Continuous Physiological Sensing Outside Controlled Environments

Doctoral Thesis

Author(s):
Meier, Manuel

Publication date:
2025

Permanent link:
https://doi.org/https://doi.org/10.3929/ethz-c-000784211

Diss. ETH No. 31250

Continuous Physiological Sensing Outside Controlled Environments

A thesis submitted to attain the degree of
DOCTOR OF SCIENCES
(Dr. sc. ETH Zurich)

presented by
MANUEL MEIER
Born on 08.12.1993

accepted on the recommendation of
Prof. Dr. Christian Holz, examiner
Prof. Dr. Christina Spengler, co-examiner
Prof. Dr. Jan Beutel, co-examiner

2025

ABSTRACT

Physiological monitoring plays a crucial role in precision medicine to enable the continuous assessment of an individual's health status through real-time physiological and behavioral data. Among the most widely monitored vital signs is heart rate, which is relevant in a broad range of contexts such as detecting cardiovascular conditions, tracking exercise intensity, or evaluating sleep quality. Other vital parameters, such as respiratory rate, also provide important insights and can serve as early indicators of adverse health events.

However, studies inside controlled laboratory or clinical settings fail to capture the complexities of human physiology that occur in free-living conditions. Wearable sensors offer a promising solution to real-world physiological data collection but face challenges such as fragmented and noisy data recording, the need for user interaction, sensor synchronization, and energy constraints that hinder long-term operation. Combined, these challenges limit the effectiveness of wearable technology for truly passive, accurate, and scalable health monitoring.

In this dissertation, we introduce an embedded wearable sensing platform and associated computational techniques for robustly monitoring physiological metrics. Our first contribution is the design of a multi-device, low-power sensing system that integrates multiple physiological and environmental sensors, which ensures long-duration, passive data collection without requiring user interventions. Our platform achieves sample-perfect synchronization across sensors and millisecond-level alignment between devices, overcoming one of the key barriers in multi-modal physiological monitoring.

To further enhance the reliability of physiological data collection, we propose a multi-modal offline synchronization method that enables precise alignment of signals across multiple body-worn sensors without requiring wireless communication or explicit user actions to support real-world scenarios that cannot guarantee continuous connectivity. We demonstrate the capabilities of the platform by introducing WildPPG, a large-scale real-world photoplethysmography (PPG) dataset collected from 16 participants over 13.5 hours in highly variable environmental and activity conditions. Unlike datasets captured inside controlled settings, WildPPG contains physiological responses to real-world activities such as physical exertion, temperature fluctuations, and varying ambient light levels. This makes our dataset a valuable resource for developing more robust wearable sensing algorithms.

Based on WildPPG, we propose a series of signal-processing techniques for robustly tracking physiological key metrics. We present a lightweight sensor fusion method that enhances heart rate estimation by combining multiple PPG signals from different body locations, mitigating the effects of motion artifacts and environmental disturbances. We also introduce a learning-based approach for generating PPG signals using multiple optical wavelengths, improving cardiac monitoring in multi-wavelength PPG sensing devices such as pulse oximetry applications. Finally, we propose a novel method for continuous respiratory rate monitoring using ultra-wideband (UWB) radar integrated into wearable devices, enabling accurate respiration tracking with reduced susceptibility to motion artifacts.

By integrating hardware innovations with computational methods, this dissertation establishes a foundation for continuous, multimodal, and unobtrusive physiological sensing in real-world environments. The wearable system and data processing techniques we propose can advance precision medicine via the passive health monitoring needed to collect representative data for the purpose of early disease detection, long-term physiological research, and, more generally, context-aware health monitoring solutions.

ZUSAMMENFASSUNG

Die kontinuierliche Erfassung physiologischer Parameter ist ein zentraler Bestandteil der personalisierten Medizin, da es eine fortlaufende Beurteilung des Gesundheitszustands in Zusammenhang mit Verhaltensdaten ermöglicht. Konventionelle Methoden, die auf Labor- oder klinische Umgebungen beschränkt sind, erfassen aber nicht alle komplexen physiologischen Prozesse, die im Alltag auftreten. Tragbare Sensorsysteme bieten zwar das Potenzial, physiologische Daten in alltäglichen Situationen zu erfassen, stehen allerdings vor fundamentalen Herausforderungen: Fragmentierte Aufzeichnungen, die Notwendigkeit von Nutzerinteraktionen, die Synchronisation mehrerer Sensoren sowie begrenzte Batterielaufzeiten erschweren deren Einsatz. Zusätzlich können externe Faktoren, etwa variierende Umweltbedingungen (z.B. Temperatur, Beleuchtung oder Luftfeuchtigkeit), fehlerhafte Platzierung der Sensoren durch die Nutzenden oder Bewegungsartefakte die Messqualität beeinträchtigen. Diese Einschränkungen reduzieren die Aussagekraft tragbarer Sensoren für ein passives, präzises und skalierbares Gesundheitsmonitoring.

In dieser Dissertation präsentieren wir eine innovative, tragbare Sensorplattform sowie robuste Methoden zur Datenanalyse. Das Ziel ist es, eine zuverlässige und energieeffiziente Erfassung physiologischer Signale zu ermöglichen, die ohne aktive Nutzerinteraktion funktioniert und eine qualitativ hochwertige Datengrundlage schafft.

Ein zentraler Beitrag dieser Arbeit ist die Entwicklung eines multimodalen, energieeffizienten Sensornetzwerks, das physiologische und umweltbezogene Daten über lange Zeiträume hinweg synchron erfassen kann, ohne dass Nutzende aktiv mit dem Gerät interagieren. Zusätzlich ermöglicht diese Plattform die exakte Synchronisation mehrerer Sensor-chips und eine Millisekunden genaue Abstimmung zwischen verschiedenen Geräten.

Zur weiteren Verbesserung der Datenqualität stellen wir eine neuartige Offline-Synchronisationsmethode für multimodale Sensorsysteme vor. Diese ermöglicht eine präzise zeitliche Ausrichtung physiologischer Signale, ohne dass Kommunikation zwischen den Geräten oder aktive Nutzerinteraktionen erforderlich sind. Dies ist sowohl in unkontrollierten Alltagsumgebungen, in denen Verbindungsprobleme auftreten können, als auch in maximal energieeffizienten Anwendungen, welche auf drahtlose Kommunikation verzichten müssen, ein zentraler Vorteil.

Ein weiterer, wesentlicher Beitrag dieser Dissertation ist der Datensatz WildPPG, welcher Photoplethysmographie (PPG) Daten aus Alltagssituationen beinhaltet. Im Gegensatz zu bestehenden PPG Datensätzen, die meist unter kontrollierten Bedingungen erfasst wurden,

enthält WildPPG zeitlich umfassende und ununterbrochene Messungen aus Alltagssituationen, die unter verschiedensten Umweltbedingungen aufgezeichnet wurden.

WildPPG beinhaltet physiologische Reaktionen auf körperliche Aktivität, Temperaturschwankungen und wechselnde Lichtverhältnisse. Damit bietet WildPPG eine einzigartige Grundlage für die Entwicklung robuster Algorithmen zur Verarbeitung von PPG-Signalen in tragbaren Sensorsystemen.

Auf Basis dieser Daten haben wir neue, signalverarbeitende Verfahren entwickelt, um die Qualität der physiologischen Messungen weiter zu verbessern. Dazu gehört ein Verfahren zur Fusion mehrerer PPG-Signale von unterschiedlichen Körperstellen, um den Einfluss von Bewegungsartefakten und Umwelteinflüssen zu reduzieren. Zudem entwickelten wir eine neue Methode, in der durch maschinelles Lernen synthetische PPG-Signale generiert werden, die aus Messungen mehrerer optischer Wellenlängen bestehen. Damit kann die Genauigkeit von kardiovaskulären Messungen, etwa in der Pulsoxymetrie, erhöht werden. Zusätzlich präsentieren wir ein innovatives Verfahren zur kontinuierlichen Atemfrequenzüberwachung. In diesem wird Ultra-Wideband (UWB) Radar in tragbare Geräte integriert und ermöglicht eine zuverlässige Erfassung der Atmung mit reduzierter Anfälligkeit für Bewegungsartefakte.

Durch die Kombination aus Innovationen im Bereich Hardware und Methoden der Datenanalyse leistet diese Dissertation einen Beitrag zur Weiterentwicklung des kontinuierlichen, multimodalen physiologischen Monitorings im Alltag. Die entwickelten tragbaren Sensorsysteme und Algorithmen bieten neue Möglichkeiten für die personalisierte Medizin und reicht von der passiven Überwachung von Gesundheitsparametern über die Möglichkeit der frühzeitigen Diagnostik bis hin zur Durchführung langfristiger physiologischer Studien.

ACKNOWLEDGMENTS

I would like to express my gratitude to Prof. Christian Holz. His belief in me and his consistent optimism regarding the outcome of my studies—at times surpassing my own—have been a great source of motivation. It has been an inspiring journey to be part of his research group from the very beginning, witnessing the team, ambitions, and research output grow over time. He invested considerable time and effort to guide me towards becoming an academic researcher without ever hindering my own curiosity.

I am also deeply thankful to my committee members. Prof. Christina Spengler, with whom I had the pleasure to collaborate on multiple occasions, generously shared her deep expertise in human physiology, which proved invaluable throughout my studies. I am grateful to Prof. Jan Beutel for kindly agreeing to join my committee as an external member, despite our limited prior interaction.

My sincere thanks extend to all my colleagues at SIPlab. I cherish the first months when we got started as a small group with Alice Mylaeus, Dr. Andreas Fender, Dr. Roman Kusche, and Dr. Hugo Romat. I am grateful to Paul Streli with whom I shared the intensity of our first submission deadlines, complete with the long nights that came with them. Thank you to Max Möbus and Björn Braun, both of whom I shared an office with and had many light-hearted and occasionally ridiculous conversations that kept morale up throughout this journey. I appreciated many interesting discussions with Rayan Armani, whom I had the pleasure of mentoring during her semester thesis project and who later joined the lab full-time. Sebastian Reidy, Berken Utku Demirel, and Jiaxi Jiang, with whom I had the pleasure to collaborate on publications. Also, thank you to Dominik Hollidt, Mohamed Kari, Dr. Christoph Gebhardt, Tiffany Luong, and Fiorella Meier.

I had the pleasure of supervising various student projects and collaborating with interns. I want to thank all of them for their patience and valuable work: Yves Kempter, Janik Baumer, Guillaume Thivolet, Rayan Armani, Christian Rubin, Emanuel Pytlik, Christoph Salzer, Cyrill Künzi, Rachel Schuchert, Oskar Pilkington, Silvan Roth, Sebastian Reidy, and Eric Weber.

I also want to thank all the participants of my studies for their time and patience, especially when the protocol required them to wear unusual-looking devices in public or complete exhausting tasks in the lab.

Finally, I want to thank my friends and family for their support. Especially Sophie, for her patience and understanding throughout all the happy and challenging phases of this journey.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACRONYMS

ADC	Analog-to-Digital Converter
AES	Advanced Encryption Standard
AMS	Acute Mountain Sickness
ASCII	American Standard Code for Information Interchange
BCG	Ballistocardiogram
BLE	Bluetooth Low Energy
BMAR	Barometric and Motion-based Alignment and Refinement for Offline Signal Synchronization (name of the method)
BP	Blood Pressure
brpm	breaths per minute
CIR	Complex Impulse Response
ECG	Electrocardiogram
EDA	Electrodermal Activity
FCC	Federal Communications Commission
FCN	Fully Convolutional Network
FFT	Fast Fourier transform
FIFO	First In, First Out
FMCW	Frequency Modulated Continuous Wave
FPC	Flexible Printed Circuit
FTSP	Flood Time Synchronization Protocol
GPS	Global Positioning System
HRV	Heart Rate Variability
HR	Heart Rate
I2C	Inter-Integrated Circuit
IBI	Inter-Beat Interval
ICA	Independent Component Analysis
IIR	Infinite Impulse Response
IMU	Inertial Measurement Unit
IoT	Internet of Things
MAE	Mean Absolute Error
MAPE	Mean Absolute Percentage Error
MLP	Multilayer Perceptron
ML	Machine Learning

OLS Ordinary Least Squares

PCB Printed Circuit Board

PPG Photoplethysmography

RBS Reference Broadcast Synchronization

RMSE Root Mean Square Error

RR Respiratory Rate

RTC Real-Time Clock

SCG Seismocardiogram

SNR Signal-to-Noise Ratio

SoC System-on-chip

SPI Serial Peripheral Interface

SpO_2 Peripheral Oxygen Saturation

TPSNs Timing-Sync Protocol for Sensor Networks

UWB Ultra-Wideband

WSN Wireless Sensor Network

1 INTRODUCTION

Physiological monitoring plays a critical role in precision medicine, which aims to tailor healthcare interventions to an individual's unique physiological and behavioral characteristics [100]. Achieving this goal requires continuous, high-resolution, multi-modal data to capture both physiological signals and behavioral context [120]. By leveraging continuous health monitoring, this approach enables a deeper understanding of an individual's health trajectory, facilitating early diagnosis, condition tracking, and personalized therapeutic strategies.

Traditionally, this type of monitoring is confined to controlled environments such as laboratories and medical facilities. However, these environments fail to capture dynamic, real-world variations in physiology and behavior, limiting their ability to inform personalized healthcare decisions.

To bridge this gap between controlled assessments and everyday health dynamics, wearable sensor technology has emerged as a promising solution, offering continuous physiological monitoring in everyday environments [11]. Advances in low-power electronics have enabled wearable devices to track physiological parameters throughout daily life, with consumer devices such as smartwatches increasingly integrating health-monitoring capabilities. However, significant technical limitations remain that prevent current wearable systems from fully realizing the goals of precision medicine [28].

1.1 CHALLENGES IN PHYSIOLOGICAL MONITORING USING WEARABLES

Continuously recorded high-quality data is one key challenge for existing physiological monitoring systems. To extend battery life and due to limitations in data transmission bandwidth, they rely on intermittent or short-duration recordings [109], which fail to capture dynamic fluctuations in vital signs and activity patterns. Without continuous data, the ability to assess long-term trends and contextualize health changes is significantly reduced.

Unobtrusive use, passive monitoring is another challenge. When devices require user interactions, such as charging, button presses, or supervision for data streaming this may affect the validity of the data capture. To not interfere with activities, passive zero-effort recording is important when capturing everyday health dynamics.

Sensor synchronization is another major challenge, particularly in multi-sensor and multi-device setups. Physiological signals from different modalities, such as motion, cardiac activity, and optical reflections, must be accurately time-aligned to enable reliable analysis [189]. Many existing wearable devices struggle with drift and latency, leading to misalignment that compromises signal fusion and subsequent data interpretation.

Long-duration standalone operation with multiple synchronized sensors that continuously log raw data in a compact wearable form factor remains challenging due to power consumption and storage constraints. Consumer-grade devices typically address this by pausing or downsampling sensing to conserve energy, but such strategies reduce data continuity. As a result, most current solutions require frequent recharging or data offloading, limiting their practicality for real-world applications.

Finally, *robust data processing* techniques are needed as real-world deployment introduces significant variability in environmental conditions, user behavior, and motion artifacts. Unlike controlled laboratory settings, everyday scenarios involve uncontrolled movements, varying lighting conditions, and diverse physical activities, all of which can degrade signal quality and reduce the reliability of physiological measurements [28]. Existing algorithms for vital sign estimation and activity recognition often fail to generalize beyond controlled environments. Many current methods assume ideal signal conditions, resulting in reduced accuracy when applied to real-world data. The challenge lies in developing robust algorithms that can extract meaningful physiological insights despite noise, motion artifacts, and environmental variability.

1.2 CONTRIBUTIONS AND DISSERTATION OUTLINE

This dissertation addresses the challenges listed above through the development and evaluation of an advanced embedded wearable platform as well as corresponding computational techniques for robust monitoring of accurate physiological metrics – in paticular Heart Rate (HR) and Respiratory Rate (RR).

Chapter 3 introduces our first contribution through the design and implementation of a multi-device, low-power sensing platform that integrates a multitude of sensors, storage, and connectivity for passive monitoring of user-related activity and physiology. Each sensor node continuously records raw sensor data in real-world conditions outside the lab without affecting the user, ensuring sample-perfect synchronization across all sensors on one node and millisecond-level synchronization across nodes. The unobtrusive devices have a 72-

hour battery life on a coin-cell battery without the need for any user interaction. They integrate Inertial Measurement Unit (IMU), Electrocardiogram (ECG), Photoplethysmography (PPG), and environmental sensors, providing synchronized multimodal data collection in real-world conditions.

In Chapter 4, we introduce a multi-modal offline synchronization method for distributed sensors that does not require active wireless communication between devices. Our synchronization method operates purely based on recorded signals and is thus suitable for offline processing. It does not require any specific user input, action, or behavior. The method operates on the data from devices worn by the same person that record barometric pressure and acceleration—inexpensive, low-power, and thus commonly included sensors in today's wearable devices. This enables accurate alignment of physiological signals across multiple body-worn sensors, allowing for comprehensive multi-site physiological analysis.

In Chapter 5, we present WildPPG, a real-world PPG dataset of long continuous recordings. The data was recorded from 16 participants during a 13-hour-long day trip from Zurich to Jungfraujoch, a mountain station at over 3500 m above sea level. Our dataset includes a Lead I ECG as ground truth reference for HR computation and was collected in real-world conditions outside of controlled environments. It includes various levels of physical activities as well as rest periods, multiple modes of transportation, and changing environmental conditions, most notably a wide range of ambient temperature and lighting conditions which are known to have an impact on PPG signal quality.

Chapter 6 presents a lightweight sensor fusion method for estimating HR that flexibly combines samples from multiple PPG sensors placed across the patient's body. In Chapter 7, we introduce a learning-based approach to generate a PPG signal based on samples at multiple optical wavelengths in a single body location as they are common in pulse oximetry. Both methods mitigate issues such as motion artifacts, changes in ambient conditions, and sensor misalignment, which often hinder downstream PPG processing methods in producing accurate measurements.

Chapter 8 presents a method to monitor the respiratory rate using a simple wearable device with less susceptibility to motion artifacts compared to other published methods. We leverage Ultra-Wideband (UWB) radios that are increasingly found in mobile and wearable devices and accurately monitor the wearer's respiration activity in everyday life. We thereby analyze the channel Complex Impulse Response (CIR) when measuring the reflection of a UWB signal transmitted into the wearer's chest.

Taken together, we establish a system for continuous, multi-modal, and unobtrusive physiological sensing in real-world environments, contributing to advancing passive health monitoring for precision medicine. The proposed platform and algorithms provide a foundation for personalized health tracking, early disease detection, and real-world physiological research, aligning with the broader goals of precision medicine to deliver individualized and context-aware healthcare solutions.

1.3 CONTRIBUTING PUBLICATIONS

As part of the research in this dissertation, we published the following peer-reviewed research papers:

Meier, M., & Holz, C. (2023). BMAR: barometric and motion-based alignment and refinement for offline signal synchronization across devices. Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies, 7(2).
Corresponds to Chapter 4.

Meier, M., Demirel B., & Holz, C. (2024). WildPPG: A Real-World PPG Dataset of Long Continuous Recordings. In 2024 Advances in Neural Information Processing Systems.
Corresponds to Chapter 5.

Meier, M., & Holz, C. (2024). Robust Heart Rate Detection via Multi-Site Photoplethysmography. In 2024 46th annual international conference of the IEEE engineering in medicine and biology society (EMBC). IEEE.
Corresponds to Chapter 6.

Meier, M., Demirel, B., & Holz, C. (2024). Tri-Spectral PPG: Robust Reflective Photoplethysmography by Fusing Multiple Wavelengths for Cardiac Monitoring. In 2024 International Conference on Body Sensor Networks. IEEE.
Corresponds to Chapter 7.

Reidy, S.*, **Meier, M.***, & Holz, C. (2024). Respiro: Continuous Respiratory Rate Monitoring During Motion via Wearable Ultra-Wideband Radar. Advances in Neural Information Processing Systems. In 2024 IEEE EMBS International Conference on Biomedical & Health Informatics (BHI). IEEE. **both authors contributed equally*
Corresponds to Chapter 8.

1.4 ADDITIONAL PUBLICATIONS

Beyond this, I was involved in the following work that was published in peer-reviewed venues throughout the duration of my PhD studies:

Meier, M., Streli, P., Fender, A., & Holz, C. (2021, March). TapID: Rapid touch interaction in virtual reality using wearable sensing. In 2021 IEEE Virtual Reality and 3D User Interfaces (VR) (pp. 519-528). IEEE.

Romat, H., Fender, A., **Meier, M.**, & Holz, C. (2021, March). Flashpen: A high-fidelity and high-precision multi-surface pen for virtual reality. In 2021 IEEE Virtual Reality and 3D User Interfaces (VR) (pp. 306-315). IEEE.

Streli, P., Jiang, J., Fender, A., **Meier, M.**, Romat, H., & Holz, C. (2022, April). TapType: Ten-finger text entry on everyday surfaces via Bayesian inference. In Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems.

Jiang, J., Streli, P., **Meier, M.**, & Holz, C. (2024, September). Egoposer: Robust real-time egocentric pose estimation from sparse and intermittent observations everywhere. In European Conference on Computer Vision (pp. 277-294). Cham: Springer Nature Switzerland.

Mylaeus, A., Vogt, S., Demirel, B., Gort, M., Meboldt, M., **Meier, M.**, & Holz, C. (2024, September). MiBOT: A head-worn robot that modulates cardiovascular responses through human-like soft massage. In 2024 IEEE International Conference on Robotics and Automation (ICRA) (pp. 13023-13029). IEEE.

2 RELATED WORK

The wearable devices and methods presented in this dissertation build on a wide range of related work that explored applications and challenges in this field. We outline the potential of physiological sensing outside controlled settings, as demonstrated in related studies, and the challenges that arise. We examine existing devices, their architectures, and their approach to design trade-offs relating to device size, energy management, and software architecture. A key aspect relevant to the interoperability of multiple devices and thus multi-modal sensing across body locations is device synchronization. In this context, we explore setups with and without wireless communication capabilities during runtime. Finally, our work focuses on the sensing of PPG and respiration, methods that we examine in more detail.

2.1 WEARABLE DEVICES FOR PHYSIOLOGICAL SENSING

Physiological sensing is among the main use cases for wearable devices [60]. There are many possible applications: They can enhance preventive care by detecting and preventing falls in elderly populations [181], monitor mental status [45], or detect unhealthy body poses and alert wearers accordingly [71]. Wearables can also be used to monitor recovering patients [144] and gain popularity in the treatment of chronically ill patients [124].

At the intersection of health monitoring and activity recognition, Strain et al. perform mortality prediction based on data collected from a wrist-worn wearable [176]. Similarly, Bari et al. combined synchronized data from separate devices recording audio, physiological, and inertial signals to detect stressful conversations [15].

More generally, wearable devices make it easier to collect multimodal long-term data and therefore enable big data methods in the space of physiological sensing [196].

However, these approaches share various challenges, including the need for high data quality in uncontrolled environments and interoperability of devices [34]. In the following sections, we discuss aspects relating to these two challenges.

2.2 DEVICE SYNCHRONIZATION

In measurement setups that rely on multiple wearable devices, time synchronization between devices is essential to match simultaneously acquired information from the devices.

2.2.1 *Wireless Synchronization*

Over recent years, Wireless Sensor Networks (WSNs) have attracted more and more interest. Time synchronization performed online or offline between the nodes of a WSN is crucial for the deployment of the Internet of Things (IoT) [9]. Currently, many time-synchronization methods for WSNs use Bluetooth Low Energy (BLE) [9, 149]. Especially in the fields of ubiquitous computing, sports science, and health applications, BLE is most commonly used since it is integrated into modern smartphones [149].

With BLE, timestamping errors of less than 50 ns are possible as demonstrated by Asgarian and Najafi with *BlueSync* [9]. *BlueSync* outperforms previous work by several magnitudes [22, 78, 79, 149, 168, 171]. While Rheinländer and Wehn achieve an accuracy of 0.9 µs [149], Sridhar et al. and Bideaux et al. achieve accuracies of 10 µs [171] and 14.9 µs [22], respectively. Somaratne et al. and Ghoshdastider et al. achieve accuracies of 19 ms [168], 37.8 ms [78], and 45.1 ms [79], respectively. While the works mentioned above primarily focus on minimizing synchronization error rather than power consumption, Makara, Tsybul'nyk, and Kurnyts'kyi aim to achieve both, reaching a synchronization error of 10 ms error at a power consumption of 133 µW when using two devices [121]. However, BLE is often limited by a lack of flexibility in network topology [103] making the estimation of the energy consumption in larger systems difficult.

There are also alternatives to BLE. This includes the work of Li, Braun, and Dimitrova who use Global Positioning System (GPS) and achieve synchronization with an accuracy well below 1 ms [113] and Seijo et al. who use Wifi to achieve sub-nanosecond time transfer accuracy [161]. However, these communication technologies share the downside of higher power consumption compared to BLE and thus have a negative impact on battery life.

2.2.2 *Data-based Synchronization*

Multiple works have introduced data-based synchronization for distributed sensor systems using acceleration data from IMUs. Most of these works focus either on synchronization between IMUs and motion-capturing systems, IMUs and video data from cameras, or between only IMUs. The latter is typically done by subjecting all IMUs to a sudden acceleration such as shaking multiple IMUs that are held together [74, 131, 150, 164, 195]. With this method, Shabani et al. obtained a minimum mean lag time bias of 25.56 ms [164]. Recently, Spilz and Munz introduced a novel data-based synchronization method for synchronizing multiple devices by subjecting them synchronously to a magnetic field impulse generated by an array of inductors causing an event that is detectable by the devices' IMUs [170] and allowing for synchronization with a maximum offset of below 2.6 ms. To synchronize IMUs and motion-capturing systems or video data, sudden movements such as a heelstroke are commonly used. The event is then visible both in the video and the acceleration data and recordings are typically manually aligned [14, 29, 98, 139, 201]. For video-to-IMU synchronization without explicit user action, Zhang et al. proposed SyncWISE, an algorithm that achieves a mean error of less than half a second [202]. When aligning recordings based on data and without distinct events, a measure for signal synchrony is required. Gashi, Di Lascio, and Santini evaluated several such measures in their work when they explored the synchrony of signals of electrodermal activity between presenters and audience members to determine engagement [76].

2.3 REFLECTIVE PHOTOPLETHYSMOGRAPHY IN WEARABLES

Reflective PPG is a common sensing modality in wearable devices to assess a wide range of physiological parameters, such as HR, Heart Rate Variability (HRV) [136], RR [51], and Peripheral Oxygen Saturation (SpO_2) [12]. Therefore, reflective PPG sensors are commonly built into devices such as fitness trackers (e.g., Fitbit, Garmin Fenix) or everyday smartwatches (e.g., Apple Watch, Pixel Watch, Galaxy Watch) to monitor these parameters passively during wear.

2.3.1 *Signal Quality*

Robust PPG sensing is a challenging task, as PPG signals are subject to artifacts such as motion or optical interference [117, 194]. Reliable PPG analysis requires good signal quality; robust signals facilitate individual pulse segmentation, such as for estimating a person's HR. They also allow analyzing characteristics of a blood volume pulse through its morphology, such as to calculate temporal features for pulse transit time [112]. Raw PPG signals are thus commonly improved with established band-pass filters [206] to aid feature extraction.

Depending on the magnitude of artifacts, however, conventional filtering can be insufficient for restoring the underlying blood volume pulses, as the signal's frequency spectrum typically overlaps with that of the artifacts. Therefore, researchers have integrated accelerometers to detect motion [178] and remove artifacts via noise cancellation [8, 165, 198].

2.3.2 *Multi-Channel Photoplethysmography and Sensor Fusion*

In response to signal quality issues, sensor designs have leveraged multiple light sources [1] or photodiodes [185] to harden PPG signal acquisition. Multi-LED designs have become commonplace in today's commercial wearables (e.g., Garmin [75]), although their processing is proprietary.

The combination of multi-channel PPG signals can be considered sensor fusion which has frequently been used for wearable physiological sensing [125]. Raiano et al. removed motion artifacts from breathing signals using multiple piezoresistive textile sensors, fusing signals via Independent Component Analysis (ICA) [143]. Combining PPG and Ballistocardiogram (BCG) has also been shown to better estimate R-R intervals [6] and HR [35]. For HR estimation, past work has also investigated sensor fusion. Bieri et al. jointly analyzed accelerometer and PPG signals for robust and uncertainty-aware heart rate estimation [23]. Wartzek et al. separately detected heartbeats in PPG and capacitive ECG signals acquired in multiple locations on a mattress and then combined them for better estimation [193]. When working with PPG signals only, using the average of multiple PPG traces is a common practice in prior work(e.g., [178]). Naturally, this reduces the impact of single-channel artifacts but the output signal is still affected – particularly when the artifacts' amplitudes are much greater than the actual signal. In an effort to combine PPG traces while taking their signal quality into account, Warren et al. switch between channels every two seconds to select the most promising for HR detection [192]. However, this method may lose information from channels that are not currently assumed to have the best signal quality which weighs especially heavily when the quality estimation is less-than-optimal. Lee et al. combined PPG signals at multiple wavelengths using ICA to reduce motion artifacts [105], whereas Holz et al. combined multiple PPG sensors across the head to derive HR [86].

2.4 MONITORING RESPIRATORY RATE WITH WEARABLES

Certain measurement techniques do not transfer well from controlled hospital or laboratory settings to everyday settings. For example, Respiratory rate measurements are routinely collected in controlled settings using capnometry and spirometry which are not suitable for long-term measurements in everyday life because they require air masks [3, 115]. Related research on respiratory rate monitoring in uncontrolled settings can be broadly classified into two categories: systems utilizing contact-based sensors that are directly attached to the human body and systems that enable remote measurement without direct physical contact.

2.4.1 *Non-contact-based RR monitoring*

Non-contact-based Methods for RR monitoring have been published using a variety of sensing modalities. Wang et al. utilized ultrasound to sense exhaled airflow in sleeping subjects, achieving a median error of less than 0.3 breaths per minute (brpm) [190]. Focussing on recovering the RR from chest movements, Bernacchia et al. used the Microsoft Kinect system, which combines a depth sensor, video camera, and microphones, reporting a standard deviation of the residual RR of 9.7% [20]. Radar technology has also been employed to observe chest movements induced by respiration. Researchers have developed various system architectures using Frequency Modulated Continuous Wave (FMCW) radar at radio and laser frequencies, where respiration-induced chest movements alter the phase and frequency of the carrier signal [3].

2.4.2 *Contact-based RR monitoring*

Contact-based methods for RR monitoring encompass an even broader range of sensing modalities compared to non-contact methods. Kumar et al. utilized wearable microphones to estimate RR [101]. Their setup, tested on subjects before, during, and after exercise, achieved a mean squared error in the RR of 0.2 brpm. Basra et al. developed a system that recovered RR using a temperature sensor on a nose clip, exploiting the temperature difference between inhaled and exhaled air [16]. Similarly, Guder et al. monitored RR by sensing humidity differences in an air mask [82].

Further exploring physiological signal variations, Heydari et al. derived RR from changes in bio-impedance measured on electrodes placed on the shoulder of participants, reporting an error of $\leqslant$ 1 brpm across different breathing patterns (slow, fast, deep, hold, and normal) [84]. Another approach leverages respiratory sinus arrhythmia, where the heart rate increases during inhalation and decreases during exhalation. This phenomenon enables RR measurement through heart rate monitoring techniques such as ECG, PPG, BCG, and Seismocardiogram (SCG) [115]. Charlton et al. evaluated 270 algorithms using ECG and PPG signals, finding that the best-performing algorithm had a 95% limit of agreement of -4.7 to 4.7 brpm [39].

Moving to inertial sensor-based methods, accelerometers, gyroscopes, and respiration belts can detect periodic thorax volume changes during breathing [115]. Rahman et al. used a single accelerometer mounted on a belt on the thorax to estimate RR, achieving successful recovery in 97% of attempts across five RRs on stationary subjects. However, this approach did not address motion artifacts, a known limitation of accelerometer-based systems [38, 142]. To mitigate motion artifacts, Bates et al. reconstructed RR from the angular velocity of the current rotation angle measured relative to the mean direction of gravity. They implemented a movement detection method to pause their system during motion, only computing RR when the subject was at rest [17].

2.4.3 *RR monitoring using Ultra-Wideband Radar*

UWB radar has garnered interest over the past two decades since the Federal Communications Commission (FCC) opened the 3.1 to 10.6 GHz spectrum for medical imaging [68]. There is substantial research on non-contact RR recovery using UWB radar [32, 89, 188], including through debris [108] or walls [42]. This remote monitoring detects chest wall motion using UWB beacons in a person's environment. However, these methods are prone to errors due to unrelated body motion and are unsuitable for tracking the RR during activities and in changing environments. UWB radar can also probe the body itself, detecting varying dielectric properties of tissues which affect UWB pulses differently [173]. Specifically, this applies to the lungs, which exhibit different dielectric properties when inflated versus deflated: More signal is reflected at the muscle-lung interface when the lung is inflated compared to its deflated state [134]. This is supported by Cavagnaro et al. who studied the propagation of UWB pulses into human tissue and proposed that effective RR recovery can be achieved by capturing the air-skin reflection signal with an antenna positioned 1 meter away from the body. They suggested that on-body antennas could potentially recover the RR from a small signal reflected by the posterior lung wall [37]. However, on-body UWB radar for RR monitoring has only been studied using laboratory equipment [137, 138], using two points of contact [48, 49], or with custom UWB antennas integrated into a seat [158], all of which were evaluated exclusively on stationary users.

2.4.3.1 *Non-contact-based RR monitoring with UWB radar*

Wang et al. conducted a comparative analysis between FMCW radar and UWB radar for non-contact vital sign extraction in varying conditions regarding distance and orientation of the subject as well as obstructing obstacles. Their findings indicated that UWB radar offers higher accuracy and a better Signal-to-Noise Ratio (SNR) than FMCW radar, showcasing the potential of UWB radar for RR monitoring [188]. Immoreev et al. demonstrated the ability to recover RR from tiny chest movements using a UWB radar mounted 2 meters above a hospital bed [89]. In more complex environments, such as disaster sites, Li et al. utilized the curvelet transform to remove antenna coupling and background clutter, followed by singular value decomposition to denoise the respiratory signal before extracting RR with a fast Fourier transform [108]. Duan et al. advanced this approach with an algorithm based on variational mode decomposition, capable of recovering RR in open spaces and through walls, achieving a correct detection rate of 95% [59]. Regev et al. further improved accuracy by fusing UWB radar signals with data from an RGB camera, resulting in a maximum error of 0.5 brpm when testing on 14 sitting participants [147]. Li et al. explored another method by using two radar modules (DWM1000, Qorvo) to recover RR from sitting subjects at different positions in a room. They investigated two approaches for computing the respiration signal from reflections at different spatial distances and concluded that a linear combination of reflections from multiple spatial distances outperformed selecting reflections solely from the most probable spatial distance [110]. This linear combination approach leveraged the advantages of capturing comprehensive respiratory patterns, thereby enhancing the robustness and accuracy of RR monitoring in varied environments.

Non-contact systems are typically constrained to stationary settings indoors. Our method, by contrast, is designed for real-world scenarios where individuals may be engaged in physical activities and not constrained to a previously equipped environment.

2.4.3.2 *Contact-based RR monitoring with UWB radar*

Pitella et al. explored single-point-of-contact antennas placed on the chest to capture individual UWB pulses reflected by the human body, concluding that the reflected signal changes with the respiration phase [137]. Building on this, they investigated various signal processing techniques for extracting RR from in-body reflected pulses, finding that mean subtraction offered favorable accuracy and fast processing [138]. Notably, their UWB radar setup could detect RR even during arm movement, unlike a piezoelectric belt which picked up the arm movement rate. However, their setup involved a vector network analyzer, making it non-wearable. Schires et al. optimized UWB antennas to maximize power transmission into the human body by embedding the antennas in the backrest of a car seat, recovering RR from phase variations in the reflected pulse [158]. Culjak et al. introduced a two-point-of-contact wearable RR monitoring system, positioning one radar module (DWM1000, Qorvo) on the front and another on the back of the thorax [49], using the CIR as input similar to Li et al. [110]. They employed mean subtraction to eliminate static components and applied a bandpass filter to the band of possible respiration frequencies in slow time. For each slow-time CIR, they identified the first four extrema in fast time and computed the difference in slow time between those fast time indexes, finally using an Fast Fourier transform (FFT) to identify the dominant frequency. They reported a relative error of 10% of the RR in the optimal configuration [49]. They later achieved a root mean squared error below

0.2 brpm using two data fusion algorithms—naive Bayes inference and Kalman filtering—to fuse UWB signals with accelerometer data [48].

3 DESIGN OF A SENSING PLATFORM

The acquisition of long-term recordings of synchronized, raw vital data from multiple on-body locations is crucial in work that investigates phenomena for which multiple sensing modalities are relevant. Fields of application range from affective computing, and stress monitoring to comprehensive sleep monitoring and detailed motion tracking. Commonly, this is only possible in controlled environments with extensive equipment. Taking this principle to low-power wearable devices, we introduce a sensing platform consisting of multiple sensor nodes. Each node is a multi-sensor wearable health device that can be attached to any body location. The devices autonomously record synchronous high-quality sensor data without the need for user interaction and with a battery life of multiple days despite a small form factor.

3.1 OVERVIEW

Live monitoring of vital signals as well as the analysis of long-term recordings is crucial to modern healthcare and research in fields such as affective computing [10], stress monitoring [154], sleep monitoring [102], or motion tracking [169]. Assessing such complex measures often requires extensive data recordings of multiple modalities per person logged over several hours or days. To support modern data analysis based on Machine Learning (ML) downstream optimally, the data should be continuous, of high resolution and synchronized across modalities and recording devices [57]. While the acquisition of such data may be technically easy in controlled environments, it is a natural goal to transfer these analyses into everyday situations outside of controlled environments to ultimately make them accessible to everyone. Therefore, we need scalable solutions for health monitoring and long-term data acquisition which is relevant to assess a person's health [64] where the general mantra is: more data is better (big data) [67, 140]. In the commercial space, devices such as the Apple Watch [7] monitor comparable sets of modalities and further amplify the relevance of this trend. Commercial manufacturers do not limit their line-up to single devices but build whole ecosystems, integrating wearables, phones, laptops, and other specialized monitoring devices such as bike computers. The main focus of these devices is the monitoring of vital signals and directly derived measures such as HR, HRV, RR, Blood Pressure (BP), or SpO_2 while optimizing for long battery life. The data acquisition usually includes compressed sensing

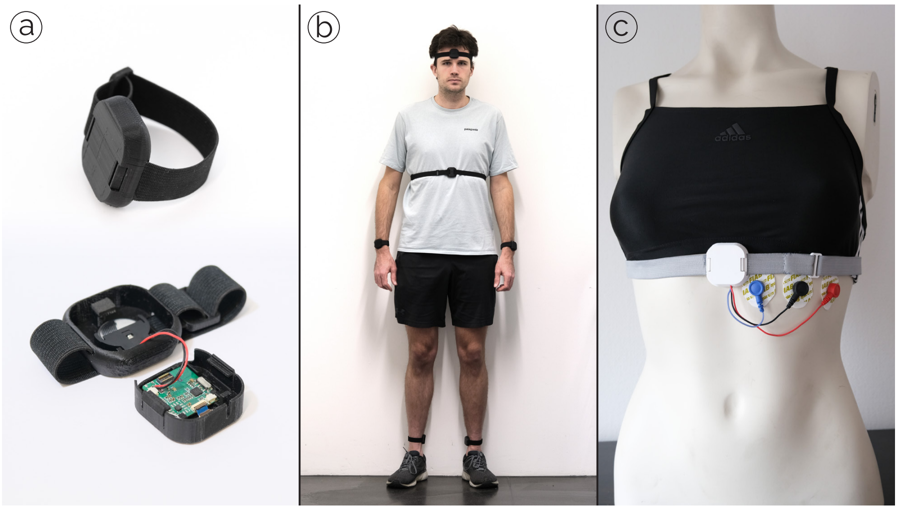

Figure 1: Our sensing platform enables the use of multiple sensing devices simultaneously. The custom electronics of each device are integrated into a 3D-printed enclosure along with the battery (a). Thanks to adjustable straps of various lengths, the devices can be placed in any body location. The proposed on-body placements include the wrists, ankles, forehead, and sternum (b). The device can be connected to gel-electrodes to record reliable Lead I ECG (c).

and on-board processing to reduce power consumption [13] which results in black box-type data available from most commercial devices and research devices which interrupt the recording regularly to save energy [197]. However, for research on complex-to-measure phenomena including affective computing, stress, comprehensive sleep monitoring, and detailed activity tracking, high-quality, continuous, and synchronized data from multiple body locations is necessary [44, 182, 186, 191]. However, this comes with the challenge of reduced battery life and throughput problems for low-power communications. We therefore present a device for the acquisition of continuous, synchronized high-quality data for research and development purposes including PPG, ECG, IMU, Electrodermal Activity (EDA), Environmental data, and Breathing. Since data streaming is energy costly and a point of failure, especially in outdoor settings involving movement, our approach is to both stream the data as far as required by the application (battery-life, bandwidth trade-off) and store it at full resolution and sampling rate on-board to be later fed into a database for further analysis. This enables both live monitoring applications and modern data analysis such as ML which benefits from large amounts of raw data.

3.2 SYSTEM ARCHITECTURE

The proposed system consists of three main components:

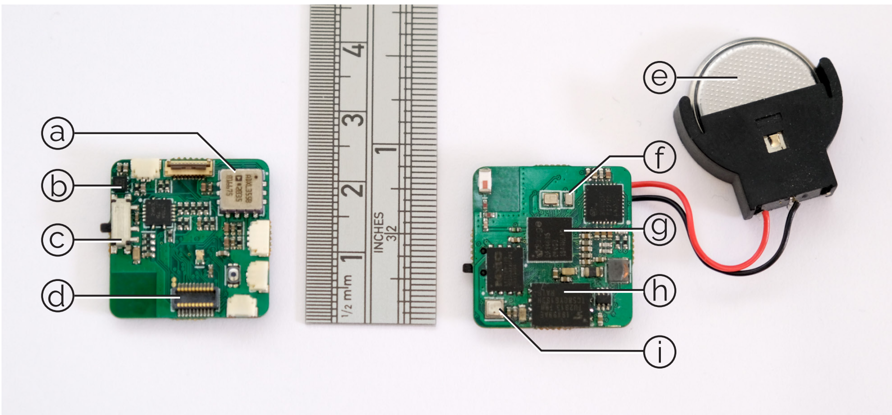

Figure 2: Top and bottom sides of the electronics of our custom wearable device, which measures 23 mm × 22 mm. Components: (a) ADXL355, high-resolution accelerometer, (b) LIS2DH, low-cost accelerometer, (c) power switch, (d) DA14695 extension connector, used for synchronization ground truth, (e) CR2032 battery, (f) Real-Time Clock (RTC) quartz crystal, (g) DA14695 System-on-a-Chip, (h) TH58CYG3S0HRAIJ 8 Gb flash memory, (i) BME280 barometric pressure sensor.

- Multiple distributed wearable sensing devices worn by one or more people simultaneously
- A BLE connected smartphone for live monitoring, time synchronization, and data streaming (gateway)
- A database backend with a full-resolution data option enabled by wired download when a device is connected to a PC

3.2.1 *Connectivity*

For the best data quality combined with full live monitoring capabilities and low power consumption, we combine both BLE live streaming and wired data download after the recording. While theoretically, BLE achieves a throughput of well over 500 kbps this can be severely reduced in both indoor and outdoor settings, especially with no line of sight, to less than 1 % of the theoretical maximum value [95]. Therefore it is not possible to reliably stream all the sensor data from our nodes at full raw resolution. Beyond this limitation, we reduce the live streaming to what is necessary for the specific application as it has a significant impact on power consumption and, hence, battery life. The live-streamed data is both down-sampled in frequency to 30 Hz and the resolution reduced to at most 16 bits per sensor. For the ECG signal, the downsampling is performed such that the maximum amplitude of the R-Peak is preserved: In case a dominant local maximum is present in the original signal, the downsampling process propagates it to the lower sampling-rate signal instead of averaging over mul-

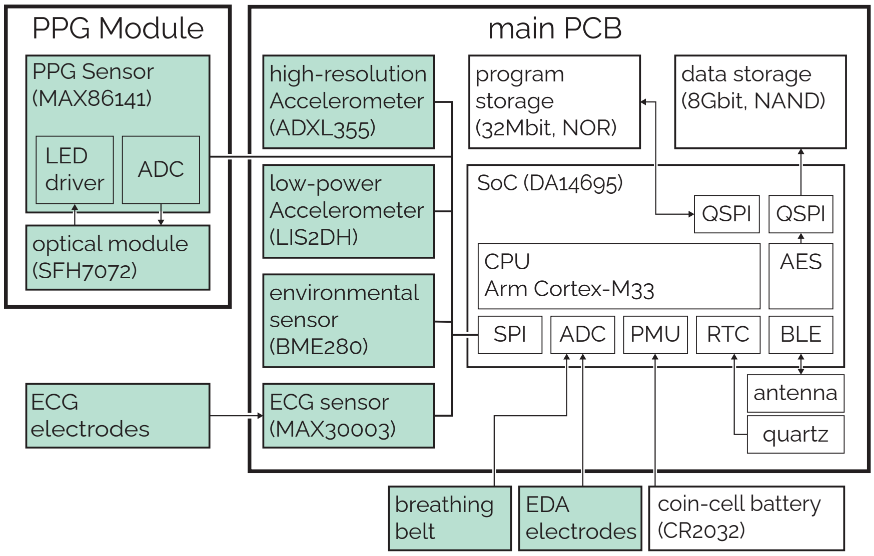

Figure 3: Each sensor node consists of a main Printed Circuit Board (PCB) which contains memory, processing unit, and sensing chips. The PPG sensing module is on a separate board to enable the placement closer to the skin. It connects through an Flexible Printed Circuit (FPC) cable to the main board. Optionally, ECG electrodes, a breathing belt, and EDA electrodes can be connected to the main board.

tiple samples. The BLE streaming contains 3 samples per sensor per transmission to reduce the number of transmissions per sensor and to simplify the interplay with the First In, First Out (FIFO) burst reads which are performed for most sensors. At the same time, the sensor data is collected in blocks equal to the flash page size, encrypted using Advanced Encryption Standard (AES)-256, and stored in the flash memory for later download. At the time of download over USB, the data is streamed page by page to a host computer, using check-sums to detect transmission errors. The data is then decoded and stored in the database for later analysis.

3.2.2 *Node platform*

The node platform is built around a low-power System-on-chip (SoC) (DA14695, Renesas) which runs FreeRTOS and handles BLE communication, power management, analog-digital conversion, communication with digital sensors and flash memory, USB communication, and timekeeping. Using the Serial Peripheral Interface (SPI) interface, data from digital sensors is read based on timer interrupts on a round-robin basis and both streamed through BLE and stored in an 8 Gbit NAND flash (TH58CYG3S0HRAIJ, Kioxia). This is in contrast to typical approaches where sensors trigger read operations by the SoC using interrupts and it dramatically reduces the number of interrupts in the system, simplifying hardware connections and reducing the number of context switches by the SoC. However, this approach requires guaranteed time-based readouts before sensor FIFO buffers overflow to avoid data loss. Therefore the temporal spacing of the readout procedure has to be shorter than the fastest possible time any sensor takes to fill its FIFO buffer. On the SoC, the data is buffered in RAM and written in blocks equal to the page size of the flash memory chip. This prevents the need for page modification operations in the flash during the runtime of the node. This is important as they require power-costly erase and multiple write operations.

The RTC of the SoC relies on a 32.768 kHz external quartz crystal (ECS-.327-7-16-C-TR, ECS Inc. International [91], Figure 2f) with a tolerance of ± 20 ppm. The SoC receives timestamps for synchronization through BLE for which the standardized current time service standard of BLE is implemented. For higher accuracy, an additional BLE time service is available which supports the transmission of one more byte to increase the sub-second accuracy from 1/256 seconds to 1/65536 seconds. The first timestamp received through BLE is used to set the RTC of the SoC. Subsequent received timestamps are used to keep track of clock skew (See Section 4.2.2). However, these timestamps are not used to set the RTC register. Instead, the clock offset is computed and logged both in the flash memory alongside the sensor data and in a globally available variable. This is done to avoid possible jitter effects caused by modifying the RTC state while still preserving the ability to resample sensor data to perfectly match sampling rates based on real-world time.

3.2.2.1 *Power*

The node is designed to be run on a CR2032 coin cell battery. While the device also supports rechargeable lithium batteries, the use of CR2032 is an additional safety feature as it has a naturally limited maximum power output and is less prone to overheating. CR2032 coin cell batteries typically have a maximum current rating of 1 mA (corresponding to 3 mW power draw), and a current draw above this mark reduces the capacity of the battery. Depending on the number of activated sensors, our device surpasses this mark which means additional power consumption has a disproportionally amplified negative effect on battery life. This highlights the need for low-power design which was a main focus both when designing hardware and software to achieve maximum battery life. Energy consumption is dominated by the power consumption of the sensors and data transmission. To maximize runtime, sampling rates can be lowered and live streaming minimized.

3.2.2.2 *Extension ports*

A miniaturized 20-pin extension port within the node enables programming, debugging, and data downloading as well as connecting additional sensors both digital and analog. Moreover, the device can receive synchronization signals through the port in an application where exact time synchronization beyond the accuracy of BLE-based synchronization between devices is needed. The accuracy of wire-based synchronization is determined by the clock cycle of the RTC-clock of 32.768 kHz which results in a synchronization tolerance of 31 µs. This capability is important to the key goal of the device to collect highly time-synchronized data. Additionally, there is a 9-pin FPC connector aimed at providing an optimal interface for PPG sensor modules which for physical alignment reasons are usually kept on a separate shuttle board. Both extension connectors provide 1.8V, direct battery voltage, and 3.3V which is suitable to run most digital and analog sensors.

3.2.3 *Sensors*

We built the node to be compatible with most modern analog and digital sensors. Digital sensors are read through SPI or Inter-Integrated Circuit (I2C) bus communication and operate in FIFO mode unless the sampling rate is low enough so that single sample read-outs are possible without excessive interrupt regularity (e.g., for environmental sensors). In our application, we included the following sensors:

- ECG (MAX30003, Analog Devices)
- 3-channel PPG (analog front-end: MAX86141, Analog Devices; optical module: SFH7072, ams-OSRAM)
- High-resolution accelerometer (ADXL355, Analog Devices)
- Standard-resolution Accelerometer (LIS2DH, STMicroelectronics)
- Electrodermal activity (analog circuit)
- Respiratory Belt (analog circuit)
- Temperature, Air pressure, Humidity (BME280, Bosch Sensortec)
- Extension possibility 1: Gyroscope sensor adapter (IIM-42652, InvenSense)
- Extension possibility 2: Two more accelerometers (LIS2DH, STMicroelectronics) on a flex PCB.

3.2.3.1 *Sampling Synchronization*

To enable sample-by-sample time synchronization, the SoC provides 128 Hz and 32768 Hz synchronization clocks which are both derived from the RTC quartz and therefore synchronized with the internal system time of the SoC. These clocks are used on all sensor chips that support external clocks, resulting in data that is perfectly synchronized with the system time of the device. Compared to a perfect time reference, this reduces the sample rate tolerance from approximately ±2% of the internal sensor clock [119] by a factor of 100 to the tolerance of the used quartz (ECS-.327-7-16-C-TR, ECS Inc., ±20ppm tolerance [145]).

Since Analog-to-Digital Converter (ADC) readouts are triggered based on the same clock, both digital and analog sensor readings are consistently synchronous. Digital sensors with no support for external clocks (e.g., the LIS2DH accelerometer) sample on their own internal clock. Deviations in sampling rate due to clock tolerances can be corrected in post-processing based on the RTC timestamps of the FIFO readouts of the sensor. This setup guarantees synchronization for the different sensing modalities acquired by a single device.

3.2.4 *Data Management*

The path of the sensor data is as follows: Upon its creation, it is cached in the FIFO buffer of the sensor from where it is read by the SoC of the device regularly. If streaming is activated for the given sensor, the relevant samples are queued for BLE streaming at this stage. Typically, this is only done for every nth sample resulting in a downsampling effect to reduce power consumption and avoid the bandwidth limitations of the radio transmission. Independently, all sensor data is stored in the volatile RAM of the SoC until enough data is accumulated to fill a complete page in the non-volatile NAND-flash memory chip (this is typically in the range of several 1000 Bytes). As soon as this is the case, the data is encrypted and stored in the flash and never modified thereafter. When the device is connected to a host computer, an image of the contents of the flash memory chip is transmitted in a page-by-page manner until a full copy of the data is on the host machine. From there it can be decrypted and saved in any secured storage medium such as a database.

3.2.4.1 *Data Format*

Upon reading data from the sensor FIFO buffer, it is reformatted into a package as shown in Figure 4 which contains a header and a body containing the data from a single FIFO readout. The header contains 4 bytes by default:

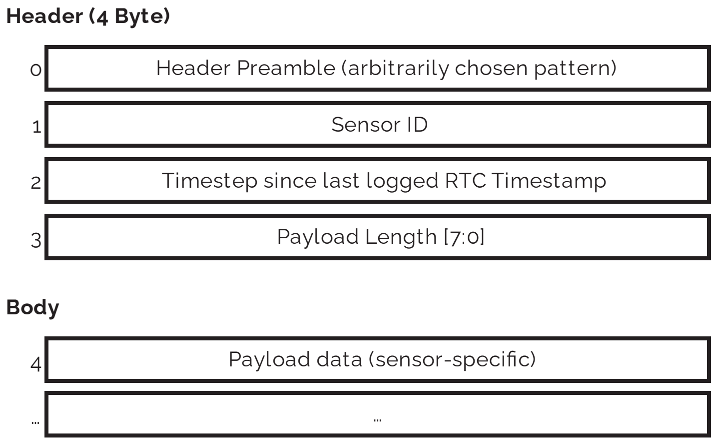

Figure 4: The package structure used for sensing and logging data on the node platform.

- **Byte 0** is a arbitrarily chosen bit-pattern $p \in [1, 254]$ marking the start of every package. Wherever p occurs in the rest of the package, it is escaped by another arbitrarily chosen escape sequence $s \in [1, 254], s \neq p$. Regular occurrences of the escape sequence s are always escaped using itself s. This approach guarantees that every package uniquely starts with an unescaped occurrence of p. For maximum memory efficiency, p and s should be chosen such that they do not naturally occur in the sensor data at a high rate.

- **Byte 1** is a sensor ID $\mathtt{id} \in [0, 254] \setminus p, s$. Each sensor ID defines the format of the payload data (see below). While p and s could be used as sensor IDs, it should be avoided to cause memory inefficiencies due to excessive escaping. $\mathtt{id} = 255$ is reserved for applications where more than 253 sensor IDs are needed. In this scenario, the first byte of the payload serves to extend the sensor ID. By recursively using this method, any number of sensor IDs can be accommodated. For maximum memory efficiency, low sensor IDs should be assigned to regularly occurring data patterns.

- **Byte 2** is a timestamp offset. See the explanation below.

- **Byte 3** is the payload length in bytes

The body of the package contains bytes representing the samples of the sensor concatenated to remove bits with no information. E.g., 4 samples of a sensor with 13-bit resolution result in a payload of 7 Bytes where only the last 4 bits of the last byte contain no information. The information about the resolution of the sensor data and whether it is signed or unsigned data is associated with the sensor ID. In scenarios where a sensor can be configured to sample at various resolutions, each possible sensor resolution is assigned a separate sensor ID.

The RTC timestamp is also logged on a regular basis using the same package format. All other packages only contain a timestamp offset to the last logged RTC timestamp which reduces memory usage. Logging of configuration details and runtime errors can also be logged using the same system by assigning each category a sensor ID and in case of unstructured data, the payload can be filled with American Standard Code for Information Interchange (ASCII) encoded strings.

3.2.5 *Scalability*

The system is designed to be scalable. Sensing devices can host more sensors by using the extension connector or modifying the PCB design. Live streaming and monitoring is limited by the maximum throughput of the BLE connection which depends on the used phone. Having one monitoring device (such as a phone) connecting to all sensing devices cannot be scalable by definition. Taking this into account, our system allows for multiple monitoring devices to be used as long as each sensing node is only connected to one monitoring device. Thanks to the full-resolution data logging in flash memory, post-measurement data download and processing are not affected by possible live-streaming interruptions or bottlenecks making the sensing system more resilient.

3.2.6 *Encasing*

We created a 3D-printed case for the sensor nodes that comprises an insert to house all the electronics and the battery as well as an outer casing with a strap for wear as shown in Figure 1a. The case supports different straps to support wear at different locations on the body. It is splash-resistant to protect the electronics from outside humidity and sweat.

3.3 POWER CONSUMPTION

In Table 1 we provide an overview of the marginal power consumption of various sensors and features of our sensing device. Power consumption for each sensor includes the overhead of writing the produced data to flash memory but not for BLE data streaming. The PPG sensor is the most power costly sensor at 1.72 mW for 3 different wavelength channels sampled at 128 Hz. The power consumption is mainly driven by the LED supply current. We minimized the power consumption by reducing the time of exposure and LED brightness and increasing photodiode sensitivity. The power consumption of all other sensors is below 1 mW. Besides the sensors, active BLE data

Table 1: Measured marginal power consumption of sensors and data streaming

Functionality / Sensor	Hz	mW
high-res IMU (ADXL355)	128	0.60
standard IMU (LIS2DH)	200	0.27
ECG (MAX30003)	128	0.49
1-ch PPG (MAX86141)	128	0.90
3-ch PPG (MAX86141)	128	1.72
Environmental (BME280)	10	0.62
Breathing (ADC)	10	0.24
EDA (ADC)	10	0.24
BLE (status, not connected)	-	0.23
BLE (status, connected)	-	0.44
BLE (active streaming 1.92 kb/s)	-	1.46

streaming is costly. This finding supports our design philosophy of reducing data streaming to the application-specific possible minimum.

3.4 CONCLUSION

In this chapter, we introduced the design and implementation of a scalable, multi-sensor wearable sensing platform aimed at high-quality physiological monitoring. The platform enables synchronized data acquisition across multiple body locations over long periods of time thanks to a focus on power efficiency.

The proposed system architecture allows for both real-time monitoring and high-resolution offline data analysis, making it suitable for data acquisition in applications including health and motion monitoring. The modular design, including extension ports and flexible sensor configurations, ensures adaptability for various applications.

In summary, this sensing platform enables long-term physiological and motion monitoring outside controlled laboratory settings without the need for any user interactions.

4 OFFLINE SYNCHRONIZATION

A requirement of cross-modal signal processing is accurate signal alignment. Though simple on a single device, accurate signal synchronization becomes challenging as soon as multiple devices are involved, such as during activity monitoring, health tracking, or motion capture—particularly outside controlled scenarios where data collection must be standalone, low-power, and support long runtimes.

In this chapter, we present a synchronization method that operates purely based on recorded signals and is thus suitable for offline processing: BMAR (Barometric and Motion-based Alignment and Refinement). BMAR needs no wireless communication between devices during runtime and does not require any specific user input, action, or behavior. BMAR operates on the data from devices worn by the same person that record barometric pressure and acceleration—inexpensive, low-power, and thus commonly included sensors in today's wearable devices. In its first stage, BMAR verifies that two recordings were acquired simultaneously and pre-aligns all data traces. In a second stage, BMAR refines the alignment using acceleration measurements while accounting for clock skew between devices. In our evaluation, three to five body-worn devices recorded signals from the wearer for up to ten hours during a series of activities. BMAR synchronized all signal recordings with a median error of 33.4 ms and reliably rejected non-overlapping signal traces. The worst-case activity was sleeping, where BMAR second stage could not exploit motion for refinement and, thus, aligned traces with a median error of 3.06 s.

4.1 OVERVIEW

Technological advances have led to a fast increase in the miniaturization of wearable technology, giving rise to many devices that can be worn throughout the day [25].

Among the main uses of wearable devices are activity tracking [4, 36, 50, 81, 94, 116, 176] and health monitoring [5, 27, 94, 122, 176, 205]. Both these types of devices incorporate multiple sensors that generate large amounts of time-dependent data that they must process in real-time, store, or transmit to another device. Oftentimes, the data is processed online to reduce storage or transmission requirements.

When raw recordings obtained on wearable devices are important, however, continuous data streaming places a large toll on battery power, and storing all data for offline processing is preferable. This is especially true of data collection experiments that enable later re-

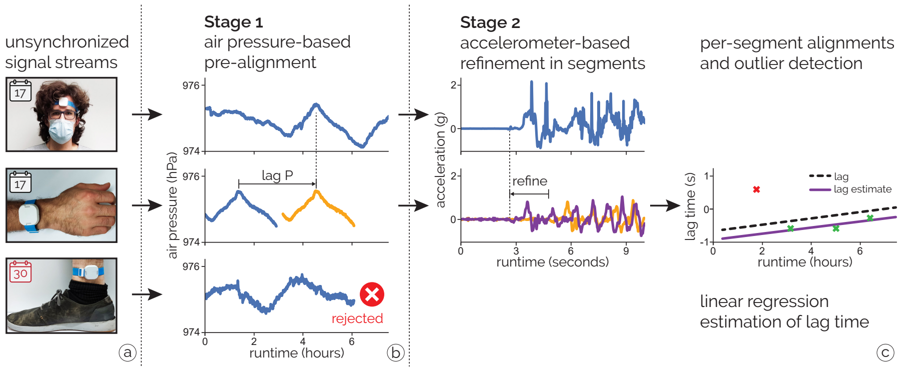

Figure 5: BMAR is a novel method to synchronize and align signals across devices without the need for specific user input, action, or explicit synchronization through wired or wireless communication (e.g., WiFi or BLE). BMAR is capable of synchronizing (a) independently recorded signals *after the fact* by (b) first pre-aligning recordings using air pressure as an inexpensive sensing modality that simultaneously allows us to reject non-simultaneous recordings. (c) In a second step, BMAR produces a refined signal alignment across sensor devices by cross-correlating accelerometer observations.

search efforts [30]. Beyond the sensor data collected by one device, data collections commonly include multiple wearable devices, either worn by one individual or across individuals, such that the collected records can later be combined for joint analysis [133]. To analyze data series stored across multiple devices, accurate temporal data synchronization is an essential requirement. Wearable sensing devices usually associate measurements with timestamps, for example, derived from an onboard RTC. Since timekeeping typically relies on a crystal oscillator with an accuracy of 10–100 ppm [180], the clocks between two devices may deviate by several seconds after hours of runtime, even if they matched at one point. Repeated synchronization is therefore essential for accurate data processing at a later point.

A common approach to cross-device time synchronization is using a central reference for timekeeping. The GPS is a popular choice because it synchronizes time at an accuracy of typically 10 ns [52] which exceeds that of an RTC and is only limited by hardware capabilities [113]. However, GPS signals are not reliable in locations with no clear view of the sky and the receiver's power consumption of several milliwatts renders them little suitable for repeated synchronization in low-power applications. Therefore, many wearables rely on a WSN that implements an online synchronization algorithm. While there are no established standards in WSN synchronization, multiple commonly used wireless synchronization protocols exist, including the Timing-Sync Protocol for Sensor Networks (TPSNs) [73], the Flood Time Synchronization Protocol (FTSP) [123], and Reference Broadcast Synchronization (RBS) [65]. However, all these approaches are designed and optimized for accuracy and are not suitable for very low-power applications [121].

A recent trend in low-power systems has been in time synchronization that optimizes for power consumption and thus limits communication and the amount of online processing [30]. For example, Makara, Tsybul'nyk, and Kurnyts'kyi's synchronization protocol tolerates average time discrepancies of 10 ms, enabling two individual wearable devices to regularly exchange timestamps through BLE at an average power consumption of just 133 μW [121]. In contrast, many other efforts have opted to perform data synchronization *entirely offline* and based on measurement data alone, which is useful for longitudinal in-the-wild collections. Spilz and Munz and Gilbert et al.'s methods build on magnetometer recordings to detect a predefined event for clock synchronization [80, 170], leveraging electromagnets built into a tray on which the devices are placed. In many applications, such as wearable health sensing or collections of activity data, synchronization is particularly needed when devices are worn by the same person to obtain a coherent dataset for each wearer. For this use case, researchers have also synchronized data series passively, detecting pre-defined events in the signals from IMUs for synchronization, such as clapping [164] or footsteps [14].

Following from previous efforts in this domain, several key challenges exist for passively synchronizing data series offline: Synchronization must be able to detect whether or not two data series were recorded at the same time as a prerequisite. Synchronization cannot rely on pre-defined activities or events that are required to be performed by a person one or more times. Synchronization cannot assume unique events to be present inside all data series that passively afford accurate alignment. Synchronization must ensure that data series are aligned with a global optimum, particularly in the presence of events that manifest in periodic patterns, such as walking, running, or cycling, and that can lead to locally optimal but overall inaccurate alignment.

In this chapter, we introduce BMAR, a novel offline synchronization method for devices worn by one person. BMAR requires no explicit input from users and accurately aligns cross-sensor signal traces without requiring users to perform particular motions or activities.

4.1.1 *Offline Signal Synchronization across Body-worn Devices*

As shown in Figure 5, BMAR synchronizes signal traces individually recorded by wearable devices by aligning them in two stages: a pre-alignment based on air pressure observations and a refinement step that processes accelerometer observations.

In Stage 1, BMAR derives a pre-alignment of signal traces based on air-pressure observations. We simultaneously leverage these measurements for the rejection of non-concurrent recordings. Both are possible because air pressure recordings are unique over time but near-identical across all devices worn on the same body. Barometers are inexpensive, small, very low-power, and not affected by motion artifacts, ideally suited for this scenario on wearable devices.

In Stage 2, BMAR refines alignment based on accelerometer measurements. Building on our pre-alignment to ensure globally optimal alignment, BMAR produces a refined alignment by cross-correlating accelerometer recordings. We compute these for patches of data, which makes our method more robust to short-term misalignments and accounts for clock skew between devices. Accelerometers are ubiquitous in wearable devices thanks to the versatility of their recordings, their low cost, and low power consumption. The distinct high-frequency events in motion across all on-body locations that these sensors observe are ideal for optimizing BMAR pre-alignment to reduce the alignment error.

To evaluate the effectiveness of our method, we conducted 10 sessions of data collection on 7 separate days. In each session, between 3 and 5 wearable devices recorded data from the same participant for up to 10 hours, resulting in 328 hours of total device runtime. The participant performed one of 3 activities: office work, cycling, or skiing to include a variety of lower and higher-intensity motions. In addition, we included sleep as a 4th activity to verify how well our method performs on recordings with little to no motion.

Our results showed that BMAR synchronizes signals with a median error of less than 40 ms during sessions with activity. This is well below the period of human activity and physiological signals [43, 118], setting our method up as a suitable complement to the increasing data monitoring efforts on multi-device platforms.

As part of BMAR, we make the following contributions:

1. BMAR, a novel method for signal synchronization that processes recordings offline through a pre-alignment and refinement stage, without relying on assumptions about signal content (e.g., specific actions or user input during recording).

2. Particular support for (wearable) low-power platforms, requiring no wireless connectivity or communication during operation, and leveraging only commonplace air pressure and accelerometer measurements.

3. An evaluation of our method across 3–5 custom-built wearable devices (with a high-resolution and a low-resolution sensor) across 4 activities against ground-truth measurements. During activities with human motion, our method aligned data traces with a median error of 33.4 ms, whereas this error was 3 s during activities without motion (sleep). Our method also reliably detected and rejected recordings that were not acquired concurrently.

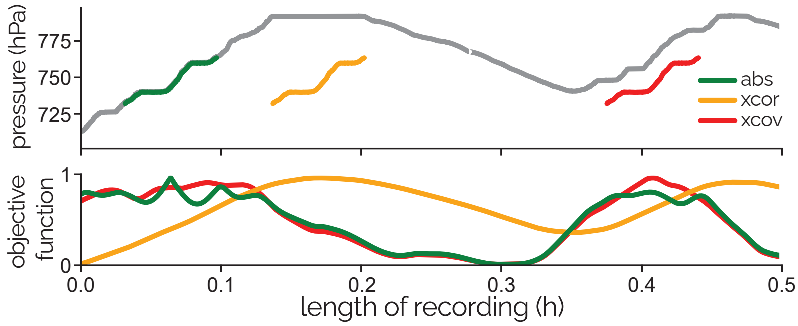

Figure 6: Example of aligning a short signal (top, orange) with a longer one (top, gray) based on cross-correlation which results in an objective function (bottom, orange) that is at its maximum where the longer signal is maximal. Similarly, the cross-covariance results in an alignment with the section of the longer signal that is the steepest (red). For 'slow' signals such as ambient pressure, alignment based on the sample-by-sample signal delta is better suited (green).

4.2 PROBLEM DEFINITION AND BACKGROUND ON SIGNAL ALIGNMENT

We now discuss preliminary considerations for our proposed method for synchronizing data series as well as the challenges involved.

4.2.1 *Temporal Alignment of Discrete Finite Signals*

The problem we are addressing is specified as follows: Given are discrete sensor recordings from two independent devices D_1 and D_2 which were acquired with temporal overlap but may be of varying duration. A recording always comprises traces of barometer, accelerometer, and optionally other sensors. The goal is to temporally align the two recordings based on the similarities in measurements contained in the sensor traces. We utilize a mixture of methods based on cross-correlation, cross-covariance, and signal delta.

4.2.1.1 *Cross-correlation and Cross-covariance to Align Discrete Finite Signals*

The approximation of the cross-correlation R_{fg} of two signals f and g of finite lengths N_f and N_g can be used to determine the best alignment of the two functions by finding the lag m_0 of g that maximizes R_{fg}:

$$R_{fg}[m] = \frac{1}{l[m]} \sum_{n=\max(0,-m)}^{\min(N_f-1, N_g-m-1)} f[n]g[n+m].$$

Here, the result is scaled relative to the length $l[m]$ of the overlap of the two signals:

$$l[m] = min(N_f, N_g - m) - max(0, -m).$$

The quality of this approximation decreases at the edges with shorter window lengths l when m is close to N_g or $-N_f$. The cross-correlation is commonly used to align data series such as gyroscope measurements [164] or accelerometer measurements [53], but it does not work well for all signal types. Most notably, signals with a characteristic timescale that is longer than the duration of the computation window (i.e., contain few distinct swings within the computation window) are prone to misalignment [187]. An example of misalignment in such a scenario is shown in Figure 6.

In some cases, this issue can be resolved by subtracting the mean values $\mu_{f'}$ and $\mu_{g'}$ from the respective signals prior to multiplication. The resulting formula is an approximation of the cross-covariance of the two signals:

$$R_{fg}[m] = \frac{1}{l[m]} \sum_{n=max(0,-m)}^{min(N_f-1, N_g-m-1)} (f[n] - \mu_f)(g[n+m] - \mu_g)$$

However, there are still scenarios in the realm of slow sensing modalities that are misaligned when using this formula for trace alignment as shown in Figure 6.

4.2.1.2 *Signal Delta to Align Discrete Finite Signals*

To align sensor signals with a long characteristic timescale in scenarios where the sensors have recorded the same underlying signal at the same time bar the impact of sensing tolerance (as is the case for air pressure measurements), the point-by-point difference of the signals can be analyzed. Considering the same two signals f and g of finite lengths N_f and N_g and the overlapping length $l[m]$ as before, a loss function, say L, can be formulated as follows:

$$L[m] = \frac{1}{l[m]} \sum_{n=max(0,-m)}^{min(N_f-1, N_g-m-1)} L_\delta(f[n] - g[n+m])$$

Where L_δ is the Huber loss function which reduces the sensitivity to measured outliers:

$$L_\delta(x) = \begin{cases} \frac{x^2}{2}, & |x| \leqslant \delta \\ \delta(|x| - \frac{\delta}{2}), & |x| > \delta \end{cases}$$

The resulting loss function $L[m]$ can be minimized to find the lag m_0 that aligns the two sensor traces the best. Alternatively, the standard deviation of the point-by-point difference of the two signals can be used as a loss function. This makes it invariant to the mean of the signal difference and effectively yields a measure of how well two signals run in parallel:

$$L[m] = \sqrt{\frac{1}{l[m]-1} \sum_{n=\max(0,-m)}^{\min(N_f-1, N_g-m-1)} |f[n]-g[n+m]-\mu|^2}$$

where μ is the mean of $f[n]-g[n+m]$:

$$\mu = \frac{1}{l[m]} \sum_{n=\max(0,-m)}^{\min(N_f-1, N_g-m-1)} f[n]-g[n+m]$$

4.2.2 *Clock Offset, Drift, and Skew*

The difference between two RTCs $\delta(t)$ can be characterized in three components [180]: (1) offset, the constant component of $\delta(t)$ caused by different starting points, (2) skew, the first derivative $\delta'(t)$ which indicates that one clock is running faster than the other resulting in a gradual de-synchronization, and (3) drift, the second derivative $\delta''(t)$ which is the variation in skew. This terminology is consistent with previous definitions as used by Mills [126] and Sundararaman, Buy, and Kshemkalyani [177]. For RTCs that rely on quartz crystals, clock drift is mainly caused by the temperature dependence of the latter which is typically characterized as $\mathrm{Err}(T) = 0.036\mathrm{ppm} \cdot (T - 25°C)$ [180]. This results in only minor clock drift for temperature differences commonly experienced by two wearable devices. Therefore, we assume clock drift to be 0 and de-synchronization effects to be linear between devices in our considerations, similar to related work [167].

4.2.3 *Measurement Errors of Air Pressure Sensors*

Air pressure measurements are subject to sensor tolerances. The resulting error can be characterized by an absolute accuracy (i.e., the absolute difference to ground truth) and a relative accuracy (i.e., the accuracy of the delta between two measurements by the same sensor). For commercially available barometric pressure sensors, the absolute accuracy is usually about an order of magnitude larger than the relative accuracy [90, 162, 163]. Therefore, BMAR uses absolute measurements to reject non-simultaneous measurements only and relies on relative measurements for trace alignment.

4.2.4 *Ensuring within-device Sensor Synchronization*

In offline settings, the RTC of a device does not provide absolute time information which could be used for cross-device synchronization but instead keeps track of the runtime since starting the device more accurately than internal sensor clocks would. With typical tolerances in the order of several percent, sensor clocks are not suitable to be used for timekeeping. Hence, a prerequisite for the alignment of multiple sensor traces across devices is the synchronization of traces with the RTC within each device.

4.2.4.1 *Alignment Offsets Due to Sensor Startup Latency*

Starting each sensor of one device at a known RTC timestamp does not guarantee synchronization. This is due to differences between the latency of each sensor for activation and starting the acquisition of its first sample. Because this latency is not commonly specified in the datasheet, it can usually not be compensated for without extensive calibration measurements. However, assuming that this latency is constant for a specific sensor type and only data from the same sensor type is compared across devices, the latencies cancel out and have no negative impact on the resulting synchronization results.

4.2.4.2 *Synchronized Sampling*

When sensors support external triggering, the RTC can be used to directly initiate measurements. This keeps sensor data traces perfectly synchronized with the RTC over arbitrarily long periods of runtime. In systems where sensors do not support external triggers from the RTC, sensors must rely on their internal clocks. In this case, the relative error in the sampling rate requires compensation.

4.3 BMAR: METHOD DESIGN FOR OFFLINE SIGNAL SYNCHRONIZATION

Optimal data synchronization across multiple wearable devices worn by a single person requires patterns in the data that are unique over time and do not depend on the body location where the device is worn. The presence of such patterns depends on user behavior and sensing modality.

Therefore, BMAR builds on barometric pressure sensors, which observe changes in the environment that affect all worn devices to a similar extent. This causes the individual on-body location of a sensor to manifest in a negligible amount and does not require our method to rely on specific user actions or activities. Barometric pressure sensors are inexpensive, low-power, and are commonly built into wearable devices. The signal they measure is influenced by changes in

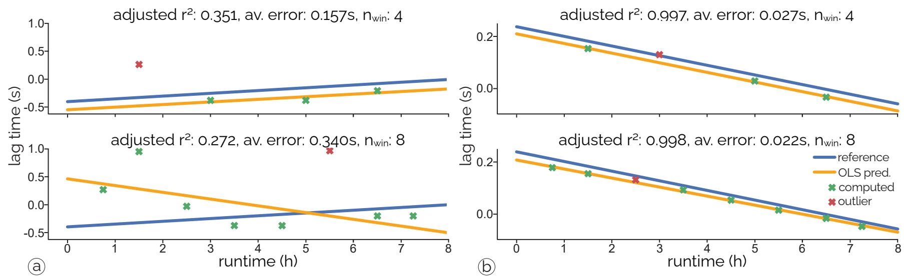

Figure 7: Approximations of the same $t_{lag}(t)$ with different n_{win} for two pairs of recordings. (a) Fewer windows lead to fewer outliers and a more accurate approximation for these recordings. (b) More windows lead to a reduction in alignment error for these recordings.

the altitude above sea level (e.g., walking up or down stairs, taking elevators, or performing any other motion not in a flat plane), meteorological changes in atmospheric pressure, or changes in ambient pressure caused by the immediate environment (e.g., opening a door of a closed room or car).

Our method aligns the data recorded by any two independent devices D_1 and D_2 or rejects the alignment in case it determines that the devices were not running at the same time or worn by the same person. The goal of this process is to determine the lag time $t_{lag}(t)$ of D2 compared to D1. $t_{lag}(t)$ itself is time-dependent due to RTC tolerances in D1 and D2 but is assumed to be linear.

4.3.1 *Rejection of Non-simultaneous Measurements*

The nature of air pressure measurements allows the rejection of two data traces that were not recorded at the same time. The difference of two temporally independent, single-sample air pressure measurements collected in the same location are expected to be normally distributed $X_1 - X_2 \sim \mathcal{N}(0, 2\sigma^2)$ with σ^2 being the variance of independent air pressure samples for the specific location. The probability of correct rejection of two such measurements can be calculated as follows:

$$p_{rej} = \frac{1}{2\sigma\sqrt{\Pi}} \int_{-2A}^{2A} e^{-\frac{x^2}{4\sigma^2}}$$

Where $\pm A$ is the range of the absolute accuracy of each barometric pressure sensor. When inserting real-world values such as $\sigma = 928\,\text{Pa}$ [33] and $A = 100\,\text{Pa}$ [162], the probability of correct rejection is 0.83. Using multiple measurements over time and incorporating a distribution over altitude above sea level, this probability is expected to increase. Here, the sensor accuracy is assumed to have strict lower and upper bounds and thus the rate of false rejection is zero when setting the acceptance tolerance to the same boundaries. Modeling the sensor accuracy as a normal distribution leads to the expected trade-off of minimizing false positives versus false negatives, however, specifications in datasheets usually do not provide enough information for such calculations.

To determine whether two devices D_1 and D_2 were running at the same time, we consider the alignment of the air pressure measurement data from both devices that minimizes the mean sample-by-sample difference and classify them as non-simultaneously recorded if this value surpasses a threshold.

4.3.2 *Stage 1: Air Pressure-based Pre-alignment*

BMAR first analyzes air pressure measurement data series to determine whether D_1 and D_2 were running at the same time and to find an approximation t_{lag}^P of $t_{lag}(t)$.

Since the error of t_{lag}^P is generally larger than the RTC clock skew effects between D_1 and D_2, it is computed as a constant value independent of time. We developed two algorithms to compute air pressure-based trace alignment. The *delta-error* method computes the mean Huber loss based on the point-by-point difference between two traces for all possible alignments (Section 4.2.1.2). The lag time corresponding to the alignment with the minimum mean Huber loss is selected as t_{lag}^P. For this method, the data from both devices is low-pass filtered to remove noise and normalized to make the alignment invariant to constant offsets due to absolute sensor accuracy tolerance.

The *delta-std* method computes the standard deviation of the point-by-point difference between two traces for all possible alignments and corresponding lag times. The lag time resulting in the minimum standard deviation is selected as t_{lag}^P. The raw data is used for this method.

4.3.3 *Stage 2: Accelerometer-based Refinement and Optimization*

Next, BMAR leverages accelerometer signals to refine t_{lag}^P and optimize the alignment. We maximize the cross-correlation of the magnitudes to find the best alignment $t_{lag}^{Accel}(t)$. To enable the time-dependent approximation of $t_{lag}(t)$ and to make the algorithm more robust to time-limited measurements that facilitate misalignment, the available data trace is split into n_{win} windows for each of which a lag time is calculated separately. The result of this is a series of n_{win} calculated lag times temporally associated with the center point of the respective window that the original trace was split into. We perform Ordinary Least Squares (OLS) regression on these lag times and remove outliers using the Bonferroni outlier test. This results in a linear approximation $t_{lag}^{Accel}(t)$.

The optimal n_{win} depends on the data. Therefore, we repeat the computation, starting with $n_{win} = 4$ and increase the number of windows by a factor of 2 until the window size drops below 1 second or $n_{win} > 512$. The result of this is a set of linear approximations of $t_{lag}(t)$. We use the r^2 adjusted for residual degrees of freedom of the OLS to select the most promising approximation. This is based on the assumption that local misalignments within some of the windows do not appear linear over the whole trace, making r^2 an indicator of alignment quality. Figure 7a shows an example where the linear approximation with a smaller n_{win} is more accurate while Figure 7b shows an example of an alignment where a larger n_{win} leads to a better result.

4.4 SYSTEM IMPLEMENTATION OF BMAR

To implement and evaluate BMAR, we used the sensing platform introduced in Chapter 3.

4.4.1 *Sensors*

For later pre-aligning of all data series, we use the digital barometric pressure sensor (BME280, Bosch Sensortec) of the devices. The sensor is operated in forced mode to enable synchronization to the RTC of the SoC, which triggers a measurement every 100 ms, leading to a sampling rate of the barometer of 10 Hz. We used no oversampling or internal Infinite Impulse Response (IIR) filter. As per the datasheet, the absolute accuracy of the measurements lies within a range of ± 1.0 hPa ($\pm 3\sigma$) and the relative accuracy is ± 0.12 hPa ($\pm 3\sigma$) [162] which are low-end characteristics compared to pressure sensors currently available commercially.

For refining synchronization during our optimization stage, we used the two available accelerometers for comparison. The ADXL355 high-resolution 3-axis digital accelerometer operates in FIFO mode, receiving triggers for measurements from the RTC of the SoC. This results in a sampling frequency of 128 Hz. We set the range to ± 2 g at a resolution of 20 bit (4 µg/LSB) [56]. The LIS2DH low-cost 3-axis digital accelerometer also operates in FIFO mode but does not support external triggers. This accelerometer was configured to sample at 200 Hz using the internal clock source, at a range of ± 2 g and a resolution of 8 bit (4 mg/LSB) [152].

4.4.2 *Synchronization*

4.4.2.1 *Buffered Sampling With Compensation*

Since the low-cost accelerometer does not support external measurement triggering from the RTC but relies on its internal clock source, its measurements require compensation to preserve the synchronization of all sensors within the device. During offline processing, the expected total data length after each FIFO readout can be calculated based on the logged RTC timestamps of the readout and the sensor activation. To account for the latency between sensor activation and data availability in the FIFO, we calculate the expected length relative to the first FIFO readout:

$$l_{exp}(t_n) = (t_n - t_0) \cdot f_{sensor} + l(t_0)$$

With the sampling frequency of the sensor f_{sensor}, the expected length l_{exp} at the RTC timestamp t_n of the n^{th} FIFO readout after the initial readout at RTC timestamp t_0 after which a total of $l(t_0)$ samples were present. Because FIFO readouts are usually not performed on a perfectly regular basis, a lack of samples after a FIFO read-out at t_n (i.e., $l(t_n) < l_{exp}(t_n)$) may automatically be compensated for by a sample which narrowly slipped into the next FIFO readout. We account for this by only adding or removing samples to match the expected length when no inverse operation is necessary in close temporal proximity. This results in a sensor data trace which is synchronized to the device's RTC with a tolerance of $1/f_{sensor}$.

4.4.2.2 *Hyperparameters*

For the rejection of non-simultaneous traces (Section 4.3.1), we set the threshold of the mean delta between two traces to the absolute tolerance of the used barometric pressure sensor of $100\,\mathrm{Pa}$ [162], trying to achieve a recall close to 1. Lowering it increases precision and lowers recall. We use the same value for the δ of the Huber loss in the delta-error method (Section 4.3.2) such that it only affects rarely occurring outliers in the measurements. We set the maximum range of the accelerometer-based refinement (Section 4.3.3) to ± 5 seconds. If the pre-alignment differs more than that, the refinement cannot fully compensate for it. On the other hand, increasing this range increases the risk of misalignment caused by this stage.

4.5 EVALUATION

To quantify the performance of BMAR, we conducted a data recording experiment during everyday activities, using between 2 and 5 devices at once with several repetitions. Table 2 shows an overview of the configurations in our evaluation. Since there is no impact of individual users' behavior on our method, all recordings were conducted by one experimenter.

To analyze the performance of our method, we consider pairs of devices that either did or did not record data at the same time while worn by the experimenter. This means that if 3 devices were worn at the same time, we considered them as 3 pairs of simultaneously

Table 2: For our evaluation, the experimenter recorded data during these activities, wearing multiple devices at the respective locations on the body for the given durations. Each activity was repeated two or more times.

activity	#	locations	runtime	rep
sleep	5	head, chest, ankle, wrist	9.9 h	4
work	3	head, chest, ankle	3.8 h	3
cycle	4-5	head, chest, ankle, wrist	1.4 h	2
ski	3	chest, wrist	8.8 h	2

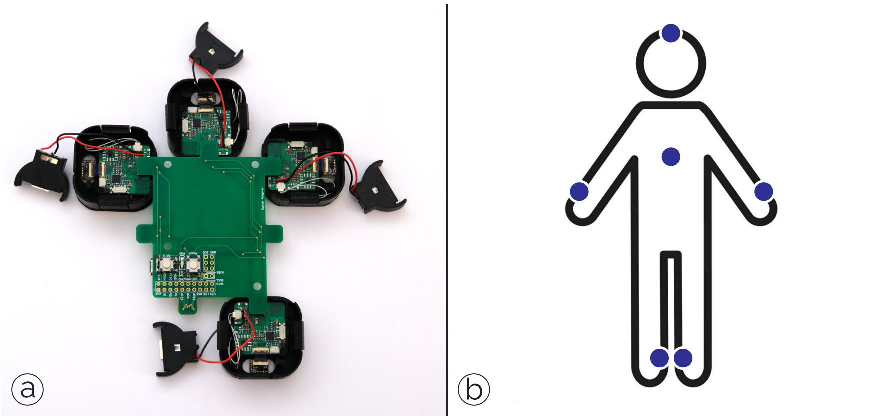

Figure 8: For the evaluation of BMAR, we adapted the sensing platform introduced in Chapter 3. We developed an additional board to connect multiple devices and trigger ground-truth synchronization events (a). For the evaluation of BMAR, sensors were placed in 6 on-body locations: On the forehead, sternum, wrists, and ankles (b).

running devices in our evaluation (D_1 and D_2, D_1 and D_3, D_2 and D_3) and also included pairings of each device with a random device from another recording which was not acquired at the same time.

4.5.1 *Procedure*

To capture as many use case scenarios as possible, we acquired data during various different activities and in multiple on-body device locations.

4.5.1.1 *Locations*

As shown in Figure 8c, devices were attached to the following locations for a recording session: strapped onto the forehead, strapped onto the sternum, on one of the wrists like a watch, or above one of the lateral malleoli.

4.5.1.2 *Activities*

The experiment included two categories of sessions: active wearer and passive wearer. For active sessions, the experimenter performed activities such as cycling, skiing, or office work over the full duration. For this session type, we collected 31 pairs of device recordings running at the same time, with an average runtime of 3.73 hours.

For passive sessions, the experimenter remained stationary over the full duration. We recorded these data during sleep, turning on the devices after getting into bed and turning them off before exiting. For this session type, we obtained 40 pairs of devices with simultaneous recordings and an average runtime of 9.9 hours. Due to the lack of activities and motion, this session type can be considered a worst-case baseline for our analysis.

4.5.2 *Ground-truth Synchronization*

We devised an apparatus to establish repeated ground-truth synchronization across all devices. The experimenter manually attached this apparatus to all devices at the beginning and the end of each session. This resulted in two ground-truth synchronization events in our logs for reference. The apparatus is shown in Figure 8a. It sends a trigger from a manual, debounced push-button to a digital pin of the SoC in each device. The tethered connection ensured accurate reference values without any impact of wireless communication.

To ensure that each device could register the precise moment of a button press, we used a timer block on the SoC to support more finely-grained timing resolution between the RTC ticks. The incoming synchronization signal was logged through the timer's event capture functionality, which logs the exact timestamp and triggers an interrupt to process the event. In theory, the accuracy of this mechanism is limited by our timer clock of 32.768 kHz, which corresponds to a tolerance of 0.03 ms. However, the tolerance of the timer event trigger capture is not specified by the SoC manufacturer. Hardware limitations (e.g., stray capacitances and tolerances in logic voltage level detection) may also increase this error.

4.5.2.1 *Accuracy of Ground-truth Synchronization*

To evaluate the accuracy in practice, we sent multiple synchronization signals to all connected devices in short succession within a few seconds. Assuming no clock skew between the first and last synchronization signal, we compared the alignment of the multiple detected events on the time axis. Across 349 such synchronization sequences with 3 signal triggers each, we found a mean error of $\mathrm{ME} = 1.2\,\mathrm{ms}$ ($\sigma = 1.1\,\mathrm{ms}$) for a single signal trigger. To improve accuracy in the

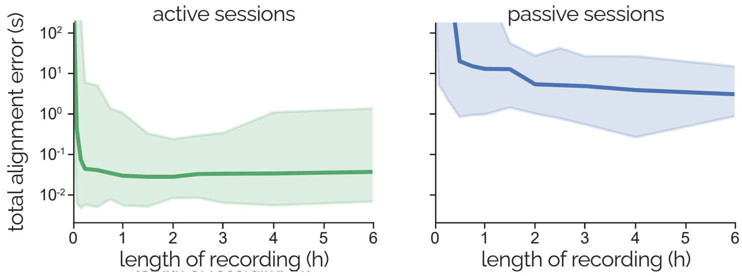

Figure 9: Median alignment error of BMAR's synchronization. The band denotes the range between 10th and 90th percentile.

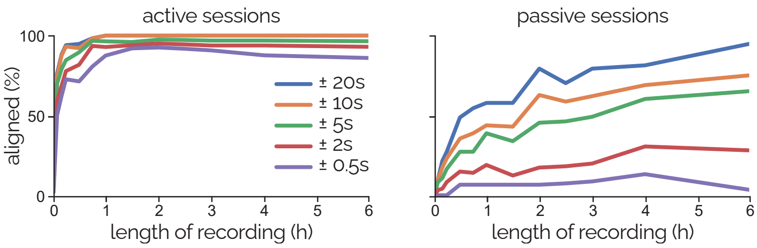

Figure 10: Percentage of traces that BMAR correctly aligned with a given error tolerance.

evaluation, we always used the mean across 3 consecutive signal triggers.

4.5.3 *Duration of Recordings*

For temporal alignment evaluations, the duration of the recording has an impact on the results (i.e., the probability of correct temporal alignment increases with longer recordings). We conducted the evaluations considering various different lengths of recordings by cropping recordings to the respective length. If a recording was multiple times the desired length, we split it and used up to 5 non-overlapping splits separately. The number of traces for different lengths is shown in Table 3. When aligning traces, we only cropped or split one of the two traces and preserved the length of the other trace. We considered the length of the shorter trace to be the minimum overlapping time of the two traces to prevent negative effects on the edges. Effectively, this means we measured the accuracy of aligning a trace of a specific length to a longer one that overlaps it.

4.5.4 *Results*

We first report the results of our complete method spanning both air pressure-based pre-alignment and accelerometer-based refinement. We then report the results of the rejection of non-simultaneously recorded data, of pre-alignment, and of refinement separately.

Table 3: Median alignment error in seconds of the complete synchronization method. The number of pairs of recordings used is stated in brackets.

	length of recording					
	5 min	10 min	15 min	30 min	1 h	⩾ 6 h
active	0.442 (155)	0.074 (155)	0.044 (155)	0.041 (107)	0.028 (73)	0.034 (31)
passive	≫60 (200)	≫60 (200)	≫60 (200)	20.2 (200)	5.42 (200)	3.06 (40)

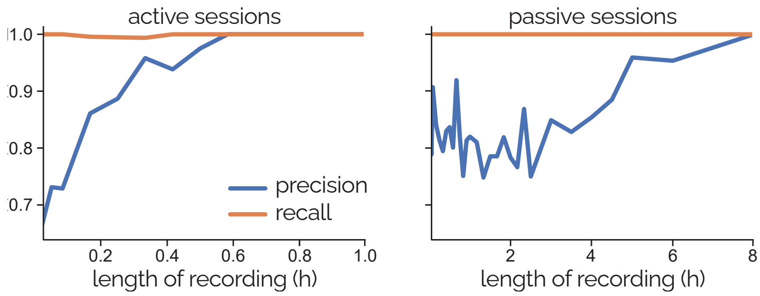

Figure 11: Precision and recall rates for BMAR detection of simultaneously acquired recordings.

4.5.4.1 *Overall Synchronization Accuracy*

As shown in Figure 9 and Table 3, the synchronization of active sessions achieved a median alignment error of 33.4 ms for recording durations of 15 min and longer, 73.6 ms for durations of 10 min, and 443 ms for durations of 5 min. The synchronization of passive sessions (i.e., a worst-case baseline without user motion or activity) resulted in a median alignment error of 3.06 s for recording durations of 6 h and longer, 5.42 s for durations of 2 h, and 20.2 s for durations of 30 min. Figure 10 shows the percentage of synchronizations that our method aligned within a given error tolerance.

We report the errors of unconstrained synchronizations with median and percentiles because the error of alignment for outliers (i.e., failed, effectively random alignments) can be up to four orders of magnitude larger than the median error of all alignments. The impact on the mean error and standard deviation makes them less informative for evaluation.

4.5.4.2 *Accuracy of Stage 1: Pressure-based Pre-alignment*

In stage 1, air pressure measurements were used to detect non-simultaneous measurements and do a first trace alignment of simultaneously acquired traces.

When using a sensor tolerance of 100 Pa as the threshold for the maximum mean difference between two traces to detect simultaneously recorded air pressure data, we reach a precision and recall of 1.0 for recordings of active sessions that contain at least 35 minutes of

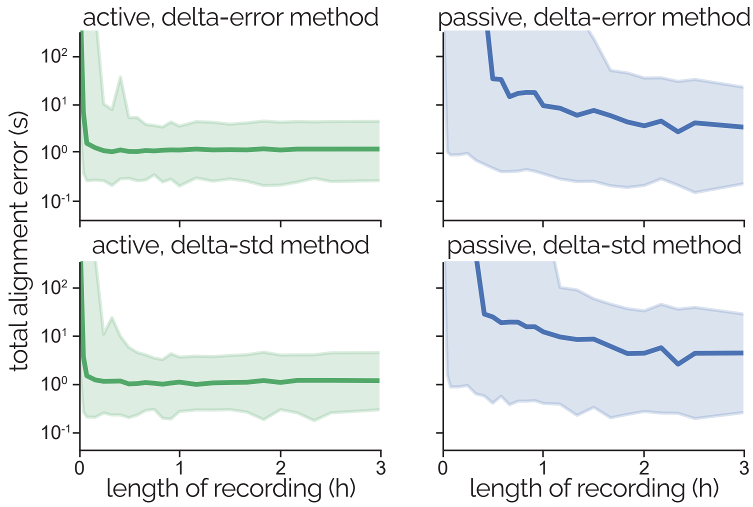

Figure 12: Median alignment error of BMAR first stage (air pressure-based pre-alignment). The band denotes the range between the 10th and 90th percentile.

data and recordings of inactive sessions of at least 8 hours of recording time. Figure 11 shows the precision and recall rates for recordings of different lengths.

These results are based on tests with the barometric pressure sensor recordings of each device D_n both paired with simultaneously and randomly selected, non-simultaneously recorded data. There is a trade-off between precision and recall: Lowering the threshold improves precision but lowers recall. However, such a lowered recall can generally not be compensated for by increasing the duration of the recording.

Aligning active sessions based on air pressure produced a median alignment error of 1.14 s for recording durations of 10 min and longer, 1.53 s for durations of 5 min, and 3.8 s for durations of 3 min, as shown in Figure 12. The synchronization of passive sessions (i.e., our worst-case baseline) based on air pressure resulted in a median alignment error of 3.36 s for recording durations of 3 h and longer, 12.4 s for durations of 1 h, and 25.3 s for durations of 30 min.

Of the two tested alignment methods, the *delta-std* method performed better than the *delta-error* method. Across all recording lengths greater than 0.5 hours (i.e., after initial convergence), the median alignment error of the *delta-std* is on average 48 ms lower for active sessions and 237 ms lower for passive sessions.

4.5.4.3 *Accuracy of Stage 2: Accelerometer-based Refinement*

Testing the accelerometer-based synchronization refinement in isolation for active sessions and under the assumption of a pre-alignment tolerance of ± 1 s, produced a median alignment error of 32.0 ms for recording durations of 15 min and longer, 40.6 ms for durations of 5 min, and 218 ms for durations of 1 min, as shown in Figure 13a. The

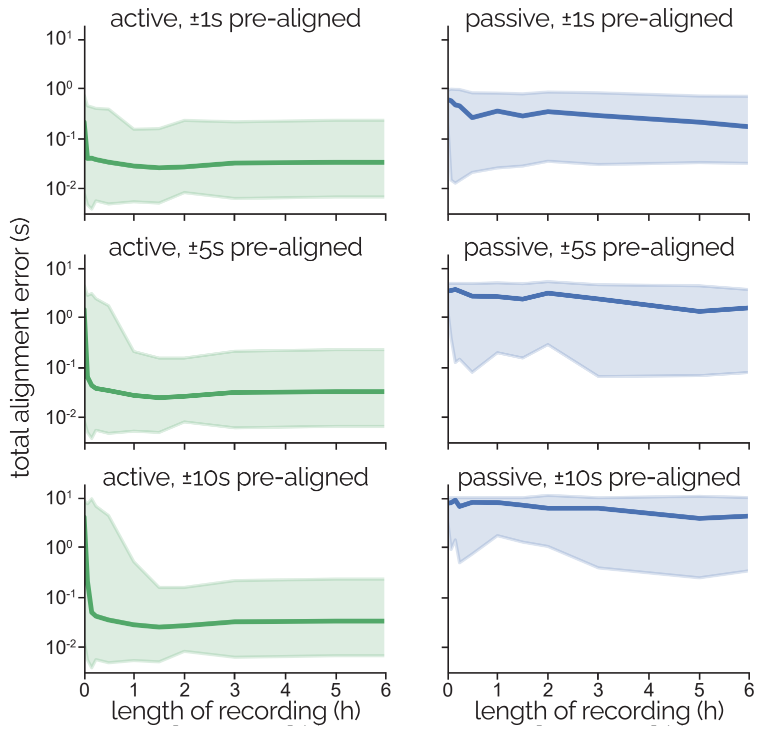

(a) high-resolution accelerometer

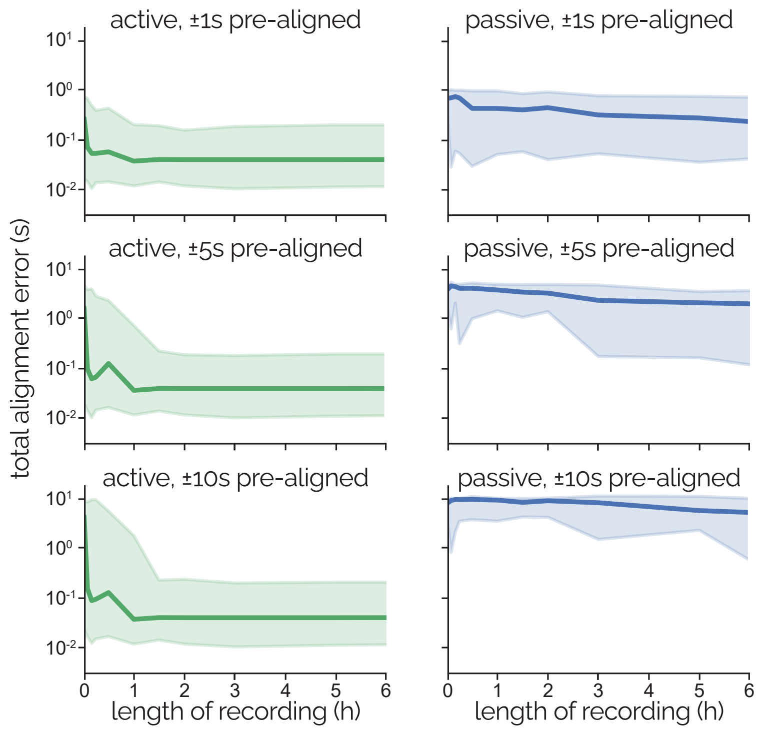

(b) standard-resolution accelerometer

Figure 13: Median alignment error after BMAR second stage (accelerometer-based refinement) for different pre-alignment errors.

refinement of passive sessions (i.e., worst-case baseline) resulted in a median alignment error of 274 ms for recording durations of 15 min and longer, 580 ms for durations of 5 min, and 606 ms for durations of 1 min. As shown in Figure 13a, the median error increases with higher pre-alignment tolerances.

We tested the refinement with data from both a high-resolution accelerometer (ADXL355, Analog Devices) and a standard-resolution accelerometer (LIS2DH, STMicroelectronics). The high-resolution accelerometer produced better results, yielding an average alignment error that was 21.5 ms lower when analyzing all recording durations, which resulted in a median alignment error below 0.1 s.

As shown in Figure 13b, the results we obtained with the standard-resolution accelerometer for active sessions and under the assumption of a pre-alignment tolerance of ± 1 s, produced a median alignment error of 44.2 ms for recording durations of 15 min and longer, 71.9 ms for durations of 5 min, and 276 ms for durations of 1 min.

The refinement of passive sessions (i.e., worst-case baseline) based on the standard-resolution accelerometer resulted in a median alignment error of 345 ms for recording durations of 15 min and longer, 685 ms for durations of 5 min, and 703 ms for durations of 1 min.

4.6 DISCUSSION, LIMITATIONS, AND FUTURE WORK

Our evaluation confirmed the robustness of our method to detect recordings that were acquired simultaneously to synchronize data stemming from multiple wearable devices. The accuracy our method achieved for recordings during active sessions is high enough to correctly match each occurrence of fast repetitive events across devices. This includes the fastest vital sign (i.e., heartbeats of a child at 3.5 Hz) and fast body motions (i.e., the cadence of a professional sprinter at 4 Hz or the catch rate of the world's fastest juggler at 9 Hz). Faster motions by humans usually occur with longer intervals in between, making cross-device analysis easier. Unsurprisingly, the results were worse for our baseline recordings, which did not include any activity other than sleeping. However, given long recordings of multiple hours, non-simultaneous recordings are still rejected reliably and the synchronization tolerance reaches 3.06 seconds which may be sufficient for the analysis of changes in sleep position.

4.6.1 *The Relevance of Pre-alignment Using Air Pressure Measurements*

Due to differences in the motion characteristics across body parts, IMU data does not lend itself well to aligning data traces without additional constraints. Using cross-correlation, for example, the resulting objective function exhibits many maxima and trace alignment without restrictive constraints on the search window or specifically performed motions by the wearer is challenging. For a more general, rough alignment, the use of a sensing modality with a longer characteristic timescale is more effective in producing a cost function that monotonically converges to a global minimum.

Temperature, perhaps the most commonly logged sensing modality, does include features common across devices worn in different body locations. Events like moving from indoors to outdoors with an abrupt change in ambient temperature could be used for synchronization, but they may be attenuated or delayed due to varying thermal insulation of the devices by clothing.

Air pressure is thus a suitable modality for cross-device synchronization: Air pressure measurements are barely affected by anything other than the geographical location, time of recording, and sensor tolerance. The former two are identical for multiple devices worn simultaneously.

4.6.2 *Impact of Sensor Accuracy*

As mentioned in Section 4.5.4.3, our standard-resolution accelerometer performed similarly well as the high-resolution accelerometer during refinement, indicating that our method does not require an expensive accelerometer for accurate signal synchronization. The range of commercially available barometric pressure sensors is smaller compared to accelerometers and the sensor used for the evaluation can be considered to be at the low end of the spectrum in terms of accuracy: In the relevant temperature and measurement range, the BME280 we evaluated BMAR with has an absolute and relative accuracy of $\pm 100\,\mathrm{Pa}$ and $\pm 12\,\mathrm{Pa}$, respectively [162]. This compares to $\pm 30\,\mathrm{Pa}$ and $\pm 6\,\mathrm{Pa}$ for the recently released BMP581 (Bosch Sensortec) [163] as well as $\pm 50\,\mathrm{Pa}$ and $\pm 3\,\mathrm{Pa}$ for the inexpensive SPL07-003 (Goertek Microelectronics Inc.) [90]. This suggests that our results likely generalize to most available pressure sensors, including inexpensive commodity barometric pressure sensors.

We expect a better absolute accuracy to improve the rejection of non-simultaneous recordings and a better relative accuracy to improve the pre-alignment of the recordings. In the future, we plan to characterize the specific impact of sensor characteristics (accuracy, sampling rate, resolution, noise, and RTC synchronization) on the synchronization error.

4.6.3 *Needed Overlap of Cross-sensor Signal Recordings*

In our current design, BMAR aligns signals of a specified duration l_1 to a recording of length l_2, which fully contains the former (Section 4.5.3). However, our method also works in scenarios where signals only *partially* overlap. In these, the effectively smaller overlap decreases the amount of data for calculating synchronization cost and may lead to a higher risk of misalignment. Considering a minimum overlap of $l' \leqslant l_1$, the expected error is in the range of our evaluation results for total recording lengths between l_1 and l'.

4.6.4 *Limitations of Accelerometer-based Refinements*

The improvement in refining synchronization based on accelerometer signals depends on the signal characteristics. For our method, we cannot give worst-case guarantees; if recordings contain no events or asynchronous patterns, our refinement may worsen the result of the pre-alignment.

To prevent potentially negative effects of our refinement, BMAR limits the range within which pre-alignment sequences may be refined as part of our synchronization method. Since the loss of the accelerometer-based alignment does not generally decrease monotonically to a correct global minimum, the refinement may worsen the result of a pre-alignment that was successful but with an error larger than the range of the refinement. Figure 13a shows how our refinement for short recordings during inactive sessions is close to a random offset within alignment tolerance and only decreases for longer recordings. Future work could improve our approach by leveraging the pre-alignment confidence and suitability of the accelerometer data to dynamically determine the range of the refinement.

4.6.5 *Devices across Multiple Users and Locations*

BMAR is currently designed to synchronize signals captured on devices worn by a single person. However, our method of aligning signal traces based on air pressure recordings could extend to synchronizing signal streams captured across multiple people and beyond the realm of wearable devices. We expect the method to produce accurate results if all devices are in the same location or move collectively (i.e., in a car, train, or elevator), but we anticipate a negative impact of movements that are spread out between devices, such as when groups of people who wear them do not transition between building floors at the same time. Depending on the setting, the use of different sensing modalities may be appropriate such as sound or temperature. We tested the latter early on in our project but found that for wearable devices, the impact of clothing and body temperature makes cross-device synchronization challenging. It may, however, be a viable option when synchronizing non-wearable devices.

4.6.6 *Limitations of BMAR's Evaluation: Single Participant vs. Breadth of Activities for More Insight*

We acknowledge that in our evaluation, multiple devices collected data on a single participant only. While such an evaluation may be uncommon, it is important to point out that BMAR is not specific to human factors or capabilities; rather, the performance of the method depends on the type of activities.Therefore, the breadth of assessed activities and environments is considerably more important in evaluating our method than the number of participants, each of which would perform activities only slightly differently.

In our evaluation, the sensor devices were affixed to five locations on the participant's body, without strict enforcement of inter-device distance, which covered a range of device constellations and variability similar to varying body heights across participants. Our results showed no notable impact of the body location on the accuracy of trace alignment.

Most importantly, our evaluation has shown that the most challenging activity for our cross-sensor synchronization approach is no activity (i.e., resting and sleeping). While we believe that our results generalize and an evaluation with multiple participants would not change the findings, we acknowledge that for assessing this conclusively, our evaluation could be extended to multiple participants in the future.

4.6.7 *Speculative Comparison with Online Synchronization*

Online synchronization through a wireless network may be a viable alternative to BMAR to reach higher accuracy (Section 2.2.1). A main advantage of BMAR is its low power consumption: At the sampling rates used in our evaluation, a barometric pressure sensor operates on less than $35\,\mu W$ [163] and the low-resolution accelerometer uses less than $70\,\mu W$ [152]. Of the commonly used wireless protocols, currently, only Bluetooth low energy reaches comparable values: Makara, Tsybul'nyk, and Kurnyts'kyi synchronized two nodes with a power consumption of 133 uW with a typical accuracy of 10 ms.

However, producing reliable implementations in practice is challenging, limiting the choice of SoC, and no prior work has shown synchronization of more than two nodes with power consumption in this order of magnitude. In such a scenario, a more powerful device such as a smartphone may be needed as central device for the other nodes to remain in the stated range of power consumption while BMAR power requirements are unaffected by the number of devices.

Generally, most existing synchronization algorithms are optimized for low error—not low power consumption, which has thus often not been evaluated in prior work. Because most wearable devices already routinely sample acceleration and/or air pressure from built-in sensors, our method imposes no additional power consumption during runtime. Additionally, our method may be used with devices that are not engineered to be compatible to be used in the same wireless network.

4.7 CONCLUSION

As part of this dissertation, we have presented BMAR, a novel method to synchronize sensor data across multiple wearable devices without the need for communication or explicit user interaction during runtime. BMAR leverages barometric pressure sensors and accelerometers, which are inexpensive, low-power, and commonly found in wearable devices. Therefore, they represent a simple solution to synchronizing wearable devices without the need for inter-device communication, allowing data recording and device operation to focus on low power consumption in miniaturized devices.

BMAR two-stage process comprises pre-alignment, which we use to avoid latching on to local but not a global optimum in alignment, and refinement. We evaluated BMAR using custom-built, low-power wearable sensing devices that recorded a series of signals across a variety of activities over the course of up to 10 hours. Our method achieved accurate synchronization with a median error of 33.4 ms, which is low enough for most practical purposes. Even during activities with little to no motion such as sleeping, our method synchronized signals with an error of 3 s, all while reliably rejecting signals that were not concurrently recorded across devices.

We believe that BMAR is a practical approach for supporting cross-device data collections that build on unobtrusive miniaturized devices, focus on duration, and ease of use.

5 REAL-WORLD PHOTOPLETHYSMOGRAPHY DATASET

In this chapter, we show that state-of-the-art HR estimation methods struggle when processing *representative* data from everyday activities in outdoor environments, likely because they rely on existing datasets that captured controlled conditions. However, PPG-based HR estimates can be substantially impacted by factors such as the wearer's activities, sensor placement, and resulting motion artifacts, as well as environmental characteristics such as temperature and ambient light – all factors that are not sufficiently represented in controlled environments.

We introduce a novel multimodal dataset and benchmark results for continuous PPG recordings during outdoor activities from 16 participants over 13.5 hours, captured from four wearable sensors, each worn at a different location on the body, totaling 216 hours. Our recordings include accelerometer, temperature, and altitude data, as well as a synchronized Lead I-based electrocardiogram for ground-truth HR references. Participants completed a round trip from Zurich to Jungfraujoch, a tall mountain in Switzerland over the course of one day. The trip included outdoor and indoor activities such as walking, hiking, stair climbing, eating, drinking, and resting at various temperatures and altitudes (up to 3,571 m above sea level) as well as using cars, trains, cable cars, and lifts for transport—all of which impacted participants' physiological dynamics. We also present a novel method that estimates HR values more robustly in such real-world scenarios than existing baselines.

Dataset & code for HR estimation is available on the project website:
https://siplab.org/projects/WildPPG

Parts of this dataset also served as a basis for further research within this dissertation. In Chapter 6 we present a method to fuse PPG signals across body locations for better HR detection and in Chapter 7 we introduce a learning-based method to extract information from multiple PPG wavelengths to generate a more robust PPG signal.

5.1 OVERVIEW

Obtaining *reliable HR estimations* from PPG signals in real-world conditions and during real-world activities is challenging. Several external factors negatively impact accurate estimation, such as motion artifacts [204] that arise from a person's movements and performed activities [148, 155], sensor misplacement [106] or slippage during wear, and environmental conditions that change over time such as low temperatures [96] or high levels of ambient light [97]. Basic implementations of HR detection have been validated on comparably clean reference datasets, which do not adequately represent the variability and noise introduced by everyday activities and conditions. Even for simple activities such as walking and running [19], recent studies have provided evidence for the resulting detrimental effects on signal quality and, thus, the performance of HR tracking implementations.

To combat the effects of motions and activity on HR detection, prior work has proposed data-driven methods that learn to recover meaningful estimates under noisy conditions, including supervised methods [46, 151, 156], and rule-based techniques [153, 179, 203, 204]. While these and other methods estimate HR with high accuracy, their robustness to real-world effects is limited by the diversity of the datasets they were trained on. Because existing datasets were dominantly captured in controlled conditions, they may not adequately represent the noisy nature under which today's wearables obtain PPG measurements in daily life and "in the wild."

As part of this dissertation, we introduce WildPPG, a novel dataset of reflective PPG measurements from four locations across the wearer's body and reference recordings from an electrode-based Lead-I ECG for HR estimation tasks from everyday activity outside controlled environments. Our dataset comprises continuous and synchronized recordings from 16 participants with an average duration of 13.5 hours during a round-trip to a tall mountain in Europe, including various daytime activities and modes of transport. As shown in Figure 14, our dataset additionally includes the continuous measurements from a 3-axis inertial sensor (accelerometer), barometric altitude sensor, and temperature sensor. Our data collection simultaneously recorded all these signals from four devices across the wearer's body (wrist, ankle, sternum, head), totaling 216 hours of synchronized multi-modal recordings that can enable a wide variety of future analyses.

We also contribute the implementation of multiple baseline methods for HR estimation and evaluate them on our dataset to identify weaknesses. Based on our analysis, we introduce a learning-based method that takes temperature as input in addition to the raw PPG signal and show that its temperature awareness allows it to produce more robust HR estimates.

Collectively, we make the following contributions:

- a multi-modal real-world dataset of continuous, noisy, and synchronized PPG signals (raw red, green, infrared) from 4 body locations (forehead, sternum, wrist, ankle) and ECG recordings for reference, continuously collected from 16 participants over an average duration of 13.5 hours (minimum duration: 12.3 hours),

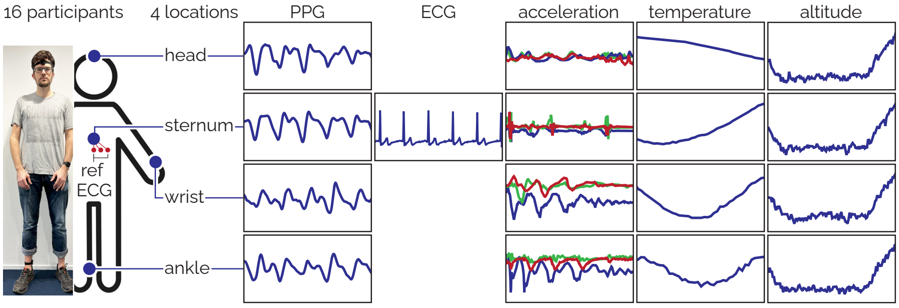

Figure 14: WildPPG comprises multi-modal signals from wearable devices at four sites on the body. Each device continuously recorded synchronized signals from a 3-channel reflective photoplethysmogram (red, green, infrared PPG), a 3-axis inertial sensor (accelerometer), a temperature, and a barometric altitude sensor. For reference, the sternum device continuously recorded a Lead-I ECG from body-mounted gel electrodes to provide ground-truth HR estimates.

- a data collection study that covered everyday outdoor activities outside controlled (laboratory) conditions: During collection, participants completed a round-trip to a mountain station at over 3500 meters above sea level using various modes of transportation (car, train, cable car, elevator) and engaging in various outdoor and indoor activities (walking, hiking, stair climbing, eating, drinking, social interactions, restroom breaks, and resting in varying conditions), including environmental conditions of low and medium temperatures and various degrees of solar radiation,

- an accuracy analysis of baseline methods that estimate HR on WildPPG based on environmental characteristics. We also contribute a learning-based method for HR estimation that takes a raw PPG signal and temperature as input and show that it produces more robust estimates, and

- additionally recorded and synchronized modalities in our dataset as the basis for future algorithm developments, including 3-axis inertial measurements (accelerometer), barometric altitude, and temperature data from each of the 4 body locations (some exposed, some covered by clothing).

5.2 RELATED DATASETS

Processing PPG signals and most relevantly robust HR estimation in real-world situations remains challenging due to the noisy nature of the PPG signals [204]. In order to develop, evaluate, and compare novel approaches, publicly available datasets are essential. However, existing public datasets are limited in terms of data size and appli-

Table 4: Qualitative comparison of WildPPG with previous datasets for HR estimation

Dataset (year)	Collection methodology	Hours (total)	Body-worn locations *	Wavelengths (PPG)	Multi-modal	In-the-wild** dataset	Temperature / Altitude changes
Ours (2024)	In the wild	216	{Head, Chest, Wrist, Ankle}	{Red, Green, Infra}	✔	✔	✔
Ear-PPG [127] (2023)	Very controlled	17	{Ear}	{Red, Green, Infra}	✔	✖	✖
Welltory [128] (2021)	Very controlled	1	{Wrist}	{Red, Green, Blue}	✖	✖	✖
DaLiA [148] (2019)	Controlled	36	{Wrist}	{Green}	✔	✔	✖
BAMI [104] (2019)	Very controlled	10	{Wrist}	{Red, Green, Infra}	✖	✖	✖
WESAD [159] (2018)	Controlled	25	{Finger}	{Green}	✔	✖	✖
BIDMC [135] (2017)	Very controlled	7	{Wrist}	—	✖	✖	✖
PPGMotion [93] (2017)	Very controlled	1.5	{Wrist}	{Green}	✖	✖	✖
IEEE SPC [203] (2015)	Very controlled	2	{Wrist}	{Red, Green}	✖	✖	✖

* *Body-worn locations* refer to the sites on the body where PPG signals are measured.

** *In-the-wild recording* is defined if the data were collected outside of (controlled) lab environments.

cability to wearable devices in real-world scenarios. For example, the IEEE Signal Processing Challenge (SPC) 2015 dataset [203], includes 5-minute PPG recordings from 12 participants while the participants were controlled in the lab environments. This dataset is divided into two sets: a set with fewer motion artifacts from 12 participants (IEEE SPC12) and a set with heavy motions from 10 participants (IEEE SPC10).

The BAMI dataset [104] was collected from 24 participants and includes resting, walking, and running activities. The experiments consisted of 2 minutes of walking as a warm-up, 3 minutes of running, 2 minutes of walking, 3 minutes of running, and 2 minutes of walking to cool down, all on a treadmill inside lab environments where even the speed was controlled (6.0 to 7.0 km/h). Similarly, the PPG-Motion dataset contains walking, running, and cycling activities of 8 participants with an average recording duration of 17 minutes [93].

Meanwhile, the WESAD dataset [203] consists of approximately 1.6 hours of recordings for each of its 15 participants while they were seated and/or standing in a lab environment. It does not contain any physical activities but includes stress and amusement stimuli as well as meditation. All aforementioned datasets were recorded in controlled indoor conditions.

To record data that is more representative of real-world conditions, i.e., outside of lab environments, the DaLiA dataset [148] includes not only indoor but also outdoor activities such as cycling, walking, and driving. It contains recordings of 15 participants with a total of 26 hours of PPG data. Although this dataset is the largest among existing datasets, it is approximately eight times smaller than our presented dataset. Moreover, while the activities in DaLiA are controlled outside of lab environments, our experiments were conducted in com-

Figure 15: WildPPG participants engaged in multiple forms of travel as well as indoor and outdoor activities with changing environmental conditions. No strict study protocol was enforced and participants completed the activities at their own preferred speed.

pletely free-living conditions, offering a more realistic representation of everyday activities. In other words, our dataset was collected "in the wild" —participants were not controlled and were free to move or engage in activities of their choice (within the framework of the experimental protocol) outside of the lab environment. It includes 216 hours of data, making it the largest dataset of its kind while including signals from multiple body locations with multiple wavelengths. Moreover, it is multi-modal. Covering the varying altitude-temperature changes, ensuring comprehensive and complete coverage of real-life conditions. Table 4 qualitatively compares WildPPG with the existing PPG datasets.

5.3 WILDPPG DATASET

The purpose of WildPPG is to enable the development of more reliable PPG-based algorithms that improve the robustness of estimating HR as well as potential other cardiovascular metrics that are of interest on wearable devices such as smartwatches or fitness trackers. Our data recording methodology was designed to capture real-world data under less-than-optimal conditions during representative outdoor activities and everyday conditions.

To achieve this, WildPPG contains long and uninterrupted recordings that our apparatus recorded from participants outside controlled environments or procedures. While our primary focus was on recording the signals crucial for HR estimation—PPG time series from a reflective green light-based sensing design across multiple body locations as well as chest ECG-based cardiac activity measurements for reference—our apparatus additionally recorded motion data from an inertial measurement unit at each on-body site, supplementary PPG signals from red and infrared reflections, as well as temperature readings and barometric altitude measurements throughout the recording for each participant.

Table 4 qualitatively compares our dataset's characteristics with commonly used datasets for HR prediction from blood volume pulse signals. We specifically focus on comparing our dataset with those collected under conditions similar to real-world conditions, where participants moved freely. BIDMC is an exception, as it was collected with critically ill adults (a subset of the MIMIC-II dataset), and we include it because of its use in learning-based HR prediction in previous work (e.g., [156]). As shown in Table 4, our dataset is ~ 8× bigger than existing datasets for HR prediction and includes multi-modal data collected under real-world conditions. Below, we detail the experimental protocol and the design of our data capture apparatus for WildPPG.

5.3.1 *Dataset Design*

5.3.1.1 *Experimental Protocol*

Participants gathered in the morning to start the study. An experimenter outfitted each participant with four recording devices and ensured PPG and ECG signal quality (20 min). The participants then took a minivan from Zurich to Grindelwald (140 min) and transitioned to a cable car and train to Jungfraujoch railway station at 3460 m above sea level (80 min). Subsequently, they spent 5 hours in public places in and around the station and engaged in various activities. While the stay at the station included the same main activities for all participants, no strict study protocol was enforced in an effort to capture real-world data while participants completed the activities at their own speed of preference. Participants were encouraged to not constrain themselves due to the study protocol and engage in social interactions, take pictures, buy and consume snacks and beverages, take restroom breaks as needed, or sit down throughout the stay.

As shown in Figure 15, the main activities included walking through the museum and exhibition area including an ice cave with below-freezing temperatures (≈ 60 min), taking a 110-meter-high elevator and walking the remaining stairs up to the observatory and outside viewing platform (≈ 60 min), sitting down for lunch (≈ 60 min), walking through the snow-covered outside area (≈ 60 min), and resting inside (≈ 60 min). Spread out throughout the stay at Jungfraujoch, the participants additionally climbed and descended a set of stairs across four floors at maximum endurable pace on three to four occasions. After the stay at the station, participants took the train and cable car back to Grindelwald (80 min) and returned to Zurich on the minivan (140 min). Upon arrival, the experimenter removed the recording devices from each participant. Throughout the study, participants were instructed to remove the device on the wrist when washing their hands. All other devices were worn continuously.

5.3.1.2 *Sensors*

The signals collected in WildPPG were acquired using the sensing platform introduced in Chapter 3. PPG measurements were obtained using the MAX86141 optical analog front-end at 128 Hz connected to the SFH7072 optical module with a green (530 nm), a red (660 nm), and an infrared (950 nm) LED as well as a broadband photodiode (410 – 1100 nm), and an infrared-cut photodiode (402 – 694 nm). Green and red PPG were acquired in combination with the infrared-cut photodiode and infrared PPG with the broadband photodiode. Accelerometer data was acquired at a sampling rate of 200 Hz using the LIS2DH MEMS digital motion sensor. For ground truth, the sternum device additionally collected the Lead I ECG at 128 Hz through the MAX30003 biopotential sensor that connected to gel electrodes placed on the chest. Temperature and barometric altitude measurements were collected at a sampling rate of 10 Hz using the BME280 environmental sensor within the device.

Sensor data was acquired synchronized within each device as described in Section 3.2.3. However, measurements across devices are still participant to clock skew and drift effects – most notably when comparing PPG measurements from the wrist, ankle, and head devices with ground truth ECG measurements acquired on the chest. With the given quartz tolerance, although low, two devices may experience sampling shifts of up to 2 seconds across a 14-hour continuous recording session. Therefore, sensor data across devices was synchronized by aligning recorded signals offline following the approach described in Chapter 4 (33 ms accuracy).

5.3.2 *Recruitment and Recording*

WildPPG's participants were 16 healthy adults, recruited on a voluntary basis without compensation beyond the coverage of all expenses incurred throughout the recording of the data and the visit of the mountain station. 7 female and 9 male participants took part, with ages ranging between 22 and 69 (mean age: 39). All participants' skin tones were within 1–3 on the Fitzpatrick scale [70]. All participants signed a consent form prior to participation. The experimenters conducted the data collection with groups of 3–5 participants, and the protocol was identical for all participant groups.

At the beginning of the recording, participants were introduced to the experimental protocol and the goal of the study. An experimenter was always with the participants throughout the recording. The recording devices worked without requiring user input, operating in a completely passive manner. The apparatus provided no feedback of any sort to participants either.

5.3.2.1 *Risks*

Participants were informed about the risk of developing symptoms of Acute Mountain Sickness (AMS) during the visit to Jungfraujoch station due to its altitude. Trains leaving to lower altitudes were available

Table 5: The synchronized time series captured in WildPPG.

Time Series	Description	depth [bits]	fps [Hz]	head	sternum	wrist	ankle
`ppg_g`	Green PPG	19	128	×	×	×	×
`ppg_ir`	Infrared PPG	19	128	×	×	×	×
`ppg_r`	Red PPG	19	128	×	×	×	×
`ecg`	Lead I ECG	18	128		×		
`accel_x`	Accelerometer X-axis	10	128	×	×	×	×
`accel_y`	Accelerometer Y-axis	10	128	×	×	×	×
`accel_z`	Accelerometer Z-axis	10	128	×	×	×	×
`altitude`	Barometric altitude	23	0.5	×	×	×	×
`temp`	Temperature	16	0.5	×	×	×	×

every 30 minutes in case a participant was feeling unwell. Additionally, a professional first aid station is available at Jungfraujoch alongside trained personnel and medical assistance if needed. The study protocol and risk assessment were approved by the ethics committee of ETH Zürich (EK 2022-N-44).

5.3.3 *Dataset Composition*

Participants' recordings are limited to the time they actually wore the sensor devices. PPG and ECG recordings are available as continuous recordings of raw sensor data. Accelerometer recordings were downsampled and synchronized to match the 128 Hz of the PPG and ECG signals. Device temperature measurements were averaged with an 8-second sliding window (2-second step size).

The conversion from barometric pressure measurements to altitude above sea level was calculated following the conversion formula in the sensor's datasheet and a normed barometric pressure measurement from the same day acquired by a weather station at Jungfraujoch station run by the Swiss Federal Office of Meteorology and Climatology [87]. Barometric altitude measurements were averaged across all 4 devices worn by a participant and averaged with an 8-second sliding window and 2-second step size. Overall, WildPPG contains 216 hours of synchronized recordings. Accounting for the four separate device locations, this corresponds to 864 hours of PPG recordings at 3 different wavelengths. An overview of all captured data is shown in Table 5.

5.3.3.1 *Ground Truth*

As manual annotation of a dataset of this size is impractical and unreliable, the dataset contains Lead I-based ECG recordings for ground truth reference. For baseline analysis, we detected R-peaks in the

ECG signal, using Pan-Tompkins [130]. To reduce the risk of corrupt ground truth values, peaks with Inter-Beat Intervals (IBIs) corresponding to HR values greater than 185 bpm or less than 35 bpm were removed, and for HR calculation, only IBIs that were part of a sequence of 4 IBIs for which $\mathrm{IBI}_{min}/\mathrm{IBI}_{max} < 0.75$ were considered. In data windows with less than two remaining IBI, no ground truth HR is computed and the window is omitted from baseline methods. Across the whole dataset, this applies to 2640 8-second windows (2.7% of total) mainly due to three participants with partially noisy ECG recordings which are responsible for 2164 of the rejected windows.

5.4 BASELINES

We computed 6 heuristic and 5 supervised baseline algorithms as well as our own method on WildPPG. To improve comparability, all baselines were computed on data recorded by wrist-worn devices. All results are shown in Table 6 and the methods are described in more detail in the following two subsections.

5.4.1 *Heuristic Methods*

MSPTD is a peak detection method suited for PPG signals that computes a local maxima scalogram (LMS), a matrix that contains information about local maxima computed at different scales by comparing samples to neighbors at varying distances in the signal. The concept was first published by Scholkmann et al [160] and expanded upon by Bishop and Ercole by combining maxima detection with minima detection and improvements in algorithm efficiency [24]. The algorithm does not require any parameter tuning and works for any periodic or quasi-periodic signal.

qppg is another peak detection algorithm built specifically for PPG signals and works by integrating the positive slope between diastolic and systolic points in the signal. Peaks are detected using adaptive thresholding [184]. HeartPy detects peaks in the PPG signal where it crosses its own moving average plus a variable offset. Multiple such offsets are dynamically tested and the one chosen which produces the most regularly spaced detected beats [77]. PWD is another PPG-specific method that relies on the detection of beats based on zero-crossings in the first derivative of the signal [107]. MSPTD, HeartPy, qppg, and PWD methods were computed using the MATLAB implementation by Charlton et al. which is released under GPL-3.0 open-source license [40]. We also included traditional Fourier transformation and autocorrelation functions, where both are used to detect periodicity [157].

Table 6: HR detection performance comparison of baselines with prior works in datasets

Method	WildPPG			SPC12*			DaLiA		
	MAE↓	RMSE↓	ρ↑	MAE↓	RMSE↓	ρ↑	MAE↓	RMSE↓	ρ↑
Heuristic									
FFT	17.62	28.68	12.11	14.83	25.81	35.15	34.98	47.13	2.72
Autocorr.	28.43	35.23	-3.06	41.16	49.81	3.945	30.59	39.09	9.18
HeartPy	18.49	30.15	14.7	19.89	24.81	20.19	12.25	18.15	74.97
MSPTD	13.30	21.19	21.37	20.35	25.31	18.22	17.20	25.11	53.95
PWD	11.72	19.57	25.19	21.16	26.32	18.75	28.26	37.48	29.51
qppgfast	19.32	29.94	21.27	17.97	21.99	32.72	13.31	20.35	66.58
Self-supervised									
SimCLR	15.75	19.81	8.12	12.42	20.96	60.41	12.01	19.46	58.31
NNCLR	14.46	18.43	11.54	13.14	18.86	69.82	12.94	20.02	51.12
BYOL	14.11	18.42	12.50	18.71	25.01	48.82	11.67	17.57	63.96
TS-TCC	12.64	17.72	18.71	11.56	18.04	68.38	8.12	14.89	67.13
TS2Vec	10.55	16.52	26.31	9.75	17.82	75.43	10.83	17.89	60.10
VICReg	15.37	19.20	9.30	13.17	20.38	59.76	14.90	21.94	45.38
Barlow Twins	16.14	20.21	9.13	13.22	20.42	64.51	18.26	23.41	23.42
Unsupervised									
UPD	23.12	25.10	7.39	9.30	16.50	77.60	27.41	31.26	18.12
Supervised									
1D ResNet	<u>8.62</u>	<u>15.00</u>	<u>42.41</u>	**5.50**	**11.16**	**84.86**	<u>4.47</u>	<u>10.03</u>	<u>85.87</u>
DCL	8.64	14.73	41.39	16.50	21.61	19.90	26.34	25.53	80.57
FCN	9.79	15.62	35.68	27.55	31.18	40.41	6.55	11.24	83.21
LSTM	9.28	15.14	36.96	18.66	25.02	56.23	5.30	11.10	78.99
Transformer	10.06	16.23	32.07	22.00	27.30	52.86	7.84	15.06	68.62
Temp-ResNet	**8.37**	**14.72**	**45.50**	—	—	—	**4.40**	**9.87**	**85.90**

* The dataset includes no temperature measurements from the wearable.

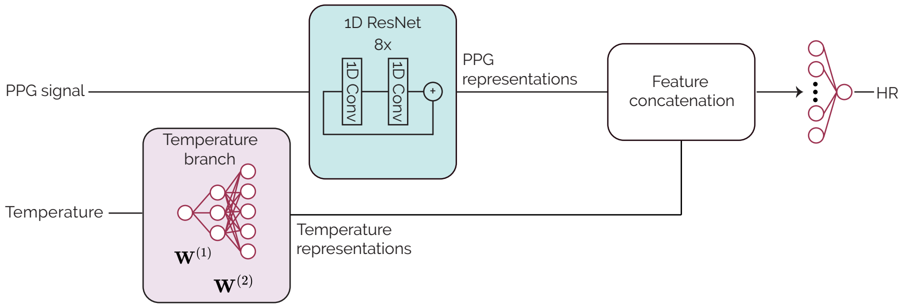

Figure 16: Architecture of Temp-ResNet. The ResNet includes batch normalization and ReLU activations after each convolution.

5.4.2 *Supervised Methods*

For each supervised baseline, we follow the original implementation of the architectures. Additionally, we search for the best hyperparameters (learning rate and its scheduler, batch size, weight decay). Specifically, we compared our method with 1D ResNet [88] which is a modified version of ResNet architecture [83] for time series with 1D filters. DCL [141] is a widely used network that combines convolutional and long-short-term memory units for HR prediction from PPG signals [26, 141]. Fully Convolutional Network (FCN) which only include convolutional blocks with ReLU non-linearities. LSTM [85] model which is designed to capture temporal dependencies in time series data through its memory cell architecture. Transformer [183] is a widely adopted deep learning model that employs self-attention mechanisms to capture long-range dependencies.

5.4.2.1 *Proposed Method: Temp-ResNet*

Considering the relation between the temperature of the measurement site and the SNR of PPG signals, we modified the 1D ResNet model by incorporating a small branch that inputs the temperature value of the segment to the overall architecture for HR estimation. Specifically, we use a Multilayer Perceptron (MLP) with one hidden layer to obtain the representation for temperature $\mathbf{z}_{\text{temp}} = \mathbf{W}^{(2)}\sigma(\mathbf{W}^{(1)}\mathbf{x}_{\text{temp}})$ where σ and $\mathbf{x}_{\text{temp}}$ are a ReLU nonlinearity and temperature value of the specific segment, respectively. The weight matrices $\mathbf{W}^{(1)}$ and $\mathbf{W}^{(2)}$ are learnable parameters of the MLP, responsible for mapping the input temperature value to a higher-dimensional space. Initially, the temperature value $\mathbf{x}_{\text{temp}}$ is expanded to a 5-dimensional vector through the first layer, $\mathbf{W}^{(1)} \in \mathbb{R}^{5\times 1}$. This 5-dimensional representation is then further expanded to a 10-dimensional vector by the second layer, $\mathbf{W}^{(2)} \in \mathbb{R}^{5\times 10}$. Then, the obtained representation is concatenated with the extracted PPG representations and fed to the final linear layer for HR estimation. The architecture of Temp-ResNet is shown in Figure 16.

5.5 LIMITATIONS

Regarding skin pigmentation, all participants were within 1–3 on the Fitzpatrick scale, and no participants rated themselves 4–6 (darker skin tones). Considering the more challenging nature of PPG measurements among people with darker skin tones [21, 66], it will be important to broaden data collection in future efforts. More generally, including 16 participants is a limitation of our dataset, and including more would allow covering as many participant-specific effects as possible. Lastly, as described in Section 5.3.3.1, 2.7% of the total dataset is affected by noisy ground truth measurements.

5.6 ETHICAL CONSIDERATIONS

Working with real-world physiological and activity data from humans requires ethical considerations. All participants were informed regarding safety and privacy both verbally and in writing. They signed a consent form that they received prior to the study for careful reading. No personal information is released with the dataset, the data is fully anonymous and does not contain personally identifiable information or offensive content. To the best of our knowledge, no potential harm or malicious use is possible using the WildPPG dataset and the provided methods. The study was approved by the ethics committee of ETH Zürich (EK 2022-N-44).

5.7 CONCLUSION

As part of this dissertation, we introduced WildPPG, a real-world dataset of long and continuous multimodal motion and physiology recordings for PPG-based HR analysis. Our experimental protocol included a wide variety of activities and environmental conditions including indoor and outdoor activities at different ambient temperatures, altitude levels, and light conditions. Participants moved, acted, and interacted naturally throughout the day and recording without being bound to a strict study procedure—for the purpose of capturing representative data that resembles the measurements taken by wearable devices in *everyday* use. Besides PPG measurements at the three common red, green, and infrared wavelengths and ECG ground-truth references, WildPPG includes acceleration, temperature, and barometric altitude measurements to suit a wide range of analysis methods and tasks. Our study recorded and synchronized measurements from multiple body locations to support exploring future estimation techniques that exceed the analysis of signals from wrist-worn devices most common on smartwatches today and to capture cardiac activity more holistically and robustly in the future.

To the best of our knowledge, WildPPG is the largest available dataset of its kind and can enable learning pipelines that benefit from longer time series during training, additional modalities, and real-world data ("in the wild"). Beyond the dataset, WildPPG includes a range of baseline algorithms—heuristic as well as supervised—and our novel method that leverages temperature readings for improved HR estimation. These methods can serve as a starting point for future work and benchmarks and our intention is to enable future work to better tailor methods to real-world situations and applications in less-than-optimal conditions. This could include multi-modal processing or analyses of additional cardiovascular metrics that are of interest in wearable devices.

6 ROBUST HEART RATE DETECTION VIA MULTI-SITE PHOTOPLETHYSMOGRAPHY

Smartwatches have become popular for monitoring physiological parameters outside clinical settings. Using reflective PPG sensors, such watches can non-invasively estimate HR in everyday environments and throughout a patient's day. However, achieving consistently high accuracy remains challenging, particularly during moments of increased motion or due to varying device placement. In this chapter, we introduce a novel sensor fusion method for estimating HR that flexibly combines samples from multiple PPG sensors placed across the patient's body, including wrist, ankle, head, and sternum (chest). Our method first estimates signal quality across all inputs to dynamically integrate them into a joint and robust PPG signal for HR estimation. We evaluate our method on a novel dataset of PPG and ECG recordings from 14 participants who engaged in real-world activities outside the laboratory over the course of a whole day. Our method achieves a mean HR error of 2.4 bpm, which is 46% lower than the mean error of the best-performing single device (4.4 bpm, head).

Code related to this chapter is available on github:
https://github.com/eth-siplab/MultiSite-PPG-Robust_HeartRate

6.1 OVERVIEW

Wearable devices, smartwatches in particular, have gained popularity for monitoring physiological parameters, most commonly HR. For this, devices typically adopt reflective PPG to optically and thus non-invasively sense blood volume pulses for HR estimation.

However, achieving high accuracy in optical HR estimations remains challenging [18], especially in settings that lead to motion artifacts [117]. Amongst athletes, who subject such sensors to considerable motion, the use of ECG-based HR monitoring is prevalent, where the electrophysiological measurements with a chest strap avoid high measurement errors. Since chest straps are more obtrusive than individual PPG sensors, researchers have sought ways to compensate for motion artifacts in PPG. Approaches range from integrating accelerometers to detect motion [178], leveraging multiple light sources [2] or photodiodes [192], which can increase HR accuracy [105, 132]. Many of today's consumer smartwatches thus embed multiple LEDs and photodiodes (e.g., Apple Watch Series 9, Fitbit Versa 4, Garmin Fenix 7).

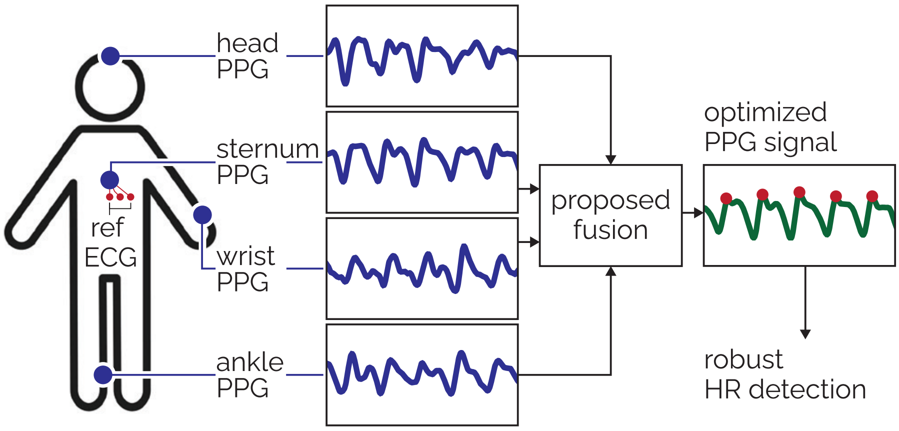

Figure 17: Our method dynamically weighs and fuses PPG signals across sensors at different body locations based on estimated signal quality, producing a continuous PPG output signal as the basis for robust HR estimation.

In this chapter, we improve HR estimation through a novel multi-site PPG signal fusion method. We build on the insight that the motion of the wearer's body affecting PPG readings is often specific to a device's location on the body, such that motion artifacts manifest to different extents in the devices worn across the body [166]. Our proposed method dynamically combines PPG signals from sensors worn across the body following our window-based signal quality estimation. For evaluation, we used the recordings of 14 participants from the dataset introduced in Chapter 5. Our signal fusion method robustly estimates each participant's HR, achieving a 46% lower error on the ground-truth ECG-based HR levels than the best PPG-based HR, which originated from the head-worn sensor.

6.2 METHODS

We propose a novel method that dynamically fuses PPG signals from multiple sensors distributed across a wearer's body based on their temporal and intermittent signal characteristics. Our method first estimates these based on assessing individual signal quality and then weighs segments to derive the combined output signal. Figure 18 illustrates our method.

6.2.1 *PPG Peak Detection and HR Measurement*

For each PPG signal input, our method first detects segments by identifying systolic peaks in each data stream. We first apply a bandpass filter to the input (Butterworth, 0.6-3.3 Hz passband) and then detect peaks whenever a PPG signal crosses its moving average plus an offset, building on van Gent et al. [77]. We determine the offset by mini-

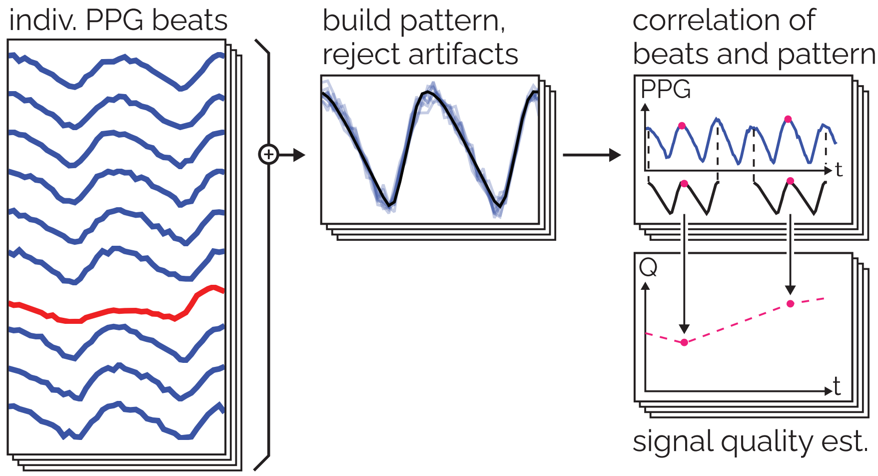

Figure 18: Our method estimates the quality of signal segments by deriving a normalized aggregate and correlating it with a leaning triangle wave to reject bad input signals during template formation. This process is separate for each sensor location to account for local differences in signal morphology.

mizing the variance in the resulting peak intervals over a one-minute window. Peaks are removed if the intervals between neighbors would result in a HR greater than 185 bpm.

Similar to the approach described in Section 5.3.3.1, we only process IBI that lie within at least five consecutive IBIs where $\frac{\min \text{IBI}}{\max \text{IBI}} > 0.51$. This reduces the impact of spurious peaks on To reduce the impact of spurious peaks on HR computation. This filtering step is relatively lenient and retains most PPG beats, even less accurate sections that more conservative filters would remove. Our method explicitly preserves these segments to ensure the presence of signals over the full time of captured data across sensors and because they are input into our fusion.

Finally, HR values are estimated from the IBI over a 30-second moving window over the data.

6.2.2 *Signal Quality Estimation*

To assess a PPG segment's signal quality on a per-beat basis, our method matches them against a template, following the approach by Warren et al. [192] as shown in Figure 18. We perform this separately for each signal to account for differences in signal morphology depending on the sensor's location on the body. For template formation, at each detected peak in the PPG signal, a section is used that is delimited by its neighboring peaks. The template therefore spans two R-R intervals, increasing the likelihood of rejecting corrupt signal segments. Sections with lengths corresponding to a HR < 40 or > 185 are ignored. Each resulting section is down-sampled to 40 samples, which compensates for changes in HR, and finally z-scored (i.e., standardized to zero mean and unit standard deviation).

We then derive the template by averaging those sections that sufficiently correlate with a leaning triangle wave to exclude corrupted segments ($r > .8$). To further lower the impact of noisy waves, we iteratively remove those segments with the lowest correlation to the template until 500 remain.

Finally, we estimate the quality of a PPG signal segment at a given systolic peak by correlating it with the template, delimiting a segment by its neighboring peaks.

6.2.3 *Derive Robust Output Signal from Fused PPG Signals*

Before integrating all signals using their individual segments, we align cross-sensor PPG segments to account for delays in blood pulse propagation through the wearer's arterial tree. Within a window of ± 150 ms, we temporally offset signals to align the previously detected systolic peaks across two segments. This window is thereby small enough to prevent misaligning heartbeats even at a HR of 180.

Next, we estimate the signal quality of PPG signals for 30-second windows of the input as the mean correlation of all contained segments with the template as outlined above. This quality estimate is then interpolated between windows to avoid abrupt changes in quality estimates, which could lead to unwanted artifacts in the fused signal.

We obtain the fused signal $s[t]$ from n PPG input signals:

$$s[t] = \frac{\sum_{i=0}^{n} ppg_i[t] \cdot \max(\delta, q_i[t])^6}{\sum_{j=0}^{n} \max(\delta, q_j[t])^6},$$

where $0 < \delta \ll 1$ bounds the quality estimate $q[t]$ and ensures positive values in the interval (0, 1], as negative correlations with the template are insignificant. Our method uses a power of six to amplify better signals in the output and account for minute differences in correlation values despite stark differences in signal quality. We normalize the weights to ensure that output amplitudes remain independent of the number of PPG input signals and their quality estimates.

6.2.4 *Experimental Protocol and Dataset*

We evaluated our method on captured PPG signals of 14 participants from the WildPPG dataset introduced in Chapter 5. The Lead I ECG acquired by the sternum device served as ground truth reference. For each of the 14 participants, ~13 hours of signals from 4 PPG sensors per participant were considered. In total, the analysis is based on 182 hours of synchronized signals and HR values.

configuration	sensor sources: head	sternum	wrist	ankle	error in bpm: mean (std)	median (std)
head	×				4.38 (1.57)	1.98 (1.16)
sternum		×			5.57 (3.35)	2.93 (2.62)
wrist			×		8.22 (3.40)	5.75 (3.19)
ankle				×	7.63 (4.02)	4.62 (4.00)
ICA						
wearable (2)	×		×		5.04 (2.22)	1.96 (1.01)
wearable (3)	×		×	×	5.72 (2.51)	1.75 (0.94)
all	×	×	×	×	5.01 (2.58)	1.53 (0.77)
Ours						
wearable (2)	×		×		3.52 (1.90)	1.47 (1.30)
wearable (3)	×		×	×	3.01 (1.44)	0.99 (1.04)
all	×	×	×	×	**2.37 (1.14)**	**0.53 (0.48)**

Table 7: Errors in HR calculated from PPG signal combinations.

6.2.5 *Processing*

Ground-truth HR was extracted from 30-second windows in the Lead I reference signal, moving in steps of 5 s and using Pan-Tompkins [130] for R peak detection. PPG-based HR was computed as outlined above for each device (i.e., signal source) separately, and the fused PPG signal following our method was computed per person as described above. As a baseline, ICA was applied to the PPG signals, and HR was extracted from the best output following related work [105].

6.3 RESULTS

Across all participants, the HR values based on the single PPG signal recorded from the head was most accurate (mean error 4.38 bpm), followed by the sternum (5.57 bpm), ankle (7.63 bpm), and wrist (8.22 bpm). Error magnitudes thereby plausibly correspond to the amount of motion experienced by each device and thus body site during the outside activities.

We evaluated the performance of the proposed fusion method and ICA in three configurations as shown in Table 7 across four, three, and two devices, respectively.

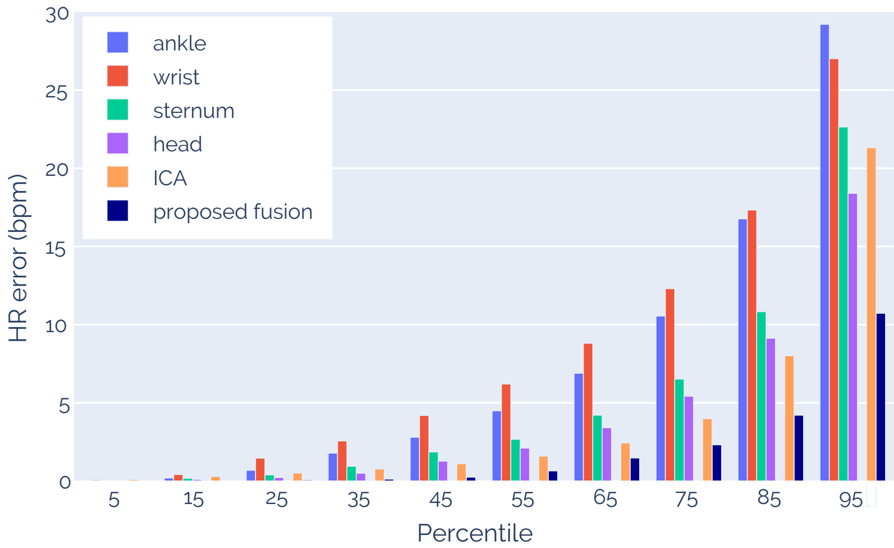

Figure 19: Percentiles of HR segments sorted by error, comparing mean error by device placements and aggregation method across all 14 participants.

6.3.1 *Cross-sensor Signal Fusion*

The proposed signal fusion method overall produced HR values with the lowest mean and median error. All three combinations resulted in a lower mean and median HR error than obtaining HR values from any of the sensor devices alone.

Compared to the HR calculated from the best-performing single device (i.e., head), the mean error decreased by 20–46%, whereas the median HR error was 26–73% lower. As shown in Figure 19, this substantial reduction in HR error is not just present during phases of high HR error in individual PPG signals (higher percentiles, right); it is also evident when the HR error in the contributing PPG signals is already low (lower percentiles, left), which highlights the potential of our method across signal quality ranges.

Figure 19 additionally juxtaposes the error reduction for multi-site signal integration (fusion, ICA) and single-site PPG-based HR calculation. This further illustrates the usefulness of multi-site PPG for robust HR calculation during motion and activity.

6.3.2 *ICA*

Comparing the best signal output of ICA with the individual device placements, in no setting it produced a lower mean error in HR than using the best-performing input PPG from the head. However, this signal from ICA-based processing did result in HR with a lower *median* error than HR from any single PPG. This suggests that it can be useful for more robust HR detection in most cases, but it also indicates that

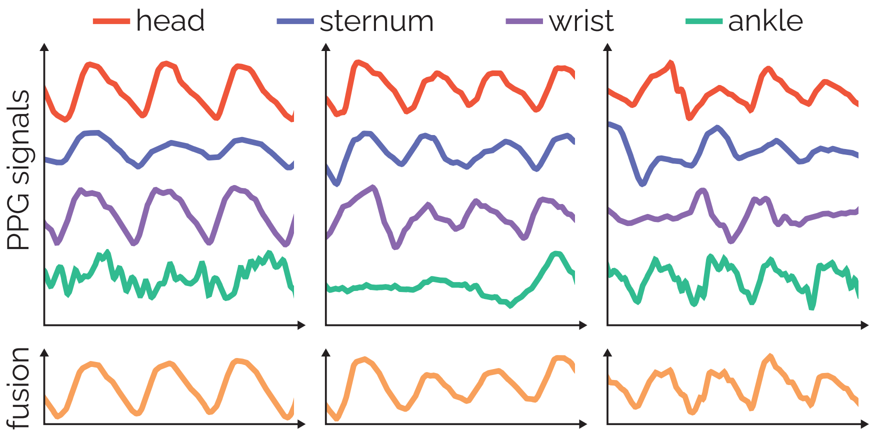

Figure 20: Three examples of four PPG signals from different body locations and the fused output signal produced by our method.

its output is more severely corrupt with higher-error HR calculation when ICA does not output a good PPG signal.

6.4 CONCLUSION

We have introduced a novel method for robustly estimating HR from PPG signals captured in real-world settings. Our method takes as input PPG from multiple sensors, distributed across the wearer's body, and dynamically assesses the quality of individual samples before synthesizing a fused signal using weighted input portions. Our method is agnostic to the number of PPG sources or body sites where sensors are worn, and its computational efficiency allows it to scale to a number of distributed sensors. Through dynamic signal integration and synthesis, our method implicitly leverages moments of lower motion and thus signal artifacts for any sensor to represent such portions more dominantly in the fused output.

We evaluate our method on a novel dataset that we captured from 14 participants over 13 hours during outdoor and mountain activities. Four small wearable devices distributed over the body continuously recorded PPG signals, synchronized with a Lead I ECG captured for ground-truth HR calculation for each participant. Our fusion method predicted HR values with a mean error that is 71% lower than PPG-based HR from the wrist (91% lower median error even), and 46% lower mean error than HR from the forehead. The fusion benefits of our method remain even for combinations of wrist+head (e.g., smartwatch+smartglasses), where the mean error is still 20% lower than on just the head (26% lower median error). Figure 20 shows a qualitative assessment of how our method recovers PPG *morphology* with examples from our dataset when individual HR estimation had a high error. We believe that our results can inspire further research on distributed wearable health monitoring systems in the context of body sensor networks and physiological sensing to further advance reliable and continuous physiological monitoring in the wild outside controlled laboratory conditions.

7 LEARNING-BASED FUSION OF MULTI-WAVELENGTH PHOTOPLETHYSMOGRAPHY

Multi-channel PPG sensors have found widespread adoption in wearable devices. Channels thereby serve different functions—whereas green is commonly used for metrics such as heart rate and heart rate variability, red and infrared are commonly used for pulse oximetry. In this chapter, we introduce a novel method that simultaneously fuses multi-channel PPG signals into a single recovered PPG signal that can be input to further processing. Via signal fusion, our learning-based method compensates for the artifacts that affect wavelengths to different extents, such as motion and ambient light changes. We evaluate our method on the multi-channel PPG recordings of 10 participants from the WildPPG dataset introduced in Chapter 5. Using the fusion PPG signal our method recovered, participants' heart rates can be calculated with a mean error of 4.5 bpm (23% lower than from green PPG signals at 5.9 bpm).

7.1 OVERVIEW

In this chapter, we introduce Tri-Spectral PPG, a novel method that dynamically fuses all reflections from a multi-channel PPG signal to recover a single and robust PPG estimate signal. As shown in Figure 21, our method integrates a U-Net convolutional neural network to infer the most likely fused PPG sequence from all input observations. For training and cross-validated evaluation, we introduce a novel dataset of 10 participants who wore a 3-channel reflective PPG sensor and Lead I ECG digitizer for 13 hours during a variety of everyday and motion-afflicted activities outside the laboratory.

Based on the PPG signals recovered with our method, HR detection had a 23% lower error than when computed based on green PPG reflections, which were best among all individual wavelengths.

7.2 METHOD

7.2.1 *Dataset Recording*

For the evaluation of our learning-based method, we used the 13-hour-long recordings of 10 Participants from the WildPPG dataset introduced in Chapter 5. For the analysis, the 3 PPG channels of the chest device were considered and the Lead I ECG served as ground

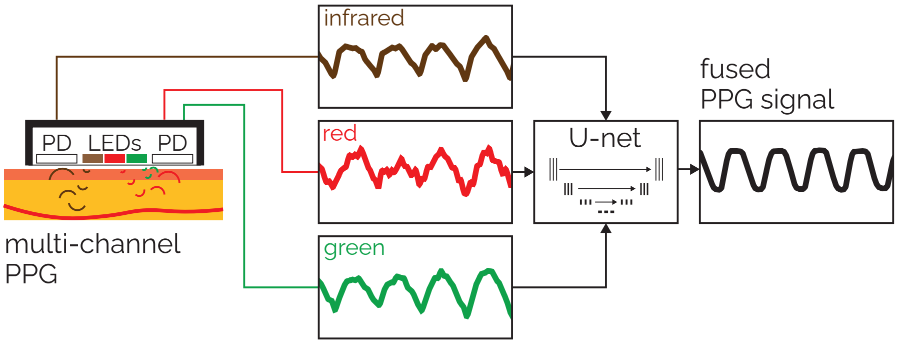

Figure 21: Our learning-based method recovers a single PPG signal from multi-wavelength input signals, thereby restoring signal morphology across them to produce an optimized output signal for subsequent assessment of cardiac dynamics, such as heart rate.

truth reference for HR computations. For each of the 10 participants, ~13 hours of signals per participant were considered. In total, the analysis is based on 130 hours of synchronized signals and HR values.

7.2.2 *Tri-Spectral PPG Signal Fusion Method*

We propose a learning-based signal fusion method that takes synchronized multi-wavelength PPG signals as input and recovers an optimized signal, thereby restoring signal morphology while removing artifacts.

7.2.2.1 *Synthesizing a PPG Reference Signal from Recordings*

To train our model, we synthesized an optimized PPG signal that was true to the recorded morphology through aggregated individual PPG waves. These patterns are resampled and concatenated to align with the R-R intervals of the reference ECG signal, using Pan-Tompkins [130] for R peak detection. We compute the PPG signal patterns following the approach by Warren et al. [192] using ensemble averaging as shown in Figure 22. To account for changes in the signal morphology over time, we compute the pattern in windows of 5 minutes.

For template formation, we split the input PPG signals using the R-peaks from the ECG signal, disregarding sections with lengths corresponding to a HR below 40 or above 185. Each resulting section is down-sampled to 100 samples to account for HR variations and is then z-scored (i.e., standardized to zero mean and unit standard deviation). The template is derived by averaging sections that sufficiently correlate with a leaning triangle wave ($r > .8$), excluding corrupted segments.

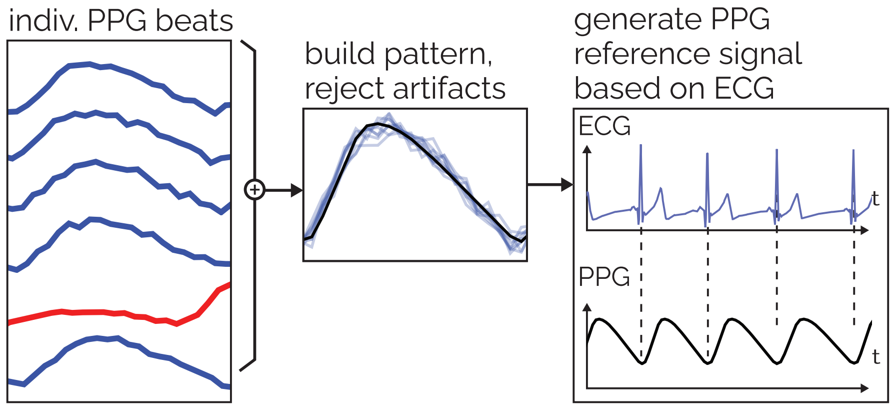

Figure 22: Synthesizing a PPG reference signal to train our learning-based method. A PPG signal pattern is repeated based on ECG R-peaks, generated in windows of 5 minutes and consisting of a normalized aggregate of the input PPG waves, which are correlated with a leaning triangle wave to reject distorted and noisy input signals.

Finally, we synthesize the reference signal by concatenating one signal pattern for each R-R-interval. The concatenated pattern is a weighted average of the two nearest neighboring computed patterns. The weighting is inversely proportional to the distance of the R-R-interval to the corresponding pattern and thus smooth transitions between different patterns are guaranteed. To match the HR and signal length, each concatenated pattern is resampled to match the corresponding interval.

7.2.2.2 *Network Architecture and Training*

For the task of signal recovery, our method adapts a U-Net architecture [175], using a three-channel PPG signal as input and generating a fused PPG signal as output. We augment the model architecture with four up-sampling and four down-sampling layers. The four up-/down- sampling layers are optimized from 32 to 256 kernels, similar to previous work on estimating waveform sources [174]. For training, we used the Adam optimizer [99] with $\beta_1 = 0.9$, $\beta_2 = 0.999$, and a mini-batch size of 80. The learning rate is initialized with 0.001 and reduced by half when the validation loss stops improving for 50 consecutive epochs. The training is terminated when 75 successive epochs show no validation performance improvements. The best model is chosen as the lowest L1 loss on the validation data.

7.2.3 *Evaluation*

To evaluate our method, we extracted 4000 evenly spaced out sections with a length of 8 seconds from the recording of each participant, which corresponds to 8.9 hours of measurements. Besides cropping

Table 8: Performance Metrics for Signal Morphology

Participant	**MAE**	**RMSE**	ρ_{ours}	ρ_G	ρ_R	ρ_{IR}
1	0.356	0.281	0.868	0.652	0.481	0.621
2	0.520	0.547	0.740	0.631	0.476	0.610
3	0.434	0.390	0.800	0.728	0.634	0.673
4	0.517	0.513	0.755	0.593	0.453	0.572
5	0.390	0.289	0.857	0.671	0.501	0.619
6	0.358	0.278	0.874	0.642	0.498	0.654
7	0.346	0.243	0.890	0.638	0.487	0.580
8	0.312	0.331	0.912	0.616	0.476	0.574
9	0.480	0.458	0.775	0.640	0.435	0.628
10	0.390	0.328	0.843	0.569	0.437	0.556
Mean	0.410	0.368	0.831	0.645	0.487	0.608

of measurements during the outfitting and removal of the sensors, sections with inconsistent or noisy ECG measurements were removed in this process while preserving a balanced data set.

To evaluate the fused PPG signal, we compare its validity for HR detection to the input signals and its waveform morphology. This is done for each section of 8 seconds separately. We first apply a bandpass filter to the input (second-order Butterworth, 0.6–3.3 Hz passband) and then detect peaks whenever a PPG signal crosses its moving average plus an offset, building on van Gent et al. [77]. We determine the offset by minimizing the variance in the resulting peak intervals. Peaks are removed if the intervals between neighbors would result in an HR greater than 185 bpm.

During model training and evaluation, we implement five-fold cross-validation across ten subjects. We test two participants in each fold. Additionally, one out of the remaining eight participants is used as validation data for early stopping and model selection. It is worth mentioning that the model is evaluated with subjects that were not used for training.

To reduce the impact of spurious peaks on HR computation, we use the same processing step as described in Section 6.2.1. This means we only process IBIs that lie within at least five consecutive IBI where $\frac{\min IBI}{\max IBI} > 0.51$ resulting in a lenient filtering step that retains most PPG beats ensuring the presence of signals in all segments which is required for the evaluation. Finally, HR values are estimated from the IBI.

Moreover, we calculate the Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and correlation (ρ) between the fused PPG signal

Table 9: HR Detection Mean and Median Error (**bpm!** (**bpm!**))

Part.	green		red		ir		ours	
1	3.4	0.5	9.4	4.8	5.1	0.7	1.9	0.7
2	6.5	1.3	17.2	11.5	11.7	6.0	5.4	1.5
3	2.6	0.3	15.4	10.7	15.2	9.8	2.2	0.7
4	4.3	1.1	11.1	7.9	8.1	5.5	2.0	0.9
5	2.1	0.3	10.7	7.2	8.9	4.7	3.2	0.9
6	8.7	3.7	15.5	10.2	13.6	8.9	5.7	1.6
7	6.2	1.7	10.3	7.6	8.8	5.3	3.6	0.9
8	6.6	3.4	15.7	11.2	12.2	7.9	10.0	6.6
9	7.2	0.5	12.9	8.5	11.5	7.0	3.6	0.9
10	9.0	4.3	11.6	8.4	10.3	6.9	5.2	1.9
Mean	5.9	1.9	13.4	9.2	11.1	6.9	**4.5**	**1.8**

and the reference signal to evaluate the accuracy of our proposed method in approximating the waveform morphology.

7.3 RESULTS

7.3.1 *Signal Morphology*

The PPG signal output by our method during inference exhibited a mean correlation of 0.831 with the synthesized reference PPG signal as shown in Table 8. This is higher than all contributing PPG channels which exhibited a mean correlation of 0.645 (green), 0.487 (red), and 0.608 (infrared).

The MAE of the signal output compared to the reference signal lies within a range of 0.312 and 0.520. The RMSE lies within a range of 0.243 and 0.547.

7.3.2 *HR Detection (as Downstream Task)*

When computing the HR based on the single-wavelength PPG trace, green performed best (MAE 5.9 bpm, median absolute error 1.9 bpm), followed by infrared (11.1 bpm, 6.9 bpm) and red (13.4 bpm, 9.2 bpm) PPG measurements as shown in Table 9. The HR based on the signal produced by our method had a MAE of 4.5 bpm and a median absolute error of 1.8 bpm. This is a 23% improvement of the MAE and a 5% of the median absolute error over the best single channel (green).

7.4 DISCUSSION AND LIMITATIONS

Our method purposely does not directly estimate metrics of cardiac activity, such as HR itself. Instead, our approach recovers the underlying morphology of the observed and possibly noisy PPG signals and can serve as input to a variety of downstream tasks, one of which is HR as used in our evaluation. We believe that our method can serve as input for several more applications of pulse wave analysis in the future. Limiting factors of this work include the small sample size of 10 Participants who all were within 1–3 on the Fitzpatrick scale, with no participants rating themselves 4–6 (darker skin tones). Considering the more challenging nature of PPG measurements among people with darker skin tones [66], it will be important to broaden data collection in future efforts and investigate the effects of different PPG wavelengths and body locations on the results of this approach.

7.5 CONCLUSION

As part of this dissertation, we have introduced Tri-Spectral PPG, a novel method that fuses PPG signals acquired at different wavelengths to retrieve a more robust fused signal which preserves the signal morphology and therefore allows further processing to retrieve arbitrary PPG-based measures. Our method takes as input PPG signals recorded at different wavelengths and leverages a U-Net to produce a single output signal. Compared to prior work, its advantage is the ability to consider information from all PPG signals, as opposed to switching between them, while also removing artifacts completely, as opposed to averaging channels. To train our model, we have generated a synthesized reference signal using an ECG signal and an aggregation method to retrieve PPG wave patterns. To evaluate our method, we contribute a novel dataset that we captured from 10 participants over 13 hours during outdoor and mountain activities. The signal produced by our method correlated better with the reference signal than any contributing signal and the HR detection on the produced signal had an error that was 23–66% lower compared to the contributing signals.

8 WEARABLE DEVICES FOR RESPIRATORY RATE MONITORING

Deviations in respiratory rate often precede abnormalities in other vital signs. However, continuously monitoring respiratory rates outside clinical settings remains challenging due to the obtrusive nature and sensitivity to body motions in existing monitoring approaches. In this chapter, we propose Respiro, a single-point-of-contact wearable device that leverages off-the-shelf, consumer-grade ultra-wideband radar to monitor respiratory rate as part of a chest strap. Our signal processing pipeline reliably extracts the wearer's respiratory signal from windowed complex channel impulse responses. In a controlled experiment, twelve participants performed various activities to evaluate the system's accuracy under motion while capturing ground-truth recordings through a spirometer. Our method extracted respiratory rates with less than 1 breath per minute deviation in 71% of all measurements, averaging 1.11 breaths per minute across all sessions and participants. Our findings underscore the potential of consumer-grade ultra-wideband radar technology in body-worn devices for unobtrusive yet effective respiratory monitoring.

8.1 OVERVIEW

Respiratory rate, along with heart rate, blood pressure, and body temperature, is a primary vital sign routinely assessed by medical professionals. RR measurements are essential in diagnosing conditions such as acidosis or pneumonia [63, 129]. Additionally, an elevated RR over the course of 24-72 hours can predict severe adverse events, including cardiac arrest and the need for intensive care unit admission [47, 69]. In everyday contexts, RR can provide insights into stressors like emotional load, heat, or physical effort [129]. Traditionally, capnometry and spirometry are considered the gold standards for respiration monitoring. However, these methods require air masks, limiting their suitability for long-term monitoring [3, 115]. Alternative methods, such as respiration belts and the use of IMUs, have emerged, but they are less accurate than spirometry or capnography and are prone to motion artifacts.

Amidst these challenges, UWB radar has garnered interest over the past two decades since the FCC opened the 3.1 to 10.6 GHz spectrum for medical imaging [68]. There is substantial research on non-contact RR recovery using UWB radar [32, 89, 188], including through debris [108] or walls [42]. This remote monitoring detects chest wall motion

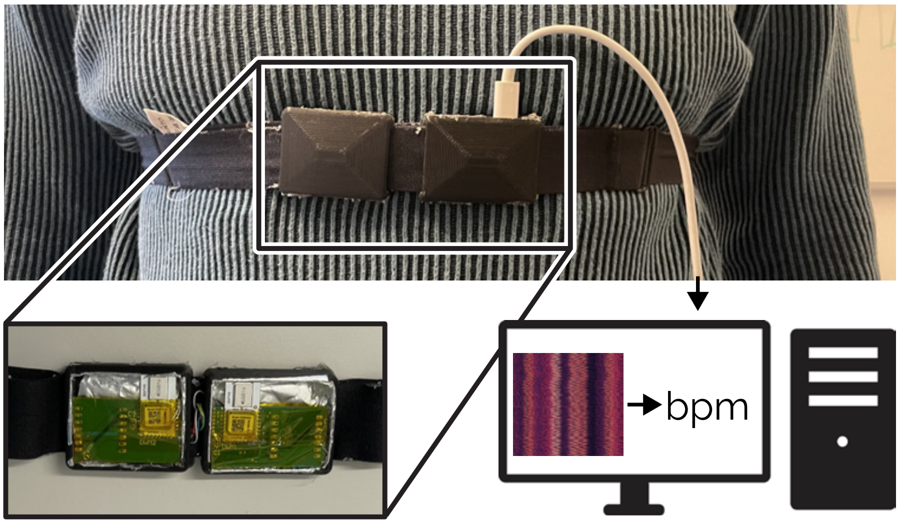

Figure 23: Our wearable device placed on the sternum uses an off-the-shelf ultra-wideband radar module to produce reliable measurements of the respiratory rate.

using UWB beacons in a person's environment. However, these methods are prone to errors due to unrelated body motion and are unsuitable for tracking the RR during activities and in changing environments. UWB radar can also probe the body itself, detecting varying dielectric properties of tissues which affect UWB pulses differently [173]. Specifically, this applies to the lungs, which exhibit different dielectric properties when inflated versus deflated: More signal is reflected at the muscle-lung interface when the lung is inflated compared to its deflated state [134]. This is supported by Cavagnaro et al. who studied the propagation of UWB pulses into human tissue and proposed that effective RR recovery can be achieved by capturing the air-skin reflection signal with an antenna positioned 1 meter away from the body. They suggested that on-body antennas could potentially recover the RR from a small signal reflected by the posterior lung wall [37]. However, on-body UWB radar for RR monitoring has only been studied using laboratory equipment [137, 138], using two points of contact [48, 49], or with custom UWB antennas integrated into a seat [158], all of which were evaluated exclusively on stationary users.

Here, we contribute Respiro, the first single-point-of-contact wearable device for RR measurement as shown in Figure 23, using consumer-grade, off-the-shelf UWB radar modules. We also present an offline data processing pipeline that computes RR from in-body reflections contained in CIR estimates. We evaluated our system in a user study where participants engaged in various activities to induce motion artifacts. This approach aligns with other studies assessing wearable devices for respiratory monitoring [38, 114]. To the best of our knowledge, this is the first study to evaluate a wearable UWB-based RR detection system on moving participants. Additionally,

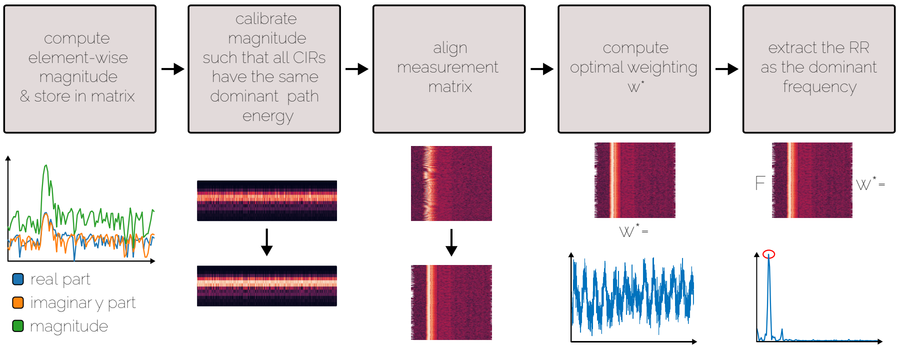

Figure 24: The processing pipeline includes (left to right): Computation of the magnitude of the CIR and storage as a row in a matrix, then the CIRs are calibrated, and aligned. Finally, an optimal weighting to maximize the energy in the relevant frequency bands is calculated, applied, and the dominant frequency computed.

we collected accelerometer data and compare our UWB-based system with the accelerometer-based systems of Bates et al. [17] and Rahman et al. [142].

8.2 METHOD

8.2.1 *Hardware*

Respiro includes two UWB radar modules (DWM3000, Qorvo), the same series as the one integrated into commercial devices such as the Google Pixel 6 Pro. Due to their half-duplex nature, two modules are required as each one cannot transmit and receive simultaneously. The radar modules are controlled by a microcontroller (ESP32-C3, Espressif Systems) through the SPI protocol, configuring them for transmitting and receiving states and acquiring measurement data. This data is then transmitted by the microcontroller to a computer for further processing via serial communication. Each radar module is soldered onto a separate PCB. The two PCBs are enclosed in 3D-printed cases lined with aluminum foil on all surfaces that do not face the wearer, reducing the direct path component of the signal in comparison to in-body reflections.

8.2.2 *Sampling*

Rather than raw reflected pulses, the radar module only provides access to a complex CIR estimate which can be used for RR measurement purposes. The CIR, represented as $h[m]$, elucidates the relationship between a probing signal $x[m]$ and its reflection $y[m]$:

$$y[m] = \sum_{k=0}^{M} h[k]x[m-k] \tag{1}$$

Physically, $h[m]$ illustrates the reflection as if the probing pulse were a Dirac delta function $x[m] = \delta[m]$. M denotes the number of samples in a CIR.

Each transmitted packet begins with an Ipatov preamble, followed by a start-of-frame delimiter. The packet body may include a physical header, followed by a payload, a scrambled timestamp sequence, or both. The radar module estimates a CIR for every packet by exploiting the perfect autocorrelation property of the Ipatov preamble. We transmit a counter value during our measurements to avoid duplicating CIR samples. The CIR preamble spans 1016 complex samples when using the radar module's default pulse repetition frequency of 64 MHz. The sampling period of the CIR equates to half the fundamental period $T_{CIR} = (2 \cdot 499.2\,\text{MHz})^{-1} \approx 1\text{ns}$.

A UWB pulse propagates through the air at a velocity of $v_{air} = \frac{c}{n_{air}} = 29.9\frac{\text{cm}}{\text{ns}}$, where c is the speed of light in a vacuum and $n_{air} \approx 1$ the refractive index of air. Various tissues within the human body have distinct dielectric properties. Prior work has used an average refractive index of $n_{thorax} = \sqrt{50}$ to model the thorax [37, 137, 138]. Under this assumption, the pulse propagation speed within the thorax is $v_{thorax} = \frac{c}{n_{thorax}} = 4.2\frac{\text{cm}}{\text{ns}}$. Consequently, each sample corresponds to a spatial distance of 4.2 cm within the human body. We extracted 100 samples, aligning with previous work where 300 and 75 samples after the direct path were used ([49, 110]). However, the direct path component commences not at index 0 of the extracted CIR: In a random sample of more than 26 000 CIR samples recorded over 14 minutes, we discovered that the average index of the direct path, i.e., the index in the CIR with the highest magnitude, is approximately 741.7 with a standard deviation of 2.4. The lowest observed direct path index was 735, and the highest was 746. Therefore, we sampled 100 indices from the CIR starting at an offset of 720. This sampling strategy ensures over 70 samples to include reflections from within the thorax, corresponding to thorax diameters up to $\frac{70}{2} \cdot 4.2\frac{\text{cm}}{\text{sample}} = 148.3\,\text{cm}$. This reduces memory and communication demands by over 90% compared to sampling the entire CIR while retaining pertinent information.

8.2.3 *Data Processing*

Figure 24 provides an overview of the data processing pipeline of Respiro. Initially, we compute the element-wise magnitude of the complex-valued CIR and store a predefined time window as a row in a measurement matrix $\mathbf{H}$. The rows of matrix $\mathbf{H}$ represent estimated reflections from a Dirac pulse at different temporal instances, while the columns correspond to reflections originating from distinct spatial distances.

Since not all direct path components possess the same energy, we adopted the dominant path approach proposed by Li et al., where they calibrated the CIR of a WIFI system to extract RR [111]. Their method involves computing a calibration coefficient based on the energy of the direct path component, which they extracted from a window around the time domain peak of the CIR. In our configuration, we utilize a 7 ns window to capture the energy of the directed path. The calibrated CIR $h'[m]$ is derived from the CIR $h[m]$ as follows:

$$h'[m] = \left(\frac{1}{2D+1} \sqrt{\sum_{|\tau-\tau^*|\leqslant D} |h[\tau]|^2} \right)^{-1} h[m] \tag{2}$$

Where $D = 3$ is the number of indices in the window left and right of the index of the CIR peak $\tau^* = \underset{\tau}{\operatorname{argmax}} |h[\tau]|$.

To address the issue that the extracted CIRs do not contain their direct path component at the same index every time and the rows of $\mathbf{H}$ are thus not perfectly aligned, we employ cross-correlation maximization to align the rows of $\mathbf{H}$. The initial CIR $h_0[m]$ in $\mathbf{H}$ serves as a reference and each subsequent CIR $h_n[m]$ is shifted by an index k_n to achieve maximum cross-correlation.

$$k_n = \underset{k}{argmax} \left| \sum_{m}^{M} h_0[m]\, h_n[m+k] \right| \tag{3}$$

Various anatomical structures within the human body, such as the anterior and posterior lung walls, act as potential sources for reflections containing RR information. We extract the RR from a linear combination of reflections across all spatial distances, denoted as $\mathbf{H}w$. The weight vector is computed using an optimization criterion introduced by Li et al. [110] when they successfully extracted respiration signals from individuals at unknown positions within a room. They maximized the energy within the band of possible respiration frequencies while maintaining the overall energy constant by computing the optimal weighting w_{opt} as follows:

$$w_{opt} = \underset{w}{\operatorname{argmax}} \, (\mathbf{F_I H}w)^H (\mathbf{F_I H}w) \tag{4}$$

$$\text{s.t. const.} = (\mathbf{FH}w)^H (\mathbf{FH}w) \tag{5}$$

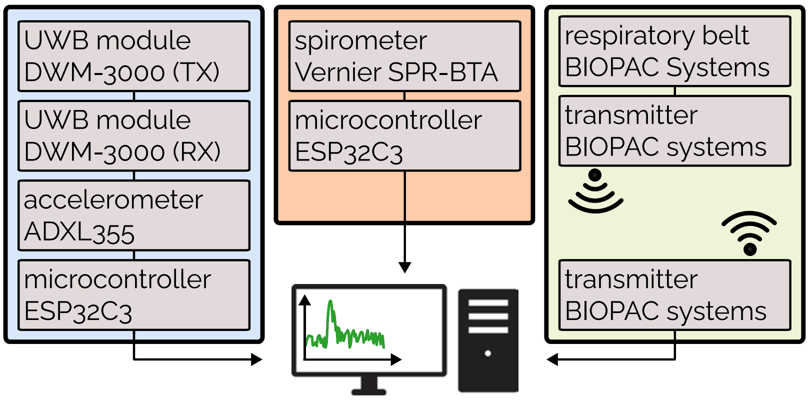

Figure 25: The data capturing system for the user study included the custom UWB device (left), a spirometer ground truth (center), and a respiratory belt (right).

where $\mathbf{F}$ is the discrete Fourier transform (DFT) matrix and $\mathbf{F_I}$ is the DFT matrix containing frequencies in the band of possible respiration frequencies. This leads to a dynamic selection of the measurement depths in the body where respiratory signals occur. We used a band of possible respiration frequencies of $f_r^l = 0.1$ Hz (6 brpm) to $f_r^h = 0.7$ Hz (42 brpm) which enables the detection of respiratory rates both at rest and during physical exercise. Ultimately, we computed the Fourier transform of $\mathbf{H}w_{\text{opt}}$ and identified the respiration frequency as the frequency exhibiting the highest energy density within the band of possible respiration frequencies.

8.3 CONTROLLED EVALUATION

8.3.1 *System Overview*

As shown in Figure 25, the study setup included the custom UWB RR detection device, a spirometer (SPR-BTA, Vernier), a respiration belt (BN-RESP-XDCR, BIOPAC Systems) with peripherals, and a computer for data collection and processing. We integrated an accelerometer (ADXL355, Analog Devices) with the UWB device to enable comparison with accelerometer-based systems proposed by Bates et al. and Rahman et al. [17, 142]. The on-board microcontroller connects to the accelerometer and the UWB radar modules and facilitates data transmission to the computer via USB-C. Similarly, the microcontroller associated with the spirometer samples the analog signal and transmits the acquired data to the computer via USB-C. The respiration belt system features an independent transmitter for wireless data transmission to the base station, which subsequently transmits the data to the computer through WIFI.

As shown in Figure 26, we positioned the UWB device on the participants' sternum to ensure proper placement over the lungs while

Figure 26: The device setup on the participant's thorax for the study included the custom UWB device which was strapped over the sternum, followed by the reference respiratory belt over the abdomen, and the transmitter of the respiratory belt.

minimizing interference with upper body movement. Following the manufacturer's guidelines, we situated the respiration belt over the abdomen. To prevent contact during movement and mitigate data artifacts, we maintained a 1–2 cm gap between the respiration belt and the UWB device. The respiration belt transmitter was placed with a sufficient gap to the respiration belt to avoid contact during participant movement. To prevent accidental sensor detachment during exercise, we secured the USB-C cable of the UWB device to both the supporting belt of the device and the belt of the respiration belt transmitter. Similarly, we affixed the microcontroller interfacing with the spirometer to the transmitter belt and secured its USB-C cable to the transmitter belt.

In the user study, we acquired CIR and accelerometer data at a sampling frequency of 32 Hz, following prior work which acquired data at rates ranging from 19.3 Hz to 65 Hz [49, 110, 158]. The analog signal from the spirometer was sampled at 50 Hz using the built-in analog-to-digital converter on the connected microcontroller. Additionally, the respiration belt signal was acquired at its default frequency of 2 kHz.

8.3.2 *Participants*

The participants were 12 healthy adults, recruited on a voluntary basis. 3 female and 9 male participants took part, with ages ranging between 21 and 29 (mean age: 25).

8.3.3 *Procedure*

Upon arrival, participants were introduced to the experimental protocol. An experimenter outfitted them with the recording devices. The participants then tested the treadmill and the experimenter adjusted the walking and running speeds if necessary, starting at 5 km/h and 10 km/h respectively. For the data acquisition, the participants then completed the following activity sequence twice:

1. two minutes of sitting on a chair
2. two minutes of standing still
3. two minutes of walking on a treadmill
4. two minutes of running on a treadmill
5. one minute break (or more at participant's request)
6. one minute of cycling on an exercise bike
7. one minute break (or more at participant's request)
8. one minute of squats
9. one minute of standing still

For the second iteration, the spirometer was removed.

The experimental procedure was reviewed and approved by the ETH Zurich Ethics Commission under the application reference EK-2023-N-183.

8.4 RESULTS

For analysis, we considered the recordings of the seven activities described in the study procedure, excluding breaks. With 12 participants each performing two iterations of the protocol, this resulted in $12 \cdot 14 = 168$ traces. Two traces were excluded from analysis: one due to a participant unintentionally breaching the study procedure during the sitting sequence, and another due to a respiration belt failure during the standing sequence.

Table 10 shows the aggregated results of the user study categorized by activity. We report the SNR, the RMSE, the Mean Absolute Percentage Error (MAPE), and the share of traces exhibiting an error of less

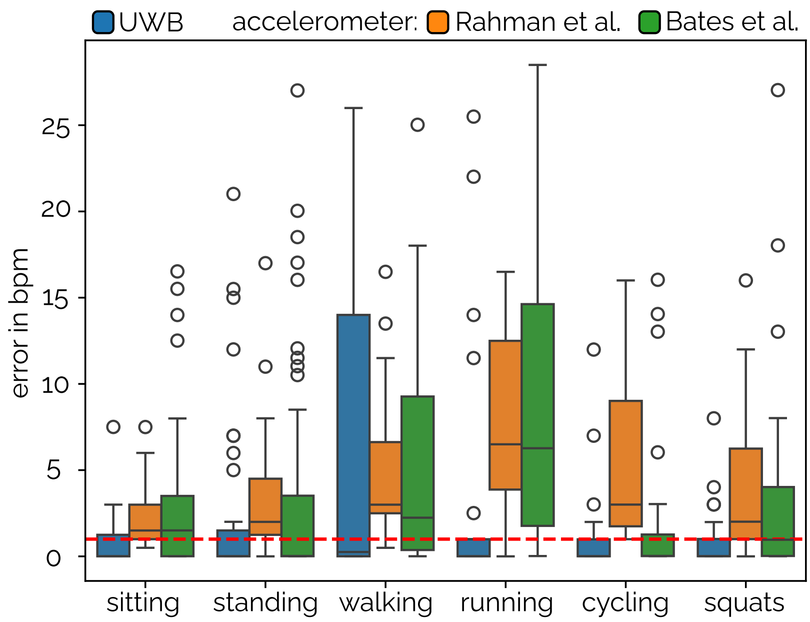

Figure 27: The error distributions of our UWB-based approach to detect respiratory rate compared to the accelerometer-based approaches of Rahman et al. and Bates et al. The red dashed line indicates the error threshold of 1 brpm used to classify successful traces.

than 1 brpm. We compute the SNR of the processed UWB signal, the respiration belt signal, and the spirometer signal using the methodology outlined by Droitcour et al. [58], wherein signal energy is determined as the energy density within a range of six brpm around the dominant frequency, while the remaining energy is considered noise. We report the RMSE and the count of successful traces, i.e., those with an error lower than one brpm, consistent with Culjak et al. and Rahman et al. [49, 142]. Furthermore, Table 10 also incorporates results from the accelerometer-based pipelines from Rahman et al. and Bates et al. [17].

The processed UWB signal generally exhibits an SNR greater than 1, except during the standing, walking, and running activities in the second round of the study. Both the spirometer and respiration belt ground truth signals consistently demonstrate an SNR greater than 1 when averaged across participants. Success rates range between 0.45 and 0.92 across activities which is better than the accelerometer-based approach by Rahman et al. across all activities and at least equal to Bates et al.'s approach in 11 of 14 cases investigated.

Figure 27 illustrates box plots representing the distributions of absolute errors for our UWB-based approach and the accelerometer-based approaches. Across all activities, the median error of our UWB-based system remains below the 1 brpm threshold and the error distribution is narrower than the ones of the accelerometer-based approaches of Bates et al. and Rahman et al. For walking traces, our

	Respiro *(ours)*				Rahman [142]		Bates [17]	
	SNR	MAPE	RMSE	error <1 brpm	RMSE	error <1 brpm	RMSE	error <1 brpm
w/ spirometer as reference								
120 s sitting	3.09	3.96	**1.08**	**0.82**	3.18	0.36	5.28	0.73
120 s standing	1.84	6.83	**2.16**	0.67	3.63	0.33	6.31	**0.75**
120 s walking	1.08	6.47	**4.49**	**0.83**	7.39	0.17	9.49	0.42
120 s running	2.10	0.43	**0.32**	**0.92**	8.84	0.08	11.00	0.33
60 s cycling	2.26	5.08	2.29	0.75	5.04	0.17	**0.65**	**0.83**
60 s squats	1.60	4.51	**2.50**	**0.75**	3.79	0.33	6.24	0.50
60 s standing	1.31	13.63	6.24	**0.83**	**3.95**	0.17	6.05	0.67
w/o spirometer, respiration belt as reference								
120 s sitting	1.75	6.45	**2.47**	**0.67**	2.80	0.42	7.48	0.08
120 s standing	0.80	15.76	6.00	0.45	**3.39**	0.09	9.26	**0.64**
120 s walking	0.49	36.57	15.29	**0.50**	**5.20**	0.17	9.20	0.25
120 s running	0.81	29.34	11.07	**0.58**	**9.25**	0.08	11.87	0.08
60 s cycling	1.16	4.36	**3.52**	**0.83**	8.15	0.25	7.48	0.50
60 s squats	1.13	4.54	**1.61**	**0.67**	6.99	0.33	8.72	**0.67**
60 s standing	1.65	10.19	**5.13**	**0.67**	6.76	0.42	8.21	**0.67**

Table 10: Aggregated user study results, comparing our method to the accelerometer-based methods of Rahman et. al. and Bates et al. RMSE is reported in brpm, MAPE as a percentage, *error < 1 brpm* as a share

UWB system demonstrates a more scattered distribution of absolute errors compared to Bates et al. and Rahman et al.'s systems.

8.4.1 *Effect of the Spirometer*

The study protocol consisted of two identical activity sequences, where a spirometer was used only during the first iteration. Everything else remained the same including the activities and sensor positioning.

To compare the results of both settings, we used the respiration belt signal from both rounds as the ground truth for this analysis. When wearing a spirometer, the MAE decreased by 2.65 brpm (3.76 brpm to 1.11 brpm, median absolute error 0.016 to 0.006), and the success rate showed an increase of 0.20. Additionally, the SNR of the UWB method increased by 70% on average. Computing the same SNR for the respiratory belt reference signal results in an increase of 98% when additionally using a spirometer compared to the procedure iteration without.

8.4.2 *Accuracy of the respiration belt*

During the first round of the user study, we captured signals from the spirometer and the respiration belt. Comparing the respiration belt to the spirometer, it achieves a success rate (< 1 brpm error) of 0.84% overall with values ranging from 0.67 (squats) to 1.00 (running). The RMSE across activities is 3.06, inflated by the walking and squats activities with RMSE of 4.54, and 6.40 respectively.

8.4.3 *Correlation of Success and SNR*

For the UWB-based method, we observe a negative correlation between the SNR and the absolute error with a correlation coefficient of -0.29. This means that traces with a high SNR in the UWB signal tend to have a lower error. Similarly, when discriminating traces at SNR = 1, the mean success rate is 0.86 among traces with SNR > 1 and 0.56 for all others with sample sizes of 84, and 82 respectively.

8.4.4 *Effect of the Window Size*

We utilized the complete duration of a trace, which spans either one or two minutes, for our analyses. Figure 28 depicts how the success rate evolves if we consider windows of shorter duration. We assessed all two-minute traces using sliding window sizes ranging from 2 to 119 seconds (3-second increment), and a step size of 10% of the window size. Across all activities assessed, a trend toward a higher success rate with longer evaluation windows is observable with a steep increase up to a window length of approximately 35 seconds and a flattening curve afterward. The breadth of the confidence interval expands with longer window sizes where the number of available windows for analysis decreases.

8.5 DISCUSSION

As part of this dissertation, we presented Respiro, a single-point-of-contact, wearable UWB radar system to monitor the RR. The method worked not only in static settings but also during activities where it was less prone to motion artifacts compared to methods based on accelerometers. Using UWB modules that are used in commercial devices such as smartphones, we were able to measure the respiratory rate not only in static settings but also during physical activities. By not requiring lab equipment or complex custom-built circuits and antennas, we show the potential of this sensing modality to be used in unobtrusive wearable devices monitoring the RR. The most closely related prior work by Culjak et al. reported a MAPE of 8.75% when analyzing a transmissive UWB signal through the chest while the partic-

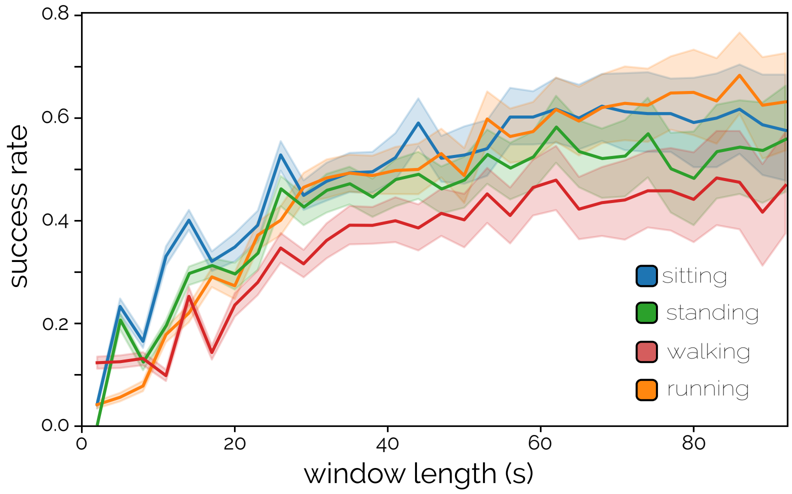

Figure 28: The success rate, i.e., error $\leqslant$ 1 brpm (95% confidence interval), depending on the window size.

ipants were sitting in a chair. In the same setting in our study (sitting participant), we observed a MAPE of 6.45%, highlighting the viability of a wearable single-point-of-contact UWB radar system in monitoring the RR, achieving accuracy comparable to that of a wearable two-point-of-contact system. Compared to the accelerometer-based systems of Rahman et al. and Bates et al., our UWB-based approach has a higher success rate by being less prone to motion artifacts. However, movement of the sensor encasement can still lead to motion artifacts in the measurements. This is particularly evident during walking activities, where the cadence range overlaps with possible respiratory rates. Limitations of our study include the relatively small sample size of 12 participants and the limited age diversity. Future research should include a larger number of participants to investigate participant-specific effects, such as variations in body types, on the results.

8.5.1 *Ground Truth with Spirometer and Respiration Belt*

The acquisition of viable ground truth respiratory rate measurements is challenging. The use of a spirometer introduces resistance to the airflow during breathing and leads to the inhalation of previously exhaled air still in the mask and the device. These effects potentially prompt subjects to engage in deeper breaths, augmenting the outcomes. We observed this effect in our study where the SNR of the respiratory belt was higher when a spirometer was used at the same time indicating deeper breathing patterns. However, the alternative use of a respiration belt for ground truth measurements is more prone to noise due to motion artifacts or movement of the belt. In our study, when comparing the RR from the respiration belt to the results from the spirometer for reference, the respiration belt yielded the correct respiratory rate in 84.3% of all traces (using 1 brpm tolerance). Both these effects lead to a lower accuracy of the RR detection in the sequence where no spirometer was used. Notably, outliers were more severe when not using the spirometer which explains the substantial difference in the reported MAE. Similar observations have been made in prior work by Pittella et al. who observed that a piezoelectric belt measured arm movements rather than respiration frequencies [138]. This underlines the importance of a well-planned ground truth acquisition in future work toward wearable, unobtrusive RR monitoring devices for everyday use.

8.5.2 *Correlation of Success and SNR*

Our processing pipeline recovers the respiration signal from which we compute the RR. The SNR estimate from Droitcour et al. serves as a metric to assess the quality of the recovered respiration signal [58]. Consequently, it is reasonable that traces with a higher SNR in the processed UWB signal correspond to a greater success rate in RR recovery. The SNR could thus be used as a predictive measure to estimate the likelihood of successful RR computation before extracting the RR from the processed UWB signal. By incorporating this into the pipeline, the system's performance could be optimized by selectively utilizing computed RR values only when the SNR exceeds a predefined threshold.

8.5.3 *Effect of the Window Size*

Figure 28 illustrates a higher success rate in RR detection for longer windows of recording data across all activities. This observation is consistent with our expectations, as the optimization problem benefits from a wider reference band. Beyond a window size of 45 seconds, no substantial gains can be observed which may be due to changes in respiratory rate throughout longer window sizes leading to ambiguities.

8.5.4 *Limitations*

All study participants were aged between 21 and 29 years. While our method adapts to different body types by targeting respiration-related frequency bands, it will be important to include people of a greater age diversity in future efforts. Additionally, our evaluation used only IMU-based baseline algorithms, as most contact-based methods require extra hardware, like custom antennas, which could introduce confounding variables.

8.6 CONCLUSION

As part of this dissertation, we presented Respiro, a method to monitor RR using off-the-shelf UWB radar modules as commonly found in devices such as smartphones. Respiro was tested with a custom-built prototype in a study on 12 healthy adults who engaged in multiple activities including walking, running, and cycling which typically produces motion artifacts. When compared with ground truth reference measurements, our method performs on par and sometimes better than more complex measuring methods while being less prone to motion artifacts than accelerometer-based approaches. This showcases the potential for unobtrusive wearable devices that monitor RR in everyday life using UWB radar.

9 CONCLUSION

In this dissertation, we have presented advancements in wearable physiological monitoring outside of controlled environments. Our main focus has been to produce high-quality measurement data and derived values. We introduced a novel method for sensor synchronization, and a data acquisition platform focused on long-term continuous signal recording with high reliability and no user interactions. By integrating hardware innovations with computational methods, this work contributes to the development of continuous, multi-modal, and unobtrusive physiological sensing solutions that improve continuous monitoring of motion, activities, breathing, and the cardiovascular system. This enables applications in precision medicine and more generally health monitoring.

9.1 SUMMARY OF CONTRIBUTIONS

We introduced a multi-device, low-power sensing system designed for passive and long-duration physiological monitoring. The system achieves sample-perfect synchronization across sensors and millisecond-level alignment between devices. To minimize power consumption, enable long battery life, and improve system robustness. The platform is designed to live-stream only as much data as necessary for the application and log full-resolution data in non-volatile memory for later analysis.

To further enhance the reliability of physiological data collection, we proposed a multi-modal offline synchronization method that enables precise alignment of signals across multiple body-worn sensors without requiring communication between the devices. This is valuable in real-world applications where radio communication during runtime cannot always be guaranteed or has to be avoided due to power restrictions.

With WildPPG, we contribute a large-scale dataset capturing real-world physiological responses under diverse environmental conditions. The multi-modal dataset includes continuous, time-synchronized PPG, ECG, motion, and environmental measurements across four different on-body locations. This dataset enables the development of more robust signal-processing techniques and machine-learning models tailored for wearable sensing applications.

Using the data of WildPPG, we introduced a lightweight sensor fusion approach for PPG signals from multiple on-body locations that enhances heart rate estimation by mitigating motion artifacts and environmental disturbances. Additionally, we developed a learning-based method to generate synthetic PPG signals based on measurements acquired at multiple optical wavelengths. This improves cardiac monitoring in applications where PPG signals are acquired at multiple wavelengths, such as pulse oximetry and is designed to enable a wide range of downstream analysis tasks by providing a continuous PPG signal as output.

Finally, we presented a novel method for continuous respiratory rate monitoring using UWB radar, demonstrating its potential for unobtrusive respiratory tracking. Compared to other wearable methods that are most commonly based on IMU measurements, our method was less prone to motion artifacts and performed better during physical activities.

9.2 OUTLOOK AND FUTURE WORK

Over the past half-century, technological advancements have led to increasingly compact and accessible devices that match or surpass the capabilities of earlier, more complex systems. Computers that once occupied entire rooms in the 1950s are now significantly outperformed by modern smartphones. This trend has continued in the realm of wearable technology, enabling new applications, as demonstrated in this dissertation. There is no indication that this trajectory will slow down. Wearables are expected to continue improving in power efficiency, computational performance, memory capacity, and form factor, further expanding their potential applications. Building on these advancements, several promising directions for future research emerge:

9.2.1 *Distributed Sensing*

This dissertation has demonstrated that fusing sensor data from multiple on-body locations enhances data quality for downstream analysis. This approach has great future potential: as wearables become smaller and require less frequent charging, the use of multiple devices will become more practical for users. Future research should focus on refining this concept by developing more efficient and reliable sensor fusion techniques. This includes the integration of diverse signal modalities, an area of ongoing investigation [23].

9.2.2 *Real-World Data Acquisition*

As shown in our research, many published algorithms struggle with data recorded in real-world applications outside of controlled environments. To enhance robustness in real-world applications, future research should prioritize the development of algorithms that perform reliably in everyday settings. Given that modern machine learning methods require large, diverse datasets, emphasis should be placed on acquiring representative data from real-world scenarios where individuals move and interact naturally. Furthermore, obtaining high-quality ground truth data in such environments is critical. This may involve capturing accurate physiological signals as well as contextual information, such as user activities and psychological states.

9.2.3 *Exploration of New Wearable Sensing Modalities*

This dissertation introduced Respiro, leveraging UWB technology as a novel sensing modality for respiratory rate monitoring. This approach was motivated by the fact that traditional methods from controlled settings such as hospitals do not translate well to wearable devices due to the need for obtrusive breathing masks. Future work should continue the exploration of alternative sensing modalities to track parameters that are challenging to measure directly in wearable devices. Besides respiratory measurements, this includes the measurement of blood pressure in the context of vital signs.

9.2.4 *Lightweight Body Pose Estimation for Health Applications*

Body pose estimation is a topic of ongoing research in human-computer interaction and presents significant potential for health-related wearable applications. The ability to track movement range and speed using minimal hardware could offer benefits for individuals recovering from medical procedures or adhering to physical therapy regimens. Future work should focus on the development of lightweight, wearable-compatible body pose estimation techniques that can provide actionable insights for rehabilitation and fitness applications.

By providing new means for data acquisition and analysis, this dissertation provides tools, data, and analysis methods to enable continuous physiological monitoring outside of controlled environments. The ability to track the physiological state of a person using small inexpensive wearable devices provides users with invaluable insights, particularly related to precision medicine applications. Wearables are a scalable solution to prevent adverse health events, improve recoveries and optimize daily routines.

A APPENDIX

In this Appendix, we provide additional information and baseline results related to WildPPG which we introduced in Chapter 5.

A.1 DATASET ACCESSIBILITY

We provide WildPPG time series in the MATLAB .mat file format. Code is provided in Matlab .m file format and Python .py source code. The python source code contains functions to load the .mat files into equivalent python structures. All data, code, and the project website are hosted on ETH Zurich servers with long-running maintenance intended for long-term availability. The dataset uses a Creative Commons CC BY-NC-SA 4.0 license [31] while the code is released under GPL-3 [72].

A.2 EXPERIMENTAL SETTINGS

This section provides more details on the implementations of the learning-based baseline methods and the dataset we tested them on.

A.2.1 *Datasets*

The datasets used alongside WildPPG are both widely employed in related work about HR estimation from PPG signals.

A.2.1.1 *IEEE SPC*

This competition provided a training dataset of 12 participants (SPC12) and a test dataset of 10 participants [204]. The IEEE SPC dataset overall has 22 recordings of 22 participants (ages 18 – 58) performing three different activities. Each recording contains data sampled at 125 Hz from a 3-axis accelerometer and 2-channel pulse oximeter sensor (green LEDs). All these recordings were captured from the wearable device placed on the wrist of each individual. Additionally, a chest ECG provides ground-truth HR estimation. During our experiments, we used the PPG channels only. We use leave-one-out-cross-validation similar to the previous setups.

A.2.1.2 *DaLiA*

PPG signals from the DaLia dataset were recorded at a sampling rate of 64 Hz. We used the whole dataset of 15 participants while following the leave-one-out-cross-validation. 10% of the training data is randomly split for validation, i.e., early stopping and saving the best model.

All PPG datasets are standardized as follows. Initially, a fourth-order Butterworth bandpass filter with a frequency range of 0.5–4 Hz is applied to the PPG signals. Subsequently, a sliding window of 8 seconds with 2-second shifts is employed for segmentation, followed by z-score normalization of each segment. Lastly, the signal is resampled to a frequency of 25 Hz for each segment.

A.2.2 *Supervised Baseline Methods*

We train the models using a batch size of 128 with a learning rate of $5e-4$. The validation loss is employed to save the optimal model and halt training to prevent overfitting. If no improvements are observed for 15 consecutive epochs, we reduce the learning rate by a factor of 10. Our experiments utilized NVIDIA GeForce RTX 4090 GPUs and involved training with three random seeds across all datasets, resulting in approximately 480 GPU hours, including ablation.

A.2.2.1 *1D ResNet*

is a modified version of ResNet for time series with 1D filters of size five across eight residual blocks. The initial block consists of 32 filters, incrementing by a factor of two for every two subsequent residual blocks. Also, batch normalization [92] is applied after each convolutional block. We apply a pooling operation after each residual with a stride of 2 while using Dropout [172] with 0.5 after each activation. Finally, a global average pooling is used before the final linear layer.

A.2.2.2 *DCL*

is a combination of convolutional blocks with the long-short term memory units [85] (LSTMs) and is widely used for time series [26, 141] as it considers the temporal relationship. Specifically, the DCL architecture has four convolutional layers with 5×1 size of 64 kernels. Then, the output is fed into the 2-layer LSTM with 128 units. The last time step of the LSTM is fed to a linear layer.

A.2.2.3 *FCN*

We also implemented a fully convolutional neural network with a 3-layer followed by ReLU activation and MaxPooling after each convolutional layer, similar to the implementation in [141]. Dropout with 0.5 is applied after the first convolutional layer. We set the kernel and padding size to 8 and 4, respectively for each convolutional layer. The number of kernels for each convolutional layer is set to 32 for the first one and 64 for the rest.

A.2.2.4 *LSTM*

We employed a unit of two layers of LSTM [85] without any feature extraction, i.e., the raw PPG data is fed to the model. The features from the final LSTM are fed to a linear layer for HR estimation.

A.2.2.5 *Transformer*

We also added transformers with positional encodings as a baseline model. Since the attention mechanism [183] is used to capture dependencies between any points within the input sequence, irrespective of their temporal distance, we used it for time series data, similar to [141]. Specifically, we used linear layers with a stack of four identical blocks. The linear layer converts the input data to embedding vectors of 128. A token of size 128 is added to the embedded input as the representation vector. Each block is made up of a multi-head self-attention layer and a fully connected feed-forward layer with residual connections around them.

A.2.3 *Self-Supervised Baseline Methods*

For our self-supervised learning experiments, we follow the same implementation setup as previous works [54, 55, 141] for self-supervised learning in time series. Specifically, we use a combination of convolutional with LSTM-based network, which is widely used in heart rate estimation from PPG signals [26, 54], as backbones for the encoder $f_\theta(.)$ where the projector is two fully connected layers. During pre-training, we use InfoNCE (for contrastive learning-based methods) as the loss function, which is optimized using Adam [99] with a learning rate of 0.003. We train the models with a batch size of 256 for 120 epochs and decay the learning rate using the cosine decay. After pre-training, we train a single linear layer classifier on features extracted from the frozen pre-trained network, i.e., linear probing.

A.2.3.1 *SimCLR*

SimCLR [41] introduces a contrastive learning framework for self-supervised visual representation learning. The method relies on maximizing agreement between differently augmented views of the same image via a contrastive loss in the latent space. We follow the previous implementations of SimCLR for time series [141] and PPG signals [54].

A.2.3.2 *NNCLR*

We follow a similar setup to SimCLR by applying two separate data augmentations, and then we use nearest neighbors in the learned representation space as the positive in contrastive losses [61]. The maximum size of the support set equals 1024.

A.2.3.3 *BYOL*

For the BYOL implementation, the exponential moving average parameter is set to 0.996 where the projector size is set to 128. We set the learning rate to 0.03 similar to other SSL techniques. Following the original implementation, we use a weight decay parameter of $1.5e-6$.

A.2.3.4 *TS-TCC*

We follow the same architecture implementation with the losses, i.e., contextual and temporal contrasting. TS-TCC [62] proposed applying two separate augmentations, one augmentation is weak (jitter-and-scale) and the other is strong (permutation-and-jitter). The authors also change the strength of the permutation window from dataset to dataset. In our experiments, we follow the previous work [55] for PPG-based augmentation.

A.2.3.5 *TS2Vec*

TS2Vec [199] is an SSL method specifically designed for time series based on contrastive (instance and temporal-wise) learning in a hierarchical way over augmented context views where the context is generated by applying timestamp masking and random cropping on the input time series. Following the original framework, we use a dilated CNN architecture with a depth of 10 and a hidden size of 64, which has a similar number of parameters to the architectures used by other SSL methods. The batch size is set to 256, and the number of epochs to 120, consistent with other SSL techniques.

A.2.3.6 *VICReg*

We follow the original implementation and set the coefficients for each loss term to 25 (λ), 25 (μ), and 1 (ν), corresponding to the invariance, variance, and covariance terms, respectively. We have not performed an additional hyper parameter search as these values are also set by previous on this application [55].

A.2.3.7 *Barlow Twins*

Barlow Twins [200] presents a function to avoid collapse for SSL by measuring the cross-correlation matrix between the outputs of two identical networks fed with augmented versions of a sample, and making it as close to the identity matrix as possible. This causes the embedding vectors of augmented versions of a sample to be similar while minimizing the redundancy between the components of these vectors. Following the original implementation, we applied batch normalization to the extracted embeddings and set the hyperparameter λ coefficient (in Equation 6) to 0.005.

$$\mathcal{L} = \sum_i (1 - C_{ii})^2 + \lambda \sum_i \sum_{j \neq i} C_{ij}^2, \tag{6}$$

where C is the cross-correlation matrix computed between the two sets of normalized embeddings.

A.2.4 *Unsupervised Baseline Method*

We also consider one of the unsupervised learning based method, unsupervised periodicity detection (UPD), as an additional baseline model.

A.2.4.1 *UPD*

UPD [55] presented two novel regularizers for detecting periodic patterns in time series data without using labels or specific augmentations. We follow the original implementation which can be found on `github`.

Since this method operates without supervision, it requires relatively clean samples to learn periodic representations. Consequently, its performance is quite low on datasets with noisy samples, such as WildPPG and DaLiA datasets.

A.3 ADDITIONAL EXPERIMENTS

Here, we present additional experiments we conducted beyond our main experiments, which focused on estimating the heart rate from green PPG recordings on the wrist as it is the most widely used measurement site and PPG wavelength for wearable devices [19].

A.3.1 *Evaluation Across Multiple Body Locations and PPG Wavelengths*

We investigated the performance of the models when we integrate PPG signals from the different measurement sites of the body as well as PPG signals recorded at different wavelengths. Table 11 presents the results for the baseline and the modified ResNet architecture performances.

From these results, we observe that incorporating additional measurement sites for HR prediction generally reduces the error, despite

Table 11: Performance comparison of baselines with prior works in datasets

Location	WildPPG		
Method	MAE↓	RMSE↓	ρ↑
Only Wrist			
1D ResNet	8.62	15.00	42.41
Temp-ResNet	8.37	14.72	45.50
Only Chest			
1D ResNet	8.55	14.95	44.30
Temp-ResNet	8.35	14.68	46.48
Wrist and Chest			
1D ResNet	8.40	14.85	44.40
Temp-ResNet	**8.28**	**14.67**	**47.55**
All			
1D ResNet	8.60	14.93	43.27
Temp-ResNet	8.33	14.70	46.50

Wavelength	WildPPG		
Method	MAE↓	RMSE↓	ρ↑
Green			
1D ResNet	8.62	15.00	42.41
Temp-ResNet	8.37	**14.72**	45.50
Red			
1D ResNet	10.25	17.63	35.21
Temp-ResNet	9.97	16.41	36.45
Infrared			
1D ResNet	8.97	15.93	36.33
Temp-ResNet	8.72	15.88	37.66
All			
1D ResNet	8.59	14.97	43.21
Temp-ResNet	**8.35**	14.68	**46.38**

increasing the number of model parameters. Furthermore, results suggest that the wrist measurement site has the lowest signal quality compared to other sites. This outcome is expected, as the wrist is particularly susceptible to motion artifacts and exposure to cold, and the results are consistent with prior work [117]. Similarly, the signal quality of alternative wavelengths is lesser compared to green PPG which is also consistent with prior work [146].

A.3.2 *Cross-Dataset Evaluation*

We evaluated model performance using a cross-dataset approach, specifically employing a leave-one-dataset-out scheme. In this setup, two datasets were used for training, while the third dataset was reserved for testing. Results are summarized in Table 12.

A.3.3 *Additional sensor modality*

In Table 13 we present the experimental results of integrating acceleration with PPG as an additional modality.

Table 12: Performance comparison of baselines in cross-dataset evaluation (one dataset left for evaluation and others used for training)

Method	Test on WildPPG			Test on SPC12			Test on DaLiA		
	MAE↓	RMSE↓	ρ↑	MAE↓	RMSE↓	ρ↑	MAE↓	RMSE↓	ρ↑
Supervision									
1D ResNet	16.32	27.33	12.33	**11.25**	**18.37**	**67.36**	17.23	22.10	67.14
DCL	19.43	29.32	10.68	15.45	26.32	55.43	20.90	25.31	60.21
FCN	16.55	26.77	11.53	14.22	24.58	50.28	24.13	24.56	65.32
LSTM	16.79	28.57	8.97	27.86	28.11	42.11	21.59	23.47	63.42
Transformer	17.25	30.98	6.72	28.45	30.38	30.85	20.31	20.17	55.21

Table 13: Performance comparison of baseline models across datasets when integrating accelerometer data with PPG as an additional modality

Method	Test on WildPPG			Test on SPC12			Test on DaLiA		
	MAE↓	RMSE↓	ρ↑	MAE↓	RMSE↓	ρ↑	MAE↓	RMSE↓	ρ↑
Supervision									
1D ResNet	8.22	14.34	44.33	**5.33**	**11.08**	**85.02**	4.43	09.88	85.96
DCL	8.69	15.03	42.01	15.30	20.98	20.78	4.88	11.73	79.12
FCN	9.65	15.16	37.42	25.95	27.87	45.13	6.12	10.92	84.04
LSTM	8.72	14.11	38.61	17.41	21.13	60.10	5.07	10.94	80.01
Transformer	9.98	17.10	31.73	21.92	25.10	50.97	7.17	14.98	68.97
Temp-ResNet	**8.13**	**13.95**	**47.48**	—	—	—	**4.27**	**9.13**	**87.31**

A.3.4 *Baseline Results on the BIDMC Dataset*

Considering the wide use of the BIDMC dataset (recorded in controlled environments), we provide baseline results of the baseline methods listed in Section A.2.2 for this dataset in Table 14.

A.4 IMPACT OF TEMPERATURE AND MOTION

The modified architecture Temp-ResNet adds the device's case temperature (measured within the encasing) as an additional input modality to improve results. Figure 29 shows the HR error of 1D ResNet based on PPG from the wrist in relation to case temperature and motion. The correlations of error and temperature/motion are supported by the investigation of signal-to-noise ratio (SNR): We have computed the SNR of the PPG signal by relating the power P_{HR} of the ground truth HR in the frequency spectrum of the PPG signal to the power $P_{noise} = P_{tot} - P_{HR}$ of the other frequency components in the relevant window of possible heart rates:

Table 14: Performance comparison of baselines when tested on the BIDMC dataset

Method	BIDMC		
	MAE↓	RMSE↓	ρ↑
Heuristic			
FFT	4.41	9.74	33.42
Supervision			
1D ResNet	**3.62**±0.22	**5.23**±0.39	85.43±0.20
DCL	4.18±0.31	5.89±0.65	83.34±1.37
FCN	4.01±0.27	5.41±0.47	84.57±1.14
LSTM	5.73±1.79	10.86±0.96	29.37±1.58
Transformer	7.71±1.40	11.03±1.35	19.21±2.20

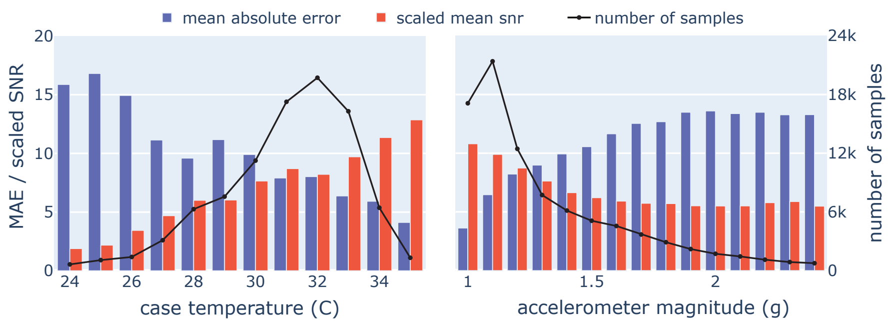

Figure 29: The error (blue) of 1D ResNet based on PPG from the wrist, dependent on the device temperature (left), and device motion (right). Inversely related is the computed SNR of the PPG signal (red).

$$\mathrm{SNR} = 10 \cdot \log \left(\frac{P_{HR}}{P_{tot} - P_{HR}} \right) \mathrm{dB}$$

While it is well understood that motion artifacts have a negative impact on PPG signal quality [204], the presented results show that the temperature within the wearable has an inverse correlation to PPG signal quality as well.

A.5 PPG DEVICE

The schematics and production data (CAD drawing, BOM, Gerber/NCDrill, pick&place) of the PPG device used in this project are available on the project website. The data is acquired using an off-the-shelf PPG LED module (SFH7072, OSRAM) and a commonly

Table 15: Register configurations of the MAX86141 PPG chip

Register Name	Addr	Value	Description
PPG_CONFIG_1	0x11	0x2B	ALC enable, 117us integr., 16uA range
PPG_CONFIG_2	0x12	0x70	128 Hz SR (ext. clock), no averaging
PPG_CONFIG_3	0x13	0xC0	12uS LED settle time
PHOTO_DIODE_BIAS	0x15	0x11	Low capacity PD
LED_RANGE_1	0x2A	0x00	LED driver range 31mA
LED1_PA	0x23	0x12	IR LED current 2.16mA
LED2_PA	0x24	0x0C	R LED current 1.44mA
LED3_PA	0x25	0x0C	G LED current 1.44mA
INTERRUPT_ENABLE_1	0x02	0x00	Disable all Interrupts
LED_SEQ_REG1	0x20	0x21	Set LED sequence (1 slot each for IR R G)
LED_SEQ_REG2	0x21	0x03	Set LED sequence (1 slot each for IR R G)
LED_SEQ_REG3	0x22	0x00	Set LED sequence (1 slot each for IR R G)

used dedicated PPG acquisition chip (MAX86141, Analog Devices). Table 15 shows the register configuration of the MAX86141 PPG chip as used during the data acquisition.

A.6 file structure

The file structure of the WildPPG dataset is shown in Figure 30

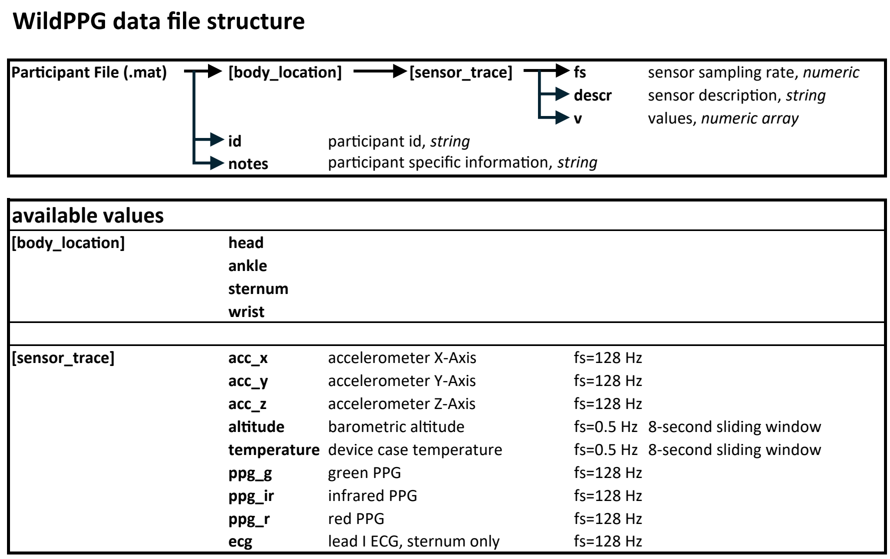

Figure 30: For each participant in WildPPG, one .mat Matlab data file is provided with the structure as illustrated.

[1] Omar Abdallah and Armin Bolz. "Adaptive Filtering by Non-Invasive Vital Signals Monitoring and Diseases Diagnosis." In: June 2011. ISBN: 978-953-307-306-4. DOI: 10.5772/16838.

[2] Omar Abdallah and Armin Bolz. "Adaptive filtering by non-invasive vital signals monitoring and diseases diagnosis." In: *Adaptive Filtering Applications* (2011), pp. 157–172.

[3] Mohamed Ali, Ali Elsayed, Arnaldo Mendez, Yvon Savaria, and Mohamad Sawan. *Contact and Remote Breathing Rate Monitoring Techniques: A Review*. 2021. DOI: 10.1109/JSEN.2021.3072607.

[4] Parastoo Alinia, Chris Cain, Ramin Fallahzadeh, Armin Shahrokni, Diane Cook, and Hassan Ghasemzadeh. "How Accurate Is Your Activity Tracker? A Comparative Study of Step Counts in Low-Intensity Physical Activities." en. In: *JMIR Mhealth Uhealth* 5.8 (Aug. 2017), e106.

[5] Byeong Wan An et al. "Smart Sensor Systems for Wearable Electronic Devices." In: *Polymers* 9.8 (2017). ISSN: 2073-4360. DOI: 10.3390/polym9080303. URL: https://www.mdpi.com/2073-4360/9/8/303.

[6] Christoph Hoog Antink, Christoph Brüser, and Steffen Leonhardt. "Multimodal sensor fusion of cardiac signals via blind deconvolution: A source-filter approach." In: *Computing in Cardiology 2014*. Sept. 2014, pp. 805–808. (Visited on 01/06/2024).

[7] *Apple Watch*. en-US. URL: https://www.apple.com/watch/ (visited on 02/28/2024).

[8] H Harry Asada, Hong-Hui Jiang, and Peter Gibbs. "Active noise cancellation using MEMS accelerometers for motion-tolerant wearable bio-sensors." In: *The 26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. Vol. 1. IEEE. 2004, pp. 2157–2160.

[9] Farzad Asgarian and Khalil Najafi. "BlueSync: Time Synchronization in Bluetooth Low Energy With Energy-Efficient Calculations." In: *IEEE Internet of Things Journal* 9.11 (2022), pp. 8633–8645. DOI: 10.1109/JIOT.2021.3116921.

[10] Rawin Assabumrungrat, Soravitt Sangnark, Thananya Charoenpattarawut, Wipamas Polpakdee, Thapanun Sudhawiyangkul, Ekkarat Boonchieng, and Theerawit Wilaiprasitporn. "Ubiquitous Affective Computing: A Review." In: *IEEE Sensors Journal* 22.3 (Feb. 2022). Conference

Name: IEEE Sensors Journal, pp. 1867–1881. ISSN: 1558-1748. DOI: 10 . 1109 / JSEN . 2021 . 3138269. URL: https : / / ieeexplore . ieee . org / abstract / document / 9662339 (visited on 07/04/2024).

[11] Mohan Babu, Ziv Lautman, Xiangping Lin, Milan HB Sobota, and Michael P Snyder. "Wearable devices: implications for precision medicine and the future of health care." In: *Annual Review of Medicine* 75.1 (2024), pp. 401–415.

[12] Sangeeta Bagha and Laxmi Shaw. "A real time analysis of PPG signal for measurement of SpO2 and pulse rate." In: *International journal of computer applications* 36.11 (2011), pp. 45–50.

[13] Pawan K. Baheti and Harinath Garudadri. "An Ultra Low Power Pulse Oximeter Sensor Based on Compressed Sensing." In: *2009 Sixth International Workshop on Wearable and Implantable Body Sensor Networks*. ISSN: 2376-8894. June 2009, pp. 144–148. DOI: 10.1109/BSN.2009.32. URL: https://ieeexplore.ieee.org/abstract/document/5226900 (visited on 05/01/2024).

[14] David Bannach, Oliver Amft, and Paul Lukowicz. "Automatic Event-Based Synchronization of Multimodal Data Streams from Wearable and Ambient Sensors." en. In: *Smart Sensing and Context*. Ed. by Payam Barnaghi, Klaus Moessner, Mirko Presser, and Stefan Meissner. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 2009, pp. 135–148. ISBN: 978-3-642-04471-7. DOI: 10.1007/978-3-642-04471-7_11.

[15] Rummana Bari, Md Mahbubur Rahman, Nazir Saleheen, Megan Battles Parsons, Eugene H Buder, and Santosh Kumar. "Automated detection of stressful conversations using wearable physiological and inertial sensors." In: *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies* 4.4 (2020), pp. 1–23.

[16] Anshul Basra, Bodhibrata Mukhopadhayay, and Subrat Kar. "Temperature sensor based ultra low cost respiration monitoring system." In: *2017 9th International Conference on Communication Systems and Networks, COMSNETS 2017*. 2017. DOI: 10.1109/COMSNETS.2017.7945448.

[17] A. Bates, M. J. Ling, J. Mann, and D. K. Arvind. "Respiratory rate and flow waveform estimation from tri-axial accelerometer data." In: *2010 International Conference on Body Sensor Networks, BSN 2010*. 2010. DOI: 10.1109/BSN.2010.50.

[18] Karim Bayoumy et al. "Smart wearable devices in cardiovascular care: where we are and how to move forward." en. In: *Nature Reviews Cardiology* 18.8 (Aug. 2021), pp. 581–599. DOI: 10.1038/s41569-021-00522-7. (Visited on 01/12/2024).

[19] Simone Benedetto, Christian Caldato, Elia Bazzan, Darren C. Greenwood, Virginia Pensabene, and Paolo Actis. "Assessment of the Fitbit Charge 2 for monitoring heart rate." In: *PLOS ONE* 13.2 (Feb. 2018), pp. 1–10. DOI: 10.1371/journal.pone.0192691. URL: https://doi.org/10.1371/journal.pone.0192691.

[20] Natascia Bernacchia, Lorenzo Scalise, Luigi Casacanditella, Ilaria Ercoli, Paolo Marchionni, and Enrico Primo Tomasini. "Non contact measurement of heart and respiration rates based on Kinect™." In: *IEEE MeMeA 2014 - IEEE International Symposium on Medical Measurements and Applications, Proceedings*. 2014. DOI: 10.1109/MeMeA.2014.6860065.

[21] Philip E Bickler, John R Feiner, and John W Severinghaus. "Effects of skin pigmentation on pulse oximeter accuracy at low saturation." In: *The Journal of the American Society of Anesthesiologists* 102.4 (2005), pp. 715–719.

[22] André Bideaux, Bernd Zimmermann, Stefan Hey, and Wilhelm Stork. "Synchronization in wireless biomedical-sensor networks with Bluetooth Low Energy." en. In: *Current Directions in Biomedical Engineering* 1.1 (Sept. 2015). De Gruyter, pp. 73–76. DOI: 10.1515/cdbme-2015-0019. URL: https://www.degruyter.com/view/journals/cdbme/1/1/article-p73.xml (visited on 05/25/2020).

[23] Valentin Bieri, Paul Streli, Berken Utku Demirel, and Christian Holz. "BeliefPPG: Uncertainty-aware Heart Rate Estimation from PPG signals via Belief Propagation." In: *Conference on Uncertainty in Artificial Intelligence (UAI)*. PMLR. 2023.

[24] Steven M. Bishop and Ari Ercole. "Multi-Scale Peak and Trough Detection Optimised for Periodic and Quasi-Periodic Neuroscience Data." en. In: *Intracranial Pressure & Neuromonitoring XVI*. Ed. by Thomas Heldt. Cham: Springer International Publishing, 2018, pp. 189–195. ISBN: 978-3-319-65798-1. DOI: 10.1007/978-3-319-65798-1_39.

[25] Rajendra Singh Bisht, Sourabh Jain, and Naveen Tewari. "Study of Wearable IoT devices in 2021: Analysis & Future Prospects." In: *2021 2nd International Conference on Intelligent Engineering and Management (ICIEM)*. Apr. 2021, pp. 577–581. DOI: 10.1109/ICIEM51511.2021.9445334.

[26] Dwaipayan Biswas et al. "CorNET: Deep Learning Framework for PPG-Based Heart Rate Estimation and Biometric Identification in Ambulant Environment." In: *IEEE Transactions on Biomedical Circuits and Systems* 13.2 (2019), pp. 282–291. DOI: 10.1109/TBCAS.2019.2892297.

[27] Richard Bloss. "Wearable sensors bring new benefits to continuous medical monitoring, real time physical activity assessment, baby monitoring and industrial applications." In: *Sensor Review* 35.2 (Jan. 2015), pp. 141–145.

[28] Sebastian Böttcher, Solveig Vieluf, Elisa Bruno, Boney Joseph, Nino Epitashvili, Andrea Biondi, Nicolas Zabler, Martin Glasstetter, Matthias Dümpelmann, Kristof Van Laerhoven, et al. "Data quality evaluation in wearable monitoring." In: *Scientific reports* 12.1 (2022), p. 21412.

[29] Alan Kevin Bourke, Espen Alexander F Ihlen, Ronny Bergquist, Per Bendik Wik, Beatrix Vereijken, and Jorunn L Helbostad. "A Physical Activity Reference Data-Set Recorded from Older Adults Using Body-Worn Inertial Sensors and Video Technology-The ADAPT Study Data-Set." en. In: *Sensors (Basel)* 17.3 (Mar. 2017).

[30] Elisa Bruno et al. "Wearable devices for seizure detection: Practical experiences and recommendations from the Wearables for Epilepsy And Research (WEAR) International Study Group." In: *Epilepsia* 62.10 (2021), pp. 2307–2321. DOI: `https://doi.org/10.1111/epi.17044`. eprint: `https://onlinelibrary.wiley.com/doi/pdf/10.1111/epi.17044`. URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/epi.17044`.

[31] *CC BY-NC-SA 4.0 Deed | Attribution-NonCommercial-ShareAlike 4.0 International | Creative Commons*. URL: `https://creativecommons.org/licenses/by-nc-sa/4.0/` (visited on 06/05/2024).

[32] Jiawei Cai, Qixiang Fu, Xuelin Yuan, Xiangwei Zhu, Huifu Lin, and Yinshen Huang. "Respiration and heartbeat signal separation algorithm using UWB radar platform." In: *ACM International Conference Proceeding Series*. 2022. DOI: `10.1145/3523286.3524556`.

[33] Dario Camuffo, Chiara Bertolin, Phil D. Jones, Richard Cornes, and Emmanuel Garnier. "The earliest daily barometric pressure readings in Italy: Pisa AD 1657-1658 and Modena AD 1694, and the weather over Europe." en. In: *The Holocene* 20.3 (May 2010). SAGE Publications Ltd, pp. 337–349. ISSN: 0959-6836. DOI: `10.1177/0959683609351900`. URL: `https://doi.org/10.1177/0959683609351900` (visited on 11/30/2022).

[34] Stefano Canali, Viola Schiaffonati, and Andrea Aliverti. "Challenges and recommendations for wearable devices in digital health: Data quality, interoperability, health equity, fairness." In: *PLOS Digital Health* 1.10 (2022), e0000104.

[35] Andrew Carek and Christian Holz. "Naptics: convenient and continuous blood pressure monitoring during sleep." In: *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2.3 (2018), pp. 1–22.

[36] Meredith A Case, Holland A Burwick, Kevin G Volpp, and Mitesh S Patel. "Accuracy of Smartphone Applications and Wearable Devices for Tracking Physical Activity Data." In: *JAMA* 313.6 (Feb. 2015), pp. 625–626.

[37] Marta Cavagnaro, Erika Pittella, and Stefano Pisa. "UWB pulse propagation into human tissues." In: *Physics in Medicine and Biology* 58.24 (2013). ISSN: 00319155. DOI: 10.1088/0031-9155/58/24/8689.

[38] Ambra Cesareo, Emilia Biffi, David Cuesta-Frau, Maria G. D'Angelo, and Andrea Aliverti. "A novel acquisition platform for long-term breathing frequency monitoring based on inertial measurement units." In: *Medical and Biological Engineering and Computing* 58.4 (2020). ISSN: 17410444. DOI: 10.1007/s11517-020-02125-9.

[39] Peter H. Charlton, Timothy Bonnici, Lionel Tarassenko, David A. Clifton, Richard Beale, and Peter J. Watkinson. "An assessment of algorithms to estimate respiratory rate from the electrocardiogram and photoplethysmogram." In: *Physiological Measurement* 37.4 (2016). ISSN: 13616579. DOI: 10.1088/0967-3334/37/4/610.

[40] Peter H. Charlton, Kevin Kotzen, Elisa Mejía-Mejía, Philip J. Aston, Karthik Budidha, Jonathan Mant, Callum Pettit, Joachim A. Behar, and Panicos A. Kyriacou. "Detecting beats in the photoplethysmogram: benchmarking open-source algorithms." en. In: *Physiological Measurement* 43.8 (Aug. 2022). Publisher: IOP Publishing, p. 085007. ISSN: 0967-3334. DOI: 10.1088/1361-6579/ac826d. URL: https://dx.doi.org/10.1088/1361-6579/ac826d (visited on 05/30/2024).

[41] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. "A Simple Framework for Contrastive Learning of Visual Representations." In: *Proceedings of the 37th International Conference on Machine Learning*. ICML'20. JMLR.org, 2020.

[42] M. Y.W. Chia, S. W. Leong, C. K. Sim, and K. M. Chan. "Through-wall UWB radar operating within FCC's mask for sensing heart beat and breathing rate." In: *EURAD 2005 Conference Proceedings - 2nd European Radar Conference*. Vol. 2005. 2005. DOI: 10.1109/EURAD.2005.1605615.

[43] A. Choi and H. Shin. "Photoplethysmography sampling frequency: pilot assessment of how low can we go to analyze pulse rate variability with reliability?" en. In: *Physiological Measurement* 38.3 (Feb. 2017). IOP Publishing, p. 586. DOI: 10.1088/1361-6579/aa5efa. URL: https://dx.doi.org/10.1088/1361-6579/aa5efa (visited on 12/09/2022).

[44] Jongyoon Choi, Beena Ahmed, and Ricardo Gutierrez-Osuna. "Development and Evaluation of an Ambulatory Stress Monitor Based on Wearable Sensors." In: *IEEE Transactions on Information Technology in Biomedicine* 16.2 (Mar. 2012). Conference Name: IEEE Transactions on Information Technology in Biomedicine, pp. 279–286. ISSN: 1558-0032. DOI: 10.1109/TITB.2011.2169804. URL: https://ieeexplore.ieee.org/abstract/document/6030940 (visited on 05/01/2024).

[45] Yerim Choi, Yu-Mi Jeon, Lin Wang, and Kwanho Kim. "A biological signal-based stress monitoring framework for children using wearable devices." In: *Sensors* 17.9 (2017), p. 1936.

[46] Sayeed Shafayet Chowdhury, Rakib Hyder, Md. Samzid Bin Hafiz, and Mohammad Ariful Haque. "Real-Time Robust Heart Rate Estimation From Wrist-Type PPG Signals Using Multiple Reference Adaptive Noise Cancellation." In: *IEEE Journal of Biomedical and Health Informatics* 22.2 (2018), pp. 450–459. DOI: 10.1109/JBHI.2016.2632201.

[47] Michelle A. Cretikos, Rinaldo Bellomo, Ken Hillman, Jack Chen, Simon Finfer, and Arthas Flabouris. "Respiratory rate: The neglected vital sign." In: *Medical Journal of Australia* 188.11 (2008). ISSN: 13265377. DOI: 10.5694/j.1326-5377.2008.tb01825.x.

[48] Ivana Culjak, Hrvoje Mihaldinec, Hrvoje Dzapo, and Mario Cifrek. "A data-fusion algorithm for respiration rate extraction based on UWB transversal propagation method." In: *I2MTC 2020 - International Instrumentation and Measurement Technology Conference, Proceedings*. 2020. DOI: 10.1109/I2MTC43012.2020.9128628.

[49] Ivana Culjak, Hrvoje Mihaldinec, Zeljka Lucev Vasic, Kresimir Friganovic, Hrvoje Dzapo, and Mario Cifrek. "A Contactless Human Respiration Rate Measurement Using UWB Transversal Propagation Method." In: *2019 International Symposium on Antennas and Propagation, ISAP 2019 - Proceedings*. 2019.

[50] Jordana Dahmen, Alyssa La Fleur, Gina Sprint, Diane Cook, and Douglas L. Weeks. "Using wrist-worn sensors to measure and compare physical activity changes for patients undergoing rehabilitation." In: *2017 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom*

Workshops). 2017, pp. 667–672. DOI: 10.1109/PERCOMW.2017.7917643.

[51] Nivedita Daimiwal, M Sundhararajan, and Revati Shriram. "Respiratory rate, heart rate and continuous measurement of BP using PPG." In: *2014 International Conference on Communication and Signal Processing*. IEEE. 2014, pp. 999–1002.

[52] Peter H Dana and Bruce M Penrod. "The role of GPS in precise time and frequency dissemination." In: *GPS World* 1.4 (1990), pp. 38–43.

[53] Waltenegus Dargie. "Analysis of time and frequency domain features of accelerometer measurements." In: *2009 Proceedings of 18th International Conference on Computer Communications and Networks*. IEEE. 2009, pp. 1–6.

[54] Berken Utku Demirel and Christian Holz. "Finding Order in Chaos: A Novel Data Augmentation Method for Time Series in Contrastive Learning." In: *Advances in Neural Information Processing Systems*. Ed. by A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine. Vol. 36. Curran Associates, Inc., 2023, pp. 30750–30783. URL: https://proceedings.neurips.cc/paper_files/paper/2023/file/61c2c6338033da68885e0226881cbe71-Paper-Conference.pdf.

[55] Berken Utku Demirel and Christian Holz. *An Unsupervised Approach for Periodic Source Detection in Time Series*. 2024. arXiv: 2406.00566 [cs.LG]. URL: https://arxiv.org/abs/2406.00566.

[56] Analog Devices. *ADXL355, Low Noise, Low Drift, Low Power, 3-Axis MEMS Accelerometers*. https://www.analog.com/media/en/technical-documentation/data-sheets/adxl354_355.pdf. Last accessed: 2022-12-09. 2018.

[57] Tanja Dorst, Yannick Robin, Sascha Eichstädt, Andreas Schütze, and Tizian Schneider. "Influence of synchronization within a sensor network on machine learning results." English. In: *Journal of Sensors and Sensor Systems* 10.2 (Aug. 2021). Publisher: Copernicus GmbH, pp. 233–245. ISSN: 2194-8771. DOI: 10.5194/jsss-10-233-2021. URL: https://jsss.copernicus.org/articles/10/233/2021/ (visited on 05/01/2024).

[58] Amy D. Droitcour, Olga Boric-Lubecke, and Gregory T.A. Kovacs. "Signal-to-noise ratio in doppler radar system for heart and respiratory rate measurements." In: *IEEE Transactions on Microwave Theory and Techniques* 57.10 (2009). ISSN: 00189480. DOI: 10.1109/TMTT.2009.2029668.

[59] Zhenzhen Duan and Jing Liang. "Non-Contact Detection of Vital Signs Using a UWB Radar Sensor." In: *IEEE Access* 7 (2019). ISSN: 21693536. DOI: 10.1109/ACCESS.2018.2886825.

[60] Jessilyn Dunn, Ryan Runge, and Michael Snyder. "Wearables and the medical revolution." In: *Personalized medicine* 15.5 (2018), pp. 429–448.

[61] D. Dwibedi, Y. Aytar, J. Tompson, P. Sermanet, and A. Zisserman. "With a Little Help from My Friends: Nearest-Neighbor Contrastive Learning of Visual Representations." In: *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. Los Alamitos, CA, USA: IEEE Computer Society, Oct. 2021, pp. 9568–9577. DOI: 10.1109/ICCV48922.2021.00945. URL: https://doi.ieeecomputersociety.org/10.1109/ICCV48922.2021.00945.

[62] Emadeldeen Eldele, Mohamed Ragab, Zhenghua Chen, Min Wu, Chee Keong Kwoh, Xiaoli Li, and Cuntai Guan. "Time-Series Representation Learning via Temporal and Contextual Contrasting." In: *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*. 2021, pp. 2352–2359.

[63] Malcolm Elliott and Alysia Coventry. "Critical care: The eight vital signs of patient monitoring." In: *British Journal of Nursing* 21.10 (2012). ISSN: 09660461. DOI: 10.12968/bjon.2012.21.10.621.

[64] Malcolm Elliott and Alysia Coventry. "Critical care: the eight vital signs of patient monitoring." In: *British Journal of Nursing* 21.10 (May 2012). Publisher: Mark Allen Group, pp. 621–625. ISSN: 0966-0461. DOI: 10.12968/bjon.2012.21.10.621. URL: https://www.magonlinelibrary.com/doi/abs/10.12968/bjon.2012.21.10.621 (visited on 02/28/2024).

[65] Jeremy Elson, Lewis Girod, and Deborah Estrin. "Fine-Grained Network Time Synchronization Using Reference Broadcasts." In: *5th Symposium on Operating Systems Design and Implementation (OSDI 02)*. USENIX Association, 2002. URL: https://www.usenix.org/conference/osdi-02/fine-grained-network-time-synchronization-using-reference-broadcasts (visited on 10/20/2022).

[66] Bennett A Fallow, Takashi Tarumi, and Hirofumi Tanaka. "Influence of skin type and wavelength on light wave reflectance." In: *Journal of clinical monitoring and computing* 27 (2013), pp. 313–317.

[67] Hua Fang, Zhaoyang Zhang, Chanpaul Jin Wang, Mahmoud Daneshmand, Chonggang Wang, and Honggang Wang. "A survey of big data research." In: *IEEE Network* 29.5 (Sept. 2015). Conference Name: IEEE Network, pp. 6–9. ISSN: 1558-156X. DOI: 10.1109/MNET.2015.7293298. URL: https://ieeexplore.ieee.org/abstract/document/7293298 (visited on 05/01/2024).

[68] Federal Communications Commission (FCC). *47 CFR Part 15 Subpart F*. Sept. 2023. URL: https://www.ecfr.gov/current/title-47/part-15/subpart-F.

[69] John F. Fieselmann, Michael S. Hendryx, Charles M. Helms, and Douglas S. Wakefield. "Respiratory rate predicts cardiopulmonary arrest for internal medicine inpatients." In: *Journal of General Internal Medicine* 8.7 (1993). ISSN: 08848734. DOI: 10.1007/BF02600071.

[70] Thomas B Fitzpatrick. "The validity and practicality of sun-reactive skin types I through VI." In: *Archives of dermatology* 124.6 (1988), pp. 869–871.

[71] Hanna A Frank, Karen Jacobs, and Hugh McLoone. "The effect of a wearable device prompting high school students aged 17-18 years to break up periods of prolonged sitting in class." In: *Work* 56.3 (2017), pp. 475–482.

[72] *GNU General Public License version 3*. en-US. Oct. 2007. URL: https://opensource.org/license/gpl-3-0 (visited on 06/05/2024).

[73] Saurabh Ganeriwal, Ram Kumar, and Mani B. Srivastava. "Timing-sync protocol for sensor networks." In: *Proceedings of the 1st international conference on Embedded networked sensor systems*. SenSys '03. New York, NY, USA: Association for Computing Machinery, Nov. 2003, pp. 138–149. ISBN: 978-1-58113-707-1. DOI: 10.1145/958491.958508. URL: https://doi.org/10.1145/958491.958508 (visited on 10/20/2022).

[74] Yan Gao, Yang Long, Yu Guan, Anna Basu, Jessica Baggaley, and Thomas Ploetz. "Towards Reliable, Automated General Movement Assessment for Perinatal Stroke Screening in Infants Using Wearable Accelerometers." In: *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 3.1 (Mar. 2019). DOI: 10.1145/3314399. URL: https://doi.org/10.1145/3314399.

[75] Garmin and Garmin Ltd or its subsidiaries. *Garmin fēnix® 7 – Standard Edition | Multisport GPS Smartwatch*. en-US. URL: https://www.garmin.com/en-US/p/735611 (visited on 02/02/2024).

[76] Shkurta Gashi, Elena Di Lascio, and Silvia Santini. "Using unobtrusive wearable sensors to measure the physiological synchrony between presenters and audience members." In: *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3.1 (2019), pp. 1–19.

[77] Paul van Gent, Haneen Farah, Nicole van Nes, and Bart van Arem. "HeartPy: A novel heart rate algorithm for the analysis of noisy signals." en. In: *Transportation Research Part F: Traffic Psychology and Behaviour* 66 (Oct. 2019), pp. 368–378.

ISSN: 1369-8478. DOI: 10.1016/j.trf.2019.09.015. URL: https://www.sciencedirect.com/science/article/pii/S1369847818306740 (visited on 11/04/2021).

[78] Unmesh Ghoshdastider, Reinhard Viga, and Michael Kraft. "Wireless time synchronization of a collaborative brain-computer-interface using bluetooth low energy." In: *SENSORS, 2014 IEEE*. 2014, pp. 2250–2254. DOI: 10.1109/ICSENS.2014.6985489.

[79] Unmesh Ghoshdastider, Reinhard Viga, and Michael Kraft. "Experimental evaluation of a pairwise broadcast synchronization in a low-power Cyber-physical system." In: *2015 IEEE Topical Conference on Wireless Sensors and Sensor Networks (WiSNet)*. 2015, pp. 50–52. DOI: 10.1109/WISNET.2015.7127399.

[80] Thomas Gilbert, Sally Day, Antonia F De C Hamilton, and Jamie Ward. "A Simple Method for Synchronising Multiple IMUs using the Magnetometer." In: *Proceedings of the 2022 ACM International Symposium on Wearable Computers*. ISWC '22. New York, NY, USA: Association for Computing Machinery, Dec. 2022, pp. 100–102. ISBN: 978-1-4503-9424-6. DOI: 10.1145/3544794.3558466. URL: https://dl.acm.org/doi/10.1145/3544794.3558466 (visited on 04/25/2023).

[81] Martin Gjoreski, Hristijan Gjoreski, Mitja Luštrek, and Matjaž Gams. "How Accurately Can Your Wrist Device Recognize Daily Activities and Detect Falls?" en. In: *Sensors (Basel)* 16.6 (June 2016).

[82] Firat Güder, Alar Ainla, Julia Redston, Bobak Mosadegh, Ana Glavan, T. J. Martin, and George M. Whitesides. "Paper-Based Electrical Respiration Sensor." In: *Angewandte Chemie - International Edition* 55.19 (2016). ISSN: 15213773. DOI: 10.1002/anie.201511805.

[83] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. *Deep Residual Learning for Image Recognition*. 2015. arXiv: 1512.03385 [cs.CV].

[84] Fatemeh Heydari, Malikeh P. Ebrahim, and Mehmet R. Yuce. "Chest-based Real-Time Pulse and Respiration Monitoring Based on Bio-Impedance." In: *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS*. Vol. 2020-July. 2020. DOI: 10.1109/EMBC44109.2020.9176348.

[85] Sepp Hochreiter and Jürgen Schmidhuber. "Long Short-Term Memory." In: *Neural Comput.* (1997).

[86] Christian Holz and Edward J. Wang. "Glabella: Continuously Sensing Blood Pressure Behavior using an Unobtrusive Wearable Device." In: *Proc. ACM IMWUT* 1.3 (Sept. 2017), 58:1–58:23. DOI: 10.1145/3132024. URL: https://doi.org/10.1145/3132024 (visited on 11/10/2022).

[87] *Homepage - MeteoSwiss*. en. URL: https://www.meteoswiss.admin.ch/ (visited on 05/31/2024).

[88] Shenda Hong, Yanbo Xu, Alind Khare, Satria Priambada, Kevin Maher, Alaa Aljiffry, Jimeng Sun, and Alexey Tumanov. "HOLMES: Health OnLine Model Ensemble Serving for Deep Learning Models in Intensive Care Units." In: *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2020, pp. 1614–1624.

[89] Igor Immoreev and Teh Ho Tao. "UWB radar for patient monitoring." In: *IEEE Aerospace and Electronic Systems Magazine*. Vol. 23. 2008. DOI: 10.1109/MAES.2008.4693985.

[90] Goertek Microelectronics Inc. *SPL07-003, Digital pressure sensor*. https://media.digikey.com/pdf/Data%20Sheets/Goertek%20Microelectronics%20PDFs/SPL07-003.pdf. Last accessed: 2023-02-14. 2022.

[91] ECS Inc International. *ECX-16, 32.768 KHz SMD Tuning Fork Crystal*. https://ecsxtal.com/store/pdf/ECX-16.pdf. Last accessed: 2022-12-09. 2021.

[92] Sergey Ioffe and Christian Szegedy. "Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift." In: *Proceedings of the 32nd International Conference on Machine Learning*. Proceedings of Machine Learning Research. PMLR, 2015. URL: https://proceedings.mlr.press/v37/ioffe15.html.

[93] Delaram Jarchi and Alexander J. Casson. "Description of a Database Containing Wrist PPG Signals Recorded during Physical Exercise with Both Accelerometer and Gyroscope Measures of Motion." In: *Data* 2.1 (2017). ISSN: 2306-5729. DOI: 10.3390/data2010001. URL: https://www.mdpi.com/2306-5729/2/1/1.

[94] Aida Kamišalić, Iztok Fister, Muhamed Turkanović, and Sašo Karakatič. "Sensors and Functionalities of Non-Invasive Wrist-Wearable Devices: A Review." In: *Sensors* 18.6 (2018). ISSN: 1424-8220. DOI: 10.3390/s18061714. URL: https://www.mdpi.com/1424-8220/18/6/1714.

[95] Heikki Karvonen, Carlos Pomalaza-Ráez, Konstantin Mikhaylov, Matti Hämäläinen, and Jari Iinatti. "Experimental Performance Evaluation of BLE 4 Versus BLE 5 in Indoors

and Outdoors Scenarios." en. In: *Advances in Body Area Networks I*. Ed. by Giancarlo Fortino and Zhelong Wang. Cham: Springer International Publishing, 2019, pp. 235–251. ISBN: 978-3-030-02819-0. DOI: 10.1007/978-3-030-02819-0_18.

[96] Musabbir Khan, Christopher G. Pretty, Alexander C. Amies, Rodney Elliott, Geoffrey M. Shaw, and J. Geoffrey Chase. "Investigating the Effects of Temperature on Photoplethysmography." In: *IFAC-PapersOnLine*. 9th IFAC Symposium on Biological and Medical Systems BMS 2015 48.20 (Jan. 2015), pp. 360–365. ISSN: 2405-8963. DOI: 10.1016/j.ifacol.2015.10.166. URL: https://www.sciencedirect.com/science/article/pii/S2405896315020571 (visited on 05/31/2024).

[97] Jongpal Kim, Takhyung Lee, Jihoon Kim, and Hyoungho Ko. "Ambient light cancellation in photoplethysmogram application using alternating sampling and charge redistribution technique." In: *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. ISSN: 1558-4615. Aug. 2015, pp. 6441–6444. DOI: 10.1109/EMBC.2015.7319867. URL: https://ieeexplore.ieee.org/abstract/document/7319867 (visited on 05/31/2024).

[98] Sunwook Kim and Maury A. Nussbaum. "Performance evaluation of a wearable inertial motion capture system for capturing physical exposures during manual material handling tasks." In: *Ergonomics* 56.2 (2013). PMID: 23231730, pp. 314–326. DOI: 10.1080/00140139.2012.742932. eprint: https://doi.org/10.1080/00140139.2012.742932. URL: https://doi.org/10.1080/00140139.2012.742932.

[99] Diederik P. Kingma and Jimmy Ba. *Adam: A Method for Stochastic Optimization*. Ed. by Yoshua Bengio and Yann LeCun. 2017. arXiv: 1412.6980 [cs.LG]. URL: http://arxiv.org/abs/1412.6980.

[100] Michael R Kosorok and Eric B Laber. "Precision medicine." In: *Annual review of statistics and its application* 6.1 (2019), pp. 263–286.

[101] Agni Kumar, Vikramjit Mitra, Carolyn Oliver, Adeeti Ullal, Matt Biddulph, and Irida Mance. "Estimating Respiratory Rate from Breath Audio Obtained Through Wearable Microphones." In: *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS*. 2021. DOI: 10.1109/EMBC46164.2021.9629661.

[102] Shinjae Kwon, Hojoong Kim, and Woon-Hong Yeo. "Recent advances in wearable sensors and portable electronics for sleep monitoring." In: *Iscience* 24.5 (2021).

[103] Mostafa Labib, Atef Ghalwash, Sarah Abdulkader, and Mohamed Elgazzar. "Networking solutions for connecting bluetooth low energy devices-a comparison." In: *MATEC web of conferences*. Vol. 292. EDP Sciences. 2019, p. 02003.

[104] H. Lee, H. Chung, and J. Lee. "Motion Artifact Cancellation in Wearable Photoplethysmography Using Gyroscope." In: *IEEE Sensors Journal* 19.3 (2019), pp. 1166–1175. DOI: 10.1109/JSEN.2018.2879970.

[105] Jongshill Lee, Minseong Kim, Hoon-Ki Park, and In Young Kim. "Motion Artifact Reduction in Wearable Photoplethysmography Based on Multi-Channel Sensors with Multiple Wavelengths." en. In: *Sensors* 20.5 (Jan. 2020), p. 1493. DOI: 10.3390/s20051493. (Visited on 01/11/2024).

[106] SeungMin Lee, HyunSoon Shin, and ChanYoung Hahm. "Effective PPG sensor placement for reflected red and green light, and infrared wristband-type photoplethysmography." In: *2016 18th International Conference on Advanced Communication Technology (ICACT)*. Jan. 2016, pp. 556–558. DOI: 10.1109/ICACT.2016.7423470. URL: https://ieeexplore.ieee.org/abstract/document/7423470 (visited on 05/31/2024).

[107] Bing Nan Li, Ming Chui Dong, and Mang I. Vai. "On an automatic delineator for arterial blood pressure waveforms." In: *Biomedical Signal Processing and Control* 5.1 (Jan. 2010), pp. 76–81. ISSN: 1746-8094. DOI: 10.1016/j.bspc.2009.06.002. URL: https://www.sciencedirect.com/science/article/pii/S1746809409000470 (visited on 05/30/2024).

[108] Jing Li, Lanbo Liu, Zhaofa Zeng, and Fengshan Liu. "Advanced signal processing for vital sign extraction with applications in UWB radar detection of trapped victims in complex environments." In: *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 7.3 (2014). ISSN: 21511535. DOI: 10.1109/JSTARS.2013.2259801.

[109] Ke Li, Cristiano Cardoso, Angel Moctezuma-Ramirez, Abdelmotagaly Elgalad, and Emerson Perin. "Heart rate variability measurement through a smart wearable device: Another breakthrough for personal health monitoring?" In: *International journal of environmental research and public health* 20.24 (2023), p. 7146.

[110] Sitian Li, Alexios Balatsoukas-Stimming, and Andreas Burg. "Band-of-Interest-based Channel Impulse Response Fusion for Breathing Rate Estimation with UWB." In: *arXiv preprint arXiv:2302.02917* (2023).

[111] Sitian Li, Andreas Toftegaard Kristensen, Andreas Burg, and Alexios Balatsoukas-Stimming. "ComplexBeat: Breathing Rate Estimation from Complex CSI." In: *IEEE Workshop on Signal Processing Systems, SiPS: Design and Implementation*. Vol. 2021-October. 2021. DOI: 10.1109/SiPS52927.2021.00046.

[112] Yanjun Li, Zengli Wang, Lin Zhang, Xianglin Yang, and Jinzhong Song. "Characters available in photoplethysmogram for blood pressure estimation: beyond the pulse transit time." In: *Australasian physical & engineering sciences in medicine* 37 (2014), pp. 367–376.

[113] Zan Li, Torsten Braun, and Desislava C. Dimitrova. "Methodology for GPS Synchronization Evaluation with High Accuracy." In: *2015 IEEE 81st Vehicular Technology Conference (VTC Spring)*. ISSN: 1550-2252. May 2015, pp. 1–6. DOI: 10.1109/VTCSpring.2015.7145929.

[114] Guan Zheng Liu, Yan Wei Guo, Qing Song Zhu, Bang Yu Huang, and Lei Wang. "Estimation of respiration rate from three-dimensional acceleration data based on body sensor network." In: *Telemedicine journal and e-health : the official journal of the American Telemedicine Association* 17.9 (2011). ISSN: 15563669. DOI: 10.1089/tmj.2011.0022.

[115] Haipeng Liu, John Allen, Dingchang Zheng, and Fei Chen. *Recent development of respiratory rate measurement technologies*. 2019. DOI: 10.1088/1361-6579/ab299e.

[116] Xi Liu, Lei Liu, Steven J. Simske, and Jerry Liu. "Human Daily Activity Recognition for Healthcare Using Wearable and Visual Sensing Data." In: *2016 IEEE International Conference on Healthcare Informatics (ICHI)*. 2016, pp. 24–31. DOI: 10.1109/ICHI.2016.100.

[117] Sally K. Longmore, Gough Y. Lui, Ganesh Naik, Paul P. Breen, Bin Jalaludin, and Gaetano D. Gargiulo. "A Comparison of Reflective Photoplethysmography for Detection of Heart Rate, Blood Oxygen Saturation, and Respiration Rate at Various Locations." en. In: *Sensors* 19.8 (Jan. 2019). Number: 8 Publisher: Multidisciplinary Digital Publishing Institute, p. 1874. ISSN: 1424-8220. DOI: 10.3390/s19081874. URL: https://www.mdpi.com/1424-8220/19/8/1874 (visited on 01/05/2024).

[118] Eduardo José da S. Luz, David Menotti, and William Robson Schwartz. "Evaluating the use of ECG signal in low frequencies as a biometry." en. In: *Expert Systems with Applications* 41.5 (Apr. 2014), pp. 2309–2315. DOI: 10.1016/j.eswa.2013.09.028. URL: https://www.sciencedirect.com/science/article/pii/S0957417413007628 (visited on 12/09/2022).

[119] *MAX86141 Datasheet and Product Info | Analog Devices*. URL: https://www.analog.com/en/products/max86141.html (visited on 06/02/2024).

[120] Sarah J MacEachern and Nils D Forkert. "Machine learning for precision medicine." In: *Genome* 64.4 (2021), pp. 416–425.

[121] Dmytro Makara, Vladyslav Tsybul'nyk, and Taras Kurnyts'kyi. "Power Efficient Clock Synchronization in Bluetooth-Based Mesh Networks." en. In: *Ambient Intelligence*. Ed. by Ioannis Chatzigiannakis, Boris De Ruyter, and Irene Mavrommati. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2019, pp. 14–26. ISBN: 978-3-030-34255-5. DOI: 10.1007/978-3-030-34255-5_2.

[122] Karandeep Malhi, Subhas Chandra Mukhopadhyay, Julia Schnepper, Mathias Haefke, and Hartmut Ewald. "A Zigbee-Based Wearable Physiological Parameters Monitoring System." In: *IEEE Sensors Journal* 12.3 (2012), pp. 423–430. DOI: 10.1109/JSEN.2010.2091719.

[123] Miklós Maróti, Branislav Kusy, Gyula Simon, and Ákos Lédeczi. "The flooding time synchronization protocol." In: *Proceedings of the 2nd international conference on Embedded networked sensor systems*. SenSys '04. New York, NY, USA: Association for Computing Machinery, Nov. 2004, pp. 39–49. ISBN: 978-1-58113-879-5. DOI: 10.1145/1031495.1031501. URL: https://doi.org/10.1145/1031495.1031501 (visited on 10/20/2022).

[124] Graeme Mattison, Oliver Canfell, Doug Forrester, Chelsea Dobbins, Daniel Smith, Juha Töyräs, and Clair Sullivan. "The influence of wearables on health care outcomes in chronic disease: systematic review." In: *Journal of Medical Internet Research* 24.7 (2022), e36690.

[125] José Jair Alves Mendes Jr., Mário Elias Marinho Vieira, Marcelo Bissi Pires, and Sergio Luiz Stevan Jr. "Sensor Fusion and Smart Sensor in Sports and Biomedical Applications." en. In: *Sensors* 16.10 (Oct. 2016), p. 1569. DOI: 10.3390/s16101569. (Visited on 01/06/2024).

[126] David Mills. *Network time protocol (version 3) specification, implementation and analysis*. Tech. rep. University of Delaware, 1992.

[127] Alessandro Montanari, Andrea Ferlini, Ananta Narayanan Balaji, Cecilia Mascolo, and Fahim Kawsar. "EarSet: A Multi-Modal Dataset for Studying the Impact of Head and Facial Movements on In-Ear PPG Signals." en. In: *Scientific Data* 10.1 (Dec. 2023), p. 850. ISSN: 2052-4463. DOI: 10.1038/s41597-023-02762-3. URL: https://www.nature.com/articles/s41597-023-02762-3 (visited on 05/30/2024).

[128] Alexander Neshitov, Konstantin Tyapochkin, Evgeniya Smorodnikova, and Pavel Pravdin. "Wavelet Analysis and Self-Similarity of Photoplethysmography Signals for HRV Estimation and Quality Assessment." In: *Sensors* 21.20 (2021). ISSN: 1424-8220. DOI: 10.3390/s21206798. URL: https://www.mdpi.com/1424-8220/21/20/6798.

[129] Andrea Nicolò, Carlo Massaroni, Emiliano Schena, and Massimo Sacchetti. *The importance of respiratory rate monitoring: From healthcare to sport and exercise*. 2020. DOI: 10.3390/s20216396.

[130] Jiapu Pan and Willis J. Tompkins. "A Real-Time QRS Detection Algorithm." In: *IEEE Transactions on Biomedical Engineering* BME-32.3 (Mar. 1985). Conference Name: IEEE Transactions on Biomedical Engineering, pp. 230–236. ISSN: 1558-2531. DOI: 10.1109/TBME.1985.325532. URL: https://ieeexplore.ieee.org/document/4122029 (visited on 01/07/2024).

[131] Stylianos Paraschiakos, Ricardo Cachucho, Matthijs Moed, Diana van Heemst, Simon Mooijaart, Eline P Slagboom, Arno Knobbe, and Marian Beekman. "Activity recognition using wearable sensors for tracking the elderly." In: *User Modeling and User-Adapted Interaction* 30.3 (July 2020), pp. 567–605.

[132] Pangi Park, Woobean Lee, and SeongHwan Cho. "An Adaptive Filter Based Motion Artifact Cancellation Technique Using Multi-Wavelength PPG for Accurate HR Estimation." In: *IEEE Transactions on Biomedical Circuits and Systems* 17.5 (Oct. 2023), pp. 1074–1083. DOI: 10.1109/TBCAS.2023.3315297. (Visited on 01/12/2024).

[133] Sungmee Park and Sundaresan Jayaraman. "The wearables revolution and Big Data: the textile lineage." In: *The Journal of The Textile Institute* 108.4 (Apr. 2017). Taylor & Francis _eprint: https://doi.org/10.1080/00405000.2016.1176632, pp. 605–614. ISSN: 0040-5000. DOI: 10.1080/00405000.2016.1176632. URL: https://doi.org/10.1080/00405000.2016.1176632 (visited on 10/19/2022).

[134] Christine N. Paulson, John T. Chang, Carlos E. Romero, Joseph Watson, Fred J. Pearce, and Nathan Levin. "Ultra-wideband radar methods and techniques of medical sensing and imaging." In: *Smart Medical and Biomedical Sensor Technology III*. Vol. 6007. 2005. DOI: 10.1117/12.630004.

[135] Marco A. F. Pimentel, Alistair E. W. Johnson, Peter H. Charlton, Drew Birrenkott, Peter J. Watkinson, Lionel Tarassenko, and David A. Clifton. "Toward a Robust Estimation of Respiratory Rate From Pulse Oximeters." In: *IEEE Transactions on Biomedical Engineering* 64.8 (2017), pp. 1914–1923. DOI: 10.1109/TBME.2016.2613124.

[136] Nuno Pinheiro, Ricardo Couceiro, Jorge Henriques, Jens Muehlsteff, Isabel Quintal, Lino Goncalves, and Paulo Carvalho. "Can PPG be used for HRV analysis?" In: *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE. 2016, pp. 2945–2949.

[137] Erika Pittella, Stefano Pisa, and Marta Cavagnaro. "Breath Activity Monitoring with Wearable UWB Radars: Measurement and Analysis of the Pulses Reflected by the Human Body." In: *IEEE Transactions on Biomedical Engineering* 63.7 (2016). ISSN: 15582531. DOI: 10.1109/TBME.2015.2496208.

[138] Erika Pittella, Blerina Zanaj, Stefano Pisa, and Marta Cavagnaro. "Measurement of Breath Frequency by Body-Worn UWB Radars: A Comparison among Different Signal Processing Techniques." In: *IEEE Sensors Journal* 17.6 (2017). ISSN: 1530437X. DOI: 10.1109/JSEN.2017.2654538.

[139] Thomas Plotz, Chen Chen, Nils Y. Hammerla, and Gregory D. Abowd. "Automatic Synchronization of Wearable Sensors and Video-Cameras for Ground Truth Annotation – A Practical Approach." In: *2012 16th International Symposium on Wearable Computers*. 2012, pp. 100–103. DOI: 10.1109/ISWC.2012.15.

[140] Carmen C. Y. Poon, Benny P. L. Lo, Mehmet Rasit Yuce, Akram Alomainy, and Yang Hao. "Body Sensor Networks: In the Era of Big Data and Beyond." In: *IEEE Reviews in Biomedical Engineering* 8 (2015). Conference Name: IEEE Reviews in Biomedical Engineering, pp. 4–16. ISSN: 1941-1189. DOI: 10.1109/RBME.2015.2427254. URL: https://ieeexplore.ieee.org/abstract/document/7096962 (visited on 02/28/2024).

[141] Hangwei Qian, Tian Tian, and Chunyan Miao. "What Makes Good Contrastive Learning on Small-Scale Wearable-Based Tasks?" In: *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. KDD '22. Washington DC, USA: Association for Computing Machinery, 2022, 3761–3771. ISBN: 9781450393850. DOI: 10.1145/3534678.3539134. URL: https://doi.org/10.1145/3534678.3539134.

[142] Mahfuzur Rahman and Bashir I. Morshed. "Estimation of Respiration Rate using an Inertial Measurement Unit Placed on Thorax-Abdomen." In: *IEEE International Conference on Electro Information Technology*. Vol. 2021-May. 2021. DOI: 10.1109/EIT51626.2021.9491900.

[143] Luigi Raiano, Joshua Di Tocco, Carlo Massaroni, Giovanni Di Pino, Emiliano Schena, and Domenico Formica. "Clean-Breathing: a Novel Sensor Fusion Algorithm Based on ICA to Remove Motion Artifacts from Breathing Signal." In: *2020 IEEE Metrology for Industry 4.0 & IoT*. June 2020, pp. 734–739.

DOI: 10.1109/MetroInd4.0IoT48571.2020.9138265. (Visited on 01/11/2024).

[144] Vincenzo Randazzo, Eros Pasero, and Silvio Navaretti. "VITAL-ECG: A portable wearable hospital." In: *2018 IEEE Sensors Applications Symposium (SAS)*. IEEE. 2018, pp. 1–6.

[145] Abbie Raphael. *ECS-.327-7-16-C-TR*. en-US. URL: https://ecsxtal.com/products/crystals/surface-mount-crystals/ecs-327-7-16-c-tr/ (visited on 06/02/2024).

[146] Daniel Ray, Tim Collins, Sandra I. Woolley, and Prasad V. S. Ponnapalli. "A Review of Wearable Multi-Wavelength Photoplethysmography." In: *IEEE Reviews in Biomedical Engineering* 16 (2023). Conference Name: IEEE Reviews in Biomedical Engineering, pp. 136–151. ISSN: 1941-1189. DOI: 10.1109/RBME.2021.3121476. URL: https://ieeexplore.ieee.org/abstract/document/9582790 (visited on 06/06/2024).

[147] Nir Regev and Dov Wulich. "Multi-modal, remote breathing monitor." In: *Sensors (Switzerland)* 20.4 (2020). ISSN: 14248220. DOI: 10.3390/s20041229.

[148] Attila Reiss, Ina Indlekofer, Philip Schmidt, and Kristof Van Laerhoven. "Deep PPG: Large-Scale Heart Rate Estimation with Convolutional Neural Networks." In: *Sensors* 19.14 (2019). ISSN: 1424-8220. DOI: 10.3390/s19143079. URL: https://www.mdpi.com/1424-8220/19/14/3079.

[149] Carl C Rheinländer and Norbert Wehn. "Precise synchronization time stamp generation for Bluetooth low energy." In: *2016 IEEE SENSORS*. IEEE. 2016, pp. 1–3.

[150] Thomas Rietveld, Riemer J. K. Vegter, Rienk M. A. van der Slikke, Aldo E. Hoekstra, Lucas H. V. van der Woude, and Sonja de Groot. "Wheelchair mobility performance of elite wheelchair tennis players during four field tests: Inter-trial reliability and construct validity." In: *PLOS ONE* 14.6 (June 2019), pp. 1–16. DOI: 10.1371/journal.pone.0217514. URL: https://doi.org/10.1371/journal.pone.0217514.

[151] Leandro Giacomini Rocha, Guilherme Paim, Dwaipayan Biswas, Sergio Bampi, Francky Catthoor, Chris Van Hoof, and Nick Van Helleputte. "LSTM-only Model for Low-complexity HR Estimation from Wrist PPG." In: *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. 2021, pp. 1068–1071. DOI: 10.1109/EMBC46164.2021.9630942.

[152] STMicroelectronics. *LIS2DH, MEMS digital output motion sensor: ultra low-power high performance 3-axis "femto" accelerometer*. https://www.st.com/resource/en/datasheet/lis2dh.pdf. Last accessed: 2022-12-09. 2011.

[153] Seyed M. A. Salehizadeh, Duy Dao, Jeffrey Bolkhovsky, Chae Cho, Yitzhak Mendelson, and Ki H. Chon. "A Novel Time-Varying Spectral Filtering Algorithm for Reconstruction of Motion Artifact Corrupted Heart Rate Signals During Intense Physical Activities Using a Wearable Photoplethysmogram Sensor." In: *Sensors* 16.1 (2016). ISSN: 1424-8220. DOI: 10.3390/s16010010. URL: https://www.mdpi.com/1424-8220/16/1/10.

[154] Cheyenne Samson and Ahyeon Koh. "Stress monitoring and recent advancements in wearable biosensors." In: *Frontiers in bioengineering and biotechnology* 8 (2020), p. 1037.

[155] Fatemeh Sarhaddi, Kianoosh Kazemi, Iman Azimi, Rui Cao, Hannakaisa Niela-Vilén, Anna Axelin, Pasi Liljeberg, and Amir M. Rahmani. "A comprehensive accuracy assessment of Samsung smartwatch heart rate and heart rate variability." In: *PLOS ONE* 17.12 (Dec. 2022), pp. 1–19. DOI: 10.1371/journal.pone.0268361. URL: https://doi.org/10.1371/journal.pone.0268361.

[156] Pritam Sarkar and Ali Etemad. "CardioGAN: Attentive Generative Adversarial Network with Dual Discriminators for Synthesis of ECG from PPG." In: *Proceedings of the AAAI Conference on Artificial Intelligence* 35.1 (May 2021), pp. 488–496. DOI: 10.1609/aaai.v35i1.16126. URL: https://ojs.aaai.org/index.php/AAAI/article/view/16126.

[157] Lawrence Saul and Jont Allen. "Periodic Component Analysis: An Eigenvalue Method for Representing Periodic Structure in Speech." In: *Advances in Neural Information Processing Systems*. Ed. by T. Leen, T. Dietterich, and V. Tresp. Vol. 13. MIT Press, 2000. URL: https://proceedings.neurips.cc/paper_files/paper/2000/file/f22e4747da1aa27e363d86d40ff442fe-Paper.pdf.

[158] Elliott Schires, Pantelis Georgiou, and Tor Sverre Lande. "Vital sign monitoring through the back using an UWB impulse radar with body coupled antennas." In: *IEEE Transactions on Biomedical Circuits and Systems* 12.2 (2018). ISSN: 19324545. DOI: 10.1109/TBCAS.2018.2799322.

[159] Philip Schmidt, Attila Reiss, Robert Duerichen, Claus Marberger, and Kristof Van Laerhoven. "Introducing WESAD, a Multimodal Dataset for Wearable Stress and Affect Detection." In: *Proceedings of the 20th ACM International Conference on Multimodal Interaction*. ICMI '18. Boulder, CO, USA: Association for Computing Machinery, 2018, 400–408. ISBN: 9781450356923. DOI: 10.1145/3242969.3242985. URL: https://doi.org/10.1145/3242969.3242985.

[160] Felix Scholkmann, Jens Boss, and Martin Wolf. "An Efficient Algorithm for Automatic Peak Detection in Noisy Periodic and Quasi-Periodic Signals." en. In: *Algorithms* 5.4 (Dec. 2012). Number: 4 Publisher: Multidisciplinary Digital Publishing Institute, pp. 588–603. ISSN: 1999-4893. DOI: 10.3390/a5040588. URL: https://www.mdpi.com/1999-4893/5/4/588 (visited on 05/30/2024).

[161] Óscar Seijo, Jesús Alberto López-Fernández, Hans-Peter Bernhard, and Iñaki Val. "Enhanced Timestamping Method for Subnanosecond Time Synchronization in IEEE 802.11 Over WLAN Standard Conditions." In: *IEEE Transactions on Industrial Informatics* 16.9 (2020), pp. 5792–5805. DOI: 10.1109/TII.2019.2959200.

[162] Bosch Sensortec. *BME280, Combined humidity and pressure sensor*. https://www.bosch-sensortec.com/media/boschsensortec/downloads/datasheets/bst-bme280-ds002.pdf. Last accessed: 2022-12-09. 2022.

[163] Bosch Sensortec. *BMP581, Barometric Pressure Sensor*. https://www.bosch-sensortec.com/media/boschsensortec/downloads/datasheets/bst-bmp581-ds004.pdf. Last accessed: 2022-12-09. 2022.

[164] Shaban Shabani, Alan K. Bourke, Amir Muaremi, Jens Praestgaard, Kate O'Keeffe, Rob Argent, Martin Brom, Celeste Scotti, Brian Caulfield, and Lorcan C. Walsh. "An Automatic Foot and Shank IMU Synchronization Algorithm: Proof-of-concept." In: *2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. ISSN: 2694-0604. July 2022, pp. 4210–4213. DOI: 10.1109/EMBC48229.2022.9871162.

[165] Takunori Shimazaki, Shinsuke Hara, Hiroyuki Okuhata, Hajime Nakamura, and Takashi Kawabata. "Cancellation of motion artifact induced by exercise for PPG-based heart rate sensing." In: *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE. 2014, pp. 3216–3219.

[166] Takunori Shimazaki, Yuzu Kuwahara, Masaya Kimoto, Shinsuke Hara, and Hiroyuki Yomo. "Effect of Position and Fastening Belt Pressure on the Accuracy of PPG-Based Heart Rate Sensor." In: *IEEE EMBC*. July 2018, pp. 4323–4326. DOI: 10.1109/EMBC.2018.8513333. (Visited on 01/12/2024).

[167] M.L. Sichitiu and C. Veerarittiphan. "Simple, accurate time synchronization for wireless sensor networks." In: *2003 IEEE Wireless Communications and Networking, 2003. WCNC 2003*. Vol. 2. ISSN: 1525-3511. Mar. 2003, 1266–1273 vol.2. DOI: 10.1109/WCNC.2003.1200555.

[168] Kasun Somaratne, F John Dian, and Amirhossein Yousefi. "Accuracy analysis of time synchronization using current consumption pattern of BLE devices." In: *2018 IEEE 8th Annual Computing and Communication Workshop and Conference (CCWC)*. IEEE. 2018, pp. 841–844.

[169] Zhipeng Song, Zhichao Cao, Zhenjiang Li, Jiliang Wang, and Yunhao Liu. "Inertial motion tracking on mobile and wearable devices: Recent advancements and challenges." In: *Tsinghua Science and Technology* 26.5 (2021), pp. 692–705.

[170] Andreas Spilz and Michael Munz. *Novel Approach To Synchronisation Of Wearable IMUs Based On Magnetometers*. arXiv:2107.03147 [cs, eess]. July 2021. DOI: 10.48550/arXiv.2107.03147. URL: http://arxiv.org/abs/2107.03147 (visited on 10/21/2022).

[171] Sabarish Sridhar, Prasant Misra, and Jay Warrior. "CheepSync: A Time Synchronization Service for Resource Constrained Bluetooth Low Energy Advertisers." In: *CoRR* abs/1501.06479 (2015). arXiv: 1501.06479. URL: http://arxiv.org/abs/1501.06479.

[172] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. "Dropout: A Simple Way to Prevent Neural Networks from Overfitting." In: *Journal of Machine Learning Research* 15.56 (2014), pp. 1929–1958. URL: http://jmlr.org/papers/v15/srivastava14a.html.

[173] Enrico M. Staderini. "UWB radars in medicine." In: *IEEE Aerospace and Electronic Systems Magazine* 17.1 (2002). ISSN: 08858985. DOI: 10.1109/62.978359.

[174] Daniel Stoller, Sebastian Ewert, and Simon Dixon. "Wave-U-Net: A Multi-Scale Neural Network for End-to-End Audio Source Separation." In: *ArXiv* abs/1806.03185 (2018). URL: https://api.semanticscholar.org/CorpusID:47015908.

[175] Daniel Stoller, Sebastian Ewert, and Simon Dixon. "Wave-u-net: A multi-scale neural network for end-to-end audio source separation." In: *arXiv preprint arXiv:1806.03185* (2018).

[176] Tessa Strain, Katrien Wijndaele, Paddy C Dempsey, Stephen J Sharp, Matthew Pearce, Justin Jeon, Tim Lindsay, Nick Wareham, and Søren Brage. "Wearable-device-measured physical activity and future health risk." In: *Nature Medicine* 26.9 (Sept. 2020), pp. 1385–1391.

[177] Bharath Sundararaman, Ugo Buy, and Ajay D. Kshemkalyani. "Clock synchronization for wireless sensor networks: a survey." en. In: *Ad Hoc Networks* 3.3 (May 2005), pp. 281–323. DOI: 10.1016/j.adhoc.2005.01.002. URL: https:

// www . sciencedirect . com / science / article / pii / S1570870505000144 (visited on 12/09/2022).

[178] Andriy Temko. "Accurate Heart Rate Monitoring During Physical Exercises Using PPG." In: *IEEE Transactions on Biomedical Engineering* 64.9 (Sept. 2017), pp. 2016–2024. DOI: 10.1109/TBME.2017.2676243. (Visited on 01/12/2024).

[179] Andriy Temko. "Accurate Heart Rate Monitoring During Physical Exercises Using PPG." In: *IEEE Transactions on Biomedical Engineering* 64.9 (2017), pp. 2016–2024. DOI: 10.1109/TBME.2017.2676243.

[180] Francisco Tirado-Andrés and Alvaro Araujo. "Performance of clock sources and their influence on time synchronization in wireless sensor networks." In: *International Journal of Distributed Sensor Networks* 15 (Sept. 2019), p. 155014771987937. DOI: 10.1177/1550147719879372.

[181] Sara Usmani, Abdul Saboor, Muhammad Haris, Muneeb A Khan, and Heemin Park. "Latest research trends in fall detection and prevention using machine learning: A systematic review." In: *Sensors* 21.15 (2021), p. 5134.

[182] Enrique Valero, Aparajithan Sivanathan, Frédéric Bosché, and Mohamed Abdel-Wahab. "Musculoskeletal disorders in construction: A review and a novel system for activity tracking with body area network." In: *Applied Ergonomics* 54 (May 2016), pp. 120–130. ISSN: 0003-6870. DOI: 10.1016/j.apergo.2015.11.020. URL: https://www.sciencedirect.com/science/article/pii/S0003687015301204 (visited on 05/01/2024).

[183] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. "Attention is All you Need." In: *Advances in Neural Information Processing Systems*. Ed. by I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett. Vol. 30. Curran Associates, Inc., 2017. URL: https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.

[184] Adriana N. Vest, Giulia Da Poian, Qiao Li, Chengyu Liu, Shamim Nemati, Amit J. Shah, and Gari D. Clifford. "An open source benchmarked toolbox for cardiovascular waveform and interval analysis." en. In: *Physiological Measurement* 39.10 (Oct. 2018). Publisher: IOP Publishing, p. 105004. ISSN: 0967-3334. DOI: 10.1088/1361-6579/aae021. URL: https://dx.doi.org/10.1088/1361-6579/aae021 (visited on 05/30/2024).

[185] R Vetter, L Rossini, A Ridolfi, J Sola, O Chetelat, M Correvon, and J Krauss. "Frequency domain SpO2 estimation based on multichannel photoplethysmographic measurements at the

sternum." In: *World Congress on Medical Physics and Biomedical Engineering, September 7-12, 2009, Munich, Germany: Vol. 25/4 Image Processing, Biosignal Processing, Modelling and Simulation, Biomechanics*. Springer. 2010, pp. 326–329.

[186] Nicolas de Vicq, Frederic Robert, Julien Penders, Bert Gyselinckx, and Tom Torfs. "Wireless Body Area Network for Sleep Staging." In: *2007 IEEE Biomedical Circuits and Systems Conference*. ISSN: 2163-4025. Nov. 2007, pp. 163–166. DOI: 10.1109/BIOCAS.2007.4463334. URL: https://ieeexplore.ieee.org/abstract/document/4463334 (visited on 05/01/2024).

[187] Roberto Vio and Willem Wamsteker. "Limits of the Cross-Correlation Function in the Analysis of Short Time Series." en. In: *Publications of the Astronomical Society of the Pacific* 113.779 (Jan. 2001). The University of Chicago Press, p. 86. ISSN: 1538-3873. DOI: 10.1086/317967. URL: https://dx.doi.org/10.1086/317967 (visited on 12/07/2022).

[188] Dingyang Wang, Sungwon Yoo, and Sung Ho Cho. "Experimental comparison of ir-uwb radar and fmcw radar for vital signs." In: *Sensors (Switzerland)* 20.22 (2020). ISSN: 14248220. DOI: 10.3390/s20226695.

[189] He Wang, Jianan Li, Benjamin E McDonald, Todd R Farrell, Xinming Huang, and Edward A Clancy. "Comparison between Two Time Synchronization and Data Alignment Methods for Multi-Channel Wearable Biosensor Systems Using BLE Protocol." In: *Sensors* 23.5 (2023), p. 2465.

[190] Tianben Wang, Daqing Zhang, Leye Wang, Yuanqing Zheng, Tao Gu, Bernadette Dorizzi, and Xingshe Zhou. "Contactless respiration monitoring using ultrasound signal with off-the-shelf audio devices." In: *IEEE Internet of Things Journal* 6.2 (2019). ISSN: 23274662. DOI: 10.1109/JIOT.2018.2877607.

[191] Yan Wang et al. "A systematic review on affective computing: emotion models, databases, and recent advances." In: *Information Fusion* 83-84 (July 2022), pp. 19–52. ISSN: 1566-2535. DOI: 10.1016/j.inffus.2022.03.009. URL: https://www.sciencedirect.com/science/article/pii/S1566253522000367 (visited on 05/01/2024).

[192] Kristen M. Warren, Joshua R. Harvey, Ki H. Chon, and Yitzhak Mendelson. "Improving Pulse Rate Measurements during Random Motion Using a Wearable Multichannel Reflectance Photoplethysmograph." en. In: *Sensors* 16.3 (Mar. 2016), p. 342. DOI: 10.3390/s16030342. (Visited on 12/21/2023).

[193] Tobias Wartzek, Christoph Brüser, Marian Walter, and Steffen Leonhardt. "Robust Sensor Fusion of Unobtrusively Measured Heart Rate." In: *IEEE Journal of Biomedical and Health Informatics*

18.2 (Mar. 2014). DOI: 10.1109/JBHI.2013.2274211. (Visited on 01/06/2024).

[194] RWCGR Wijshoff, Jeroen Veen, Alexander M Van der Lee, Lars Mulder, Marco Stijnen, Sjoerd Van Tuijl, and Ronald M Aarts. "PPG motion artifact handling using a self-mixing interferometric sensor." In: *Proc. SPIE*. Vol. 7894. 2011, 78940F.

[195] Harry J Witchel et al. "Thigh-Derived Inertial Sensor Metrics to Assess the Sit-to-Stand and Stand-to-Sit Transitions in the Timed Up and Go (TUG) Task for Quantifying Mobility Impairment in Multiple Sclerosis." en. In: *Front Neurol* 9 (Sept. 2018), p. 684.

[196] Jing Wu, He Li, Sherri Cheng, and Zhangxi Lin. "The promising future of healthcare services: When big data analytics meets wearable technology." In: *Information & Management* 53.8 (2016), pp. 1020–1033.

[197] Taiyang Wu, Fan Wu, Chunkai Qiu, Jean-Michel Redouté, and Mehmet Rasit Yuce. "A Rigid-Flex Wearable Health Monitoring Sensor Patch for IoT-Connected Healthcare Applications." In: *IEEE Internet of Things Journal* 7.8 (Aug. 2020). Conference Name: IEEE Internet of Things Journal, pp. 6932–6945. ISSN: 2327-4662. DOI: 10.1109/JIOT.2020.2977164. URL: https://ieeexplore.ieee.org/document/9018016 (visited on 02/22/2024).

[198] Rasoul Yousefi, Mehrdad Nourani, Sarah Ostadabbas, and Issa Panahi. "A motion-tolerant adaptive algorithm for wearable photoplethysmographic biosensors." In: *IEEE journal of biomedical and health informatics* 18.2 (2013), pp. 670–681.

[199] Zhihan Yue, Yujing Wang, Juanyong Duan, Tianmeng Yang, Congrui Huang, Yunhai Tong, and Bixiong Xu. "TS2Vec: Towards Universal Representation of Time Series." In: *Proceedings of the AAAI Conference on Artificial Intelligence* 36.8 (2022), pp. 8980–8987. DOI: 10.1609/aaai.v36i8.20881. URL: https://ojs.aaai.org/index.php/AAAI/article/view/20881.

[200] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stephane Deny. "Barlow Twins: Self-Supervised Learning via Redundancy Reduction." In: *Proceedings of the 38th International Conference on Machine Learning*. Ed. by Marina Meila and Tong Zhang. Vol. 139. Proceedings of Machine Learning Research. PMLR, July 2021, pp. 12310–12320. URL: https://proceedings.mlr.press/v139/zbontar21a.html.

[201] Yun C. Zhang, Shibo Zhang, Miao Liu, Elyse Daly, Samuel Battalio, Santosh Kumar, Bonnie Spring, James M. Rehg, and Nabil Alshurafa. "SyncWISE: Window Induced Shift Estimation for Synchronization of Video and Accelerometry from

Wearable Sensors." In: *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 4.3 (Sept. 2020). DOI: 10.1145/3411824. URL: https://doi.org/10.1145/3411824.

[202] Yun C Zhang, Shibo Zhang, Miao Liu, Elyse Daly, Samuel Battalio, Santosh Kumar, Bonnie Spring, James M Rehg, and Nabil Alshurafa. "Syncwise: Window induced shift estimation for synchronization of video and accelerometry from wearable sensors." In: *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4.3 (2020), pp. 1–26.

[203] Zhilin Zhang. "Photoplethysmography-Based Heart Rate Monitoring in Physical Activities via Joint Sparse Spectrum Reconstruction." In: *IEEE Transactions on Biomedical Engineering* 62.8 (2015), pp. 1902–1910. DOI: 10.1109/TBME.2015.2406332.

[204] Zhilin Zhang, Zhouyue Pi, and Benyuan Liu. "TROIKA: A General Framework for Heart Rate Monitoring Using Wrist-Type Photoplethysmographic Signals During Intensive Physical Exercise." In: *IEEE Transactions on Biomedical Engineering* 62.2 (2015), pp. 522–531. DOI: 10.1109/TBME.2014.2359372.

[205] Ya-Li Zheng, Xiao-Rong Ding, Carmen Chung Yan Poon, Benny Ping Lai Lo, Heye Zhang, Xiao-Lin Zhou, Guang-Zhong Yang, Ni Zhao, and Yuan-Ting Zhang. "Unobtrusive sensing and wearable devices for health informatics." en. In: *IEEE Trans Biomed Eng* 61.5 (May 2014), pp. 1538–1554.

[206] R Cassani et al. "Optimal filter characterization for photoplethysmography-based pulse rate and pulse power spectrum estimation." In: *2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. 2020.