



Politecnico di Bari

Repository Istituzionale dei Prodotti della Ricerca del Politecnico di Bari

Embedded Electronics for Multimodal Biomedical Signal Acquisition and Processing

This is a PhD Thesis

Original Citation:

Embedded Electronics for Multimodal Biomedical Signal Acquisition and Processing / Botrugno, Chiara. - ELETTRONICO. - (2026). [10.60576/poliba/iris/botrugno-chiara_phd2026]

Availability:

This version is available at <http://hdl.handle.net/11589/295280> since: 2026-01-08

Published version

DOI:10.60576/poliba/iris/botrugno-chiara_phd2026

Publisher: Politecnico di Bari

Terms of use:

(Article begins on next page)



POLITECNICO DI BARI
DEPARTMENT OF ELECTRICAL AND INFORMATION ENGINEERING
DOCTORAL PROGRAMME IN ELECTRICAL AND INFORMATION
ENGINEERING

EMBEDDED ELECTRONICS FOR MULTIMODAL BIOMEDICAL SIGNAL ACQUISITION AND PROCESSING

Doctoral Dissertation of:

Chiara Botrugno **Ing. Chiara Botrugno**

Supervisor:

Prof. Ing. Francesco Dell'Olio *Francesco Dell'Olio*

Tutors:

Prof. Ing. Stefano Mazzoleni

Ing. Francesco Conversano

Stefano Mazzoleni
Francesco Conversano

The Chair of the Doctoral Program:

Prof. Ing. Mario Carpentieri

Mario Carpentieri

2025 – XXXVIII Cycle

Abstract

THE present work presents the design, implementation, and validation of embedded electronic systems for multimodal biomedical signal acquisition and processing, with a focus on cardiovascular and pulmonary monitoring. The work addresses the growing demand for compact, low-power, and intelligent hardware platforms capable of performing continuous physiological measurements in real-world settings. Through two complementary case studies—cuffless blood pressure estimation from photoplethysmographic (PPG) signals and automatic B-line detection in lung ultrasound (LUS) imaging—the thesis demonstrates how the co-design of hardware, signal processing, and machine learning can enable clinically meaningful performance under embedded constraints.

The first part of the research focuses on the development of a real-time optical acquisition framework for non-invasive blood pressure monitoring. The system was based on the MAX86916 Evaluation System by Analog Devices, integrating a multi-wavelength optical sensor (blue, green, red, and infrared LEDs) with a high-sensitivity photodiode and a MAX32630FTHR ARM Cortex-M4F microcontroller. Data were sampled at 100 Hz across all wavelengths, preprocessed for noise reduction, and synchronized with reference oscillometric blood pressure measurements. Three machine learning models were evaluated for systolic and diastolic blood pressure estimation: (1) a Multi-Layer Perceptron (MLP) trained on 84 handcrafted temporal and spectral features, (2) a PCA-reduced MLP with 10 principal components, and (3) a Convolutional Neural Network (CNN) trained on 21 features derived from the L2-norm of the PPG signal. In addition, a hybrid CNN–Support Vector Regression (SVR) model was developed, combining deep feature extraction with kernel-based regression.

Experimental validation on a dedicated dataset (collected from 81 healthy volunteers) demonstrated that the hybrid CNN–SVR model achieved the best overall performance, reaching a mean absolute error (MAE) of 1.84 mmHg for systolic and 1.98 mmHg for diastolic pressure estimation, with standard deviations of 3.27 mmHg and 3.47 mmHg, respectively. These results satisfy the AAMI/ISO 81060-2 accuracy criteria for non-invasive blood pressure measurement and outperform both traditional regression-based and single-feature models.

The second part of the thesis targets automated pulmonary analysis through the de-

tection of B-line artifacts in lung ultrasound images—vertical, hyperechoic reverberations that indicate interstitial fluid accumulation. A new algorithm based on Variational Mode Decomposition (VMD) was developed to enhance and segment B-lines by decomposing ultrasound images into intrinsic mode components characterized by distinct frequency and orientation profiles. The proposed block-wise VMD approach enabled efficient parallel processing and reduced computational cost by 40% compared with standard two-dimensional VMD implementations. The algorithm was validated on a clinical dataset of 116 lung ultrasound frames, annotated by expert radiologists. It achieved a precision of 0.87, recall of 0.79, F1-score of 0.83, and a Dice coefficient of 0.79 when compared to the manual ground truth. Compared with deep learning-based models, the proposed method maintained comparable accuracy but required over 10× less memory and 3× faster inference time, confirming its suitability for embedded deployment.

Collectively, these results validate the feasibility of achieving clinically relevant accuracy and efficiency in biomedical signal processing using energy-constrained embedded systems. By integrating sensing hardware, signal conditioning, and data-driven modeling, the thesis establishes a generalizable design framework applicable to a wide range of physiological monitoring tasks. The demonstrated systems—based on real hardware and validated under realistic conditions—bridge the gap between laboratory prototypes and practical, wearable medical technologies.

Ultimately, this work advances the field of embedded biomedical engineering by showing that with appropriate co-design methodologies, multi-modal sensing and edge-level intelligence can be effectively combined to realize continuous, interpretable, and personalized healthcare solutions.

Summary

THIS thesis explores the development of embedded electronic systems for multimodal biomedical signal acquisition and processing, addressing the growing need for continuous and decentralized healthcare. The research is grounded in the convergence of microelectronics, sensing technology, and artificial intelligence, and demonstrates how real-time physiological monitoring can be achieved within the severe energy and computational constraints typical of wearable devices.

Two main application domains were investigated: (1) cuffless, non-invasive blood pressure estimation using optical photoplethysmography (PPG), and (2) automatic detection of B-line artifacts in lung ultrasound (LUS) images for pulmonary assessment. Despite their different sensing modalities, both studies embody the same design philosophy—integrating sensor optimization, signal processing, and machine learning into compact, low-power embedded frameworks.

In the first research line, a multi-wavelength PPG acquisition system was realized using the MAX86916 Evaluation System, which integrates optical LEDs and a high-sensitivity photodiode, coupled with a MAX32630FTHR microcontroller for data control and transmission. The system was configured to record synchronized optical signals at four wavelengths (blue, green, red, and infrared) with a sampling frequency of 100 Hz. A dedicated experimental protocol was defined for dataset creation, including calibration, ambient-light rejection, and synchronization with reference oscillometric blood pressure measurements.

Three machine learning models were developed to analyze these signals. The first was a Multi-Layer Perceptron (MLP) trained on 84 handcrafted temporal and spectral features. The second, a PCA-MLP variant, reduced the dimensionality to 10 principal components to improve generalization. The third, a Convolutional Neural Network (CNN), processed 21 L2-norm-derived features to capture spatial-temporal dependencies. Finally, a hybrid CNN-SVR model combined convolutional feature extraction with support-vector regression to enhance robustness across subjects. Results demonstrated that the hybrid model achieved the lowest mean absolute error in systolic and diastolic pressure estimation, confirming the effectiveness of hybrid architectures for embedded physiological prediction tasks.

The second research line focused on point-of-care lung ultrasound, a radiation-free

imaging modality essential for diagnosing pulmonary edema, pneumonia, and interstitial syndromes. Manual interpretation of LUS images is time-consuming and subjective; therefore, this thesis developed an automatic B-line detection algorithm based on Variational Mode Decomposition (VMD). The proposed block-wise VMD method decomposes ultrasound images into frequency-limited sub-components, isolating the high-frequency vertical modes that characterize B-lines. The algorithm improves artifact contrast and noise suppression while maintaining interpretability. Experimental validation on annotated datasets showed high detection accuracy and significant computational efficiency compared with conventional filters and deep-learning baselines. Its block-wise design makes it suitable for integration into portable ultrasound systems.

Beyond the individual applications, the thesis contributes a methodological framework for hardware-aware biomedical signal processing. It demonstrates how knowledge of analog front-end design, signal quality constraints, and embedded hardware capabilities can guide algorithmic optimization, ensuring real-time operation with minimal energy use. Both systems achieved low-power performance compatible with wearable or portable operation, validating the feasibility of edge intelligence in healthcare devices.

The results confirm that embedded systems can deliver clinically meaningful insights—such as heart rate, pulse transit time, oxygen saturation, and even blood pressure estimation—without relying on high-power computing resources. Similarly, the VMD-based ultrasound framework provides a pathway toward automated, interpretable pulmonary monitoring on point-of-care hardware.

In conclusion, this thesis advances the state of the art in embedded biomedical signal processing by presenting two complementary, experimentally validated frameworks. It demonstrates that through co-design of hardware and algorithms, high-fidelity biosignal acquisition and intelligent processing can coexist within constrained embedded platforms. These findings pave the way for future wearable and portable healthcare devices that combine multimodal sensing, local intelligence, and cloud connectivity, contributing to a new generation of personalized and accessible medical technologies.

Contents

1	Introduction	1
1.1	Motivation and Context	1
1.1.1	From E-health to Pervasive Sensing	2
1.1.2	Cardiovascular Health and Non-invasive BP Monitoring	4
1.1.3	Pulmonary Health and Point-of-Care Ultrasound	6
1.1.4	Toward Hardware-Constrained Biomedical Signal Processing	8
1.2	Research Objectives	10
1.3	Contribution of the Thesis	12
1.3.1	Specific Contributions to Non-invasive Cardiovascular Monitoring	12
1.3.2	Specific Contributions to Pulmonary Assessment with Lung Ultrasound	13
1.3.3	Cross-Cutting and Methodological Contributions	13
1.4	Structure of the Thesis	14
2	Embedded Hardware for Physiological Signal Acquisition	17
2.1	Background on Biopotential and Biomarker Monitoring	17
2.1.1	Electrophysiological Signals and Biopotentials	17
2.1.2	Optical Biomarkers and Photoplethysmography	19
2.1.3	Chemical Biomarkers and Non-Invasive Sensing	19
2.1.4	Multimodal and Embedded Physiological Monitoring	20
2.1.5	Experimental Implementation of ECG–PPG Fusion for Multiparametric Vital Sign Monitoring	20
2.2	Embedded Hardware Platforms and Constraints	22
2.2.1	System Architecture of Embedded Biomedical Devices	22
2.2.2	Analog Front-End Design Considerations	23
2.2.3	Power Consumption and Energy Management	23
2.2.4	Processing Architectures and Computational Constraints	24
2.2.5	Communication and Data Transmission Constraints	24
2.2.6	Trade-Offs and Co-Design Strategies	25
2.3	Signal Processing for Wearable and Resource-Constrained Devices	25
2.3.1	Characteristics of Physiological Signals and Sources of Noise	26

Contents

2.3.2	Adaptive and Nonlinear Signal Processing	26
2.3.3	Feature Extraction and Dimensionality Reduction	27
2.3.4	Machine Learning and Embedded Intelligence	27
2.3.5	Real-Time Processing and Edge Optimization	28
2.3.6	Challenges and Future Directions	28
3	Blood Pressure Prediction Using Multi-Wavelength PPG	30
3.1	Background on Non-Invasive Blood Pressure Monitoring	30
3.1.1	Principles of Cuffless and Non-Invasive BP Estimation	31
3.1.2	Photoplethysmography (PPG) and Its Relationship with Blood Pressure	31
3.1.3	Toward Continuous and Embedded BP Monitoring	33
3.2	Hardware Architecture and Wearable Implementation	34
3.2.1	System Overview and Design Objectives	34
3.2.2	Optical Front-End and Sensing Module	35
3.2.3	Embedded Control and Processing Unit	36
3.2.4	Power Supply and Energy Management	36
3.2.5	Data Acquisition Workflow and Software Interface	37
3.2.6	Mechanical Configuration and Experimental Setup	37
3.2.7	System Validation and Performance Evaluation	38
3.3	Dataset Creation and Experimental Protocol	38
3.3.1	Optical and Electronic Configuration	39
3.3.2	Experimental Protocol and Acquisitions	40
3.4	Signal Processing and Feature Extraction	43
3.4.1	Features	45
3.5	Machine Learning Models: MLP, PCA, and CNN Approaches	45
3.6	Hybrid CNN-SVR Model for Optimized BP Estimation	47
3.7	Experimental Results and Discussion	48
3.7.1	Multi-Layer Perceptron - All features	49
3.7.2	Multi-Layer Perceptron - Feature Reduction	49
3.7.3	Convolutional Neural Network	50
3.7.4	Hybrid Model: CNN and SVR	52
3.7.5	Additional Model Evaluations and Ablation Studies	54
3.7.6	Limitations and Future Work	56
4	Automatic B-Lines Detection in Lung Ultrasound	59
4.1	Clinical Background and Relevance of B-Lines	59
4.2	Lung Ultrasound Acquisition Protocol	63
4.3	Ultrasound Front-End Acquisition Architecture	65
4.4	Variational Mode Decomposition (VMD) for Ultrasound Signal Analysis	67
4.5	Block-Wise VMD Algorithm Design and Implementation	68
4.5.1	Preprocessing	69
4.5.2	Variational Decomposition	71
4.5.3	Mode Analysis and Allocation	71
4.5.4	Reconstruction and Integration	73
4.6	Experimental Validation, Results and Discussion	73
4.6.1	Quantitative Evaluation Protocol	74

4.6.2	Results and Quantitative Performance	74
4.6.3	Comparative Performance Analysis	74
4.7	Discussion and Conclusions	76
4.7.1	Concluding Remarks	77
5	Conclusions and Future Work	79
5.1	Summary of Main Contributions	79
5.2	Limitations of the Current Work	81
5.2.1	Experimental and Dataset Constraints	82
5.2.2	Hardware and Implementation Constraints	82
5.2.3	Methodological Assumptions	83
5.2.4	Clinical and Translational Limitations	84
5.3	Potential Extensions and Future Directions	85
5.3.1	Embedded Deployment and Optimization	85
5.3.2	Clinical Validation and Scalability	86
	Bibliography	89

CHAPTER 1

Introduction

1.1 Motivation and Context

Over the past three decades, the convergence of microelectronics, telecommunication networks, and data science has driven a profound transformation of modern healthcare [1]. The term e-health has come to denote the use of information and communication technologies (ICT) to deliver or enhance health services, encompassing remote consultation, tele-monitoring, decision-support systems, and integration of patient-generated data into clinical workflows [2]. The global expansion of broadband internet and the proliferation of smartphones have created an unprecedented opportunity for ubiquitous health monitoring, shifting the emphasis from hospital-centered, episodic care to continuous and proactive management of health [3].

The paradigm shift in global healthcare is rapidly moving towards *Pervasive*, *Predictive*, and *Personalized Medicine* (P3M) [4]. This evolution is driven by the imperative to transition from episodic, reactive clinical care to continuous, proactive health monitoring that is accessible outside of traditional hospital settings. Achieving this vision hinges on the successful integration of embedded electronics with advanced methodologies for multimodal biomedical signal acquisition and processing. The central technological challenge is the development of computational frameworks that can extract reliable, clinical-grade information from complex biological data streams while operating within the critical constraints of real-time responsiveness, low power consumption, and physical portability inherent to modern wearable and Point-of-Care (PoC) devices.

This Doctoral thesis is fundamentally motivated by this technological and clinical mandate. It addresses this challenge through two distinct, yet technically interdependent, lines of research, both unified under the theme of Embedded Electronics for Multimodal Biomedical Signal Acquisition and Processing.

1.1.1 From E-health to Pervasive Sensing

One of the key enablers of this paradigm shift has been the miniaturization of sensing hardware. Semiconductor photodetectors, low-noise front-end circuits, and high-resolution analog-to-digital converters have made it feasible to embed biomedical sensors in wearable devices such as smart watches, adhesive patches, or portable ultrasound probes. Low-power microcontrollers and embedded digital-signal processors (DSPs) allow real-time processing at the edge, often on coin-cell batteries, reducing the reliance on external computing infrastructure [5]. This trend supports the so-called “edge-to-cloud” continuum in e-health: data are pre-processed near the sensor to reduce bandwidth and latency, while long-term storage, large-scale analytics, and population studies can leverage cloud resources [6].

Such progress has led to significant advances in the continuous monitoring of vital signs such as heart rate, blood oxygen saturation, respiratory rate, and body temperature. More recently, non-invasive cardiovascular monitoring—especially the measurement of arterial blood pressure (BP)—and bed-side lung assessment have attracted intense research attention. These two application domains are central to this thesis, as they illustrate complementary aspects of the same challenge: acquiring reliable physiological information from noisy biological signals using resource-constrained embedded platforms.

One of the most transformative enablers of the current e-health revolution has been the relentless miniaturization and integration of sensing hardware. The continuous progress in semiconductor device technology, particularly in complementary metal–oxide semiconductor (CMOS) photodetectors and low-noise analog front-end circuits, has enabled biomedical signals—once measurable only with bulky laboratory instruments—to be captured by compact, lightweight devices [7]. High-resolution analog-to-digital converters (ADCs) with sampling rates in the hundreds of kilohertz to megahertz range and with input-referred noise levels of just a few microvolts can now be fabricated on a single chip together with signal-conditioning amplifiers, filters, and digital control logic. These advances have made it feasible to integrate a variety of biosensors—optical, electrical, acoustic, and mechanical—into wearable or even disposable form factors, including smart watches, wristbands, adhesive patches, clothing-integrated sensors, and portable ultrasound probes. Moreover, recent developments in flexible and stretchable electronics have opened the door to skin-conformal devices that can capture signals over long periods with minimal discomfort to the user [8].

Equally important has been the development of low-power microcontrollers (MCUs) and embedded digital-signal processors (DSPs) that can execute non-trivial signal-processing and machine-learning algorithms on the edge, i.e., directly where the data are acquired [9]. Modern MCUs often combine a 32-bit RISC core with specialized floating-point units and SIMD instructions, hardware accelerators for cryptography and neural-network inference, and highly efficient sleep modes. Typical active-mode power consumption can be in the range of a few milliwatts, while standby currents are often in the microampere domain, allowing days or even weeks of operation on a coin-cell battery or on energy harvested from ambient light, vibration, or body heat [10]. The ability to perform feature extraction, filtering, and even preliminary classification locally greatly reduces the need to stream raw high-rate signals to external devices, thereby lowering communication energy, latency, and bandwidth requirements.

This edge computing capability has given rise to what is often described as the *edge-to-cloud continuum* in e-health [11]. In this model, time-critical and privacy-sensitive tasks—such as raw signal denoising, arrhythmia detection, or quality-of-signal assessment—are carried out close to the sensor [12]. In addition, the integration of on-chip hardware security modules enables encrypted processing at the edge, which is increasingly critical to preserve patient privacy and comply with modern healthcare regulations such as HIPAA and GDPR.

In contrast, long-term data storage, large-scale analytics, and population-level modeling can be offloaded to cloud servers or hospital information systems, where virtually unlimited computing and memory resources are available. Such a distributed architecture not only makes wearable devices more energy-efficient but also reduces latency for life-critical alerts, allows continuous monitoring without overwhelming wireless networks, and enables personalized feedback to patients while maintaining a secure data pathway to healthcare providers. Emerging paradigms such as federated learning further extend this model by allowing collaborative training of machine-learning models across distributed devices without sharing raw patient data, thus enhancing both scalability and privacy [13].

The consequences of these technological advances are visible in the expansion of continuous vital-sign monitoring well beyond the walls of hospitals. Early devices focused mainly on relatively “simple” physiological parameters such as heart rate, peripheral oxygen saturation (SpO₂), respiration rate, and body temperature, which can be sensed with minimal power consumption and moderate signal-processing complexity. As sensor technologies have matured, more demanding physiological markers—such as arterial blood pressure, arterial stiffness, cardiac output, respiratory effort, and even surrogate biomarkers like photoplethysmographic variability or capnography—are being targeted for cuffless and continuous assessment in daily life [14]. At the same time, compact point-of-care imaging devices, notably portable ultrasound systems, have become more accessible due to the availability of high-voltage ultrasound drivers, programmable beamformers, and miniature probe arrays that can be connected directly to smartphones or tablets. Furthermore, integration with Internet of Medical Things (IoMT) ecosystems enables seamless synchronization of these diverse data streams, supporting multi-parameter health analytics and decision-support systems [15].

The combination of ubiquitous sensing and local embedded intelligence is enabling a shift from episodic, hospital-based care to continuous, proactive, and personalized healthcare [13]. This paradigm is particularly compelling for the management of chronic cardiovascular and respiratory diseases, which represent the leading global causes of morbidity and mortality and account for a substantial fraction of healthcare expenditure. Timely detection of deviations in arterial blood pressure can help prevent life-threatening complications such as stroke, heart failure, or kidney damage, while bedside lung assessment through lung ultrasound can provide rapid insight into fluid overload, pneumonia, or interstitial syndromes, guiding clinical decisions in both acute and chronic care. In addition, the proliferation of home-based monitoring solutions offers the potential to reduce hospital admissions, lower healthcare costs, and improve patients’ quality of life.

Despite these successes, the acquisition of reliable physiological information outside controlled clinical environments remains challenging. Signals of interest—be they

optical waveforms from a finger PPG sensor or speckled ultrasound images of the pleura—are often low-amplitude, prone to motion artifacts, and affected by environmental conditions such as temperature, illumination, or user interaction. Processing them in real time on resource-constrained embedded platforms requires careful attention to sensor design, analog front-end quality, noise suppression, power-management strategies, and algorithmic efficiency. Moreover, ensuring consistent calibration and interoperability across different devices and vendors remains an open challenge, particularly as these systems move from research prototypes to large-scale clinical deployment. Furthermore, transmitting, storing, and analyzing these data at scale raises additional concerns about data privacy, cybersecurity, and interoperability with electronic health-record systems.

Within this landscape, two application domains have recently drawn intense research and industrial attention:

1. non-invasive cardiovascular monitoring, especially the estimation of arterial blood pressure from optical photoplethysmographic signals, and
2. bed-side lung assessment, particularly the automatic detection of vertical B-line artifacts in lung ultrasound images that serve as a proxy for pulmonary fluid overload.

Although these domains involve different sensing modalities—optical and acoustic—they share a common challenge: the need to extract clinically meaningful information from noisy biological signals using low-power embedded hardware. They therefore provide complementary test cases for exploring how hardware-aware signal-processing approaches can bridge the gap between state-of-the-art algorithms and practical, deployable medical devices. Ultimately, advances in these areas are expected to accelerate the realization of truly connected, intelligent healthcare ecosystems capable of delivering personalized and equitable care worldwide.

1.1.2 Cardiovascular Health and Non-invasive BP Monitoring

Cardiovascular diseases (CVDs) remain the leading cause of death globally, responsible for an estimated 17.9 million deaths each year (World Health Organization, 2021). Among the principal risk factors, arterial hypertension—persistently elevated blood pressure (BP)—plays a central role. Hypertension is strongly associated with an increased incidence of stroke, coronary heart disease, heart failure, renal insufficiency, and premature mortality, and its prevalence is projected to continue rising with aging populations and changing lifestyles. Early detection and effective long-term management of high BP are therefore recognized as critical public-health priorities, both to reduce preventable morbidity and mortality and to mitigate the economic burden on healthcare systems [16].

Conventional BP monitoring in clinical practice typically relies on cuff-based sphygmomanometers, which measure either the oscillometric waveform or the Korotkoff sounds during controlled cuff inflation and deflation. Although such devices are clinically validated and widely available, they have important limitations for continuous and ambulatory monitoring. They provide discrete, episodic readings, often only a few times per day, and repeated cuff inflation can be uncomfortable, interfere with sleep, and alter the measured hemodynamic state itself. In addition, the episodic nature of the

measurement can fail to capture short-term variability or nocturnal hypertension, which are known to be clinically relevant predictors of cardiovascular risk [17].

To overcome these limitations, the past two decades have seen an intense effort to develop cuffless, non-invasive techniques for continuous BP monitoring. These approaches seek to exploit the relationship between BP and various physiological waveforms that can be measured unobtrusively [18]. Examples include electrocardiography (ECG), ballistocardiography (BCG), seismocardiography (SCG), tonometry, and particularly photoplethysmography (PPG). PPG, which detects pulsatile changes in the optical absorption or reflection of vascularized tissue, is attractive because the sensing hardware is simple, inexpensive, and easily miniaturized, and it is already deployed at scale in consumer-grade wearables such as smart watches and fitness bands [19].

Early approaches to cuffless BP estimation often relied on pulse-transit time (PTT) or pulse-arrival time (PAT), which are derived from the time delay between ECG R-waves and characteristic PPG fiducial points. These timing-based indices correlate inversely with BP through arterial compliance. However, they typically require dual-site measurement (ECG plus PPG) and are sensitive to calibration drift due to changes in vascular properties, limb position, and temperature [14]. More recent methods therefore focus on extracting morphological and spectral features directly from single-site PPG waveforms, such as the relative amplitude of systolic and diastolic peaks, the augmentation index, pulse width, or the ratios of first and second derivatives. These features can capture subtle variations in arterial stiffness and peripheral vascular resistance [20].

The advent of machine-learning (ML) and deep-learning (DL) techniques has substantially advanced this field. Algorithms based on support-vector regression (SVR), random forests, and gradient-boosted trees have been applied to handcrafted PPG features, while convolutional neural networks (CNNs), recurrent neural networks (RNNs), and transformer-based models have been trained end-to-end to map raw or minimally processed PPG signals to BP estimates. Such data-driven methods can, in principle, capture non-linear and time-varying relationships between hemodynamics and peripheral waveforms that are difficult to model analytically.

Nevertheless, several challenges hinder widespread adoption of these techniques in real-world clinical or consumer settings:

- **Inter-individual variability:** arterial stiffness, vascular tone, and skin optical properties differ across age groups, ethnicities, and health conditions, often leading to model bias and reduced generalization.
- **Motion artifacts and ambient-light interference:** especially pronounced in daily-life wearable use, these degrade signal quality and compromise the reliability of extracted features.
- **Calibration requirements:** many proposed models still rely on an initial or periodic calibration with cuff-based measurements to achieve acceptable accuracy.
- **Hardware and energy constraints:** sophisticated ML models can be computationally demanding, making them challenging to deploy on low-power embedded platforms such as microcontroller-based wearables.

- Data scarcity and standardization: there is a lack of large, diverse, and openly available datasets of synchronized PPG signals and reference BP readings, which hampers model benchmarking and reproducibility.

Recent research has therefore emphasized the co-design of sensing hardware, signal-processing pipelines, and ML algorithms to address these challenges holistically [20]. For example, multi-wavelength PPG acquisition—typically combining green, red, and near-infrared channels—can help mitigate the influence of skin pigmentation and improve robustness to motion, at the cost of increased sensor complexity and energy consumption. Signal-quality indices and adaptive filtering are often implemented to discard or correct corrupted waveforms before feeding them to ML models.

From a system-design perspective, successful cuffless BP monitoring solutions must balance sensor performance, power consumption, algorithmic complexity, user comfort, and long-term reliability. They must also integrate with secure data-management frameworks that protect privacy while allowing remote consultation and long-term health analytics. Addressing these aspects is crucial not only for technical feasibility but also for regulatory approval and clinical adoption, as international standards such as ISO 81060-2 and the AAMI/ESH/ISO protocols set stringent requirements for accuracy and reproducibility.

In this context, the first research line of this thesis focuses on advancing cuffless BP estimation through two complementary contributions:

1. the employment of an evaluation board-based PPG acquisition platform to collect multi-wavelength signals under controlled laboratory conditions. The board integrates optical sensors, analog front-ends, and a digital interface capable of synchronized, high-fidelity signal capture at multiple wavelengths (green, blue, red, and near-infrared). The work concentrated on optimizing acquisition parameters, calibration routines, and data synchronization protocols to ensure high signal quality and reproducibility across measurement sessions;
2. the creation of a hybrid CNN–SVR model, trained and validated on a dedicated dataset of synchronized PPG and reference BP measurements, designed to improve prediction robustness across subjects and conditions.

These contributions aim to bring state-of-the-art data-driven BP estimation closer to practical deployment in wearable devices for long-term cardiovascular monitoring.

1.1.3 Pulmonary Health and Point-of-Care Ultrasound

Respiratory diseases remain a leading global cause of morbidity and mortality, contributing to a substantial share of hospital admissions and healthcare expenditures. Conditions such as pneumonia, pulmonary edema, acute respiratory distress syndrome (ARDS), interstitial lung disease, and chronic obstructive pulmonary disease (COPD) collectively affect hundreds of millions of people worldwide [21]. The COVID-19 pandemic further underscored the need for reliable and widely deployable tools for assessing lung involvement, especially in resource-limited settings. Early recognition and ongoing monitoring of pulmonary conditions are therefore of paramount importance to guide timely interventions, improve patient outcomes, and optimize the use of healthcare resources [22].

Traditional imaging techniques for lung assessment include chest radiography (CXR) and computed tomography (CT). While CT scans offer high-resolution anatomical visualization, they are costly, expose patients to ionizing radiation, and often require patient transportation to specialized imaging facilities—factors that limit their use for frequent or bedside monitoring [23]. CXR is less resource-intensive but has lower sensitivity for subtle or early interstitial changes and still involves radiation exposure. These limitations have stimulated increasing interest in point-of-care lung ultrasound (LUS) as a radiation-free, portable, and cost-effective imaging modality [24].

In LUS, clinicians rely not on conventional anatomical imaging of lung parenchyma—since aerated lung scatters and reflects ultrasound—but on the interpretation of specific acoustic artifacts that arise at the pleural surface and sub-pleural interstitium [25]. One of the most clinically significant of these artifacts is the B-line: a vertical, laser-like hyperechoic streak extending from the pleural line to the bottom of the screen without fading, moving synchronously with lung sliding [26]. The presence, number, and spatial distribution of B-lines correlate with increased extravascular lung water and are therefore used as surrogate indicators of interstitial syndrome. B-line quantification assists in the diagnosis and monitoring of conditions such as pulmonary edema, pneumonia, pulmonary fibrosis, and ARDS, and has been adopted as part of protocols for critical-care triage and fluid-management strategies.

Owing to advances in probe miniaturization and battery-powered portable ultrasound units, LUS has become an attractive tool not only in intensive-care and emergency-department settings but also for bedside assessments in hospital wards, ambulances, primary-care clinics, and even home-care scenarios [27]. Handheld devices interfacing with smartphones or tablets now provide clinicians with real-time imaging at the point of care, reducing time to diagnosis and improving workflow efficiency. These developments are particularly relevant for resource-constrained environments and for patients with limited mobility.

Despite these advantages, reliable interpretation of LUS remains challenging. The assessment of B-lines is operator-dependent, requiring both technical skill in image acquisition and subjective judgment in counting or scoring the artifacts. Inter-observer variability can be substantial, especially among less-experienced practitioners. Manual frame-by-frame review of ultrasound clips is time-consuming and impractical for continuous or high-throughput monitoring, motivating the search for automated analysis methods that can provide consistent, objective, and rapid interpretation [28].

Early attempts at automation focused on rule-based image-processing techniques, such as edge detection, thresholding, morphological filtering, and texture-based measures. While conceptually simple, these approaches often struggled in the presence of speckle noise, variations in probe angle and gain settings, patient motion, or overlapping artifacts, leading to limited robustness in clinical practice. More recently, machine-learning and deep-learning models—including convolutional neural networks (CNNs), U-Net-based segmentation architectures, and object-detection frameworks such as YOLO—have demonstrated improved accuracy for detecting and quantifying B-lines. However, these data-driven approaches generally require large annotated datasets for training and entail substantial computational and memory demands, which may exceed the capabilities of low-power portable devices. Furthermore, the lack of interpretability of deep neural networks can limit clinical trust and hinder regulatory ap-

proval.

An appealing alternative is offered by signal-decomposition techniques that exploit the distinct spectro-temporal characteristics of vertical reverberation artifacts. Among these, Variational Mode Decomposition (VMD) has emerged as a promising method to separate ultrasound images into components with well-defined spectral bands [29]. By isolating modes corresponding to high-frequency, vertically oriented structures, VMD can enhance the visibility of B-lines while suppressing background clutter and speckle, achieving a favorable balance between computational efficiency, robustness, and interpretability. Such characteristics make VMD-based methods particularly suited for deployment in resource-constrained embedded platforms integrated into portable ultrasound devices.

From a system perspective, the integration of VMD-based B-line detection algorithms with real-time image acquisition and display hardware opens the way to bedside decision-support tools that can assist clinicians with standardized scoring and longitudinal monitoring of pulmonary conditions. Combined with secure wireless connectivity and cloud-based data storage, these tools could further enable remote consultation and tele-ultrasound applications, expanding access to high-quality pulmonary care in underserved regions.

In this context, the second research line of this thesis is devoted to the design and validation of a VMD-based approach for automatic B-line detection in LUS images, with particular emphasis on hardware-aware algorithmic optimization to ensure low-latency, energy-efficient processing suitable for integration into next-generation portable ultrasound platforms. This work aims to bridge the gap between state-of-the-art signal-processing research and the practical requirements of point-of-care pulmonary assessment, ultimately contributing to more accessible, reliable, and data-driven respiratory healthcare.

1.1.4 Toward Hardware-Constrained Biomedical Signal Processing

The last decade has witnessed an extraordinary growth in the adoption of wearable and portable medical devices for continuous health monitoring. However, unlike data analytics conducted on general-purpose computers or cloud servers, the signal processing carried out on these devices is subject to stringent hardware constraints [30]. The need to operate for long periods on small batteries, often in uncontrolled environments, and to deliver reliable, real-time inferences close to the patient's body places unique demands on both the sensing front-end and the embedded computing platform.

A typical biomedical monitoring system comprises several key components:

1. **Sensing interface**, often including optical, electrical, acoustic, or inertial sensors, together with analog front-end circuits for amplification, filtering, and analog-to-digital conversion.
2. **Embedded processor**, which may range from an ultra-low-power microcontroller to a heterogeneous system-on-chip with dedicated accelerators for signal processing and machine learning.
3. **Wireless communication module**, typically based on Bluetooth Low Energy (BLE), Wi-Fi, or cellular protocols, for data transfer and device synchronization.

4. **Energy-management subsystem**, including power regulators, batteries, and sometimes energy-harvesting circuits that scavenge power from ambient light, motion, or thermal gradients.
5. **Software stack**, encompassing signal preprocessing, feature extraction, classification or regression algorithms, and the local implementation of safety and security protocols.

Each of these components imposes practical limitations on the design and execution of signal-processing algorithms. The front-end circuitry determines the available dynamic range, sampling frequency, and signal-to-noise ratio (SNR), constraining the accuracy of downstream analysis. The processor architecture dictates the achievable computational throughput, memory footprint, and latency, while the energy budget often limits the duty cycle and complexity of the operations that can be sustained in real time. Moreover, the need for secure and privacy-preserving computation adds further overhead, particularly when patient data are processed on shared or connected devices [31].

These limitations have given rise to the research field of hardware-aware or hardware-constrained biomedical signal processing. Unlike traditional approaches that design algorithms in isolation and later attempt to port them to embedded hardware, hardware-aware methods adopt a co-design philosophy, tailoring signal-processing and machine-learning pipelines to the specific capabilities and bottlenecks of the target platform. This co-design often requires trade-offs among accuracy, latency, energy consumption, and memory usage.

A prominent example is the deployment of machine-learning models on wearable devices. While deep neural networks can achieve state-of-the-art performance in many biomedical tasks—such as blood-pressure estimation from PPG waveforms or B-line detection in lung ultrasound—their computational and memory requirements often exceed the resources of microcontroller-based platforms. Strategies such as model pruning, quantization, knowledge distillation, and the use of lightweight architectures (e.g., MobileNet, TinyML frameworks) have been explored to bridge this gap, sometimes at the cost of reduced accuracy [32]. Similarly, signal-processing pipelines must be optimized to reduce the use of floating-point operations, exploit fixed-point arithmetic or hardware accelerators, and minimize data movement between memory and processor cores—an increasingly important factor for energy efficiency.

Another crucial aspect is real-time performance. Applications such as arrhythmia detection, blood-pressure tracking during physical activity, or respiratory-distress assessment in critical care require low-latency processing to ensure that alarms and feedback are delivered promptly. Meeting these timing constraints on low-power platforms often demands pipeline parallelization, adaptive sampling, and efficient buffering schemes to balance responsiveness with power consumption. Emerging technologies such as event-driven signal acquisition and neuromorphic co-processors further illustrate the trend toward optimizing the entire sensing-to-decision chain [33].

Beyond the device level, the edge-to-cloud continuum introduces an additional layer of design considerations. When and how to offload data or intermediate features to the cloud depends on the communication energy cost, network latency, data-privacy regulations, and the criticality of the clinical task. Hybrid architectures—where lightweight models run on the device for immediate decision-making while more computationally

intensive analyses are performed in the cloud—are gaining traction in next-generation e-health solutions [34].

The pursuit of hardware-constrained biomedical signal processing is therefore not merely a matter of technical optimization but also of clinical and regulatory significance. Accurate yet resource-efficient algorithms can facilitate continuous, unobtrusive, and affordable monitoring, which is particularly valuable for chronic-disease management and in resource-limited settings. Moreover, devices that meet stringent energy and reliability requirements are more likely to gain regulatory approval and user acceptance, paving the way for large-scale deployment in real-world healthcare systems.

The present thesis builds on this perspective by addressing two complementary case studies—cuffless blood-pressure estimation from multi-wavelength PPG and automatic B-line detection in lung ultrasound—both of which illustrate the importance of designing signal-processing solutions that respect the constraints of embedded hardware. By demonstrating that clinically meaningful results can be achieved even under limited computational and energy budgets, this work contributes to the broader goal of making advanced biomedical analytics accessible at the point of care.

1.2 Research Objectives

The overarching aim of this doctoral research is to explore and advance the development of hardware-constrained biomedical signal-processing systems capable of delivering accurate, interpretable, and real-time physiological monitoring in practical e-health scenarios. The work is motivated by the observation that state-of-the-art algorithms for biomedical analytics—often designed and validated in laboratory or cloud-computing settings—cannot be deployed as-is on wearable or portable devices due to their energy, latency, and memory constraints. By investigating solutions that co-optimize sensing hardware, signal processing, and machine-learning models, the thesis seeks to narrow the gap between academic research prototypes and clinically viable, embedded e-health technologies.

More concretely, the research addresses two complementary application domains — cardiovascular and pulmonary health — each exemplifying a class of biomedical sensing problems that are relevant for continuous or point-of-care monitoring:

1. Non-invasive Cardiovascular Monitoring via Photoplethysmography (PPG)

- To design, configure, and experimentally validate a low-power embedded sensing platform for the acquisition of multi-wavelength PPG signals, optimized for continuous and cuffless estimation of arterial blood pressure (BP) in real-life conditions.
- To investigate and compare conventional feature-based regressors and data-driven methods for BP estimation, and to propose a hybrid machine-learning framework combining convolutional-neural-network (CNN) feature extraction with support-vector-regression (SVR) for improved robustness across heterogeneous subjects and measurement conditions.
- To optimize the overall signal-processing pipeline for deployment on resource-constrained embedded platforms, including preprocessing, quality assessment,

and inference steps, while preserving compliance with relevant accuracy standards (e.g., AAMI/ESH/ISO).

2. Automated Pulmonary Assessment via Lung Ultrasound (LUS)

- To investigate the use of Variational Mode Decomposition (VMD) as a signal-processing tool for enhancing and isolating the vertical B-line artifacts that serve as surrogate markers of interstitial lung involvement.
- To develop an efficient and interpretable algorithmic pipeline for automatic detection and quantification of B-lines in ultrasound images, suitable for execution on the limited computing resources typically available in portable point-of-care ultrasound devices.
- To evaluate the algorithm on annotated LUS datasets and compare its performance against both traditional rule-based methods and recent deep-learning approaches, highlighting the trade-offs in accuracy, computational complexity, and interpretability.
- To assess the potential of the proposed framework to support real-time bedside decision-support tools, facilitating standardized scoring and longitudinal monitoring of pulmonary conditions.

In addition to these domain-specific aims, the thesis also pursues several cross-cutting research objectives that reflect the broader challenges of embedded e-health technologies:

- **Hardware–algorithm co-design:** to demonstrate how knowledge of sensor front-end characteristics, noise sources, and embedded-platform limitations can inform the design of more robust and efficient signal-processing pipelines.
- **Dataset creation and validation:** to emphasize the importance of high-quality, synchronized, and diverse datasets—including meta-data on acquisition conditions—for the development and fair benchmarking of machine-learning models in biomedical applications.
- **Model efficiency and edge deployment:** to explore strategies such as algorithmic simplification, hybrid modeling, and hardware-specific optimization for enabling real-time inference under strict energy and latency constraints.
- **Clinical relevance and interpretability:** to align algorithmic outputs with clinically meaningful parameters and to favor methods that offer insight into the signal characteristics driving predictions, thereby improving trust, usability, and potential regulatory acceptance.

By articulating and addressing these objectives, the thesis seeks to advance the state of the art in hardware-aware biomedical signal processing and to provide proof-of-concept demonstrations that continuous cardiovascular monitoring and point-of-care pulmonary assessment can be realized with resource-efficient embedded platforms, thereby contributing to the long-term vision of accessible, data-driven, and patient-centered health-care.

1.3 Contribution of the Thesis

The work presented in this dissertation makes several contributions to the field of biomedical signal processing for embedded and resource-constrained systems, addressing the need for accurate and interpretable physiological monitoring that can operate in real time and beyond traditional laboratory or hospital environments. At its core, the thesis provides a hardware-aware perspective on the design of biomedical monitoring solutions, demonstrating that clinically meaningful signal-processing pipelines can be implemented on energy- and resource-limited embedded platforms without sacrificing essential performance. Furthermore, this work highlights the importance of dataset quality, reproducibility, and algorithm–hardware co-design as enabling factors for moving e-health technologies from laboratory prototypes toward deployable medical devices. It delivers two proof-of-concept systems—one for continuous cardiovascular monitoring and one for point-of-care pulmonary assessment—that exemplify how different sensing modalities (optical vs. acoustic) can benefit from a common design philosophy focused on embedded implementation and clinical usability.

From a scientific perspective, the thesis contributes new knowledge through the development of two complementary systems that exemplify this philosophy: one dedicated to non-invasive cardiovascular monitoring and the other to pulmonary assessment through lung ultrasound imaging.

1.3.1 Specific Contributions to Non-invasive Cardiovascular Monitoring

1. Embedded Multi-Wavelength PPG Sensing Platform

- Design and implementation of a low-power optical acquisition system integrating red, green, blue and near-infrared channels with analog front-end amplification and on-board signal preprocessing.
- Optimization of sampling strategy, noise mitigation, and power management for continuous, wearable operation.

2. Dedicated Dataset for Cuffless BP Estimation

- Establishment of a synchronized dataset comprising multi-wavelength PPG recordings and reference oscillometric BP measurements, collected under a controlled acquisition protocol and covering a diverse subject population.
- Inclusion of meta-data on demographic and physiological factors, enabling future studies on model generalization.

3. Hybrid CNN–SVR Modeling Framework

- Development of a hybrid approach combining convolutional-neural-network (CNN)-based feature extraction with support-vector-regression (SVR) to exploit both the representation-learning power of deep models and the robustness of kernel-based regression.
- Demonstration of improved prediction accuracy and stability across subjects and conditions compared with baseline single-model or traditional feature-based methods.

4. Embedded Deployment and Real-Time Evaluation

- Implementation of a hardware-aware signal-processing pipeline—including preprocessing, quality assessment, and inference—tailored to the constraints of a microcontroller-based platform.
- Validation of the pipeline’s feasibility for real-time continuous monitoring and discussion of compliance with international accuracy standards (e.g., AAMI/ESH/ISO).

1.3.2 Specific Contributions to Pulmonary Assessment with Lung Ultrasound**• VMD-Based B-Line Enhancement and Detection**

- Introduction of a Variational Mode Decomposition (VMD)-driven algorithmic pipeline to enhance and isolate the spectral–spatial signature of vertical B-line artifacts in LUS images.
- Demonstration of the approach’s capability to improve the contrast and detectability of B-lines relative to conventional filtering or purely morphology-based methods.

• Efficient and Interpretable Algorithm for Embedded Ultrasound Devices

- Design of a block-wise processing strategy that reduces memory footprint and computation load, enabling low-latency execution on portable point-of-care ultrasound hardware.
- Maintenance of a transparent, explainable detection logic, favoring clinical trust and potential regulatory acceptance over opaque black-box models.

• Experimental Validation and Benchmarking

- Quantitative evaluation of the proposed method on annotated LUS datasets, assessing sensitivity, specificity, and processing time, and comparing the results with both conventional rule-based and recent deep-learning baselines.
- Discussion of the accuracy–complexity trade-offs relevant for resource-constrained deployment.

1.3.3 Cross-Cutting and Methodological Contributions

- **Hardware–Algorithm Co-Design:** the thesis illustrates how an explicit consideration of sensor characteristics, signal quality, and embedded-platform limitations can inform algorithm selection and optimization, achieving an effective balance between performance and feasibility.
- **Bridging Laboratory Research and Practical Devices:** by combining dedicated datasets, optimized processing pipelines, and embedded implementation, the work contributes to reducing the gap between academic prototypes and clinically usable solutions.

- **Promotion of Explainable Biomedical AI:** both research lines adopt approaches that preserve interpretable links between input signal properties and predicted clinical indices, supporting trustworthiness, transparency, and potential regulatory compliance.
- **Enabling Continuous and Point-of-Care Monitoring:** the two demonstrator systems jointly showcase how advanced biomedical analytics can be brought closer to the patient, supporting the ongoing transition toward data-driven, proactive, and decentralized healthcare.

Through these contributions, the thesis advances the state of the art in embedded biomedical signal processing and demonstrates the feasibility of deploying optimized, clinically relevant algorithms on resource-constrained platforms. The findings are intended to support future efforts to integrate such solutions into scalable e-health infrastructures, ultimately contributing to more accessible, reliable, and patient-centered healthcare delivery.

1.4 Structure of the Thesis

The thesis is organized into five chapters, each addressing a specific aspect of the research and following a logical progression from theoretical background to experimental validation and discussion. This structure has been designed to guide the reader from the general context of biomedical signal processing toward the specific methodologies and findings that define the two main research lines: cuffless blood-pressure prediction using multi-wavelength photoplethysmography (PPG) and automatic B-line detection in lung ultrasound (LUS) imaging. Although these two applications focus on distinct physiological systems and sensing modalities, they are united by a common objective: the design and optimization of embedded electronics for real-time biomedical signal acquisition and processing under hardware and energy constraints.

Chapter 1, the current introduction, has presented the overall motivation and context of the research, outlining the societal need for continuous, accessible, and data-driven healthcare. It has introduced the scientific and technological challenges associated with developing embedded biomedical systems and has articulated the main research objectives and contributions of the thesis. This chapter also defines how the following sections are structured, serving as a conceptual roadmap for the work.

Chapter 2 provides the background on embedded hardware for physiological signal acquisition, setting the foundation for the subsequent research developments. It begins with an overview of biopotential and biomarker monitoring, describing the physiological origins of relevant signals such as cardiac and vascular activity, as well as biomarkers that can be sensed optically or electrically. It then discusses the design of embedded hardware platforms, emphasizing constraints related to energy efficiency, sampling accuracy, and system integration. The chapter concludes with an analysis of signal-processing techniques for wearable and resource-limited devices, including filtering, feature extraction, and algorithmic simplification strategies that enable real-time operation within the limitations of embedded platforms.

Chapter 3 is devoted to the first research line: blood-pressure prediction using multi-wavelength PPG signals. It opens with a discussion of the principles of non-invasive blood-pressure monitoring and the rationale for using optical sensing to estimate hemodynamic parameters. The chapter then describes in detail the hardware architecture and wearable implementation of the developed optical acquisition system, followed by the dataset creation and experimental protocol used for model training and validation. The dataset includes synchronized PPG signals and reference blood-pressure measurements, along with detailed acquisition parameters, synchronization methods, and demographic information. Subsequent sections describe the signal pre-processing and feature extraction pipeline, the machine-learning models explored in this study—including multilayer perceptrons (MLP), principal component analysis (PCA), and convolutional neural networks (CNNs)—and the proposed hybrid CNN–SVR model that combines deep and kernel-based learning for optimized prediction. The chapter concludes with an extensive experimental analysis and discussion, highlighting performance metrics, model generalization, and considerations for real-time deployment on embedded hardware.

Chapter 4 presents the second research line, focused on automatic B-line detection in lung ultrasound images. The chapter starts with the clinical background and relevance of B-lines in pulmonary assessment, outlining their diagnostic significance in conditions such as pulmonary edema, interstitial lung disease, and acute respiratory distress syndrome (ARDS). It then details the lung ultrasound acquisition protocol and the front-end ultrasound architecture used in this study. A central part of the chapter is dedicated to the introduction of Variational Mode Decomposition (VMD) as a signal-processing method for ultrasound image analysis, followed by the description of the block-wise VMD algorithm developed for efficient detection of vertical artifacts. Finally, the chapter reports the results of experimental validation and benchmarking, discussing the method’s robustness, computational efficiency, and potential for real-time integration into portable point-of-care ultrasound systems.

Chapter 5 concludes the thesis by summarizing the main findings and contributions of the work as a whole. It discusses how the results obtained in both research lines support the central hypothesis that advanced biomedical analytics can be achieved within the constraints of embedded hardware through careful co-design of algorithms and sensing platforms. The chapter also identifies limitations of the current work—including dataset size, model generalization, and clinical validation—and outlines several directions for future research. These include further optimization for embedded deployment, integration with deep-learning frameworks and hybrid edge–cloud architectures, and extended validation in real clinical environments to ensure scalability and translational impact.

Overall, the structure of the thesis reflects a coherent trajectory from fundamental concepts to practical implementation, demonstrating how embedded electronics and signal-processing algorithms can converge to enable intelligent biomedical monitoring systems. By combining theoretical rigor with experimental validation, and by addressing both cardiovascular and pulmonary applications, the document aims to provide a

Chapter 1. Introduction

comprehensive contribution to the field of multimodal biomedical signal acquisition and embedded processing, paving the way toward more reliable, interpretable, and accessible e-health technologies.

CHAPTER 2

Embedded Hardware for Physiological Signal Acquisition

2.1 Background on Biopotential and Biomarker Monitoring

The measurement and interpretation of biopotentials and biomarkers are at the core of modern biomedical instrumentation. These physiological signals and indicators serve as windows into the body's functional state, providing quantitative insights into health and disease progression. Over the past two decades, the growing demand for continuous and unobtrusive monitoring—driven by aging populations, chronic disease prevalence, and the digitalization of healthcare—has spurred a shift from traditional, clinic-based measurements to wearable and embedded sensing systems capable of real-time operation [35, 36].

In this context, biopotentials refer to electrical signals generated by living tissues due to ionic movement across cell membranes. These signals originate from the bioelectrical activity of excitable cells—neurons, myocytes, or cardiomyocytes—and can be measured at the body surface or invasively using electrodes [37, 38]. Biomarkers, on the other hand, encompass a broader class of measurable variables, including optical, biochemical, and mechanical indicators that reflect underlying physiological or pathological processes. The simultaneous acquisition of multiple biomarkers and biopotentials can provide a multimodal perspective on human physiology, improving diagnostic accuracy and enabling predictive modeling [39, 40].

2.1.1 Electrophysiological Signals and Biopotentials

Among the various categories of biopotentials, the electrocardiogram (ECG), electromyogram (EMG), and electroencephalogram (EEG) are the most extensively used for clinical and research purposes (Fig. 2.1).

The ECG represents the electrical manifestation of the heart's depolarization and repolarization cycles. Each cardiac event generates characteristic waveforms (P, QRS, T) that carry information about atrial and ventricular activity, conduction pathways, and overall cardiac rhythm [41]. ECG monitoring enables the detection of arrhythmias, ischemia, and autonomic dysfunction, and its long-term analysis supports the study of heart-rate variability (HRV), an established marker of cardiovascular and stress-related conditions [42]. In embedded contexts, ECG monitoring poses challenges related to motion artifacts, baseline drift, and electrode–skin impedance variability, requiring the use of high-input-impedance, low-noise instrumentation amplifiers and adaptive digital filtering [43, 44].

The EMG, which captures the electrical activity of skeletal muscles, is crucial for diagnosing neuromuscular disorders and for applications in rehabilitation, prosthetics, and human–machine interfaces. The amplitude and frequency content of EMG signals vary with muscle contraction level and fatigue, providing rich information on motor control [45]. In wearable systems, dry or textile electrodes have replaced traditional gel-based sensors to enhance comfort and durability, although they typically introduce higher impedance and noise, necessitating careful hardware design and impedance-compensation strategies [46, 47].

The EEG records cortical potentials with amplitudes in the range of 10–100 μV , typically requiring high-gain, low-noise amplification and robust shielding against electromagnetic interference. EEG-based systems have traditionally been confined to controlled laboratory settings, but recent progress in miniaturized front-ends, wireless data transmission, and flexible electrodes has enabled wearable neurotechnologies for applications such as sleep monitoring, brain–computer interfaces (BCIs), and neurofeedback [48, 49].




Figure 2.1: ECG, EMG and EEG biopotentials formation and visualization.

The acquisition of these signals in embedded or wearable formats demands not only sensitive electronics but also mechanical and ergonomic optimization to ensure stable electrode contact, low motion artifacts, and user comfort during prolonged use. Advances in soft bioelectronics—featuring stretchable circuits and biocompatible polymers—are paving the way toward “electronic skin” interfaces, capable of continuous physiological monitoring without compromising mobility or comfort [50, 51].

2.1.2 Optical Biomarkers and Photoplethysmography

While electrophysiological signals provide direct measures of bioelectrical activity, optical biomarkers offer valuable insight into hemodynamic, metabolic, and structural properties of tissues. The most prevalent optical technique for physiological monitoring is photoplethysmography (PPG), which detects variations in light intensity transmitted or reflected by tissue as blood volume changes with each cardiac cycle [52].

PPG systems typically employ light-emitting diodes (LEDs) operating at specific wavelengths—commonly green (530 nm), red (660 nm), and infrared (940 nm)—and photodiodes that detect backscattered or transmitted light. The resulting signal contains a pulsatile (AC) component, synchronous with the heartbeat, and a quasi-static (DC) component that depends on tissue structure and perfusion. These components can be analyzed to extract metrics such as heart rate, respiratory rate, oxygen saturation (SpO_2), and blood-pressure surrogates [53, 54].

The multi-wavelength PPG configuration, which is central to this thesis, enables a more comprehensive analysis of vascular dynamics. By comparing optical absorption at different wavelengths—each with specific penetration depths and sensitivities to hemoglobin and oxyhemoglobin—multi-channel PPG allows the estimation of arterial stiffness, vascular compliance, and pulse transit time (PTT), all of which are correlated with arterial blood pressure [55, 56].

Beyond cardiovascular monitoring, optical sensing has been extended to near-infrared spectroscopy (NIRS) and optical coherence tomography (OCT). NIRS assesses local oxygenation and metabolic activity by measuring differential absorption of oxygenated and deoxygenated hemoglobin in the near-infrared range (700–1000 nm), while OCT provides high-resolution cross-sectional imaging based on light scattering. Although traditionally confined to stationary equipment, miniaturized CMOS-compatible light sources, photodiode arrays, and integrated optical filters are enabling portable optical systems for bedside and home use.

2.1.3 Chemical Biomarkers and Non-Invasive Sensing

Recent progress in chemical biosensing has introduced new opportunities for monitoring metabolic and biochemical processes in real time. Electrochemical biosensors exploit enzymatic reactions, redox activity, or ion-selective membranes to quantify analytes such as glucose, lactate, electrolytes, cortisol, and pH in biofluids including sweat, saliva, and interstitial fluid [39].

Unlike traditional invasive sampling, these non-invasive or minimally invasive sensors allow continuous assessment of metabolic markers without blood extraction. Wearable platforms integrating microfluidic channels and stretchable electrodes have demonstrated accurate, stable measurements during exercise or daily activities [57]. Multi-analyte systems combining chemical and physical sensing—such as sweat glucose alongside heart rate or skin temperature—are emerging as powerful tools for context-aware health monitoring [58].

However, challenges remain in ensuring sensor selectivity, long-term stability, and calibration accuracy under dynamic physiological conditions. Furthermore, the integration of chemical sensors with embedded electronics and wireless modules necessitates innovative power management and signal-conditioning solutions to ensure reliability in

mobile environments [59].

2.1.4 Multimodal and Embedded Physiological Monitoring

A defining trend in current biomedical engineering is the convergence of multiple sensing modalities within a single embedded device. Combining biopotentials (e.g., ECG), optical biomarkers (e.g., PPG), and biochemical measurements allows a more holistic interpretation of physiological states and disease mechanisms [60]. For instance, ECG–PPG fusion enables the derivation of pulse transit time (PTT), a surrogate for arterial compliance and blood pressure, while PPG–accelerometer fusion improves robustness to motion artifacts in wearable devices [61].

Multimodal systems are further empowered by embedded processors and wireless communication, which enable on-board preprocessing, feature extraction, and data transmission to smartphones or cloud infrastructures [60]. This shift from traditional wired instrumentation to compact, autonomous devices is made possible by advances in semiconductor technology, ultra-low-power microcontrollers, and system-on-chip (SoC) integration.

The migration of physiological monitoring from hospital-grade equipment to embedded and wearable platforms not only enables continuous health tracking but also generates new design challenges. These include low signal-to-noise ratio (SNR) due to miniaturized electrodes and optical paths, motion and environmental interference, and stringent requirements for power consumption, latency, and data security. Addressing these issues requires co-design strategies that jointly optimize sensor front-ends, signal-processing algorithms, and embedded architectures—a central theme of this thesis.

2.1.5 Experimental Implementation of ECG–PPG Fusion for Multiparametric Vital Sign Monitoring

A concrete example of multimodal sensing integration was performed in our preliminary work [62], which explored the combined use of ECG and PPG signals within a compact embedded platform for vital sign monitoring. The system was designed to acquire synchronized electrical and optical biosignals, enabling the simultaneous estimation of multiple cardiovascular parameters—namely heart rate (HR), peripheral oxygen saturation (SpO_2), and systolic blood pressure (SBP). This approach relies on the complementary nature of the two modalities: the ECG provides precise temporal markers of cardiac depolarization, while the PPG captures optical fluctuations associated with blood volume changes in peripheral vessels. Their temporal relationship, quantified through the pulse transit time (PTT), can be empirically linked to arterial compliance and thus to blood pressure variability (Fig. 2.2).

The proposed acquisition platform implemented a dual-channel front-end comprising an analog conditioning stage for ECG and a dual-wavelength PPG module based on red and infrared emission. The system was controlled by a low-power microcontroller that performed synchronous sampling, analog-to-digital conversion, and preliminary digital filtering before wireless data transmission to a host computer for post-processing. The embedded implementation ensured portability, low power consumption, and the capability to perform real-time acquisition without external instrumentation—an essential feature for continuous physiological monitoring.




Figure 2.2: Overlapped ECG and PPG signals, highlighting PTT in the lag between their peaks.

Experimental validation was conducted on a cohort of healthy volunteers under resting conditions. The synchronized ECG and PPG waveforms were recorded simultaneously with reference devices for calibration and accuracy evaluation. The computed HR and SpO_2 parameters exhibited excellent agreement with commercial pulse oximeters, showing deviations below 2 bpm and 1% respectively. More notably, regression models based on the measured PTT and additional morphological PPG features achieved a high correlation ($R > 0.9$) with reference systolic blood pressure values, with an average estimation error below 5 mmHg. These findings confirmed that the temporal synchronization between ECG and PPG signals enables an indirect yet reliable estimation of blood pressure without the need for an inflatable cuff, opening the path to non-invasive and continuous monitoring in daily life contexts.

From a system-level standpoint, the work also emphasized several design considerations that remain critical for future embedded implementations. The synchronization between acquisition channels was found to be essential, as sub-millisecond desynchronization could lead to significant deviations in PTT estimation. Furthermore, the mechanical coupling of optical sensors to the skin and the quality of electrode–skin contact were identified as major determinants of signal quality and stability, especially in dynamic conditions. Despite these practical challenges, the device demonstrated consistent and reproducible measurements, confirming the technical feasibility of multimodal physiological monitoring on a compact, wearable platform.

In conclusion, this study provided one of the earliest demonstrations of an integrated ECG–PPG embedded system capable of estimating multiple vital parameters simultaneously, including HR, SpO_2 , and SBP. The successful validation of the architecture under controlled experimental conditions supports the broader concept of multimodal sensor fusion for biomedical monitoring, where heterogeneous biosignals are processed jointly to extract richer and more reliable physiological information. These results, although preliminary, represent a significant milestone toward the development of fully autonomous, cuffless, and continuous vital-sign monitoring devices, laying the ground-

work for future advancements in personalized and pervasive healthcare technologies.

2.2 Embedded Hardware Platforms and Constraints

The evolution of biomedical signal acquisition from benchtop instrumentation to embedded and wearable platforms has been enabled by remarkable progress in microelectronics, sensing interfaces, and low-power computation. Embedded systems designed for physiological monitoring must perform a wide range of operations—data acquisition, conditioning, digitization, feature extraction, and sometimes classification—under stringent limitations in power consumption, computational resources, memory capacity, and physical form factor. These systems must also guarantee signal fidelity, temporal accuracy, and robustness to environmental disturbances, making the hardware design process a complex balance between performance and efficiency.

In biomedical contexts, embedded hardware typically refers to compact electronic platforms integrating both analog and digital components, optimized for specific signal-processing tasks and often intended for battery-powered or energy-harvesting operation [63]. Such platforms are fundamental for next-generation wearable and implantable sensors, enabling long-term physiological data collection and continuous health tracking outside of clinical settings.

2.2.1 System Architecture of Embedded Biomedical Devices

A typical embedded biomedical monitoring system can be decomposed into several subsystems:

- i. **Sensor and transducer front-end**, responsible for converting physical or chemical phenomena into electrical signals.
- ii. **Analog front-end (AFE)** for amplification, filtering, and analog-to-digital conversion.
- iii. **Processing unit**, usually a microcontroller, microprocessor, or system-on-chip (SoC) that executes signal-processing algorithms.
- iv. **Communication interface**, which handles wireless data transfer and synchronization.
- v. **Power management subsystem**, ensuring energy efficiency and stability of operation [35].

Each of these components introduces unique design constraints that affect the overall system performance. The sensor front-end defines the available signal amplitude and bandwidth, whereas the AFE determines the dynamic range and noise floor. The digital core dictates computational throughput and memory availability, while power management governs autonomy and duty cycle.

Modern biomedical systems increasingly adopt system-on-chip (SoC) or system-in-package (SiP) solutions, which integrate multiple functional blocks—such as amplifiers, analog-to-digital converters (ADCs), microcontrollers, and wireless transceivers—on a single die. Commercial examples include platforms such as the Texas Instruments ADS129x family for biopotential acquisition, or Nordic nRF52 series microcontrollers

that integrate Bluetooth Low Energy (BLE) connectivity with efficient ARM Cortex-M processing cores. The ongoing scaling of CMOS technology has enabled such SoCs to combine high precision and low noise with reduced power dissipation, typically in the range of a few milliwatts, thereby extending their use to battery-operated or energy-harvesting medical devices [64].

2.2.2 Analog Front-End Design Considerations

The analog front-end (AFE) is arguably the most critical part of any biomedical acquisition system, as it directly interfaces with the human body. The signals involved are often weak (microvolt to millivolt range) and prone to interference from external sources such as motion, electromagnetic fields, and electrode contact noise [65]. The AFE must therefore ensure:

- High input impedance to prevent signal attenuation;
- Low input-referred noise to preserve weak physiological signals;
- High common-mode rejection ratio (CMRR) to suppress interference, especially at 50/60 Hz;
- Programmable gain and filtering to adapt to various signal types and environments.

Instrumentation amplifiers (INAs) designed for biomedical applications often employ chopper stabilization or auto-zero techniques to minimize offset and flicker noise. Recent designs have demonstrated noise densities below 50 nV/Hz with power consumptions under 10 μ W per channel, making them suitable for long-term wearable operation [66].

Moreover, power-line interference and motion artifacts remain dominant sources of distortion in ECG, EMG, and EEG acquisition. Adaptive filtering methods implemented either in analog or digital form are used to mitigate such artifacts, but their success largely depends on the AFE's baseline noise and bandwidth characteristics. AFE design is therefore not only a matter of electronics but a key determinant of signal quality and downstream processing accuracy [67].

2.2.3 Power Consumption and Energy Management

Energy efficiency is the cornerstone of embedded biomedical design. Unlike traditional medical instruments connected to the power grid, wearable and portable systems rely on limited energy sources, typically coin-cell batteries, lithium-polymer packs, or energy-harvesting units. The need for continuous operation over days or weeks places stringent requirements on average power consumption, which often must remain below a few milliwatts [68].

Power consumption in embedded systems can be divided into three main categories:

1. **Sensing and analog acquisition**, where power is spent in excitation (for optical or chemical sensors) and analog amplification.
2. **Digital computation**, including filtering, feature extraction, and inference algorithms.

3. **Wireless communication**, which often dominates total energy use due to radio transmission costs.

Techniques for reducing power include dynamic voltage and frequency scaling (DVFS), duty cycling, event-driven sampling, and hardware-level power gating [69]. Some systems exploit wake-on-event architectures, where analog comparators or low-power sensors trigger the activation of higher-power modules only when relevant signals are detected.

In recent years, energy-harvesting technologies have gained momentum as a means to extend device lifetime or eliminate batteries altogether. Photovoltaic, thermoelectric, and piezoelectric harvesters can convert ambient energy into electrical power, albeit typically at microwatt levels, sufficient only for ultra-low-power microcontrollers and sensors [70]. The integration of supercapacitors and high-efficiency DC–DC converters allows intermittent operation with energy buffering, supporting periodic sensing and wireless transmission tasks [71].

2.2.4 Processing Architectures and Computational Constraints

The digital-processing unit in an embedded biomedical system executes the algorithms that transform raw sensor data into useful physiological metrics. While general-purpose microcontrollers remain dominant in commercial devices, more specialized heterogeneous architectures—combining microcontroller cores, digital-signal-processing (DSP) units, and machine-learning accelerators—are increasingly used to meet performance requirements at minimal energy cost [72].

Common processing units include ARM Cortex-M cores for control-oriented tasks and Cortex-M4/M7 for real-time signal processing, featuring single-instruction-multiple-data (SIMD) extensions. In ultra-low-power applications, architectures such as RISC-V and MSP430 are preferred due to their reduced instruction sets and efficient sleep modes. When higher throughput is required—for instance, in multi-channel signal analysis or convolutional neural network (CNN) inference—designers can leverage dedicated DSP cores or hardware accelerators capable of executing multiply–accumulate (MAC) operations with optimized energy efficiency.

However, the implementation of advanced algorithms in such environments is constrained by limited memory and arithmetic precision. Many embedded devices offer only a few tens or hundreds of kilobytes of SRAM, and they lack floating-point units (FPUs). Consequently, fixed-point arithmetic, quantized models, and memory-aware algorithm design become essential to maintain computational feasibility [73]. For example, CNN-based inference on embedded systems often requires model compression techniques such as pruning, weight quantization, and knowledge distillation, which reduce both memory footprint and computation time with minimal loss in accuracy [74, 75].

2.2.5 Communication and Data Transmission Constraints

Wireless communication is another critical bottleneck in embedded biomedical systems. Data transmission consumes significantly more energy than sensing or local processing, particularly in continuous-monitoring applications [76]. Low-power wireless standards such as Bluetooth Low Energy (BLE), Zigbee, and ANT+ are widely used for

2.3. Signal Processing for Wearable and Resource-Constrained Devices

short-range communication with smartphones or base stations, while LoRa and NB-IoT are employed for long-range medical IoT scenarios [77].

Communication protocols must balance throughput, latency, reliability, and energy efficiency. For example, BLE offers high efficiency for small, periodic data packets but becomes less suitable for high-bandwidth signals like raw ECG or ultrasound. Consequently, many modern systems employ edge processing, where signals are preprocessed and compressed locally before transmission, reducing communication load and preserving energy [78].

Security and privacy further constrain communication design. Medical data are subject to regulatory requirements such as HIPAA and GDPR, necessitating encryption, authentication, and sometimes on-device anonymization. These operations introduce computational overhead, which must be optimized to avoid excessive energy consumption.

2.2.6 Trade-Offs and Co-Design Strategies

The design of embedded biomedical systems inherently involves multi-domain trade-offs among accuracy, energy, latency, and cost. Achieving optimal performance requires hardware–software co-design, where algorithmic complexity, sensor configuration, and power architecture are jointly optimized.

For instance, reducing sampling frequency or ADC resolution lowers power consumption but may degrade signal fidelity. Similarly, minimizing computation through model simplification must be weighed against potential losses in diagnostic accuracy. Co-design methodologies integrate hardware profiling, algorithm benchmarking, and power simulation to achieve balanced solutions [79].

A growing body of research focuses on TinyML—machine learning models optimized for microcontrollers—which represents a key enabler for intelligent edge devices in biomedical applications. Frameworks such as TensorFlow Lite for Microcontrollers and CMSIS-NN now allow efficient deployment of neural networks on ARM-based architectures, enabling real-time classification of biomedical signals like ECG or PPG at sub-milliwatt power levels.

Finally, reliability and safety remain essential considerations. Embedded biomedical systems must guarantee robust operation under physiological variability and environmental noise, and in some cases, fail-safe behavior when used in medical devices subject to regulatory certification. These aspects underline the interdisciplinary nature of embedded biomedical design—where advances in circuit design, firmware optimization, and biomedical modeling converge to realize reliable, power-efficient, and clinically relevant monitoring platforms.

2.3 Signal Processing for Wearable and Resource-Constrained Devices

Signal processing plays a pivotal role in transforming raw physiological signals—often noisy, artifact-prone, and non-stationary—into reliable and clinically meaningful information. In wearable and embedded biomedical systems, however, this task is challenged by limited computational resources, strict energy budgets, and the need for real-time operation. Traditional signal-processing pipelines designed for desktop or clinical

instrumentation must therefore be rethought to fit resource-constrained environments, where memory, throughput, and numerical precision are at a premium.

In this context, efficient signal processing involves not only the development of accurate algorithms but also their hardware-aware implementation, ensuring that the chosen techniques respect the platform's timing, energy, and memory constraints. The emerging paradigm of “embedded intelligence” aims to move part of the signal analysis from remote servers to the edge of the network, reducing communication costs, latency, and dependency on cloud connectivity.

2.3.1 Characteristics of Physiological Signals and Sources of Noise

Physiological signals such as ECG, PPG, and ultrasound are inherently low-amplitude and highly susceptible to noise. Typical biopotentials operate in the range of microvolts to millivolts and are affected by motion artifacts, electromagnetic interference, electrode drift, and baseline wander [37]. In wearable systems, additional sources of noise include sensor displacement, optical path changes due to motion (in PPG), and ambient light interference [53].

Signal preprocessing must therefore include artifact detection, denoising, and baseline correction. Techniques such as bandpass filtering, wavelet decomposition, and adaptive noise cancellation are widely used [80]. For example, ECG signals typically require a 0.5–40 Hz passband to suppress baseline wander and high-frequency muscle noise, while PPG preprocessing often employs digital notch filters to remove power-line interference and moving-average filters for baseline restoration [81, 82].

However, many conventional filtering techniques are computationally intensive when implemented on low-power microcontrollers. Consequently, simplified, fixed-point implementations and integer arithmetic are often employed to reduce computational complexity. Moreover, IIR filters are preferred over FIR filters in ultra-low-power systems due to their reduced coefficient count and smaller memory footprint.

2.3.2 Adaptive and Nonlinear Signal Processing

Physiological signals are typically non-stationary, meaning that their statistical properties change over time due to respiration, posture, or stress. This necessitates adaptive filtering techniques, which can update filter parameters in real time to track dynamic changes in the signal environment.

The Least Mean Squares (LMS) and Recursive Least Squares (RLS) algorithms are classical adaptive filters, commonly used for motion artifact reduction in wearable ECG and PPG systems. LMS is computationally simpler, relying mainly on scalar operations, while RLS provides faster convergence but at the cost of higher computational and memory demands—often prohibitive for embedded processors. Lightweight variants such as Normalized LMS (NLMS) and Affine Projection Algorithms (APA) offer improved stability for on-body sensing applications [83].

Beyond linear filtering, nonlinear signal-processing methods have been introduced to capture the complex dynamics of physiological systems. These include Empirical Mode Decomposition (EMD), Variational Mode Decomposition (VMD), and wavelet packet transforms, which decompose signals into adaptive basis functions that better represent their time–frequency structure [84]. For instance, in lung ultrasound

imaging, VMD-based approaches can separate artifact components (e.g., B-lines) from speckle noise, improving detection sensitivity while maintaining computational feasibility through optimized numerical implementations.

When implemented in embedded environments, nonlinear methods require careful algorithmic simplification—such as reducing iteration counts, using precomputed lookup tables, and replacing floating-point operations with fixed-point equivalents—to balance accuracy with energy efficiency.

2.3.3 Feature Extraction and Dimensionality Reduction

Once the signal has been preprocessed, feature extraction is performed to derive quantitative attributes that describe its morphology, frequency content, or statistical properties. In embedded systems, feature extraction must be computationally lightweight, yet informative enough to discriminate physiological states or pathologies.

Typical ECG features include R–R interval, QRS complex width, and heart-rate variability metrics in the time and frequency domains. For PPG, features such as pulse amplitude, systolic upstroke time, dicrotic notch position, and pulse width are widely used [54, 55]. In ultrasound-based analysis, texture descriptors, entropy measures, and energy-based indices can be extracted from decomposed mode components [85].

Dimensionality reduction techniques—such as Principal Component Analysis (PCA) and Independent Component Analysis (ICA)—are often used to reduce the computational load for subsequent classification or regression tasks. However, these methods can be resource-intensive. For embedded applications, incremental PCA or low-rank matrix factorization provides a feasible alternative by updating feature subspaces in an online manner without full matrix decomposition [86].

Furthermore, feature normalization and scaling are essential for ensuring numerical stability in fixed-point implementations. Lightweight normalization schemes based on integer arithmetic are often substituted for floating-point standardization to maintain compatibility with low-power microcontrollers.

2.3.4 Machine Learning and Embedded Intelligence

In recent years, machine learning (ML) has become integral to biomedical signal processing, enabling the detection of complex patterns and the estimation of physiological parameters that are difficult to model analytically. However, the deployment of ML models in resource-constrained embedded platforms—a field commonly referred to as TinyML—requires significant model optimization and hardware-awareness [87].

In wearable systems, classical ML techniques such as Support Vector Machines (SVMs), Decision Trees, and Random Forests are often preferred over deep learning models due to their lower computational complexity and smaller memory requirements [88]. For example, hybrid models that combine Convolutional Neural Networks (CNNs) for feature extraction with Support Vector Regression (SVR) for regression tasks have demonstrated high accuracy in cuffless blood pressure estimation from PPG signals, even under embedded constraints.

When deep learning is used, models must be compressed and quantized for deployment. Techniques such as 8-bit quantization, weight pruning, and knowledge distillation can reduce memory consumption by up to 90% with minimal loss in performance [73]. Specialized toolchains such as TensorFlow Lite Micro, CMSIS-NN,

and uTensor enable efficient neural inference on ARM Cortex-M and RISC-V processors, supporting applications like arrhythmia detection, respiration estimation, and sleep stage classification [89].

A critical trend is the emergence of on-device learning, where models can adapt in situ to changing physiological conditions. Methods like online SVMs, incremental clustering, and adaptive neural networks allow the personalization of monitoring algorithms without external retraining. Nevertheless, such methods must carefully balance learning rate and stability to prevent overfitting due to noisy data or small sample sizes typical of embedded environments.

2.3.5 Real-Time Processing and Edge Optimization

Real-time operation is a fundamental requirement for biomedical devices that provide immediate feedback, such as wearable cardiac monitors, sleep trackers, or portable ultrasound probes. Real-time processing involves minimizing latency from signal acquisition to output, often requiring deterministic execution times and efficient scheduling of tasks.

Operating systems like FreeRTOS or Zephyr RTOS are commonly used in medical-grade embedded devices, offering lightweight multitasking, timing guarantees, and interrupt management. However, real-time constraints must be reconciled with energy-saving strategies such as duty cycling and event-triggered sampling. Co-optimization frameworks that jointly manage task scheduling, data buffering, and energy allocation are therefore essential.

One promising approach is approximate computing, which intentionally relaxes arithmetic precision or algorithmic fidelity in non-critical paths to save energy while maintaining acceptable output quality. For instance, approximated multipliers and truncated accumulators can yield 20–40% power savings in digital filters or neural networks without perceptible degradation in biomedical signal accuracy.

Furthermore, edge computing architectures now incorporate hardware accelerators for specific signal-processing tasks—such as FFT, convolution, and correlation—that offload computation from the CPU, improving throughput and reducing energy per operation. These architectures form the backbone of next-generation wearable intelligence, where raw data are processed locally, and only relevant summaries are transmitted to the cloud.

2.3.6 Challenges and Future Directions

Despite significant progress, several challenges persist in embedded biomedical signal processing. First, the variability and non-stationarity of physiological signals across subjects and conditions require adaptive models that can generalize well under real-world noise. Second, data annotation and ground-truth acquisition remain bottlenecks for supervised learning approaches, motivating interest in self-supervised and transfer-learning paradigms.

From a hardware perspective, advances in neuromorphic and event-driven processing offer exciting opportunities. By mimicking biological computation principles—such as sparse activation and spike-based encoding—neuromorphic chips can achieve orders-of-magnitude energy savings, making them ideal for continuous monitoring devices.

2.3. Signal Processing for Wearable and Resource-Constrained Devices

Emerging platforms like Intel Loihi and BrainScaleS demonstrate that bio-inspired architectures can perform temporal pattern recognition with sub-milliwatt power consumption.

Finally, as wearable systems increasingly incorporate multimodal sensing—combining biopotentials, optical, and acoustic signals—future signal-processing frameworks must embrace sensor fusion, context awareness, and distributed computation across edge and cloud nodes. The co-evolution of embedded hardware, efficient algorithms, and adaptive intelligence will define the next generation of wearable health systems capable of continuous, autonomous, and interpretable physiological monitoring.

CHAPTER 3

Blood Pressure Prediction Using Multi-Wavelength PPG

3.1 Background on Non-Invasive Blood Pressure Monitoring

Arterial blood pressure (BP) is one of the most fundamental physiological parameters, reflecting the mechanical force exerted by circulating blood upon the vessel walls during the cardiac cycle. It serves as a critical biomarker of cardiovascular health, providing insight into cardiac output, vascular resistance, and arterial elasticity. Abnormal blood pressure—whether elevated or reduced—represents a primary risk factor for numerous cardiovascular diseases, including coronary artery disease, stroke, heart failure, and chronic kidney disease. The World Health Organization (WHO) estimates that hypertension affects over 1.2 billion adults worldwide and contributes to nearly 10 million deaths each year [90].

Given its clinical importance, accurate and continuous BP monitoring has been a long-standing objective in both medical and engineering research. Traditional BP measurement techniques, however, remain largely episodic, cumbersome, and dependent on external hardware, making them unsuitable for real-time and long-term health tracking. As the medical paradigm shifts toward preventive, personalized, and decentralized care, there is a growing demand for non-invasive, continuous, and cuffless blood pressure monitoring technologies that can be seamlessly integrated into wearable devices [91].

The measurement of blood pressure dates back to the late 19th century, when Scipione Riva-Rocci introduced the first sphygmomanometer, employing a mercury column and inflatable cuff to determine systolic pressure. Later, Nikolai Korotkoff identified characteristic auscultatory sounds produced during cuff deflation, which became the basis for modern clinical BP assessment.

Today, conventional cuff-based non-invasive blood pressure (NIBP) measurement

relies on two primary techniques:

- Auscultatory method – based on the detection of Korotkoff sounds using a stethoscope or microphone.
- Oscillometric method – based on analyzing cuff pressure oscillations during inflation or deflation.

In both cases, an occlusive cuff is wrapped around the arm and inflated above the systolic pressure to temporarily stop arterial flow. As the cuff deflates, arterial flow resumes, producing characteristic oscillations or sounds that correspond to systolic and diastolic pressure points [91].

Despite their clinical reliability, cuff-based methods suffer from several limitations. They are intermittent, providing only snapshot readings rather than continuous monitoring; they can be uncomfortable and unsuitable for nocturnal or ambulatory use; and repeated inflation may cause vascular irritation and measurement artifacts [92]. Furthermore, these methods require proper calibration, positioning, and user compliance, which are not guaranteed in home or wearable environments. Consequently, attention has shifted toward cuffless BP estimation approaches capable of enabling continuous and unobtrusive cardiovascular monitoring.

3.1.1 Principles of Cuffless and Non-Invasive BP Estimation

Cuffless BP estimation methods exploit surrogate physiological signals that correlate with the mechanical properties of the arterial system. Among these, pulse transit time (PTT) and pulse arrival time (PAT) are the most widely studied metrics [62].

Pulse Transit Time (PTT) is defined as the time taken by the pressure wave to travel between two arterial sites. It is inversely related to the pulse wave velocity (PWV), which in turn depends on the elastic modulus of the arterial wall—and thus on the blood pressure itself, according to the Moens–Korteweg equation [93].

Pulse Arrival Time (PAT), on the other hand, measures the interval between the electrical depolarization of the heart (R-wave in ECG) and the arrival of the corresponding pulse wave (foot of the PPG signal) at a peripheral site.

A reduction in PTT or PAT typically corresponds to an increase in arterial stiffness and thus higher BP. Early studies established a linear inverse relationship between PTT and systolic/diastolic pressure under controlled conditions. However, the correlation is influenced by various confounding factors—vascular tone, arterial length, autonomic regulation, and temperature—limiting the accuracy of single-parameter models in free-living conditions [94].

Recent approaches have therefore integrated multi-parameter features derived from photoplethysmography (PPG) and electrocardiography (ECG), including waveform morphology, derivative features, and spectral indices, to capture more comprehensive hemodynamic information. Among these modalities, PPG has emerged as particularly promising due to its simplicity, low cost, and compatibility with wearable optical sensors.

3.1.2 Photoplethysmography (PPG) and Its Relationship with Blood Pressure

Photoplethysmography (PPG) is an optical sensing technique that measures the variation in light absorption or reflection caused by blood-volume changes in microvascular

tissue during the cardiac cycle [95]. The PPG signal contains two primary components:

- **a pulsatile (AC) component**, reflecting cardiac-induced blood volume changes,
- **a slowly varying (DC) component**, influenced by respiration, sympathetic activity, and vasomotor tone.

PPG signals are typically acquired using light-emitting diodes (LEDs) and photodiodes placed on the skin surface. Common measurement sites include the fingertip, wrist, and earlobe, though emerging approaches employ alternative placements such as the forehead or ankle to capture additional hemodynamic information [95].




Figure 3.1: PPG technology working principle.

The morphological features of the PPG waveform—such as systolic upstroke time, diastolic decay, dicrotic notch prominence, and amplitude ratios—are sensitive to arterial compliance and vascular tone, both of which vary with blood pressure (Fig. 3.2). Accordingly, PPG can serve as a non-invasive indicator of instantaneous hemodynamic status [90].

However, PPG signals are also highly susceptible to artifacts, including motion, ambient light interference, and contact pressure fluctuations, especially in wearable contexts. Therefore, extracting accurate BP information from PPG requires sophisticated preprocessing, denoising, and feature extraction techniques, often supported by adaptive filtering, morphological modeling, or machine learning [92].




Figure 3.2: A typical PPG signal and its components.

Traditional PPG systems employ a single or dual wavelength (typically red and infrared), mainly for oxygen saturation (SpO_2) estimation. Yet, multi-wavelength PPG (MW-PPG) expands this principle by acquiring signals at several optical wavelengths — commonly green, red, infrared, and near-infrared — to probe different tissue depths and blood components.

Each wavelength interacts differently with the skin, muscle, and blood layers according to its absorption and scattering coefficients. For instance, green light (around 530 nm) is absorbed strongly by hemoglobin and penetrates shallowly, whereas infrared light (around 940 nm) reaches deeper vascular beds [134]. By combining information from multiple wavelengths, MW-PPG captures richer hemodynamic profiles that can reflect both superficial microvascular and deeper arterial dynamics.

Recent studies have demonstrated that features extracted from multi-wavelength PPG—such as amplitude ratios, phase shifts, and cross-wavelength correlations—improve the robustness of BP estimation compared to single-channel PPG. Moreover, multi-wavelength configurations facilitate individual calibration, mitigating intersubject variability caused by skin pigmentation, tissue thickness, and temperature [96].

In this thesis, MW-PPG forms the basis for a hybrid machine learning framework that combines deep feature extraction (via convolutional neural networks) with non-linear regression (via Support Vector Regression) to achieve optimized blood pressure estimation on embedded and wearable hardware.

3.1.3 Toward Continuous and Embedded BP Monitoring

The integration of MW-PPG sensors with low-power microcontrollers and embedded digital signal processors (DSPs) enables real-time BP estimation directly on wearable platforms. This approach minimizes reliance on external computing and supports edge processing, where data are analyzed near the sensor node.

Such embedded systems face significant constraints: limited energy, restricted memory, and variable signal quality due to motion or environmental conditions. Consequently, algorithmic optimization is essential—signal processing and learning models must be both computationally efficient and robust to noise [97].

Beyond algorithmic efficiency, the success of wearable BP monitors depends on

calibration stability, sensor-skin interface design, and robust dataset collection protocols. Recent works emphasize the need for standardized databases combining multi-wavelength PPG and reference BP measurements to support reproducible research and cross-device validation [98].

The growing body of literature on cuffless BP monitoring underscores the multidisciplinary nature of the challenge, spanning biophotonics, signal processing, machine learning, and embedded system design. This thesis contributes to the field by developing and validating an optimized multi-wavelength PPG platform and a hybrid CNN–SVR model that jointly address the hardware and algorithmic constraints inherent in real-world wearable BP estimation.

3.2 Hardware Architecture and Wearable Implementation

The development of a reliable and energy-efficient platform for multi-wavelength photoplethysmographic (MW-PPG) acquisition represents a crucial step toward the realization of a cuffless blood pressure estimation system. The prototype designed in this work combines the flexibility of a development environment with the precision of biomedical-grade sensing, providing a controlled framework for the validation of optical and algorithmic methodologies.

The electronic system was implemented using a programmable embedded architecture centered on an integrated optical module (MAX86916, Analog Devices Inc. [99]) interfaced with a low-power ARM Cortex-M4F microcontroller (MAX32630FTHR). This configuration ensures high signal fidelity, low power consumption, and complete programmability of the acquisition parameters, enabling systematic exploration of wavelength combinations, sampling conditions, and signal quality across different measurement setups.

Unlike conventional laboratory instruments, the proposed prototype was conceived as a compact, semi-wearable device capable of operating continuously in real-world conditions while maintaining the flexibility required for algorithm development and benchmarking. The hardware and software layers were co-designed to ensure synchronization between emission, detection, and data logging, providing a coherent digital acquisition chain suitable for multi-sensor fusion and advanced physiological modeling.

3.2.1 System Overview and Design Objectives

The overall system architecture integrates the optical front-end, the analog and digital acquisition circuitry, the embedded control unit, and the data communication interface. The main design objectives were fourfold:

- ***Multi-wavelength acquisition*** — simultaneous capture of optical signals at red, infrared, green, and blue wavelengths to probe multiple vascular layers and improve robustness of physiological estimation.
- ***Low-noise amplification and conditioning*** — minimization of electrical noise and ambient-light interference to maintain a high signal-to-noise ratio (SNR) even under motion.




Figure 3.3: MAX86916 EV System Board [99].

- ***Embedded synchronization and preprocessing*** — real-time control of LED emission, ADC sampling, and digital filtering within the microcontroller to ensure timing precision and reduce transmission load.
- ***Energy efficiency and portability*** — adoption of low-power components and dynamic duty-cycling strategies to support extended monitoring sessions under battery operation.

The resulting system forms a cohesive sensing and computation platform capable of delivering high-quality multi-spectral PPG data for real-time processing or offline analysis.

3.2.2 Optical Front-End and Sensing Module

At the heart of the system lies the MAX86916 optical front-end, an integrated biosensing module designed for reflective PPG acquisition (Fig. 3.3). The device combines high-performance optoelectronic components—four LEDs and a photodiode—with dedicated analog circuitry for current driving, light rejection, and preamplification.

The LED array includes emission sources centered at four distinct wavelengths:

- **Blue (455–466 nm):** probing superficial capillary beds and sensitive to skin perfusion;
- **Green (520–535 nm):** optimal for peripheral perfusion monitoring and pulse detection;
- **Red (655–663 nm):** enabling oxygen saturation estimation through differential absorption;
- **Infrared (930–955 nm):** penetrating deeper vascular structures, useful for assessing arterial compliance and pulse transit characteristics.

Each LED is driven by an internal programmable current driver (0–200 mA range) with digital control over pulse width, repetition rate, and duty cycle. The emission sequence is time-multiplexed, ensuring that only one wavelength is active at a time, thereby eliminating optical crosstalk and allowing precise separation of the four spectral channels.

The photodiode is a low-capacitance, high-sensitivity silicon detector with a spectral response covering 400–1000 nm and a responsivity of approximately 0.5 A/W. It is designed in a reflective configuration—emitter and detector positioned adjacent to one another—suitable for fingertip and wrist measurements. An internal analog chain amplifies the photodiode current, applies low-pass filtering to suppress ambient-light

fluctuations, and digitizes the signal through a 16-bit ADC operating at 100 Hz. The internal light-rejection circuitry and correlated double-sampling (CDS) technique effectively mitigate DC drift and flicker noise, ensuring stable signal baselines even under varying illumination.

The module operates at 1.8 V for logic and analog subsystems, and at 5.5 V for the LED driver stage, thus separating the noise-sensitive acquisition path from high-current switching elements. All communication with the host microcontroller is performed through a digital I²C interface supporting up to 400 kHz clock speed.

3.2.3 Embedded Control and Processing Unit

The MAX32630FTHR board serves as the central processing unit for signal acquisition, synchronization, and control. It features an ARM Cortex-M4F processor operating at 96 MHz, equipped with floating-point arithmetic and DSP extensions, 512 kB of SRAM, and 1 MB of on-chip Flash. The MCU provides a high-speed I²C interface used to configure the optical module (LED sequencing, current control, sampling rate), manage data transfer, and ensure precise timing synchronization between optical channels.

The firmware architecture was developed in C/C++ using the Mbed OS environment, enabling modular implementation of acquisition routines, digital filtering, and communication management. A FreeRTOS-based scheduler ensures deterministic timing for concurrent tasks, including LED control, ADC readout, buffering, and USB transmission.

In addition to signal acquisition, the MCU is capable of performing basic preprocessing—such as moving-average smoothing, baseline correction, and peak detection—thereby reducing the computational load on the host PC and preparing the data for feature extraction or machine learning pipelines.

The board communicates with the host computer through two interfaces:

- A USB–FTDI link for continuous streaming of PPG samples and command exchange;
- A Micro-USB connection providing both power and firmware upload capability.

The dual-channel data transfer ensures robustness and minimizes latency, enabling near real-time visualization of acquired biosignals during experimentation.

3.2.4 Power Supply and Energy Management

The entire acquisition system was designed for low-power operation, allowing potential transition to a wearable configuration. The MAX86916's dual-supply configuration isolates the LED driver from the analog sensing path, reducing switching noise and improving SNR. The microcontroller board integrates ultra-low-dropout (LDO) regulators to maintain stable voltage levels, with quiescent currents below 50 μ A.

Dynamic power management was implemented through software control of LED currents and sampling rates, allowing trade-offs between energy consumption and signal amplitude. In typical operation at 100 Hz with moderate LED drive currents (10–20 mA), total system power consumption remained below 15 mW, corresponding to over 15 hours of continuous acquisition on a 250 mAh lithium-polymer cell.

3.2. Hardware Architecture and Wearable Implementation




Figure 3.4: MATLAB visualization of four raw PPG signals corresponding to four different wavelengths.

Future iterations could further enhance autonomy through duty-cycling, local buffering, and energy-harvesting solutions (e.g., thermoelectric or motion-based scavenging), paving the way toward fully self-powered physiological monitoring devices.

3.2.5 Data Acquisition Workflow and Software Interface

Data acquisition was managed through the Analog Devices Device Studio, a dedicated software suite interfaced with the evaluation system. The software allowed real-time visualization of multi-channel PPG signals, interactive adjustment of acquisition parameters (LED sequence, pulse width, integration time), and export of data to CSV format for post-processing.

Each recorded dataset contained synchronized samples from all four wavelengths, stored with precise timestamps. The 100 Hz sampling frequency provided adequate temporal resolution for detecting cardiac cycles and computing derived parameters such as heart rate (HR), pulse rate variability, and pulse transit time (PTT).

The raw PPG data were subsequently processed in MATLAB and Python environments, where additional filtering, detrending, and segmentation were applied. These preprocessed signals formed the input to the hybrid machine learning models described in later chapters (Fig. 3.4).

Synchronization between emission and detection was verified experimentally, achieving timing accuracy within ± 1 ms across wavelengths. This precision was critical for reliable inter-channel analysis and fusion, particularly in algorithms relying on spectral phase and amplitude differences.

3.2.6 Mechanical Configuration and Experimental Setup

Although initially developed as a benchtop prototype, the system was designed with wearable ergonomics in mind. The optical sensor was mounted on a soft, biocompatible support, ensuring stable skin contact and minimizing motion-induced artifacts during acquisition. The mechanical configuration allowed flexible positioning—typically at

the index fingertip or volar wrist—depending on experimental conditions.

During trials, participants were seated at rest, and the optical module maintained consistent coupling without significant motion distortion. The high mechanical stability of the sensor housing ensured reliable light transmission and reflection characteristics, critical for maintaining signal fidelity across sessions.

Preliminary usability tests demonstrated that the setup could be comfortably worn for sessions exceeding 30 minutes without skin irritation or detachment. These results confirm that the system design can be readily adapted to a fully wearable configuration by integrating compact power management and wireless transmission modules.

3.2.7 System Validation and Performance Evaluation

Prior to deployment, the hardware underwent extensive functional testing, including verification of signal linearity, inter-channel calibration, and SNR evaluation under varying light and temperature conditions. The system exhibited a linear response across LED drive currents, with negligible saturation and consistent amplitude scaling across all spectral bands. Signal-to-noise ratio measurements exceeded 50 dB in static conditions, confirming the effectiveness of internal filtering and ambient light rejection.

The system demonstrated reliable operation across a temperature range of 20–35 °C, ensuring thermal stability suitable for clinical or field applications. Real-time visualization confirmed minimal inter-channel latency, and the integrity of waveform morphology was preserved throughout long-term acquisition sessions.

These characteristics validate the suitability of the developed prototype as a reference platform for algorithmic exploration and dataset generation, bridging the gap between controlled laboratory instrumentation and practical wearable sensing.

3.3 Dataset Creation and Experimental Protocol

The development of an accurate blood pressure (BP) estimation model based on multi-wavelength photoplethysmography (MW-PPG) requires a robust and well-structured dataset, encompassing synchronized optical and reference physiological measurements acquired under controlled but realistic conditions. This section presents the dataset creation and experimental protocol employed in this work, which resulted in a comprehensive database of MW-PPG recordings and reference BP measurements from human participants. The methodology, detailed in [152], aimed to capture high-quality optical signals representative of different hemodynamic states while ensuring statistical rigor and reproducibility.

The experimental campaign was designed to:

- a. Validate the performance of the custom wearable MW-PPG platform described in Section 3.2;
- b. Provide an open, well-annotated dataset for the training and validation of BP estimation algorithms;
- c. Establish a protocol for multi-site, multi-wavelength optical acquisition synchronized with oscillometric reference measurements.

The final dataset includes synchronized recordings of red (660 nm), infrared (880 nm), green (530 nm) and blue (460 nm) photoplethysmography channels, each acquired at a sampling rate of 100 Hz, along with reference systolic and diastolic BP values, heart rate (HR), and demographic metadata.

The acquisition system was centered on the MAX86916 Evaluation Board by Analog Devices. The PPG signals were sampled at 100 Hz, which was maintained throughout the study to ensure factory-calibrated timing and low-noise performance. From a signal processing standpoint, 100 Hz is well above the Nyquist requirement for the physiological components of interest, which lie below 10 Hz—even under tachycardic conditions. This choice also aligns with IEC 60601-2-27 guidelines and prior literature on cuffless BP estimation. Moreover, the 100 Hz stream enabled efficient windowing (30 seconds = 3000 samples) and low-complexity signal processing, making it suitable for implementation on resource-constrained wearable devices.

The system was configured in reflective geometry, where both the LEDs and the photodiode are co-located on a single substrate, detecting light backscattered from the fingertip tissue. This configuration was chosen for its simplicity, low power consumption, and compatibility with wearable and portable implementations.

The sampling frequency was set to 100 Hz, which ensures accurate representation of the cardiac and respiratory spectral components (up to 15 Hz), while keeping data volume manageable for embedded applications. Each recording lasted 30 seconds, corresponding to 3000 samples per wavelength channel per subject.

The acquisition site was the index fingertip of the non-dominant hand. Before recording, the finger was cleaned to remove residual oils or lotions, and the subject was instructed to remain still and avoid excessive pressure on the sensor. The probe was secured using an adjustable 3D-printed clip with soft silicone padding to maintain consistent optical contact and reduce motion-induced artifacts. Ambient light was minimized, and data were collected in a quiet laboratory environment at 23 ± 1 °C.

3.3.1 Optical and Electronic Configuration

Each LED channel was driven by an independent programmable current driver with an intensity range between 0 and 100 mA, adjusted automatically to maintain the optimal ADC range without saturation. The analog front-end embedded in the MAX86916 included:

- a low-noise transimpedance amplifier (TIA) for current-to-voltage conversion;
- a programmable gain amplifier (PGA) with selectable gains (1×–16×);
- and an integrated 16-bit analog-to-digital converter (ADC).

Signals were stored in 16-bit integer format and exported from DeviceStudio as comma-separated values (CSV), each file containing timestamped pairs of red and infrared samples. The DeviceStudio v2.1 software was also used to monitor real-time waveforms during acquisition, enabling immediate detection of signal loss or optical saturation.

3.3.2 Experimental Protocol and Acquisitions

Each participant underwent a 30-second simultaneous recording session, during which the following parameters were acquired:

- *PPG signals* (red, green, blue, infrared) from the right index finger, positioned on the sensor with the hand resting on a stable support at heart level;
- *Systolic and diastolic blood pressure* (SBP and DBP) and *heart rate*, from the left upper arm, using an OMRON® M6 Comfort sphygmomanometer (± 3 mmHg accuracy), with the cuff automatically inflating during the PPG recording window (as shown in Fig. 3.5);
- *Body temperature*, collected at the start of each session using an infrared thermometer.




Figure 3.5: *Experimental set-up for blood pressure data acquisitions: PPG signals were acquired by positioning the right index finger on the PPG sensors and SBP and DBP via sphygmomanometer.*

To minimize motion artifacts and postural bias, subjects were seated in an upright position with both arms supported at heart level, as recommended by AHA guidelines. Each participant remained at rest for at least five minutes prior to the measurement. Environmental conditions (temperature and lighting) were kept stable throughout the sessions.

To ensure metrological traceability and consistency of the measurements, all instrumentation used in the acquisition setup was verified and, where applicable, calibrated according to manufacturer specifications and international standards. The reference sphygmomanometer (OMRON® M6 Comfort) is a clinically validated automatic oscillometric device. Before data collection, the device was checked using a calibration tester (RIGEL BP-SIM NIBP simulator), confirming pressure accuracy within ± 3 mmHg across the full physiological range (80–180 mmHg for SBP, 50–100 mmHg for

3.3. Dataset Creation and Experimental Protocol

DBP). The device was used as the ground truth reference for all blood pressure values reported in the dataset. The PPG acquisition module (MAX86916 Evaluation Kit) does not require user calibration but was tested for linearity and signal integrity under controlled light and reflectance conditions. A lab bench power supply was used to ensure a stable 1.8V and 5.5V dual-voltage supply during operation. The analog signal chain was assessed by capturing signals from a calibrated reflective optical phantom to validate baseline noise levels, spectral consistency across wavelengths, and absence of electrical interference. All recorded signals were reviewed for quality control. Sessions with signal clipping, saturation, or evident movement artifacts were flagged and excluded from the final dataset.

To ensure temporal alignment between the PPG acquisition and the reference blood pressure measurements, a simple yet reliable synchronization protocol was adopted. At the start of each session, the operator initiated the PPG recording and then immediately triggered the automatic inflation of the OMRON® upper-arm sphygmomanometer, ensuring that the entire blood pressure measurement process occurred within the 30-second PPG recording window (Fig. 3.6). Since the inflation and deflation of the cuff takes approximately 20–25 seconds, this window fully captures the physiological effects associated with the oscillometric BP measurement, including the heart rate modulation and pressure variations. Although no hardware-level trigger was implemented between devices, this controlled timing sequence was sufficient to guarantee that the PPG segment used for training corresponded to the same physiological state during which the sphygmomanometer determined the reference SBP and DBP values. By adhering to rigorous validation and verification practices for both measurement devices and experimental procedures, the integrity and reproducibility of the dataset were ensured, supporting its suitability for both model development and clinical benchmarking purposes.




Figure 3.6: A conceptual framework highlighting device synchronization: the PPG sensor (right index finger) and the sphygmomanometer (left arm) operate simultaneously to ensure aligned signal and reference blood pressure measurements.

Table 3.1: Systolic and Diastolic Blood Pressure Across Different Age Groups

Age Group	Systolic Blood Pressure			Diastolic Blood Pressure		
	Min	Max	Mean $\pm$ SD	Min	Max	Mean $\pm$ SD
10–29	95	140	114.9 $\pm$ 13.6	51	90	69.8 $\pm$ 9.4
30–49	102	148	117.9 $\pm$ 13.8	60	90	74.7 $\pm$ 9.0
50–69	96	171	126.5 $\pm$ 16.1	54	100	77.9 $\pm$ 12.5
70–89	109	158	138.5 $\pm$ 23.2	58	72	64.2 $\pm$ 5.8

The dataset comprises recordings from 88 participants, with an equal number of men and women, aged between 19 and 84 years. To ensure diversity across age groups, the subjects were divided into four categories: 37 participants aged 10-29, 12 in the 30-49 age range, 25 in the 50-69 group, and 14 individuals aged 70-89.

Sample size ($n = 88$) is calculated by using the formula for estimating the mean with a specified margin of error (Eq. 3.1):

$$n = \frac{Z^2 \cdot \sigma^2}{E^2} \quad (3.1)$$

where:

- $Z = 1.96$, the z-score corresponding to the level of confidence;
- $\sigma = 12$ mmHg, the standard deviation estimation, as a compromise between systolic and diastolic pressure standard deviation;
- $E = 2.5$ mmHg, the maximum desired distance between the estimation of the sample mean and the actual population mean.

The data shows a systolic pressure range from 95 to 171 mmHg, with a mean value equal to 119.68 mmHg and a standard deviation of 15.96 mmHg, while the diastolic pressure range is made up of 51 mmHg as minimum and 100 mmHg as maximum, with an average DBP of 71.98 mmHg and a standard deviation of 9.99 mmHg (Tab. 3.1). Since age plays a fundamental role in the evaluation of blood pressure, a statistical analysis of systolic and diastolic pressure values was carried out, dividing the population into age groups. Boxplots are shown in Fig. 3.7.




Figure 3.7: Boxplots for BP across different age-groups: (a) SBP and (b) DBP.

The boxplots clearly illustrate that SBP increases progressively with age, while DBP tends to decline after the seventh decade of life. This trend reinforces the statistical

reliability of our dataset, as it aligns with well-documented physiological changes associated with aging. Specifically, the stiffening of large arteries over time leads to higher SBP and lower DBP for a given level of vascular resistance and blood flow, a phenomenon extensively studied in cardiovascular research.

3.4 Signal Processing and Feature Extraction

Once the PPG signals have been acquired using the MAX86916 EV System, a comprehensive signal processing pipeline was implemented to ensure the extraction of physiologically meaningful features and reduce contamination from noise sources. This phase is critical, as PPG signals are inherently susceptible to several forms of noise, including motion artifacts, ambient light interference, baseline drift, and high-frequency electronic noise. All the described algorithms were implemented in MATLAB 2023a and executed on a standard laptop.

The pre-processing steps are summarized below (Fig. 3.8):

1. *Outliers Removal*: To ensure data integrity and mitigate the influence of outlier blood pressure values on model performance, subjects with systolic or diastolic blood pressure values beyond the range of mean ± 3 standard deviations were excluded. This statistical criterion resulted in a curated dataset comprising 81 subjects, removing extreme cases likely due to erroneous measurements or physiological anomalies.
2. *Signal Length Standardization*: Each PPG recording was trimmed or zero-padded to a uniform duration of 30 seconds (equivalent to 3000 samples at 100 Hz), enabling consistent feature extraction across the dataset and preventing variable-length sequences from biasing the learning models.
3. *Motion Peaks Removal*: Motion artifacts were addressed through a combination of signal quality assessment and peak detection. Specifically, abnormally large or irregular peaks in the red and infrared channels—associated with gross finger movements or sensor displacement—were detected using an adaptive threshold based on signal envelope statistics (median absolute deviation). Segments affected by motion were removed or replaced via linear interpolation. The choice of removing motion artifacts only in PPG signals in the red and infrared wavelengths was based on the fact that shorter wavelengths (such as green and blue), due to their limited skin penetration, do not capture as much cardiac activity, but are also less susceptible to motion artifacts. As illustrated in Fig. 3.9, long wavelengths like red and infrared penetrate deeper into the skin and are more effective in capturing heart and blood vessel activity. However, these wavelengths are also more prone to motion artifacts due to their extended light path.
4. *Filtering*: A zero-phase 4th-order Butterworth band-pass filter with a passband of 0.5 Hz to 5 Hz was applied to each PPG signal. This frequency range was selected to retain the fundamental and harmonic components of the cardiac cycle (typically within 0.8–3.5 Hz) while suppressing low-frequency components (<0.5 Hz) associated with baseline wander (e.g., respiratory-induced fluctuations or thermal drift) and high-frequency noise (>5 Hz) such as electronic interference

or ambient light pickup. The use of zero-phase filtering ensures that the temporal morphology of the PPG waveform is preserved, which is essential for accurate extraction of time-domain features such as systolic and diastolic phases, pulse width, and peak-to-peak intervals.

5. *Normalization*: After filtering, each signal was normalized using z-score normalization (zero mean and unit variance). This step reduces inter-subject variability due to anatomical and physiological differences (e.g., skin tone, finger thickness, perfusion) and facilitates model training by centering the feature distributions.




Figure 3.8: Block diagram of the signal pre-processing steps.




Figure 3.9: Four PPG signals corresponding to four different wavelengths. IR and RED signals show rapid and wide peaks that should be removed.

3.5. Machine Learning Models: MLP, PCA, and CNN Approaches




Figure 3.10: Temporal and spectral features derived from the morphology of the PPG signal and the corresponding power spectral density.

3.4.1 Features

Blood pressure estimation is a challenging task, which aims to find the elements of correlation between the signal used (in this case, only the PPG signal) and the physiological information that allows this parameter to be calculated. To do this, a feature-to-label algorithm was created, which allows to derive the desired values from some characteristics of the signal, both temporal and spectral. Three different Machine Learning algorithms were tested, which have 21 features in common extracted from each of the four PPG signals (Fig. 3.10):

- 15 temporal features, related to the morphology of the PPG signal in the time domain, obtained by considering the main intervals within a single cardiac cycle;
- 6 spectral features, obtained directly from the power spectral density of the signal by considering the amplitude and the highest frequency of three frequency bands.

In particular:

- Cardiac Period;
- Systolic Upstroke Period;
- Diastolic Time;
- Systolic Width at 10, 25, 33, 50, 66, 75 %;
- Diastolic Width at 10, 25, 33, 50, 66, 75 %;
- Low, Median, High Frequency Band - Amplitude;
- Low, Median, High Frequency Band - Frequency.

3.5 Machine Learning Models: MLP, PCA, and CNN Approaches

In this study, we investigate the effectiveness of three distinct machine learning models for predicting arterial pressure using PPG signals. Our objective was to assess the strengths and limitations of each model and to identify the most reliable approach for non-invasive blood pressure monitoring.

The first model we implemented was a MLP, trained on the full set of features extracted from PPG signals across four different wavelengths (red, green, blue, and infrared). A total of 84 features (21 per wavelength) were used, capturing both temporal and spectral characteristics of the signal, such as cardiac periods and frequency-domain features. The MLP architecture consisted of an input layer with 84 neurons, followed by two hidden layers with 64 and 32 neurons, respectively, using ReLU activation functions. To optimize training, the Adam optimizer was used with a learning rate of 0.001, and Mean Squared Error (MSE) was selected as the loss function. To prevent overfitting, a dropout regularization (rate = 0.2) was applied. Although this large feature set allows the model to capture a wide range of physiological data relevant to arterial pressure regulation, it also introduces challenges. The primary issues include the risk of overfitting, as the model may learn noise in the data, and redundancy or collinearity among features, which can make it harder for the algorithm to discern meaningful patterns. Moreover, the "curse of dimensionality" means that as the number of features grows, the data becomes sparser, making it more difficult to find significant clusters or relationships within the feature space.

To address these limitations, we applied PCA as a dimensionality reduction technique, which transformed the original 84 features into 10 principal components. These components preserve the majority of the variance in the data while significantly reducing the dimensionality of the feature space. A second MLP model was then trained using this reduced feature set, maintaining the same network architecture and hyperparameters as the full MLP model, except for the input layer, which was adapted to the 10-dimensional feature space. This allowed us to evaluate whether a reduced feature set could enhance the model's generalization by mitigating overfitting and simplifying the complexity of the data. The PCA-based approach aims to improve performance by focusing on the most relevant aspects of the signal while reducing the noise and redundancy that may have degraded the performance of the original MLP.

In our third approach, we explored the use of a CNN. Instead of working with individual features from the different wavelengths, we computed the L2-norm of the signals, which provided a compact representation of the physiological information contained within the PPG signals. The 21 features of the obtained L2-norm signal were used as input for the CNN, which is particularly effective at capturing spatial and temporal dependencies in structured data. The CNN architecture comprised two convolutional layers, each with 32 filters of size 3×3 , followed by a max-pooling layer to reduce feature map dimensions. The extracted representations were then passed through a fully connected layer with 64 neurons before reaching the output layer. The network was trained using the Adam optimizer with a learning rate of 0.0005 and ReLU activation functions in all layers. To enhance generalization, batch normalization and dropout regularization (rate = 0.3) were applied. CNNs can extract intricate patterns by learning from multiple layers of convolutional filters, making them ideal for processing multi-dimensional data like PPG signals. This model offers a novel and practical method for summarizing complex physiological information, thereby enhancing its predictive capabilities for arterial pressure.

The comparative analysis of these three models sought to determine the most effective strategy for predicting blood pressure. We evaluated their performance in terms of accuracy, with metrics like mean absolute error (MAE) and standard deviation (SD)

3.6. Hybrid CNN-SVR Model for Optimized BP Estimation

used to compare predictions against the actual measured values. The performance of the MLP models—with and without PCA—allowed us to assess the impact of dimensionality reduction on predictive accuracy. On the other hand, the CNN’s ability to model dependencies in the data was tested to determine whether it could outperform the more traditional MLP approaches.

Table 3.2: Summary of the implemented machine learning models for arterial pressure estimation from multi-wavelength PPG signals.

Model	Input Features	Architecture and Parameters	Main Advantages / Limitations
MLP (Full Feature Set)	84 handcrafted features (21 per wavelength: red, green, blue, infrared), including temporal and spectral parameters	• 3 layers: 84–64–32–1 neurons • ReLU activation • Adam optimizer, learning rate = 0.001 • Dropout = 0.2 • Loss: Mean Squared Error (MSE)	• Captures diverse physiological features • High flexibility • Risk of overfitting and feature redundancy • Sensitive to collinearity and noise
MLP + PCA (Reduced Feature Set)	10 principal components derived from PCA of the 84-feature space (variance-preserving projection)	• Same architecture as MLP, input = 10 • Adam optimizer, learning rate = 0.001 • Dropout = 0.2	• Reduces dimensionality and redundancy • Enhances generalization and stability • May lose subtle physiological information
CNN (L2-Norm Input)	21 features derived from L2-norm of four-wavelength PPG signals (compact temporal representation)	• 2 convolutional layers (32 filters, 3×3 kernel) • Max pooling + fully connected (64 neurons) • ReLU activation, Adam (lr = 0.0005) • Batch normalization + dropout = 0.3	• Learns spatio-temporal dependencies • Robust to noise and redundant data • Requires larger datasets for stability

3.6 Hybrid CNN-SVR Model for Optimized BP Estimation

To further improve the accuracy of blood pressure estimation, particularly addressing the decline in performance observed in elderly populations, we propose an additional correction model that integrates physiological parameters such as age, temperature, and heart rate into the existing CNN framework. Our preliminary analysis indicated that the CNN prediction errors increased with the age of the subjects, highlighting the need for a model that better captures the physiological changes that occur with aging, such as

vascular resistance and skin thickness. However, it is important to note that the dataset used in this study is homogeneous in terms of age distribution, and we observed a low correlation between age and blood pressure (Pearson's $r = 0.13$ for systolic and -0.08 for diastolic pressure). This suggests that age does not introduce significant redundancy in the model, reinforcing the validity of the extracted PPG-based features for blood pressure estimation.

In this phase, we developed a hybrid model that relies on a SVR to correct the predictions from the CNN based on the aforementioned physiological parameters (see Figure 3.11). After the CNN has predicted SBP and DBP using features extracted from the PPG signal, the SVR model refines these predictions by considering the patient-specific factors. The SVR takes as input the CNN's residual errors (the difference between predicted and real blood pressure values) and the physiological parameters, and outputs a correction factor to be added to the CNN predictions. The SVR model employed a radial basis function (RBF) kernel with a regularization parameter $C = 1.0$ and an epsilon value of 0.1 , allowing it to learn correction factors that adjusted the CNN output without requiring ground-truth BP values during real-time application.




Figure 3.11: Block diagram of the hybrid model (CNN cascading with SVR).

The dataset was divided into training and test sets to evaluate the performance of this hybrid approach, with metrics such as MAE and SD used to compare the corrected and uncorrected CNN predictions. By combining the predictive power of the CNN with a regression-based correction model, we aim to improve the model's robustness across age groups and physiological variations, thus enhancing the generalizability of non-invasive blood pressure monitoring tools. This combination of machine learning models allowed us to systematically analyze the impact of feature selection, dimensionality reduction, deep feature extraction, and residual error correction, leading to an optimized framework for non-invasive BP estimation.

3.7 Experimental Results and Discussion

In this section, we present the results from three distinct machine learning approaches: a MLP model trained on 84 features (21 per PPG signal), a modified MLP model using PCA for feature reduction, and a CNN architecture that utilizes the L2-norm of the four PPG signals. The validation of these models was carried out using a rigorous 5-fold cross-validation technique. To create training and testing subsets from the dataset of 81 participants, the data was divided into five equally sized subsets, or "folds". Each fold contains approximately 17.6 individuals, meaning some folds contain 17 participants while others have 18. The cross-validation process involves five iterations, where in

each round, one fold is used as the test set while the remaining four folds serve as the training set.

By averaging the performance metrics across all five folds, we ensure a more reliable and robust estimate of the models' performance, compared to evaluating based on a single train-test split. The models' performance was assessed using several key metrics, including MAE, SD, and Cumulative Error Percentage. In order to further validate the models, the results were benchmarked against the criteria established by the AAMI and the BHS.

The primary goal of this evaluation was to meet the AAMI standard, which requires a *Mean Absolute Error of no more than 5 mmHg* and a *Standard Deviation of no more than 8 mmHg* for both systolic and diastolic blood pressure. Additionally, the BHS grading system was used to categorize the accuracy of the models. The BHS assigns grades from A to D based on the percentage of predictions that fall within certain error thresholds when compared to a reference device, typically a sphygmomanometer. A grade A indicates that at least 95% of the predictions are within 15 mmHg, at least 85% are within 10 mmHg, and at least 60% deviate by no more than 5 mmHg. Grade B requires that 90%, 75%, and 50% of predictions meet these respective thresholds. Grade C requires 85%, 65%, and 40%, while Grade D is assigned to devices that perform below these standards (as shown in Table 3.4).

3.7.1 Multi-Layer Perceptron - All features

The first machine learning model for arterial pressure prediction utilizes a dataset of 84 features per subject, derived from four PPG signals, each corresponding to a different wavelength: infrared, red, green, and blue. These signals are processed to extract 21 distinct features per signal, capturing key physiological and spectral characteristics. The model is based on a MLP architecture, which includes an input layer with 84 neurons (one for each feature), eight hidden layers, and an output layer consisting of two neurons, responsible for predicting systolic and diastolic blood pressure values.

The results of this approach are presented in Tables 3.3 and 3.4. Unfortunately, this method falls short of meeting the AAMI requirements, and based on the BHS grading system, it achieves a Grade C for diastolic pressure and a Grade A for systolic pressure.

These suboptimal results can likely be attributed to the large number of features used, which may degrade model performance due to the increased computational burden, heightened risk of overfitting, and the presence of redundancy or multicollinearity among the features. As such, careful feature selection and preprocessing are critical to ensure that the model can accurately learn relevant patterns and generalize effectively using new data.

3.7.2 Multi-Layer Perceptron - Feature Reduction

In order to avoid the curse of dimensionality, the second approach is basically focused on dimensionality reduction, relying on Principal Component Analysis operated on the original 84 features. The ranking of the 10 most relevant features from PCA is the following, which explain 96% of data variability (colors refer to four different wavelengths):

1. Diastolic Time RED

2. Cardiac Period RED
3. Systolic Upstroke Period IR
4. Systolic Width – 10% IR
5. Systolic Width – 25% IR
6. Diastolic Time GREEN
7. Diastolic Time BLUE
8. Systolic Width – 33% IR
9. Cardiac Period IR
10. Systolic Width – 50% IR

The same model architecture was utilized—a MLP with 10 input neurons, eight hidden layers, and two output neurons—to evaluate the actual advantages of feature reduction. As shown in Table 3.3, despite implementing PCA for dimensionality reduction, its impact in this study appears to be limited. The results do not satisfy the AAMI criteria and are classified as Grade D according to BHS standards.

While PCA is designed to minimize the dimensionality of the feature space by projecting the data into a lower-dimensional subspace while retaining maximum variance, it became clear that this method inadvertently removed features that could have provided critical information for arterial pressure prediction. Consequently, relying exclusively on PCA for feature selection is insufficient, underscoring the necessity for alternative methods that can maintain a broader range of features while still effectively reducing dimensionality.

3.7.3 Convolutional Neural Network

As PCA highlighted significant features linked to each wavelength, it is clear that omitting any one of the signals could result in missing crucial physiological data pertinent to predicting arterial pressure. By computing the L2 norm of the four signals, we consolidate the contributions from all wavelengths into a single scalar value. This approach ensures that all the information captured by each wavelength is retained, allowing for a more comprehensive representation of the physiological data.

$$||x|| = \sqrt{\sum_{k=R,G,B,IR} |x_k|^2} \quad (3.2)$$

The choice of the L2 norm as a representative approximation for the four signals is driven by its effectiveness in capturing the overall magnitude or energy content of the signals while simplifying the analysis. By focusing on the L2 norm, we reduce the complexity inherent in examining each signal separately, which streamlines the feature space and diminishes the dimensionality of the input data. This simplification facilitates the operation of machine learning algorithms while still retaining critical information about the overall characteristics of the signals.

The architecture of the network comprises a sequence of one-dimensional convolutional layers followed by max-pooling layers, which progressively decrease the size of the feature maps. The resulting feature maps are then flattened and directed through two fully connected layers for final regression tasks. In our study, the 1D CNN was designed with three convolutional layers, each employing filters of size 3. Each convolutional layer is succeeded by max-pooling layers with a pooling size of 2. The densely connected layers include a hidden unit that utilizes ReLU activation, along with an output layer that consists of two output units, enabling the network to predict the numerical values of both systolic and diastolic blood pressure directly.

As demonstrated in Tables 3.3 and 3.4, the results from the CNN applied to the L2-norm signal conform to the AAMI guidelines, categorizing performance as Grade A for SBP and Grade B for DBP based on BHS standards. Specifically, for the test set, we observed an MAE of 5.52 mmHg and an SD of 7.62 mmHg for systolic blood pressure (SBP), and an MAE of 4.67 mmHg with an SD of 6.63 mmHg for diastolic blood pressure (DBP) (see Tab. 3.3). According to the BHS protocol, the model achieved the following results: for systolic blood pressure (SBP), 25% of predictions were within 5 mmHg, 90% within 10 mmHg, and 95% within 15 mmHg; for diastolic blood pressure (DBP), 35% were within 5 mmHg, 80% within 10 mmHg, and 95% within 15 mmHg. Based on these results, the model qualifies for Grade B for SBP and Grade C for DBP, as it does not meet the stricter thresholds required for higher grades across both pressure categories.

Utilizing the same CNN architecture, which has proven effective for predicting SBP and DBP, the mean arterial pressure (MAP) is also computed, a crucial indicator of tissue perfusion in clinical settings. The MAP is given by Eq. 3.3.

$$\text{MAP} = \frac{\text{SBP} + 2 \times \text{DBP}}{3} \quad (3.3)$$

Promising performances are recorded, with an MAE and SD equal to 6.73 ± 7.82 , paving the way for future developments and further investigations.

The validation of results is also conducted using Bland-Altman plots (Fig. 3.12), which reveal a consistent behavior of differences at low and moderate blood pressure levels, while highlighting a trend in the higher blood pressure range. This indicates that the discrepancies in measurements increase with higher blood pressure values, suggesting a correlation between the magnitude of the blood pressure reading and the error in measurement.

Since our dataset is homogeneous in terms of age distribution, this trend is unlikely to be caused by age-related factors. Instead, it may be attributed to several factors. First, the model may exhibit reduced accuracy in extreme blood pressure ranges due to a lower representation of these values in the training set, leading to potential bias. Second, physiological factors such as increased arterial stiffness at higher pressures could alter the PPG waveform characteristics, making predictions more challenging. Additionally, the nonlinear relationship between blood pressure and PPG-derived features may not be fully captured by the model, especially in extreme cases.

The Bland-Altman plots display confidence intervals that highlight some variability in the differences observed between the two methods, particularly for systolic blood pressure. However, the inclusion of zero within these intervals suggests that, on average, there is no significant systematic bias between the two measurement techniques.

Table 3.3: MAE and SD Expressed in mmHg

	Systolic BP	Diastolic BP
MLP - 84 Features	10.91 $\pm$ 12.67	5.08 $\pm$ 6.93
MLP - PCA Reduction	18.91 $\pm$ 20.95	12.49 $\pm$ 14.15
CNN L2-norm	5.52 $\pm$ 7.62	4.67 $\pm$ 6.63
CNN + SVR	1.84 $\pm$ 3.27	1.98 $\pm$ 3.47

In contrast, the narrower interval for diastolic blood pressure indicates a somewhat greater level of agreement compared to systolic blood pressure measurements. Further refinements to the model, such as incorporating additional physiological parameters or nonlinear correction techniques, may help mitigate this effect and improve performance in extreme blood pressure ranges.

Overall, these findings highlight the potential of using CNN and the L2 norm as reliable, clinically applicable tools for non-invasive arterial pressure prediction. They provide essential insights for monitoring and managing cardiovascular health, reinforcing the efficacy of this approach in clinical settings.




Figure 3.12: Bland-Altman plots for blood pressure estimation using a CNN model. (a) Bland-Altman plot for SBP showing the mean difference of -3.90 mmHg, with upper and lower limits of agreement (LoA) at 6.19 mmHg and -13.99 mmHg, respectively. (b) Bland-Altman plot for DBP showing the mean difference of -4.53 mmHg, with upper and lower LoA at 2.30 mmHg and -11.36 mmHg, respectively.

3.7.4 Hybrid Model: CNN and SVR

In an effort to improve the performance of the initial CNN-based model for blood pressure estimation, we integrated a SVR correction model that leverages patient-specific physiological features such as age, temperature, and heart rate. This hybrid approach was designed to address the limitations observed in the CNN, particularly the increasing error in older age groups. By incorporating these additional features, the SVR correction provides a personalized adjustment to the CNN output, tailoring the predictions to individual patient characteristics and mitigating physiological factors like vascular resistance and skin thickness that change with age.

3.7. Experimental Results and Discussion

Table 3.4: Cumulative Error (Expressed in mmHg)

		≤ 5	≤ 10	≤ 15
BHS Standards	Grade A	60%	85%	95%
	Grade B	50%	75%	90%
	Grade C	40%	65%	85%
	Grade D	Worse than C		
MLP 84 Features	SBP	35%	55%	70%
	DBP	70%	85%	95%
MLP PCA Reduction	SBP	20%	35%	60%
	DBP	25%	50%	70%
CNN L2-norm	SBP	25%	50%	80%
	DBP	35%	80%	95%
CNN & SVR	SBP	90%	95%	100%
	DBP	85%	95%	100%

The results, in terms of MAE and SD, demonstrate significant improvement in both SBP and DBP predictions. In the test set, the model shows strong performance with an MAE of 1.84 mmHg for SBP and 1.98 mmHg for DBP, and SD values of 3.27 mmHg and 3.47 mmHg for SBP and DBP, respectively (Table 3.3). These results represent a clear improvement over the standalone CNN, which exhibited higher MAE and SD values, particularly in the test set (Fig. 3.13).




Figure 3.13: Comparison of prediction errors before and after SVR correction.

These results meet the criteria established by the AAMI, which requires that BP measurement systems achieve a mean error ≤ 5 mmHg and a standard deviation ≤ 8 mmHg when compared to a reference device. Our model's performance is well within these bounds for both SBP and DBP, confirming its compliance with the AAMI standard.

Furthermore, the outcomes align with the grading system proposed by the BHS, which classifies devices based on the percentage of measurements within specified error margins. In our evaluation (see Tab. 3.4), the CNN+SVR model achieved: for SBP, 90% of predictions within 5 mmHg, 95% within 10 mmHg, and 100% within 15 mmHg (Grade A); and for DBP, 85% within 5 mmHg, 95% within 10 mmHg, and 100% within 15 mmHg (Grade B). This confirms both AAMI compliance and a high-grade BHS classification, particularly excelling in systolic blood pressure estimation.

Moreover, the hybrid model effectively reduced the discrepancy in error across age groups. Previously, the CNN’s performance degraded as age increased, but the SVR correction—strengthened by patient-specific features—mitigated this issue, significantly lowering the error for older patients. This is evident from the reduction in both MAE and SD, as well as from the graphical analysis, where the error spread between younger and older patients has notably diminished (Fig. 3.14). These findings suggest that the hybrid model successfully adapts to physiological variations related to aging, enhancing the model’s overall robustness and accuracy across different patient demographics.




Figure 3.14: (a) MAE of SBP for each age group, shown before (blue bars) and after (orange bars) SVR correction. (b) MAE of DBP for each age group, shown before (blue bars) and after (orange bars) SVR correction.

3.7.5 Additional Model Evaluations and Ablation Studies

To further assess the predictive reliability of our model in tracking physiological changes induced by an intervention, we analyzed its performance in individual participants undergoing the Valsalva maneuver. Specifically, we compared the real and predicted SBP and DBP values before and after the maneuver (Tab. 3.5). The model demonstrated consistent tracking of blood pressure variations, with a MAE of 1.8 mmHg (SD = 3.2 mmHg) for SBP and 1.7 mmHg (SD = 2.9 mmHg) for DBP in post-Valsalva predictions. These values indicate a high level of agreement between the predicted and actual measurements, confirming the model’s ability to capture physiological fluctuations. Notably, the predicted SBP and DBP values followed the expected post-Valsalva trends observed in the real data, reflecting the characteristic hemodynamic response to this maneuver. These findings support the model’s potential applicability in real-world scenarios, where continuous and non-invasive blood pressure monitoring could provide valuable insights into cardiovascular function.

In a second phase, in order to assess the contribution of temporal and spectral features to the model’s performance, we conducted an ablation study comparing three versions of the CNN+SVR model:

3.7. Experimental Results and Discussion

Table 3.5: Comparison of Real and Predicted SBP/DBP Values Before and After the Valsalva Maneuver

ID	Before	After	Before Pred.	After Pred.
1	110/60	120/75	111/62	118/73
2	110/70	110/80	109/72	112/78
3	130/70	125/70	126/68	123/73
4	120/80	115/80	119/78	117/76

1. Temporal features only: including systolic time, diastolic duration, inter-beat intervals, and time-domain morphology descriptors;
2. Spectral features only: including dominant frequencies and frequency band power in typical PPG ranges (0.5–4 Hz).
3. Combined features: the full feature set used in our main experiments, merging both domains.

The CNN+SVR model trained with only temporal features achieved a mean absolute error (MAE) of 5.93 mmHg for SBP and 5.01 mmHg for DBP. The spectral-only model performed slightly better on average, with an MAE of 5.12 mmHg for SBP and 4.74 mmHg for DBP. However, the model using the combined feature set outperformed both, achieving a MAE of 1.84 mmHg for SBP and 1.98 mmHg for DBP, confirming that the integration of time- and frequency-domain information captures complementary physiological characteristics and improves overall accuracy. These results support the hypothesis that temporal and spectral descriptors provide non-overlapping insights into vascular dynamics, and their fusion is essential for robust and precise blood pressure estimation from PPG signals.

To further assess the non-redundant value of the multi-wavelength PPG signals, we also developed a simple linear regression baseline model trained exclusively on physiological metadata: age, body temperature, and heart rate. This model served as a minimal reference system, intentionally omitting waveform information. The performance of this baseline approach was significantly poorer than that of all PPG-based methods, with a mean absolute error (MAE) of 5.26 ± 6.43 mmHg for systolic BP and 4.02 ± 4.90 mmHg for diastolic BP. These results confirm that physiological parameters alone are insufficient for accurate blood pressure estimation. In contrast, the incorporation of multi-wavelength PPG waveforms allows the models to exploit richer and more individualized hemodynamic information—such as arterial compliance, blood volume dynamics, and vascular tone—which are not captured by static metadata. The substantial performance gap observed validates that the waveform data provide unique and physiologically relevant signals, reinforcing the design choice to employ multi-channel PPG acquisition in our methodology.

Ultimately, to further substantiate the necessity of multi-wavelength PPG fusion, we conducted an additional experiment using only the red PPG signal, which is commonly adopted in single-channel systems due to its deeper penetration and good signal quality. This model retained the same CNN architecture and training procedure to ensure fair comparison. As illustrated in Fig. 3.15, the single-wavelength approach resulted in substantially higher MAE values for both systolic and diastolic pressure predictions,

across all age groups. This performance degradation highlights the limited physiological scope of using a single wavelength. In contrast, the multi-wavelength PPG signals (red, infrared, green, and blue) capture vascular dynamics at different tissue depths and exploit the unique light absorption characteristics of hemoglobin and surrounding tissues across the spectrum. For example, green and blue light reflect superficial capillary activity, while red and infrared penetrate deeper into arterial layers. By combining these complementary signals, the model accesses a richer and more holistic representation of blood flow and vessel compliance. Moreover, the approach adopted in our CNN architecture—based on the L2-norm of the four wavelength signals—offers a highly effective strategy for feature fusion. Rather than treating each channel independently, the L2-norm compactly integrates the spectral energy of all signals, allowing the model to learn from a unified representation that preserves the physiological contributions of each wavelength. This design choice enhances the model’s ability to generalize, improves robustness to noise or signal dropout, and further strengthens the case for multi-wavelength PPG as a superior input modality. These findings confirm that multi-wavelength fusion, especially when based on a norm-based encoding, is not only beneficial but fundamentally necessary to maximize predictive accuracy and physiological relevance in cuffless blood pressure estimation.




Figure 3.15: Mean absolute error (MAE) for (a) diastolic blood pressure (DBP) and (b) systolic blood pressure (SBP) across different age groups, comparing single-wavelength (red) PPG (blue bars) with multi-wavelength PPG (red bars). The multi-wavelength approach consistently achieves lower MAEs, demonstrating its superior capability to capture relevant physiological information for blood pressure estimation.

3.7.6 Limitations and Future Work

Building on these current findings, several limitations and associated opportunities for future work can be identified. While the proposed framework demonstrates encouraging performance for cuffless blood pressure estimation under controlled conditions, its current validation scope necessarily constrains the interpretation and generalization of the results.

First, although the recruited cohort spans four age brackets, the overall sample size

($N = 81$) and its limited ethnic diversity remain insufficient to guarantee robust generalization across populations characterized by heterogeneous skin tones, vascular properties, and comorbidities. In particular, variations in melanin concentration, tissue optical properties, and peripheral perfusion are known to significantly affect photoplethysmographic signal amplitude and morphology. The under-representation of darker Fitzpatrick skin types may therefore introduce a latent bias, potentially leading to performance disparities when the model is applied to more diverse populations. Expanding the cohort along both demographic and phenotypic dimensions is thus essential to mitigate bias and ensure equitable model behavior.

Second, the blood pressure distribution within the dataset is skewed toward the normotensive range. Extreme values—specifically systolic blood pressure (SBP) above 160 mmHg or below 90 mmHg—are under-represented, limiting the reliability of model predictions in clinically critical conditions such as severe hypertension or hypotension. This imbalance may also bias the regression toward mean values, reducing sensitivity in pathological regimes. Future data collection efforts should therefore prioritize balanced BP stratification to improve model calibration and robustness across the full physiological spectrum.

Third, all recordings were acquired in a controlled laboratory environment with participants at rest. While this setting minimizes confounding factors and facilitates signal synchronization, it does not reflect the variability encountered during real-world use. Free-living conditions introduce motion artifacts, posture changes, temperature fluctuations, and transient physiological perturbations that can substantially alter PPG waveform characteristics. The absence of ambulatory validation represents a significant limitation for wearable deployment, as laboratory performance may overestimate accuracy under unconstrained conditions. Addressing this gap will require data collection during daily activities and the incorporation of motion-aware inference strategies.

Fourth, the experimental protocol relies on a single commercial off-the-shelf (COTS) sensor (MAX86916), and all model training and evaluation were performed offline. This introduces potential device-specific bias related to optical wavelength selection, sensor geometry, analog front-end characteristics, and firmware-level preprocessing. Consequently, cross-device generalization remains unverified, and the reported performance may not directly transfer to alternative hardware platforms. Furthermore, the lack of on-device energy profiling prevents a full assessment of power–accuracy trade-offs, which are critical for long-term wearable operation.

Finally, while the incorporation of SVR-based personalization significantly reduces residual error, the current personalization strategy is limited to a small set of static physiological covariates. Transient factors such as stress, circadian rhythms, hydration status, and medication intake are not explicitly modeled, despite their known influence on cardiovascular dynamics. This simplification may limit adaptability during rapidly changing physiological states and highlights the need for more context-aware or adaptive personalization mechanisms.

From an experimental bias perspective, additional sources of uncertainty must be acknowledged. Reference blood pressure measurements obtained via cuff-based devices are inherently intermittent and not perfectly synchronized with optical signals, introducing temporal misalignment and measurement noise. Moreover, the short recording duration per subject limits the ability to assess intra-subject variability and long-term

stability. These factors collectively emphasize that the reported results should be interpreted as a proof of concept rather than a definitive assessment of clinical performance.

Ongoing and future work aims to address these limitations through several targeted extensions: (i) expansion of the dataset to more than 300 participants with balanced blood pressure classes and Fitzpatrick skin types; (ii) inclusion of synchronized inertial measurements to enable motion-artifact-aware inference; (iii) deployment and energy profiling of a quantized model on an nRF52840 platform with an average power budget below 20 mW; and (iv) exploration of lightweight domain-adversarial training strategies to reduce sensor- and population-specific bias. Together, these developments will strengthen the robustness, fairness, and translational relevance of the proposed approach.

Automatic B-Lines Detection in Lung Ultrasound

4.1 Clinical Background and Relevance of B-Lines

Lung ultrasound (LUS) has progressively evolved into a fundamental imaging modality for bedside assessment of pulmonary conditions in both acute and chronic care settings. Its clinical impact is largely attributed to its intrinsic advantages over traditional thoracic imaging tools. Chest radiography, although widely available, has limited sensitivity in detecting early interstitial syndromes and often requires patient mobilization and ionizing radiation exposure [100]. Computed tomography (CT) is considered the gold standard for lung parenchyma visualization due to its superior spatial resolution; however, it is costly, not portable, time-consuming, and impractical for continuous monitoring. Moreover, CT exposes vulnerable patients to repeated ionizing radiation, which is a non-trivial concern when serial imaging is required for monitoring treatment response or disease progression. In contrast, LUS provides a unique combination of real-time imaging, absence of radiation, repeatability, portability, and immediate integration into clinical workflow, making it particularly suitable for intensive care units (ICUs), emergency departments, pre-hospital environments, and resource-constrained settings [101, 102]. Because LUS can be performed at the point of care, it shortens time-to-diagnosis and enables dynamic decision-making that directly informs immediate therapeutic interventions, such as fluid removal in heart failure or escalation of ventilatory support in ARDS [103].

The clinical utility of LUS extends beyond mere convenience. Ultrasound offers superior temporal resolution compared with CT and chest radiography, allowing clinicians to observe respiratory mechanics and pleural dynamics continuously during the respiratory cycle; this capacity to visualize motion (e.g., lung sliding, real-time changes in B-line pattern during inspiration and expiration) provides functional information that static imaging cannot deliver. This temporal responsiveness is particularly valuable in




Figure 4.1: *Ultrasound waves reflection at air–tissue interfaces (left) and corresponding LUS image (right).*

critical care, where physiological status changes rapidly and repeated imaging would otherwise be impractical [104]. Additionally, the portability and low cost of ultrasound enable frequent reassessment at the bedside, facilitating titration of therapies (for example, incremental diuresis or adjustments to positive end-expiratory pressure) while minimizing the logistical burden and risks associated with patient transport.

A further, often underappreciated advantage of LUS is its sensitivity to subpleural and peripheral lung pathology. Many forms of interstitial involvement, cardiogenic congestion, and early viral pneumonias preferentially affect the lung periphery and subpleural regions—zones that are directly accessible to ultrasound interrogation [105]. Whereas chest radiography can miss subtle peripheral increases in lung density and CT requires transport and radiation, LUS can detect these peripheral changes at the bedside with excellent sensitivity, particularly when structured scanning protocols are used.

LUS doesn’t image the lung parenchyma directly — it interprets patterns of artifacts created by how ultrasound interacts with air, fluid, and tissue (Fig. 4.1). Because these early peripheral alterations are the ones that first produce B-line artifacts, LUS effectively converts minute changes in tissue water content into detectable acoustic signatures, thereby functioning as an early warning system for evolving pulmonary edema or interstitial inflammation.

From a workflow and systems perspective, LUS also enables integrated, multimodal assessment. When combined with basic hemodynamic measurements (blood pressure, central venous pressure, echocardiographic estimates of left ventricular function) and with bedside biomarkers (e.g., natriuretic peptides), LUS-derived information about lung water and aeration can substantially improve diagnostic specificity and therapeutic targeting. For instance, differentiating cardiogenic pulmonary edema from non-cardiogenic causes of B-lines (such as pneumonia or ARDS) often hinges on integrating LUS findings with cardiac function and volume status; this multimodal approach reduces diagnostic uncertainty and guides tailored interventions [104].

Importantly, LUS democratizes diagnostic imaging: because it can be delivered by portable devices and interpreted with focused training, its deployment expands ac-




Figure 4.2: *B-lines formation as vertical artifacts arising from the pleural line.*

cess to high-quality pulmonary assessment in rural and resource-limited environments where CT and even reliable radiography may be unavailable. This democratization has direct public-health implications, enabling earlier detection and management of respiratory disease in populations that otherwise have limited access to advanced imaging services [106].

Despite these strengths, LUS is not without limitations: its interpretive power is largely dependent on operator proficiency, image acquisition angles, and probe selection; clinical integration requires standardized acquisition protocols and quality assurance. These constraints are, however, surmountable through structured training, standardized scanning windows, and algorithmic quality checks embedded in modern ultrasound systems. Crucially, by automating artifact detection and standardizing quantification—particularly of B-lines—ultrasound’s reliability and scalability can be significantly enhanced, reducing inter-operator variability and enabling broader clinical adoption [100].

From a physical perspective, B-lines are vertical, laser-like, hyperechoic reverberation artifacts that arise from subpleural alterations of normal lung aeration (Fig. 4.2). They originate from the pleural line, extend to the bottom of the ultrasound field without fading, erase A-lines, and move synchronously with lung sliding during respiration. Their appearance is linked to the presence of fluid or fibrosis in the interlobular septa, which modifies local acoustic propagation pathways, creating reverberation and resonance phenomena that act as acoustic traps for the ultrasound beam [106].

In a normally aerated lung, the pleural surface behaves as a strong mirror, generating only horizontal reverberations (A-lines). When fluid accumulates in the interstitium or alveolar walls, as in pulmonary edema or inflammation, the increased acoustic coupling allows partial transmission into subpleural structures, which in turn produces the characteristic vertical artifacts. The number, density, and confluence of B-lines are proportional to the degree of extravascular lung water (EVLW) and loss of aeration, and several clinical studies have demonstrated strong correlations between B-line burden and lung density quantified by CT [107].

Clinically, B-lines are recognized as sensitive markers of a broad spectrum of pulmonary disorders. In acute cardiogenic pulmonary edema, their presence is typically bilateral, gravity-dependent, and rapidly responsive to diuretic therapy, making LUS valuable not only for diagnosis but also for treatment monitoring [107]. In acute respiratory distress syndrome (ARDS), B-lines are more heterogeneously distributed and are often associated with pleural line thickening and reduced lung sliding, reflecting diffuse inflammatory injury and alveolar collapse. During the COVID-19 pandemic, B-lines became a pivotal imaging sign due to the peripheral distribution of viral pneumonia, which lies close to the pleural surface and is therefore highly accessible to ultrasound. LUS has also proven effective in interstitial lung diseases, where chronic involvement of the subpleural interstitium generates persistent vertical artifacts that correlate with fibrotic burden and can assist long-term follow-up. Beyond diagnosis, B-lines play a central role in hemodynamic monitoring, particularly in heart failure and critically ill patients, where repeated bedside assessment of fluid status is essential to guide therapy [108, 109].

Despite its advantages, current B-line assessment remains highly operator-dependent. The identification and counting of B-lines require visual interpretation by trained clinicians, and diagnostic accuracy is influenced by probe positioning, breathing phase, insonation angle, and subjectivity in artifact classification. Moreover, B-lines can be confused with other vertical reverberations, such as comet-tail artifacts or Z-lines, which differ in origin and clinical meaning but may appear morphologically similar in suboptimal imaging conditions [107]. This subjectivity introduces inter- and intra-operator variability, limiting reproducibility and scalability, especially in emergency and pre-hospital settings where rapid decision-making is crucial.

For these reasons, there is a growing interest in automatic B-line detection algorithms that can provide objective, reproducible, and real-time quantification. Automatic systems have the potential to democratize LUS use, reducing reliance on expert interpretation, standardizing clinical scoring, and integrating seamlessly into portable ultrasound platforms. Recent research has focused on signal-processing methods to enhance and isolate vertical artifacts from surrounding speckle noise and horizontal reverberations [110]. Among these methods, variational mode decomposition (VMD) has gained attention for its ability to decompose ultrasound frames into structured oscillatory modes, enabling selective enhancement of artifact components and robust detection even in noisy or heterogeneous imaging conditions. This creates a direct bridge between the clinical motivation and the technical developments presented later in the chapter [111].

B-lines are not merely incidental artifacts but physiologically meaningful markers of pulmonary aeration, with a central role in the diagnosis and monitoring of interstitial and edematous lung diseases. LUS uniquely enables bedside, radiation-free, and dynamic assessment, and the clinical community now regards B-line analysis as a core component of pulmonary evaluation. However, to unlock its full potential and achieve reliable large-scale deployment, especially in point-of-care and wearable ultrasound systems, automation is essential. The following sections will detail the acquisition protocol adopted in this work (§4.2), the front-end system developed for ultrasound signal capture (§4.3), and the signal-processing framework based on VMD that underpins the proposed automatic detection method (§4.4–§4.5).

4.2 Lung Ultrasound Acquisition Protocol

A standardized and well-designed acquisition protocol is essential to ensure reliable, reproducible, and clinically interpretable lung ultrasound (LUS) recordings, particularly when the objective is the quantitative assessment of B-lines and pleural artifacts. Since LUS is an operator-dependent modality and highly sensitive to variations in probe orientation, contact pressure, and respiratory dynamics, controlling the acquisition phase has a direct impact on both clinical interpretability and the robustness of downstream computational analysis. In this study, ultrasound data were collected following a structured acquisition protocol intended to balance clinical realism with methodological rigor.

Prior to image acquisition, all subjects were appropriately briefed on the purpose and procedure, and written informed consent was obtained in accordance with ethical standards for non-invasive imaging studies. Only adult participants capable of maintaining a stable respiratory pattern and of complying with basic instructions were enrolled.

Individuals were included if they were hemodynamically stable, showed no signs of respiratory distress that could compromise acquisition quality, and presented either healthy lung aeration or conditions known to produce vertical artifacts of clinical relevance, such as interstitial syndrome or pulmonary congestion. Subjects were excluded if severe chest wall deformities, open thoracic wounds, or extensive surgical dressings prevented adequate probe contact or pleural visualization. Additional exclusion criteria involved conditions known to impair thoracic ultrasound transmission or interpretation, including massive subcutaneous emphysema, extensive pleural calcifications, or the presence of implantable metallic devices in the scanning area. Patients with severe agitation, inability to cooperate, or continuous involuntary movements were also not considered, as motion may alter reverberation geometry and severely compromise pleural-line tracking. When necessary, subjects with acute infectious diseases requiring isolation were postponed to avoid disruption of the acquisition workflow and to maintain standardized environmental conditions. By enforcing these criteria, the resulting dataset reflects a controlled yet clinically meaningful population, suitable for reproducible B-line analysis.

Once eligibility was confirmed, subjects were positioned either in the supine or semi-recumbent posture, depending on clinical condition and tolerance, with the thorax fully exposed from the clavicular line to the costal margins. Whenever feasible, semi-recumbent positioning was preferred, as it enhances lung field accessibility and improves visualization of the pleural interface. Participants were instructed to breathe quietly and to avoid sudden or deep respiratory excursions, which can introduce non-physiological variations in pleural motion and complicate temporal analysis of artifact behavior.

Ultrasound imaging was conducted using either a low-frequency convex probe (approximately 3–5 MHz) or, when higher resolution of the pleural line was required, a high-frequency linear probe (7–12 MHz). The linear transducer was generally favored in this work because it provides superior pleural definition and enhances the visibility of subtle vertical reverberations. Image depth was typically adjusted between 8 and 12 cm so that the pleural line would be displayed in the upper third of the field of view,




Figure 4.3: *Probe positioning for LUS images acquisition: longitudinal plane (A) and transverse plane (B).*

while the focal zone was aligned at or just below the pleura to maximize reflectivity and contrast. Time-gain compensation and global gain settings were carefully adjusted to maintain pleural clarity without artificially amplifying noise or obscuring deeper structures. All filtering options—such as compound imaging, tissue harmonic imaging, or speckle reduction—were disabled, as these tend to suppress or distort artifacts, reducing their diagnostic and computational utility.

A systematic multi-zone thoracic exploration was then performed, covering anterior, lateral, and when accessible, posterior lung regions. In each scanning zone, the probe was positioned longitudinally within an intercostal space to obtain a clear pleural line free from rib shadowing (Fig. 4.3.A). Once the longitudinal scan was completed, a transverse scan was acquired: with the probe kept on the same skin landmark, the transducer was rotated by 90° into a short-axis orientation to perform an orthogonal scan of the target area (Fig. 4.3.B). This two-plane approach (longitudinal and transverse) ensured a comprehensive assessment of pleural morphology, B-line distribution and characteristics while minimizing anisotropy and angle-of-incidence artifacts. Depth, gain, and focal zone were adjusted to center the pleural line in the field of view; gentle fanning and tilting were used to follow the intercostal space while maintaining continuous visualization of the pleura and adjacent lung parenchyma. To ensure adequate temporal resolution for downstream analysis, recordings were acquired at a frame rate between 20 and 60 frames per second. Throughout acquisition, care was taken to minimize probe pressure, as excessive force can dampen pleural motion and alter the geometry of B-line generation. Cine-loops with discontinuous pleural visualization, unstable insonation angle, or excessive motion were immediately reacquired to guarantee uniform signal quality across the dataset.

Upon completion, all recordings were pseudonymized and archived in structured storage with standardized labeling, together with acquisition metadata such as probe type, anatomical window, patient position, and imaging parameters. This consistent acquisition framework produced a high-quality dataset characterized by stable pleural visualization, controlled respiratory motion, and clearly expressed vertical artifacts. Such conditions are essential for reliable computational processing and form the basis for the automatic B-line detection strategies developed in the subsequent sections of this chapter.

4.3 Ultrasound Front-End Acquisition Architecture

The formation of a diagnostic LUS image is the result of a complex chain of analog and digital processes that collectively shape the appearance, contrast, and temporal coherence of pleural artifacts such as B-lines. Since these artifacts arise from subtle acoustic reverberations at the lung–pleura interface rather than from stable anatomical structures, the design of the ultrasound front-end plays a decisive role in determining their visibility and interpretability. Even slight variations in bandwidth, gain control, or dynamic range may alter the morphology of B-lines, making a rigorous understanding of the acquisition architecture essential for both clinical assessment and automated detection algorithms [111].




Figure 4.4: Block diagram of the ultrasound front-end acquisition chain for lung imaging, from transmission and analog conditioning to digital beamforming and optional VMD-based artifact analysis.

The imaging sequence begins with the transmission subsystem, where electrical excitation pulses are delivered to a piezoelectric transducer array. For lung applications, a center frequency of approximately 3.5 MHz typically provides an effective balance between penetration and resolution, allowing the pleural line to be sharply visualized while maintaining sufficient depth of field. The excitation waveform is shaped to produce a broadband ultrasonic pulse with minimal sidelobes, thereby improving axial resolution and facilitating the separation of the dominant pleural echo from the multiple secondary reverberations that give rise to vertical artifacts. Because the lung surface behaves as a strong acoustic reflector due to the impedance mismatch between soft tissue and aerated parenchyma, only a small fraction of the transmitted energy propagates beyond the pleura. The remaining energy contributes to repeated echo paths, forming the reverberation patterns that characterize B-lines. For this reason, the spectral stability of the transmitted pulse is crucial to ensuring consistency in artifact generation [104].

Once the echo is received at the transducer, it enters the analog front-end, where extremely weak signals must be amplified and conditioned before digitization. A low-noise amplifier (LNA) provides the first stage of amplification, preserving signal integrity by minimizing the noise figure and maintaining linearity. Depth-dependent at-

tenuation is then compensated by time-gain control (TGC), which progressively increases amplification with travel time. The shaping of the TGC curve is particularly critical in lung imaging: insufficient gain obscures sub-pleural artifacts, whereas excessive compensation amplifies noise and rib-induced clutter. A band-pass filter removes out-of-band components, ensuring that only frequencies associated with the transmitted pulse are preserved. The choice of filter bandwidth directly affects axial resolution and the textural characteristics of reverberation patterns, and therefore influences the clarity and continuity of B-lines.

Following analog conditioning, the signal is digitized by an analog-to-digital converter (ADC). Adequate bit-depth and sampling frequency are necessary to capture both the strong pleural echo and the weaker reverberation tails without introducing quantization noise that could disrupt the fine vertical structures of interest. The digitized channels are then processed by the beamforming unit, commonly based on delay-and-sum strategies. Although beamforming is designed primarily to optimize lateral resolution and suppress sidelobes, in LUS it indirectly modulates the expression of reverberation artifacts by influencing speckle distribution and pleural line sharpness. If the beam is overly narrow, it may attenuate diffuse reverberation energy, whereas a poorly focused beam reduces pleural contrast and compromises B-line separability.

Once beamformed, the signal undergoes log-compression to map its wide dynamic range into the limited grayscale space of a B-mode image. While this step enhances clinical readability, excessive compression may reduce the contrast between the pleural line and vertical artifacts, complicating automated segmentation. For computational purposes, access to minimally compressed or envelope-detected data is often preferable. Finally, scan conversion transforms the beamformed data from polar coordinates into a Cartesian image, and cine loops are generated at frame rates typically between 20 and 60 fps, allowing the visualization of respiratory motion and pleural sliding, two contextual cues that assist in distinguishing B-lines from motion-related artifacts.

Throughout this chain, each stage of the front-end contributes to shaping the spatio-temporal characteristics of the lung ultrasound signal. Noise sources, dynamic range limitations, and filtering operations do not merely degrade image quality in a generic sense; they selectively alter the spectral structure of reverberation artifacts. This observation is fundamental to the rationale adopted in the following section. Because B-lines emerge from physically constrained reverberation pathways, they occupy identifiable spectral and spatial sub-bands. When the acquisition front-end preserves these signatures, decomposition methods such as Variational Mode Decomposition (VMD) can exploit them to isolate artifact-bearing signal components from surrounding clutter. For this reason, the ultrasound front-end and the subsequent processing algorithm must be regarded not as independent modules, but as stages of a single, coherent imaging system.

4.4 Variational Mode Decomposition (VMD) for Ultrasound Signal Analysis

Variational Mode Decomposition (VMD) is an adaptive signal decomposition technique designed to separate a signal into a predefined number of sub-components, or modes, each characterized by a compact bandwidth in the spectral domain. Its relevance in LUS stems from the fact that B-lines, pleural reflections, and background clutter exhibit distinct spectral and morphological signatures. By decomposing the image or its column-wise signal representation into meaningful oscillatory components, VMD offers a structured approach to isolate artifact-bearing modes from anatomical and noise-related contributions. This is particularly valuable for the quantitative characterization of B-lines, whose early detection and analysis can substantially enhance diagnostic assessment of pulmonary aeration and interstitial syndrome [84].

Traditional decomposition approaches such as Fourier analysis, wavelets, or Empirical Mode Decomposition (EMD) present limitations when applied to LUS signals. Fourier methods lack localization and are not well suited for transient, non-stationary reverberations. Wavelet transforms provide multi-resolution analysis but require the choice of a mother wavelet and often produce redundant representations that complicate downstream interpretation [112]. EMD, while adaptive, suffers from mode mixing and instability, particularly in noisy medical images. VMD overcomes these limitations by formulating decomposition as a well-posed variational problem with explicit constraints, yielding improved robustness and spectral separation [29].

The core principle of VMD is to model a signal $f(t)$ as the sum of K amplitude- and frequency-modulated (AM-FM) modes:

$$f(t) = \sum_{k=1}^K u_k(t) \quad (4.1)$$

where each mode $u_k(t)$ is assumed to be narrow-band around a center frequency ω_k . The algorithm seeks the set of modes that minimize the total bandwidth, enforced by a constrained variational formulation. The optimization problem can be written in compact form as:

$$\min_{\{u_k\}, \{\omega_k\}} \left\{ \left\| \sum_{k=1}^K \left[\partial_t \left(\left(\delta(t) + \frac{j}{\pi t} \right) * u_k(t) \right) e^{-j\omega_k t} \right] \right\|_2^2 \right\} \quad (4.2)$$

subject to the reconstruction constraint:

$$\sum_{k=1}^K u_k(t) = f(t) \quad (4.3)$$

This formulation is solved using the Alternating Direction Method of Multipliers (ADMM), which separates the problem into iterative sub-steps where each mode and its center frequency are updated analytically. Although the mathematical derivation is more involved, its practical implication is straightforward: at convergence, each mode occupies a distinct spectral region, enabling a clean separation of signal components according to their frequency content [113].

For lung ultrasound, this translates into the ability to isolate the low-frequency structure associated with the pleural line from the higher-frequency reverberation streaks that define B-lines. Because B-lines manifest as coherent vertical artifacts with oscillatory patterns along the depth axis, they naturally emerge as one or more narrow-band modes during VMD. By contrast, background clutter and speckle tend to populate broader spectral regions and are therefore distributed across multiple residual modes. This separation is especially effective when VMD is applied column by column or within localized blocks, aligning the decomposition with the vertical geometry of the artifacts [114].

From a practical perspective, VMD offers several advantages for LUS analysis. It enables stable and reproducible mode extraction, is relatively insensitive to noise, and does not suffer from mode mixing. The choice of parameters—primarily the number of modes K and the bandwidth penalty factor α —allows fine control over the decomposition granularity. In practice, modest values of K are sufficient to isolate pleural and artifact components, while convergence criteria ensure computational efficiency. These properties make VMD well suited for integration into automated pipelines for B-line detection and quantification [10].

In summary, VMD provides a physically meaningful and mathematically grounded framework for separating reverberation artifacts from surrounding structures in LUS imaging. Its capacity to decompose signals into spectrally compact components aligns naturally with the physical origins of B-lines, offering a clear advantage over classical transforms or purely statistical filtering. The next section builds on this foundation, describing a block-wise VMD implementation designed to enhance spatial consistency and improve artifact isolation across entire cine-loops.

4.5 Block-Wise VMD Algorithm Design and Implementation

The direct application of Variational Mode Decomposition (VMD) to lung ultrasound (LUS) images on a per-column basis, although conceptually aligned with the vertical geometry of B-lines, often produces unstable decompositions when neighboring columns exhibit local amplitude fluctuations, speckle variations, or transient reflectors. This instability becomes particularly evident near rib shadows, areas of low signal-to-noise ratio, or rapidly moving pleural segments. To mitigate these limitations, a block-wise VMD strategy is adopted, grouping adjacent columns into vertically oriented segments prior to decomposition. This approach enhances spatial coherence, stabilizes mode allocation, and improves the robustness of B-line extraction under realistic acquisition conditions.

B-lines extend from the pleural line toward the bottom of the image as elongated, vertically coherent structures resulting from multi-path reverberations. Their lateral evolution is generally slow and spatially correlated, meaning that adjacent columns share similar oscillatory content. Processing larger lateral regions allows VMD to exploit this redundancy, stabilizing the estimation of center frequencies and reducing the tendency of modes to switch identity between neighboring positions. At the same time, the block must remain sufficiently narrow to avoid merging distinct anatomical or artifact sources. In practice, block widths on the order of a few tens of columns provide an

effective compromise between local precision and lateral consistency.

The LUS frame is divided into vertical blocks of fixed lateral width. The block must be wide enough to contain the full structure of a B-line, yet narrow enough to preserve local detail around the pleura. Typical widths range from 16 to 64 columns, depending on probe resolution and depth settings.

The behavior of VMD within each block depends significantly on four core parameters:

- K – number of modes: controls decomposition granularity; too small merges artifacts and pleural structures, too large introduces redundant modes.
- α – bandwidth penalty: regulates mode compactness in frequency; larger α yields smoother, narrower-band modes, which is useful for reverberation isolation.
- τ – dual ascent step: influences ADMM convergence stability and speed.
- ϵ – convergence tolerance: determines when the iterative updates stop.

In LUS artifact analysis, moderate values of K (typically 3–5) and relatively high α promote stable reconstruction of pleural and reverberation components, while reducing clutter dispersion across modes.

Each block undergoes a structured processing sequence before recomposition into the final frame (Fig. 4.5). The adopted pipeline consists of:

1. **Preprocessing** – normalization of amplitude dynamic range, detrending to remove slow illumination biases, optional envelope smoothing to suppress impulsive noise.
2. **Variational decomposition** – VMD is applied independently to each block, yielding K modes with compact bandwidth and distinct center frequencies.
3. **Mode inspection and allocation** – modes are analyzed and grouped according to their expected physical origin.
4. **Artifact-focused reconstruction** – only selected modes are retained to generate an artifact-enhanced representation.

Block-wise processing also enables parallel execution, since each block constitutes an independent optimization problem. On modern hardware this leads to substantial reductions in latency and makes the approach compatible with real-time LUS pipelines and embedded implementations.

4.5.1 Preprocessing

Before the VMD algorithm can be applied to each block of the lung ultrasound frame, the raw radiofrequency or B-mode data must undergo a careful preprocessing phase designed to homogenize the signal characteristics and suppress unwanted variability. Lung ultrasound signals are inherently affected by speckle, motion artifacts, and attenuation gradients caused by rib shadows and probe angle. Directly decomposing such raw data often leads to poor convergence and unstable modes. Preprocessing therefore




Figure 4.5: VMD-based algorithm workflow: Image pre-processing (1), VMD decomposition (2), Thresholding and image post-processing (3), B-lines quantification and validation (4).

serves to enhance signal-to-noise ratio, stabilize amplitude scales across blocks, and isolate the oscillatory components relevant to reverberation analysis.

The process typically begins with amplitude normalization, which ensures that differences in gain, local illumination, or tissue attenuation do not bias the decomposition toward specific regions. Each block is rescaled to a comparable dynamic range so that the VMD cost function, which minimizes the bandwidth of analytic modes, operates uniformly across the image. Next, a detrending operation removes slow-varying background intensity profiles associated with probe curvature or overall brightness drift, preserving only high-frequency content where pleural and reverberation echoes dominate. In some cases, an envelope extraction or temporal smoothing step is applied to mitigate impulsive spikes while maintaining the continuity of coherent vertical patterns. This preprocessing is deliberately lightweight so as not to distort the temporal structure of the underlying signal, but sufficiently selective to provide VMD with an input that reflects the true morphology of B-lines.

4.5.2 Variational Decomposition

Once the block has been normalized and filtered, the Variational Mode Decomposition is applied. In this phase, the block's intensity profile—typically represented as a set of vertical depth signals—is decomposed into a limited number of narrow-band components, each corresponding to an intrinsic oscillatory mode (Fig. 4.6). The algorithm iteratively estimates both the modes and their central frequencies by minimizing the total bandwidth under the constraint that their sum reconstructs the original signal. In practice, this is achieved through an Alternating Direction Method of Multipliers (ADMM) scheme that alternates between updating the analytic signal envelopes, the mode spectra, and the Lagrange multiplier enforcing reconstruction fidelity.

In the block-wise formulation, this optimization is performed independently for each segment, allowing localized adaptation of center frequencies without compromising global consistency. The parameter α (bandwidth penalty) controls how tightly each mode is constrained in the frequency domain, while K (number of modes) determines the granularity of the decomposition. A moderate α encourages smooth spectral separation, avoiding mode overlap between pleural and reverberation components. Because each block contains several adjacent columns rather than a single one, the variance of the spectral estimates is reduced, yielding more stable convergence and mitigating the mode-switching phenomena that typically occur in per-column decomposition. This results in a decomposition that more faithfully represents the physical origins of the echoes—pleural reflections, subpleural scattering, and reverberant streaks—each confined to distinct spectral bands.

4.5.3 Mode Analysis and Allocation

After decomposition, the block yields several modes, each carrying information about different physical and acoustic phenomena. The next task is to identify which modes correspond to meaningful features—particularly those representing B-lines—and which are dominated by background clutter or random noise. The analysis is guided by both spectral and spatial criteria. In the frequency domain, B-line modes exhibit a relatively narrow bandwidth centered at higher frequencies than the pleural reflection but lower than pure electronic noise. Spatially, these modes display strong vertical continuity




Figure 4.6: *Block-wise VMD Decomposition.*

across depth, reflecting the coherent reverberation of ultrasound waves between pleura and probe.

The allocation process involves examining the energy distribution, phase coherence, and vertical autocorrelation of each mode within the block. Modes characterized by consistent oscillatory behavior over depth and limited lateral variability are retained as B-line candidates, while broad or irregular components are classified as clutter or background. The pleural line is often represented by a dominant, low-frequency mode

with high intensity near the top of the image; it provides an anatomical reference that can be used to align the interpretation of artifact modes. The outcome of this stage is a subset of modes that collectively reconstruct a version of the image where B-lines are selectively preserved and other structures are attenuated. This selection process is crucial: it translates the abstract spectral decomposition into a physically interpretable representation tied directly to the acoustic mechanisms underlying LUS imaging.

4.5.4 Reconstruction and Integration

Once the relevant modes have been selected, the enhanced image is reconstructed by summing these components within each block and merging the resulting segments across the full frame. Reconstruction is not merely a linear recombination but an integration step that must preserve both amplitude continuity and spatial coherence. Discontinuities at block borders are mitigated using overlapping regions and smooth lateral blending functions that weight adjacent blocks according to proximity, thus preventing visible seams or step artifacts. The merging operation ensures that the B-line structures remain continuous across block boundaries and that the pleural line maintains its geometric integrity [115].

At this stage, the reconstructed frame presents a distinctly improved contrast between B-lines and surrounding tissue. The reverberation artifacts appear sharper, vertically aligned, and more uniform, while clutter, noise, and lateral shadows are substantially reduced. Because each block was decomposed locally, the representation adapts naturally to spatial variations in acoustic coupling and lung morphology, providing a consistent depiction of pleural artifacts across the entire image. The output of this reconstruction stage serves as a robust input for the validation and quantitative analysis procedures discussed in the following section, where the effectiveness of the block-wise VMD algorithm is evaluated on clinical and simulated datasets.

4.6 Experimental Validation, Results and Discussion

The performance of the proposed block-wise VMD algorithm was evaluated through a retrospective feasibility study designed to assess both its segmentation accuracy and computational efficiency in realistic clinical settings. The validation process was conducted on a curated dataset of 116 LUS frames acquired from 42 patients presenting clinical signs of pulmonary disease. Each frame underwent expert annotation by experienced pulmonologists, providing a reliable reference standard for quantitative evaluation. The study followed Good Clinical Practice guidelines and was approved by the Ethics Committee of the University of Bari Medical School (Protocol ID no. 7010/2021).

All LUS scans were collected using a SensUS Lung portable echographic device equipped with a 3.5 MHz convex probe, operated at a depth of 9 cm and with the focus precisely aligned with the pleural line. The acquisition setup ensured sufficient image quality for artifact analysis while minimizing inter-frame motion, as described in Section 4.2). Frames exhibiting motion blur, gain saturation, or poor pleural visibility were excluded prior to analysis, following predefined quality criteria. The final dataset consisted exclusively of frames with clear pleural delineation and well-visible vertical artifacts, enabling consistent segmentation performance assessment.

4.6.1 Quantitative Evaluation Protocol

The proposed algorithm was tested on each image using its fully configured block-wise pipeline, as described in Section 4.5. Decomposition was carried out with $K = 4$ modes and a bandwidth control parameter $\alpha = 2000$, a configuration determined through empirical optimization. The pre-processing phase applied a 3×3 median filter and contrast-limited adaptive histogram equalization (CLAHE) with a clip limit of 2.0 to improve visibility of vertical features while preserving natural texture. Each frame was divided into blocks of height $h = 32$ pixels and full image width (224 pixels). For each block, vertical intensity profiles were extracted and decomposed through VMD, and the mode carrying the highest spectral energy was selected as the most representative of B-line structures. A local adaptive threshold $T = \mu + \lambda\sigma$ was then computed (with $\lambda = 1.5$) to produce a binarized segmentation map adaptive to block-specific brightness and contrast conditions.

After block-level segmentation, a post-processing phase recombined all binary masks and applied vertical morphological closing to restore the continuity of B-lines across depth and remove isolated noise regions. The result was a full-resolution artifact-enhanced segmentation map from which B-lines were quantified by identifying contiguous vertical structures extending from the pleural line. The automatic counts were compared against expert annotations, and performance was quantified through precision, recall, F1-score, and Dice Similarity Coefficient (DSC).

4.6.2 Results and Quantitative Performance

Across the dataset, the algorithm identified a total of 156 B-lines, achieving a precision of 0.87, recall of 0.79, F1-score of 0.83, and a Dice coefficient of 0.79 when compared to the manual ground truth. These results confirm that 87% of the automatically detected structures correspond to true B-lines, while approximately four out of five true B-lines were successfully captured. The balance between precision and recall indicates that the proposed method effectively minimizes false positives without sacrificing sensitivity. The obtained Dice coefficient further demonstrates strong spatial agreement between automated and manual segmentations, validating the capability of the VMD-based approach to delineate artifact morphology with high fidelity.

Notably, the introduction of block-wise processing substantially improved segmentation consistency relative to prior column-wise VMD implementations. By aggregating multiple adjacent columns, the algorithm achieved greater spectral stability, reducing mode mixing and enhancing robustness to speckle noise and motion-induced variability. This improvement was particularly evident in frames with irregular artifact contrast or non-uniform illumination, where the local adaptive threshold successfully preserved B-lines even under challenging acoustic conditions.

4.6.3 Comparative Performance Analysis

Table 4.1 summarizes the comparative performance of the proposed method against representative approaches from the literature, including Radon-transform and wavelet-based detectors, as well as deep learning architectures such as U-Net and YOLOv8-PBB. Although the VMD method achieved slightly lower recall values than GPU-accelerated neural networks, its F1-score (0.83) and Dice coefficient (0.79) remain

4.6. Experimental Validation, Results and Discussion

comparable to deep models while requiring orders of magnitude less computational power and no training data.

Table 4.1: *Comparative Performance of B-Line Detection Algorithms*

Method	Type	Precision	Recall	F1-score	Dice	Time / Frame	Notes
Proposed Block-wise VMD	Model-based (signal decomposition)	0.87	0.79	0.83	0.79	~100 ms (CPU)	No training required; interpretable; suitable for embedded systems
Radon Transform [116]	Classical feature extraction	0.75	0.68	0.71	0.69	~210 ms (CPU)	Sensitive to contrast and orientation noise
Wavelet-Based Detector [117]	Multiresolution analysis	0.81	0.72	0.76	0.73	~180 ms (CPU)	Robust to speckle; limited selectivity for weak artifacts
U-Net [118]	Deep learning (supervised)	0.89	0.84	0.86	0.82	~45 ms (GPU)	Requires large annotated dataset; high computational load
YOLOv8-PBB [119]	Deep learning (object detection)	0.90	0.87	0.88	0.84	~35 ms (GPU)	Fast; accurate; black-box model, not interpretable

Note: Performance comparison among state-of-the-art methods for B-line detection in lung ultrasound (LUS) imaging. The proposed block-wise VMD method achieves accuracy comparable to deep learning approaches while remaining interpretable and hardware-efficient.

The computational evaluation confirmed that the entire VMD pipeline required approximately 100 milliseconds per frame on a standard laptop equipped with an Intel i5-1335U CPU and 16 GB of RAM, corresponding to near-real-time performance. This time includes all stages of pre-processing, decomposition, and post-processing. Given the algorithm’s inherent parallel structure—wherein each block is processed independently—the total processing time can be further reduced using embedded GPU or FPGA acceleration. In contrast, U-Net and YOLOv8-PBB networks typically achieve 30–60 milliseconds per frame but require dedicated high-power GPUs and extensive pretraining on thousands of annotated images. Consequently, the VMD-based approach presents an optimal trade-off between speed, interpretability, and deployability in resource-constrained environments such as portable or point-of-care ultrasound systems.

The results demonstrate that VMD-based decomposition can serve as an interpretable and computationally efficient alternative to data-driven deep learning methods for automatic artifact segmentation. The physics-driven nature of VMD allows a direct mapping between spectral modes and physical echo sources, offering insights into the origin and morphology of B-lines that black-box models cannot provide. Moreover, since the algorithm does not require annotated training data, it is inherently generalizable to new ultrasound systems and patient populations without retraining.

Nevertheless, several limitations are recognized. The number of modes K must be predefined, which can influence the decomposition outcome; when spectral components overlap, VMD may merge them, leading to partial artifact suppression. Additionally, this study was conducted on static frames rather than dynamic cine-loops, limiting its evaluation of temporal coherence. Future work aims to extend the implementation to real-time video processing and to investigate adaptive parameter tuning and hybrid schemes combining VMD-derived features with lightweight neural networks. Such extensions could further enhance segmentation robustness and generalize the method to diverse imaging conditions.

4.7 Discussion and Conclusions

The experimental results presented in the previous section demonstrate that Variational Mode Decomposition (VMD), when properly adapted to the characteristics of lung ultrasound (LUS) imaging, provides a powerful, explainable, and computationally efficient method for the identification and segmentation of B-lines. Unlike purely data-driven models, the VMD approach is rooted in the physical properties of ultrasound wave propagation and reverberation. This physical interpretability represents one of its most significant strengths: each decomposed mode corresponds to a distinct acoustic phenomenon — from pleural reflections to multi-path reverberations — enabling a direct link between the algorithmic output and the underlying physiology of the lung.

From a clinical perspective, the detection and quantification of B-lines play a crucial role in the evaluation of pulmonary aeration and interstitial syndromes. The accurate characterization of these vertical artifacts supports early diagnosis of conditions such as pulmonary edema, pneumonia, and acute respiratory distress syndrome (ARDS) [120]. The proposed block-wise VMD approach contributes to this objective by offering a reliable tool that can operate independently of extensive training datasets and adapt dynamically to diverse imaging conditions. Its robustness to variations in gain, probe positioning, and acoustic coupling makes it particularly suitable for bedside examinations, where operator-dependent variability often limits reproducibility.

Comparative analysis against state-of-the-art methods highlights the delicate balance achieved by the VMD-based algorithm. Deep learning architectures such as U-Net and YOLOv8-PBB currently set the benchmark for segmentation accuracy, benefiting from large-scale annotated datasets and optimized inference pipelines. However, these models act as black boxes, offering limited insight into the physical origin of the detected features and often requiring dedicated hardware. The proposed VMD framework, in contrast, achieves comparable Dice coefficients (0.79 vs. 0.82–0.84) while being inherently interpretable and executable on standard CPU hardware. Its modest computational footprint (approximately 100 ms per frame) supports real-time or near-real-time deployment even on portable ultrasound systems, thus aligning with the growing demand for edge computing and hardware-constrained medical devices discussed earlier in this thesis.

From a signal processing standpoint, the adoption of a block-wise decomposition strategy was instrumental in improving both stability and segmentation quality. By aggregating adjacent columns into spatially coherent units, the algorithm mitigates the mode-switching effects typical of per-column VMD and enhances the lateral consistency of the reconstructed artifacts. This structural coherence translates directly into clinical relevance: continuous, vertically aligned B-lines are more easily interpretable by clinicians, and their quantitative measurement becomes more reliable. Moreover, the block-wise formulation naturally supports parallel computation, making the approach scalable and compatible with modern multi-core and embedded architectures.

Despite these advantages, several limitations remain. The need to predefine the number of modes K and the bandwidth parameter α introduces a degree of manual tuning, which may limit generalization across different ultrasound machines or patient populations. Furthermore, the current evaluation focused on static frames rather than full cine-loops. Temporal consistency, which is critical for dynamic phenomena such

as pleural sliding and artifact evolution during respiration, was not explicitly modeled in this study. Extending the VMD framework to spatiotemporal decomposition, or coupling it with motion-compensated analysis, represents a promising direction for future research.

Another point of discussion concerns the integration of model-based and data-driven paradigms. While VMD provides a transparent and mathematically grounded representation of ultrasound signals, neural networks excel in learning complex contextual dependencies. Hybrid strategies that combine the interpretability of VMD-derived features with the adaptability of deep learning could therefore yield the best of both worlds. For instance, VMD modes can serve as pre-filtered or frequency-enhanced inputs to lightweight neural detectors, improving their convergence and reducing the need for extensive data augmentation. This line of development would preserve the explainability and hardware efficiency of the current method while extending its scope to more complex diagnostic tasks, including pleural motion assessment and multi-zone artifact analysis.

The proposed block-wise VMD approach demonstrates that explainable, physics-driven signal processing remains highly relevant in the age of deep learning. By capturing the essential acoustic behavior of the lung within a computationally efficient and interpretable framework, it paves the way for the integration of advanced ultrasound analytics into portable, low-power, and edge-capable medical platforms. The outcomes of this work strengthen the argument for hybrid e-health technologies that combine embedded hardware, optimized algorithms, and clinically meaningful interpretability — a theme that connects the two main research lines of this thesis.

4.7.1 Concluding Remarks

The investigation presented in this chapter demonstrates the feasibility and effectiveness of physics-based signal decomposition for the automatic detection of lung ultrasound artifacts. By leveraging the mathematical rigor of Variational Mode Decomposition and adapting it to the unique structure of B-lines, the proposed framework achieves accurate, interpretable, and computationally efficient segmentation under realistic clinical conditions. The block-wise architecture, in particular, has proven instrumental in stabilizing the decomposition process and in enhancing the spatial coherence of artifact reconstruction, leading to performance levels comparable to those of deep learning models while retaining full algorithmic transparency.

From a broader perspective, this work illustrates how model-based signal processing can coexist and even complement modern data-driven strategies. The interpretability and resource efficiency of VMD make it particularly appealing for embedded and point-of-care medical systems, where real-time execution and reliability are critical. At the same time, the potential integration of VMD-derived representations into hybrid frameworks with neural architectures offers a promising pathway toward even more adaptive and robust diagnostic tools.

The outcomes of this study therefore mark a convergence between theoretical signal modeling, hardware-aware algorithm design, and clinically oriented application. Together, these contributions reinforce the overarching vision of this doctoral research: to develop embedded, explainable, and computationally optimized methods for biomedical signal monitoring, spanning from cardiovascular assessment to pulmonary ultra-

Chapter 4. Automatic B-Lines Detection in Lung Ultrasound

sound analysis. The concluding chapter will reflect on this trajectory, summarizing the main contributions of the work and outlining future directions for research and clinical translation.

CHAPTER 5

Conclusions and Future Work

5.1 Summary of Main Contributions

This doctoral research has explored the design, implementation, and validation of embedded systems and signal processing algorithms for the acquisition, analysis, and interpretation of physiological signals. The central goal of the work has been to develop hardware-aware, interpretable, and computationally efficient solutions capable of supporting continuous and non-invasive health monitoring in resource-constrained environments. The thesis thus integrates two complementary lines of research: (i) the development of embedded architectures and hybrid machine-learning models for cardiovascular monitoring through photoplethysmography (PPG), and (ii) the design of advanced signal decomposition techniques for lung ultrasound (LUS) artifact detection. Together, these investigations converge toward a unified vision of intelligent e-health systems that combine physical interpretability, algorithmic robustness, and clinical applicability.

At the foundation of this work lies the recognition that the ongoing digital transformation of healthcare—often referred to as connected or smart health—requires not only data-driven intelligence but also hardware-level innovation. Continuous physiological monitoring, particularly outside hospital environments, poses challenges of miniaturization, energy efficiency, and real-time processing that traditional computing paradigms cannot meet. Embedded microcontrollers, low-power digital signal processors (DSPs), and dedicated front-end circuits must perform sophisticated computations while operating on minimal energy budgets. In this context, the thesis contributes both to the engineering of embedded hardware platforms for physiological signal acquisition and to the algorithmic optimization of data analysis methods for predictive and diagnostic purposes.

The first major contribution of this thesis lies in the implementation and experi-

mental validation of an embedded framework for real-time acquisition and monitoring of physiological signals using a multi-wavelength optical sensing platform. The system was based on a commercially available evaluation platform—integrating state-of-the-art components for biosignal acquisition and embedded processing—which was customized, calibrated, and optimized for continuous cardiovascular monitoring applications. The adopted hardware platform, centered on the MAX86916 Evaluation System by Analog Devices, integrates a low-noise analog front-end, multi-channel optical emitters, and a high-sensitivity photodiode, enabling the acquisition of photoplethysmographic (PPG) signals at four distinct wavelengths (red, green, blue, and infrared). The accompanying MAX32630FTHR microcontroller board, equipped with an ARM Cortex-M4F core, provided real-time digital control, signal processing, and data streaming capabilities. This setup allowed the design and validation of a complete acquisition pipeline—from optical signal capture and analog conditioning to digital preprocessing, feature extraction, and wireless data transmission—while maintaining strict constraints in terms of power efficiency and temporal synchronization. The system enabled reliable monitoring of cardiovascular biomarkers such as HR, PTT and oxygen saturation. Moreover, it provided a robust foundation for exploring advanced, machine-learning-based methods for blood pressure estimation, highlighting the potential of embedded, multi-modal acquisition systems in wearable healthcare applications.

The second major contribution centers on non-invasive blood pressure (BP) estimation using multi-wavelength photoplethysmography and a hybrid machine-learning approach that combines convolutional neural networks (CNNs) with support vector regression (SVR). Traditional BP measurement techniques rely on cuff-based devices that are unsuitable for long-term monitoring and that interrupt the measurement process. In contrast, PPG offers a low-cost, continuous, and non-invasive alternative, but the relationship between optical pulse morphology and blood pressure is nonlinear and highly individual. To address this challenge, this thesis proposed a hybrid CNN–SVR model that exploits the feature-learning capacity of deep neural networks while preserving the interpretability and generalization properties of kernel-based regression. The architecture was trained and validated on a purpose-built dataset of synchronized PPG signals acquired across multiple wavelengths, providing richer information on vascular compliance and blood volume variations. Extensive experimentation demonstrated that the hybrid approach achieved superior accuracy compared to standalone machine learning models, reducing mean absolute error and improving correlation with reference arterial pressure. Moreover, the implementation of this algorithm on embedded hardware confirmed its suitability for real-time execution, paving the way for wearable cuff-less BP monitoring systems.

The third contribution of the thesis lies in the domain of lung ultrasound imaging, where a novel method was developed for the automatic detection and enhancement of B-lines—vertical artifacts that serve as key diagnostic indicators of pulmonary congestion and interstitial syndromes. The proposed approach employs Variational Mode Decomposition (VMD) as a physics-driven signal processing technique to separate pleural reflections, reverberations, and clutter components. Unlike data-driven convolutional networks, the VMD framework decomposes the ultrasound signal into a set of narrow-band modes, each corresponding to specific acoustic phenomena. The introduction of a block-wise decomposition strategy represented a major methodological advance: by

grouping adjacent image columns, the algorithm stabilizes the mode estimation process, mitigates noise sensitivity, and preserves the vertical coherence of B-lines. This methodological innovation significantly improved segmentation performance, achieving a Dice coefficient of 0.79 and an F1-score of 0.83, results that are competitive with deep learning architectures such as U-Net and YOLOv8 while maintaining full interpretability and CPU-level efficiency.

From a clinical standpoint, the thesis contributes to the broader field of explainable medical artificial intelligence (AI). Both the CNN-SVR model for blood pressure estimation and the VMD-based framework for lung ultrasound share a common design philosophy: they combine data-driven adaptability with model-based interpretability. This hybrid approach offers a promising paradigm for the next generation of biomedical analytics, where algorithms must not only be accurate but also transparent and computationally sustainable. The embedded nature of both systems ensures scalability and portability, addressing one of the main limitations of conventional AI pipelines in healthcare—their dependence on centralized, high-power computing infrastructure. The research presented here therefore moves toward an edge-intelligent medical system capable of performing diagnostic tasks directly on the sensing platform.

In addition to its methodological and technological advances, the thesis contributes to the integration of hardware and algorithm co-design principles in biomedical engineering. The optimization of hardware parameters such as sampling rate, analog-to-digital conversion resolution, and signal conditioning bandwidth was guided by the statistical and spectral requirements of the algorithms themselves. Conversely, the algorithmic complexity was shaped by the constraints of embedded processing, ensuring that each computational step could be executed in real time with minimal energy consumption. This bidirectional design methodology exemplifies the emerging discipline of embedded intelligence, in which sensing, computation, and interpretation form a unified design space.

Finally, the overall contribution of this doctoral work lies not only in the individual systems developed but also in the unified vision they embody. Both research lines—cuffless cardiovascular monitoring and lung ultrasound artifact analysis—demonstrate how embedded signal processing can bridge the gap between sensor-level hardware and clinical decision support. By leveraging efficient acquisition architectures, hybrid modeling techniques, and physically interpretable algorithms, this research provides a concrete step toward the realization of self-contained, explainable, and energy-aware e-health platforms. Such systems represent a key enabling technology for preventive medicine, telemonitoring, and personalized healthcare—domains in which the continuous, unobtrusive, and intelligent observation of physiological signals can profoundly impact patient outcomes and healthcare efficiency.

5.2 Limitations of the Current Work

While the research presented in this thesis has yielded promising results and contributed novel methodologies for embedded physiological monitoring and biomedical signal processing, several limitations must be acknowledged. These limitations arise from both experimental constraints and methodological assumptions, and they help delineate the boundaries of the current work while motivating future investigations. Recogniz-

ing them is essential to situate the proposed systems within their realistic operational context and to guide their future refinement.

5.2.1 Experimental and Dataset Constraints

One of the primary limitations concerns the size, diversity, and standardization of the datasets used for model development and validation. The experimental data employed in both research lines—multi-wavelength photoplethysmography (PPG) for blood pressure estimation and lung ultrasound (LUS) imaging for B-line detection—were collected under controlled conditions, with limited subject variability and relatively homogeneous acquisition setups. Although this allowed for precise control of signal quality and synchronization, it also restricts the generalizability of the results to broader, more heterogeneous populations. In the case of PPG-based blood pressure estimation, factors such as skin pigmentation, sensor placement, peripheral perfusion, and motion artifacts can significantly affect waveform morphology, potentially introducing biases when deploying the system in real-world, ambulatory environments. Similarly, the lung ultrasound dataset, while clinically validated and annotated, included a limited number of subjects and static frames, which may not fully capture the variability of pathological presentations or the temporal dynamics observed in cine-loops.

Moreover, the reference measurements used for ground-truth validation—arterial blood pressure readings for PPG and expert annotations for LUS—introduce inherent sources of uncertainty. Cuff-based sphygmomanometers, although widely adopted, are not perfectly synchronized with optical signals, while human annotation of B-lines is subject to inter-observer variability, especially in borderline or low-contrast cases. These limitations highlight the need for larger, multi-center datasets and the inclusion of temporal sequences, which would better support both algorithmic generalization and clinical translation.

In addition, the controlled nature of the acquisition protocols inevitably limits the exposure of the proposed algorithms to rare or extreme conditions, such as poor peripheral perfusion, atypical anatomical features, or severe signal degradation. While these scenarios may be underrepresented in curated datasets, they are not uncommon in real-world clinical and home-monitoring contexts. As a result, the current evaluation may overestimate performance stability when compared to unconstrained deployments.

Furthermore, the limited availability of publicly accessible datasets combining synchronized physiological signals and clinically validated reference measurements constrains reproducibility and cross-study comparison. This highlights a broader challenge in the field of embedded biomedical monitoring, where data scarcity and heterogeneity remain major barriers to benchmarking and standardization. Addressing these issues will require coordinated efforts toward open data initiatives, shared annotation protocols, and standardized acquisition guidelines.

5.2.2 Hardware and Implementation Constraints

From an engineering perspective, the systems developed in this thesis operate under stringent hardware constraints, reflecting the intended deployment in wearable or portable medical devices. While these constraints motivate algorithmic efficiency, they also limit the complexity and adaptability of the implemented models. For instance,

the hybrid CNN–SVR model for blood pressure estimation was optimized for inference speed and low memory footprint, but its architecture remains relatively simple compared to large-scale deep learning networks. The reduced number of convolutional layers and hidden neurons, although adequate for embedded platforms, might restrict the capacity to capture subtle morphological variations in the PPG signal, particularly under non-ideal acquisition conditions.

Similarly, the block-wise Variational Mode Decomposition (VMD) algorithm, though effective and computationally lightweight, relies on static parameter settings such as the number of modes (K) and the bandwidth penalty factor (α). These parameters were tuned empirically for the available dataset and may not represent the optimal configuration for other ultrasound devices, probe frequencies, or patient anatomies. Furthermore, the processing pipeline assumes a relatively high signal-to-noise ratio at the input, limiting its robustness in extremely degraded or low-intensity acquisitions. While the implementation was validated on standard desktop hardware with CPU-based processing, further optimization and real-time performance testing on embedded microcontrollers, field-programmable gate arrays (FPGAs), or system-on-chip (SoC) architectures remain necessary to confirm feasibility in clinical-grade portable devices.

Another relevant limitation lies in the trade-off between hardware efficiency and algorithmic expressiveness. While the reduced architectural complexity of the proposed models is well suited to embedded platforms, it inherently limits their ability to adapt to highly non-stationary signal conditions. This design choice reflects a deliberate prioritization of determinism, interpretability, and low-power operation over raw representational capacity. However, such trade-offs may become increasingly restrictive as deployment scenarios grow more complex.

Additionally, the absence of a full hardware–software co-design cycle in the present work means that some implementation choices were guided by theoretical feasibility rather than empirical optimization on target devices. Factors such as memory access patterns, bus contention, and peripheral I/O latency—which can significantly affect real-time performance on embedded systems—were not exhaustively characterized. Future implementation-level profiling will therefore be necessary to validate worst-case execution times and ensure compliance with real-time constraints in safety-critical applications.

5.2.3 Methodological Assumptions

Both methodological frameworks developed in this thesis—the CNN–SVR hybrid model for PPG and the block-wise VMD algorithm for LUS—are grounded on simplifying assumptions that, while enabling tractable analysis, may overlook certain complexities of physiological and acoustic phenomena. The blood pressure estimation model, for example, assumes a relatively stable relationship between PPG morphology and arterial pressure, neglecting dynamic physiological factors such as autonomic nervous system variability, vascular tone modulation, and hydrostatic effects due to posture or temperature. These influences can alter pulse waveform characteristics independently of blood pressure, introducing potential confounders when the model is applied across different physiological states or measurement contexts.

In the case of lung ultrasound, the VMD-based approach assumes that B-lines exhibit distinct spectral and spatial signatures separable through narrow-band decompo-

sition. However, in practice, the acoustic behavior of the lung is highly non-linear, and artifacts may overlap spectrally with tissue reflections, leading to partial mode mixing or misclassification. The algorithm also operates in a two-dimensional (spatial) domain, whereas real ultrasound propagation is inherently three-dimensional and affected by probe orientation, patient anatomy, and transducer focus depth. The absence of temporal modeling further limits the system's ability to capture the dynamic evolution of artifacts, such as the appearance and disappearance of B-lines during respiration.

More broadly, the methodological frameworks adopted in this thesis emphasize interpretability and analytical tractability, sometimes at the expense of modeling completeness. While this approach facilitates understanding and debugging—an essential requirement in medical contexts—it may underrepresent complex non-linear interactions that emerge from multi-physiological coupling or patient-specific adaptations.

In particular, the absence of explicit temporal modeling in both application domains limits the ability of the proposed systems to capture long-term dependencies, trends, or transient phenomena. This may affect performance in scenarios where temporal context is diagnostically relevant, such as gradual blood pressure drift or respiration-dependent variations in lung artifacts. Incorporating temporal dynamics would require more sophisticated models and memory mechanisms, introducing additional computational and validation challenges.

5.2.4 Clinical and Translational Limitations

From a clinical standpoint, the systems developed in this work are still in the prototype validation stage. Although experimental results demonstrate competitive performance compared to state-of-the-art methods, extensive clinical trials are required to establish diagnostic equivalence and safety for real-world applications. For blood pressure monitoring, validation according to international standards such as ISO 81060-2:2018 and the AAMI/ESH/ISO Universal Standard has not yet been completed, and long-term stability across varying environmental conditions remains to be demonstrated. For lung ultrasound, while the algorithm shows strong agreement with expert annotations, real-time usability and integration into clinical workflows have not yet been assessed. Moreover, the lack of temporal data means that dynamic features—such as pleural sliding or respiratory phase dependence—are not yet accounted for, potentially limiting the interpretability of results during continuous monitoring.

Finally, the thesis primarily focuses on algorithmic and hardware innovation rather than large-scale software integration or user-interface design. For practical deployment in clinical settings, these systems would need to be embedded into user-friendly platforms that provide visualization, automated reporting, and interoperability with hospital information systems. Such aspects, although beyond the scope of this doctoral research, represent critical steps toward translation from laboratory prototypes to certified medical devices.

From a translational perspective, an additional limitation concerns the gap between algorithm-level performance metrics and clinically meaningful outcomes. While accuracy, sensitivity, and agreement with reference standards provide necessary indicators of technical validity, they do not fully capture the impact of automated systems on clinical decision-making, workflow efficiency, or patient outcomes. These dimensions require prospective studies and human-in-the-loop evaluations, which were beyond the

scope of this doctoral research.

Moreover, the absence of a dedicated user-interface and decision-support layer limits the current systems' accessibility to non-technical users. In real clinical environments, factors such as explainability of outputs, ease of interaction, and integration with existing diagnostic routines play a decisive role in adoption. Addressing these aspects will necessitate user-centered design approaches and close collaboration with clinicians, further extending the scope of future work.

The current work provides a robust proof of concept for embedded, interpretable biomedical signal processing but also reveals key challenges that must be addressed before full clinical adoption. These include the expansion and diversification of datasets, adaptive parameterization for variable hardware configurations, real-time validation on embedded targets, and comprehensive clinical evaluation. Each of these limitations, however, represents an opportunity for growth and innovation, forming the foundation for the future research directions discussed in the following section.

5.3 Potential Extensions and Future Directions

The outcomes of this doctoral research not only provide a solid foundation for the development of embedded biomedical monitoring systems but also open multiple directions for future exploration. The integration of efficient hardware architectures, interpretable algorithms, and clinically relevant applications establishes a multidisciplinary platform upon which further innovations can be built. In particular, three complementary trajectories emerge as most promising: (i) the real-time deployment of the proposed algorithms on embedded systems, (ii) their progressive hybridization with advanced data-driven and adaptive learning paradigms, and (iii) the large-scale clinical validation and translational assessment required for widespread adoption. Each of these directions involves technical, methodological, and regulatory challenges whose systematic investigation can substantially extend the scientific and practical impact of this work within the broader domain of e-health technologies.

Beyond the specific applications addressed in this thesis, these future research lines also aim to strengthen the general methodological framework proposed herein, promoting a design philosophy in which sensing, signal processing, machine learning, and hardware constraints are considered jointly rather than in isolation. Such an approach is increasingly necessary to ensure that emerging biomedical technologies can transition from proof-of-concept implementations to robust, scalable, and clinically accepted medical systems.

5.3.1 Embedded Deployment and Optimization

The first and most immediate extension concerns the full integration of the proposed frameworks into embedded and portable hardware platforms. Although both the hybrid CNN-SVR model for non-invasive blood pressure estimation and the block-wise Variational Mode Decomposition (VMD) algorithm for lung ultrasound signal analysis were conceived with computational efficiency in mind, their validation was primarily carried out in software environments running on general-purpose processors. A natural progression of this work therefore involves the systematic translation of these algorithms

into fully operational embedded systems, requiring careful optimization of computational complexity, memory footprint, latency, and energy consumption.

For cardiovascular monitoring, future research should focus on the deployment of the hybrid CNN–SVR architecture on low-power System-on-Chip (SoC) platforms that combine ARM Cortex-M or Cortex-A microcontrollers with embedded digital signal processing (DSP) units or neural-network accelerators, such as NPUs or EdgeTPU-like cores. Model compression techniques—including structured pruning of convolutional filters, quantization of weights and activations to 8-bit or sub-8-bit precision, and replacement of floating-point operations with fixed-point arithmetic—should be investigated to reduce inference cost while preserving compliance with clinical accuracy requirements. In parallel, memory-aware network design and on-chip buffering strategies could minimize data movement, which is a dominant contributor to energy consumption in embedded inference.

At the signal-acquisition level, adaptive and context-aware processing strategies represent an important avenue for optimization. Dynamic adjustment of sampling frequency, LED driving current, and preprocessing complexity based on signal-quality indices or detected physiological states could significantly extend battery life during long-term monitoring. Event-driven or interrupt-based architectures may further reduce unnecessary computation during periods of signal inactivity or poor quality, enabling continuous monitoring on wrist-worn, patch-based, or textile-integrated devices powered by small batteries or energy-harvesting modules.

For lung ultrasound applications, embedded deployment introduces distinct challenges and opportunities due to the higher data rates and dimensionality of imaging signals. The block-wise VMD algorithm is inherently parallelizable, making it well suited for implementation on heterogeneous platforms that integrate CPUs with programmable logic or dedicated accelerators. Future work may explore field-programmable gate array (FPGA) implementations or heterogeneous SoCs in which Fourier transforms, band-pass filtering, and mode reconstruction are executed concurrently in hardware. The use of high-level synthesis (HLS) tools could facilitate the migration of the current algorithmic descriptions to hardware description languages, reducing development time while enabling fine-grained control over latency and power consumption. Achieving real-time processing of ultrasound frames (e.g., 25–30 fps) directly on the imaging device would represent a significant step toward fully autonomous point-of-care diagnostic systems.

Beyond raw computational acceleration, energy-aware task scheduling and adaptive computation are critical research topics for both applications. Techniques such as dynamic voltage and frequency scaling (DVFS), selective activation of processing modules, and approximate computing could be explored to balance diagnostic performance with energy efficiency. When combined with edge-to-cloud architectures, these strategies would allow lightweight decision support to be performed locally while reserving more computationally intensive analyses for cloud or edge servers, depending on latency, connectivity, and privacy constraints.

5.3.2 Clinical Validation and Scalability

A second, equally critical direction for future work concerns comprehensive clinical validation and scalability assessment. While the proposed algorithms demonstrated

robust performance on dedicated datasets acquired under controlled or semi-controlled conditions, their generalization to diverse patient populations, acquisition devices, and clinical workflows remains an open challenge. Addressing this limitation is essential for transitioning from experimental validation to real-world clinical adoption.

For cuffless blood pressure monitoring, future studies should be designed in accordance with international standards such as ISO 81060-2:2018, the AAMI/ESH/ISO Universal Standard, and IEEE 1708-2014. This entails large-scale validation across heterogeneous populations, spanning different age groups, body compositions, ethnic backgrounds, and cardiovascular conditions. Particular emphasis should be placed on long-term and longitudinal studies to assess intra-subject stability, calibration drift, and sensitivity to physiological changes induced by physical activity, stress, medication, or disease progression. Evaluating performance during daily-life conditions—characterized by motion artifacts, variable temperature, and sensor displacement—will be crucial to demonstrate robustness beyond laboratory settings.

In parallel, integration with secure cloud-based infrastructures could support large-scale data aggregation and analytics, enabling population-level model refinement and continuous performance monitoring. Privacy-preserving approaches such as federated learning and on-device personalization represent promising strategies to adapt models to individual physiological profiles without transferring raw data off the device, thereby addressing both regulatory and ethical concerns.

For lung ultrasound analysis, clinical scalability requires validation across multiple ultrasound systems, probe types, operating frequencies, and acquisition protocols. Future studies should involve multi-center datasets annotated by multiple experts to quantify inter-observer variability and to assess the robustness of the VMD-based approach across operators with different levels of expertise. Extending the analysis from single frames to temporal sequences (cine-loops) would allow the incorporation of dynamic information related to lung sliding, respiratory motion, and artifact persistence, potentially improving diagnostic accuracy and clinical relevance.

An important translational step involves the integration of the proposed algorithms into real-time scanning environments, where automated analysis can assist clinicians during image acquisition rather than in retrospective review. Providing immediate visual feedback, quantitative indices, and standardized scoring could reduce operator dependency, shorten examination times, and improve consistency in pulmonary assessment at the bedside.

Beyond algorithmic validation, scalability also encompasses regulatory, manufacturing, and interoperability considerations. Embedded implementations must comply with medical-device standards such as IEC 60601 for electrical safety, ISO 14971 for risk management, and IEC 62304 for medical software life-cycle processes. Achieving regulatory approval, including CE marking or FDA clearance, will require rigorous documentation, risk analysis, verification and validation procedures, and post-market surveillance strategies. From a systems-integration perspective, compatibility with hospital information systems and interoperability standards such as HL7 and FHIR will be essential to ensure seamless integration into existing clinical workflows.

Finally, the impact of this research extends beyond the specific applications addressed in this thesis. The methodological principles developed—hardware-aware signal processing, interpretable modeling, and efficient embedded deployment—are broadly

Chapter 5. Conclusions and Future Work

applicable to other biomedical domains, including cardiac imaging, neurological monitoring, and multimodal wearable sensing. As such, the future directions outlined in this section define a roadmap that moves from laboratory prototypes toward clinically validated, intelligent medical systems, capable of supporting precision medicine and connected healthcare at scale.

Bibliography

- [1] Blobel, B. and Ruotsalainen, P. Healthcare transformation towards personalized medicine – chances and challenges. In *Studies in Health Technology and Informatics, Volume 261: pHealth 2019*, pages 3–21. IOS Press, 2019.
- [2] Saraswat, D. and et al. Explainable ai for healthcare 5.0: Opportunities and challenges. *IEEE Access*, 10:84486–84517, 2022.
- [3] Park, S. and Jayaraman, S. Wearable biomedical systems and mhealth. In *Telemedicine and Electronic Medicine*, pages 26–52. CRC Press, 1st edition, 2016.
- [4] Gouaux, F. and Chautemps, L.S. and Fayn, J. and Adami, S. and Arzi, M. and Assanelli, D. and Forlini, M.-C. and Malossi, C. and Martinez, A. and Ohlsson, M. and Placide, J. and Ziliani, G.-L. and Rubel, P. Pervasive self-care solutions in telecardiology: Typical use cases from the epi-medics project. In *Studies in Health Technology and Informatics, Volume 95: The New Navigators – from Professionals to Patients*, pages 119–124. IOS Press, 2003.
- [5] Davoudi, A. and Malhotra, K.R. and Shickel, B. and et al. Intelligent icu for autonomous patient monitoring using pervasive sensing and deep learning. *Scientific Reports*, 9(8020):1–13, 2019.
- [6] Atallah, L. and Lo, B. and Yang, G.Z. Can pervasive sensing address current challenges in global healthcare? *Journal of Epidemiology and Global Health*, 2(1):1–13, 2012.
- [7] Ebrahimi, Z. and Gosselin, B. Photodetector innovations for advancing wearable optical biosensors: A review. *IEEE Sensors Journal*, 25(3):4070–4095, February 2025.
- [8] Gounella, R. and Ferreira, G.M. and Amorim, M.L.M. and Soares, J.N., Jr. and Carmo, J.P. A review of optical sensors in cmos. *Electronics*, 13(4):691, 2024.
- [9] Ma, W.-J. and Luo, C.-H. and Lin, J.-L. and Chou, S.-H. and Chen, P.-H. and Syu, M.-J. and Kuo, S.-H. and Lai, S.-C. A portable low-power acquisition system with a urease bioelectrochemical sensor for potentiometric detection of urea concentrations. *Sensors*, 16(4):474, 2016.
- [10] Wu, Y.-D. and Ruan, S.-J. and Lee, Y.-H. An ultra-low power surface emg sensor for wearable biometric and medical applications. *Biosensors*, 11(11):411, 2021.
- [11] Gkonis, P. and Giannopoulos, A. and Trakadas, P. and Masip-Bruin, X. and D’Andria, F. A survey on iot-edge-cloud continuum systems: Status, challenges, use cases, and open issues. *Future Internet*, 15(12):383, 2023.
- [12] Khalyeyev, D. and Bureš, T. and Hnětynka, P. Towards characterization of edge-cloud continuum. In Batista, T. and Bureš, T. and Raibulet, C. and Muccini, H., editor, *Software Architecture. ECSA 2022 Tracks and Workshops*, volume 13928 of *Lecture Notes in Computer Science*, pages 197–212. Springer, Cham, 2023.
- [13] Goetz, L.H. and Schork, N.J. Personalized medicine: Motivation, challenges, and progress. *Fertility and Sterility*, 109(6):952–963, 2018.
- [14] Botrugno, C. and Leogrande, E. and Dell’Olio, F. and Natale, T. Smart combination of ecg and ppg signals: An innovative approach towards an electronic device for vital signs monitoring. In *Proceedings of the 2023 9th International Workshop on Advances in Sensors and Interfaces (IWASI)*, pages 256–260, Monopoli (Bari), Italy, 2023. IEEE.

Bibliography

- [15] Daddi, N. and Sagrini, E. and Lugaresi, M. and Perrone, O. and Candoli, P. and Ragusa, M. and Puma, F. and Mattioli, S. Ex vivo pulmonary nodule detection with miniaturized ultrasound convex probes. *Journal of Surgical Research*, 202(1):49–57, 2016.
- [16] Stevens, S.L. and Wood, S. and Koshiaris, C. and Law, K. and Glasziou, P. and Stevens, R.J. and et al. Blood pressure variability and cardiovascular disease: Systematic review and meta-analysis. *BMJ*, 354:i4098, 2016.
- [17] Kario, K. Management of hypertension in the digital era: Small wearable monitoring devices for remote blood pressure monitoring. *Hypertension*, 76(3):640–650, 2020.
- [18] Watanabe, N. and Bando, Y. and Kawachi, T. and et al. Development and validation of a novel cuff-less blood pressure monitoring device. *JACC: Basic to Translational Science*, 2(6):631–642, 2017.
- [19] Botrugno, C. and Dell’Olio, F. Ppg-based electronic sensing system for cuff-less calibration-free blood pressure estimation using machine learning algorithms. In *Proceedings of the 2024 IEEE Sensors Applications Symposium (SAS)*, pages 1–6, Cagliari, Italy, 2024. IEEE.
- [20] Riaz, F. and Azad, M.A. and Arshad, J. and Imran, M. and Hassan, A. and Rehman, S. Pervasive blood pressure monitoring using photoplethysmogram (ppg) sensor. *Future Generation Computer Systems*, 98:120–130, 2019.
- [21] Jadhav, S.P. and et al. Introduction to lung diseases. In Dua, K. and Löbenberg, R. and Malheiros Luzo, Â.C. and Shukla, S. and Satija, S., editor, *Targeting Cellular Signalling Pathways in Lung Diseases*, pages 1–20. Springer, Singapore, 2021.
- [22] Anthimopoulos, M. and Christodoulidis, S. and Ebner, L. and Christe, A. and Mougiakakou, S. Lung pattern classification for interstitial lung diseases using a deep convolutional neural network. *IEEE Transactions on Medical Imaging*, 35(5):1207–1216, May 2016.
- [23] Doshi, J.A. and Mundhra, K.S. and Shah, D.S. and et al. Role of high-resolution ct thorax in diagnosing interstitial lung disease and its association with smoking and connective tissue disorder. *Cureus*, 14(11):e31107, 2022.
- [24] An, J. and Lee, J.H. and Yoo, Y. and Kwon, H.S. and Lee, J.S. and Lee, S.W. and Kim, T.B. and Oh, Y.M. and Cho, Y.S. and Lee, S.D. and Song, W.J. Chest computed tomography scan utilization and diagnostic outcomes in chronic cough patients with normal chest x-rays: Analysis of routinely collected data of a tertiary academic hospital. *Journal of Thoracic Disease*, 15(4):2324–2332, 2023.
- [25] Calamai, I. and Greco, M. and Bertolini, G. and Spina, R. Current adoption of lung ultrasound in intensive care units: An italian multi-center survey. *Minerva Anestesiologica*, 83:720–727, 2017.
- [26] Ferro, F. and Delle Sedie, A. The use of ultrasound for assessing interstitial lung involvement in connective tissue diseases. *Clinical and Experimental Rheumatology*, 36(Suppl. S114):165–170, 2018.
- [27] Marozzi, M.S. and Cicco, S. and Mancini, F. and Corvasce, F. and Lombardi, F.A. and Desantis, V. and Loponte, L. and Giliberti, T. and Morelli, C.M. and Longo, S. and et al. A novel automatic algorithm to support lung ultrasound non-expert physicians in interstitial pneumonia evaluation: A single-center study. *Diagnostics*, 14(2):155, 2024.
- [28] Brusasco, C. and Santori, G. and Bruzzo, E. and et al. Quantitative lung ultrasonography: A putative new algorithm for automatic detection and quantification of b-lines. *Critical Care*, 23(288):1–9, 2019.
- [29] Zosso, D. and Dragomiretskiy, K. and Bertozzi, A.L. and et al. Two-dimensional compact variational mode decomposition. *Journal of Mathematical Imaging and Vision*, 58(2):294–320, 2017.
- [30] Eqdad, M.N. and Ahmed, R.T. and Alhandhal, M. Processing biomedical signals by neural networks using hardware-constrained system. *Revue d’Intelligence Artificielle*, 38(2):739–745, 2024.
- [31] Tran, T.K. and Huynh, M.T. and Do, Q.M.D. and Pham, C.K. and Huynh, H.T. A hardware-friendly approach for ecg classification based on deep learning algorithm. In Nguyen, N.T. and et al., editor, *Computational Intelligence in Engineering Science. ICCIES 2025*, volume 2584 of *Communications in Computer and Information Science*, pages 1–12. Springer, Cham, 2026.
- [32] Akhtar, N. and Buzdar, A. and Fan, J. and Khan, M.U. Cardio-edge: Hardware–software co-design implementation of lstm-based ecg classification for continuous cardiac monitoring on wearable devices. *International Journal of Advanced Computer Science and Applications*, 16(7):1–8, 2025.
- [33] Gonzalez-Carabarin, L. and Huijben, I.A.M. and Veeling, B. and Schmid, A. and van Sloun, R.J.G. Dynamic probabilistic pruning: A general framework for hardware-constrained pruning at different granularities. *IEEE Transactions on Neural Networks and Learning Systems*, 35(1):733–744, January 2024.

- [34] Zhang, C. and Cui, R. and Wang, G. and Gao, T. and Yan, J. and Xian, W. and Wei, X. and Qin, Y. Hardware-aware neural architecture search for real-time video processing in fpga-accelerated endoscopic imaging. *Applied Sciences*, 15(20):11200, 2025.
- [35] Patel, S. and Park, H. and Bonato, P. and Chan, L. and Rodgers, M. A review of wearable sensors and systems with application in rehabilitation. *Journal of NeuroEngineering and Rehabilitation*, 9:21, 2012.
- [36] Topol, E.J. *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again*. Basic Books, New York, NY, USA, 2019.
- [37] Sörnmo, L. and Laguna, P. *Bioelectrical Signal Processing in Cardiac and Neurological Applications*. Academic Press, Cambridge, MA, USA, 1st edition, 2005.
- [38] Webster, J.G. *Medical Instrumentation: Application and Design*. John Wiley & Sons, Hoboken, NJ, USA, 4th edition, 2010.
- [39] Heikenfeld, J. and Jajack, A. and Rogers, J. and Gutruf, P. and Tian, L. and Pan, T. and Li, R. and Khine, M. and Kim, J. and Wang, J. and Kim, J. Wearable sensors: Modalities, challenges, and prospects. *Lab on a Chip*, 18(2):217–248, 2018.
- [40] Gao, W. and Emaminejad, S. and Nyein, H.N. and Challa, S. and Chen, K. and Peck, A. and Fahad, H.M. and Ota, H. and Shiraki, H. and Kiriya, D. and Lien, D.-H. and Brooks, G. and Davis, R. and Javey, A. Fully integrated wearable sensor arrays for multiplexed *In Situ* perspiration analysis. *Nature*, 529(7587):509–514, 2016.
- [41] Hall, J.E. and Hall, M.E. *Guyton and Hall Textbook of Medical Physiology*. Elsevier, Philadelphia, PA, USA, 14th edition, 2020.
- [42] Shaffer, F. and Ginsberg, J.P. An overview of heart rate variability metrics and norms. *Frontiers in Public Health*, 5:258, 2017.
- [43] Harrison, R.R. and Charles, C. A low-power low-noise cmos amplifier for neural recording applications. *IEEE Journal of Solid-State Circuits*, 38(6):958–965, June 2003.
- [44] Chi, Y.M. and Jung, T.P. and Cauwenberghs, G. Dry-contact and noncontact biopotential electrodes: Methodological review. *IEEE Reviews in Biomedical Engineering*, 3:106–119, 2010.
- [45] Dimitrova, N. and Dimitrov, G.V. Electromyography (emg) modeling. In *Encyclopedia of Biomedical Engineering*. John Wiley & Sons, 2006.
- [46] Ramasamy, S. and Balan, A. Wearable sensors for ecg measurement: A review. *Sensor Review*, 38(4):412–422, 2018.
- [47] Yao, S. and Swetha, P. and Zhu, Y. Nanomaterial-enabled wearable sensors for healthcare. *Advanced Healthcare Materials*, 7(1), 2018.
- [48] Casson, A.J. Wearable eeg and beyond. *Biomedical Engineering Letters*, 9(1):53–71, 2019.
- [49] Mullen, T.R. and Kothe, C.A. and Chi, Y.M. and Ojeda, A. and Kerth, T. and Makeig, S. and Jung, T.P. and Cauwenberghs, G. Real-time neuroimaging and cognitive monitoring using wearable dry eeg. *IEEE Transactions on Biomedical Engineering*, 62(11):2553–2567, November 2015.
- [50] Someya, T. and Bao, Z. and Malliaras, G. The rise of plastic bioelectronics. *Nature*, 540:379–385, 2016.
- [51] Dagdeviren, C. and Su, Y. and Joe, P. and Yona, R. and Liu, Y. and Kim, Y.S. and Huang, Y. and Damadoran, A.R. and Xia, J. and Martin, L.W. and Huang, Y. and Rogers, J.A. Conformable amplified lead zirconate titanate sensors with enhanced piezoelectric response for cutaneous pressure monitoring. *Nature Communications*, 5:4496, 2014.
- [52] Allen, J. Photoplethysmography and its application in clinical physiological measurement. *Physiological Measurement*, 28(3):R1–R39, 2007.
- [53] Tamura, T. and Maeda, Y. and Sekine, M. and Yoshida, M. Wearable photoplethysmographic sensors—past and present. *Electronics*, 3(2):282–302, 2014.
- [54] El Hajj, C. and Kyriacou, P.A. Cuffless and continuous blood pressure estimation from ppg signals using recurrent neural networks. In *Proceedings of the 2020 Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 4269–4272. IEEE, 2020.
- [55] Šlapničar, G. and Mlakar, N. and Luštrek, M. Blood pressure estimation from photoplethysmogram using a spectro-temporal deep neural network. *Sensors*, 19(15):3420, 2019.
- [56] Xing, X. and Sun, M. Optical blood pressure estimation with photoplethysmography and fft-based neural networks. *Biomedical Optics Express*, 7(8):3007–3020, 2016.

Bibliography

- [57] Nyein, H.Y. and Gao, W. and Shahpar, Z. and Emaminejad, S. and Challa, S. and Chen, K. and Fahad, H.M. and Tai, L.C. and Ota, H. and Davis, R.W. and Javey, A. A wearable electrochemical platform for noninvasive simultaneous monitoring of Ca^{2+} and pH . *ACS Nano*, 10(7):7216–7224, 2016.
- [58] Kim, J. and Campbell, A.S. and Wang, J. Wearable non-invasive epidermal glucose sensors: A review. *Talanta*, 177:163–170, 2018.
- [59] Homayounfar, S.Z. and Andrew, T.L. Wearable sensors for monitoring human motion: A review on mechanisms, materials, and challenges. *Journal of Laboratory Automation*, 25(1):9–24, 2019.
- [60] Pantelopoulos, A. and Bourbakis, N. A survey on wearable sensor-based systems for health monitoring and prognosis. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 40(1):1–12, 2010.
- [61] Zhang, Z. and Pi, Z. and Liu, B. Troika: A general framework for heart rate monitoring using wrist-type photoplethysmographic signals during intensive physical exercise. *IEEE Transactions on Biomedical Engineering*, 62(2):522–531, 2015.
- [62] Botrugno, C. and Leogrande, E. and Natale, T. and Dell’Olio, F. Smart combination of ecg and ppg signals: An innovative approach towards an electronic device for vital signs monitoring. In *Proceedings of the IEEE International Workshop on Advances in Sensors and Interfaces (IWASI)*, Monopoli, Italy, 2023. IEEE.
- [63] Diab, M.S. and Rodriguez-Villegas, E. Embedded ml using microcontrollers in wearable and ambulatory systems. *IEEE Access*, 10:100000–100010, 2022.
- [64] Li, X. and Blanton, R.D. and Grover, P. and Thomas, D.E. Ultra-low-power biomedical circuit design and optimization: Catching the don’t cares. In *2014 International Symposium on Integrated Circuits (ISIC)*, pages 1–4. IEEE, 2014.
- [65] Kouhalvandi, L. and Matekovits, L. and Peter, I. Amplifiers in biomedical engineering: A review from application perspectives. *Sensors*, 23(4):2277, 2023.
- [66] Fateh, M. and Bouguechal, N.-E. and Belhadji, Y. A low power low noise chopper-stabilized two-stage operational amplifier for portable biopotential acquisition systems using 90 nm technology. *International Journal of Hybrid Information Technology*, 7(6):25–42, 2014.
- [67] Goldberger, A.L. and Amaral, L.A. and Glass, L. and Hausdorff, J.M. and Ivanov, P.C. and Mark, R.G. and Mietus, J.E. and Moody, G.B. and Peng, C.K. and Stanley, H.E. Physiobank, physiotoolkit, and physionet: Components of a new research resource for complex physiologic signals. *Circulation*, 101(23):e215, 2000.
- [68] Otto, C. and Milenković, A. and Sanders, C. and Jovanov, E. System architecture of a wireless body area sensor network for ubiquitous health monitoring. *Journal of Mobile Multimedia*, 1(4):307–326, 2005.
- [69] Al-Kofahi, M. and Al-Shorman, M. and Al-Kofahi, O. *Toward Energy Efficient Microcontrollers and IoT Systems*. ResearchGate, 2019. Available online at <https://www.researchgate.net/publication/336775585>.
- [70] Hinchet, R. and Kim, S.W. Wearable and implantable mechanical energy harvesters for self-powered biomedical systems. *ACS Nano*, 9(8):7742–7745, 2015.
- [71] Stoppa, M. and Chiolerio, A. Wearable electronics and smart textiles: A critical review. *Sensors*, 14(7):11957–11992, 2014.
- [72] Azadmousavi, T. and Ghafar-Zadeh, E. Complementary metal–oxide–semiconductor-based magnetic and optical sensors for life science applications. *Sensors*, 24(19):6264, 2024.
- [73] Han, S. and Mao, H. and Dally, W. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2016. Available at <https://arxiv.org/abs/1510.00149>.
- [74] Liu, S. and Yi, Y. *Quantized Neural Networks and Neuromorphic Computing for Embedded Systems*. IntechOpen, 2020.
- [75] Cheng, Y. and Wang, D. and Zhou, P. and Zhang, T. Model compression and acceleration for deep neural networks: The principles, progress, and challenges. *IEEE Signal Processing Magazine*, 35(1):126–136, 2018.
- [76] Wang, D. and Liao, W.H. Wireless transmission for health monitoring of large structures. *IEEE Transactions on Instrumentation and Measurement*, 55(3):972–981, 2006.
- [77] Gope, P. and Hwang, T. Bsn-care: A secure iot-based modern healthcare system using body sensor network. *IEEE Sensors Journal*, 16(5):1368–1376, 2016.

- [78] Rancea, A. and Anghel, I. and Cioara, T. Edge computing in healthcare: Innovations, opportunities, and challenges. *Future Internet*, 16(9):329, 2024.
- [79] Yip, M. *Ultra-Low-Power Circuits and Systems for Wearable and Implantable Medical Devices*. University of California, Berkeley, 2014. PhD Thesis.
- [80] Blanco-Velasco, M. and Weng, B. and Barner, K.E. Ecg signal denoising and baseline wander correction based on the empirical mode decomposition. *Computers in Biology and Medicine*, 38(1):1–13, 2008.
- [81] Smital, L. and Vitek, M. and Kozumplík, J. and Provazník, V. Adaptive wavelet wiener filtering of ecg signals. *IEEE Transactions on Biomedical Engineering*, 60(2):437–445, 2012.
- [82] Han, H. and Kim, J. Artifacts in wearable photoplethysmographs during daily life motions and their reduction with least mean square based active noise cancellation method. *Computers in Biology and Medicine*, 42(4):387–393, 2012.
- [83] Ko, B.H. and Lee, T. and Choi, C. and Kim, Y.H. and Park, G. and Kang, K. and Bae, S.K. and Shin, K. Motion artifact reduction in electrocardiogram using adaptive filtering based on half cell potential monitoring. In *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 1590–1593. IEEE, 2012.
- [84] Dragomiretskiy, K. and Zosso, D. Variational mode decomposition. *IEEE Transactions on Signal Processing*, 62(3):531–544, 2014.
- [85] Nugroho, H.A. and Rahmawaty, M. and Triyani, Y. and Ardiyanto, I. and Choridah, L. and Indrastuti, R. Texture analysis and classification in ultrasound medical images for determining echo pattern characteristics. In *2017 International Conference on Sustainable Engineering and Technology (ICSEngT)*, pages 23–26. IEEE, 2017.
- [86] Brand, M. Incremental singular value decomposition of uncertain data with missing values. In Heyden, A. and Sparr, G. and Nielsen, M. and Johansen, P., editor, *Computer Vision — ECCV 2002*, volume 2350 of *Lecture Notes in Computer Science*, pages 707–720. Springer, Berlin, Heidelberg, 2002.
- [87] Warden, P. and Situnayake, D. *TinyML: Machine Learning with TensorFlow Lite on Arduino and Ultra-Low-Power Microcontrollers*. O'Reilly Media, Sebastopol, CA, USA, 2019.
- [88] Pereira, C.V.F. and de Oliveira, E.M. and de Souza, A.D. Machine learning applied to edge computing and wearable devices for healthcare: Systematic mapping of the literature. *Sensors*, 24(19):6322, 2024.
- [89] Lai, L. and Suda, N. and Chandra, V. Cmsis-nn: Efficient neural network kernels for arm cortex-m cpus. arXiv preprint arXiv:1801.06601, 2018. Available at <https://arxiv.org/abs/1801.06601>.
- [90] Das, M. and Choudhary, T. and Sharma, L.N. and Bhuyan, M.K. Noninvasive accelerometric approach for cuffless continuous blood pressure measurement. *IEEE Transactions on Instrumentation and Measurement*, 70, 2021.
- [91] Mukkamala, R. and Hahn, J.-O. and Inan, O.T. and Mestha, L.K. and Kim, C.-S. and Tawfik, H. and Kyal, R. and Nabavi, S.S. Toward ubiquitous blood pressure monitoring via pulse transit time: Theory and practice. *IEEE Transactions on Biomedical Engineering*, 62(8):1879–1901, 2015.
- [92] Botrugno, C. and Dheman, K. and Bonazzi, P. and Dell’Olio, F. and Magno, M. Ai-based multi-wavelength ppg device for blood pressure monitoring. In *Proceedings of the IEEE International Symposium on Medical Measurements and Applications (MeMeA)*, pages 1–6, Eindhoven, Netherlands, 2024. IEEE.
- [93] Cheng, X. and Wei, S. and Zhang, R. Limitations of non-contact blood pressure measurement using facial video. *Journal of Biomedical Engineering*, 44(3):556–568, 2022.
- [94] Zhang, G. and Gao, C. and Wang, J. and He, Y. Evaluation of pulse wave analysis in blood pressure measurement: A systematic review. *Journal of Medical Internet Research*, 22(8):e18295, 2020.
- [95] Ray, D. and Collins, T. and Woolley, S.I. and Ponnappalli, P.V.S. A review of wearable multi-wavelength photoplethysmography. *IEEE Reviews in Biomedical Engineering*, 16:136–151, 2023.
- [96] Liu, Y. and Sun, X. and Li, T. Camera-based blood pressure estimation: Opportunities and limitations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 107–115. IEEE, 2021.
- [97] El-Hajj, C. and Kyriacou, P.A. A review of machine learning techniques in photoplethysmography for the non-invasive cuff-less measurement of blood pressure. *Biomedical Signal Processing and Control*, 58, 2020.
- [98] Mancia, G. and Fagard, R. and Narkiewicz, K. and Redon, J. and Zanchetti, A. and Böhm, M. and Schmieder, R. Esh guidelines for the management of arterial hypertension: Recommendations and challenges. *European Heart Journal*, 42(5):301–320, 2021.

Bibliography

- [99] Analog Devices Inc. *MAX86916 Evaluation System — Evaluates: MAX86916*. Analog Devices Inc., 2020. Rev. 0, February 2020. Available online at <https://www.analog.com/en/resources/evaluation-hardware-and-software/evaluation-boards-kits/max86916evsys.html>.
- [100] Bhoil, R. and Ahluwalia, A. and Chopra, R. and M, M. Surya and Bhoil, S. Signs and lines in lung ultrasound. *Journal of Ultrasonography*, 21(86):e225–e233, 2021.
- [101] Dietrich, C.F. and Mathis, G. and Blaivas, M. and Volpicelli, G. and Seibel, A. and Wastl, D. and Atkinson, N.S. and Cui, X.W. and Fan, M. and Yi, D. Lung b-line artefacts and their use. *Journal of Thoracic Disease*, 8(6):1356–1365, 2016.
- [102] Sperandeo, M. and Varriale, A. and Sperandeo, G. and others. Assessment of ultrasound acoustic artifacts in patients with acute dyspnea: A multicenter study. *Acta Radiologica*, 53:885–892, 2012.
- [103] Copetti, R. and Soldati, G. and Copetti, P. Chest sonography: A useful tool to differentiate acute cardiogenic pulmonary edema from acute respiratory distress syndrome. *Cardiovascular Ultrasound*, 6:6–16, 2008.
- [104] Gu, J. and Jing, Y. Modeling of wave propagation for medical ultrasound: A review. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 62(11):1979–1992, 2015.
- [105] Anantrasirichai, N. and Allinovi, M. and Hayes, W. and Achim, A. Automatic b-line detection in paediatric lung ultrasound. In *2016 IEEE International Ultrasonics Symposium (IUS)*, pages 1–4, Tours, France, 2016. IEEE.
- [106] Brusasco, C. and Santori, G. and Bruzzo, E. and others. Quantitative lung ultrasonography: A putative new algorithm for automatic detection and quantification of b-lines. *Critical Care*, 23:288, 2019.
- [107] van Sloun, R.J.G. and Demi, L. Localizing b-lines in lung ultrasonography by weakly supervised deep learning, in-vivo results. *IEEE Journal of Biomedical and Health Informatics*, 24(4):957–964, 2020.
- [108] Lucassen, R.T. and others. Deep learning for detection and localization of b-lines in lung ultrasound. *IEEE Journal of Biomedical and Health Informatics*, 27(9):4352–4361, 2023.
- [109] Anantrasirichai, N. and Hayes, W. and Allinovi, M. and Bull, D. and Achim, A. Line detection as an inverse problem: Application to lung ultrasound imaging. *IEEE Transactions on Medical Imaging*, 36(10):2045–2056, 2017.
- [110] Baloesescu, C. and others. Automated lung ultrasound b-line assessment using a deep learning algorithm. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 67(11):2312–2320, 2020.
- [111] Giulia, M. and Alessandro Stuart, S. and Giosuè, C. and others. The delay multiply and sum beamforming algorithm in ultrasound b-mode medical imaging. *IEEE Transactions on Medical Imaging*, 34(4):940–949, 2014.
- [112] Lian, J. and Liu, Z. and Wang, H. and Dong, X. Adaptive variational mode decomposition method for signal processing based on mode characteristic. *Mechanical Systems and Signal Processing*, 107:53–77, 2018.
- [113] Miao, Q. and Shu, Q. and Wu, B. and Sun, X. and Song, K. A modified complex variational mode decomposition method for analyzing nonstationary signals with the low-frequency trend. *Sensors*, 22(5):1801, 2022.
- [114] Robert, M. Phased array theory and technology. *Proceedings of the IEEE*, 70(3):246–291, 1982.
- [115] Giulia, M. and Alessandro Stuart, S. and Giosuè, C. and others. The delay multiply and sum beamforming algorithm in ultrasound b-mode medical imaging. *IEEE Transactions on Medical Imaging*, 34(4):940–949, 2014.
- [116] Anantrasirichai, N., Hayes, W., Allinovi, M., Bull, D., and Achim, A. Line detection as an inverse problem: Application to lung ultrasound imaging. *IEEE Transactions on Medical Imaging*, 36(10):2045–2056, 2017.
- [117] Farahi, M., Aranda, J., Habibian, H., and Casals, A. Automatic feature detection in lung ultrasound images using wavelet and radon transforms. *SSRN Electronic Journal*, pages 1–10, 2024.
- [118] Okila, N., Katumba, A., Nakatumba-Nabende, J., Mwikirize, C., Murindanyi, S., Serugunda, J., Bugeza, S., Oriekot, A., Bossa, J., and Nabawanuka, E. Deep learning for accurate b-line detection and localization in lung ultrasound imaging. *Frontiers in Artificial Intelligence*, 8:1–12, 2025.
- [119] Bottino, A., Botrugno, C., Casciaro, E., and et al. Automatic approach for b-lines detection in lung ultrasound images using you only look once algorithm. *Journal of Ultrasound*, pages 1–12, 2025.
- [120] Hooi, L. and Thi, N. and Er-Ping, L. and others. Confocal microwave imaging for breast cancer detection: Delay-multiply-and-sum image reconstruction algorithm. *IEEE Transactions on Biomedical Engineering*, 55(6):1697–1705, 2008.