Med2ECG: Medical-Guided BCG-To-ECG Reconstruction for Diverse Populations

Lin Chen
HKUST (GZ)
Guangzhou, China
lchen297@connect.hkust-gz.edu.cn

Yandao Huang
HKUST
Clear Water Bay, Hong Kong
yhuangfg@connect.ust.hk

Chenggao Li
HKUST
Clear Water Bay, Hong Kong
clicx@connect.ust.hk

Jun Chen
HKUST (GZ)
Guangzhou, China
jchen512@connect.hkust-gz.edu.cn

Shuxin Zhong*
HKUST (GZ)
Guangzhou, China
shuxinzhong@hkust-gz.edu.cn

Minghui Qiu
HKUST (GZ)
Guangzhou, China
mqiu585@connect.hkust-gz.edu.cn

Chunzhen Guo
HKUST (GZ)
Guangzhou, China
cguo546@connect.hkust-gz.edu.cn

Yi Wang
HKUST (GZ)
Guangzhou, China
ywang183@connect.hkust-gz.edu.cn

Qian Zhang
HKUST
Clear Water Bay, Hong Kong
qianzh@cse.ust.hk

Kaishun Wu*
HKUST (GZ)
Guangzhou, China
wuks@hkust-gz.edu.cn

ABSTRACT

Continuous ECG monitoring is vital for early detection of arrhythmias and other cardiac abnormalities—especially during sleep, when symptoms often go unnoticed—yet existing solutions remain expensive, obtrusive, and impractical for long-term daily use. Ballistocardiography (BCG)—a passive, contactless modality that captures cardiac-induced body motion—offers a compelling alternative. However, prior efforts treat ECG reconstruction as waveform regression, leading to overfitting to individual-specific or posture-dependent artifacts. Inspired by the fact that ECG and BCG reflect parallel structures in cardiac event sequences (e.g., P/QRS/T-waves vs. I/J-waves), we design Med2ECG: a structure-aligned system that reconstructs ECG from high-fidelity BCG by explicitly aligning their latent physiological events. Med2ECG incorporates three key designs: (i) *Multi-Scale Feature Extractor* captures hierarchical temporal dynamics, preserving clinically relevant fine-grained features; (ii) *Shared-Personalized Experts* employs a Mixture-of-Experts (MoE) to adaptively disentangle signal variations due to individual and environmental factors; (iii) *Medical-Informed Strategies* introduces a diagnostic-driven multi-objective loss, integrating structural alignment, morphological fidelity, and landmark-aware supervision to preserve clinically critical intervals. Experiments across public and self-collected in-hospital datasets (20 healthy individuals and 10 patients with diverse cardiovascular conditions), Med2ECG achieves > 0.92 Pearson correlation, < 15% amplitude error, and precise PR/QRS/QT/RR interval estimation within 5–20 ms—demonstrating strong generalization across subjects, postures, and environments.

CCS CONCEPTS

• **Human-centered computing → Ubiquitous and mobile computing**.

KEYWORDS

Healthcare, ECG Waveform Reconstruction, Ballistocardiogram

ACM Reference Format:
Lin Chen, Yandao Huang, Chenggao Li, Jun Chen, Shuxin Zhong, Minghui Qiu, Chunzhen Guo, Yi Wang, Qian Zhang, and Kaishun Wu. 2026. Med2ECG: Medical-Guided BCG-To-ECG Reconstruction for Diverse Populations. In *ACM/IEEE International Conference on Embedded Artificial Intelligence and Sensing Systems (SenSys '26), May 11–14, 2026, Saint Malo, France.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3774906.3800463

1 INTRODUCTION

Cardiovascular diseases (CVDs) are the leading global cause of death, claiming over 17.9 million lives annually [39]. For individuals with obstructive sleep apnea (OSA)—a common comorbidity—up to 46% of sudden cardiac deaths (SCDs) occur during sleep [14], when symptoms typically go undetected. These silent events underscore the need for continuous, unobtrusive cardiac monitoring that remains effective in unsupervised conditions such as nighttime rest. Electrocardiography (ECG) is the clinical gold standard for assessing cardiac electrical activity [28]. However, current solutions—such

*Corresponding authors.

SenSys '26, May 11–14, 2026, Saint Malo, France

ACM ISBN 979-8-4007-2309-4/26/05
https://doi.org/10.1145/3774906.3800463

as Holter monitors [6] and implantable loop recorders [3]—are expensive and obtrusive [34], making them unsuitable for long-term deployment in everyday settings [15, 44, 45].

To expand accessibility, recent efforts have focused on reconstructing ECG from more ubiquitous signals such as photoplethysmography (PPG) [25, 31, 33, 46] or inertial data [4] captured via consumer-grade wearables [35]. While more affordable and widely deployable, these approaches remain contact-dependent and are vulnerable to motion artifacts and adherence issues in free-living conditions [32, 45]. In response, *contactless sensing* has emerged as a compelling paradigm. For instance, mmWave radar tracks chest-wall micro-movements non-invasively [40, 41], yet suffers from line-of-sight dependency and environmental interference [43]. Ballistocardiography (BCG), by contrast, measures cardiac-induced body recoil via ambient physical contact with surfaces like beds, chairs, or cushions [13, 29, 36]. This allows for *naturally embedded, alignment-free, and zero-effort sensing*—making BCG particularly suited for scalable, long-term deployment in home environments.

Early work [29] applied wavelet decomposition to uncover coarse correlations between BCG and ECG signals. Subsequent approaches leveraged deep neural networks to enhance waveform fidelity—RecAUNet [36] used an attention-augmented UNet to reconstruct fine-grained ECG subwaves (P, QRS, T), while [44] employed a BiLSTM to model temporal dynamics across cardiac cycles. Despite promising, these methods treat ECG reconstruction as a *point-wise regression task*, mapping BCG signals directly to ECG waveforms. This formulation assumes a stable, modality-consistent mapping—overlooking the fundamental physiological gap between the two: ECG reflects electrical conduction at the skin surface, while BCG captures mechanical recoil transmitted through body tissues and external environments. Such mismatch makes models prone to overfitting subject-specific waveform features or postural variations—ultimately compromising generalization in real-world settings.

To move beyond this limitation, we propose a structure-aligned formulation grounded in the shared physiological events underlying both signals. Although ECG and BCG originate from distinct modalities—electrical conduction vs. mechanical recoil—they reflect the same canonical sequence of cardiac events (e.g., P-wave, QRS complex, T-wave), with temporally aligned mechanical counterparts (e.g., H-, I-, J-waves). This shift not only narrows the modality gap but also enhances robustness across subjects and postures.

This intuitive idea confronts three fundamental challenges:

- **C1: Multi-scale temporal hierarchy.** BCG signals exhibit a rich temporal hierarchy, reflecting layered cardiac dynamics across multiple timescales. At the fine scale (tens of milliseconds), sharp transients—such as J-waves and oscillations within the J–K segment—encode rapid mechanical impulses critical for heartbeat localization and QRS alignment. At the mid scale (200–500 ms), the I–J–K complex captures the full mechanical response of a cardiac cycle, embedding electromechanical coupling. At the long scale (seconds), heart rate variability (HRV) and respiratory modulation reflect inter-cycle coherence—essential for rhythm-aware interpretation and downstream tasks like arrhythmia screening.
- **C2: Morphological variability.** BCG morphology varies sharply across users, postures, and sensor placement. The same cardiac event may appear time-shifted, attenuated, or morphologically deformed—breaking temporal alignment and semantic correspondence across conditions. Such distortions undermine the model's ability to learn consistent mappings, posing a major obstacle to generalization.
- **C3: Diagnostic fidelity.** A clinically usable ECG is not merely plausible in appearance—it must preserve high-resolution structures such as P-wave onset, QRS steepness, and T-wave recovery. These features carry critical diagnostic value, and any loss in morphological integrity can render the signal inadequate for real-world cardiac screening or interval-based clinical decision-making.

To address those challenges, we introduce `Med2ECG`, a generalizable, structure-aligned system for ECG reconstruction from high-fidelity BCG signals collected via self-designed, low-cost optical fiber sensors. Unlike prior waveform-mimicking approaches, `Med2ECG` performs *physiology-grounded inference* — enabling robust, interpretable, and clinically meaningful ECG reconstruction in free-living, non-contact scenarios. Specifically, to address **C1**, Multi-Scale Feature Extractor preserves high-resolution detail by encoding the temporal hierarchy of BCG signals, which captures sharp J-wave transients (milliseconds), full-cycle I–J–K complexes (hundreds of milliseconds), and long-range modulations like HRV and respiration (seconds). To tackle **C2**, Shared-Personalized Experts improves generalization across subjects by adopting a Mixture-of-Experts (MoE) architecture, which isolates invariant physiological patterns from nuisance factors. To ensure **C3**, Medical-Informed Strategies enforces diagnostic fidelity via three physiologically informed objectives: (1) *Structure-Aligned Loss* corrects temporal jitter, ensuring morphological consistency across shifts; (2) *Wave-Aware Loss* redistributes learning focus toward diagnostically underrepresented subwaves; (3) *Diagnostic Anchor Loss* preserves temporal landmarks (e.g., onset, peak, offset) essential for computing diagnostic intervals such as PR and QT. Our contributions are summarized as follows:

- We are the first one to shift BCG-to-ECG reconstruction from direct waveform regression to a *physiology-informed alignment* of latent cardiac event structures (e.g., I–J–K vs. P–QRS–T). This structural reframing captures the shared temporal semantics between modalities and supports robust generalization across users and contexts.
- We develop `Med2ECG`, a generalizable ECG reconstruction system through three technical designs: (i) Multi-Scale Feature Extractor captures the temporal hierarchy of cardiac activity—from millisecond-scale J-waves to full-cycle I–J–K complexes and longer-range rhythms such as HRV and respiration—ensuring morphological detail is preserved across scales; (ii) Shared-Personalized Experts employs an MoE architecture to disentangle latent physiological signals from subject-specific and environmental noise, enhancing generalization; and (iii) Medical-Informed Strategies enforces diagnostic consistency through a tri-objective loss that integrates structure alignment, waveform sensitivity, and landmark-aware supervision-bridging morphological fidelity with clinical interpretability.
- We conduct extensive experiments on both public and in-hospital datasets. On the public dataset, `Med2ECG` achieves 0.916 correlation and 14.72% amplitude error, while preserving

PR/QRS/QT/RR intervals within 5–20 ms. On a self-collected in-hospital dataset (20 healthy individuals and 10 patients with diverse CVDs), Med2ECG maintains 0.927 correlation and 12.82% error under subject-independent conditions — validating strong generalizability across clinical and ambient settings.

To support the community, we make our source code publicly available at https://github.com/leenchen0/Med2ECG.

2 PRELIMINARIES

2.1 BCG: Contactless Mechanical Tracing

BCG captures the subtle mechanical recoil of the body caused by cardiac ejection. As the heart contracts and relaxes, blood accelerates through the aorta, generating reaction forces that induce minute body displacements. These movements are transduced into a waveform, BCG, which reflects the timing and strength of cardiac activity, enabling non-invasive cardiovascular monitoring via everyday surfaces (e.g., beds, chairs) [10]. To formalize the physical origin of the BCG signal, Kim et al. [17] modeled the aorta as two connected tubes—ascending and descending—driven by blood pressure gradients (shown in Figure 1). The net force transmitted to the body can be expressed as:

$$F_{BCG}(t) = A_D[P_1(t) - P_2(t)] - A_A[P_0(t) - P_1(t)], \quad (1)$$

where $P_0(t)$, $P_1(t)$, and $P_2(t)$ are pressure waveforms at key vascular junctions, and A_A, A_D are cross-sectional areas of the aortic segments. The resulting BCG waveform exhibits distinct peaks (H/I/J/K/L/M/N), each tied to physiological events in the cardiac cycle. For example, the I wave corresponds to the initial impulse from blood acceleration, while the prominent J wave aligns with peak ventricular ejection. This sequence encodes rich temporal semantics, serving as *a mechanical proxy for cardiac phases.*

2.2 ECG: Canonical Electrical Reference

Electrocardiography (ECG) is the clinical cornerstone for evaluating cardiac rhythm and conduction. By placing electrodes at key body sites, it records the heart's electrical impulses during depolarization and repolarization cycles. The resulting waveform, shown in Figure 1, comprises distinct peaks—P, QRS, and T—corresponding to atrial depolarization, ventricular contraction, and relaxation. These fine-grained morphological patterns serve as *critical biomarkers for diagnosing cardiac abnormalities and guiding treatment.*

2.3 BCG-to-ECG: From Correlation to Structured Mapping

BCG and ECG encode distinct yet temporally aligned aspects of the cardiac cycle—mechanical recoil and electrical conduction. Figure 1 shows that key phases in the ECG map to corresponding deflections in the BCG: the P wave (atrial contraction) aligns with the H wave, the QRS complex (ventricular depolarization) with the I wave, and the J wave reflects the peak of ventricular ejection during the ST segment. Subsequent waves (K–M) capture mechanical responses to late ejection, relaxation, and rapid filling—highlighting a shared physiological sequence underlying cross-modal learning.

Although BCG and ECG exhibit aligned physiological structures, they are not interchangeable. BCG captures the mechanical consequences of cardiac cycles—exhibiting rich temporal dynamics—but lacks direct access to the heart's electrical activity. As a result, it cannot reveal electrophysiological markers such as P-wave onset, QRS steepness, or QT interval, which are essential for diagnosing arrhythmias, conduction disorders, and ischemia. Nevertheless, unlike ECG, BCG can be collected passively and unobtrusively in daily environments, offering a practical pathway toward effortless, long-term monitoring.

This motivates our goal: to *bridge the **physiological** structure of BCG with the diagnostic fidelity of ECG*, enabling **clinically meaningful, zero-effort cardiac monitoring** that scales to real-world environments without compromising interpretability. While prior works [22, 36, 41] have mainly exploited BCG–ECG correlations through statistical waveform mappings, they overlook the structured physiological alignment that governs cardiac dynamics. In contrast, **Med2ECG advances from mere correlation to structure-informed physiological mapping**, explicitly aligning latent BCG events (I–J–K) with diagnostic ECG landmarks (P–QRS–T). This alignment serves as a physiological scaffold that guides representation learning to preserve both temporal precision and morphological fidelity, ensuring that reconstructed signals remain diagnostically interpretable rather than superficially similar.

3 THE DESIGN OF MED2ECG

As shown in Figure 2, we introduce Med2ECG, a structure-aligned and generalizable system for ECG reconstruction from high-fidelity BCG signals. The system comprises five key components that span *sensing, signal conditioning, representation learning, and physiologically grounded supervision*:

- **Optical Microfiber Sensors** enable effective cardiac monitoring through their high sensitivity, spatial resolution, and immunity to electromagnetic interference (Sec. 3.1).
- **Data Preprocessing** mitigates noise and motion artifacts via a two-stage pipeline: *Bandpass Filtering* suppresses non-cardiac frequencies (e.g., respiration, high-frequency noise), while *Motion Artifact Removal* eliminates corrupted segments to ensure clean inputs (Sec. 3.2).
- **Multi-Scale Feature Extractor** addresses the multi-scale nature of cardiac signals by extracting features across a broad temporal hierarchy—from sharp J-wave transients to long-range modulations such as HRV—preserving both local detail and global rhythm (Sec. 3.3).
- **Shared-Personalized Experts** tackles inter-subject variability via an MoE architecture that dynamically routes inputs to specialized experts, enabling domain-adaptive modeling across diverse physiological profiles (Sec. 3.4).
- **Medical-Informed Strategies** bridges morphological fidelity and clinical interpretability by using three physiologically informed objectives: structure alignment, wave-aware weighting, and landmark-guided supervision (Sec. 3.5).

3.1 Optical Microfiber Sensors

To enable high-fidelity, non-intrusive BCG monitoring, we design and implement a novel fiber-optic sensing system based on polydimethylsiloxane (PDMS) flexible optical fibers, as shown in Figure3. Compared to conventional electronic sensors, optical fiber sensors

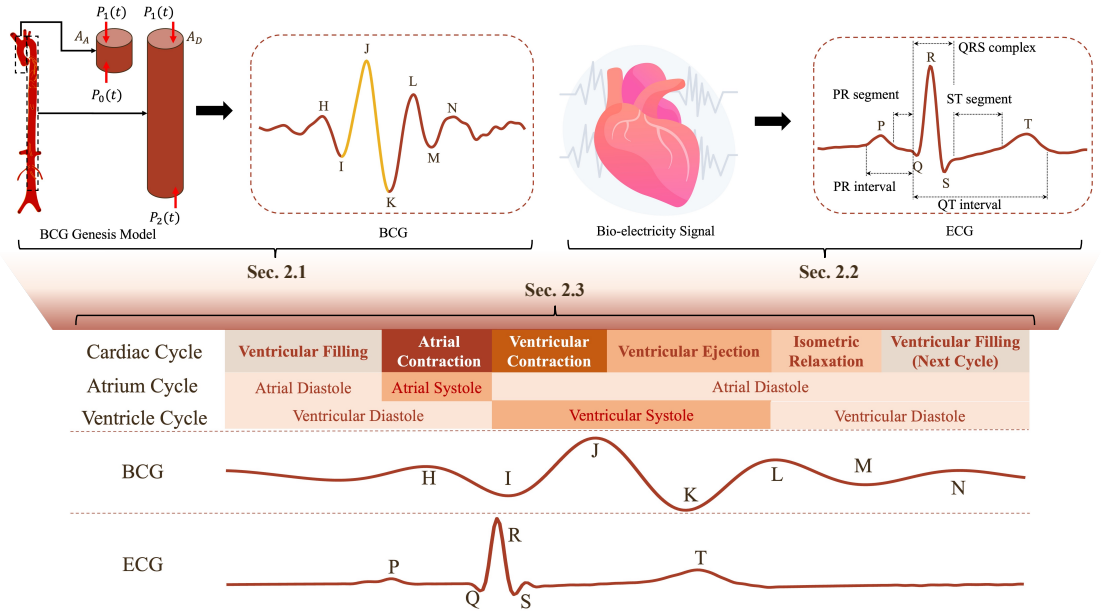

Figure 1: The relationship between BCG and ECG.

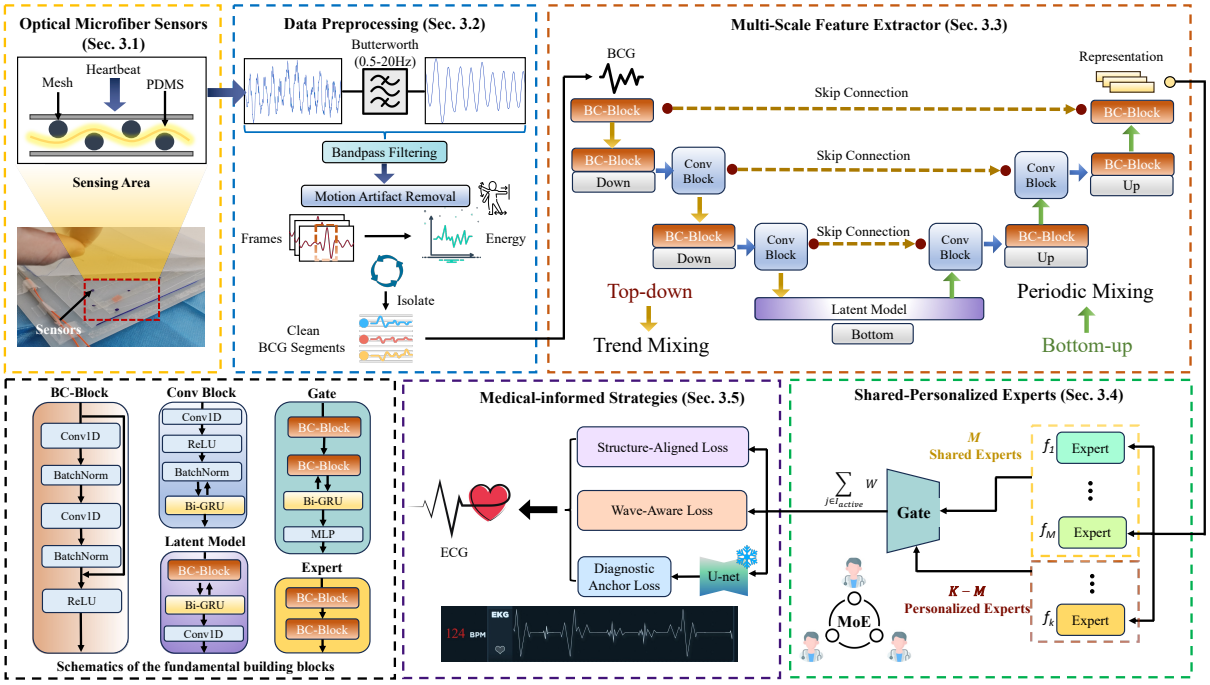

Figure 2: The framework of Med2ECG.

offer several key advantages: low-cost, strong immunity to electromagnetic interference, excellent biocompatibility, high sensitivity, small form factor, and passive operation without electrical contact. These characteristics make them particularly suitable for long-term physiological monitoring with minimal user burden.

The core sensing component consists of a flexible, stretchable multimode optical fiber embedded in PDMS, which ensures mechanical compliance and strong fabric adhesion[42]. This configuration allows the sensor to effectively capture tiny mechanical perturbations induced by cardiac activity, enabling high-resolution BCG signal extraction. Unlike rigid or fragile silica-based fibers, the use

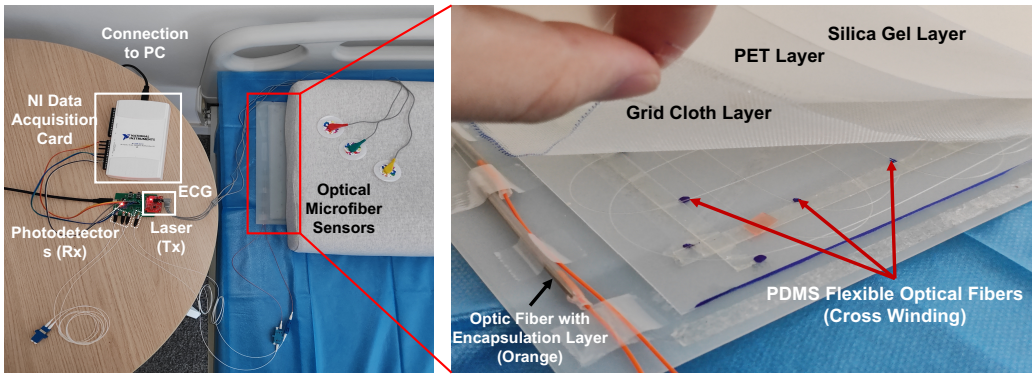

Figure 3: The prototype of Med2ECG, as a standalone device inserted under the pillow.

of soft polymer optical fibers significantly enhances comfort, safety, and robustness in dynamic human-interaction scenarios. Even in the case of mechanical deformation or damage, the sensor remains safe for daily use and does not pose any injury risk. The sensor can be embedded in a variety of objects used in daily life, such as pillows, cushions, mattresses and other facilities. The deployment method is convenient and does not affect the user's normal use or resting state. The sensors can work stably whether the user is lying or sitting.

The sensing mechanism is based on the microbend loss effect, wherein subtle mechanical perturbations induced by heartbeat-driven body movements cause localized deformations in the optical fiber, resulting in attenuation of light transmission. According to the macrobend loss model and Marcuse's curvature loss theory [21], the relationship between the transmitted optical power and the macrobend radius R for a specific propagation mode is given by:

$$\frac{P_{\text{out},\omega}}{P_{\text{in},\omega}} = k_1(\omega) \cdot \frac{1}{\sqrt{R}}\, e^{k_2(\omega)R}, \tag{2}$$

where $P_{\text{in},\omega}$ and $P_{\text{out},\omega}$ denote the input and output light intensity at angular frequency ω, and $k_1(\omega)$, $k_2(\omega)$ are mode-dependent constants. The macrobend radius R is further related to the bending angle θ by the geometric relationship $\theta = \frac{l}{R}$, where l is the arc length subtended by angle θ. Combining this with Eq. (2), we can derive the dependency of output intensity on the bending angle for a single mode:

$$P_{\text{out},\omega} = P_{\text{in},\omega}\, k_1(\omega)\, \sqrt{\theta}\, e^{k_2(\omega)\,\theta^{-1}}, \tag{3}$$

Since a multi-mode optical fiber is used in our system, the total transmitted intensity is the sum of modal contributions:

$$P_{\text{out}} = \sum_i P_{\text{out}}^{(i)} = \sum_i P_{\text{in},\omega}^{(i)}\, k_1^{(i)}(\omega)\, \sqrt{\theta}\, e^{k_2^{(i)}(\omega)\,\theta^{-1}}, \tag{4}$$

By continuously and precisely tracking these light intensity variations across multiple modes, the system enables *real-time extraction of the BCG waveform*, effectively capturing subtle cardiovascular dynamics through optical modulation.

The hardware system comprises an optical transceiver module and a multilayer flexible sensor mat. The optical circuit includes a 1550nm laser source and three high-sensitivity photodetectors covering the 1100–1650nm spectral range. In this work, only one photodetector is actively used; the remaining two are reserved for future multi-channel extensions. The sensing mat measures 50 cm × 30 cm and utilises a multi-layer composite structure design. The outermost layer is composed of a soft silicone material, which ensures the comfort of the skin during use. The intermediate layer contains PET and mesh cloth, which are used to fix the distribution of the flexible fiber optic sensors. The core layer consists of approximately 7 m of flexible PDMS optical fiber. Unlike previous sensors, we have also adopted a cross-winding sensing structure to to enhance the bending degree of the optical fiber and improve the response sensitivity[12, 18].These fibers are connected to the lasers and detectors through LC interfaces to form a complete signal perception link.

To validate the accuracy of the BCG signals, we simultaneously record single-lead ECG signals using an AD8232-based ECG monitoring sensor [1] which is widely used for ECG measurement [9, 27] and ground truth reference [4, 36]. Both the BCG and ECG sensors are synchronized via a shared NI 6210 data acquisition card, allowing time-aligned data collection. The sampling rate for both signals is set to 1000 Hz and later downsampled to 250 Hz for further analysis, as the relevant cardiac information mainly lies below 50 Hz and downsampling helps reduce computational cost and noise.

3.2 Data Preprocessing

We implement a lightweight two-stage pre-processing pipeline to retain clean BCG signals by filtering non-cardiac frequencies and removing motion-contaminated segments.

3.2.1 Bandpass Filtering. BCG sensors capture a wide range of mechanical vibrations, including not only fine-grained cardiac-induced micro-movements but also lower-frequency respiratory artifacts [19, 20]. To isolate cardiac components, we apply a third-order Butterworth bandpass filter (0.5–20 Hz), which matches the canonical frequency range of ECG-correlated mechanical activity [13, 36]. Formally, the filtered segment $\tilde{\mathbf{x}}_{\text{bf}}$ is computed as:

$$\tilde{\mathbf{x}}_{\text{bf}} = \text{Butterworth}^{(3)}_{[0.5\text{-}20\,\text{Hz}]}(\mathbf{x}), \tag{5}$$

where $\mathbf{x} = [x_1, \ldots, x_T]$ is the raw BCG segment.

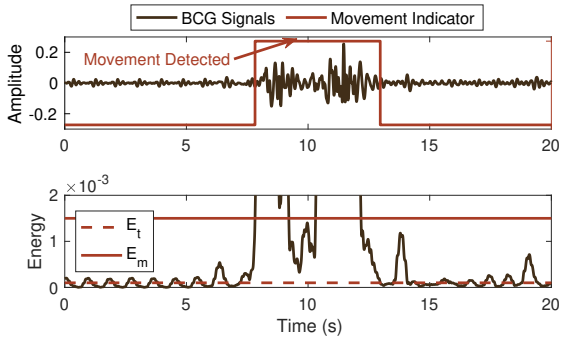

Figure 4: Movement detection.

3.2.2 Motion Artifact Removal. Spontaneous movements during sleep or rest introduce abrupt mechanical artifacts into BCG recordings. These bursts often overlap with cardiac frequencies, distorting temporal structures and degrading subsequent ECG reconstruction. To address this, we segment the filtered signal into short frames (0.1 s, 1-point shift) and compute per-frame energy:

$$E(f) = \frac{1}{N_f} \sum_{i=1}^{N_f} \tilde{x}_{\text{bf},i}^2. \tag{6}$$

Frames are discarded if (i) the maximum energy exceeds $E_m = 0.0015$, or (ii) the energy exceeds $E_t = 0.0001$ continuously for more than $T_t = 1.0$s—criteria derived from the temporal footprint of motion bursts. This dual-threshold strategy robustly isolates clean BCG segments (see Figure 4), forming a reliable basis for downstream ECG reconstruction. The remaining frames are concatenated to yield the final motion-cleaned BCG signal:

$$\tilde{\mathbf{x}}_{\text{mr}} = \text{MotionRemoval}(\tilde{\mathbf{x}}_{\text{bf}}). \tag{7}$$

3.3 Multi-Scale Feature Extractor

To capture BCG's hierarchical temporal complexity, we design a decomposable multi-scale architecture with integrated temporal modeling (Figure 2). It preserves both short-range oscillations (e.g., J-waves) and long-range modulations (e.g., HRV, respiration) , enabling scale-aware representation learning [37] while remaining lightweight for edge deployment. The encoder first processes the filtered BCG signal $\tilde{\mathbf{x}}_{\text{mr}} \in \mathbb{R}^{T\times 1}$ through a stack of 1D convolutions:

$$\mathbf{h}^{\text{conv}} = \text{Conv1D}(\tilde{\mathbf{x}}_{\text{mr}}), \tag{8}$$

followed by max pooling for temporal downsampling:

$$\mathbf{h}^{\text{pool}} = \text{MaxPool1D}(\mathbf{h}^{\text{conv}}). \tag{9}$$

To capture long-range dependencies, we use a bidirectional GRU (Bi-GRU), encoding forward and backward context:

$$\mathbf{h}_t^{\text{Bi-GRU}} = \text{GRU}_{\rightarrow}(\mathbf{h}^{\text{pool}}) || \text{GRU}_{\leftarrow}(\mathbf{h}^{\text{pool}}). \tag{10}$$

The decoder mirrors the encoder using transposed convolutions (Deconv1D) to recover the original resolution. Skip connections between symmetric encoder and decoder layers preserve timing cues critical for waveform morphology and phase alignment. The final output $\mathbf{e} \in \mathbb{R}^{T\times f_{\text{out}}}$ encodes multi-resolution cardiac features, including sharp transients (e.g., J-waves), composite structures (e.g., I–J–K complexes), and global rhythm patterns (e.g., HRV cycles).

3.4 Shared-Personalized Experts

To address inter-subject variability, we adopt Shared-Personalized Experts, a lightweight Mixture-of-Experts (MoE) module that disentangles generic cardiovascular patterns from individual-specific dynamics by adaptively routing feature sequences through shared and personalized expert branches (Figure 2).

The input $\mathbf{e} \in \mathbb{R}^{T\times f_{\text{out}}}$ is fed into K parallel experts $f_1, \ldots, f_K$, where the first M model generic cardiovascular patterns (e.g., HRV, QRS morphology), and the remaining $K - M$ specialize in user-specific dynamics such as timing shifts and waveform idiosyncrasies. A gating network $g(\cdot)$ serves as a learned router, dynamically weighting and selecting the top-k most informative branches for each input sequence. The final representation is a sparse, weighted combination:

$$\mathbf{c} = \sum_{j\in I_{\text{active}}} g(\mathbf{e}) f_j(\mathbf{e}), \tag{11}$$

where $I_{\text{active}} \subseteq \{1, \cdots, K\}$ denotes the selected expert indices and $g(\mathbf{e})$ provides normalized gating weights. This mechanism allows `Med2ECG` to generalize robustly to unseen users while still providing personalized adaptation when sufficient subject-specific data is available.

The resulting representation $\mathbf{c} \in \mathbb{R}^{T\times f}$ captures both population-level invariants and subject-specific features. To reconstruct the ECG waveform, we apply a lightweight decoder comprising 1D convolutional layers:

$$\hat{\mathbf{y}}_{\text{ECG}} = \text{Conv1D}(\mathbf{c}). \tag{12}$$

3.5 Medical-Informed Strategies

Accurate ECG reconstruction demands more than point-wise accuracy—it requires preserving the physiological morphology and diagnostic landmarks across multiple temporal scales. To this end, we design a composite loss that integrates three physiologically informed objectives: *Structure-Aligned Loss*, *Wave-Aware Loss*, and *Diagnostic Anchor Loss*.

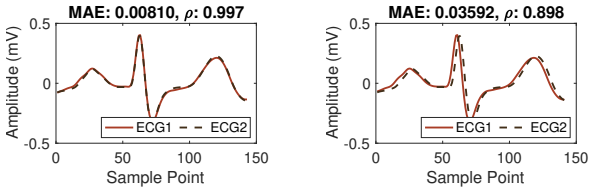

(a) Well-aligned ECG signals **(b) With a temporal shift of 8 ms**

Figure 5: ECG Loss with different temporal shifts.

3.5.1 Structure-Aligned Loss. Shape similarity is fundamental to preserving the clinical interpretability of reconstructed waveforms. However, conventional metrics—such as cosine similarity or Pearson correlation—are brittle to slight temporal misalignments. Even a subtle 8 ms shift, induced by device jitter or physiological variability, can cause the correlation to drop from 0.997 to 0.898 (Figure 5a vs. Figure 5b), despite the waveforms appearing visually identical. This fragility underscores the need for a delay-tolerant objective that prioritizes morphological fidelity over strict time alignment.

Thus, we introduce a delay-invariant formulation that selects the best alignment within a bounded temporal window:

$$\tau^* = \arg \min_{\tau \in [-\Delta, \Delta]} \mathcal{L}_{\text{sim}}(y_{\text{ECG}}, \hat{y}_{\text{ECG}}^{(\tau)}), \tag{13}$$

where $\hat{y}_{\text{ECG}}^{(\tau)}$ denotes the predicted results shifted by τ time steps, and $\mathcal{L}_{\text{sim}}$ represents the cosine similarity.

3.5.2 Wave-Aware Loss. While shape similarity captures morphological alignment, it lacks stable point-wise supervision. Standard Mean Absolute Error (MAE), in contrast, provides robust convergence by uniformly penalizing per-timestep deviations. However, this uniformity also leads to systematic underfitting of clinically critical yet low-amplitude regions, such as the P and T waves. Medical studies have shown that these subtle components—particularly P wave morphology and T wave symmetry—carry essential diagnostic information for atrial conduction disorders, electrolyte imbalance, and early ischemic changes [16].

Thus, we introduce *Wave-Aware Loss* that dynamically reweights each region's contribution based on its relative signal energy. Specifically, each labeled ECG waveform is segmented into four regions—P wave, QRS complex, T wave, and non-wave—and assigned a weight:

$$w_i = 2 - \frac{\sum_{j=1}^{N_i} \left| s_j^i \right|}{\sum_{j=1}^{N} \left| x_j \right|} \quad \text{for } i \in \{\text{P}, \text{QRS}, \text{T}\}, \tag{14}$$

where s_j^i is the j-th sample in subwave i, and x_j is the j-th sample of the full ECG. Each time point inherits the weight of its subwave membership, forming a temporal weight profile w_i The non-wave region is assigned a constant weight of 1. Finally, the wave-aware loss is computed as:

$$\mathcal{L}_{\text{wav}} = \frac{1}{N} \sum_{i=1}^{N} w(i) \cdot |\hat{y}_i^{\text{ECG}} - y_i^{\text{ECG}}|. \tag{15}$$

3.5.3 Diagnostic Anchor Loss. Although waveform-level alignment captures the overall morphology, it does not take into account the diagnostic anchors that delineate critical clinical intervals, such as PR and QT. These anchors include the onset of P waves, the onset and offset of QRS waves, and the offset of T waves. These intervals are susceptible to even minor misalignments, where millisecond-level deviations have the potential to obscure early indications of arrhythmia or ischemia[26].

We therefore introduce *Diagnostic Anchor Loss* to explicitly align these landmarks, grounding the reconstruction in clinically actionable structure. Specifically, a U-Net-based network is trained to generate a binary mask y that highlights the key landmarks of the ECG signal, where peak time steps are assigned a value of 1 and all others are 0.

To counter extreme foreground-background imbalance, we employ the Tversky loss [24]:

$$\mathcal{L}_{\text{dal}} = 1 - \frac{TP}{TP + FP + FN} \tag{16}$$

where TP, FP, FN quantify prediction quality on annotated cardiac landmarks vs. other timepoints.

3.5.4 Final Objective. The final objective combines all three losses into a physiologically grounded composite:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{sim}} + \lambda_1 \mathcal{L}_{\text{wav}} + \lambda_2 \mathcal{L}_{\text{dal}} \tag{17}$$

where λ_1, λ_2 balance waveform and landmark fidelity. This unified formulation ensures delay-robust, subwave-sensitive, and landmark-aware reconstruction—bridging morphological accuracy and clinical utility.

4 EVALUATION

4.1 Datasets

We evaluate the performance of Med2ECG on two datasets: a public benchmark from Kansas State University [5] and a self-collected clinical dataset.

- **Public Dataset.** The dataset contains synchronized multimodal recordings, including BCG, ECG, and PPG, collected via four electromechanical film sensors and four load cells placed under the mattress. We select the load cell channel with the highest signal variance as the BCG input due to its superior signal quality. ECG signals are recorded via a GE Datex Ohmeda CardioCap 5 monitor. All signals are digitized through a National Instruments 9220 analog input module at 1000 Hz, ensuring precise temporal alignment.
- **Clinical Dataset.** To evaluate real-world generalization, we collect a complementary dataset from 30 participants (14 female, aged 22–84, mean age 35.0 ± 18.7 years). The cohort includes 20 healthy individuals and 10 cardiac patients diagnosed with arrhythmia, angina, heart failure, myocardial infarction, or coronary artery disease. Participants wear ECG patches while a BCG sensor is placed beneath the pillow. The AD8232 module used for ground truth acquisition provides a single-lead ECG that closely approximates the standard Lead II configuration. They rest supine for 10 minutes with minimal movement to mitigate motion artifacts.

4.2 Evaluation Metrics

To holistically evaluate ECG reconstruction quality, we organize the metrics into two functional categories: (1) *morphological fidelity*, assessed via Amplitude Error and Pearson Correlation Coefficient (PCC); and (2) *temporal precision*, measured by RR, QRS, PR, and QT intervals.

- **Amplitude Error** quantifies the relative discrepancy in signal magnitude, capturing the fidelity of waveform morphology: Amplitude Error = $\frac{1}{N} \sum_{i=1}^{N} |\frac{\hat{y}_i - y_i}{y_i}| \times 100\%$.
- **Pearson Correlation Coefficient** measures global waveform similarity via linear alignment between predicted and ground-truth ECGs: PCC = $\frac{\sum_i^N (\hat{y}_i - \bar{\hat{y}})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^N (\hat{y}_i - \bar{\hat{y}})^2} \sqrt{\sum_{i=1}^N (y_i - \bar{y})^2}}$.
- **RR Interval Error** reflects heartbeat periodicity and directly related to heart rate variability (HRV): RR Error = $|\hat{t}_R^{(i+1)} - \hat{t}_R^{(i)} - (t_R^{(i+1)} - t_R^{(i)})|$.
- **QRS Interval Error** measures the accuracy in reconstructing ventricular depolarization duration, which is critical for detecting conduction abnormalities: QRS Error = $|\hat{t}_S - \hat{t}_Q - (t_S - t_Q)|$.
- **PR Interval Error** evaluates atrioventricular conduction timing fidelity: PR Error = $|\hat{t}_R - \hat{t}_P - (t_R - t_P)|$.

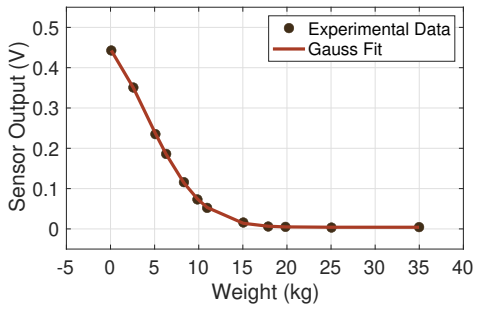

Figure 6: Relationship between applied weight and optical fiber sensor output.

- **QT Interval Error** reflects the model's ability to recover the full cycle of ventricular electrical activity: QT Error = $|\hat{t}_T - \hat{t}_Q - (t_T - t_Q)|$.

4.3 Implementation Details

All models are trained on an NVIDIA A40 GPU (40GB memory) using the Adam optimizer with a learning rate of 0.001 and a batch size of 24. Training runs for 120 epochs. The convolutional backbone adopts a kernel size of 7 in the multi-scale feature extractor and 3 in all remaining convolutional layers. The number of experts is set to 16, where one of them is a shared expert. The top two experts are selected. We set the loss balancing coefficients λ_1 and λ_2 to 1 and 0.1. Following standard practice, the model takes as input non-overlapping 10-second BCG segments sampled at 250 Hz, resulting in input sequences of 2500 timepoints. This window length covers multiple cardiac cycles and balances temporal context with computational efficiency.

To evaluate generalization across subjects, we adopt a leave-one-subject-out (LOSO) validation scheme: the model is trained on all but one subject and evaluated on the held-out individual. The auxiliary *Diagnostic Anchor Loss* is trained solely on ECG signals from the training subjects in each fold.

4.4 Sensor Feasibility for BCG Monitoring

To evaluate the feasibility of optical microfiber sensors, we designed a strain-based experiment. To simulate different weight conditions, Adelberg HPA hand-force digital force tester was employed, with the applied force being distributed evenly over an area of 0.05 square metres on the sensor pad. The system was designed to detect optical intensity, and the corresponding voltage signals were recorded as the applied weight varied, as illustrated in Figure 6. The findings suggest that light attenuation within the sensor follows an exponential decay, which is consistent with the description provided in Eq. (4). The sensor undergoes deformation under different weights, which in turn affects the transmitted optical power. In accordance with this principle, the sensor can be employed to monitor strain variations induced by heartbeat-induced micro-vibrations in the sensing region.

4.5 Overall Performance

We benchmark `Med2ECG` against three state-of-the-art baselines— Multi-level ECG-Recovery Model (MEM) [35], Rec-AUNet (RAU) [36], and BiLSTM [44]— on both the **Clinical Dataset** and **Public Dataset** (Figure 7). To clarify evaluation focus, we report metrics in two domains: (i) *morphological fidelity*, assessed by Amplitude Error and PCC; and (ii) *temporal precision*, evaluated via RR, QRS, PR, and QT intervals.

- Across both datasets, `Med2ECG` consistently outperforms baselines in morphology-related metrics. In the self-collected dataset, it reduces the amplitude error from 21.27% (RAU) to 12.82% and increases PCC from 0.818 to 0.927. On the public dataset, amplitude error drops from 21.08% to 14.72%, with PCC improving from 0.805 to 0.916. These results underscore `Med2ECG`'s superior capability in reconstructing waveform morphology under varied conditions.
- In terms of temporal metrics, `Med2ECG` achieves the lowest MAE across all cardiac intervals. For instance, RR/QRS/PR/QT errors are 7.31/15.90/12.10/13.75 ms in the self-collected dataset, and 4.77/19.27/17.07/15.55 ms on the public dataset. While RAU shows relatively stronger performance on amplitude error and RR interval than MEM, it still underperforms `Med2ECG` across the board.

A sample result of reconstructed ECG from BCG is shown in Figure 8. Taken together, the results demonstrate `Med2ECG`*'s robustness in both structural waveform recovery and clinically relevant timing estimation, validating its generalizability in leave-one-subject-out evaluation settings.*

4.6 Ablation Study

To quantify the contribution of each design choice in `Med2ECG`, we conduct an ablation study by systematically removing key modules from our full model, which we denote as `All`. The study includes five variants (Figure 9):

- `w/o SAL` removes the similarity constraints (*Structure-Aligned Loss*), which encourages waveform consistency;
- `w/o WAL` disables subwave-aware error reweighting (*Wave-Aware Loss*), treating all timepoints equally;
- `w/o DAL` removes the temporal landmark alignment module (*Diagnostic Anchor Loss*), forgoing the explicit alignment of clinically critical landmarks;
- `w/o SPE` eliminates generalization module (*Shared-Personalized Experts*), preventing the model from adapting to subject-specific ECG patterns;
- `GM` denotes a general model trained and evaluated on all subjects, serving as an upper-bound reference. However, this setting is unrealistic in practical deployments, where `GM` is expected to generalize to unseen users without requiring subject-specific BCG and ECG recordings for training.

The results reveal three overarching insights:

- `w/o SAL` leads to the largest drop in PCC (a drop of 8.9%) and the highest RR interval error (increases of 3.25 ms), confirming its importance for global waveform alignment and rhythm preservation.
- `w/o SPE` significantly degrades QRS interval accuracy (increases of 3.37 ms) and overall waveform correlation (a drop of 3.3%), underscoring its critical role in preserving ventricular depolarization morphology.

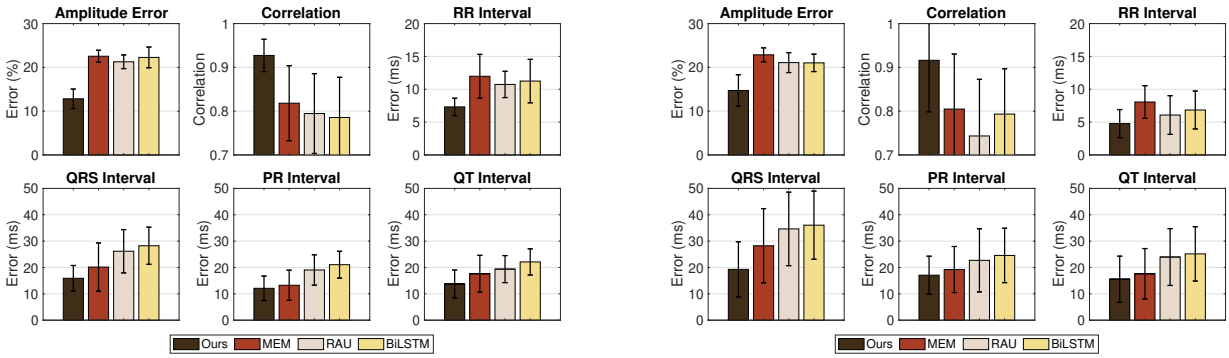

(a) Clinical dataset. (b) Public dataset.

Figure 7: The overall performance of Med2ECG.

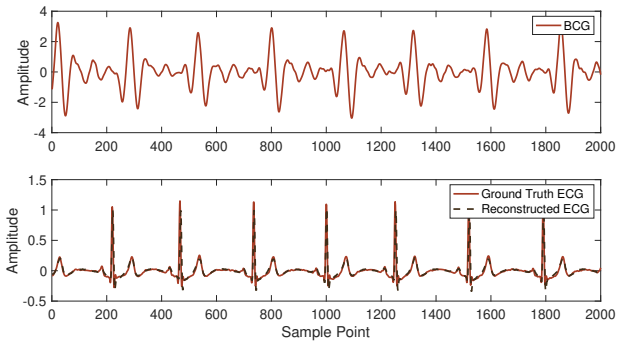

Figure 8: Sample result: BCG to ECG.

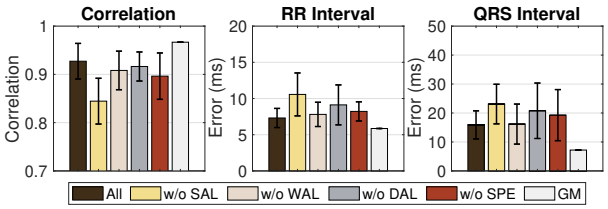

Figure 9: The performance of ablation study.

- `w/o DAL` and `w/o WAL` cause noticeable performance declines in interval-based metrics and correlation, validating the necessity of explicitly aligning the landmarks and energy-aware gradient emphasis.

4.7 Sensitivity Analysis

4.7.1 The Number of Experts. We investigate the impact of the total number of experts, K, while fixing the number of Shared Experts at $M = 1$ and the number of selected Top-K experts at $k = 2$.

As shown in Figure 10, increasing K from 11 to 21 resulted in a consistent enhancement across all performance metrics—the PCC increased, while the RR and QRS Interval Errors decreased significantly. This trend confirms that a larger pool of experts provides the necessary capacity to model the diverse individual- and environment-specific variabilities inherent in the data. However, performance gains plateau beyond $K = 16$. Consequently, we select $K = 16$ as the final configuration, which offers the best trade-off between model complexity and reconstruction fidelity. We also examined the effect of the number of selected Top-k experts ($k = 1$–3), where $k = 2$ yielded the best overall performance.

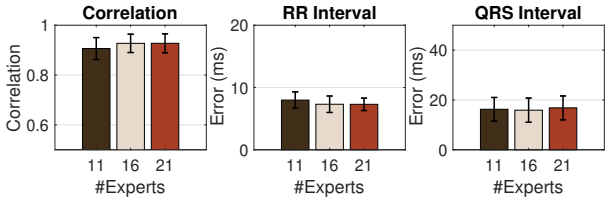

Figure 10: The performance at different number of experts.

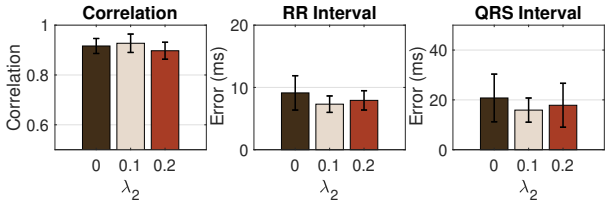

Figure 11: The performance at different loss weight of DAL.

4.7.2 The Loss Weight of DAL. We analyze the sensitivity of the model's performance to the weight λ_2 of $\mathcal{L}_{DAL}$. Figure 11 demonstrates that setting λ_2 to 0 (disabling the loss) yielded the poorest results, underscoring the critical importance of $\mathcal{L}_{DAL}$ for clinical accuracy. As λ_2 increased, performance improved rapidly, achieving peak metrics—highest PCC and lowest interval errors—at $\lambda_2 = 0.1$. Further increasing λ_2 led to a slight performance degradation, suggesting that overemphasizing the sparse landmark anchors may compromise the overall morphological integrity guided by the similarity losses. Thus, $\lambda_2 = 0.1$ is chosen as the optimal weight to ensure a robust balance between waveform similarity and temporal precision of diagnostic intervals.

4.8 Physiological Robustness

To assess the robustness of Med2ECG, we executed a comprehensive suite of experiments under different settings.

4.8.1 Heart Rate Sensitivity. We first evaluate the performance of Med2ECG under different heart rate (HR) ranges. Participants perform light physical activity for one minute to elevate heart rates, followed by data collection. Each heartbeat is categorized into one of four HR groups based on local RR intervals: ≤ 60 bpm, (60, 80] bpm, (80, 100) bpm, and ≥ 100 bpm, where instantaneous HR is estimated using the average RR intervals of the current and adjacent beats.

As illustrated in Figure 12, the system maintains strong performance across all HR bins, consistently achieving PCC > 0.9. Interestingly, performance improves from bradycardic to moderate tachycardic ranges (60–100 bpm), likely due to reduced inter-beat variability and shorter RR intervals enhancing temporal coherence. However, a slight decline is observed beyond 100 bpm, possibly due to elevated motion artifacts and signal instability in post-exercise states, which reduce signal-to-noise ratio.

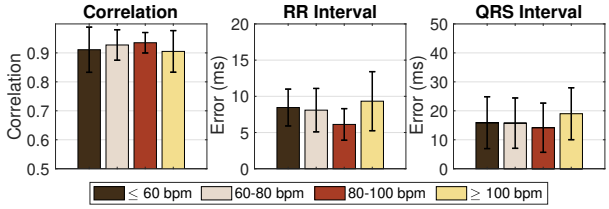

Figure 12: The performance at different heart rates.

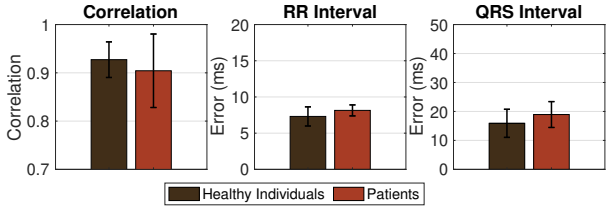

Figure 13: Healthy vs. Clinical Individuals.

4.8.2 Evaluation on Cardiovascular Patients. Beyond healthy individuals, we further evaluate Med2ECG's applicability to 10 patients (**Clinical dataset**) with diverse cardiovascular conditions (e.g., arrhythmia, angina, heart failure, myocardial infarction, or coronary heart disease). As shown in Figure 13, Med2ECG maintains strong reconstruction performance, achieving a mean PCC of 0.904, RR interval error of 8.15 ms, and QRS error of 18.93 ms—demonstrating robust generalization to pathological cases. To illustrate condition fidelity, Figure 14 highlights three representative cases: (a) tachycardia (HR > 100 bpm), (b) bradycardia (HR < 60 bpm), and (c) myocardial infarction, marked by elevated T-waves.

4.8.3 Sleeping Posture Sensitivity. To assess robustness against real-world variability, we evaluate system performance under four common sleeping postures. Each participant lies in 3 non-supine positions—left lateral, right lateral, and prone—for three minutes each, following a baseline supine posture.

As shown in Figure 15, performance degrades modestly across non-supine postures due to signal attenuation and altered pressure distribution. Nevertheless, the system maintains strong reconstruction fidelity, achieving PCC scores of 0.84 (prone), 0.87 (left), and 0.87 (right), compared to 0.92 in the supine condition. These results demonstrate that Med2ECG generalizes well to challenging sleep scenarios, offering consistent ECG estimation even in postures that typically induce motion artifacts and sensor decoupling.

4.8.4 Longer-Duration Validation. We conducted a 30-minute study with 10 participants to assess the robustness of Med2ECG under extended monitoring conditions. This duration was sufficient to capture multiple spontaneous posture transitions—such as supine, lateral, and prone turns—which represent the dominant sources of signal disruption during overnight sleep.

As illustrated in Figure 16, Med2ECG maintained stable and high-quality ECG reconstruction throughout the entire session. The mean PCC remained consistently high across all three segments (0.910, 0.882, and 0.895, respectively), while the RR interval error exhibited minimal variation, confirming the system's ability to continuously and accurately track heart rate variability. Minor fluctuations between segments were primarily associated with transient motion artifacts caused by posture shifts. Quantitatively, the differences in average metrics across segments were negligible, reinforcing that Med2ECG is resilient to cumulative disturbances and capable of sustaining reliable performance during realistic, long-duration use.

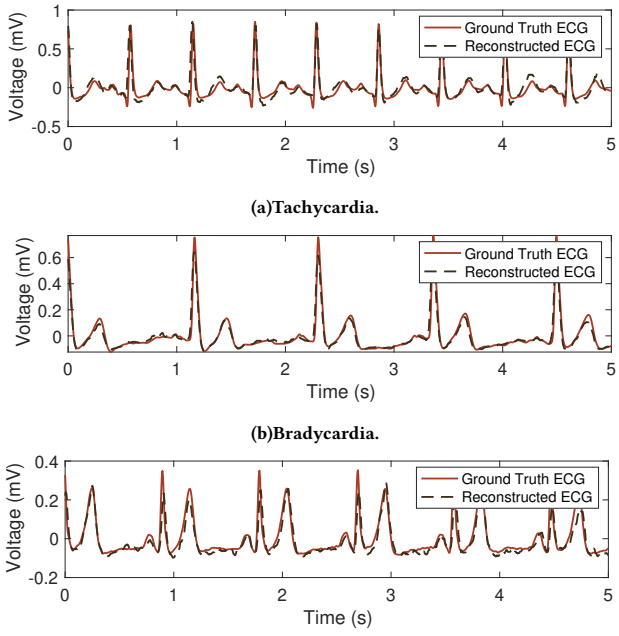

(a)Tachycardia.

(b)Bradycardia.

(c)Myocardial infarction.

Figure 14: The sample results on patients' data.

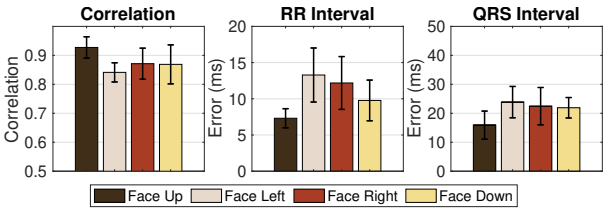

Figure 15: Results for different sleeping postures.

4.8.5 Robustness to Environmental Vibration. To assess the robustness of Med2ECG in real-world environments, we simulate common ambient mechanical disturbances, e.g., foot traffic near the bed, by having an individual continuously walk within a one-meter radius during data collection. Figure 17 shows reconstruction accuracy remains stable across all metrics, indicating Med2ECG's resilience to pedestrian-induced vibrations typical of home or clinical settings.

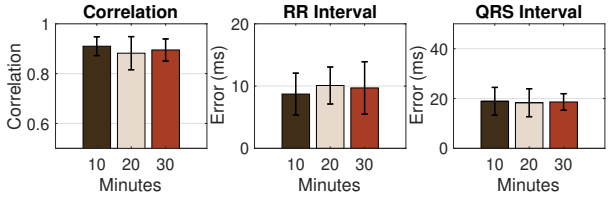

Figure 16: The performance changes over the initial 30 minutes.

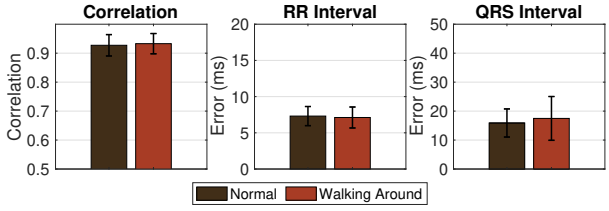

Figure 17: The performance under external vibration.

4.9 Real-Time Inference Efficiency

We deploy Med2ECG on Raspberry Pi 4 Model B [2] and evaluate its latency running on the CPU cores only. For each 10 s BCG segment (2500 samples at 250 Hz), inference takes 33.25 ms for filtering, 23.36 ms for artifact removal, and 1342.10 ms for reconstruction—totaling 1398.71 ms per window. This yields a real-time factor (RTF) of 0.14 ×, confirming its suitability for edge deployment.

4.10 Clinical Context of Reconstruction Errors

To assess the clinical significance of reconstruction errors, we benchmarked the observed deviations against cardiology guidelines from the American Heart Association and the European Society of Cardiology, which define acceptable tolerances of ±20 ms for QRS and PR intervals and ±30 ms for QT intervals [23, 30]. When benchmarked against these standards, the reconstructed signals from Med2ECG remain well within clinically safe margins. The mean absolute errors for PR, QRS, and QT intervals are all below their respective tolerance thresholds, indicating that the model faithfully preserves key diagnostic temporal landmarks. Importantly, morphological variations—such as QRS widening which are critical indicators of arrhythmia or ischemic injury—are reconstructed with sufficiently low deviations to prevent misinterpretation in both automated and clinician-based analyses [11]. Although minor waveform-level discrepancies persist, they do not compromise the interpretability or diagnostic reliability of the reconstructed ECGs. Therefore, Med2ECG demonstrates not only high numerical accuracy but also clinical robustness, ensuring that diagnostic conclusions derived from reconstructed signals remain trustworthy under real-world deployment conditions.

5 RELATED WORKS

5.1 Contact-based ECG Reconstruction

Traditional ECG acquisition relies on skin-contact electrode patches, which are obtrusive and impractical for continuous, long-term monitoring [7]. To address this limitation, recent research has explored reconstructing ECG from more accessible modalities such as photoplethysmography (PPG) [31, 33, 46] or cardiac-induced mechanical vibrations [4, 22]. For instance, CardioGAN [25] leverages a generative adversarial framework to translate PPG into ECG waveforms, enabling more precise heart rate estimation than raw PPG.

5.2 Contactless ECG Reconstruction

Recent advances in wireless sensing, including millimeter-wave (mmWave) and ultra-wideband (UWB) technologies, have enabled non-contact estimation of ECG signals. For example, Toda et al. [32] use convolutional neural networks to map Frequency Modulated Continuous Wave (FMCW) radar data to ECG, while mmECG [41] captures sub-millimeter chest wall vibrations in vehicles for ECG reconstruction. CardiacWave [40] leverages radar polarization effects for improved signal capture, and RF-ECG [38] applies conditional GAN on UWB-derived heartbeat signals. Recently, AirECG [45] employs a diffusion model guided by calibration ECGs to boost reconstruction fidelity. While effective, these radar-based solutions suffer from high cost, regulatory overhead, and vulnerability to environmental interference.

In contrast, our system—Med2ECG—leverages cost-effective, flexible optical fiber sensors to acquire high-fidelity BCG signals, enabling robust, contactless ECG reconstruction. Compared to vibration-based approaches like HeartQuake [22], which rely on geophones to capture low-resolution bed vibrations, our sensor delivers finer-grained signal quality and requires no structural modification, making it practical for scalable deployment. While prior BCG-to-ECG methods [29, 36] adopt pointwise regression to map BCG to ECG waveforms, they fail to generalize across subjects due to their purely data-driven design. Med2ECG *addresses this limitation by introducing a structure-aligned formulation grounded in shared physiological events—such as the temporal correspondence between J-waves in BCG and R-peaks in ECG—allowing for semantically meaningful and more generalizable reconstructions.*

6 DISCUSSION

6.1 Insights and Lessons Learned

Our comprehensive analysis yields two key insights:

- **Physiology-aware structure enables robust and generalizable ECG reconstruction.** Combining a Mixture-of-Experts architecture with physiologically informed objectives (structure-, wave-, and landmark-aware losses) allows Med2ECG to disentangle individual-specific variations from core cardiac dynamics. This design yields high-fidelity reconstructions (supported by Figure 8) and consistently generalizes across *unseen users* (supported by Figure 7), *postural variations* (supported by Figure 15), and *clinically diverse patient populations* (supported by Figure 13)—a nontrivial achievement in BCG-based sensing.

- **BCG-to-ECG transformation enables robust cardiovascular monitoring.** By successfully recovering the R-peak and other key cardiac landmarks, our framework provides a standardized signal format that is significantly more amenable to standard clinical toolsets than raw BCG. This facilitates more reliable estimation of cardiovascular metrics, such as heart rate variability and timing-based indices, effectively bridging the gap between consumer-grade sensing and clinical-grade insights.

6.2 Limitations and Future Work

Med2ECG performs well, yet leaves room for improvement:

- Although our optical fiber sensor supports flexible deployment, signal quality may degrade under high-load conditions. To enhance the generalizability of the sensor, further optimization of the fiber's geometry—such as core thickness and wiring layout—is needed to better match different pressure characteristics. Furthermore, environmental factors such as transmission medium thickness and postural changes can impact BCG signal quality and subsequent reconstruction. While user displacement may cause temporary signal interruptions, mattress-integrated sensing presents a feasible strategy to mitigate this issue[18]. Although consistent across subjects, systematically investigating the effects of medium thickness and posture remains a priority for future work.
- We collect ground-truth ECG using the AD8232 sensor and evaluate Med2ECG in resting states. This work represents an initial step towards single-lead ECG reconstruction, and it is imperative that future endeavours focus on extending validation to encompass long-term overnight monitoring. Furthermore, a comparison with clinical-grade 12-lead ECGs will be essential to substantiate high-stakes healthcare use cases. Additionally, more advanced motion artifact mitigation strategies are needed to ensure reliability in fully unconstrained environments.

6.3 Practicality and Cost Considerations

A key strength of Med2ECG is its cost-effectiveness and accessibility. The cost for our optical microfiber sensors is estimated to be under USD 20, including the light source, photodiode and signal acquisition board. By contrast, conventional Holter ECG systems typically cost hundreds to thousands of dollars. Meanwhile, consumer ECG patches remain priced at USD 100–300 per unit and are only reusable a limited number of times [8]. Med2ECG dramatically reduces the economic barrier to large-scale or long-term monitoring. This affordability, combined with the system's fully non-contact design, highlights its potential for clinical deployment in settings with limited resources and for continuous at-home cardiovascular monitoring.

6.4 Implication and Generalization

Med2ECG's design holds significant implications for ubiquitous health monitoring. First, Med2ECG enables passive cardiac monitoring in everyday environments (e.g., homes, hospitals, and elderly care facilities) without requiring electrode placement or user compliance, by reconstructing diagnostic-grade ECG from unobtrusive, contactless BCG signals collected via low-cost optical fiber sensors. Second, Med2ECG's structure-aligned, physiology-grounded design generalizes well across users, postures, and conditions, including patients with cardiac abnormalities and varying heart rates, demonstrating effectiveness in real-world scenarios.

Beyond ECG reconstruction, our learning paradigm, anchored on temporal hierarchy encoding, context-aware routing, and landmark-aware supervision, offers a reusable template for translating non-invasive, coarse-grained biosignals into clinically meaningful representations. This opens pathways to generalize Med2ECG toward other physiological signals (e.g., PPG, respiration, blood pressure) and fosters future cross-modal modeling efforts in health sensing, ambient intelligence, and context-aware interaction systems.

6.5 Ethics and Privacy

All experiments involving human participants complied with institutional review board (IRB) protocols and received formal ethical approval. Prior to data collection, participants for the main experiments and downstream applications were informed of the study objectives, procedures, and potential risks, and provided written informed consent. All personally identifiable information (PII) was anonymized, and physiological signals, particularly ECG data, were securely stored on encrypted servers with restricted access. All data were used solely for research and will not be publicly released without further anonymization and consent.

7 CONCLUSION

We present Med2ECG, a generalizable and contactless ECG reconstruction system that uses high-fidelity BCG signals to enable early cardiovascular risk detection in unobtrusive monitoring during periods of rest. Departing from conventional waveform regression, Med2ECG introduces a structure-aligned formulation that models the shared physiological sequence of cardiac events (e.g., P-wave, QRS complex, T-wave vs. I-, J-waves), anchoring learning on semantically meaningful cardiac structure. It consists of three dedicated components: i) Multi-Scale Feature Extractor captures hierarchical cardiac dynamics across fine-grained transients, full-cycle morphology, and long-range rhythms, ii) Shared-Personalized Experts disentangles physiological patterns from subject- and posture-induced variability, and iii) Medical-Informed Strategies preserves diagnostic fidelity via waveform-aware weighting, alignment-consistent similarity, and landmark-anchored supervision. Experiments across public and in-hospital datasets (20 healthy people and 10 patients with diverse CVDs) demonstrate strong generalization across subjects and environments, achieving >0.92 correlation, <15% amplitude error, and accurate PR/QRS/QT/RR interval estimation within 5–20 ms. *By bridging physiological structure with robust machine learning,* Med2ECG *lays the foundation for scalable and clinically meaningful cardiac screening.*

ACKNOWLEDGMENTS

This research is partially supported by China NSFC Grant (62472366), the Project of DEGP (No.2023KCXTD042, 2024GCZX003), Guangdong Provincial Key Lab of Integrated Communication, Sensing and Computation for Ubiquitous Internet of Things (No.2023B1212010007), "111 Center (No.D25008)", Shenzhen Science and Technology Foundation (ZDSYS20190902092853047).

REFERENCES

[1] 2014. AD8232. https://learn.sparkfun.com/tutorials/ad8232-heart-rate-monitor-hookup-guide/all. Last accessed October 5, 2023.

[2] 2025. Raspberry Pi computer hardware - Raspberry Pi Documentation. https://www.raspberrypi.com/documentation/computers/raspberry-pi.html#raspberry-pi-4-model-b. Last accessed October 1, 2025.

[3] Michele Brignole, Panos Vardas, Ellen Hoffman, Heikki Huikuri, Angel Moya, Renato Ricci, Neil Sulke, Wouter Wieling, Angelo Auricchio, Gregory YH Lip, et al. 2009. Indications for the use of diagnostic implantable and external ECG loop recorders. *Europace* 11, 5 (2009), 671–687. https://doi.org/10.1093/europace/eup097

[4] Yetong Cao, Fan Li, Huijie Chen, Xiaochen liu, Li Zhang, and Yu Wang. 2022. Guard Your Heart Silently: Continuous Electrocardiogram Waveform Monitoring with Wrist-Worn Motion Sensor. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 3 (2022), 1–29. https://doi.org/10.1145/3550307

[5] Charles Carlson, Vanessa-Rose Turpin, Ahmad Suliman, Carl Ade, Steve Warren, and David E. Thompson. 2020. Bed-based ballistocardiography: Dataset and ability to track cardiovascular parameters. *Sensors* 21, 1 (2020), 156. https://doi.org/10.3390/s21010156

[6] John P DiMarco and John T Philbrick. 1990. Use of ambulatory electrocardiographic (Holter) monitoring. *Annals of internal medicine* 113, 1 (1990), 53–68. https://doi.org/10.7326/0003-4819-113-1-53

[7] Heal Force. 2023. PC-08B. http://www.healforce.com/en/html/products/portableecgmonitors/healthcare-equipment-portable-ecg-monitors-pc-80b.html. Last accessed September 5, 2021.

[8] Erik Fung, Marjo-Riitta Järvelin, Rahul N Doshi, Jerold S Shinbane, Steven K Carlson, Luanda P Grazette, Philip M Chang, Rajbir S Sangha, Heikki V Huikuri, and Nicholas S Peters. 2015. Electrocardiographic patch devices and contemporary wireless cardiac monitoring. *Frontiers in physiology* 6 (2015), 149.

[9] Muhammad Wildan Gifari, Hasballah Zakaria, and Richard Mengko. 2015. Design of ECG Homecare: 12-lead ECG acquisition using single channel ECG device developed on AD8232 analog front end. In *2015 International Conference on Electrical Engineering and Informatics (ICEEI)*. IEEE, 371–376. https://doi.org/10.1109/ICEEI.2015.7352529

[10] Laurent Giovangrandi, Omer T. Inan, Dipanjan Banerjee, and Gregory T.A. Kovacs. 2012. Preliminary results from BCG and ECG measurements in the heart failure clinic. In *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, 3780–3783. https://doi.org/10.1109/EMBC.2012.6346790

[11] Hilma Holm, Daniel F Gudbjartsson, David O Arnar, Gudmar Thorleifsson, Gudmundur Thorgeirsson, Hrafnhildur Stefansdottir, Sigurjon A Gudjonsson, Aslaug Jonasdottir, Ellisiv B Mathiesen, Inger Njølstad, et al. 2010. Several common variants modulate heart rate, PR interval and QRS duration. *Nature genetics* 42, 2 (2010), 117–122.

[12] Yandao Huang, Lin Chen, Chenggao Li, Junyao Peng, Qingyong Hu, Yu Sun, Hao Ren, Weimin Lyu, Wen Jin, Junzhang Tian, et al. 2024. AI-driven system for non-contact continuous nocturnal blood pressure monitoring using fiber optic ballistocardiography. *Communications Engineering* 3, 1 (2024), 183. https://doi.org/10.1038/s44172-024-00326-w

[13] Yandao Huang, Minghui Qiu, Lin Chen, Zhencan Peng, Qian Zhang, and Kaishun Wu. 2023. NF-Heart: A Near-field Non-contact Continuous User Authentication System via Ballistocardiogram. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 7, 1 (2023), 1–24. https://doi.org/10.1145/3580851

[14] Vita N Jaspan, Garred S Greenberg, Siddhant Parihar, Christine M Park, Virend K Somers, Michael D Shapiro, Carl J Lavie, Salim S Virani, and Leandro Slipczuk. 2024. The role of sleep in cardiovascular disease. *Current atherosclerosis reports* 26, 7 (2024), 249–262. https://doi.org/10.1007/s11883-024-01207-5

[15] Hui Ji and Pengfei Zhou. 2024. Advancing PPG-Based Continuous Blood Pressure Monitoring from a Generative Perspective. In *Proceedings of the 22nd ACM Conference on Embedded Networked Sensor Systems*. 661–674. https://doi.org/10.1145/3666025.3699365

[16] Mark E Josephson. 2008. *Clinical cardiac electrophysiology: techniques and interpretations*. Lippincott Williams & Wilkins.

[17] Chang-Sei Kim, Stephanie L. Ober, M. Sean McMurtry, Barry A. Finegan, Omer T. Inan, Ramakrishna Mukkamala, and Jin-Oh Hahn. 2016. Ballistocardiogram: Mechanism and potential for unobtrusive cardiovascular health monitoring. *Scientific reports* 6, 1 (2016), 31297. https://doi.org/10.1038/srep31297

[18] Linqing Li, Chengwei Yang, Zhuo Wang, Kun Xiao, and Rui Min. 2024. Stretchable polymer optical fiber embedded in the mattress for respiratory and heart rate monitoring. *Optics & Laser Technology* 171 (2024), 110356.

[19] Onno Linschmann, Steffen Leonhardt, Antti Vehkaoja, and Christoph Hoog Antink. 2022. Estimation of the respiratory rate from ballistocardiograms using the Hilbert transform. *BioMedical Engineering OnLine* 21, 1 (2022), 54. https://doi.org/10.1186/s12938-022-01024-4

[20] David C. Mack, James T. Patrie, Paul M. Suratt, Robin A. Felder, and Majd Alwan. 2008. Development and preliminary validation of heart rate and breathing rate detection using a passive, ballistocardiography-based sleep monitoring system. *IEEE Transactions on information technology in biomedicine* 13, 1 (2008), 111–120. https://doi.org/10.1109/TITB.2008.2007194

[21] Dietrich Marcuse. 1976. Curvature loss formula for optical fibers. *Journal of the Optical Society of America* 66, 3 (1976), 216–220. https://doi.org/10.1364/JOSA.66.000216

[22] Jaeyeon Park, Hyeon Cho, Rajesh Krishna Balan, and JeongGil Ko. 2020. Heartquake: Accurate low-cost non-invasive ecg monitoring using bed-mounted geophones. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 3 (2020), 1–28. https://doi.org/10.1145/3411843

[23] Pentti M Rautaharju, Borys Surawicz, and Leonard S Gettes. 2009. AHA/ACCF/HRS recommendations for the standardization and interpretation of the electrocardiogram: part IV: the ST segment, T and U waves, and the QT interval a scientific statement from the American Heart Association Electrocardiography and Arrhythmias Committee, Council on Clinical Cardiology; the American College of Cardiology Foundation; and the Heart Rhythm Society Endorsed by the International Society for Computerized Electrocardiology. *Journal of the American College of Cardiology* 53, 11 (2009), 982–991.

[24] Seyed Sadegh Mohseni Salehi, Deniz Erdogmus, and Ali Gholipour. 2017. Tversky loss function for image segmentation using 3D fully convolutional deep networks. In *International workshop on machine learning in medical imaging*. Springer, 379–387. https://doi.org/10.48550/arXiv.1706.05721

[25] Pritam Sarkar and Ali Etemad. 2021. Cardiogan: Attentive generative adversarial network with dual discriminators for synthesis of ecg from ppg. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 488–496. https://doi.org/10.48550/arXiv.2010.00104

[26] Peter J Schwartz, Michael J Ackerman, Charles Antzelevitch, Connie R Bezzina, Martin Borggrefe, Bettina F Cuneo, and Arthur AM Wilde. 2020. Inherited cardiac arrhythmias. *Nature reviews Disease primers* 6, 1 (2020), 58.

[27] Jonathan Eprilio Soaduon Simanjuntak, Masayu Leylia Khodra, and Martin Clinton Tosima Manullang. 2020. Design methods of detecting atrial fibrillation using the recurrent neural network algorithm on the Arduino AD8232 ECG module. In *Iop conference series: Earth and environmental science*, Vol. 537. IOP Publishing, 012022. https://doi.org/10.1088/1755-1315/537/1/012022

[28] Konstantinos C. Siontis, Peter A. Noseworthy, Zachi I. Attia, and Paul A. Friedman. 2021. Artificial intelligence-enhanced electrocardiography in cardiovascular disease management. *Nature Reviews Cardiology* 18, 7 (2021), 465–478. https://doi.org/10.1038/s41569-020-00503-2

[29] Aravind Srinivasan, Zhang Haihong, Lin Zhiping, Jit Biswas, and Chen Zhihao. 2015. Towards numerical temporal-frequency system modelling of associations between electrocardiogram and ballistocardiogram. In *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. 394–397. https://doi.org/10.1109/EMBC.2015.7318382

[30] Charlotte Stephansen, Mads Brix Kronborg, Christoffer Tobias Witt, Jens Kristensen, Christian Gerdes, Anders Sommer, Jesper Møller Jensen, and Jens Cosedis Nielsen. 2019. Reproducibility of measuring QRS duration and implications for optimization of interventricular pacing delay in cardiac resynchronization therapy. *Annals of Noninvasive Electrocardiology* 24, 3 (2019), e12621.

[31] Qunfeng Tang, Zhencheng Chen, Rabab Ward, Carlo Menon, and Mohamed Elgendi. 2023. PPG2ECGps: An End-to-End Subject-Specific Deep Neural Network Model for Electrocardiogram Reconstruction from Photoplethysmography Signals without Pulse Arrival Time Adjustments. *Bioengineering* 10, 6 (2023), 630. https://doi.org/10.3390/bioengineering10060630

[32] Daiki Toda, Ren Anzai, Koichi Ichige, Ryo Saito, and Daichi Ueki. 2021. ECG signal reconstruction using FMCW radar and convolutional neural network. In *2021 20th International Symposium on Communications and Information Technologies (ISCIT)*. IEEE, 176–181. https://doi.org/10.1109/ISCIT52804.2021.9590627

[33] Khuong Vo, Emad Kasaeyan Naeini, Amir Naderi, Daniel Jilani, Amir M. Rahmani, Nikil Dutt, and Hung Cao. 2021. P2E-WGAN: ECG waveform synthesis from PPG with conditional wasserstein generative adversarial networks. In *Proceedings of the 36th Annual ACM Symposium on Applied Computing*. 1030–1036. https://doi.org/10.1145/3412841.3441979

[34] Joseph A Walsh III, Eric J Topol, and Steven R Steinhubl. 2014. Novel wireless devices for cardiac monitoring. *Circulation* 130, 7 (2014), 573–581. https://doi.org/10.1161/CIRCULATIONAHA.114.009024

[35] Lei Wang, Xingwei Wang, Dalin Zhang, Xiaolei Ma, Yong Zhang, Haipeng Dai, Chenren Xu, Zhijun Li, and Tao Gu. 2023. Knowing Your Heart Condition Anytime: User-Independent ECG Measurement Using Commercial Mobile Phones. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 7, 3 (2023), 1–28. https://doi.org/10.1145/3610871

[36] Peng Wang, Chuanqi Han, Fang Yu, Ye Zheng, Xi Huang, and Li Cui. 2022. Rec-AUNet: Attentive UNet for Reconstruction of ECG from BCG. In *2022 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. IEEE, 1241–1248. https://doi.org/10.1109/BIBM55620.2022.9994933

[37] Shiyu Wang, Haixu Wu, Xiaoming Shi, Tengge Hu, Huakun Luo, Lintao Ma, James Y. Zhang, and JUN ZHOU. 2024. TimeMixer: Decomposable Multiscale Mixing for Time Series Forecasting. In *The Twelfth International Conference on Learning Representations*. https://openreview.net/forum?id=7oLshfEIC2

[38] Zhi Wang, Beihong Jin, Siheng Li, Fusang Zhang, and Wenbo Zhang. 2023. ECG-grained Cardiac Monitoring Using UWB Signals. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 4 (2023), 1–25. https://doi.org/10.1145/3569503

[39] WHO. 2025. Cardiovascular diseases. https://www.who.int/health-topics/cardiovascular-diseases. Last accessed March 12, 2024.

[40] Chenhan Xu, Huining Li, Zhengxiong Li, Hanbin Zhang, Aditya Singh Rathore, Xingyu Chen, Kun Wang, Ming-chun Huang, and Wenyao Xu. 2021. Cardiacwave: A mmwave-based scheme of non-contact and high-definition heart activity computing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 3 (2021), 1–26. https://doi.org/10.1145/3478127

[41] Xiangyu Xu, Jiadi Yu, Chengguang Ma, Yanzhi Ren, Hongbo Liu, Yanmin Zhu, Yi-Chao Chen, and Feilong Tang. 2022. mmecg: Monitoring human cardiac cycle in driving environments leveraging millimeter wave. In *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*. IEEE, 90–99. https://doi.org/10.1109/INFOCOM48880.2022.9796912

[42] Haoyue Zhang, Zhuo Wang, Chuanxin Teng, Santosh Kumar, Xiaoli Li, and Rui Min. 2023. Wearable cardiorespiratory sensor for real-time monitoring with smartphone integration. *IEEE Transactions on Instrumentation and Measurement* 73 (2023), 1–10. https://doi.org/10.1109/TIM.2023.3338709

[43] Jia Zhang, Rui Xi, Yuan He, Yimiao Sun, Xiuzhen Guo, Weiguo Wang, Xin Na, Yunhao Liu, Zhenguo Shi, and Tao Gu. 2023. A Survey of mmWave-Based Human Sensing: Technology, Platforms and Applications. *IEEE Communications Surveys & Tutorials* 25, 4 (2023), 2052–2087. https://doi.org/10.1109/COMST.2023.3298300

[44] Lingyin Zhang. 2024. Reconstructing ECG from BCG for Non-Contact Cardiac Monitoring Using Bidirectional LSTM. In *Proceedings of the 2024 7th International Conference on Signal Processing and Machine Learning*. 242–251. https://doi.org/10.1145/3686490.368652

[45] Langcheng Zhao, Rui Lyu, Hang Lei, Qi Lin, Anfu Zhou, Huadong Ma, Jingjia Wang, Xiangbin Meng, Chunli Shao, Yida Tang, et al. 2024. AirECG: Contactless Electrocardiogram for Cardiac Disease Monitoring via mmWave Sensing and Cross-domain Diffusion Model. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 8, 3 (2024), 1–27. https://doi.org/10.1145/3678550

[46] Qiang Zhu, Xin Tian, Chau-Wai Wong, and Min Wu. 2021. Learning your heart actions from pulse: ECG waveform reconstruction from PPG. *IEEE Internet of Things Journal* 8, 23 (2021), 16734–16748. https://doi.org/10.1109/JIOT.2021.3097946