two temporal resolutions. Specifically, $\mathbf{x}1 \in \mathbb{R}^{T\times D\text{model}}$ preserves the original temporal resolution after embedding, whereas $\mathbf{x}2 \in \mathbb{R}^{\frac{T}{2}\times D\text{model}}$ captures the slow-varying information obtained through average pooling followed by embedding. The two representations are defined as follows:

$$\mathbf{x}_1 = \mathrm{Embed}(\mathbf{X}), \quad \mathbf{x}_2 = \mathrm{Embed}(\mathrm{AvgPool}(\mathbf{X})), \tag{1}$$

where $\mathrm{Embed}(\cdot)$ denotes the projection of inputs into the D_{model} dimensional latent space. To eliminate the low-frequency trends from the high-frequency dynamics, we employ a residual-based approach. Given that $\mathbf{x}_2$ captures the low-frequency information, we first project it back to the original temporal resolution T to obtain the underlying trend, denoted as $\hat{\mathbf{x}}_2$. This is achieved by transforming $\mathbf{x}_2$ into the frequency domain via Fast Fourier Transform (FFT), zero-padding the spectral sequence to the target length, and reconstructing the signal via Inverse Fast Fourier Transform (IFFT)

$$\hat{\mathbf{x}}_2 = \mathrm{IFFT}(\mathrm{Pad}_T(\mathrm{FFT}(\mathbf{x}_2))), \qquad \mathbf{x}_1 \leftarrow \mathbf{x}_1 - \hat{\mathbf{x}}_2. \tag{2}$$

where $\mathrm{Pad}_T(\cdot)$ zero-pads the spectrum to the target length T, and $\hat{\mathbf{x}}_2 \in \mathbb{R}^{T\times D\text{model}}$ denotes the reconstructed slow-varying component. The residual subtraction suppresses slow-varying trends in the full-resolution stream, allowing $\mathbf{x}_1$ to emphasize rapid variations and fine-grained temporal patterns while $\mathbf{x}_2$ preserves the slower temporal structure. As pooling and Fourier interpolation are parameter-free, this decomposition introduces little additional computational overhead.

3.2 Decomposition-Aware Dual-Order KAN

Following the decomposition in Sec. 3.1, $\mathbf{x}_1$ and $\mathbf{x}_2$ denote the fast residual and slow trend, respectively. The fast residual preserves sharp beat-level variations and morphology-related details, whereas the slow trend varies more smoothly. Using the same functional capacity for both components may therefore underfit the fast branch or unnecessarily over-parameterize the slow branch.

To address this asymmetry, we employ Chebyshev KAN (cKAN) (SS et al., 2024) and use its polynomial order to control branch capacity. Specifically, we assign $D_1 = b + 1$ to the fast branch and $D_2 = b$ to the slow branch, where b is the base polynomial order. This decomposition-aware order allocation matches model capacity to the temporal complexity of each component without uniformly increasing the backbone width or depth.

At layer l, each component is processed by an order-specific cKAN branch and a parallel convolutional branch, whose outputs are combined through element-wise addition:

$$\mathbf{x}_j^{l+1} = \mathrm{cKAN}\big(\mathbf{x}_j^l; D_j\big) + \mathrm{Conv1D}\big(\mathrm{BN}\big(\mathbf{x}_j^l\big)\big), \qquad j \in \{1, 2\}, \tag{3}$$

where $j = 1$ and $j = 2$ index the fast and slow branches, respectively. The cKAN branch provides component-specific nonlinear capacity, while the convolutional branch captures local patterns. The

key design is the coupling of multi-scale decomposition with branch-specific cKAN orders, rather than the use of cKAN alone.

After L stacked layers, the slow representation is restored to the original temporal resolution using the same Fourier interpolation as in Sec. 3.1. It is then added to the fast representation to obtain the fused temporal feature:

$$\mathbf{x}^L = \mathbf{x}_1^L + \mathrm{IFFT}\left(\mathrm{Pad}_T\left(\mathrm{FFT}(\mathbf{x}_2^L)\right)\right). \tag{4}$$

The fused representation $\mathbf{x}^L$ combines the fast temporal variations with the reconstructed slow trend and is subsequently passed to the Physics-Informed KAN.

3.3 Physics-Informed KAN

As shown in Fig. 2(c), the Physics-Informed KAN (PIKAN) integrates the learned representation $\mathbf{x}_i^L$ with explicit cardiovascular descriptors. For each signal segment $\mathbf{X}_i$, we extract and temporally align $\mathbf{u}_i = [\Delta A_i, v_i^{\mathrm{PWV}}, HR_i]^\top$ from the corresponding raw waveform, where the three components denote amplitude change, a PWV proxy, and heart rate, respectively. Detailed definitions and extraction procedures are provided in Appendix C.1. The BP prediction is then given by

$$\hat{y}_i = f_\theta(\mathbf{X}_i, \mathbf{u}_i) = \mathrm{KAN}\big([\mathbf{x}_i^L; \mathbf{u}_i]\big). \tag{5}$$

Following prior work (Sel et al., 2023a; Wang et al., 2026), we regularize the model using a first-order Taylor approximation along the physiological-descriptor dimensions. For two consecutive segments from the same subject and session,

$$\widetilde{f}_\theta(\mathbf{X}, \mathbf{u}) = f_\theta(\mathbf{X}_i, \mathbf{u}_i) + \sum_{k=1}^{3} \frac{\partial f_\theta(\mathbf{X}_i, \mathbf{u}_i)}{\partial u_i^k}(u_{i+1}^k - u_i^k). \tag{6}$$

The physics-informed loss penalizes the discrepancy between this local approximation and the prediction on the subsequent segment:

$$\mathcal{L}_{\mathrm{physics}} = \frac{1}{N-1}\sum_{i=1}^{N-1}\left[\widetilde{f}_\theta(\mathbf{X_{i+1}}, \mathbf{u}_{i+1}) - f_\theta(\mathbf{X}_{i+1}, \mathbf{u}_{i+1})\right]^2. \tag{7}$$

3.4 Training Objective

We optimize DPIKAN using a supervised regression loss together with the physics-informed regularization introduced above. The losses are defined as

$$\mathcal{L}_{\mathrm{data}} = \frac{1}{S}\sum_{i=1}^{S}(\hat{y}_i - y_i)^2, \quad \mathcal{L}_{\mathrm{total}} = \mathcal{L}_{\mathrm{data}} + \gamma\mathcal{L}_{\mathrm{physics}}, \tag{8}$$

where $S \le N$ denotes the number of labeled training samples, and γ controls the contribution of the physics-informed regularization. This objective balances supervised BP prediction with local consistency under variations in the cardiovascular descriptors.

4 Ring-HGCPT: A Controlled Wearable Benchmark

We construct Ring-HGCPT, a wearable physiological dataset designed to evaluate cuffless BP estimation under controlled cardiovascular variations and subject-independent settings. Existing cuffless BP datasets cover diverse sensing technologies and acquisition protocols, but practical wearable sensing, controlled physiological perturbations, and strict unseen-subject evaluation are not always jointly available. Ring-HGCPT combines continuous smart-ring PPG sensing with standardized Handgrip (HG) and Cold Pressor Tests (CPT), together with sparse cuff-based BP references, providing a complementary setting for evaluating physiology-informed representation learning under changing cardiovascular states.

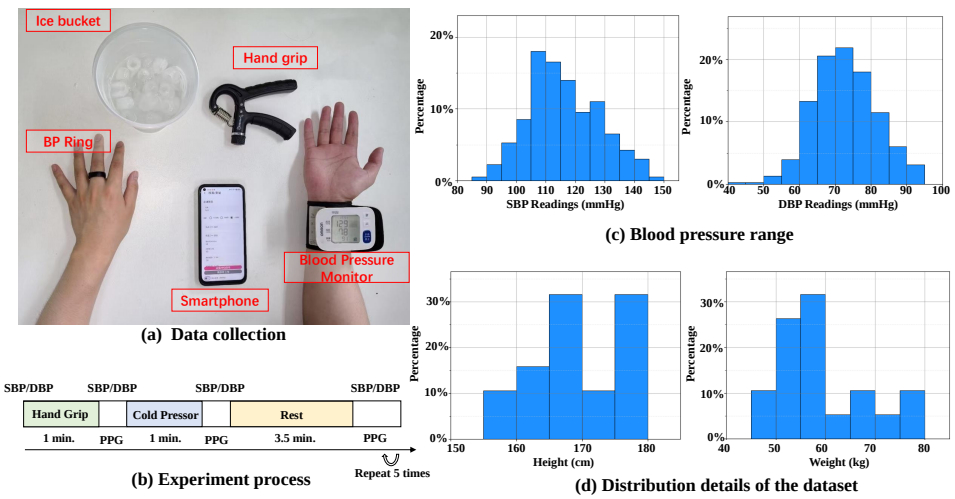

Figure 3: Overview of the Ring-HGCPT dataset. (a) Wearable PPG acquisition with a smart ring and cuff-based BP reference. (b) Repeated Handgrip, Cold Pressor, and Rest protocol for inducing cardiovascular variations. (c) Distribution of collected SBP and DBP measurements. (d) Distribution of participants' height and weights.

Data Collection Platform. Ring-HGCPT uses a portable smart-ring sensor to continuously acquire PPG signals, while a commercial cuff-based blood pressure monitor provides intermittent SBP and DBP references. The ring-type PPG sensor follows the wearable sensing platform used in prior RingBP studies, enabling continuous physiological signal acquisition in a compact wearable form factor. A total of 20 participants were recruited [1], and PPG signals were recorded at 200 Hz. Detailed device specifications are provided in Appendix B.4.

Data Collection Protocol. To induce controlled within-subject cardiovascular variations, each participant undergoes a repeated protocol consisting of Handgrip, Cold Pressor, and Rest phases. Each trial begins with a resting baseline measurement, followed by 1 min of handgrip, 1 min of cold pressor stimulation, and 3.5 min of rest. PPG is acquired continuously throughout the protocol, while cuff-based SBP/DBP measurements are collected at baseline and after each of the three phases. The complete procedure is repeated five times.

Dataset Statistics. Ring-HGCPT comprises 400 synchronized PPG–BP pairs. Across participants, SBP and DBP span 89–148 mmHg and 43–106 mmHg, respectively. The mean maximum within-subject BP differences are 21.0±5.2 mmHg for SBP and 22.0±6.0 mmHg for DBP, showing that the protocol captures appreciable changes within individual participants. BP and participant height and weight distributions are shown in Fig. 3(c,d), with further cohort details in Appendix B.3.

5 Experiments

5.1 Experimental Setup

Datasets. We evaluate DPIKAN on four datasets spanning heterogeneous physiological sensing modalities: Graphene-HGCPT (Kireev et al., 2022), Ring-CPT (Sel et al., 2023b), Blumio (Gomes et al., 2021), and our Ring-HGCPT benchmark. Graphene-HGCPT and Ring-CPT provide wearable BioZ signals with continuous beat-to-beat BP references, while Blumio contains PPG and mmWave measurements from 115 subjects. Ring-HGCPT contains smart-ring PPG recordings paired with intermittent cuff-based BP measurements collected during controlled Handgrip and Cold Pressor Tests. Together, these datasets cover BioZ, PPG, and mmWave sensing conditions. Following prior work (Wang et al., 2026), we adopt dataset-specific evaluation protocols. For the three public datasets, we use subject-specific training following their established settings. For Ring-HGCPT,

[1]The study was approved by the Institutional Review Board, and all participants provided written informed consent before data collection.

Table 1: Performance comparison across public datasets and sensing modalities. Results are reported as SBP/DBP, with RMSE in mmHg. ↑/↓ indicate higher/lower is better. **Bold** and underlined values denote the best and second-best performance.

Methods	Blumio (PPG)		Blumio (mmWave)		Graphene-HGCPT (BioZ)		Ring-CPT (BioZ)	
	SBP/DBP		SBP/DBP		SBP/DBP		SBP/DBP	
	Corr. ↑	RMSE ↓	Corr. ↑	RMSE ↓	Corr. ↑	RMSE ↓	Corr. ↑	RMSE ↓
General time-series baselines								
TimesNet [ICLR 23]	0.22/0.27	4.3/3.7	0.22/0.28	4.1/3.6	0.32/0.26	14.4/13.1	0.56/0.49	9.1/7.3
iTransformer [ICLR 24]	0.28/0.32	7.5/5.7	0.27/0.27	6.6/5.2	0.38/0.34	19.3/18.4	0.55/0.49	11.3/10.0
TimeMixer [ICLR 24]	0.30/0.29	4.2/4.1	0.33/0.29	4.4/3.8	0.25/0.30	15.0/13.3	0.57/0.46	8.9/7.8
TimeKAN [ICLR 25]	0.19/0.22	4.5/4.0	0.14/0.20	4.0/3.6	0.01/-0.02	14.5/12.6	0.17/0.13	10.7/7.4
Time-TK [WWW 26]	0.26/0.28	4.0/3.8	0.32/0.29	4.0/3.5	0.28/0.35	13.7/11.7	0.60/0.46	8.6/6.3
BP-specific baselines								
ResNet1D [Sci. Data 23]	0.28/0.35	7.2/5.9	0.29/0.30	6.7/6.0	0.23/0.29	17.6/14.8	0.59/0.56	10.4/9.3
BpBLS [JBHI 26]	0.32/0.37	6.9/5.3	0.41/0.34	5.7/4.9	0.42/0.38	17.1/16.0	0.63/0.56	9.5/9.2
Physics-Informed baselines								
PINN [npj Digit. Med. 23]	0.41/0.40	4.0/3.7	0.45/0.37	3.8/3.5	0.48/0.50	12.8/11.1	0.66/0.67	8.7/6.0
PINNMamba [ICML 25]	0.35/0.36	4.1/3.6	0.46/0.39	3.9/3.4	0.37/0.38	13.2/11.6	0.56/0.47	8.4/6.6
PITN (base) [TMC 26]	**0.42**/0.40	4.2/3.8	**0.47**/0.40	3.8/3.4	0.60/0.56	11.4/10.3	0.74/**0.69**	7.0/5.5
DPIKAN (Ours)	**0.42/0.41**	**3.5/3.3**	**0.47/0.41**	**3.4/3.2**	**0.61/0.57**	**11.1/9.9**	**0.75**/0.66	**6.8/5.3**

we use leave-one-subject-out (LOSO) cross-validation to evaluate generalization to unseen subjects, with each held-out subject used exclusively for testing. Detailed dataset statistics, preprocessing and evaluation protocols are described in Appendix B.

Baselines. We compare DPIKAN with three groups of representative baselines: (1) general time-series models, including TimesNet (Wu et al., 2023), iTransformer (Liu et al., 2024), TimeMixer (Wang et al., 2024a), TimeKAN (Huang et al., 2025), and Time-TK (Zhang et al., 2026); (2) BP-specific models, including ResNet1D (González et al., 2023) and BpBLS (Huang et al., 2026); and (3) physics-informed models, including PINN (Sel et al., 2023a), PINNMamba (Xu et al., 2025), and PITN (Wang et al., 2026). All baselines are evaluated using the same dataset splits and preprocessing protocols as DPIKAN.

Implementation Details. All models are implemented in PyTorch and optimized with Adam using a learning rate of 10^{-3} on a single NVIDIA RTX 3090 GPU. Computational efficiency is evaluated using parameter count, FLOPs, and inference latency, with latency measured on a Raspberry Pi 4 Model B (4GB RAM) under a unified batch-size-one setting. Detailed training configurations and efficiency benchmarking protocols are provided in Appendix C.

5.2 Quantitative Results

Results on Public Datasets. Table 1 summarizes the results across four evaluation settings spanning three sensing modalities: PPG, mmWave, and BioZ. DPIKAN achieves consistently strong performance across all evaluated datasets. Compared with the physics-informed PITN baseline, DPIKAN reduces SBP/DBP RMSE in all four settings, with particularly clear improvements on the PPG and mmWave subsets of Blumio, while maintaining competitive correlations across datasets. These results demonstrate that the proposed physiology-informed representation framework can be applied consistently across heterogeneous sensing mechanisms without modality-specific architectural redesign.

Table 2: Performance comparison on Ring-HGCPT. Higher correlation and lower RMSE indicate better performance.

Methods	SBP		DBP	
	Corr. ↑	RMSE ↓	Corr. ↑	RMSE ↓
General time-series baselines				
TimesNet [ICLR 23]	0.89	5.7	0.77	6.5
iTransformer [ICLR 24]	0.28	14.5	0.20	11.7
TimeMixer [ICLR 24]	0.29	13.8	0.15	12.5
TimeKAN [ICLR 25]	0.88	5.8	0.82	5.9
Time-TK [WWW 26]	0.50	11.1	0.28	9.8
BP-specific baselines				
ResNet1D [Sci. Data 23]	0.56	11.6	0.57	9.6
BpBLS [JBHI 26]	0.67	9.3	0.48	9.0
Physics-Informed baselines				
PINN [npj Digit. Med. 23]	0.83	5.6	0.80	5.7
PINNMamba [ICML 25]	0.46	10.7	0.52	8.8
PITN [TMC 26]	0.23	11.9	0.21	10.1
DPIKAN (Ours)	**0.90**	**5.5**	**0.88**	**4.8**

Results on Ring-HGCPT. We further evaluate DPIKAN on Ring-HGCPT under leave-one-subject-out (LOSO) cross-validation to assess subject-independent performance under controlled physiological perturbations. As shown in Table 2, DPIKAN achieves the strongest overall performance, with SBP/DBP correlations of 0.90/0.88 and RMSEs of 5.5/4.8 mmHg, respectively. PITN performs less favorably under LOSO evaluation, which may be because it was originally designed and evaluated in subject-specific settings and therefore relies more heavily on subject-dependent physiological relationships. The gains are particularly clear for DBP, where DPIKAN improves both correlation and prediction error over the strongest baselines. These results provide evidence that the proposed framework remains effective for subjects unseen during training and under induced cardiovascular variations.

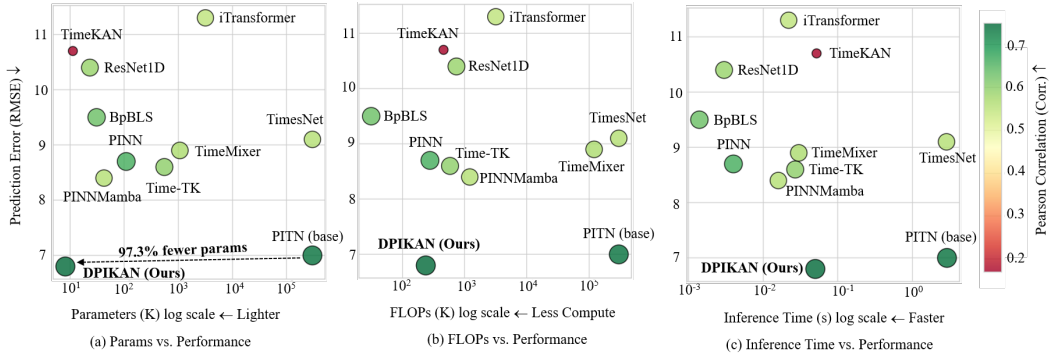

Figure 4: Performance–efficiency comparison in terms of (a) parameters, (b) FLOPs, and (c) inference time. Lower-left positions indicate lower prediction error and computational cost, while marker color and size encode Pearson correlation.

Computational Efficiency Figure 4 compares parameter count, FLOPs, and inference latency. DPIKAN occupies a favorable low-error and low-cost region across all three comparisons. Relative to PITN, it uses 97.3% fewer parameters while requiring substantially fewer FLOPs and lower inference latency. Although PINN is faster, DPIKAN achieves lower RMSE and higher correlation, providing a stronger overall balance between predictive performance and computational efficiency.

5.3 Ablation Study

We ablate DPIKAN from two perspectives: multi-scale waveform modeling and cardiovascular descriptor guidance. As shown in Table 3, removing multi-scale decomposition or replacing the branch-specific cKAN orders with uniform orders or an MLP backbone consistently degrades performance. Removing the cardiovascular descriptors, and consequently the associated physics-informed regularization, also causes substantial degradation, while removing only the regularization retains the descriptors as model inputs but still reduces performance. The full DPIKAN achieves SBP/DBP correlations of 0.75/0.66 and RMSEs of 6.8/5.3 mmHg on Ring-CPT, supporting the contributions of multi-scale decomposition, branch-specific capacity allocation, and cardiovascular descriptor guidance.

Table 3: Ablation of the main components in DPIKAN. Results are reported as SBP/DBP.

Variants	Graphene-HGCPT SBP/DBP		Ring-CPT SBP/DBP	
	Corr. ↑	RMSE ↓	Corr. ↑	RMSE ↓
Representation modeling				
w/o Multi-scale Modeling	0.20/0.18	13.6/11.7	0.60/0.47	7.0/6.9
Fixed low-order KANs	0.39/0.34	13.3/11.4	0.56/0.49	9.3/6.6
Fixed high-order KANs	0.42/0.39	12.8/11.4	0.58/0.49	9.1/6.7
MLP backbone	0.40/0.38	13.1/11.2	0.65/0.47	8.7/6.7
Physics-Informed learning				
w/o Cardiovascular descriptors	0.22/0.21	13.5/11.8	0.29/0.10	10.5/7.5
w/o Physics Loss	0.24/0.17	13.1/11.9	0.60/0.37	8.5/6.9
DPIKAN (Ours)	**0.61/0.57**	**11.1/9.9**	**0.75/0.66**	**6.8/5.3**

5.4 Qualitative Results

Figure 5 presents the agreement between DPIKAN's predictions and reference measurements on Ring-HGCPT. The Bland–Altman plots in Fig. 5(a) show mean errors of 1.08/0.09 mmHg and error

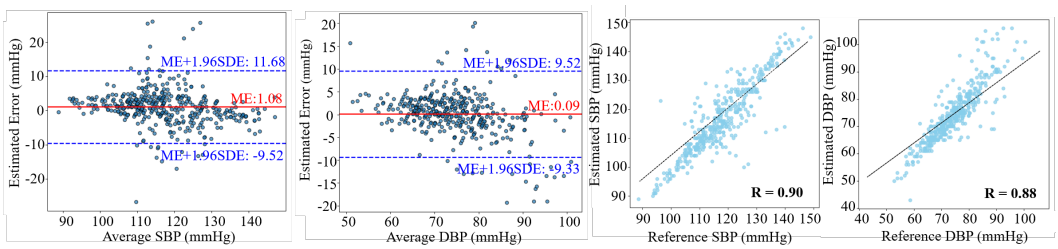

(a) Bland-Altman Analysis (b) Pearson Correlation

Figure 5: Bland–Altman and correlation analyses of SBP and DBP predictions on Ring-HGCPT.

standard deviations of 5.41/4.81 mmHg for SBP/DBP, where the solid red and dashed blue lines indicate the mean errors and 95% limits of agreement, respectively. In Fig. 5(b), the predicted and reference values achieve Pearson correlation coefficients of 0.90/0.88 for SBP/DBP, with the diagonal lines indicating identity.

5.5 Representation Analysis

To examine how DPIKAN models the decomposed physiological signals, we visualize the temporal components and their learned representations in Figure 6. Figure 6(a) shows that the multi-scale decomposition separates the input into slow-varying and fast-varying components. As shown in Fig. 6(b), the corresponding representations exhibit substantial overlap before differentiated KAN modeling, but become more clearly separated after the Dual-Order KAN. This observation suggests that decomposition alone does not fully differentiate the two temporal components in the latent space. Together with the ablation results, the visualization provides qualitative support for assigning different functional capacities to heterogeneous temporal dynamics.

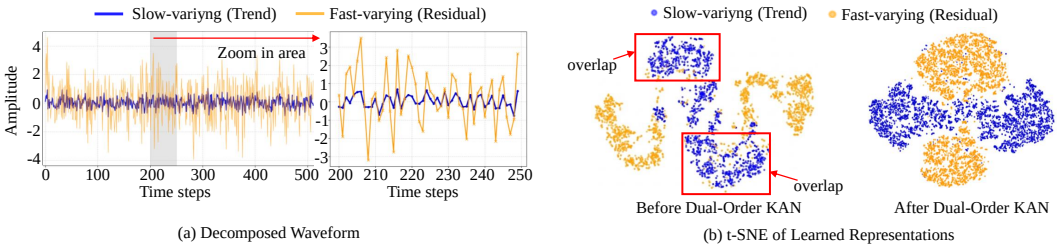

(a) Decomposed Waveform (b) t-SNE of Learned Representations

Figure 6: Visualization of multi-scale decomposition and branch representations. (a) The input signal is decomposed into a slow-varying trend and a fast-varying residual. (b) t-SNE projections of the branch representations at the input and output of the Dual-Order KAN, showing reduced overlap and clearer branch-specific organization after differentiated modeling.

6 Conclusion

We presented DPIKAN, a compact physiology-informed framework that combines multi-scale waveform modeling with cardiovascular descriptor guidance to learn representations from heterogeneous physiological signals. We also introduced Ring-HGCPT, a smart-ring PPG benchmark collected during controlled Handgrip and Cold Pressor Tests for evaluating representation robustness and generalization to unseen subjects through leave-one-subject-out evaluation. Experiments across PPG, BioZ, and mmWave sensing settings, together with ablation and representation analyses, demonstrate consistent performance improvements and validate the proposed representation-learning design. Overall, this work advances compact physiological representation learning toward continuous cardiovascular health monitoring across wearable and contactless sensing platforms.

AI USE STATEMENT

Generative AI tools were used to improve the clarity, organization, and presentation of the manuscript

ETHICS STATEMENT

This study involves human physiological data collected under a protocol approved by the Institutional Review Board. All participants provided informed consent prior to data collection, and all collected data were anonymized before analysis. This work is intended solely for research purposes and should not be considered a clinical diagnostic system.

REPRODUCIBILITY STATEMENT

To support reproducibility, the model architecture and training objective are described in Section 3. Dataset configurations, preprocessing procedures, training settings, and evaluation protocols are documented in Section 5 and Appendices B and C.

REFERENCES

Yetong Cao, Huijie Chen, Fan Li, and Yu Wang. Crisp-bp: Continuous wrist ppg-based blood pressure measurement. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*, pp. 378–391, 2021.

Feiyi Fan, Yang Gu, Jianfei Shen, Fan Dong, and Yiqiang Chen. Fewshotbp: Towards personalized ubiquitous continuous blood pressure measurement. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 7(3):1–39, 2023.

Taoran Fang, Tianhong Gao, Chunping Wang, Yihao Shang, Wei Chow, Lei Chen, and Yang Yang. Kaa: Kolmogorov-arnold attention for enhancing attentive graph neural networks. In *ICLR*, 2025.

Alexander C Flint, Carol Conell, Xiushui Ren, Nader M Banki, Sheila L Chan, Vivek A Rao, Ronald B Melles, and Deepak L Bhatt. Effect of systolic and diastolic blood pressure on cardiovascular outcomes. *New England Journal of Medicine*, 381(3):243–251, 2019.

Elizabeth Gomes, Catherine Liao, Oliver Shay, and Nikhil Bikhchandani. A dataset of synchronized signals from wearable cardiovascular monitoring sensors, 2021.

Sergio González, Wan-Ting Hsieh, and Trista Pei-Chun Chen. A benchmark for machine-learning based non-invasive blood pressure estimation using photoplethysmogram. *Scientific Data*, 10(1): 149–164, 2023.

Qingyong Hu, Qian Zhang, Hao Lu, Shun Wu, Yuxuan Zhou, Qianyi Huang, Huangxun Chen, Ying-Cong Chen, and Ni Zhao. Contactless arterial blood pressure waveform monitoring with mmwave radar. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 8 (4):1–29, 2024.

Songtao Huang, Zhen Zhao, Can Li, and Lei Bai. Timekan: Kan-based frequency decomposition learning architecture for long-term time series forecasting. In *ICLR*, 2025.

Zi-Xuan Huang, Zeng-Ding Liu, Ye Li, CL Philip Chen, Hai-Liang Wang, and Fen Miao. Bpbls: a knowledge-embedded bi-incremental broad learning system for wearable cuffless blood pressure estimation. *IEEE Journal of Biomedical and Health Informatics*, 30(6):5489–5502, 2026.

Dmitry Kireev, Kaan Sel, Bassem Ibrahim, Neelotpala Kumar, Ali Akbari, Roozbeh Jafari, and Deji Akinwande. Continuous cuffless monitoring of arterial blood pressure via graphene bioimpedance tattoos. *Nature Nanotechnology*, 17(8):864–870, 2022.

Simon Lee, Cyrus Tanade, Hao Zhou, Juhyeon Lee, Megha Thukral, Md Sazzad Hissain Khan, Keum San Chun, Baiying Lu, Migyeong Gwak, Mehrab Bin Morshed, et al. Himae: Hierarchical masked autoencoders discover resolution-specific structure in wearable time series. In *ICLR*, volume 2026, pp. 64859–64895, 2026.

Yumeng Liang, Anfu Zhou, Xinzhe Wen, Wei Huang, Pu Shi, Lingyu Pu, Huanhuan Zhang, and Huadong Ma. airbp: Monitor your blood pressure with millimeter-wave in the air. *ACM Transactions on Internet of Things*, 4(4):1–32, 2023.

Yong Liu, Tengge Hu, Haoran Zhang, Haixu Wu, Shiyu Wang, Lintao Ma, and Mingsheng Long. itransformer: inverted transformers are effective for time series forecasting. In *ICLR*, 2024.

Ziming Liu, Yixuan Wang, Sachin Vaidya, Fabian Ruehle, James Halverson, Marin Soljačić, Thomas Y Hou, and Max Tegmark. Kan: Kolmogorov-arnold networks. In *ICLR*, 2025.

Chenbin Ma, Peng Zhang, Fan Song, Yangyang Sun, Guangda Fan, Tianyi Zhang, Youdan Feng, and Guanglei Zhang. Kd-informer: A cuff-less continuous blood pressure waveform estimation approach based on single photoplethysmography. *IEEE Journal of Biomedical and Health Informatics*, 27(5):2219–2230, 2022.

Arvind Pillai, Dimitris Spathis, Fahim Kawsar, and Mohammad Malekzadeh. Papagei: Open foundation models for optical physiological signals. In *ICLR*, 2025.

Ivandro Sanches, Victor V Gomes, Carlos Caetano, Lizeth SB Cabrera, Vinicius H Cene, Thomas Beltrame, Wonkyu Lee, Sanghyun Baek, and Otávio AB Penatti. Mimic-bp: A curated dataset for blood pressure estimation. *Scientific Data*, 11(1):1233, 2024.

Kaan Sel, Amirmohammad Mohammadi, Roderic I Pettigrew, and Roozbeh Jafari. Physics-informed neural networks for modeling physiological time series for cuffless blood pressure estimation. *npj Digital Medicine*, 6(1):110–125, 2023a.

Kaan Sel, Deen Osman, Noah Huerta, Arabella Edgar, Roderic I Pettigrew, and Roozbeh Jafari. Continuous cuffless blood pressure monitoring with a wearable ring bioimpedance device. *npj Digital Medicine*, 6(1):59–74, 2023b.

Khemraj Shukla, Juan Diego Toscano, Zhicheng Wang, Zongren Zou, and George Em Karniadakis. A comprehensive and fair comparison between mlp and kan representations for differential equations and operator networks. *Computer Methods in Applied Mechanics and Engineering*, 431: 117290, 2024.

Sidharth SS, Keerthana AR, Anas KP, et al. Chebyshev polynomial-based kolmogorov-arnold networks: An efficient architecture for nonlinear function approximation. *arXiv preprint arXiv:2405.07200*, 2024.

George S Stergiou, Bruce Alpert, Stephan Mieke, Roland Asmar, Neil Atkins, Siegfried Eckert, Gerhard Frick, Bruce Friedman, Thomas Graßl, Tsutomu Ichikawa, et al. A universal standard for the validation of blood pressure measuring devices: Association for the advancement of medical instrumentation/european society of hypertension/international organization for standardization (aami/esh/iso) collaboration statement. *Hypertension*, 71(3):368–374, 2018.

Rui Wang, Mengshi Qi, Yingxia Shao, Anfu Zhou, and Huadong Ma. Pitn: Physics-informed temporal networks for cuffless blood pressure estimation. *IEEE Transactions on Mobile Computing*, 25(4):4876–4889, 2026.

Shiyu Wang, Haixu Wu, Xiaoming Shi, Tengge Hu, Huakun Luo, Lintao Ma, James Y Zhang, and Jun Zhou. Timemixer: Decomposable multiscale mixing for time series forecasting. In *ICLR*, 2024a.

Yizheng Wang, Jia Sun, Jinshuai Bai, Cosmin Anitescu, Mohammad Sadegh Eshaghi, Xiaoying Zhuang, Timon Rabczuk, and Yinghua Liu. Kolmogorov–arnold-informed neural network: A physics-informed deep learning framework for solving forward and inverse problems based on kolmogorov–arnold networks. *Computer Methods in Applied Mechanics and Engineering*, 433: 117518, 2025.

Zhi Wang, Beihong Jin, Fusang Zhang, Siheng Li, and Junqi Ma. Uwb-enabled sensing for fast and effortless blood pressure monitoring. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 8(2):1–26, 2024b.

Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long. Timesnet: Temporal 2d-variation modeling for general time series analysis. In *ICLR*, 2023.

Chenhui Xu, Dancheng Liu, Yuting Hu, Jiajie Li, Ruiyang Qin, Qingxiao Zheng, and Jinjun Xiong. Sub-sequential physics-informed learning with state space model. In *ICML*, 2025.

Xingyi Yang and Xinchao Wang. Kolmogorov-arnold transformer. In *ICLR*, 2025.

Hyungjun Yoon, Seungjoo Lee, Yu Wu, Xiaomeng Chen, Taiting Lu, Freddy Liu, Taeckyung Lee, Hyeongheon Cha, Haochen Zhao, Gaoteng Zhao, et al. Beyond hearing: Learning task-agnostic exg representations from earphones via physiology-informed tokenization. In *ICLR*, 2026.

Fan Zhang, Shiming Fan, and Hua Wang. Time-tk: A multi-offset temporal interaction framework combining transformer and kolmogorov-arnold networks for time series forecasting. In *Proceedings of the ACM Web Conference*, pp. 7495–7506, 2026.

Zixin Zheng, Yumeng Liang, Rui Lyu, Junjie Bao, Yiwen Huang, Anfu Zhou, Huadong Ma, Jingjia Wang, Xiangbin Meng, Chunli Shao, et al. Bp3: Improving cuff-less blood pressure monitoring performance by fusing mmwave pulse wave sensing and physiological factors. In *Proceedings of the 22nd ACM Conference on Embedded Networked Sensor Systems*, 2024.

Zixin Zheng, He Hao, Yiwen Huang, Rui Lyu, Huadong Ma, Yuqing Zhang, Jing Liu, Chunli Shao, and Anfu Zhou. Advances in principles and technologies of non-mechanical blood pressure monitoring. *npj Cardiovascular Health*, 3(1):7, 2026.

Quan Zhou, Changhua Pei, Fei Sun, Jing Han, Zhengwei Gao, Dan Pei, Haiming Zhang, Gaogang Xie, and Jianhui Li. Kan-ad: Time series anomaly detection with kolmogorov-arnold networks. In *ICML*, 2025.

Guanzhou Zhu, Dong Zhao, Yixuan Song, Zilong Wang, Anfu Zhou, Huadong Ma, Xiangbin Meng, Jingjia Wang, Chunli Shao, and Yida Tang. Ringbp: Towards continuous, comfortable, and generalized blood pressure monitoring using a smart ring. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 9(2):1–24, 2025.

A Reproducibility Details

To support independent reimplementation, this appendix documents the dataset configurations, preprocessing procedures, model architecture, training hyperparameters, and evaluation protocols used in our experiments. Dataset characteristics and preprocessing procedures are provided in Appendix B, while implementation details, cardiovascular descriptor extraction, evaluation metrics, and efficiency benchmarking settings are described in Appendix C. Together with the model formulation and training objective in Section 3, these details specify the complete experimental pipeline used to obtain the reported results.

B Dataset Details

B.1 Comparison with Existing Benchmarks

To contextualize the characteristics of Ring-HGCPT, Table 4 compares it with existing datasets for cuffless BP estimation in terms of the BP perturbation protocol, reference device, sensing modality, cohort size, BP coverage, intra-subject BP variation, and recording duration.

Table 4: Comparison of Ring-HGCPT with existing datasets for cuffless BP estimation. HG and CPT denote the handgrip and cold pressor tests, respectively, while Osc. denotes an oscillometric cuff-based device. BP ranges and mean maximum within-subject differences are reported as SBP/DBP in mmHg. Recording durations are expressed in hours, with values in parentheses indicating valid waveforms retained after quality control. N/A indicates that the corresponding information was not reported.

Dataset	Graphene-HGCPT	Ring-CPT	airBP	RingBP	mmBP+	Ubicon-BP	Ring-HGCPT (Ours)
	2022	2023	2023	2025	2025	2025	2026
BP perturbation	HG & CPT	CPT	Walking & cycling	Hydrostatic pressure	N/A	Running	HG & CPT
Reference device	Finapres NOVA	Finapres NOVA	Cuff-based (Osc.)	Cuff-based (Osc.)	Cuff-based (Osc.)	Cuff-based (Osc.)	Cuff-based (Osc.)
Sensing device	Graphene tattoos	Graphene tattoos	mmWave radar	Smart ring	mmWave radar	Smartphone	Smart ring
Subjects	15	10	41	85	33	50	20
BP range (mmHg)	100–180/50–120	89–213/42–122	N/A	N/A	N/A	N/A	89–148/43–106
Avg. max difference (mmHg)	N/A	N/A	N/A	24/16	N/A	N/A	21/22
Recording duration (h)	15 (4)	N/A	N/A	N/A	N/A	50	20 (4)

As shown in Table 4, existing datasets differ substantially in their sensing modalities, reference devices, and BP perturbation protocols. Graphene-HGCPT and Ring-CPT use wearable graphene sensors with Finapres NOVA as the reference device, whereas more recent datasets primarily use oscillometric cuff measurements together with radar, smart-ring, or smartphone sensing. Ring-HGCPT combines handgrip and cold pressor tests within a single standardized protocol. These controlled interventions induce measurable BP changes within each participant, enabling the evaluation of models across within-subject cardiovascular variations.

Ring-HGCPT covers an SBP/DBP range of 89–148/43–106 mmHg, with mean maximum within-subject differences of 21/22 mmHg. It contains approximately 20 hours of recordings from 20 participants, of which approximately 4 hours were retained as valid waveform data after quality control. In contrast to activity-based protocols involving walking, cycling, or running, Ring-HGCPT uses two distinct controlled BP perturbations. Because the compared datasets differ in their reference devices, sensing modalities, acquisition protocols, and signal-quality criteria, the statistics are intended to provide descriptive context rather than a direct ranking of dataset quality.

B.2 Data Processing and Inference Pipeline

Figure 7 summarizes the end-to-end processing and inference pipeline for Ring-HGCPT. The raw smart-ring PPG signal is segmented into fixed-length waveform windows, from which temporally aligned cardiovascular descriptors are extracted. The waveform segments are processed through multi-scale decomposition and the Dual-Order KAN, while the associated descriptors are incorporated by the Physics-Informed KAN. The resulting model output provides the estimated SBP or DBP value for each input segment.

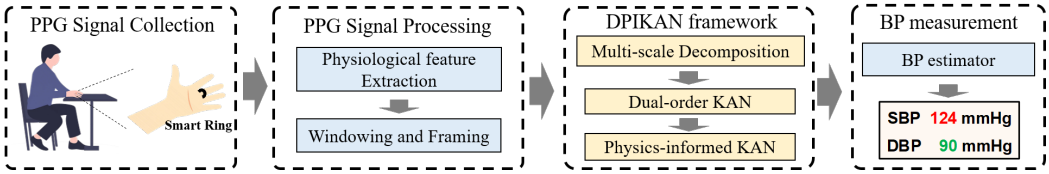

Figure 7: Overview of the DPIKAN data processing and inference pipeline. The workflow begins with raw PPG signal collection via the smart ring, followed by preprocessing stage for physiological feature extraction and temporal framing. These synchronized inputs are then processed by the core DPIKAN architecture to estimate the final Systolic (SBP) and Diastolic (DBP) blood pressure values.

B.3 Participants and Data Statistics

A total of 20 healthy volunteers (10 male and 10 female participants) were recruited from the university campus. All participants self-reported that they were non-smokers with no known history of cardiovascular disease. The cohort had a mean age of 25.4±2.1 years (range: 20–30), a mean height of 170.3±6.5 cm, and a mean weight of 61.6±10.1 kg. Detailed demographic distributions are shown in Fig. 3(d).

To induce controlled within-subject BP variations, participants completed a protocol involving Handgrip (HG) and Cold Pressor Test (CPT) interventions, as detailed in Figure 3. In a post-study questionnaire using a five-point scale, the mean tolerability ratings exceeded 4.6, and the willingness-to-reuse rating was 4.8±0.4, indicating high self-reported acceptability of the protocol.

For each participant, we calculated the difference between the maximum and minimum BP values across all measurements. The resulting mean within-subject differences were 21.0±5.2 mmHg for SBP and 22.0±6.0 mmHg for DBP. Across all participants, SBP ranged from 89 to 148 mmHg and DBP ranged from 43 to 106 mmHg, as shown in Fig. 3(c). PPG signals were sampled at 200 Hz, and the resulting dataset comprises 400 synchronized PPG–BP pairs for evaluating models under controlled within-subject BP variations.

B.4 Experimental Devices

For PPG signal acquisition, we utilized the ring-type sensor previously developed for RingBP study (Zhu et al., 2025). The device, worn on the index finger, is built around the Maxim Integrated MAX30101 module, which includes a green LED (527 nm) and a red LED (660 nm). To account for motion artifacts, the ring also integrates a LIS2DH12TR 3-axis accelerometer. An STM32F103C8T6 microcontroller (MCU) processes the sensor data and transmits it at a 200Hz sampling rate via Bluetooth. Reference BP was measured at the right wrist using a commercial automatic oscillometric wrist blood pressure monitor (Omron). All data from the ring and the cuff-based BP readings were collected and synchronized through a custom-developed Android application on a smartphone. Additionally, a hand dynamometer and an ice water bucket were used to induce blood pressure fluctuations during the experiment, as per the protocol.

B.5 Experiment Protocol

This section outlines our experimental protocol, which was approved by the Institutional Review Board (IRB). Each subject completed a session comprising five identical trials. As detailed in Fig. 3(b), each trial was designed to induce blood pressure variations through a specific sequence of activities. A single trial began with a 1-minute hand grip test, immediately followed by a 1-minute cold pressor test, and concluded with a 3.5-minute rest period. Throughout this process, photoplethysmography (PPG) signals were recorded, and reference Systolic/Diastolic Blood Pressure (SBP/DBP) values were measured using a cuff-based monitor after each activity phase. Prior to the session, participants were prepared for data acquisition as shown in Fig. 3(a), with sensors placed on their fingers and a blood pressure cuff on their wrist.

B.6 Signal Processing

The data acquisition and preparation pipeline is illustrated in Fig. 7. The workflow commences with PPG Signal Collection module, where raw physiological waveforms are captured via the smart ring's optical sensors. Subsequently, the raw data enters the PPG Signal Processing stage, where it undergoes a bifurcated pre-processing strategy to construct the model inputs. First, the Physiological Feature Extraction module derives explicit hemodynamic priors specifically, Heart Rate (HR), Pulse Wave Velocity proxy (PWV), and amplitude metrics from the raw waveforms. Simultaneously, the Windowing and Framing module applies a sliding window technique to segment the continuous time series into fixed-length sequences. These two distinct data streams the explicit physical features and the windowed deep signal frames are then aligned to serve as inputs for the subsequent inference network.

C Implementation Details

C.1 Cardiovascular Descriptor Extraction

For each signal segment $\mathbf{X}_i$, let $s_i(t)$ denote the corresponding preprocessed and polarity-aligned physiological waveform. We extract three label-free cardiovascular descriptors, $\mathbf{u}_i = [u_i^1, u_i^2, u_i^3]^\top$, as defined below.

Pulse-amplitude change. The pulse-amplitude descriptor is defined as

$$u_i^1 = s_{i,\max} - s_{i,\min}. \tag{9a}$$

It measures the peak-to-trough excursion within the segment and serves as a proxy for the extent of arterial expansion. For BioZ measurements, $s_{i,\max}$ and $s_{i,\min}$ denote the maximum and minimum values of the polarity-aligned impedance waveform. The same definition is applied to the preprocessed PPG and mmWave waveforms.

PWV proxy. The morphology-derived PWV proxy is defined as

$$u_i^2 = \frac{1}{|t_{F,i} - t_{B,i}|}. \tag{9b}$$

Here, $t_{F,i}$ denotes the trough landmark in the first-derivative waveform, and $t_{B,i}$ denotes the zero-crossing landmark associated with the ascending transition. This reciprocal interval characterizes the temporal separation between waveform landmarks associated with the forward-traveling and reflected pulse components and is used as a proxy rather than a direct measurement of PWV.

Heart rate. The heart-rate descriptor is computed as

$$u_i^3 = \frac{60}{|t_{J,i} - t_{A,i}|}. \tag{9c}$$

Here, $t_{J,i}$ denotes the peak landmark in the first-derivative waveform, and $t_{A,i}$ denotes the zero-crossing landmark associated with the descending transition. When the landmark times are expressed in seconds, the factor of 60 converts the resulting rate to beats per minute.

The three descriptors are temporally aligned with their corresponding signal segments before being provided to the Physics-Informed KAN.

C.2 Metrics

To evaluate the model's performance, we follow the standard practices in blood pressure estimation and utilize four key metrics: the Root Mean Square Error (RMSE), Pearson's correlation coefficient (r), Mean Error (ME), and the Standard Deviation of the Error (SDE).

Let K denote the number of evaluated subjects and M_k denote the number of held-out test samples for subject k. The subject-level RMSE is defined as:

$$\mathrm{RMSE}_k = \sqrt{\frac{1}{M_k}\sum_{i=1}^{M_k}(\hat{y}_{k,i} - y_{k,i})^2}, \tag{10}$$

where $\hat{y}_{k,i}$ and $y_{k,i}$ are the predicted and reference BP values of the i-th test sample from subject k, respectively. The reported RMSE is the macro-average across subjects:

$$\mathrm{RMSE} = \frac{1}{K}\sum_{k=1}^{K}\mathrm{RMSE}_k. \tag{11}$$

Similarly, Pearson's correlation coefficient is first calculated separately for each subject:

$$r_k = \frac{\sum_{i=1}^{M_k}(\hat{y}_{k,i} - \bar{\hat{y}}_k)(y_{k,i} - \bar{y}_k)}{\sqrt{\sum_{i=1}^{M_k}(\hat{y}_{k,i} - \bar{\hat{y}}_k)^2}\sqrt{\sum_{i=1}^{M_k}(y_{k,i} - \bar{y}_k)^2}}, \tag{12}$$

where $\bar{\hat{y}}_k$ and $\bar{y}_k$ are the subject-level means. The reported correlation is the arithmetic macro-average:

$$\mathrm{Corr.} = \frac{1}{K}\sum_{k=1}^{K} r_k. \tag{13}$$

For subject-specific evaluation on the public datasets, $M_k = N_k - S_k$, where N_k and S_k denote the total number of labeled samples and the number of training samples for subject k, respectively. Under LOSO evaluation, M_k is the number of test samples from the participant held out in the k-th fold.

Finally, we report the Mean Error (ME) and Standard Deviation of the Error (SDE), two error statistics commonly used in blood pressure device evaluation (Stergiou et al., 2018). Let $\mathcal{A}$ denote the set of all aligned prediction–reference pairs in the test data. For each aligned point $j \in \mathcal{A}$, the prediction error is defined as

$$e_j = \hat{y}_j - y_j, \tag{14}$$

where $\hat{y}_j$ and y_j are the predicted and reference BP values, respectively.

ME is calculated over all aligned test points as

$$\mathrm{ME} = \frac{1}{|\mathcal{A}|}\sum_{j\in\mathcal{A}} e_j. \tag{15}$$

SDE characterizes the dispersion of these errors around the overall mean error:

$$\mathrm{SDE} = \sqrt{\frac{1}{|\mathcal{A}|}\sum_{j\in\mathcal{A}}(e_j - \mathrm{ME})^2}. \tag{16}$$

Thus, RMSE and Pearson's correlation coefficient are first calculated separately for each subject and then macro-averaged, whereas ME and SDE are calculated directly from all aligned prediction–reference pairs in the test data.

For descriptive context, Criterion 1 of the AAMI/ESH/ISO framework commonly uses numerical criteria of $|\mathrm{ME}| \leq 5$ mmHg and $\mathrm{SDE} \leq 8$ mmHg. We use these thresholds only as contextual references; the present evaluation does not constitute formal AAMI/ESH/ISO device validation.

D Additional Results

D.1 Mean Error and Error Variability

We additionally report the mean error (ME) and standard deviation of errors (SDE) to characterize prediction bias and error variability across datasets. For descriptive context, we compare these

Table 5: Mean error (ME) and standard deviation of errors (SDE) across datasets. The AAMI/ESH/ISO numerical thresholds are included for descriptive reference and do not constitute formal standards validation.

Dataset	Type	ME (mmHg)	SDE (mmHg)
Ring-HGCPT	SBP	1.08	5.41
	DBP	0.09	4.81
Ring-CPT	SBP	0.02	7.68
	DBP	0.59	6.98
Graphene-HGCPT	SBP	-0.09	7.21
	DBP	-0.64	7.77
Blumio (PPG)	SBP	-0.08	5.07
	DBP	-0.23	4.21
Blumio (mmWave)	SBP	-1.13	6.50
	DBP	0.29	4.31

Table 6: Comparison with representative baselines on Ring-HGCPT in terms of mean error (ME) and standard deviation of errors (SDE).

Models	SBP (mmHg)		DBP (mmHg)	
	ME	SDE	ME	SDE
iTransformer	-6.67	12.94	-4.45	10.92
TimesNet	-0.71	5.67	0.24	6.51
TimeMixer	-1.33	11.85	0.33	10.08
TimeKAN	0.17	5.81	0.91	5.87
ResNet1D	3.22	11.14	2.64	9.27
PINN	-0.17	5.68	-0.27	5.71
PITN	-0.22	18.42	1.00	14.35
DPIKAN (Ours)	1.08	**5.41**	**0.09**	**4.81**

results with the aggregate numerical criteria commonly reported in AAMI/ESH/ISO evaluations, namely $|\mathrm{ME}| \leq 5$ mmHg and $\mathrm{SDE} \leq 8$ mmHg. Because our datasets and evaluation protocols were not designed as formal device-validation studies, this comparison should not be interpreted as establishing compliance with the standard.

Table 5 summarizes the results across the evaluated datasets and sensing modalities. DPIKAN falls within the referenced numerical thresholds for both SBP and DBP in all settings, indicating low average bias and controlled error variability across BioZ, PPG, and mmWave measurements. These results should nevertheless be interpreted in light of the different reference devices, participant cohorts, and evaluation protocols used by the datasets.

To further compare error variability across methods, Table 5 reports ME and SDE on Ring-HGCPT under leave-one-subject-out evaluation. Several baselines exhibit larger error variability when applied to participants unseen during training, whereas DPIKAN obtains the lowest SDE for both SBP and DBP among the evaluated methods. Together with the RMSE and correlation results, these findings suggest that decomposition-aware capacity allocation and cardiovascular descriptor guidance contribute to more consistent BP estimation under subject-independent evaluation.

D.2 Qualitative Results

We provide a qualitative analysis of our model's performance in Fig. 8, using a representative subject from the graphene-HGCPT dataset. The figure presents a beat-to-beat comparison of the SBP and DBP estimations from DPIKAN against those from several baseline models. The visualizations clearly show that DPIKAN's predictions track the ground truth BP values with high fidelity, accurately capturing the temporal variations in the subject's cardiovascular dynamics. In contrast, the

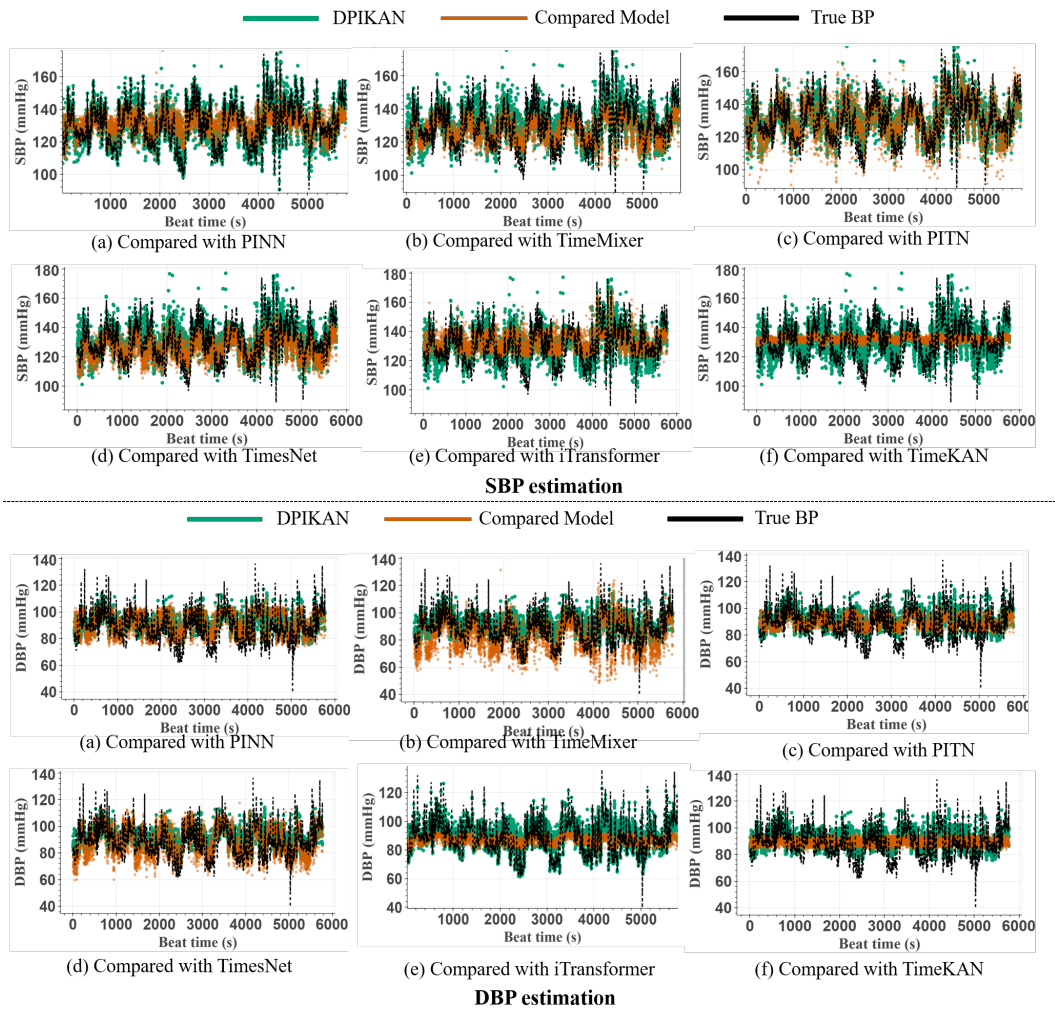

Figure 8: Qualitative beat-to-beat BP tracking comparison on a representative subject. DPIKAN predictions are shown in green, baseline predictions in orange, and reference BP in black. The top and bottom panels show SBP and DBP estimation, respectively.

predictions from the baseline models exhibit more significant deviations and errors, often failing to consistently model the BP trends. This visual evidence reinforces the quantitative results presented earlier, demonstrating the superior ability of our proposed method to accurately model beat-to-beat blood pressure.

To interpret the internal efficacy of our proposed framework, we visualize the feature evolution at distinct stages in Fig. 6 using t-SNE. First, Fig. 6(a) confirms that the Multi-Scale Decomposition module effectively separates the raw signal into a Low-Frequency Trend (blue) and a High-Frequency Residual (orange) based on temporal scales. However, the latent visualization in Fig. 6(b) reveals a critical insight: despite this signal-level separation, the learned representations initially remain entangled (highlighted by the overlap regions). This indicates that temporal decomposition alone acts only as a preliminary step and is insufficient to fully distinguish the underlying physiological semantics in the high-dimensional feature space. In contrast, Fig. 6(c) demonstrates the indispensable role of the Dual-Order KAN. By assigning distinct polynomial orders to process the decomposed streams, the network successfully resolves this entanglement, projecting the components into distinct, orthogonal clusters. Crucially, this transition from the entangled state in (b) to the clear separation in (c) confirms the synergistic necessity of our two-stage design. This explicit

feature disentanglement minimizes mutual interference between rapid cardiac dynamics and slow hemodynamic trends, directly contributing to the superior estimation accuracy reported in Table 1.

D.3 PRE-PROCESSING

For consistency with previous work in graphene-HGCPT, ring-CPT, and Blumio dataset, our experiments adhere to a minimal training criterion (Sel et al., 2023a; Wang et al., 2026). We first synchronized the bioimpedance and reference blood pressure data. A second-order low-pass Butterworth filter with a 6 Hz cutoff frequency was then applied to the raw bioimpedance stream. This step effectively isolates blood volumetric changes tied to the heart rate while removing extraneous noise. Next, we segmented the filtered signal into individual heartbeats by identifying fiducial points. Specifically, we located the consecutive maximum slope points of the systolic wave by detecting the peak amplitudes in the signal's first inverted derivative. A validation step ensured that the detected peaks corresponded to the most prominent slopes in the original signal and maintained a plausible periodicity around the subject's heart rate. This pre-processing procedure is in consistency with prior work (Wang et al., 2026). We trained separate, personalized models for each subject to ensure precise predictions. The training data for each model was selected using the minimal training criterion outlined in (Sel et al., 2023a). This involves partitioning the labeled data into 1 mmHg-wide bins based on the full range of SBP and DBP values. The initial training set is then formed by randomly selecting one data point from each bin. Finally, each model's performance was assessed against a separate test set containing reference blood pressure values that were not used during the training process. For consistency, the same data partitioning strategy and experimental configurations were maintained for all models.

For our proposed Ring-HGCPT dataset, we adopt a direct input strategy with minimal preprocessing. The raw data together with the extracted physiological features from all four channels are used as model inputs. To evaluate subject-independent performance, we employ Leave-One-Subject-Out (LOSO) cross-validation, where the model is trained on all subjects except one and evaluated on the held-out subject; this procedure is repeated until each participant has served as the test subject. The evaluation protocol is adapted to the reference BP acquisition characteristics of each dataset. Public datasets such as Graphene-HGCPT provide dense, continuous beat-to-beat reference BP measurements, which support subject-specific personalized training following prior work. In contrast, Ring-HGCPT uses a commercial cuff-based monitor and therefore provides sparse and intermittent reference BP labels for each participant. Under this setting, reliable subject-specific training is difficult due to the limited number of labeled samples per individual. We therefore adopt LOSO evaluation for Ring-HGCPT to assess model performance on subjects whose data are entirely unseen during training.

E LIMITATIONS AND BROADER IMPACT

Ring-HGCPT currently includes a relatively small cohort of young, healthy participants under controlled cardiovascular perturbations. Further evaluation on larger and more diverse populations, including individuals with cardiovascular conditions and recordings collected in free-living environments, is therefore needed. Nevertheless, DPIKAN provides a compact approach to physiological representation learning across heterogeneous sensing modalities, with potential value for wearable and continuous health monitoring. The current system is intended for research rather than clinical diagnosis, and clinical deployment would require additional validation.