Edge-AI and Wearable Sensor Telemetry for Early Signal Detection of Cardiovascular Decompensation in Rural and Underserved U.S. Health Networks

Authors

Bob Cuddeeback, Sunday Oladele, Anderson Parmer, Steven Gerhart, Henk Nijaland, Eric Tesson, Dennis Jose

Date; July 21, 2026

Abstract

Cardiovascular disease remains the leading cause of mortality in the United States, with rural and underserved populations experiencing disproportionately higher rates of adverse outcomes due to limited access to continuous physiological monitoring and specialist care. Traditional approaches to detecting cardiovascular decompensation rely on episodic clinical assessments that fail to capture the gradual physiological deterioration preceding acute events. This study addresses the critical gap in early warning capabilities by developing and validating a hybrid Edge-AI framework for real-time detection of cardiovascular decompensation using wearable sensor telemetry. The proposed system integrates multimodal physiological data—including heart rate variability, photoplethysmography waveforms, respiratory rate, and oxygen saturation—processed through a lightweight gradient-boosting machine learning pipeline optimized for edge deployment. Using retrospective data from 1,247 patients with heart failure across rural health networks and prospective simulation validation, the framework achieved 89.4% accuracy (AUC = 0.93) in predicting decompensation events 4.2 hours before clinical recognition, with a false alarm rate of 12.6%. The system demonstrated robust performance on resource-constrained hardware (latency < 100ms, power consumption < 50mW), making it suitable for deployment in low-infrastructure settings. These findings suggest that Edge-AI-enabled wearable telemetry can provide clinically meaningful early warnings of cardiovascular deterioration, potentially reducing emergency admissions and improving outcomes in underserved communities. The framework offers a scalable, privacy-preserving solution for extending advanced cardiac monitoring capabilities to populations currently underserved by traditional healthcare infrastructure.

Keywords: Edge Artificial Intelligence, Wearable Sensors, Cardiovascular Decompensation, Rural Health, Remote Patient Monitoring, Predictive Analytics

1. Introduction

1.1 Background

Cardiovascular disease (CVD) remains the leading cause of morbidity and mortality in the United States, accounting for approximately 19.8 million deaths annually worldwide. Within the U.S., rural populations face significantly worse cardiovascular outcomes compared to their urban counterparts, driven by a confluence of factors including limited access to specialty care, longer emergency response times, and lower rates of preventive health engagement. The American Heart Association has identified rural-urban disparities in cardiovascular mortality as a critical priority area, noting that residents of rural counties have a 20% higher age-adjusted death rate from heart disease than those in metropolitan areas.

Acute decompensated heart failure (ADHF) represents one of the most common and costly manifestations of cardiovascular deterioration. Patients typically experience gradual physiological changes—fluid accumulation, worsening dyspnea, declining heart rate variability, and subtle hemodynamic shifts—over days to weeks before reaching a crisis point requiring emergency hospitalization. Traditional monitoring approaches, which rely on episodic clinical visits and patient self-reporting of symptoms, consistently fail to detect these early warning signs. The result is that many patients present to emergency departments in advanced stages of decompensation, leading to prolonged hospitalizations, higher mortality rates, and substantial healthcare costs exceeding $30 billion annually in the U.S. alone.

Recent technological advances have created new possibilities for addressing this monitoring gap. Wearable sensor technologies—including photoplethysmography (PPG), electrocardiography (ECG), and inertial measurement units—have become increasingly miniaturized and affordable, enabling continuous physiological data collection outside clinical settings. Concurrently, artificial intelligence (AI) has emerged as a transformative tool for analyzing the complex, high-dimensional data streams generated by these devices. The convergence of wearable technology and AI represents a paradigm shift in cardiovascular healthcare, enabling continuous real-time monitoring beyond conventional clinical settings .

1.2 Problem Statement

Despite the significant potential of AI-enhanced wearable monitoring, several critical barriers have prevented its effective deployment in rural and underserved U.S. health networks. First, existing AI models for cardiovascular decompensation prediction have been predominantly developed and validated using high-quality data collected in academic medical centers, limiting their generalizability to populations with different demographic characteristics, comorbidity profiles, and environmental exposures. Second, these models typically require substantial computational resources and cloud connectivity, making them impractical in settings with limited broadband infrastructure. Third, concerns regarding data privacy and transmission security have created regulatory and patient acceptance barriers to cloud-based monitoring solutions. Finally, the implementation science challenges of integrating novel monitoring technologies into existing rural healthcare workflows remain largely unexplored.

Current research has demonstrated the technical feasibility of detecting unstable heart failure using machine learning approaches on physiological data . Projects such as HeartFM have shown that algorithms can identify patterns in physiological data that precede clinical decompensation using implantable cardioverter-defibrillator data. However, these approaches have not been specifically designed for the constraints of rural and underserved settings. The hardware requirements for edge deployment, the computational efficiency demands of continuous real-time inference, and the specific physiological signatures of decompensation in diverse patient populations remain inadequately addressed by the existing literature.

The specific gap addressed by this study is the absence of a validated, resource-constrained predictive framework that combines multimodal wearable sensor data with Edge-AI processing to detect early signals of cardiovascular decompensation in rural and underserved U.S. populations. While prior work has established the potential of machine learning for prediction of heart failure events , no comprehensive framework has been developed that addresses the unique technical, clinical, and implementation challenges of deployment in low-resource settings. Furthermore, the optimal combination of sensor modalities, algorithmic approaches, and lead-time thresholds for practical clinical utility in rural health networks has not been systematically determined.

1.3 Objectives of the Study

General objective:
To develop and validate an Edge-AI framework that integrates multimodal wearable sensor telemetry for early detection of cardiovascular decompensation in rural and underserved U.S. health networks.

Specific objectives:

1. To identify the most predictive physiological signal combinations from multimodal wearable sensors for early detection of cardiovascular decompensation.
2. To design and implement a lightweight hybrid machine learning pipeline optimized for edge deployment on resource-constrained hardware.
3. To validate the predictive performance and real-time feasibility of the framework using retrospective clinical data and prospective simulation.
4. To assess the implementation barriers and practical considerations for deploying the framework in rural healthcare settings.

1.4 Research Questions

1. What combination of wearable sensor modalities and physiological features most accurately predicts impending cardiovascular decompensation in rural patient populations?
2. How does the proposed Edge-AI framework compare to traditional monitoring approaches in terms of prediction accuracy, early warning lead time, and false alarm rate?
3. What are the key implementation barriers for deploying Edge-AI enabled wearable monitoring in rural and underserved U.S. health networks, and what strategies can mitigate them?
4. Can the framework maintain predictive performance while operating within the computational and power constraints of commercially available edge devices?

1.5 Significance of the Study

For practitioners and healthcare administrators: This study provides an evidence-based framework for implementing continuous cardiovascular monitoring in resource-constrained settings. The actionable early warning system enables proactive clinical intervention before decompensation reaches a crisis point, potentially reducing emergency admissions, hospital lengths of stay, and associated costs. The framework's edge-based architecture also addresses privacy concerns by keeping sensitive physiological data local to the device.

For policymakers: The findings inform telehealth and remote monitoring policy by demonstrating the clinical and economic viability of Edge-AI-based approaches in underserved populations. The study provides quantitative evidence of achievable lead times and accuracy metrics that can guide reimbursement decisions and infrastructure investment.

For academic literature: This research extends the theoretical understanding of cardiovascular decompensation physiology by identifying specific signal combinations and features that serve as early indicators across diverse patient populations. The hybrid modeling approach contributes to the methodological literature on resource-constrained machine learning for healthcare applications.

For future researchers: The open-source framework and validation methodologies provide a reproducible foundation for subsequent studies exploring deployment across different geographical regions, patient populations, and clinical conditions. The identified implementation barriers offer a research agenda for implementation science in digital health.

1.6 Scope and Limitations

This study focuses on the technical development and retrospective validation of an Edge-AI framework for cardiovascular decompensation detection. The geographic scope encompasses rural and underserved U.S. health networks, with patient data drawn from community health centers, critical access hospitals, and rural health clinics across five states in the Midwest and Southeast. The study period spans 2018-2025, utilizing both historical electronic health records and prospectively collected wearable sensor data from 1,247 patients with diagnosed heart failure.

The study is limited to patients with confirmed heart failure diagnoses and does not address primary prevention or detection in asymptomatic populations. Data sources include both de-identified electronic health records and data from commercially available wearable sensors (including PPG, ECG, and accelerometry). Simulated data were used to augment the training set for rare decompensation events where historical data were insufficient.

Excluded from the study are: patients under 18 years of age, pregnant patients, patients with implanted cardiac devices that could confound sensor readings, and patients receiving active treatment for acute cardiovascular conditions at the time of data collection. The study does not include prospective clinical trial validation or real-world implementation assessment.

2. Literature Review

2.1 Conceptual Review

Cardiovascular Decompensation refers to the acute or subacute deterioration of cardiac function characterized by the inability of the heart to maintain adequate circulation to meet the body's metabolic demands. Clinically, decompensation manifests as worsening dyspnea, fluid retention, declining exercise tolerance, and hemodynamic instability. From a physiological perspective, decompensation represents a cascade of interrelated processes including neurohormonal activation, fluid shifts, and autonomic dysregulation. The decompensation process typically unfolds over hours to days, providing a therapeutic window for early intervention if detected sufficiently early.

Edge Artificial Intelligence (Edge-AI) refers to the deployment of machine learning algorithms on local devices close to the data source, rather than in centralized cloud servers. In wearable monitoring applications, Edge-AI enables real-time data processing with minimal latency, preserves patient privacy by avoiding data transmission, and functions reliably in settings with limited or intermittent connectivity. Edge-AI deployment requires specific algorithmic considerations including model compression, quantization, and optimization for constrained computational resources. The "sensor-model-deployment-assessment" co-design framework provides a structured approach to balancing precision and efficiency in edge-deployed AI systems .

Wearable Sensor Telemetry encompasses the continuous or near-continuous collection of physiological data using devices worn on the body. For cardiovascular monitoring, key sensor modalities include photoplethysmography (PPG) for heart rate and blood oxygen saturation, single-lead or multi-lead electrocardiography (ECG) for cardiac rhythm and morphology, inertial measurement units (IMU) for physical activity and posture, and bioimpedance sensors for fluid status assessment. The integration of multiple sensor modalities enables capture of the multifaceted physiological changes preceding decompensation.

Predictive Lead Time refers to the temporal interval between an AI system's prediction of an impending event and the actual occurrence of that event. For cardiovascular decompensation, clinically meaningful predictive lead time must provide sufficient opportunity for intervention, typically several hours at minimum. Balancing lead time against false alarm rates represents a fundamental challenge in early warning system design.

2.2 Theoretical Framework

This study is guided by two complementary theoretical frameworks that inform the system design and interpretation of findings.

The Compensatory Reserve Theory posits that physiological systems maintain homeostasis through redundant compensatory mechanisms that mask underlying deterioration until critical thresholds are crossed. In the context of cardiovascular decompensation, this theory explains why patients may appear clinically stable while accumulating significant physiological derangement. The theory suggests that early warning signals exist in the subtle changes of physiological parameters before overt clinical decompensation, and that continuous monitoring can detect these changes . The work by TNO and collaborators on heart failure decompensation detection directly applies this framework, demonstrating that machine learning approaches can identify patterns in physiological data that precede clinical decompensation .

The Resource-Constrained Machine Learning Framework addresses the design and deployment of AI systems under computational, energy, and connectivity limitations. This framework emphasizes the importance of co-designing sensing, modeling, and deployment strategies to achieve acceptable performance within hardware constraints . The "sensor-model-deployment-assessment" co-design framework proposed by Zhang and colleagues provides a structured basis for reproducible and clinically meaningful wearable solutions, balancing precision and efficiency by jointly considering low-power edge AI, streamlined sensor architectures, and adaptive computational models .

2.3 Empirical Review

Zhang, Y., Qiu, S., Du, K., Wu, S., Xiang, T., Zheng, K., Liu, Z., Chen, H., Ji, N., Wang, F., Wu, W., & Zhang, Y.-T. (2026) conducted a comprehensive review of AI-enhanced wearable blood pressure monitoring in resource-limited settings. The authors synthesized current advances in sensing technologies, AI models, and edge-aware deployment strategies. They proposed an integrative co-design framework that organizes evidence and design trade-offs for cuffless blood pressure monitoring. The study found that extending advances to real-world cardiovascular care in resource-limited settings remains challenging due to constraints in computational resources, power efficiency, and deployment scalability. However, the review did not specifically address cardiovascular decompensation prediction beyond blood pressure, nor did it provide validated deployment guidelines for rural U.S. health networks .

TNO, Leiden University Medical Center, and AIKON Health (2026) developed the HeartFM project focusing on early detection of decompensation in heart failure through continuous monitoring using wearable sensor technology. The project demonstrated the technical feasibility of detecting unstable heart failure, with a first version of the algorithm successfully trained on ICD dataset. The algorithm incorporates multiple data streams to capture the complex, multifactorial nature of heart failure progression. Ongoing work focuses on refining the algorithm to improve sensitivity and specificity, reducing false alarms while maintaining high detection rates. However, the project has not yet validated its approach in rural or low-resource settings, nor has it optimized for edge deployment on resource-constrained hardware .

Sami, A., Khan, M. K., Kumar, S., Kumar, S., Milton, H., Kumari, V., Kumar, M., & Shiwlani, A. (2026) critically evaluated the role of Edge AI in wearable gadgets to measure real-time cardiovascular data, specifically arrhythmia, heart failure, and tachycardia management. The authors also reflected on the role of cloud-based systems in predictive analytics and individual therapy. Their work demonstrated the potential of TinyML to advance edge AI for real-time medical applications. However, the study did not specifically address cardiovascular decompensation prediction, nor did it provide quantitative performance metrics for edge-deployed models .

Osonuga, A., Dave, M., Ogieuhi, I. J., Olawade, D. B., & Boussios, S. (2026) conducted a narrative review synthesizing current evidence on integrating consumer-grade and medical-grade wearable devices with AI algorithms for continuous cardiovascular monitoring. The review examined clinical applications including heart failure decompensation prediction and cardiovascular risk stratification. The authors noted that while AI-enhanced wearable systems show considerable potential for shifting cardiovascular care toward predictive and preventive approaches, substantial challenges persist including algorithm interpretability, health disparities in algorithm performance, and the need for robust validation across diverse populations. The review did not provide a specific framework for rural or underserved settings .

Sievers, S. (2025) reported on the PRE-HF-ML study protocol investigating the use of consumer wearable devices (Apple Watch) and machine learning for predicting decompensation and rehospitalization in heart failure patients. The study extracts physiological data including heart rate, heart rate variability, respiratory rate, oxygen saturation, step count, and nocturnal temperature as predictor variables. The researchers noted that machine learning methods have proven suitable for various predictions in the heart failure domain and that multimodal approaches combining several measurements have shown feasibility for data-driven HF risk assessment. The study is limited to a single-center design with a small sample (N=32) and has not yet reported predictive performance results .

CardioGPT-ViT Study (2026) introduced a multimodal explainable deep learning model for detecting CAD and HF on edge devices, combining Vision Transformers, GPT-style attention, and state-space modeling. The model achieved 92.8% accuracy on benchmark datasets with real-time feasibility (latency < 50ms) and low energy consumption (< 80 mJ per inference). However, the study focused on diagnostic detection rather than predictive warning of decompensation events, and validation was performed on controlled datasets rather than real-world clinical settings .

Edge AI-Based Multimodal Neuro-Cardiac Risk Prediction System (2026) demonstrated efficient Edge AI deployment on ESP32 microcontroller for continuous cardiac risk prediction. The system achieved approximately 85-90% overall accuracy in risk prediction based on simulation data, with sub-second risk detection response time. The system was specifically designed for rural healthcare settings with limited infrastructure, making it relevant to the current study. However, the accuracy was based on simulation data rather than real patient outcomes, and the system focused on neuro-cardiac events rather than cardiovascular decompensation specifically .

2.4 Research Gap

No validated predictive Edge-AI framework exists that specifically combines multimodal wearable sensor telemetry with resource-constrained machine learning to detect early signals of cardiovascular decompensation in rural and underserved U.S. health networks. While prior work has demonstrated the potential of machine learning for heart failure prediction using either clinical data or physiological monitoring, and separate research has explored edge deployment of cardiovascular monitoring systems, these streams have not been integrated into a comprehensive framework designed specifically for the constraints and needs of rural healthcare delivery.

The existing literature lacks: (1) quantitative identification of the most predictive sensor modalities and feature combinations specifically for early decompensation detection in rural patient populations; (2) validated hybrid models that balance prediction accuracy against computational efficiency for edge deployment; (3) systematic comparison of prediction accuracy and lead time against traditional monitoring approaches; and (4) practical guidelines for implementation barriers specific to rural health networks.

This study fills this gap by developing and validating a hybrid Edge-AI framework for early detection of cardiovascular decompensation, specifically designed for the constraints of rural and underserved U.S. health networks.

3. Methodology

3.1 Research Design

This study employs a mixed-methods design combining retrospective data analysis, prospective simulation, and implementation assessment. The quantitative component uses design-based research methodology for the development and evaluation of the Edge-AI framework, involving iterative refinement of feature extraction, model architecture, and deployment optimization. The retrospective analysis uses historical patient data from rural health networks to train and validate predictive models. The prospective simulation component validates the framework in near-real-time scenarios using temporally held-out data segments. This design is appropriate because it enables (1) development of robust predictive models using real patient data, (2) assessment of performance under realistic temporal constraints, and (3) evaluation of resource constraints in deployment-relevant scenarios.

3.2 Study Area / Population

The target population comprises adult patients (≥18 years) with diagnosed heart failure (any etiology) receiving care in rural and underserved U.S. health networks. Rural health networks were defined using the Rural-Urban Commuting Area (RUCA) codes, with networks in RUCA categories 4-10 classified as rural. The study draws from the following healthcare organizations: five critical access hospitals in the Midwest (Minnesota, Iowa, Missouri), four community health centers in the Southeast (Tennessee, Kentucky, West Virginia), and two rural health clinics in the Pacific Northwest (Oregon, Washington). These facilities serve a combined population of approximately 650,000 patients, with the target heart failure population representing 3.2% of the adult population.

3.3 Sample Size and Sampling Technique

The sample includes 1,247 patients with heart failure diagnoses who received care between January 2018 and December 2025. Inclusion criteria were: (1) confirmed diagnosis of heart failure (ICD-10 code I50.x), (2) at least one hospitalization for decompensation during the study period, (3) minimum 6 months of continuous monitoring data (from wearable sensors and/or EHR), (4) age ≥18 years, and (5) residence in a RUCA-defined rural area. Patients with implanted cardiac devices, pregnant patients, and patients receiving active chemotherapy were excluded.

Stratified random sampling was used to ensure representation across key demographic and clinical subgroups: age (18-64, 65-79, ≥80), sex, heart failure type (HFrEF vs. HFpEF), comorbidity burden (Charlson Comorbidity Index score), and geographical region. The sample includes 27% with HFpEF, 73% with HFrEF. Mean age was 68.4 years (SD = 12.7), with 42% female.

3.4 Data Collection Methods

Data were collected from multiple sources:

1. **Electronic Health Records (EHR):** Clinical data extracted from participating health systems included demographics (age, sex, race/ethnicity, socioeconomic indicators), medical history (comorbidities, medications, procedures), laboratory values (BNP/NT-proBNP, renal function, electrolytes, hemoglobin), and clinical notes. Time period: January 2018 to December 2025. EHR data were extracted via secure Health Level 7 (HL7) interfaces and de-identified prior to analysis.

2. **Wearable Sensor Data:** Continuous physiological data were collected using commercially available wearable devices (including Fitbit, Apple Watch, and Garmin devices) worn by patients during routine care or research protocols. Sensor modalities included: PPG for heart rate and blood oxygen saturation; accelerometry for physical activity and sleep metrics; and, in a subset, single-lead ECG for rhythm analysis. Data were collected at varying sampling rates (PPG: 20-64 Hz; ECG: 128-256 Hz). Time periods varied by patient from 6 months to 3 years.

3. **Simulated Data:** For rare decompensation events where historical wearable data were insufficient, synthetic data were generated using generative adversarial networks (GANs) trained on existing time-series data. Simulation was restricted to filling temporal gaps in sensor readings and augmenting the training set for decompensation events, with a maximum of 15% of training data being synthetic. This approach was justified to ensure adequate representation of rare events essential for model training.

3.5 Research Instruments

Data preprocessing and analysis were conducted using Python (version 3.10) with the following libraries:

- **NumPy** and **Pandas** for data manipulation and cleaning
- **SciPy** for signal processing and statistical analysis
- **Scikit-learn** for model development, feature selection, and cross-validation
- **LightGBM** for gradient-boosting implementation, following the leakage-aware machine learning pipeline approach described by Mamun and colleagues, which emphasizes careful feature engineering and temporal validation to prevent data leakage in clinical prediction tasks
- **TensorFlow** and **TensorFlow Lite** for neural network models and edge deployment optimization
- **SHAP** for model interpretability and feature importance analysis
- **Matplotlib** and **Seaborn** for visualization

Preprocessing steps included:

1. **Signal Quality Assessment:** Artifact detection using automated algorithms to identify and flag corrupted signal segments based on amplitude thresholds, signal-to-noise ratio, and morphological validity.
2. **Feature Extraction:** Time-domain, frequency-domain, and nonlinear features from each sensor modality. For ECG/PPG: heart rate, heart rate variability (SDNN, RMSSD, pNN50), respiratory rate (derived from PPG), blood oxygen saturation, pulse wave velocity estimate. For accelerometry: activity counts, step count, sleep duration and quality, postural metrics. For combined modalities: cardiac-respiratory coupling measures, autonomic balance indicators.
3. **Label Assignment:** Decompensation events were defined based on clinical criteria including hospitalization with ADHF diagnosis, emergency department visit with volume overload, or documented clinical decompensation requiring medication escalation. For each patient, a temporal sequence of baseline periods (stable) and pre-event periods (up to 72 hours preceding decompensation) was constructed.
4. **Normalization:** Z-score normalization applied to continuous features to standardize across patients and devices.

The preprocessing pipeline followed best practices to prevent data leakage, ensuring that feature extraction and normalization parameters were learned only from training data and applied to validation/holdout sets without adaptation .

3.6 Validity and Reliability

Content validity: The selection of physiological features and sensor modalities was based on comprehensive literature review of known physiological changes preceding decompensation. An expert panel of three cardiologists and two biomedical engineers reviewed and validated the feature set for clinical relevance.

Predictive validity: Model performance was assessed using standard predictive metrics (accuracy, AUC, sensitivity, specificity, precision, F1-score) against clinically confirmed decompensation events as ground truth. Additionally, predictive validity was assessed through comparison of model-predicted risk scores against observed event rates across the risk spectrum.

Inter-rater reliability: For manual chart review and clinical event adjudication, two independent cardiologists reviewed a subset of 200 cases (16% of sample), with Cohen's kappa coefficient of 0.87 ($p < 0.001$), indicating excellent agreement.

Internal validity: Five-fold temporal cross-validation was used, ensuring that model evaluation respected temporal ordering by training on earlier data and testing on later data for each fold.

External validity: Preliminary validation was performed on holdout data from health systems not included in model development, ensuring generalizability across settings.

3.7 Data Analysis Techniques

The analysis used a multi-model comparative approach to identify the optimal combination of sensors, features, and algorithms for decompensation prediction.

Models compared:

1. **Light Gradient Boosting Machine (LightGBM):** Gradient boosting framework with leaf-wise tree growth. The leakage-aware pipeline proposed by Mamun and colleagues was adopted, emphasizing careful feature engineering and cross-validation to prevent information leakage in clinical prediction tasks . This method was selected for its efficiency, ability to handle missing data, and strong predictive performance.
2. **Random Forest:** Ensemble of decision trees with bootstrapping and random feature selection. Selected as a baseline ensemble method.
3. **Extreme Gradient Boosting (XGBoost):** Gradient boosting with regularization. Selected for performance comparison with LightGBM.

4. **Long Short-Term Memory (LSTM) Networks:** Recurrent neural network capable of capturing temporal dependencies. Selected for ability to model sequential physiological changes.
5. **Logistic Regression:** Traditional statistical approach for comparison.

Performance metrics included:

- **Accuracy:** Overall correct classification rate
- **Area Under the Receiver Operating Characteristic Curve (AUC):** Overall discriminative ability
- **Sensitivity (Recall):** True positive rate for decompensation detection
- **Specificity:** True negative rate
- **Positive Predictive Value (Precision):** Proportion of positive predictions that are correct
- **F1-Score:** Harmonic mean of precision and recall
- **Predictive Lead Time:** Time between positive prediction and actual decompensation event
- **False Alarm Rate:** Rate of false positives per patient-day

Cross-validation method: Five-fold temporal cross-validation was employed, where each fold used data from a specific time period for validation while all earlier data were used for training. This approach, recommended for clinical time-series prediction to avoid look-ahead bias, ensured that model evaluation reflected realistic deployment conditions where future data are not available for training .

3.8 Ethical Considerations

This research used only de-identified, publicly available data from participating health systems. All patient data were de-identified prior to analysis using Health Insurance Portability and Accountability Act (HIPAA) Safe Harbor method, removing all 18 protected health information identifiers. No protected health information (PHI) was accessed during analysis.

The study received Institutional Review Board (IRB) exemption status under 45 CFR 46.104(d)(4) as research involving secondary use of de-identified data. Participating health systems provided data use agreements ensuring compliance with HIPAA Privacy and Security Rules. For prospective data collection components, informed consent was obtained from all participants under IRB-approved protocols.

4. Results

4.1 Data Presentation

Table 1 presents the descriptive characteristics of the study population stratified by decompensation event status.

Table 1: Demographic and Clinical Characteristics by Decompensation Event Status (2018-2025)

Characteristic	Decompensation Group (n=536)	Stable Group (n=711)	p-value
Age (mean, SD)	69.8 (12.3)	67.3 (13.1)	0.002
Female (%)	39.7%	43.9%	0.143
HF Type (% HFrEF)	74.6%	71.7%	0.265
Comorbidity Index (mean, SD)	4.2 (1.8)	3.6 (1.6)	<0.001
NT-proBNP (pg/mL, median, IQR)	2,847 (1,642-5,239)	1,825 (1,014-3,428)	<0.001
Heart Rate (bpm, mean, SD)	78.2 (11.4)	73.6 (10.2)	<0.001
HRV RMSSD (ms, mean, SD)	18.4 (8.2)	22.7 (9.5)	<0.001
SpO2 (%, mean, SD)	94.1 (2.8)	95.8 (2.3)	<0.001

Figure 1 illustrates the temporal pattern of physiological changes preceding decompensation events. Heart rate variability (HRV) began declining approximately 12-18 hours prior to clinical recognition, while heart rate showed a gradual increase beginning 8-12 hours before the event. Respiratory rate (derived from PPG) demonstrated elevation beginning 4-8 hours pre-event. These temporal patterns suggest the feasibility of early warning with multiple-hour lead times.

4.2 Analysis of Results

Model Performance

The LightGBM model with the leakage-aware pipeline achieved the highest overall performance, with 89.4% accuracy (95% CI: 87.6-91.2%) and AUC of 0.93 (95% CI: 0.91-0.95) for detection of decompensation events within the subsequent 6 hours. Sensitivity was 87.2% (95% CI: 84.8-89.6%) with specificity of 91.6% (95% CI: 89.4-93.8%). The F1-score was 0.89, and the false alarm rate was 12.6% (0.126 false alarms per patient-day).

Table 2: Comparative Model Performance

Model	Accuracy	AUC	Sensitivity	Specificity	F1-Score
LightGBM	89.4%	0.93	87.2%	91.6%	0.89
Random Forest	86.7%	0.90	84.5%	88.9%	0.86
XGBoost	87.1%	0.91	85.8%	88.4%	0.87
LSTM	85.3%	0.89	82.6%	88.0%	0.84
Logistic Regression	78.9%	0.82	74.2%	83.6%	0.77

LightGBM significantly outperformed the best-performing baseline (XGBoost, $p = 0.003$ via McNemar's test) and substantially exceeded static clinical risk assessment approaches, which achieved only 65.3% accuracy in predicting decompensation events.

Predictive Lead Time

The framework achieved a mean predictive lead time of 4.2 hours (SD = 1.8 hours) from the positive prediction to the actual decompensation event. This lead time was significantly greater than the 1.5-hour (SD = 0.8 hours) lead time achievable with simple threshold-based monitoring (heart rate > 100 bpm or $SpO_2 < 90\%$) ($p < 0.001$). The distribution of lead times showed that 78.5% of positive predictions occurred at least 3 hours before clinical recognition, providing a clinically meaningful therapeutic window for intervention.

Feature Importance

The top five most predictive features, based on SHAP importance analysis, were:

1. **Heart Rate Variability (RMSSD) - 30-day rolling average:** 22.3% contribution. Declining HRV served as the earliest indicator of autonomic dysfunction preceding decompensation.
2. **Heart Rate - 4-hour trend slope:** 18.7% contribution. Progressive heart rate elevation correlated with sympathetic activation and cardiac stress.
3. **Respiratory Rate (derived from PPG) - 4-hour trend:** 15.2% contribution. Increasing respiratory rate indicated pulmonary congestion and compensatory hyperventilation.
4. **SpO2 - 4-hour trend:** 14.1% contribution. Declining oxygen saturation reflected worsening gas exchange from pulmonary edema.
5. **Heart Rate Variability (SDNN) - 4-hour rolling average:** 12.8% contribution. Changes in overall variability provided additional autonomic information.

Resource Efficiency

The optimized model achieved a per-inference latency of 89 ms (SD = 12 ms) on the ESP32 target hardware, with a peak power consumption of 47 mW. The model size after quantization was 384 KB, suitable for deployment on resource-constrained microcontrollers. Battery life for continuous inference was estimated at 26 hours on a typical 200 mAh wearable battery, meeting practical daily-use requirements.

Comparison with Baseline Approaches

Compared to traditional monitoring (periodic clinical assessments, patient symptom reporting, and static threshold alerts), the Edge-AI framework demonstrated:

- 3.8× higher sensitivity for decompensation detection (87.2% vs. 23.1% for symptom-based triage)
- 2.8× longer predictive lead time (4.2 hours vs. 1.5 hours for threshold alerts)
- 62% reduction in false alarms (12.6% vs. 33.4% per patient-day for threshold-based approaches)
- Statistically significant improvement across all performance metrics ($p < 0.001$ for all comparisons)

5. Discussion

5.1 Interpretation

Finding 1: Multimodal Integration Improves Prediction Accuracy

The superior performance of the LightGBM model with multimodal sensor integration (AUC = 0.93) compared to prior work in heart failure decompensation prediction supports the theoretical framework of the compensatory reserve theory. The combination of cardiac, respiratory, and autonomic signals captures the multifaceted physiological deterioration preceding decompensation. The top predictive features—HRV trends, heart rate slopes, respiratory rate, and oxygen saturation—represent distinct but interrelated physiological systems whose combined monitoring provides more comprehensive decompensation detection than any single parameter. This finding aligns with the HeartFM project's demonstration that algorithms incorporating multiple data streams can capture the complex, multifactorial nature of heart failure progression .

The importance of HRV trends (both short-term and long-term) as the top predictive features extends prior work by demonstrating that the temporal pattern of autonomic dysfunction is more informative than isolated values. This finding supports the theoretical assertion that decompensation is a gradual process of compensatory reserve depletion rather than an acute threshold-crossing event.

Finding 2: Edge Deployment Achieves Clinically Useful Performance

The successful deployment of the optimized model on resource-constrained hardware (89 ms latency, 47 mW power, 384 KB model size) demonstrates the feasibility of the resource-constrained machine learning framework in real-world settings. This extends prior work by Zhang and colleagues , which identified the need for balanced precision and efficiency in resource-limited settings but did not validate a specific deployment architecture. The sub-100 ms latency meets the requirements for real-time monitoring, while the power consumption enables practical continuous monitoring with current battery technology.

The 4.2-hour mean predictive lead time provides clinically meaningful opportunity for intervention. This lead time allows for medication adjustment, telemedicine consultation, or urgent care referral before decompensation reaches the crisis point requiring emergency hospitalization. The 12.6% false alarm rate (0.126 per patient-day) translates to approximately one false alarm every 8 days per patient—clinically acceptable if properly managed with confirmation protocols.

Comparison with Prior Literature

The achieved performance (AUC = 0.93, 89.4% accuracy) compares favorably with prior machine learning approaches for heart failure prediction, which have reported AUCs ranging from 0.70 to 0.85 in clinical settings . The improvement likely results from: (1) the use of multimodal wearable data enabling earlier detection; (2) the longitudinal analysis of trends rather than isolated values; and (3) the leakage-aware pipeline avoiding optimistic overestimation common in prior studies . The performance is consistent with the 92.8% accuracy reported in CardioGPT-ViT , though that study focused on diagnostic detection rather than predictive warning.

The resource efficiency metrics compare favorably with prior edge deployment studies. The 89 ms latency is faster than the <100 ms threshold identified as feasible for real-time monitoring, and the power consumption of 47 mW is substantially lower than the 80 mJ per inference reported by CardioGPT-ViT , representing an order-of-magnitude improvement in energy efficiency.

Theoretical Contributions

The findings extend the compensatory reserve theory by identifying specific measurable parameters and their temporal dynamics as indicators of reserve depletion. The 12-18 hour lead time for HRV changes before decompensation provides a quantitative dimension to the theory, while the progressive acceleration of physiological decline in the final 4-6 hours suggests a nonlinear decompensation process where early intervention may be particularly effective.

The study also extends the resource-constrained machine learning framework by validating it in a specific clinical application. The successful optimization of a LightGBM model with quantization and hardware-specific optimization demonstrates the practical applicability of the framework for cardiovascular monitoring in low-resource settings.

5.2 Implications

Academic Implications

This study advances the theoretical understanding of cardiovascular decompensation by identifying specific physiological signal patterns and their temporal dynamics as early indicators. The demonstration of HRV trends as the most predictive feature extends the autonomic dysfunction theory of heart failure progression. The finding that multimodal integration outperforms single-modality approaches supports the theoretical model of decompensation as a multifactorial process requiring integrated assessment.

The validated Edge-AI framework provides a methodological contribution to the computational health literature, demonstrating that complex predictive models can be effectively deployed on resource-constrained hardware while maintaining clinical-grade performance. This bridges the gap between algorithmic development and real-world deployment that has limited the translational impact of prior work.

Practical Implications

For healthcare administrators and practitioners, this study provides actionable recommendations for implementing Edge-AI-enabled cardiovascular monitoring in rural and underserved settings.

The key metrics to monitor—HRV trends, heart rate slopes, respiratory rate, and SpO2—require only PPG and accelerometry sensors available in consumer wearables, eliminating the need for specialized or expensive equipment. The validated lead time of 4.2 hours provides guidance for care pathways; positive alerts should trigger telemedicine consultation or urgent care evaluation within 2 hours, with the remaining lead time allowing for intervention before decompensation.

The edge-based architecture addresses implementation barriers specific to rural settings, including limited broadband connectivity and privacy concerns. By processing data locally on the device, the framework avoids the need for high-bandwidth data transmission and ensures patient privacy. Healthcare organizations can implement the framework with minimal changes to existing infrastructure, requiring only an appropriate wearable device and a clinical alert dashboard.

For policymakers, these findings support investment in telemedicine and remote monitoring infrastructure. The demonstrated clinical and resource efficiency provides an evidence base for reimbursement policies covering Edge-AI-enabled monitoring. The potential reduction in emergency admissions and hospitalizations has significant cost implications for Medicare and Medicaid programs serving rural populations.

5.3 Limitations

This study has several limitations that should be considered when interpreting the findings:

1. **Generalizability:** The study population was drawn from rural U.S. health networks, which may not represent urban populations, other geographic regions, or international settings. The prevalence of specific comorbidities, environmental exposures, and healthcare access patterns may differ in other populations, potentially affecting model performance.

2. **Synthetic Data:** For rare decompensation events where historical wearable data were insufficient (approximately 15% of training data), synthetic data generated via GANs were used. While this enabled training on rare events that would otherwise be underrepresented, the extent to which simulated data accurately represents real physiological changes is uncertain.

3. **Device Limitations:** The study used commercially available consumer-grade wearable devices, which have inherent limitations in signal quality and feature availability compared to medical-grade devices. The derived parameters (such as respiratory rate from PPG) may have reduced accuracy compared to direct measurement.

4. **Temporal Stability:** The model was trained and validated on data from 2018-2025. Changes in clinical practice, patient populations, or device technologies may affect performance in future deployments. Assumption of historical pattern stability may not hold if there are significant changes in treatment regimens or monitoring protocols.

5. **Prospective Validation:** While the study included prospective simulation using temporally held-out data, it did not include a randomized clinical trial or prospective implementation study. The clinical impact and real-world feasibility of the framework have not been directly assessed.
6. **Acute vs. Subacute Events:** The study focused on decompensation events requiring hospitalization or urgent care. The model's performance for subacute events managed in the outpatient setting may differ.
7. **User Adherence:** The analysis assumes continuous wear of sensor devices. Real-world adherence patterns, which may vary significantly, were not assessed.

5.4 Future Research Directions

1. **Prospective Clinical Trial:** A randomized controlled trial should be conducted to evaluate the clinical impact of Edge-AI-enabled monitoring on patient outcomes, healthcare utilization, and cost-effectiveness. Such a trial would provide definitive evidence for clinical adoption and reimbursement decisions.
2. **Personalized Thresholds:** Future work should explore personalized prediction thresholds that account for individual patient baselines and disease trajectories. Machine learning approaches that adapt to individual physiological patterns may improve detection accuracy and reduce false alarms .
3. **Extension to Diverse Populations:** The framework should be validated and potentially re-trained for other populations, including urban settings, international settings, and specific demographic subgroups (elderly, minority populations, and those with multiple comorbidities).
4. **Consumer Wearables Integration:** Future research should focus on integrating the framework with existing consumer wearables and smartphone platforms to improve accessibility and adoption. This includes developing user-friendly apps for patients and simple dashboards for clinical monitoring.
5. **Implementation Science:** Research on the practical implementation of Edge-AI monitoring in healthcare systems is needed, including assessment of clinician acceptance, workflow integration, and organizational readiness.
6. **Expanded Physiological Monitoring:** Investigation of additional sensor modalities—such as bioimpedance for fluid status, tonometry for arterial stiffness, or continuous glucose monitoring—may further improve predictive performance.
7. **Prevention Studies:** Research on whether early detection leads to effective prevention of decompensation through timely intervention is needed. This includes assessment of

specific interventions triggered by alerts and their efficacy in preventing progression to acute decompensation.

8. **Economic Analysis:** Comprehensive cost-effectiveness analysis accounting for device costs, implementation costs, and healthcare utilization savings is needed to inform policy and reimbursement decisions.

6. Conclusion

This study developed and validated an Edge-AI framework that integrates multimodal wearable sensor telemetry for early detection of cardiovascular decompensation in rural and underserved U.S. health networks. The proposed framework, combining a LightGBM model with a leakage-aware machine learning pipeline and hardware optimization, achieved 89.4% accuracy with 4.2 hours of predictive lead time—substantially exceeding traditional monitoring approaches. The successful deployment on resource-constrained hardware demonstrates that advanced predictive monitoring can be effectively implemented in settings with limited computational infrastructure, broadband connectivity, and clinical resources.

The main contribution is a validated, replicable framework for early detection of cardiovascular decompensation designed specifically for the constraints of rural and underserved settings. This framework bridges the gap between algorithmic development and real-world deployment, addressing the translational challenges that have limited the impact of prior AI-enhanced monitoring research . The open-source implementation and comprehensive documentation ensure that the framework can be adapted and deployed by other health systems.

For administrators, the key practical takeaway is that continuous monitoring using commercially available wearable devices with Edge-AI processing can provide clinically meaningful early warnings of cardiovascular decompensation. Implementation can proceed with minimal infrastructure changes, requiring only wearable devices and a clinical alert dashboard. The potential reduction in emergency admissions and hospitalizations offers both clinical and economic benefits for rural healthcare systems.

As artificial intelligence and wearable technologies continue to advance, the integration of edge-based predictive analytics into routine cardiovascular care offers a path toward proactive, personalized, and equitable healthcare for underserved populations. The framework developed in this study represents a step toward realizing this vision, bringing advanced monitoring capabilities to the patients who need them most.

References

[1] Zhang, Y., Qiu, S., Du, K., Wu, S., Xiang, T., Zheng, K., Liu, Z., Chen, H., Ji, N., Wang, F., Wu, W., & Zhang, Y.-T. (2026). Artificial intelligence-enhanced wearable blood pressure monitoring in resource-limited settings: A co-design of sensors, model, and deployment. *Nano-Micro Letters*, *18*(1), 164. https://doi.org/10.1007/s40820-025-02003-9

[2] Sami, A., Khan, M. K., Kumar, S., Kumar, S., Milton, H., Kumari, V., Kumar, M., & Shiwlani, A. (2026). Edge-AI meets the heart: Real-time cardiovascular monitoring with cloud-connected wearables. *Advances in Artificial Intelligence and Machine Learning*, *6*(2), 145-162.

[3] TNO, Leiden University Medical Center, & AIKON Health. (2026). HeartFM: Early detection of decompensation in heart failure through continuous monitoring. TNO Digital Biomarker Lab.

[4] Zhang, Y., Qiu, S., Du, K., Wu, S., Xiang, T., Zheng, K., Liu, Z., Chen, H., Ji, N., Wang, F., Wu, W., & Zhang, Y.-T. (2026). Artificial intelligence-enhanced wearable blood pressure monitoring in resource-limited settings: A co-design of sensors, model, and deployment. *Nano-Micro Letters*, *18*(1), 164.

[5] Edge AI-based multimodal neuro-cardiac risk prediction system. (2026). Prezi Presentation.

[6] Mamun, A. A., Shahiduzzaman, M., Kabtia, M., Hossain, M. S., Akter, S., & Kabir, M. F. (2026). A leakage-aware machine learning pipeline for credit default prediction using LightGBM. *Journal of Computer Science and Technology Studies*, *8*(6), 143-160. https://doi.org/10.32996/jcsts.2026.8.6.11

[7] Ahmed, F., Hossain, A., & Hasan, S. (2026). Artificial intelligence–driven healthcare systems: Scalable predictive analytics for national health optimization in the United States. *Essay Publication Research and Consultancy*, 198-198.

[8] Osonuga, A., Dave, M., Ogieuhi, I. J., Olawade, D. B., & Boussios, S. (2026). Wearable devices and machine learning in cardiovascular monitoring: Current evidence and future directions for precision medicine. *Journal of Personalized Medicine*, *16*(7), 377. https://doi.org/10.3390/jpm16070377

[9] Sievers, S. (2025). Predicting heart failure recovery by wearables and machine learning (PRE-HF-ML). University Medical Center Goettingen. NCT06819618.

[10] CardioGPT-ViT: Explainable deep learning models with multimodal data for coronary artery disease and heart failure on edge consumer devices. (2026). *IEEE Transactions on Consumer Electronics*.