

LubDubDecoder: Bringing Micro-Mechanical Cardiac Monitoring to Hearables

Siqi Zhang*

Carnegie Mellon University, USA
siqiz2@andrew.cmu.edu

Xiyuxing Zhang*

Tsinghua University, China
zxyx22@mails.tsinghua.edu.cn

Duc Vu*

Michigan State University, USA
vuduc2@msu.edu

Tao Qiang*

Shanghai Jiao Tong University, China
riderdecade@sjtu.edu.cn

Clara Palacios

Comillas Pontifical University, Spain
clara.palacios.spain@gmail.com

Jiangyifei Zhu

Carnegie Mellon University, USA
jiangyiz@andrew.cmu.edu

Yuntao Wang[†]

Tsinghua University, China
yuntaowang@tsinghua.edu.cn

Mayank Goel

Carnegie Mellon University, USA
mayankgoel@cmu.edu

Justin Chan[†]

Carnegie Mellon University, USA
justinchan@cmu.edu





Figure 1: LubDubDecoderSystem overview. (a) *LubDubDecoder* reconstructs the heart’s micro-mechanical signals from coarse-grained heart sounds recorded at the ear using microphones and speakers across a range of hearables. It identifies the timing of key micro-cardiac events, and enables hands-free monitoring of cardiovascular health in everyday scenarios. (b) Smartphone app user interface with hearables.

Abstract

We present *LubDubDecoder*, a system that enables fine-grained monitoring of micro-cardiac vibrations associated with the opening and closing of heart valves across a range of hearables. Our system transforms the built-in speaker, the only transducer common to all hearables, into an acoustic sensor that captures the coarse “lub-dub” heart sounds, leverages their shared temporal and spectral structure to reconstruct the subtle seismocardiography (SCG) and

gyrocardiography (GCG) waveforms, and extract the timing of key micro-cardiac events. In an IRB-approved feasibility study with 25 users, our system achieves correlations of 0.88–0.95 compared to chest-mounted reference measurements in within-user and cross-user evaluations, and generalizes to unseen hearables using a zero-effort adaptation scheme with a correlation of 0.91. Our system is robust across remounting sessions and music playback.

*Co-primary authors

[†]Co-corresponding authors



This work is licensed under a Creative Commons Attribution 4.0 International License.
CHI '26, Barcelona, Spain

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2278-3/2026/04

<https://doi.org/10.1145/3772318.3790445>

CCS Concepts

• Human-centered computing → Ubiquitous and mobile computing systems and tools.

Keywords

Hearables, Cardiovascular health, Wearable health monitoring, Well-being, Seismocardiography

ACM Reference Format:

Siqi Zhang, Xiyuxing Zhang, Duc Vu, Tao Qiang, Clara Palacios, Jiangyifei Zhu, Yuntao Wang, Mayank Goel, and Justin Chan. 2026. *LubDubDecoder: Bringing Micro-Mechanical Cardiac Monitoring to Hearables*. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 25 pages. <https://doi.org/10.1145/3772318.3790445>

1 Introduction

Fine-grained micro-mechanical vibrations of the chest reveal the timing of key cardiac events, such as the opening and closing of heart valves, and can be used for screening heart conditions including heart failure, ischemia, and atrial fibrillation [42, 52, 63, 75]. These vibrations are measured using seismocardiography (SCG) and gyrocardiography (GCG), which capture chest motion through linear acceleration and rotation, respectively [47, 77, 83, 98].

Collecting these signals, however, typically requires a clinical setting in which the patient lies down, removes their shirt, and is instrumented with accelerometers and gyroscopes [98]. Furthermore, recordings are typically limited to a few minutes due to time constraints in busy clinics, as well as patient discomfort during prolonged sessions.

The ability to enable unobtrusive monitoring of these micro-cardiac events outside the clinic allows for measurements in a naturalistic setting free from clinical stressors such as the “white coat effect” that can spuriously elevate cardiovascular measurements [30]. Such monitoring can support timely interventions including lifestyle changes to physical activity and diet [14, 35, 66, 76, 88]. It is also particularly useful for populations with mobility restrictions, such as older adults who are at highest risk of cardiac abnormalities [11, 64].

Given cardiac events send a pressure wave through the body that causes body parts to vibrate, the use of smartphone inertial measurement units (IMUs) appear at first to be an accessible way to capture these signals. However, the measured waveforms are highly sensitive to minute variations in placement: small shifts of just a few centimeters on the chest can substantially alter the recordings due to differences in local tissue composition and device coupling, making it difficult to compare results across sessions [73]. This challenge is amplified when measurements are performed by lay users outside clinical settings, where consistent and precise placement can be difficult to ensure.

Hearables, by contrast, occupy a fixed anatomical position in the ear, minimizing placement variability across removal and re-mounting sessions. Given that micro-cardiac vibrations propagate from the chest to the ear canal, hearables present themselves as a promising platform for convenient everyday monitoring of these signals.

We present *LubDubDecoder*, the first hearable system that bring micro-mechanical cardiac monitoring out of the clinic and enables monitoring in everyday scenarios across a range hearables (Fig. 1). While prior work [31] has reconstructed SCG using IMU-equipped hearables, IMUs are typically only found on true wireless stereo earbuds [31] or higher-end earphones such as those supporting active noise cancellation [17]. Other approaches measure SCG using IMUs on smartphones [40, 53, 91] and contactless mmWave radar [39], however, these require strapping or holding the phone against the

chest, or constraining users to remain in a fixed location and orientation. In contrast, our approach is broadly compatible across hearables including low-cost earphones and enables monitoring across environments, including during music playback.

However, there are several challenges in building such a system.

- (1) **Measuring micro-cardiac events at the ear.** Cardiac micro-vibrations are conventionally measured using chest-mounted IMUs because they are mechanical vibrations. However, IMUs are not available on all hearables. Further, these mechanical vibrations are attenuated and distorted as they propagate through the body channel consisting of bone, muscle, fat, and skin [17], making them difficult to sense at the ear.

Our key insight is that the familiar “lub-dub” sounds of the heart beating can be measured from the ear [15, 17], and leveraged to reconstruct the fine-grained SCG and GCG waveforms, enabling accurate extraction of the timing of micro-cardiac events. The intuition here is that because the heart sounds and vibratory SCG and GCG signals both originate from the same heart pumping mechanics, they carry shared temporal and spectral information [89] which can be learned to derive a cross-modal mapping between the signals.

To achieve this, we transform the speaker, *the only transducer common to all hearables*, into a microphone to capture and amplify the soft, coarse-grained heart sounds from the ear and reconstruct the SCG and GCG signal. Our system works across hearables with different designs and varied price points including bone-conduction earphones that contact the outer ear, over-ear headphones that enclose the ears, as well as both wired and wireless earbuds.

- (2) **Adapting to individual physiology.** As individuals differ in physiology, the channel between the heart and ear varies from person to person. To account for these effects, our system requires only a brief calibration lasting a few seconds, where the user places their smartphone against the chest while recording simultaneously from the hearable and the phone’s IMU. This lightweight human-in-the-loop process is used to finetune a pretrained model to preserve the timing of micro-cardiac events.

We note that periodic calibration is standard practice across both clinical and consumer cardiac monitoring due to drift in physiological signals from changes in body state (e.g. BMI) and skin properties. Clinic-grade cardiac monitoring devices [49], and commercial smart devices such as the CE-marked Aktiia wristband [90] and Samsung Galaxy Watch [6] require monthly recalibration to maintain accuracy. Our approach follows this established industry practice.

- (3) **Generalizing across hearables.** Measured heart sounds at the ear will vary across hearables due to sensor variability. As a result physiological signal mappings learned from one device cannot automatically translate to another device. To generalize to new hearable devices not seen during training, we apply a zero-effort normalization strategy that does not require any explicit calibration effort by the user. Specifically, ear-based cardiac sounds from the new device are normalized against recordings from a previously seen device, which serves as a reference to compute a set of weights. These weights are then applied to all subsequent recordings from the new device, enabling zero-effort adaptation.

This work contributes to the HCI community through:

- A novel sensing approach that repurposes the in-ear speakers for cardiac sensing, transforming everyday hearables into an equitable health monitoring platform that remains accessible to individuals with limited income or low technology literacy.
- An automated labeling pipeline for high-throughput segmentation and annotation of cardiac signals that can be applied to other cardiac health monitoring contexts.
- Experimental protocols and a motion artifact detection algorithm for assessing feasibility in free living deployments across everyday static and dynamic activities and during music playback, which other researchers can build on.

Below we summarize the main technical contributions of our system:

- *LubDubDecoder* reconstructs the fine-grained SCG and GCG waveforms from ear-based heart sounds and extracts the timing of key micro-cardiac events in an IRB-approved feasibility study with 25 subjects, achieving Pearson correlations between 0.88–0.95 compared to conventional chest-mounted measurements in both within-user and cross-user evaluations.
- Our approach generalizes across different hearable designs in hardware by leveraging built-in speakers and microphones, and in software through a zero-effort adaptation scheme that achieves waveform correlations of 0.91 in cross-device testing.
- Our system maintains robust performance across remounting sessions without recalibration, and in the presence of music playback from the hearable.
- Our motion artifact removal pipeline detects and removes unwanted motion events with $97.7 \pm 2.2\%$ accuracy.
- We perform a user experience survey demonstrating high levels of ease of use (4.6 of 5) and system trustworthiness (4.3 of 5).

Contribution to hearables and cardiac sensing research in HCI. Conceptually, to the best of our knowledge, we are the first HCI system that is focused around using the *built-in speaker*, the only transducer common to all hearables, to perform cardiac monitoring across broader range of hearables of varying price points. In contrast, prior works use dedicated IMUs or in-ear microphones which are not present on most classes of hearables in particular low-cost earphones [31]. *We show that despite weaker acoustic output, speaker-based sensing can enable robust cardiac monitoring at comparable performance to in-ear mics (Fig. 13).* This specific design choice is what allows us to have the benefits of greater ubiquity without trading off system performance.

We focus on extracting SCG/GCG signals as it is unique amongst cardiac signals in its ability to measure *sub-audible mechanical events not measurable by PCG* and can assess cardiac conditions [41, 62, 63, 70, 71, 78, 81, 82, 84, 85, 97], and in some cases even outperforming ECG electrical signals [95]. This particular set of cardiac signals is under-explored in comparison to prior works which have focused on extracting ECG [15], PCG [17], respiration volume monitoring [58], and arterial pressure waveform [21].

Enabling practical SCG/GCG monitoring requires contributions from a system design perspective:

- (1) *Cross condition characterization.* We perform a detailed characterization of micro-cardiac waveform variability in **Sec. 3** for intra-session, inter-session, inter-user and inter-device scenarios which motivates the design of our system, and reveals (a) the susceptibility of conventional smartphone-based SCG/GCG monitoring to placement variability (b) the comparative stability of in-ear SCG/GCG signals even after remounting (c) the need for per-person and per-device calibration.
- (2) *Novel calibration framework.* Based on these findings, we introduce in **Sec. 3.4.5 (i)** a human-in-the-loop adaptation procedure for new individuals which incorporates a brief simultaneous measurement using the smartphone on the chest and the hearable, and (ii) a zero-effort normalization strategy for new hearables without any explicit effort on the part of the user that normalizes the frequency response of a new device with recordings from a previously seen device.
- (3) *Micro-cardiac event timing reconstruction.* With this distinct setup design and formulation, we are able to reconstruct the timing of fiducial points from the ear in cross-session, cross-user and cross-device scenarios (**Sec. 5.3**) which is not shown in prior works [31].

Ethics and consent. This work includes (1) a feasibility study evaluating the cardiac waveform reconstruction performance of our system, (2) a post study user experience survey, and (3) recordings to characterize the cardiac signals and train a motion artifact detection classifier. For each of these components, informed written consent was obtained prior to participation and prior to any data collection. All these procedures were conducted under IRB approval (STUDY2025_00000155) and complied with relevant ethical regulations.

2 Background and related work

Cardiac signals at the ear. When the heart beats, it produces sounds in particular the characteristic “lub-dub” sound of heart valves closing which can be measured at the chest using a stethoscope in a process known as phonocardiography (PCG). Beyond these audible sounds, the heart’s mechanical activity generates minute chest wall motions that can be measured through seismocardiography (SCG) [98] and gyrocardiography (GCG) [47].

SCG captures fine-grained linear vibrations of the chest wall, while GCG provides complementary information on rotational dynamics. Unlike PCG, which is limited to audible valve closures, SCG and GCG sense both audible and sub-audible mechanical events of the heart. These include clinically important but acoustically silent events such as aortic valve opening, isovolumetric moment, or rapid ejection. Together, SCG and GCG extend the capabilities of health assessments beyond PCG and can be used in the assessment of arrhythmias, myocardial infarction, ischemia, and hemorrhage [41, 62, 63, 70, 71, 78, 81, 82, 84, 85, 97], and in some cases have even outperformed the electrical signals of the heart (ECG) in detecting coronary heart disease during exercise [95].

Prior work [16] has shown that ear-based cardiac sounds can be used to reconstruct the PCG signal. As the PCG signals align closely with the SCG and GCG signal in magnitude ratios and temporal trends (Fig. 2), this supports the view [89] that the two are coupled views of the same underlying cardiac mechanics. This also suggests

Reference	Sensor	Micro-mechanical cardiac monitoring	Ubiquitous sensors	Setup-free across environments	Generalizes across hearables
Ha et al. (2020) [39]	mmWave radar	✓	✗	✗	N/A
Fan et al. (2021) [29]	Hearable mic	✗	✓	✓	✓
Gilliam et al. (2022) [34]	Hearable infrasonic mic	✗	✗	✓	✗
Fan et al. (2023) [28]	Hearable mic and speaker	✗	✓	✓	✗
Butkow et al. (2023) [15]	Hearable in-ear mic	✗	✓	✓	✗
Cao et al. (2024) [16]	Hearable mic	✗	✓	✓	✗
Chen et al. (2024) [17]	Over-ear headphone speaker	✗	✓	✓	✓
Li et al. (2024) [57]	Surface acoustic wave microphone	✗	✗	✓	N/A
Fu et al. (2025) [31]	Hearable IMU	✓	✗	✓	✗
Islam et al. (2025) [44]	Hearable IMU and PPG sensor	✗	✗	✓	✗
Liu et al. (2025) [58]	Hearable mic	✗	✓	✓	✗
Christofferson et al. (2025) [21]	Hearable mic	✗	✓	✓	✗
LubDubDecoder (ours)	Hearable mics and speakers	✓	✓	✓	✓

Table 1: Comparison of representative cardiac sensing systems. Representative works from the past five years that focus on biosignal reconstruction for mobile and wireless devices. *LubDubDecoder* is designed to extract micro-mechanical cardiac vibrations (SCG and GCG) across a range of hearables.




Figure 2: Cardiac signal timing diagram. Each heartbeat begins with the electrical depolarization of the ventricles, measurable by the ECG. This is followed by the mitral valve closing which creates the “lub” (S1) sound of the heart, which generates an acoustic signal detectable at the ear. Later in the cardiac cycle, the aortic valve closes, producing the “dub” (S2) sound. The heart’s mechanical motion can be captured at using SCG and GCG.

that cardiac sounds recorded at the ear can be used to reconstruct the SCG and GCG signal.

Prior work on cardiac monitoring with mobile devices. While prior work on hearables have been designed for cardiac monitoring [15–17, 21, 28, 29, 34, 57, 58] they are mostly focused on reconstructing PCG, ECG and audioplethysmography (APG, which is

similar to PPG) signals, arterial pressure waveforms, and respiration volume from in-ear cardiac sounds, we present the first system to reconstruct micro-mechanical cardiac signals across a range of hearables, from ear-based cardiac sounds and extract the timing of micro-cardiac events.

Recent IMU-based approaches including BallistoBud [44] and TWSCardio [31] have demonstrated the feasibility of capturing cardiac motion using earbud accelerometers. A limitation of this approach is that IMUs are specialized sensors found mainly in true wireless stereo earbuds [31] or high-end hearables [17]. In contrast, our approach reconstructs fine-grained cardiac waveforms from widely available in-ear microphones and speakers, which is an approach that applies more broadly across even low-end hearables.

Smartphone IMU-based approaches [40, 53, 91] have measured SCG by strapping the phone to the chest or holding it in place. However, this approach is uncomfortable and this method is not practical for long-term monitoring across multiple hours.

Radar-based systems [39] have leveraged mmWave to capture chest vibrations corresponding to SCG. This approach requires users to remain directly in front of the radar within an operational range of 25–50 cm, and relies on comparatively expensive devices ($\approx$ \$350 [2]). Because the radar is typically mounted in a fixed position, it requires users to carry and reposition it wherever they want to take a measurement. In contrast, our approach is inherently mobile and can be worn throughout daily life across different environments, and with greater flexibility on user position.

Relation to active sonar approaches. Prior work has explored active sonar techniques using inaudible ultrasonic signals. APG [28] captures a PPG-like signal from earbuds by transmitting tones in the 30–39 kHz range. However, a limitation of this approach is that most commercial earbuds cannot transmit above 30 kHz. A related system, EarMonitor [80], uses a 16–21 kHz chirp for heart rate monitoring, but the lower end of this range remains audible to young adults [48], potentially causing irritation.




Figure 3: Challenge of conventional IMU-based micro-cardiac measurements. Differences in sensor placement lead to variations in waveforms, making comparisons across repeated measurements challenging. Precise and consistent placement is difficult to ensure when measurements are performed by lay users outside clinical settings. Each waveform corresponds to a cycle of 800 ms, and amplitudes are normalized to their own maximum.

We note that an advantage of ultrasonic active sonar approaches is that they can operate during music playback, as their probing signals lie outside the frequency range of typical music. However, our passive sensing system is also able to operate during music playback as SCG and GCG signals occupy low frequencies (< 50 Hz) that lie outside the dominant energy bands of music and environmental sound.

3 System design

3.1 Designing for accessible, everyday use

Our design was guided by principles of accessibility, comfort, and everyday usability. We prioritized a hardware-agnostic approach that works across a wide range of hearables including low-end ones to reduce adoption barriers and cost. Leveraging the built-in speaker and microphone enables passive monitoring without requiring users to wear dedicated sensors or require active user participation, which lowers the barrier to adoption particularly for older adults who may have cognitive impairment or have lowered technology literacy rates. The ear location was chosen for its consistent coupling and comfort during long-term wear. These considerations align with prior HCI work emphasizing unobtrusive, low-effort health monitoring and sustainable engagement [50, 59].

3.2 Conventional approach of measuring micro-cardiac signals

Conventionally, the SCG and GCG signal is measured using standalone or smartphone IMUs [40, 98]. In a clinical setting, the optimal placement has been shown to be between the third and fifth intercostal space [73, 89]. However, it is challenging to ensure that lay users can perform these precise readings repeatedly.

To underscore the importance of placement accuracy, we investigate the impact of smartphone position on SCG waveform variation. To do this, we placed the smartphone at nine different locations on the torso as shown in Fig. 3. Placements 1 and 2 were on the right lateral ribcage, placements 7 and 8 on the left lateral ribcage, placements 4 and 5 on the sternum, and placements 3, 6, and 9 on the lower abdomen. More specifically, placements 2, 5, and 8 were located between the 3rd and 5th intercostal spaces near the sternum.

Our results show that the SCG signal varies markedly across these nine positions, with an average Pearson correlation of 0.37 ± 0.34 between mean cardiac cycles. While most locations (8 out of 9) produced visually recognizable SCG signals, their morphologies differed. At placement 3, no clear cardiac events could be identified, likely because it was the site farthest from the heart.

We then considered five key fiducial points reflecting micro-cardiac events from medical literature [23, 37]: mitral valve closure (MC), isovolumetric moment (IM), aortic valve opening (AO), maximal acceleration (MA), and rapid ejection (RE). We observed that at sternum locations (4 and 5) and upper torso sites (1, 4, 7), events such as AO and IM appeared sharp, crisp, and stereotyped, exhibiting well-defined peaks. In contrast, at the abdomen or lower ribs, waveforms became broader and more attenuated; some cardiac events appeared smeared, with delayed or ambiguous onsets. These variations are due to differences in channel propagation where tissue, muscle, and fat create time delays and frequency-dependent attenuation. These results are supported by prior work [73] which have shown that IMU-based SCG measurements are highly sensitive to placement.

This analysis reveals a limitation of current approaches for micro-mechanical cardiac measurements outside the clinic, and points to the potential advantage of hearables, which maintain a relatively fixed position in the ear and are therefore less sensitive to placement variation. We examine this benefit in further detail in the next subsection.

3.3 Characterizing cardiac signals measured on mobile devices.

The goal of this initial set of experiments is to assess the robustness of ear-based cardiac sounds, SCG, and GCG signals to three different real-world conditions: (1) repeated remounting of devices, (2) inter-user differences, and (3) inter-device differences. We quantify the variability introduced by these factors and describe calibration strategies to mitigate their impact.

3.3.1 Hearables for capturing cardiac sounds. Our system measures ear-based cardiac sounds through different hearables which can be divided into three different form factors (1) **in-ear earbuds**, (2) **over-ear headphones**, and (3) **bone conduction headphones**, that each receive the heart sounds in different ways.

- **In-ear earbuds.** The in-ear microphones and speakers on earphones are able to detect cardiac sounds due to the occlusion effect [79], a phenomenon that amplifies low-frequency body sounds when the ear canal is sealed. The rubber ear tips of commercial earbuds naturally create this seal and have been shown to capture sounds from the heart.

	Ear sounds	SCG	GCG
Intra-session	0.88 ± 0.08	0.89 ± 0.06	0.87 ± 0.06
Inter-session	0.82 ± 0.09	0.73 ± 0.19	0.79 ± 0.07
Inter-user	0.31 ± 0.19	0.03 ± 0.24	-0.03 ± 0.36
Inter-device	0.15 ± 0.36	0.22 ± 0.19	0.42 ± 0.29

Table 2: Cardiac waveform variability across different experimental conditions. Waveform variability is measured using Pearson correlation coefficient (mean $\pm$ std). Ear sounds are recorded from the in-ear speaker. SCG and GCG are recorded from the smartphone IMU against the chest.




Figure 4: Dataset collection setup. (a) Ear-based cardiac sounds are measured using the microphone or speaker of a hearable; mechanical cardiac vibrations are measured at the left lower sternal border around the heart using a smartphone IMU. (b) Hearables used for data collection span a range of device types.

- **Over-ear headphones.** For the over-ear headphones which typically only contain a speaker, we leverage the principle of acoustic reciprocity which shows that the speaker diaphragm can work in reverse and be used to act as a microphone to measure the cardiac sounds [17]. It additionally takes advantage of the passive noise isolation of the ear cup design to pick up the subtle signals.
- **Bone conduction earphones.** The bone conduction earphones are positioned at the tragus of the ear and measure the lower-frequency cardiac sounds from pulse waves which causes tissue compression and subtle deformations of the skin. We note that these sounds have been picked up by surface acoustic wave microphones at the carotid artery (neck area) [56].

To learn the mapping between the ear-based cardiac sounds, and the micro-mechanical cardiac events at the heart, our system collects paired cardiac sounds at the ear through a hearable, and the micro-mechanical signal through a smartphone IMU. Details of the hearables and smartphone used and their data acquisition settings are described in **Sec. 4**.

		Intra session	Inter session	Inter user	Inter device
SCG	AO-MC	2.6	6.7	29.3	31.2
	AO-IM	1.4	2.6	18.8	19.7
	AO-MA	5.0	3.5	13.1	14.6
	AO-RE	7.6	9.1	26.2	27.6
GCG	AO-MC	2.2	8.5	27.8	31.7
	AO-IM	3.9	5.5	13.8	18.4
	AO-MA	6.8	1.7	24.8	12.4
	AO-RE	7.3	7.2	29.7	23.9

Table 3: Timing variability of fiducial point timing across different experimental conditions. Standard deviations (ms) for timing from aortic opening (AO) to four fiducial points in SCG and GCG from the smartphone IMU against the chest.

3.3.2 Measuring waveform variation. We characterize waveform variation using two metrics:

- (1) **Waveform similarity.** We use the Pearson correlation coefficient to quantify waveform similarity, computing it across pairwise comparisons of cardiac cycles within a stream, across sessions, and across users or devices. The coefficient ranges from -1 to 1 , where -1 indicates a perfect negative correlation and 1 indicates a perfect positive correlation.
- (2) **Stability of fiducial timing in SCG and GCG signals.** For SCG and GCG, we use the AO as the anchor point as it is often the most distinct micro-cardiac event, and we measured the temporal distances to the other four events and assessed their variability across cycles.

Our algorithms for signal segmentation, waveform reconstruction, and fiducial point labeling are described in greater detail in **Sec. 3.4.2**.

Intra-session. This analysis provides a baseline for cardiac cycle variability over time within an individual. To characterize this, we had a single participant wear the wireless earbuds and place the smartphone on their chest without removal or re-mounting over the course of 2 minutes. We then segmented the ear-based cardiac sounds, SCG, and GCG signals into individual cardiac cycles of fixed duration (800 ms) corresponding to the typical length of a single cardiac cycle [94] (further description of algorithm in **Sec. 3.4.2**). We note that this participant was a member of the study team that provided informed written consent prior to any data collection.

We plot these signals visually in Fig. 5 which shows that these three cardiac signals retain consistent shape. We quantify this in Table 2, and show that the mean correlation across cycles was 0.87 ± 0.11 for ear-based cardiac sounds, 0.86 ± 0.07 for SCG, and 0.87 ± 0.06 for GCG. For the timing of the micro-cardiac events in the SCG signal, the mean standard deviations between the AO and the other fiducial points ranged from 1.4 to 7.6 ms (Table 3).

Inter-session. We next consider how much the signals change when the device is removed and re-mounted. This is an important practical consideration as hearables are frequently taken off and worn again in daily life. To evaluate this, we recruited two participants as before and perform seven remounting trials. In each trial,




Figure 5: Effect of device remounting on cardiac signals. Within a single session, cardiac signals show similar morphology across cycles ($n = 60$ cycles). After remounting the hearable and smartphone, ear-based cardiac sounds and micro-mechanical signals maintain comparable waveform shapes, with a modest increase in variability across cycles. Colored opaque line is mean signal across cardiac cycles, all cardiac cycles are overlaid in translucent gray. Each waveform corresponds to a cycle of 800 ms, and amplitudes are normalized to their own maximum.




Figure 6: Effect of individual physiology on cardiac signal variability. Waveform variability is presented for a random subset of $n = 6$ subjects from our human subjects study showing differences in ear-based cardiac sounds, SCG, and GCG signals ($n = 60$ cycles). Solid opaque lines represents the mean across all cardiac cycles within each subject, all cardiac cycles are overlaid in translucent color. Each waveform corresponds to a cycle of 800 ms, and amplitudes are normalized to their own maximum.




Figure 7: Effect of hardware differences on cardiac signals. Hardware differences across (a) hearables create differences in the measured ear-based cardiac sounds ($n = 60$ cycles) and (b) across smartphones create differences in the measured SCG and GCG signals ($n = 15$ cycles). These differences motivate the need for our zero-effort calibration procedure upon use of a new device. Blue line is the mean signal across cardiac cycles, shaded blue region represents one standard deviation from the mean, all cardiac cycles are overlaid in translucent gray. Each waveform corresponds to a cycle of 800 ms, and amplitudes are normalized to their own maximum.

the hearable was removed and re-mounted, as well as the smartphone for SCG and GCG data collection, followed by a 60-second recording session. We note that these two participants are members of the study team that provided informed written consent prior to any data collection.

As shown in Fig. 5 all three cardiac signals retain a similar shape for a single user both intra-session and across remounted sessions. This is quantified across both recruited users with the mean Pearson correlation of the ear-based cardiac sounds across all remounts as 0.82 ± 0.09 which is slightly lower, but still comparable to the intra-session correlation of 0.87 ± 0.11 (Table 2). This stability is likely because hearables occupy a fixed anatomical position at the ear or head, with limited room for placement variability. The SCG signal was more sensitive to remounting. When the smartphone was repositioned on the participant's body at approximately the same location, the waveform similarity decreased to 0.73 ± 0.19 from the intra-session similarity of 0.86 ± 0.07 .

When repeating the fiducial point timing analysis, the variation in temporal distances between AO and the other fiducial points ranged from 2.6 to 9.1 ms, which is on average 1.3 ms higher than the intra-session variability reported earlier (Table 3). This increased variability likely arises from small but unavoidable differences in device placement, orientation, and coupling to the body.

We evaluate the end-to-end effect of remounting variation on system performance in the evaluation in Sec. 5.7 and show that the system is able to operate effectively across remounting sessions.

Inter-user. Beyond temporal and remounting variability, another critical factor is how cardiac signals vary across different individuals. Inter-user variability reflects intrinsic physiological differences

that affect how hearables and smartphones capture the signals. To assess this, we analyzed recordings collected from 6 randomly selected subjects from our human subjects study using the same hearable device (in-ear mic from the wireless earbuds). We describe in Sec. 4.2 details of the population and note that all participants provided informed written consent prior to any data collection.

As shown in Fig. 6 across all three cardiac signals, there is substantial waveform morphological diversity. Quantitatively, the waveform similarity exhibited low correlations 0.31 for the ear-based cardiac sounds, and essentially no similarity of 0.03 ± 0.24 and -0.03 ± 0.36 for the SCG and GCG waveforms respectively (Table 2). The fiducial points' distance standard deviations for both the SCG and GCG were in the range 13-30 ms, which is larger than that exhibited within a single individual which was 1-8 ms (Table 3). This suggests that mechanical event timing is highly individualized, reflecting differences in cardiac physiology.

Given intrinsic differences across individual physiologies which affect the heart-to-ear channel, these results suggest that calibration is needed to adapt to physiological differences for new individuals.

Inter-device. Finally, we examined the influence of device hardware. Even when worn by the same user, different hearable devices introduce variability due to differences in microphone and speaker characteristics in particular sensitivity, and frequency response. Similarly, hardware differences in smartphone IMUs result in variability in the measured SCG and GCG waveform.

To evaluate device-related variability, we recorded ear-based cardiac sounds from three participants, each using five different hearable devices (Fig. 4), and the SCG and GCG signal across four different smartphones (iPhone 12, Motorola Edge, Google Pixel 7,

and Samsung Galaxy S9). We note that these three participants are members of the study team that provided informed written consent prior to any data collection. We show the variation in the physiological signals for one of the participants in Fig. 7 which illustrates that the ear-based cardiac sounds exhibited substantial differences across devices. We then show the quantitative variation for this inter-device variation averaged across three users and show the Pearson correlation for the ear-based cardiac sounds, SCG, and GCG signal was 0.15 ± 0.36 , 0.22 ± 0.19 , and 0.42 ± 0.29 respectively (Table 2), which is lower than the intra-session and inter-session similarities. The fiducial point timing variability is also higher, ranging from 12 to 31 ms for the SCG and GCG signals, in comparison to the intra-session variability of 1 to 8 ms (Table 3).

These differences in hardware motivate the need for a device-specific normalization procedure when deploying our method on unseen hardware. Without such normalization, the learned transformation between signals may not generalize, since the same underlying cardiac activity may manifest differently across devices.

3.4 Reconstructing micro-cardiac signals

We designed a signal processing and machine learning pipeline to reconstruct the SCG and GCG signals from ear-based cardiac signals (Fig. 8).




Figure 8: LubDubDecoder system pipeline.

3.4.1 Motion artifact removal pipeline. In practical use, the in-ear cardiac sounds can be corrupted by the subject’s body and head movements, which introduce motion artifacts into the audio recording. The challenge with detecting head and body motions such as chewing, nodding, or walking is that they generate low-frequency components, which can only be partially filtered out in the frequency domain. Furthermore, these motion signals are often non-periodic, and their amplitude can be larger than the cardiac signals.

To address this, we designed a data-driven, temporal segmentation approach to identify and discard segments affected by user movement, rather than relying on band-pass filtering alone. We

opted to classify and discard 10-second segments that contain motion which avoids the risk of leaving behind short clean fragments that may be too sparse or unreliable to use.

We recorded a dataset containing six types of motion events: brow raising, chewing, nodding, talking, and walking. All recordings were collected from a single subject using in-ear microphones. This subject was a member of the study team who provided informed written consent prior to any data collection. For each acoustic event, we recorded 10-second segments, a duration that is in line with standard practice in large scale audio datasets such as AudioSet [33], which uses this duration to capture sufficient acoustic context for a diverse range of events for use in machine learning models. This choice does not limit how frequently our system can output results because (1) a real-time implementation can use a sliding window with a small overlap between windows and (2) our waveform reconstruction model can run in real time and complete inference on a 10 second clip in 19 ms on a smartphone.

We also recorded a set of recordings when the user was static and when there was no external interference or music playback (Fig. 9a). Fig. 9b shows that the various motion types produced distinct signal distortions that were clearly visible and would obscure the smaller heart sounds.




Figure 9: Effect of user motions on in-ear cardiac sounds. (a) Dataset composition used to train our motion artifact removal classifier, showing the number of 10-second audio recordings in each class. (b) Motion artifacts produce signals with much higher amplitude than heart sounds, obscuring them. Our system automatically detects and discards segments affected by such artifacts.

To distinguish between segments with and without motion, for each audio segment, we preprocess them using Mel-Frequency Cepstral Coefficients (MFCCs), a widely used feature representation in audio analysis [36]. MFCCs capture the short-term power spectrum of a signal in a way that reflects how humans perceive sound. In our

context, they provided a compact and informative summary of each segment’s audio characteristics, which can allow a downstream model to differentiate between silent and motion-contaminated recordings. We evaluate the performance of our motion artifact removal pipeline in the evaluation **Sec. 5.7** and demonstrate it can detect and remove time windows with unwanted motion artifacts.

3.4.2 Signal processing. We describe the various signal processing steps below to condition the paired ear-based cardiac sounds, SCG and GCG signals for training our biosignal reconstruction model.

Signal conditioning. To ensure that all three cardiac signal streams have consistent dimensions, we first resampled them to 500 Hz. We then applied a 4th-order Butterworth bandpass filter (5–45 Hz) to the signals. This filter choice is consistent with prior literature [67] which notes that the energy of micro-cardiac signals primarily resides below 50 Hz. The lower cutoff at 5 Hz reduces the effects of respiration artifacts, and other low frequency noise. Finally, each segment was z-score normalized (zero mean, unit variance) to minimize the effects of amplitude variability across recordings and improve model training stability.

Cardiac cycle segmentation. To prepare the signals for the model, our goal is to segment the ear-based cardiac sounds and the target SCG and GCG signals into individual cardiac cycles. Segmentation was anchored to the most prominent peak of each modality: the S1 peak for ear-based cardiac sounds and the AO peak for SCG and GCG signals.

We first identified the S1 and AO anchor points by applying a peak-finding algorithm across the signal while enforcing a minimum spacing of 0.55 s between peaks, which corresponds to a maximum plausible heart rate of 110 bpm, slightly above the upper bound of resting heart rate in healthy adults. For each valid S1 or AO, we then extracted a fixed 800 ms window (200 ms before to 600 ms after the peak), which approximates the duration of a cardiac cycle at a resting heart rate of 75 bpm.

We additionally apply heuristics specific to each cardiac signal type: (1) *Ear-based cardiac sounds.* Each heartbeat produces two heart sounds, S1 and S2. While S1 is generally more prominent, in some users S2 can appear equally strong, leading to ambiguity in peak detection. To differentiate between the two, we compared the average energy in two 400 ms windows placed before and after each detected peak. For true S1 peaks, the right window after the peak would contain the S2 sound and would have higher energy compared to the left window before the peak which would contain background noise. The opposite pattern occurs if S2 peaks were initially detected.

(2) *SCG and GCG signals.* In SCG and GCG recordings, the AO peak is typically the most prominent feature of each cycle, but inter-subject variability can cause neighboring peaks such as MC or RE to appear equally strong. To improve robustness, we first apply the general peak-finding procedure to identify an initial AO candidate. We then examine all peaks within a narrow 200 ms window centered on this candidate. For each peak, we compute a local prominence score defined as the sum of its amplitude and those of its immediate left and right neighbors. The peak with the highest score is then selected as the definitive AO peak.

Multi-cycle training. While single-cycle inputs allow the model to learn detailed intra-beat morphology, they provide limited information about the inter-beat features across consecutive heartbeats. To address this, we incorporate a multi-cycle training strategy in which the continuous signal is segmented into 3-second windows that contain multiple cardiac cycles. These longer segments enable the model to capture inter-beat dynamics especially when the user’s heart rate changes across time, such as after exercise, or during sleeping. Combined with single-cycle training, this hybrid approach allows the model to learn both fine-grained intra-beat features and broader inter-beat patterns, improving overall reconstruction fidelity.

To help the model learn inter-beat patterns across a wide range of heart rates, we apply a stretching-based data augmentation strategy [39]. Specifically, each segment is time-stretched by factors ranging from 0.8 to 1.3, effectively simulating slower and faster heart rates. This augmentation improves the model’s ability to generalize to different heart rates during real-world use.

SNR-guided filtering. After segmentation, to ensure that only high-quality ear-based cardiac sounds are used for reconstruction, we computed the signal-to-noise ratio for each cardiac cycle and discarded those below a threshold of 7 dB. Because each cardiac cycle contains two distinct heart sounds (S1 and S2), we define a 400 ms window around the S1 peak that captures both S1 and S2 as the signal region, while the remainder of the cycle is treated as background noise. Over a continuous recording, this produces alternating segments of signal and noise. We then compute the power of the signal segments and compare it with that of the noise segments to derive the overall SNR. This measure allows us to evaluate the reliability of each ear’s recording and to select the channel with the higher SNR for subsequent reconstruction tasks.

For each region, the average power was computed as

$$P = \frac{1}{N} \sum_{i=1}^N x[i]^2,$$

where $x[i]$ are the amplitudes and N is the number of samples in the region. We denote P_{signal} and P_{noise} as the power within the signal and noise regions, respectively. The SNR is then computed as

$$\text{SNR} = 10 \log_{10} \left(\frac{P_{\text{signal}}}{P_{\text{noise}}} \right).$$

3.4.3 Model architecture. Similar to prior works on biosignal reconstruction [16, 57], we designed a temporal autoencoder with a multi-branch convolutional architecture to learn a cross-modal mapping from in-ear cardiac signals to SCG waveforms (Fig. 8). Our design directly uses raw time-domain signals as input as this preserves the full temporal resolution and allows the model to learn latent structure without compression artifacts associated with transformations like spectrograms and MFCCs.

Autoencoder design. Our model consists of an encoder-decoder structure tailored for cross-modal reconstruction. The encoder transforms the input into a lower-dimensional latent space that captures temporal and physiological features shared between in-ear cardiac sounds and SCG and GCG signals.

The encoder employs a two-branch convolutional neural network with progressively increasing channel dimensionality (64 →

128 $\rightarrow$ 256 $\rightarrow$ 512). The local branch uses dilated 1D convolutions with small kernel sizes (3×1) to capture fine-grained temporal patterns within individual cardiac cycles, while the global branch employs larger receptive fields (48×1) to model inter-beat context and periodic structure. The outputs from both branches are concatenated and refined through a temporal self-attention module that dynamically reweights feature contributions over time. Each dual-path block is followed by MaxPooling for temporal downsampling and Dropout for regularization. Residual connections and batch normalization layers are integrated throughout to stabilize training and ensure smooth gradient flow.

Following the encoder, three additional convolutional layers with 512 channels and batch normalization further refine the latent representation, preserving temporal continuity while enhancing local feature selectivity. The decoder mirrors the encoder through four upsampling blocks, each consisting of an upsampling layer (scale factor of 2) followed by a convolutional layer (kernel size 3), batch normalization, and LeakyReLU activation. Channel dimensionality is progressively reduced from $512 \rightarrow 256 \rightarrow 12 \rightarrow 64 \rightarrow 32$ to reconstruct the temporal resolution of the original signal. Attention modules are integrated within both the encoder's dual-path blocks and decoder to enhance temporal coherence and focus on salient cardiac events. Finally, a 1D convolution layer projects the intermediate feature maps into a single output channel representing the reconstructed SCG or GCG waveform.

We train the model using the Adam optimizer and a mean squared error loss. To improve convergence, a learning rate scheduler is used during training.

3.4.4 Fiducial point labeling algorithm. We develop an automated pipeline to label fiducial points in SCG and GCG signals. Each cardiac cycle contains five micro-cardiac events that occur in a fixed order: mitral valve closure (MC), isovolumetric moment (IM), aortic valve opening (AO), maximal blood acceleration (MA), and rapid ejection (RE).

After segmenting the waveform into individual cycles, the global maximum within each cycle is identified as the AO peak, which serves as the anchor for locating the remaining fiducial points. The MC point is defined as the maximum peak to the left of AO, and the RE point as the maximum peak to the right. The IM point is selected as the minimum amplitude within the [MC, AO] interval, while the MA point is selected as the minimum amplitude within the [AO, RE] interval.

3.4.5 Calibration. As shown in Table 2 and 3, the ear-based cardiac signals, SCG, and GCG waveforms vary across users due to intrinsic physiological differences, and across devices due to hardware variability. This highlights the importance of incorporating calibration strategies to enable more accurate and robust signal reconstruction. **Cross-user calibration.** For cross-user generalization, we first train a base model across data from all users in our training cohort. When encountering a new user, we collect a short calibration recording of five cardiac cycles (approximately four seconds) using the smartphone's IMU which is then used to fine-tune the pre-trained model, allowing it to adapt to the new user's physiology with minimal effort. We show in our evaluation in Fig. 12 that five cycles are sufficient to enable strong cross-user performance while keeping calibration effort minimal.

Cross-device normalization. We leveraged frequency-domain equalization to normalize signals across devices. *For this normalization step, only data from the hearable is needed, not the smartphone. In this way, no explicit user effort is required for this normalization* For each device, we first extracted ten consecutive cardiac cycles (approximately 8 seconds) and computed their mean waveform. Let $x_{\text{ref}}(t)$ and $x_{\text{tgt}}(t)$ denote the mean cycles for the reference and target devices, respectively. Their Fourier transforms are computed using the FFT $\mathcal{F}$:

$$X_{\text{ref}}(f) = \mathcal{F}\{x_{\text{ref}}(t)\}, \quad X_{\text{tgt}}(f) = \mathcal{F}\{x_{\text{tgt}}(t)\}. \quad (1)$$

We derive a frequency-domain mapping function $H(f)$ by taking the ratio of the two spectra:

$$H(f) = \frac{X_{\text{ref}}(f) \cdot X_{\text{tgt}}^*(f)}{|X_{\text{tgt}}(f)|^2 + \epsilon}, \quad (2)$$

where $(\cdot)^*$ denotes the complex conjugate and ϵ is a small regularization constant to ensure numerical stability.

Given a new cardiac cycle $x_{\text{tgt}}^{(i)}(t)$ from the target device, we compute its Fourier transform $X_{\text{tgt}}^{(i)}(f)$, apply the mapping $H(f)$, and transform back to the time domain:

$$\hat{x}_{\text{tgt} \rightarrow \text{ref}}^{(i)}(t) = \mathcal{F}^{-1}\{H(f) \cdot X_{\text{tgt}}^{(i)}(f)\}. \quad (3)$$

Finally, we apply an energy normalization step to match the amplitude of the reference mean cycle using the L2 norm of the reference and new cardiac cycle:

$$\hat{x}_{\text{norm}}^{(i)}(t) = \alpha \cdot \hat{x}_{\text{tgt} \rightarrow \text{ref}}^{(i)}(t), \quad \alpha = \frac{\|x_{\text{ref}}(t)\|_2}{\|\hat{x}_{\text{tgt} \rightarrow \text{ref}}^{(i)}(t)\|_2}. \quad (4)$$

4 Feasibility study

4.1 Hardware setup

We designed a data collection platform (Fig. 4) that measures ear-based cardiac sounds across five hearables. Two used **in-ear microphones** and three used **in-ear speakers**:

- (1) Wireless earbuds (Honor Earbuds 3¹, **in-ear microphone**, obtained through a collaborative research agreement with Honor Device Co., Ltd.)
- (2) Custom wired earbuds (POM-2730L-HD-R², **in-ear microphone**)
- (3) Wired earbuds (Sephia SP3060³, **in-ear speaker**)
- (4) Over-ear headphones (Audio-Technica ATH-M30x⁴, **speaker**)
- (5) Wired bone conduction earphones (FSC Wired Bone Conduction⁵, **speaker**)

Audio was sampled at 16 kHz with 16-bit resolution. Audio from the wireless earbuds were streamed in real time via Bluetooth to a host computer. The Babyface AD/DA converter [3] was used to acquire signals from the speaker of the hearables for the purposes of demonstrating proof-of-concept. Subsequent miniaturization

¹<https://www.honor.com/global/audio/honor-earbuds-3-pro/>

²<https://www.digikey.com/en/products/detail/pui-audio-inc/POM-2730L-HD-R/7898330>

³<https://www.amazon.com/sephia-SP3060-Earbuds-Lightweight-Tangle-Free/dp/B0170RB9J9Q>

⁴<https://www.amazon.com/dp/B00HVLUQW8>

⁵<https://www.amazon.com/dp/B0915DB6DN>

efforts could focus on a low-cost integrated PCB design. For convenience, we collected data from a single random ear, however bilateral data collection is possible with the earbuds. A custom smartphone Android app was used to collect the subject's corresponding SCG and GCG data using the onboard IMU. We used the Google Pixel 7 with a sampling rate of 454 Hz. The smartphone was placed at the lower left sternal border between the third and fifth intercostal space where the SCG and GCG signal is the strongest [73, 89].

4.2 Dataset collection protocol

This study was approved by our Institutional Review Board (STUDY2025_00000155). All studies complied with relevant ethical regulations. Participants were recruited by word of mouth through our university campus. Informed written consent was obtained for human subjects participating in the study before any study procedures began. Randomization was not applicable and investigators were not blinded.

Participants above the age of 18 were eligible for the study. Exclusionary criteria include pregnant women, prisoners, adults with cognitive impairments, individuals with allergies or sensitivities to adhesives or sensors, individuals with implanted medical devices, such as pacemakers, individuals with significant mobility impairments. Individuals were screened via inclusion and exclusion criteria prior to the study. We recruited 25 adults with a female-to-male ratio of 0.39, of whom 19 had complete demographic records: mean age 36.4 ± 17.9 years, height 176.6 ± 7.8 cm, weight 72.0 ± 13.7 kg. All of whom were healthy and without history of cardiovascular disease (Fig. 10).

We collected data in a quiet room from using the hearables to record ear-based cardiac sounds, and the smartphone IMU for SCG and GCG data. To synchronize the cardiac signals between the hearable and the phone, we tapped both together twice. Participants then remained in a static seated position for 10 minutes as data was recorded. After each measurement session, we aligned the recordings based on the tap events. We collected a total of 710 minutes (11.8 hours) of synchronized data from the participants. In a deployment of our system, this synchronization could similarly be achieved by having the smartphone transmit an inaudible acoustic chirp (18–20 kHz) that is simultaneously received by the smartphone and hearable microphones.

For all 25 participants, we collected recordings from the in-ear speaker of wired earbuds, and for 15 of them, also from the in-ear microphone of wireless earbuds. In addition, for a subset of 3 participants recordings were obtained across all five hearables (Fig. 4).

5 Evaluation

5.1 Evaluation setup

We evaluate the performance of our system through three protocols:

- (1) *Within-user evaluation*, where we apply five-fold cross-validation on the waveforms collected for each subject. This essentially captures the upper bound of system performance using a model trained and tested only on data from a single individual.
- (2) *Cross-user evaluation*, where a model is pretrained on data from

all but one user, fine-tuned on a small number of synchronized cardiac cycles from the hearable and smartphone from the held-out user, and evaluated on the remaining cycles. This calibration step accounts for inter-individual physiological variability that must be adapted for during deployment.

(3) *Cross-device evaluation*, where the model is trained and tested from data measured from the same user, but across five different devices. We first train a model from one reference hearable device. For each unseen hearable device, we use a small number of cardiac cycles to compute normalization weights relative to the reference device's ear-based heart sounds. These weights are then applied to subsequent cycles from the unseen device, which are passed into the reconstruction model to generate the micro-cardiac signals. We note here, that we only need recordings for the hearable, not the smartphone for the normalization.

For our within-user and cross-user evaluation, we use signals collected from the in-ear speaker of wired earbuds and in-ear microphone of wireless earbuds (device 3 and 1 in Fig. 4 respectively).

5.2 Waveform reconstruction performance

Within-user. We first evaluate the reconstruction performance of SCG and GCG waveforms under a within-user setting, where both training and testing are performed on data from the same participant. When using the *in-ear speaker of wired earbuds*, across 25 subjects, the reconstructed signals closely match the ground truth, achieving average Pearson correlation scores of 0.94 ± 0.04 for SCG and 0.93 ± 0.04 for GCG (Fig. 11a).

In comparison, when using the *in-ear microphone of wireless earbuds*, the reconstruction performance was similar achieving 0.95 ± 0.05 for SCG and 0.95 ± 0.04 for GCG. These results demonstrate that we are able to achieve comparable reconstruction results from both in-ear speakers and microphones on two different devices.

Cross-user. We then perform cross-user evaluation by training a model on all but one user, and fine-tuning it on five calibration cardiac cycles from the held-out user. Under this protocol, the system using the *in-ear speaker* achieves mean Pearson correlation scores of 0.88 ± 0.07 for SCG and 0.88 ± 0.05 for GCG (Fig. 11b). Similarly, when using the *in-ear microphone*, the system achieved scores of 0.84 ± 0.13 for SCG and 0.85 ± 0.09 for GCG.

We note that Pearson similarity scores of 0.90 is considered very high in the context of medical research [61]. Prior work using mmWave sensors to reconstruct SCG achieved a correlation of 0.72 [39], while IMU-based earbuds reported a similarity of 0.92 [31]. Our approach, based on ear-based speakers and microphones, provides a complementary modality that expands micro-cardiac monitoring across a wide range of hearables.

Effect of number of calibration cycles. Fig. 12 shows how average Pearson correlation score varies with the number of calibration cycles. Performance increases from 83% with a single cycle to 88% with five cycles, after which performance gains plateau. Based on this trend, we select five calibration cycles for fine-tuning across all cross-user evaluations. Assuming an average adult heartbeat of 75 bpm, five calibration cycles corresponds to a brief 4-second calibration period.

Cross-device. When a users buys a new hearable, we apply a zero-effort normalization strategy that removes the need for users to




Figure 10: Demographic summary of participants in the human subjects study. Complete demographic data was collected for 19 of 25 participants. Gender data was collected for all 25 participants




Figure 11: Waveform reconstruction performance across 25 subjects. Results are shown for within-user and cross-user evaluations.

explicit perform a calibration step (see **Sec. 3.4.5**). Here, the in-ear speaker of wired earbuds served as the reference device, while the remaining four hearables (in-ear microphone of wired earbuds, in-ear microphone of wireless earbuds, microphone of over-ear headphones, or bone-conduction microphones) are target devices for testing.

We performed frequency-domain normalization on the ear-based cardiac sounds by using the mean of ten cardiac cycles between the reference and target device to compute equalization weights. Across data from four hearables and three users, the normalized waveforms achieved a mean Pearson correlation of 0.94 ± 0.02 when compared to the reference device.

We then trained a user-specific model for SCG reconstruction for each of the three users. Next, we applied the computed weights

to produce normalized cardiac sounds data from the hearables. We then applied the models to the normalized waveforms, to obtain the reconstructed SCG waveforms. Across three users and four devices, the system achieved an average Pearson correlation of 0.91 ± 0.04 (Fig. 13a), which is comparable to our within-user and cross-user results, demonstrating that our normalization approach enables cross-device generalization without explicit user calibration effort.

5.3 Timing accuracy of micro-cardiac events

We next evaluate the timing error of key micro-cardiac fiducial points (MC, IM, AO, MA, RE) between the reconstructed and reference waveforms. For each point, Table 4 reports the median and 95th percentile absolute errors with 95% confidence intervals (CI), along with a significance test across participants.




Figure 12: Effect of number of cardiac cycles used for calibration on cross-user SCG reconstruction performance.




Figure 13: Cross-device performance. (a) SCG waveform reconstruction (b) fiducial point timing error. The ear-based heart sounds from different hearables are normalized to a single reference device (in-ear speaker of wired earbuds), and inference is performed on a model trained on that reference. Data is presented for three different subjects.

Within-user. Fig. 14a summarizes within-user performance, and Table 4 reports the corresponding timing errors for SCG and GCG. Using a Wilcoxon signed-rank test across participants, we found no significant difference in the timing errors for any fiducial point (SCG: all $p > 0.14$; GCG: all $p > 0.38$). These values translate to relative median timing errors of 0.0–0.5% and 95th percentile errors of 0.5–3.9%, assuming an 800 ms cardiac cycle typical of a healthy adult heartbeat at 75 bpm.

Cross-user. As illustrated in Fig. 14b, Table 4 reports cross-user timing errors for SCG and GCG. We again found no significant

differences across participants (SCG: all $p > 0.23$; GCG: all $p > 0.31$). These correspond to relative median timing errors of 0.0–0.5% and 95th percentile errors of 0.5–3.8%.

We note that prior work [39] using mmWave for SCG reconstruction obtains similar median timing errors of 0.3–1.6%. Furthermore, the within-session variability of fiducial point timing is often higher, ranging from 1–8 ms for SCG and GCG (Table 3).

Cross-device. We evaluated the timing accuracy of reconstructed SCG fiducial points after applying cross-device normalization. As shown in Fig. 13b and Table 4, the system achieved comparable




Figure 14: Timing accuracy of micro-cardiac events. System performance is shown for SCG and GCG timing accuracy as collected on the in-ear speaker on earbuds.

timing accuracy to within-user and cross-user settings. Using a Wilcoxon signed-rank test across participants, we found no significant difference in timing errors for any fiducial point (all $p > 0.16$). These correspond to relative median errors of 0.0–0.5% and 95th percentile errors of 0.75–3.5% of the cardiac cycle.

5.4 Subgroup analysis

We performed a subgroup analysis on the cross-user waveform reconstruction results to evaluate system performance across different demographic dimensions (Fig. 15).

Body mass index. Among the 25 participants that participated in our study, only 19 of them shared their BMI data. 11% ($n = 2$) of subjects were underweight ($BMI < 18.5$), 63% ($n = 12$) had normal weight ($18.5 \leq BMI < 24.9$), 21% ($n = 4$) were overweight ($25.0 \leq BMI < 29.9$), and 5% ($n = 1$) were obese ($BMI \geq 30$). For the underweight group, the system achieved average reconstruction

correlation of 0.93 ± 0.02 (SCG) and 0.82 ± 0.09 (GCG). In the normal weight group, similarities were 0.85 ± 0.08 (SCG) and 0.88 ± 0.25 (GCG). Among overweight subjects, results were 0.87 ± 0.09 (SCG) and 0.87 ± 0.01 (GCG). Finally, in the obese group, the system reached 0.90 ± 0.09 (SCG) and 0.89 ± 0.10 (GCG).

Sex. In our study, 28% ($n = 7$) of subjects were female and 72% ($n = 18$) of subjects were male. When comparing system performance, female subjects had an SCG reconstruction similarity of 0.90 ± 0.03 and a GCG similarity of 0.88 ± 0.07 , while male subjects had an SCG similarity of 0.87 ± 0.07 and a GCG similarity of 0.88 ± 0.05 .

Age. Among the 19 participants whom we have full demographic information in our study, 37% ($n = 7$) of them were among the age of 18–24, 21% ($n = 4$) were between 25–32, 32% ($n = 6$) were 41–55, and 11% ($n = 2$) were 70+. For the 18–24 age group, our system achieved average reconstruction correlation of 0.88 ± 0.05 (SCG) and 0.89 ± 0.04 (GCG). In the 25–32 age group, our system achieved




Figure 15: Subgroup analysis on cross-user waveform reconstruction performance by (a) BMI, (b) gender, (c) age.

Within-user (SCG)					
	MC	IM	AO	MA	RE
Median	2.0[2.0,2.0]	2.0[2.0,2.0]	0.0[0.0,0.0]	2.0[2.0,2.0]	2.0[2.0,4.0]
95th	18.0[16.0,20.0]	8.0[6.0,12.0]	2.0[2.0,2.0]	18.0[16.0,22.0]	20.0[15.1,27.1]
Within-user (GCG)					
	MC	IM	AO	MA	RE
Median	4.0[2.0,4.0]	2.0[2.0,2.0]	0.0[0.0,0.0]	2.0[2.0,2.0]	4.0[2.0,4.0]
95th	30.6[26.0,37.2]	14.0[10.0,20.0]	2.0[2.0,4.0]	16.0[14.0,20.0]	18.0[16.0,20.6]
Cross-user (SCG)					
	MC	IM	AO	MA	RE
Median	4.0[4.0,4.0]	2.0[2.0,2.0]	0.0[0.0,0.0]	4.0[4.0,4.0]	4.0[4.0,4.0]
95th	26.0[24.0,26.0]	24.0[22.0,24.0]	2.0[2.0,2.0]	22.0[20.0,22.0]	22.0[22.0,24.0]
Cross-user (GCG)					
	MC	IM	AO	MA	RE
Median	4.0[4.0,4.0]	2.0[2.0,4.0]	2.0[2.0,2.0]	4.0[4.0,4.0]	4.0[4.0,4.0]
95th	28.0[26.0,30.0]	18.0[18.0,20.0]	4.0[4.0,6.0]	16.0[16.0,18.0]	30.0[29.8,34.0]
Cross-device (SCG)					
	MC	IM	AO	MA	RE
Median	2.0[2.0,4.0]	2.0[2.0,2.0]	0.0[0.0,0.0]	4.0[2.0,4.0]	4.0[4.0,4.0]
95th	20.0[18.0,22.0]	20.0[16.0,22.5]	6.0[4.5,8.0]	18.0[16.0,18.0]	28.0[26.0,30.0]

Table 4: Timing error (ms) of micro-cardiac fiducial points between reconstructed and reference waveforms. Values are median and 95th percentile absolute error with 95% confidence intervals.

reconstruction correlation of 0.87 ± 0.09 (SCG) and 0.92 ± 0.03 (GCG). In the 41-55 age group, our system achieved reconstruction correlation of 0.86 ± 0.04 (SCG) and 0.85 ± 0.04 (GCG). For the 70+ age group, our system achieved reconstruction correlation of 0.82 ± 0.09 (SCG) and 0.78 ± 0.05 (GCG). The results show that our system continues to achieve reliable performance in older adults.

Demographic classification. We perform an exploratory analysis to assess whether the reconstructed SCG signals could be used to infer demographic attributes, and focus on the task of sex classification. From the reconstructed signals, we extracted features including linear predictive coding coefficients, mean and standard deviation. Using signals from the in-ear speaker of earphones, the binary classification accuracy reached 89%. We did not perform a similar analysis for BMI, as our dataset does not have a sufficiently

balanced distribution across BMI categories. Our classification result aligns with prior work [86] which report that sex-related differences in body composition influence SCG characteristics and can be used for classification.

5.5 Generalization across real-world conditions

Heart rate fluctuates naturally across days, increases after physical activity, and decreases during sleep. To evaluate whether our system can generalize across these real-world variations, we conducted three sets of experiments on two participants: (1) longitudinal testing across multiple days and weeks, (2) evaluation while lying down during a nap, and (3) evaluation after exercise.

We first trained the model on each user’s calm, static data using our multi-cycle training strategy and stretching-based augmentation. We then tested the model on data collected on different days, during sleep-like conditions, and after exercise. To quantify inter-beat accuracy, we report the inter-beat interval (IBI) percentage error, defined as the absolute timing error normalized by the true inter-beat interval:

$$IBI_{err} = 100 \times \frac{|IBI_{pred} - IBI_{true}|}{IBI_{true}} \quad (5)$$

Longitudinal performance across time. Our system maintains consistently high performance across 1-day, 5-day, and 2-week spans for both earphones and headphones (Fig. 16). For earphones, SCG IBI percentage errors remained stable over time, with median errors of 0.5%, 0.5%, and 0.5% at 1 day, 5 days, and 2 weeks, respectively, and corresponding 95th-percentile errors of 1.7%, 2.0%, and 1.5%. For GCG, median errors were 0.8%, 0.6%, and 0.5%, with 95th-percentile errors of 2.0%, 6.0%, and 1.1%. For headphones, performance was similarly stable, with SCG median errors of 0.6%, 0.5%, and 0.9% at 1 day, 5 days, and 2 weeks, and 95th-percentile errors of 4.1%, 1.8%, and 2.6%. GCG median errors were 0.7%, 0.8%, and 0.9%, with 95th-percentile errors of 2.1%, 3.5%, and 2.1%. These results demonstrate that once trained on calm data, the model generalizes across natural day-to-day and week-to-week heart rate changes.

Lying down while sleeping. Here we had the participants take a one hour nap, during which their heart rate slowed, under these conditions our system continued to perform robustly (Fig. 17a). When using earphones, SCG median error was 0.5% (95th percentile




Figure 16: Longitudinal effect of time on system performance. CDFs show the system’s IBI percentage error for SCG and GCG reconstruction across one-day, five-day, and two-week intervals for both the (a) wired earphones and (b) headphones. Performance remains stable over time, demonstrating that the model generalizes across natural day-to-day variations in heart rate.




Figure 17: Generalization across scenarios that perturb heartrate: (a) sleeping and (b) after exercise. CDFs illustrate the system’s IBI percentage for reduced heart rate during a nap and elevated heart rate following exercise. Results show that the model maintains accurate inter-beat timing across both reduced and increased heart rate conditions.

1.5%), and GCG median error was 0.5% (95th percentile 1.1%). For headphones, SCG had a median of 0.9% (95th percentile 2.6%), and GCG a median of 0.9% (95th percentile 2.1%). Performance remains in line with results measured when the user is awake, indicating robustness to naturally reduced heart rates during rest and sleep-like conditions.

Exercise. We had participants jog on the spot for 10 minutes and measured system performance after this exercise session (Fig. 17b). For earphones, SCG median errors were 1.6% (95th percentile 4.5%),

and GCG median error increased to 1.5% (95th percentile 4.4%). These errors are modest given the rapid heart rate changes post-exercise.

5.6 Motion artifact removal

We evaluate the performance of our motion artifact removal pipeline on an in-ear microphone audio dataset from the wireless earbuds (Fig. 9a) collected when the user is static and moving (brow raising, chewing, nodding, talking, and walking), which we describe




Figure 18: Motion artifact removal pipeline performance. Evaluation is performed on in-ear audio segments when the user is static and moving. (a) PCA projection of the first two components of MFCC audio features illustrating separability between the two classes. (b) ROC curve from five-fold cross validation. (c) Confusion matrix indicating optimal operating point on the ROC curve.

in Sec. 3.4.1. To assess the separability between these two classes, we applied principal component analysis (PCA) on the MFCCs of the audio and visualized the projection of the first two components (Fig. 18a) which shows separability between the classes.

We trained a Random Forest classifier to distinguish between the static and motion recordings using 10-second segments as input. The model was evaluated using 5-fold cross-validation, achieving a ROC AUC of 0.994 ± 0.011 (Fig. 18b) and a mean accuracy of $97.7 \pm 2.2\%$ (Fig. 18c).

Free-living deployment. To evaluate our motion artifact removal pipeline in realistic conditions, we conducted a one-hour recording session in which a participant wore the *in-ear speaker of wired earbuds* while going about daily activities. As shown in Fig. 19, the system accurately detects and discards all motion-induced artifacts throughout the recording. In contrast, during stationary activities such as desk work our system consistently captures clean ear-based cardiac signals. Importantly, our system captures the in-ear cardiac signals during brief stationary intervals between dynamic activities, and even short pauses between walking are sufficient to extract a clean cardiac signal. Finally, the system is also able to correctly suppress transient noise that occurs during stationary activities such as snoring when the participant is taking a nap. Overall, this real-world evaluation demonstrates that our motion removal pipeline can reliably identify clean cardiac segments across a wide range of everyday scenarios.

5.7 Benchmarks

Effect of device remounting. To evaluate the effect of device-remounting on end-to-end system performance, we first trained a within-session model between paired hearable and smartphone recordings as a baseline. We then trained our model on remounted recordings, where both the hearable and the smartphone were independently removed and remounted, and tested on unseen remounted data. The baseline model achieved an average similarity score of 0.95 ± 0.02 , while models evaluated on unseen remounted data achieved an average similarity score of 0.89 ± 0.02 . To put this in context, the intrinsic variability of the SCG waveform within a single user is 0.86 (Table 2). *This suggests that device-remounting has only a modest effect on end-to-end system performance.*

Effect of music playback. We evaluated the effect of music playback by having a subject wear wireless earbuds while music was played through the in-ear speaker at volume levels of +6, +9, +12, +15, +18, and +21 dB above ambient noise, with recordings collected from the in-ear microphone.

As shown in Fig. 20, the energy from music and the SCG signal largely reside in non-overlapping frequency ranges with music being concentrated in higher frequencies (>45 Hz). As a result, the overall shape of the SCG waveform and the fiducial points can still be observed both with and without a 5–45 Hz band-pass filter in the time domain. We evaluated end-to-end system performance by assessing waveform-reconstruction similarity for the different sound levels. Across all sound levels, waveform similarity remains high, with an average Pearson correlation of 0.95 ± 0.01 with frequency filtering and 0.94 ± 0.01 without. *This shows that our system continues to perform even in the presence of music playback.*

Effect of ear tip diameter and material. We selected one foam ear tip and three rubber ear tips with diameters of 11, 12, 13, and 14 mm, that were coupled to the in-ear earphone (Device 3 in Fig. 4b). A single participant wore the device with each ear tip, and ear-based cardiac signals were recorded from the built-in speaker. Fig. 21 shows the waveforms recorded from each of the ear tips. The average Pearson correlation across the four eartips was 0.88 ± 0.06 , suggesting that for the common ear tip sizes and materials used for adults, it has minimal effect on the morphology of the measured cardiac signals.

To further examine whether ear tips influence system performance, we evaluated cross-ear tip reconstruction accuracy. Specifically, we trained the model using data recorded with one ear tip and tested it on signals recorded with the others. Reconstruction similarity remained consistently high across ear tips, with an average Pearson correlation of 0.89 ± 0.05 . These results indicate that variations in ear tip size or material have little impact on overall system performance.

System latency. We evaluated the inference latency of our autoencoder model on a HUAWEI Pura 70 Ultra smartphone by exporting our PyTorch model to the ONNX format for mobile deployment. The system was run on in-ear audio chunks of 10000 ms, yielding a mean and standard deviation inference latency for each segment




Figure 19: Free-living deployment across different stationary and dynamic activities. The figure shows the in-ear audio signal from the speakers of wired earbuds with activity labels. Red shaded segments denote periods flagged by our motion detection algorithm. All dynamic activities are correctly flagged, no stationary activities were flagged, and the cardiac signal is visible during these stationary periods.




Figure 20: Effect of music playback. As micro-cardiac signals and music largely occupy non-overlapping frequency bands, the overall shape of the cardiac cycle and fiducial points are visible even without a 5–45 Hz bandpass filter.




Figure 21: Effect of ear tip on ear-based cardiac sounds. Mean cardiac cycles recorded from one participant using four different ear-tip types.

of 19.4 ± 2.4 ms, satisfying the timing requirements for real-time execution.

5.8 User experience survey

Beyond evaluating waveform reconstruction performance, we sought to understand aspects of the user experience that influence real-world adoption rates and long-term engagement in everyday contexts. Specifically, we seek to understand user comfort, perceived unobtrusiveness, and attitudes towards in-ear cardiac sensing. We evaluated user experience of our system through a short survey. 23 of the 25 participants from the human subjects study consented to a post-study user experience survey that was sent over email, the remaining 2 participants did not respond. Informed written consent was obtained before the survey began. Participants were sent a form containing the following questions:

- (1) Is the system easy to use? (1 = Not at all, 5 = Very easy)
- (2) Can you trust our system's heart sounds results? (1 = Not at all, 5 = Completely)
- (3) How much would you pay for this system? (USD)
- (4) Would this system make you more conscious of your cardiac health?
- (5) In what scenarios would you envision using this system?

Fig. 22 summarizes the results of the survey. Participants rated the system as easy to use (mean 4.6 ± 0.6) and trustworthy (mean 4.3 ± 0.9). The average amount participants would be willing to pay for the system was 57.7 ± 58.9 , though this value likely reflects the socioeconomic status of our recruitment population, which primarily included individuals from our institution and surrounding community. In terms of whether the system would make one more conscious of cardiac health, 70% ($n = 16$) said yes, 22% ($n = 5$) said maybe, and 9% ($n = 2$) said no.

Participants envisioned a wide range of scenarios for potential use, with two recurring themes from these results: *First*, participants were interested in the system's value to provide continuous monitoring during daily life and possibly even mitigate white coat effects of elevated cardiac measurements that can be found during clinical monitoring. One participant noted its possible value for cardiac disease prediction broadly beyond SCG and GCG monitoring given its potential to detect emergent conditions like cardiac arrest: *When I am very old and alone, it could call an ambulance*




Figure 22: User experience survey ($n = 23$). Histograms summarize user perceptions across (a) ease of use, (b) system trustworthiness, (c) amount willing to pay, and (d) whether the system made them more conscious of their cardiac health.

automatically if my heart stops. We note that related emergency alert systems are integrated into commercial products, including smartwatches from Apple [8], Samsung [9], and Google [4], which can automatically contact emergency services in the event of a fall or car crash.

Second, participants were interested in the system’s potential to monitor the effects of aerobic exercise, such as running, on overall cardiac activity. In the context of micro-cardiac monitoring, prior work [52, 72] has noted that comparing SCG signals before and after exercise can be useful for detecting coronary artery disease.

6 Discussion

Frequency and ease of recalibration. The frequency at which the system needs to be recalibrated for an individual depends on changes in body state and skin properties over time which affect the body propagation channel. Such recalibration is common in both clinic-grade cardiac monitoring devices [49] and wrist-worn smart devices for cardiac monitoring from Samsung [6] and Aktia [90], the latter two of which recommend monthly recalibration.

Our calibration procedure to account for physiological variation is lightweight: it only requires the user to wear their hearable and place their phone against their chest for approximately four seconds. Synchronization between the two devices can be achieved via an inaudible chirp transmitted from the phone and received by both devices, which is a technique used for standardizing across sensor streams [91].

Privacy considerations. While our system is designed for continuous operation, it is privacy-preserving by design and does not capture intelligible human speech. This is because human speech occupies frequencies > 80 Hz, with the majority of speech content between 300 and 3400 Hz [13], which our system removes in software and can also suppress through an equivalent analog filter in hardware. Specifically, our design uses a bandpass filter to isolate ear-based cardiac signals < 50 Hz [67]. Furthermore, because hearables sit in the ear canal with rubber eartips or ear cups, they provide passive acoustic isolation that attenuates ambient sounds. Finally, all audio data are processed entirely on device in real time, and no data leave the user’s local devices.

Battery life. We measured the battery consumption of our neural network model running in real time on a Samsung Galaxy S9 smartphone using the ONNX runtime framework with NNAPI hardware

acceleration to use the device’s neural processing unit. When the model processes 800 ms chunks of audio data continuously, the battery drains at a rate of 8.2% an hour, which equates to over 12 hours of uninterrupted monitoring. In comparison, when the model is not running and the phone remains idle, the battery drains at a rate of 7.6% an hour. This shows that our inference pipeline only introduces a small overhead and is practical for continuous use.

Extended hearable use. We note that our system does not require wearing hearables for extended durations to have health benefits. As it performs brief, opportunistic measurements during everyday moments when individuals wear hearables, such as commuting, it can provide a richer picture of cardiac health than episodic clinic visits.

6.1 Positioning and future directions in HCI

We position *LubDubDecoder* as a HCI contribution at the intersection of earable physiological sensing and personal informatics.

Cardiac sensing with hearables. While there is rich prior HCI work that has explored cardiac monitoring from the ear [21, 44, 58], our key sensing approach is novel in its use of the *built-in speaker* which uniquely enables generalization across a broader range of devices, given that it is the only transducer common to all hearables. Despite speakers producing lower amplitude signals than microphones, we are able to achieve comparable results between the two transducers (Fig. 13). *This distinction motivates the need for carefully designed interfaces that take into account the known attitudes, behavior, and psychology of users of mobile health technology [32] when designing trustworthy user interfaces.*

Micro mechanical cardiac monitoring with SCG and GCG remains under-explored in earables, and to the best of our knowledge no prior earable work has focused on constructing the timing of fiducial events corresponding to valve opening and closing. This focus matters because SCG and GCG encode cardiac mechanical timing, which complements prior work measuring electrical, acoustic and volumetric signals [15–17, 21, 28, 29, 34, 57, 58]. While many cardiac signals are coupled through shared cardiovascular hemodynamics, they differ in physical origin, conventional sensing modality.

Personal informatics. A popular area in the HCI literature that our work can be categorized under is personal informatics. Work in this space has examined why people abandon self tracking tools and has produced multi-stage models of how tracking unfolds over

Common failure modes in personal health informatics	Design response
Erosion of trust due to inconsistent readings	Characterization of increased variability over conventional chest based IMU methods
Context switching and competing activities	Evaluation during music playback; evaluation while lying down; motion artifact detection to enable automatic signal quality screening
Lapses in data quality after interruptions	Ear worn placement supports repeatable remounting with minimal variability
Device replacement and turnover	Zero effort cross device calibration enables continuity when the hearable is replaced or upgraded
Longitudinal drift and reduced comparability over time	Stability evaluated across time; design of a periodic lightweight human in the loop calibration to maintain longitudinal comparability

Table 5: Mapping common failure modes in personal health informatics systems identified by Epstein et al. [26, 27] to corresponding design responses in our system.

time [19, 26, 54, 55], including iterative stages spanning “preparation, collection, integration, reflection, and action”, often noting that whether users are “goal-oriented” versus “curiosity-oriented” can affect how they approach these stages and how much they care about the underlying accuracy of the system, and the amount of data that is presented to them [69].

Across this literature, a consistent theme is that self tracking breaks down when it depends on careful or repeated manual setup, when routines are disrupted, or when use becomes intermittent due to friction, device replacement, or unmet information needs [26].

We view the micro-mechanical cardiac monitoring system through a personal informatics lens, and not just as a sensing lens. Work by Epstein et al. [26, 27] underscores the importance of developing systems that function “in the wild reality” with context switching, interruptions and inconsistency which are central to everyday self tracking. They further note that designing for these realities is central to a HCI contributions.

Here, we design techniques that take into account different real-world issues that target specific breakdown points identified in the HCI personal informatics literature that prevent users from reaching “reflection” and “action” stages which are critical for behavior change. We show in Table 5 various such breakdowns and how we mitigate them. *CHI learns that viewing physiological sensing through a personal informatics lens translates to design implications from the start to mitigate interruptions, device turnover, and trust breakdowns, which can pave the way for adoption.*

6.1.1 Future directions in HCI. Future directions within the HCI community fall under three categories:

Self-tracking and behavior change. TummyTrials [50] observed that those who attempt diagnostic self-tracking often lack knowledge to identify triggers that cause negative health conditions and thus may not reach their goals of managing their health conditions. Future iterations of our work can enable diagnostic self-tracking [69] by mitigating common pitfalls identified from a community of self trackers [20] through two approaches:

(1) We can implement lightweight activity recognition classifiers [60] on the recorded audio to capture the user’s context and behavior, instead of simply discarding it as a motion event such as in the current system implementation. This can help to draw correlations between changes in cardiac rhythm and the user’s activity to identify triggers that could worsen cardiac conditions such as the trigger of arrhythmias. *Concretely, this shifts the system from producing isolated measurements to producing context-tagged episodes that better support “reflection” and “action” in personal informatics terms.*

This also aligns with the methodology Drunk User Interfaces [59] which focused around the need to control for confounding effects such as user fatigue which would affect measurements of the end-point of interest. In our case measuring pertinent variables such as posture, location, time of day, and producing a stratified performance reports to the user can help to ameliorate experimental confounds.

(2) We can further build-upon self-experimentation frameworks [51] that can identify causal relationships between potential triggers and cardiac abnormalities through single-case experimental designs that generate a randomized study protocol that dictate the user’s level of physical activity on a given day.

Clinical workflow integration. A recurring theme in HCI research [93] on patient generated data is that the value of self tracking does not automatically translate to value in clinical care. This is because data must cross boundaries between patient and clinician who view data with a different lens. Prior work [93] shows that breakdowns often occur when patient and clinician goals are misaligned.

HCI literature has framed interpretation of patient generated data as a collaborative sensemaking problem, emphasizing designs that reduce translation burden and support shared interpretation rather than expecting either party to independently derive meaning from raw data [22, 45, 65, 68, 93].

Building on this literature, an important future direction is to move beyond reliable sensing and explicitly support clinician mediated interpretation through interfaces that are aware of clinician workflows. We systematically provide future considerations and recommendations in Table 6 using the barriers identified by West et al. [93].

Human-AI partnerships. Recent work has centered around the use of human-AI partnerships [87] to inform personal health decisions by acting as collaborative peers that could proactively and automatically surface health issues. In the context of our system this could be accomplished through automated interpretation of continuous time series data [25] and integration with other digital phenotypes [43] to extract behavioral patterns to surface results for review by users or clinical professionals.

Barrier	Implication for interfaces and outputs
Structure and relevance	Generate visit ready summaries using familiar clinical conventions.
Reliability and completeness	Attach quality and coverage metadata that indicates when data was corrupted due to motion or low signal quality.
Context	Future work here could include an activity recognition module to correlate cardiac events with activities performed by the user as well as user posture.
Time burden, overload, and interoperability	Provide clinician dashboards with “triage views” that compress high volume time series into a small set of markers of interest, and support export formats that integrate with documentation workflows.

Table 6: Design implications aligned with barriers for patient generated data use in clinical care identified by West et al. [93].

This is an exciting future direction with unique challenges at the technical and user-centered design perspective. For instance, literature has shown that there is a danger to being overly trusting of AI decisions compared to those mediated by medical professionals [38]. Future work could focus on how such systems should incorporate known information about the user’s decision-making styles [74], the user’s health knowledge [24], pre-existing views about trust in AI [46] to design a collaborative experience between users and clinicians.

7 Limitations

Open-ear earbuds. We note that it is challenging to reliably record ear-based cardiac sounds from air-conduction open-ear earbuds [1]. All of the hearables evaluated in our study provide some form of isolation against noise, either through a rubber ear tip seal, a passive ear cup, or a bone-conduction transducer. While open-ear devices comprise only a modest share of the hearables market, estimated at about 2% [7], future designs could address these limitations by integrating active noise cancellation that compensate for the lack of passive isolation.

Usage during motion. Our cardiac sensing system is currently designed to operate during periods of stationary activity including sitting, standing, or lying down (e.g. during rest or sleep), and to ignore periods of motion, which overwhelm the cardiac signals measured at the ear. However, our system can still extract cardiac signals during brief pauses between movements such as short breaks during exercise, which capture the effect of exercise intensity on cardiac dynamics. We note while that motion artifact removal algorithms have been developed to extract heart rate from in-ear cardiac sounds using U-Net based architectures [15], future work would be required to evaluate whether such techniques can be robustly applied to measure the micro-mechanical cardiac signals and the timing information of the fiducial points.

8 Future Directions

Evaluation on subjects with cardiovascular conditions. We present a proof-of-concept study to evaluate how SCG and GCG measurements could be captured across different hearable devices and users. Future studies are needed to perform a clinical validation of this technology in patients with or at risk of cardiovascular disease. These studies would be needed to assess if this technology can be used for the detection or diagnosis of cardiac conditions. These studies could build on prior clinical protocols [72] that measure SCG before and after exercise to detect changes in cardiac function that could be used to detect coronary artery disease.

Joint reconstruction of micro-cardiac waveforms. Although the SCG and GCG signals originate from the same underlying cardiac mechanics, they capture different motion components. In practice, one modality, often the GCG can have a lower SNR because rotational motion can be more subtle and more easily masked by sensor noise. Future work could explore a joint reconstruction framework that exploits complementary information between the two signals and leverages the stronger modality to guide reconstruction of the weaker signal. Prior work [96] has shown that gyroscopic readings can improve automatic annotation of SCG fiducial points and used to better estimate heart rate variability and cardiac time intervals.

Hearing aids. Deployment of our system on hearing aids poses an interesting opportunity, particularly because these devices are widely adopted by older adults [18] which is the same demographic at elevated risk for both cardiac disease and hearing loss [12, 92].

Traditional hearing aids such as those from Oticon [5] place in-ear speakers deep in the ear canal for sound amplification. In principle, these speakers would be able to capture the ear-based cardiac sounds robustly. The main challenge however, is that the hearing aid speakers continuously amplify ambient speech and sounds, and conventional speakers cannot simultaneously transmit and receive audio at high fidelity.

More recent devices such as the Apple AirPods Pro, which is licensed as a clinical-grade hearing aid [10], incorporate both an in-ear microphone and speaker that would in principle be able to perform duplex operation. Here, the microphone could capture the cardiac sounds, while the speaker amplifies surrounding sounds. Given our evaluation (Sec. 5.7) showing that the system works reliably even during music playback due to the non-overlapping frequencies of music and the cardiac sounds, having the system work on the Apple AirPods Pro remains a potentially exciting area of future work.

References

- [1] 2025. Avantree Resolve – Wired Open-Ear Earbuds. <https://www.amazon.com/dp/B0B4JV5K4B>
- [2] 2025. AWR1443BOOST. <https://www.digikey.com/en/products/detail/texas-instruments/AWR1443BOOST/9860052>
- [3] 2025. Babyface Pro FS. <https://rme-audio.de/babyface-pro-fs.html>
- [4] 2025. Get help in an emergency with Google Pixel Watch safety features. <https://support.google.com/googlepixelwatch/answer/12663810>
- [5] 2025. Hearing aid brands & models. <https://www.hearinglife.com/hearing-aids/models-and-brands>
- [6] 2025. How Do I Recalibrate the Galaxy Watch? <https://www.samsung.com/sg/support/apps-services/how-do-i-recalibrate-the-galaxy-watch/>
- [7] 2025. The rise of open earbuds: challenges and opportunities. <https://www.canalys.com/insights/rise-of-open-earbuds-challenges-and-opportunities>

- [8] 2025. Use Fall Detection with Apple Watch. <https://support.apple.com/en-us/108896>
- [9] 2025. Use the Detect fall feature on your Samsung smart watch. <https://www.samsung.com/us/support/answer/ANS10003423/>
- [10] 2025. Use the Hearing Aid feature on your AirPods Pro 2. <https://support.apple.com/en-us/120992>
- [11] Wafik Farah Andrawes, Caroline Bussy, and Joël Belmin. 2005. Prevention of cardiovascular events in elderly people. *Drugs & aging* 22, 10 (2005), 859–876.
- [12] Rachael R Baiduc, Joshua W Sun, Caitlin M Berry, Melinda Anderson, and Eric A Vance. 2023. Relationship of cardiovascular disease risk and hearing loss in a clinical population. *Scientific reports* 13, 1 (2023), 1642.
- [13] Ronald J Baken. 1987. Clinical measurement of speech and voice. (No Title) (1987).
- [14] Pierre Bouteuyrie, Patrick Lacolley, Marie Briet, Véronique Regnault, Alice Stanton, Stéphane Laurent, and Azra Mahmud. 2011. Pharmacological modulation of arterial stiffness. *Drugs* 71 (2011), 1689–1701.
- [15] Kayla-Jade Butkow, Ting Dang, Andrea Ferlini, Dong Ma, and Cecilia Mascolo. 2023. hEART: Motion-resilient heart rate monitoring with in-ear microphones. In *2023 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, 200–209.
- [16] Alvin Cao, Ken Christofferson, Parker Ruth, Naveed Rabbani, Yuanchun Shi, Alex Mariakakis, Yuntao Wang, and Shwetak Patel. 2024. EarSteth: Cardiac Auscultation Audio Reconstruction Using Earbuds. In *2024 46th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 1–4.
- [17] Tao Chen, Yongjie Yang, Xiaoran Fan, Xiuzhen Guo, Jie Xiong, and Longfei Shangguan. 2024. Exploring the feasibility of remote cardiac auscultation using earphones. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*. 357–372.
- [18] Wade Chien and Frank R Lin. 2012. Prevalence of hearing aid use among older adults in the United States. *Archives of internal medicine* 172, 3 (2012), 292–293.
- [19] Eun Kyoung Choe. 2014. *Designing self-monitoring technology to promote data capture and reflection*. Ph. D. Dissertation.
- [20] Eun Kyoung Choe, Nicole B Lee, Bongshin Lee, Wanda Pratt, and Julie A Kientz. 2014. Understanding quantified-selfers' practices in collecting and exploring personal data. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1143–1152.
- [21] Kenneth Christofferson, Michelle Lin, Joseph A Cafazzo, and Alex Mariakakis. 2025. ArtEARial: Arterial Pressure Waveform Reconstruction Using Earbud Audio. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 9, 4 (2025), 1–24.
- [22] Chia-Fang Chung, Kristin Dew, Allison Cole, Jasmine Zia, James Fogarty, Julie A Kientz, and Sean A Munson. 2016. Boundary negotiating artifacts in personal informatics: patient-provider collaboration with patient-generated data. In *Proceedings of the 19th ACM conference on computer-supported cooperative work & social computing*. 770–786.
- [23] M Di Rienzo, E Vaini, P Castiglioni, G Merati, P Meriggi, G Parati, A Faini, and F Rizzo. 2013. Wearable seismocardiography: Towards a beat-by-beat assessment of cardiac mechanics in ambulant subjects. *Autonomic Neuroscience* 178, 1-2 (2013), 50–59.
- [24] Murat Dikmen and Catherine Burns. 2022. The effects of domain knowledge on trust in explainable AI and task performance: A case of peer-to-peer lending. *International Journal of Human-Computer Studies* 162 (2022), 102792.
- [25] Zachary Englhardt, Chengqian Ma, Margaret E Morris, Chun-Cheng Chang, Xuhai" Orson" Xu, Lianhui Qin, Daniel McDuff, Xin Liu, Shwetak Patel, and Vikram Iyer. 2024. From classification to clinical insights: Towards analyzing and reasoning about mobile and behavioral health data with large language models. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 8, 2 (2024), 1–25.
- [26] Daniel A Epstein, Monica Caraway, Chuck Johnston, An Ping, James Fogarty, and Sean A Munson. 2016. Beyond abandonment to next steps: understanding and designing for life after personal informatics tool use. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. 1109–1113.
- [27] Daniel A Epstein, An Ping, James Fogarty, and Sean A Munson. 2015. A lived informatics model of personal informatics. In *Proceedings of the 2015 ACM international joint conference on pervasive and ubiquitous computing*. 731–742.
- [28] Xiaoran Fan, David Pearl, Richard Howard, Longfei Shangguan, and Trausti Thormundsson. 2023. APG: Audioplethysmography for cardiac monitoring in hearables. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*. 1–15.
- [29] Xiaoran Fan, Longfei Shangguan, Siddharth Rupavatharam, Yanyong Zhang, Jie Xiong, Yunfei Ma, and Richard Howard. 2021. HeadFi: bringing intelligence to all headphones. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*. 147–159.
- [30] Stanley S Franklin, Lutgarde Thijs, Tine W Hansen, Eoin O'brien, and Jan A Staessen. 2013. White-coat hypertension: new insights from recent studies. *Hypertension* 62, 6 (2013), 982–987.
- [31] Yongjian Fu, Ke Sun, Ruyao Wang, Xinyi Li, Ju Ren, Yaoxue Zhang, and Xinyu Zhang. 2025. Enabling Cardiac Monitoring using In-ear Ballistocardiogram on COTS Wireless Earbuds. *arXiv preprint arXiv:2501.06744* (2025).
- [32] Sandra Gabriele and Sonia Chiasson. 2020. Understanding fitness tracker users' security and privacy knowledge, attitudes and behaviours. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [33] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter. 2017. Audio set: An ontology and human-labeled dataset for audio events. In *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 776–780.
- [34] Francis Roosevelt Gilliam III, Robert Ciesielski, Karlen Shahinyan, Pratistha Shakya, John Cunsolo, Jal Mahendra Panchal, Bartłomiej Król-Józaga, Monika Król, Olivia Kierul, Charles Bridges, et al. 2022. In-ear infrasonic hemodynography with a digital health device for cardiovascular monitoring using the human audiome. *NPJ Digital Medicine* 5, 1 (2022), 189.
- [35] Guilherme Veiga Guimarães, Emmanuel Gomes Ciolac, Vitor Oliveira Carvalho, Veridiana Moraes D'Ávila, Luiz Aparecido Bortolotto, and Edimar Alcides Bocchi. 2010. Effects of continuous vs. interval exercise training on blood pressure and arterial stiffness in treated hypertension. *Hypertension Research* 33, 6 (2010), 627–632.
- [36] Shikha Gupta, Jafreezal Jaafar, WF Wan Ahmad, and Arpit Bansal. 2013. Feature extraction using MFCC. *Signal & Image Processing: An International Journal* 4, 4 (2013), 101–108.
- [37] Viatcheslav Gurev, Kouhyar Tavakolian, Jason Constantino, Bozena Kaminska, Andrew P Blaber, and Natalia A Trayanova. 2012. Mechanisms underlying isovolumic contraction and ejection peaks in seismocardiogram morphology. *Journal of medical and biological engineering* 32, 2 (2012), 103.
- [38] Andrea L Guzman and Seth C Lewis. 2020. Artificial intelligence and communication: A human-machine communication research agenda. *New media & society* 22, 1 (2020), 70–86.
- [39] Unsoo Ha, Salah Assana, and Fadel Adib. 2020. Contactless seismocardiography via deep learning radars. In *Proceedings of the 26th annual international conference on mobile computing and networking*. 1–14.
- [40] Amin Hossein, Elza Abdessater, Paniz Balali, Elliot Cosneau, Damien Gorlier, Jérémy Rabineau, Alexandre Almorad, Vitalie Faoro, and Philippe Van De Borne. 2024. Smartphone-Derived Seismocardiography: Robust Approach for Accurate Cardiac Energy Assessment in Patients with Various Cardiovascular Conditions. *Sensors* 24, 7 (2024), 2139.
- [41] Zuhair Ifrikhar, Olli Lahdenoja, Mojtaba Jafari Tadi, Tero Hurnanen, Tuija Vasankari, Tuomas Kiviniemi, Juhani Airaksinen, Tero Koivisto, and Mikko Pänkäälä. 2018. Multiclass classifier based cardiovascular condition detection using smartphone mechanocardiography. *Scientific reports* 8, 1 (2018), 9344.
- [42] Omer T Inan, Maziyar Baran Pouyan, Abdul Q Javaid, Sean Dowling, Mozziyar Etemadi, Alexis Dorier, J Alex Heller, A Ozan Bicen, Shuvo Roy, Teresa De Marco, et al. 2018. Novel wearable seismocardiography and machine learning algorithms can assess clinical status of heart failure patients. *Circulation: Heart Failure* 11, 1 (2018), e004313.
- [43] Thomas R Insel. 2017. Digital phenotyping: technology for a new science of behavior. *Jama* 318, 13 (2017), 1215–1216.
- [44] Md Saiful Islam, Md Mahbubur Rahman, Mehrab Bin Morshed, David J Lin, Yunzhi Li, Hao Zhou, Wendy Berry Mendes, and Jilong Kuang. 2025. BallistoBud: Heart Rate Variability Monitoring using Earbud Accelerometry for Stress Assessment. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 110, 21 pages. <https://doi.org/10.1145/3706598.3714029>
- [45] Anna Sigridur Islind, Tomas Lindroth, Johan Lundin, and Gunnar Steineck. 2019. Shift in translations: Data work with patient-generated health data in clinical practice. *Health informatics journal* 25, 3 (2019), 577–586.
- [46] Alon Jacovi, Ana Marasović, Tim Miller, and Yoav Goldberg. 2021. Formalizing trust in artificial intelligence: Prerequisites, causes and goals of human trust in AI. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*. 624–635.
- [47] Mojtaba Jafari Tadi, Eero Lehtonen, Antti Saraste, Jarno Tuominen, Juho Koskinen, Mika Teräs, Juhani Airaksinen, Mikko Pänkäälä, and Tero Koivisto. 2017. Gyrocadiography: A new non-invasive monitoring method for the assessment of cardiac mechanics and the estimation of hemodynamic variables. *Scientific reports* 7, 1 (2017), 6823.
- [48] Milan Jilek, Daniel Suta, and Josef Syka. 2014. Reference hearing thresholds in an extended frequency range as a function of age. *The Journal of the Acoustical Society of America* 136, 4 (2014), 1821–1830.
- [49] Daniel W Jones, Lawrence J Appel, Sheldon G Sheps, Edward J Roccella, and Claude Lenfant. 2003. Measuring blood pressure accurately: new and persistent challenges. *Jama* 289, 8 (2003), 1027–1030.
- [50] Ravi Karkar, Jessica Schroeder, Daniel A Epstein, Laura R Pina, Jeffrey Scofield, James Fogarty, Julie A Kientz, Sean A Munson, Roger Vilaridaga, and Jasmine Zia. 2017. Tummytrials: a feasibility study of using self-experimentation to detect individualized food triggers. In *Proceedings of the 2017 CHI conference on human*

- factors in computing systems. 6850–6863.
- [51] Ravi Karkar, Jasmine Zia, Roger Vilardaga, Sonali R Mishra, James Fogarty, Sean A Munson, and Julie A Kientz. 2016. A framework for self-experimentation in personalized health. *Journal of the American Medical Informatics Association* 23, 3 (2016), 440–448.
 - [52] Iwona Korzeniowska-Kubacka, Maria Bilińska, and Ryszard Piotrowicz. 2005. Usefulness of seismocardiography for the diagnosis of ischemia in patients with coronary artery disease. *Annals of noninvasive electrocardiology* 10, 3 (2005), 281–287.
 - [53] Hyoung Youn Lee, Yong Hun Jung, Kyung Woon Jeung, Dong Hun Lee, Byung Kook Lee, Geuk Young Jang, Tong In Oh, Najmiddin Mamadjonov, and Tag Heo. 2021. Discrimination between the presence and absence of spontaneous circulation using smartphone seismocardiography: A preliminary investigation. *Resuscitation* 166 (2021), 66–73.
 - [54] Ian Li, Anind Dey, and Jodi Forlizzi. 2010. A stage-based model of personal informatics systems. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 557–566.
 - [55] Ian Li, Anind K Dey, and Jodi Forlizzi. 2011. Understanding my data, myself: supporting self-reflection with ubicomp technologies. In *Proceedings of the 13th international conference on Ubiquitous computing*. 405–414.
 - [56] Kaylee Yaxuan Li, Yasha Iravanchi, Hyunmin Park, Yiming Liu, and Alanson Sample. [n. d.]. ECG Signal Construction From Heart Sounds via Single Node, Surface Acoustic Sensing. In *EMBC '24*.
 - [57] Kaylee Yaxuan Li, Yasha Iravanchi, Hyunmin Park, Yiming Liu, and Alanson Sample. 2024. ECG Signal Construction From Heart Sounds via Single Node, Surface Acoustic Sensing. In *2024 46th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 1–4.
 - [58] Yang Liu, Qiang Yang, Kayla-Jade Butkow, Jake Stuchbury-Wass, Dong Ma, and Cecilia Mascolo. 2025. EarMeter: Continuous Respiration Volume Monitoring with Earables. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 9, 4 (2025), 1–29.
 - [59] Alex Mariakakis, Sayna Parsi, Shwetak N Patel, and Jacob O Wobbrock. 2018. Drunk user interfaces: Determining blood alcohol level through everyday smartphone tasks. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–13.
 - [60] Alessandro Montanari, Andrea Ferlini, Ananta Narayanan Balaji, Cecilia Mascolo, and Fahim Kawsar. 2023. Earset: A multi-modal dataset for studying the impact of head and facial movements on in-ear ppg signals. *Scientific data* 10, 1 (2023), 850.
 - [61] Mavuto M Mukaka. 2012. A guide to appropriate use of correlation coefficient in medical research. *Malawi medical journal* 24, 3 (2012), 69–71.
 - [62] Keya Pandia, Omer T Inan, Gregory TA Kovacs, and Laurent Giovangrandi. 2012. Extracting respiratory information from seismocardiogram signals acquired on the chest using a miniature accelerometer. *Physiological measurement* 33, 10 (2012), 1643.
 - [63] Mikko Pänkäälä, Tero Koivisto, Olli Lahdenoja, Tuomas Kiviniemi, Antti Saraste, Tuija Vasankari, and Juhani Airaksinen. 2016. Detection of atrial fibrillation with seismocardiography. In *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 4369–4374.
 - [64] Jonathan P Piccini, Bradley G Hammill, Moritz F Sinner, Adrian F Hernandez, Allan J Walkey, Emelia J Benjamin, Lesley H Curtis, and Susan R Heckbert. 2014. Clinical course of atrial fibrillation in older adults: the importance of cardiovascular events beyond stroke. *European heart journal* 35, 4 (2014), 250–256.
 - [65] Adrienne Pichon, Kayla Schiffer, Emma Horan, Bria Massey, Suzanne Bakken, Lena Mamykina, and Noemie Elhadad. 2021. Divided we stand: the collaborative work of patients and providers in an enigmatic chronic disease. *Proceedings of the ACM on human-computer interaction* 4, CSCW3 (2021), 1–24.
 - [66] Thomas G Pickering, William Gerin, and Amy R Schwartz. 2002. What is the white-coat effect and how should it be measured? *Blood pressure monitoring* 7, 6 (2002), 293–300.
 - [67] Deepak Rai, Hiren Kumar Thakkar, Shyam Singh Rajput, Jose Santamaria, Chintan Bhatt, and Francisco Roca. 2021. A comprehensive review on seismocardiogram: current advancements on acquisition, annotation, and applications. *Mathematics* 9, 18 (2021), 2243.
 - [68] Shriti Raj, Mark W Newman, Joyce M Lee, and Mark S Ackerman. 2017. Understanding individual and collaborative problem-solving with patient-generated data: challenges and opportunities. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW (2017), 1–18.
 - [69] John Rooksby, Mattias Rost, Alistair Morrison, and Matthew Chalmers. 2014. Personal tracking as lived informatics. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1163–1172.
 - [70] Prasan Kumar Sahoo, Hiren Kumar Thakkar, Wen-Yen Lin, Po-Cheng Chang, and Ming-Yih Lee. 2018. On the design of an efficient cardiac health monitoring system through combined analysis of ECG and SCG signals. *Sensors* 18, 2 (2018), 379.
 - [71] David M Salerno and John Zanetti. 1991. Seismocardiography for monitoring changes in left ventricular function during ischemia. *Chest* 100, 4 (1991), 991–993.
 - [72] David M Salerno, John M Zanetti, Liviu C Poliac, Richard S Crow, Peter J Hannan, Kyuhyun Wang, Irvin F Goldenberg, and Robert A Van Tassel. 1992. Exercise seismocardiography for detection of coronary artery disease. *American journal of noninvasive cardiology* 6, 5 (1992), 321–330.
 - [73] Richard H Sandler, Md Khushidul Azad, John D'Angelo, Peshala Gamage, Nirav Y Raval, Robert J Mentz, and Hansen A Mansy. 2020. Documenting spatial variation of SCG signals for optimal sensor placement. *Journal of Cardiac Failure* 26, 10 (2020), S92.
 - [74] Susanne G Scott and Reginald A Bruce. 1995. Decision-making style: The development and assessment of a new measure. *Educational and psychological measurement* 55, 5 (1995), 818–831.
 - [75] Mobashir Md Hasan Shandhi, Joanna Fan, J Alex Heller, Mozziyar Etemadi, Omer T Inan, and Liviu Klein. 2019. Seismocardiography and machine learning algorithms to assess clinical status of patients with heart failure in cardiopulmonary exercise testing. *Journal of Cardiac Failure* 25, 8 (2019), S64–S65.
 - [76] Shigeki Shibata, Naoki Fujimoto, Jeffrey L Hastings, Graeme Carrick-Ranson, Paul S Bhella, Christopher M Heaton Jr, and Benjamin D Levine. 2018. The effect of lifelong exercise frequency on arterial stiffness. *The Journal of physiology* 596, 14 (2018), 2783–2795.
 - [77] Szymon Sieciński, Paweł S Kostka, and Ewaryst J Tkacz. 2020. Gyrocardiography: A review of the definition, history, waveform description, and applications. *Sensors* 20, 22 (2020), 6675.
 - [78] Brian E Solar, Amirtaha Taebi, and Hansen A Mansy. 2017. Classification of seismocardiographic cycles into lung volume phases. In *2017 IEEE Signal Processing in Medicine and Biology Symposium (SPMB)*. IEEE, 1–2.
 - [79] Michael A Stone, Anna M Paul, Patrick Axon, and Brian CJ Moore. 2014. A technique for estimating the occlusion effect for frequencies below 125 Hz. *Ear and hearing* 35, 1 (2014), 49–55.
 - [80] Xue Sun, Jie Xiong, Chao Feng, Wenwen Deng, Xudong Wei, Dingyi Fang, and Xiaojian Chen. 2023. EarMonitor: In-ear motion-resilient acoustic sensing using commodity earphones. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 4 (2023), 1–22.
 - [81] Mojtaba Jafari Tadi, Eero Lehtonen, Tero Hurmanen, Juho Koskinen, Jonas Eriksson, Mikko Pänkäälä, Mika Teräs, and Tero Koivisto. 2016. A real-time approach for heart rate monitoring using a Hilbert transform in seismocardiograms. *Physiological measurement* 37, 11 (2016), 1885.
 - [82] Amirtaha Taebi and Hansen A Mansy. 2017. Grouping similar seismocardiographic signals using respiratory information. In *2017 IEEE signal processing in medicine and biology symposium (SPMB)*. IEEE, 1–6.
 - [83] Amirtaha Taebi, Brian E Solar, Andrew J Bomar, Richard H Sandler, and Hansen A Mansy. 2019. Recent advances in seismocardiography. *Vibration* 2, 1 (2019), 64–86.
 - [84] Kouhyar Tavakolian. 2010. Characterization and analysis of seismocardiogram for estimation of hemodynamic parameters. (2010).
 - [85] Kouhyar Tavakolian, Andrew P Blaber, Alireza Akhbardeh, Brandon Ngai, and Bozena Kaminska. 2010. Estimating cardiac stroke volume from the seismocardiogram signal. *CMBES Proceedings* 33 (2010).
 - [86] Fadime Tokmak and Beren Semiz. 2023. Investigating the effect of body composition differences on seismocardiogram characteristics. In *2023 IEEE 36th International Symposium on Computer-Based Medical Systems (CBMS)*. IEEE, 323–328.
 - [87] Sterre Van Arum, Hüseyin Uğur Genç, Dennis Reidsma, and Armağan Karahanoglu. 2025. Selective Trust: Understanding Human-AI Partnerships in Personal Health Decision-Making Process. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–21.
 - [88] Paolo Verdecchia, Giuseppe Schillaci, Claudia Borgioni, Antonella Ciucci, Ivano Zampi, Roberto Gattobigio, Nicola Sacchi, and Carlo Porcellati. 1995. White coat hypertension and white coat effect similarities and differences. *American journal of hypertension* 8, 8 (1995), 790–798.
 - [89] Anna C. Voyatzoglou. 2022. An Introduction to the Comparison of Seismocardiography and Phonocardiography. <https://stars.library.ucf.edu/honorstheses/1217> Honors Undergraduate Theses, No. 1217, University of Central Florida.
 - [90] Anna Vybornova, Erietta Polychronopoulou, Arlene Wurzner-Ghajarzadeh, Sibylle Fallet, Josep Sola, and Gregoire Wuerzner. 2021. Blood pressure from the optical Aktiia Bracelet: a 1-month validation study using an extended ISO81060-2 protocol adapted for a cuffless wrist device. *Blood pressure monitoring* 26, 4 (2021), 305–311.
 - [91] Edward Jay Wang, Junyi Zhu, Mohit Jain, Tien-Jui Lee, Elliot Saba, Lama Nachman, and Shwetak N Patel. 2018. Seismo: Blood pressure monitoring using built-in smartphone accelerometer and camera. In *Proceedings of the 2018 CHI conference on human factors in computing Systems*. 1–9.
 - [92] Kapil Wattamwar, Z Jason Qian, Jenna Otter, Matthew J Leskowitz, Francesco F Caruana, Barbara Siedlecki, Jaclyn B Spitzer, and Anil K Lalwani. 2018. Association of cardiovascular comorbidities with hearing loss in the older old. *JAMA Otolaryngology-Head & Neck Surgery* 144, 7 (2018), 623–629.
 - [93] Peter West, Max Van Kleek, Richard Giordano, Mark J Weal, and Nigel Shadbolt. 2018. Common barriers to the use of patient-generated data across clinical settings. In *proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13.

- [94] Eric Widmaier, Hershel Raff, and Kevin T Strang. 2022. *Vander's human physiology*. McGraw-Hill US Higher Ed USE.
- [95] Richard A Wilson, Virinderjit S Bamrah, Joseph Lindsay Jr, Markus Schwaiger, and Joel Morganroth. 1993. Diagnostic accuracy of seismocardiography compared with electrocardiography for the anatomic and physiologic diagnosis of coronary artery disease during exercise testing. *The American journal of cardiology* 71, 7 (1993), 536–545.
- [96] Chenxi Yang, Sunli Tang, and Negar Tavassolian. 2016. Annotation of seismocardiogram using gyroscopic recordings. In *2016 IEEE Biomedical Circuits and Systems Conference (BioCAS)*. IEEE, 204–207.
- [97] Vahid Zakeri, Alireza Akhbardeh, Nasim Alamdari, Reza Fazel-Rezai, Mikko Paukkunen, and Kouhyar Tavakolian. 2016. Analyzing seismocardiogram cycles to identify the respiratory phases. *IEEE Transactions on Biomedical Engineering* 64, 8 (2016), 1786–1792.
- [98] John M Zanetti and Kouhyar Tavakolian. 2013. Seismocardiography: Past, present and future. In *2013 35th annual international conference of the IEEE engineering in medicine and biology society (EMBC)*. IEEE, 7004–7007.