Integrating Multimodal Deep Learning Architectures and Wearable Biosensor Telemetry for Real-Time Prediction of Adverse Cardiovascular Outcomes in Post-Discharge Ischemic Heart Disease Patients

Authors

Scott Satter Lee, Jack Richard, Mario Alverdin, Abilly Elly

Date: June 22, 2026

Abstract

Ischemic heart disease (IHD) remains a leading cause of global morbidity and mortality, with post-discharge patients facing a persistently high risk of adverse cardiovascular outcomes. Current risk stratification tools, predominantly reliant on episodic clinical assessments, fail to capture the dynamic physiological changes that precede acute events. This prospective cohort study addresses this critical gap by integrating multimodal deep learning architectures with continuous wearable biosensor telemetry to enable real-time prediction of major adverse cardiovascular events (MACE) in post-discharge IHD patients. A hybrid deep learning

framework combining Convolutional Neural Networks (CNNs) for spatial feature extraction from electrocardiogram (ECG) and photoplethysmography (PPG) waveforms with a Temporal Attention-Gated Recurrent Unit (TA-GRU) for capturing temporal dependencies was developed and validated on a prospective cohort of 3,247 patients monitored over eight months post-discharge . The proposed framework achieved exceptional predictive performance with an area under the curve (AUC) of 0.89, surpassing traditional machine learning models (AUC = 0.86) and conventional risk scores (AUC = 0.81) . The model demonstrated 89.4% sensitivity at a specificity of 85.2%, with a mean early warning lead time of 4.7 hours prior to clinical manifestation. This research provides a validated framework for transforming continuous physiological data into actionable clinical intelligence, supporting the paradigm shift from reactive to predictive cardiovascular care.

Keywords: Multimodal Deep Learning, Wearable Biosensors, Cardiovascular Prediction, Ischemic Heart Disease, Real-Time Monitoring, Edge Computing

1. Introduction

1.1 Background

Cardiovascular diseases (CVDs) remain the leading cause of mortality worldwide, accounting for an estimated 17.9 million deaths annually . Ischemic heart disease (IHD), resulting from reduced blood supply to the heart muscle, represents the most prevalent form of CVD and imposes a substantial burden on healthcare systems globally. Patients discharged following an acute IHD event, such as myocardial infarction or unstable angina, face a particularly elevated risk of recurrent adverse outcomes, including subsequent myocardial infarction, heart failure exacerbation, and sudden cardiac death. The post-discharge period, often characterized by physiological instability and medication adjustments, represents a critical window for intervention and risk mitigation.

Traditional approaches to post-discharge risk stratification rely on clinical assessments conducted during follow-up visits, utilizing tools such as the Global Registry of Acute Coronary Events (GRACE) risk score or the Thrombolysis in Myocardial Infarction (TIMI) risk score. These tools, while valuable, suffer from fundamental limitations: they capture only episodic, static snapshots of patient health and cannot account for the dynamic, rapidly evolving physiological changes that precede acute events. The clinical need for continuous, real-time monitoring capable of detecting early warning signs of decompensation has driven interest in wearable biosensor technologies and artificial intelligence (AI).

Recent advances in wearable biosensor technology have enabled the non-invasive, continuous collection of multiple physiological signals, including electrocardiography (ECG), photoplethysmography (PPG), bioimpedance spectroscopy, and accelerometry. These multimodal data streams encode complementary and overlapping information about cardiovascular function . Simultaneously, deep learning has emerged as a transformative approach for analyzing complex physiological signals, demonstrating superior performance in pattern recognition and predictive modeling across diverse cardiovascular applications . The convergence of these technological advances creates an unprecedented opportunity to develop intelligent monitoring systems capable of real-time prediction of adverse cardiovascular outcomes.

1.2 Problem Statement

Despite significant progress in cardiovascular care, a critical gap exists in the ability to predict adverse outcomes in post-discharge IHD patients. Current clinical practice relies on periodic assessments that cannot capture the subtle, time-varying physiological patterns that indicate impending decompensation. Several challenges have hindered the development of effective predictive monitoring systems:

First, traditional risk stratification models are limited by their reliance on a narrow set of clinical variables measured at a single time point. These models fail to incorporate the rich temporal dynamics present in continuous physiological data, resulting in limited predictive accuracy and high false alarm rates . Second, while wearable devices can collect extensive physiological data, current systems predominantly employ simple threshold-based algorithms that generate numerous false positives and cannot leverage the complex patterns present in multimodal data. Third, existing deep learning approaches for cardiovascular prediction have predominantly focused on single-modality analysis, missing the synergistic information available when multiple physiological signals are jointly analyzed. Research has demonstrated that multimodal AI approaches can identify approximately 13.0-19.3% more significant features than unimodal approaches . Fourth, the computational demands of deep learning models have limited their deployment on resource-constrained wearable devices, necessitating cloud processing that introduces unacceptable latency and privacy concerns .

The specific unsolved issue addressed by this research is: **How can multimodal deep learning architectures be integrated with continuous wearable biosensor telemetry to enable accurate, low-latency, real-time prediction of adverse cardiovascular outcomes in post-discharge IHD patients, while maintaining operational feasibility on resource-constrained edge devices?**

1.3 Objectives of the Study

General Objective:
To develop and validate a multimodal deep learning framework that integrates wearable

biosensor telemetry for real-time prediction of major adverse cardiovascular events (MACE) in post-discharge ischemic heart disease patients.

Specific Objectives:

1. To identify the most predictive multimodal physiological features (ECG, PPG, heart rate variability, and activity patterns) for early detection of adverse cardiovascular outcomes in post-discharge IHD patients.
2. To design a hybrid deep learning architecture combining convolutional neural networks (CNNs) and temporal attention-based recurrent networks that can effectively process multimodal time-series physiological data for real-time risk prediction.
3. To validate the predictive performance, lead time, and operational feasibility of the proposed framework on a prospective cohort of 3,247 post-discharge IHD patients.
4. To compare the proposed multimodal deep learning framework against existing risk prediction methods, including traditional clinical risk scores (GRACE, TIMI) and unimodal machine learning models.

1.4 Research Questions

Research Question 1: What combination of physiological modalities and temporal features most accurately predicts impending adverse cardiovascular outcomes in post-discharge IHD patients?

Research Question 2: How does the proposed multimodal deep learning framework (CNN + TA-GRU) compare to traditional risk scores and unimodal models in terms of prediction accuracy, sensitivity, specificity, and early warning lead time?

Research Question 3: What are the key implementation barriers and design considerations for deploying the proposed framework on resource-constrained wearable devices in real-world clinical settings?

1.5 Significance of the Study

This research addresses a critical clinical need with substantial implications for multiple stakeholders. The study's significance is fourfold:

For Clinicians and Healthcare Administrators: The proposed framework offers an objective, continuous risk stratification tool that can alert clinicians to impending deterioration hours before clinical manifestation. This early warning capability enables timely interventions, potentially reducing hospital readmissions, intensive care unit (ICU) admissions, and mortality. The integration with wearable devices supports remote patient management, reducing the burden on healthcare facilities and enabling personalized care delivery.

For Policymakers: This research provides evidence to support the integration of AI-powered remote monitoring into national healthcare strategies. The demonstrated reduction in preventable readmissions and adverse events addresses key healthcare quality metrics and cost-containment priorities. Findings can inform regulatory frameworks for AI-enabled medical devices and reimbursement policies for remote patient monitoring.

For Academic Literature: The study contributes to the growing body of knowledge on multimodal deep learning for healthcare applications by providing a validated framework specifically designed for continuous cardiovascular monitoring. The comparative analysis against traditional methods establishes a benchmark for future research in this domain.

For Future Researchers: The study provides a reproducible framework, including model architecture specifications, preprocessing pipelines, and evaluation protocols, that can be extended or adapted for other physiological monitoring applications. The identification of key predictor variables and temporal patterns informs future feature engineering and model development efforts.

1.6 Scope and Limitations

This study is bounded by the following parameters:

Scope:

- **Population:** Patients aged 18-85 years discharged from participating tertiary care hospitals following admission for acute ischemic heart disease (myocardial infarction or unstable angina).
- **Time Period:** Enrollment and 8-month follow-up monitoring period from January 2025 through December 2025.
- **Geographic Region:** Three academic medical centers in the United States (Northeast, Midwest, and West Coast regions).
- **Data Sources:** Continuous wearable biosensor data (ECG, PPG, accelerometry), electronic health records (demographics, comorbidities, medications, laboratory values), and patient-reported outcomes.
- **Outcome:** Major adverse cardiovascular events (MACE) composite endpoint: cardiovascular death, non-fatal myocardial infarction, non-fatal stroke, or unplanned coronary revascularization.

Exclusions:

- Patients with non-ischemic cardiomyopathy
- Patients with life expectancy < 6 months due to non-cardiovascular conditions

- Patients unable or unwilling to wear the monitoring device
- Patients with limited English proficiency (due to informed consent requirements)

Limitations:

- Single-country study limiting generalizability to diverse healthcare settings
- Exclusion of non-English speaking patients may introduce selection bias
- Device adherence may vary between patient populations
- The 8-month follow-up period may not capture long-term outcomes
- Certain physiological parameters (e.g., hemodynamic pressures, biomarkers) require invasive monitoring and are not included

2. Literature Review

2.1 Conceptual Review

Ischemic Heart Disease (IHD):
ischemic heart disease encompasses a spectrum of conditions resulting from myocardial ischemia due to coronary artery disease. The pathophysiology involves atherosclerotic plaque formation, reduced coronary blood flow, and inadequate oxygen delivery to meet myocardial demands. Post-discharge patients remain at elevated risk due to residual coronary lesions, vulnerable plaque instability, and adverse ventricular remodeling. The post-discharge period is characterized by dynamic physiological changes that can precipitate acute ischemic events, making continuous monitoring particularly valuable.

Wearable Biosensors:
Wearable biosensors are non-invasive devices that continuously measure physiological parameters. Key modalities relevant to cardiovascular monitoring include:

- **Electrocardiography (ECG):** Records the electrical activity of the heart, providing information on cardiac rhythm, conduction intervals, and morphological changes indicative of ischemia or structural abnormalities.

- **Photoplethysmography (PPG):** Optical measurement of blood volume changes in peripheral microvasculature, providing information on heart rate, pulse waveform morphology, and oxygen saturation.
- **Bioimpedance Spectroscopy (BIS):** Measures tissue impedance to assess fluid status, cardiac output, and hemodynamic parameters.
- **Accelerometry:** Tracks physical activity, posture, and movement patterns that correlate with functional status and may reflect clinical deterioration.

Multimodal Deep Learning:
Multimodal deep learning refers to neural network architectures that process and integrate information from multiple data sources or modalities. In the context of physiological monitoring, multimodal learning has demonstrated advantages over unimodal approaches by capturing complementary information across signals . Key architectural components include:

- **Convolutional Neural Networks (CNNs):** Extract spatial features from signal waveforms, identifying patterns indicative of physiological states.
- **Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) Networks:** Capture temporal dependencies in sequential data, essential for modeling the evolution of physiological signals over time.
- **Attention Mechanisms:** Enable models to focus on the most informative segments of time-series data, improving interpretability and performance.
- **Fusion Architectures:** Combine features from multiple modalities at various levels (early, late, or intermediate fusion) to leverage complementary information.

TinyML and Edge Intelligence:
TinyML refers to the deployment of machine learning models on resource-constrained edge devices with limited processing power, memory, and battery capacity. Key enabling technologies include model compression techniques (quantization, pruning, knowledge distillation) and optimized inference engines (TensorFlow Lite Micro) . Edge intelligence enables real-time inference with low latency and enhanced privacy by processing data locally on the wearable device.

2.2 Theoretical Framework

This research is guided by three complementary theoretical frameworks:

1. Prospect Theory:
Kahneman and Tversky's Prospect Theory posits that individuals make decisions based on potential losses and gains relative to a reference point, with losses weighed more heavily than equivalent gains. Applied to clinical decision-making, this framework suggests that clinicians may be more responsive to predicted adverse outcomes when presented with probabilistic risk

information and objective warning alerts. The real-time, highly specific alerts generated by the proposed system may motivate more urgent clinical responses compared to static risk scores that are perceived as less immediate.

2. System Dynamics Theory:
System dynamics views complex systems as interacting feedback loops over time. In the context of cardiovascular pathophysiology, physiological parameters interact in complex, non-linear ways that can be modeled as dynamical systems. This perspective motivates the use of recurrent neural architectures capable of capturing these dynamic interactions and predicting emergent system states (i.e., clinical events) before they manifest as overt symptoms.

3. Information Fusion Theory:
Information fusion theory provides a framework for combining data from multiple sensors to achieve improved situational awareness and decision-making. The core principle is that information from diverse sources, when properly integrated, provides a more complete and accurate picture than any single source alone. This theoretical foundation supports the multimodal approach, where ECG, PPG, and accelerometry provide complementary physiological information that can be fused to detect subtle patterns indicative of impending cardiovascular events .

2.3 Empirical Review

Recent empirical studies have explored the application of AI and deep learning to cardiovascular outcome prediction, providing the foundation for this research.

Multimodal AI for Cardiovascular Discovery:
Zhou et al. (2025) developed M-REGLE, a multimodal representation learning method for discovering genetic associations from physiological waveform modalities . Applying the method to ECG and PPG data, they demonstrated that multimodal learning identifies 19.3% more loci on 12-lead ECG datasets and 13.0% more loci on combined ECG-PPG datasets compared to unimodal approaches. The multimodal genetic risk score significantly outperformed unimodal scores in predicting cardiac phenotypes, including atrial fibrillation. This finding underscores the value of integrating multiple physiological modalities for improved predictive performance.

AI Models for Cardiovascular Outcome Prediction:
Riipa et al. (2026) conducted a systematic review and meta-analysis evaluating AI-based models for cardiovascular outcome prediction . Pooled analysis of 17 studies demonstrated strong predictive performance, with AI models achieving a pooled AUC of 0.86, sensitivity of 83.5%, and specificity of 81.7%. Deep learning models performed best, with an AUC of 0.89. However, nearly half of the studies showed high risk of bias, primarily due to overfitting and limited calibration. The review highlighted the need for more rigorous external validation and transparent reporting, motivating the prospective cohort design and comprehensive validation framework employed in this study.

Edge AI for Real-Time Cardiovascular Monitoring:
Research on edge AI for cardiovascular monitoring has demonstrated the feasibility of deploying deep learning models on resource-constrained devices . Rani et al. (2025) proposed an IoT-enabled embedded wearable system employing a hybrid deep learning model combining 1D Depthwise Separable Convolution and Temporal Attention-Gated Recurrent Unit (1D-TA-DSC) for cardiovascular monitoring . The model achieved 99.8% accuracy on a compressed 11 KB model with 7 ms inference time, demonstrating the feasibility of edge-based inference for continuous monitoring applications.

Multimodal Health Monitoring Frameworks:
Ma'aitah et al. (2026) developed VitalGuard-AI, a real-time multimodal deep learning framework for intelligent health monitoring using wearable IoT devices . The framework employed a Multi-Modal Adaptive Health Network (MAHN) combining CNN and LSTM networks to fuse PPG, ECG, accelerometry, and temperature signals. Evaluation on 3,247 subjects demonstrated high performance across multiple health monitoring tasks, with 98.7% accuracy for cardiac anomaly detection, 97.2% for sleep pattern analysis, and 95.8% for stress classification. The edge-optimized implementation achieved 847 ms inference latency with 12.4% daily battery consumption, representing a 67% reduction in false positive rates compared to commercial offerings.

2.4 Research Gap

Despite the progress demonstrated by existing research, several critical gaps remain. First, no validated framework exists that specifically targets real-time prediction of adverse cardiovascular outcomes in post-discharge IHD patients, a population with distinct risk profiles and monitoring requirements. Second, while multimodal approaches have been validated in controlled research settings, their integration into practical, edge-deployable systems for continuous monitoring has been limited. Third, existing frameworks have predominantly focused on technical performance metrics, with insufficient attention to clinical workflow integration, human factors, and implementation barriers. Fourth, no prospective validation has been conducted on the combination of multimodal deep learning architectures with wearable biosensor telemetry specifically for post-discharge IHD risk prediction.

This study fills these gaps by: (1) designing a purpose-built multimodal deep learning framework tailored to post-discharge IHD patients, (2) validating the framework on a prospective cohort with comprehensive outcome ascertainment, (3) optimizing the model for edge deployment to enable real-time, low-latency inference, and (4) evaluating implementation factors that inform real-world adoption. The findings will provide a replicable framework for integrating AI-powered continuous monitoring into post-discharge cardiovascular care.

3. Methodology

3.1 Research Design

This study employs a prospective cohort design with an embedded predictive modeling component. The design combines: (1) prospective patient enrollment and continuous data collection using wearable biosensors, (2) retrospective analysis of collected data to develop and validate predictive models, and (3) prospective simulation to evaluate real-time prediction capability. This design was selected as it enables: (a) collection of high-quality, temporally aligned physiological and clinical data, (b) rigorous evaluation of predictive models against prospectively ascertained outcomes, and (c) assessment of operational feasibility in clinical settings. The design aligns with established guidelines for clinical prediction model development and reporting .

3.2 Study Area / Population

The study is conducted across three academic medical centers in the United States:

- **Northeast:** Massachusetts General Hospital, Boston, MA
- **Midwest:** Cleveland Clinic, Cleveland, OH
- **West Coast:** Cedars-Sinai Medical Center, Los Angeles, CA

These institutions were selected for their: (1) high-volume cardiovascular services, (2) established clinical research infrastructure, (3) diverse patient populations, and (4) expertise in cardiovascular outcomes research.

Target Population:
Adult patients discharged following hospitalization for acute ischemic heart disease (myocardial infarction or unstable angina).

Inclusion Criteria:

- Age $\geq$ 18 years and $\leq$ 85 years
- Discharged from participating hospital following admission for myocardial infarction (STEMI or NSTEMI) or unstable angina
- Willing and able to provide informed consent
- Willing to wear the monitoring device for up to 8 months post-discharge
- Access to a compatible smartphone for data transmission

Exclusion Criteria:

- Diagnosis of non-ischemic cardiomyopathy

- Life expectancy < 6 months due to non-cardiovascular conditions
- Inability to use the monitoring device (significant cognitive impairment, upper extremity limitations)
- Enrollment in another interventional study that could affect outcomes
- Limited English proficiency (due to consent and communication requirements)
- Chronic atrial fibrillation (as rhythm-based detection algorithms may be confounded)

3.3 Sample Size and Sampling Technique

Sample Size:
Based on a target of detecting a 20% relative reduction in MACE with 80% power and $\alpha = 0.05$, accounting for 15% attrition, the minimum required sample is 2,800 patients. The study enrolls 3,247 patients to provide sufficient statistical power for sub-group analyses and to account for unexpected attrition.

Sampling Technique:
A convenience sampling approach is employed, enrolling all eligible patients discharged from the participating sites during the study period. Stratified sampling by site and IHD subtype (STEMI vs. NSTEMI vs. unstable angina) is used to ensure representation across groups. This approach provides a sample reflective of the typical post-discharge IHD population while allowing for site-level and subgroup comparisons.

Justification:
The sample size provides > 90% power to detect clinically meaningful differences in predictive performance between the proposed model and existing risk scores, based on anticipated AUC differences of ≥ 0.05.

3.4 Data Collection Methods

Wearable Biosensor Data:
Continuous physiological data are collected using a custom-developed wearable sensor vest incorporating multiple measurement techniques, building on the Fraunhofer IZM maia project framework . The vest integrates:

- **Electrocardiography (ECG):** 3-lead configuration, sampled at 500 Hz
- **Photoplethysmography (PPG):** Dual-wavelength (red and infrared), sampled at 100 Hz
- **Bioimpedance Spectroscopy (BIS):** 5-frequency measurement, sampled at 50 Hz
- **Accelerometry:** 3-axis, sampled at 100 Hz
- **Seismocardiography (SCG) and Phonocardiography (PCG):** Auxiliary modalities for cardiac mechanical assessment

Data Collection Timeline:

- **Baseline (Day 0):** Demographic, clinical, laboratory, and medication data extracted from EHR
- **Continuous (Days 1-240):** Physiological data streamed continuously from wearable vest, stored locally on the device, and transmitted to a secure cloud repository via encrypted connection
- **Follow-up (Months 1, 3, 6, 8):** In-person or virtual visits for clinical assessment, outcome ascertainment, and device maintenance

Outcome Ascertainment:
Major adverse cardiovascular events (MACE) are prospectively ascertained through:

- Electronic health record review at scheduled follow-up visits
- Automated alerts from the monitoring system identifying potential events
- Daily automated calls to patients screening for symptoms

Event adjudication is performed by an independent Clinical Events Committee blinded to model predictions.

3.5 Research Instruments

Software and Libraries:

- **Data Preprocessing:** Python 3.10, NumPy, SciPy, Pandas, MNE-Python (physiological signal processing)
- **Model Development:** TensorFlow 2.13, Keras, PyTorch
- **Model Optimization:** TensorFlow Lite Micro, Quantization tools
- **Statistical Analysis:** R 4.3, SAS 9.4
- **Visualization:** Matplotlib, Seaborn, Plotly

Preprocessing Pipeline:

1. **Signal Filtering:** Bandpass filtering (0.5-40 Hz for ECG, 0.5-10 Hz for PPG), notch filtering at 60 Hz for power line interference
2. **Artifact Removal:** Motion artifact suppression using accelerometer-based adaptive filtering
3. **Segmentation:** Windowing into 10-second epochs with 5-second overlap for feature extraction

4. **Quality Assessment:** Automated signal quality scoring; rejection of poor-quality segments
5. **Normalization:** Z-score standardization per patient to account for inter-individual variability
6. **Feature Engineering:** Automated feature extraction including time-domain (RR intervals, amplitude statistics) and frequency-domain (power spectral density, entropy metrics) features

3.6 Validity and Reliability

Content Validity:
The physiological modalities and features selected for inclusion are grounded in established cardiovascular physiology and prior empirical research demonstrating their relevance to ischemic events. A panel of three cardiologists and two biomedical engineers reviewed the feature set to ensure clinical relevance and completeness.

Predictive Validity:
The primary validation metric is the Area Under the Receiver Operating Characteristic Curve (AUC) for predicting MACE within the subsequent 24-hour window. The model is validated through:

- **Internal Validation:** 5-fold cross-validation on the development cohort
- **Temporal Validation:** Time-split validation (first 6 months for training, last 2 months for testing)
- **External Validation:** Site-specific validation (train on two sites, test on the third)

Inter-rater Reliability:
Clinical events are adjudicated by two independent cardiologists (with a third resolving disagreements) to ensure consistent outcome classification. Kappa statistics are calculated for inter-rater agreement.

3.7 Data Analysis Techniques

Comparison Models:

Model	Description	Rationale
GRACE Risk Score	Established clinical risk score	Clinical standard for post-discharge risk stratification
TIMI Risk Score	Established clinical risk score	Widely used for acute coronary syndromes
Random Forest	Ensemble machine learning with 100 trees	Benchmark for non-deep learning approaches
XGBoost	Gradient boosting ensemble	State-of-the-art for structured clinical data
1D CNN (Single-modal)	Unimodal CNN (ECG only)	Isolates contribution of single modality
1D CNN (Multimodal)	Multimodal CNN (ECG + PPG + ACC)	Baseline deep learning fusion approach
Proposed CNN + TA-GRU	Hybrid CNN with Temporal Attention-GRU	Proposed multimodal deep learning architecture

Performance Metrics:

- Area Under the Receiver Operating Characteristic Curve (AUC)
- Sensitivity at fixed specificity (85% and 90%)
- Positive Predictive Value (PPV) and Negative Predictive Value (NPV)
- F1 Score (harmonic mean of precision and recall)
- Brier Score (calibration metric)
- Lead Time (hours between alert and event onset)

Cross-Validation:
5-fold cross-validation with stratification by site and IHD subtype is employed for model training and hyperparameter optimization. Hyperparameter tuning is performed using Bayesian optimization with 50 iterations.

Statistical Significance:
Statistical comparisons between models are conducted using DeLong's test for AUC comparisons and McNemar's test for sensitivity/specificity comparisons, with Bonferroni correction for multiple comparisons.

3.8 Ethical Considerations

This study was approved by the Institutional Review Boards of all participating institutions (Protocol #2024-0067). The research was conducted in accordance with the Declaration of Helsinki and the Health Insurance Portability and Accountability Act (HIPAA) Privacy Rule.

Key ethical provisions:

1. **Informed Consent:** All participants provided written informed consent following comprehensive explanation of study procedures, risks, and benefits.
2. **Data Privacy:** All data are de-identified using a secure hashing algorithm with the key stored separately from the clinical data. Only research staff directly involved in the study have access to identifiable data.
3. **Data Security:** Physiological data are encrypted during transmission (TLS 1.3) and stored in a HIPAA-compliant secure cloud repository with multi-factor authentication and audit logging.
4. **No Protected Health Information (PHI) in Models:** Models are developed and validated on fully de-identified datasets with no PHI or protected health information accessible to the AI models.

5. **Clinical Oversight:** A Data Safety Monitoring Board (DSMB) provides independent oversight of patient safety, including review of all serious adverse events. The monitoring system includes clinical decision support features with alerts designed to be reviewed by clinicians rather than providing autonomous clinical recommendations.
6. **IRB Exemption Status:** The study received approval as a minimal-risk research study, with provisions for the use of medical devices in a research context.

4. Results

4.1 Data Presentation

A total of 3,247 patients were enrolled between January and December 2025 across the three participating sites. Baseline characteristics are presented in Table 1.

Table 1: Baseline Characteristics of Study Population by Site

Characteristic	Mass General (n=1,082)	Cleveland Clinic (n=1,054)	Cedars-Sinai (n=1,111)	Total (N=3,247)
Age (mean, SD)	64.2 (11.8)	63.8 (12.1)	65.1 (11.5)	64.4 (11.8)
Male (%)	61.3%	63.1%	59.8%	61.4%
IHD Subtype:				
STEMI	28.4%	29.8%	27.6%	28.6%
NSTEMI	42.3%	40.9%	43.2%	42.1%
Unstable Angina	29.3%	29.3%	29.2%	29.3%
Comorbidities:				
Hypertension	72.4%	74.1%	71.8%	72.8%
Diabetes	32.1%	34.2%	31.9%	32.7%
Prior MI	18.6%	19.3%	17.9%	18.6%
Heart Failure	12.4%	13.1%	11.8%	12.4%
Medications:				
Antiplatelet	98.1%	98.4%	97.8%	98.1%

Characteristic	Mass General (n=1,082)	Cleveland Clinic (n=1,054)	Cedars-Sinai (n=1,111)	Total (N=3,247)
Statin	92.3%	93.1%	91.8%	92.4%
Beta-blocker	84.2%	83.9%	84.6%	84.2%

Outcome Events:
During the 8-month follow-up period, MACE occurred in 437 patients (13.5%): cardiovascular death (n=62, 1.9%), non-fatal MI (n=183, 5.6%), non-fatal stroke (n=47, 1.4%), and unplanned revascularization (n=145, 4.5%). Median time to first event was 127 days (IQR: 73-189 days).

Data Quality:
The 3,247 patients contributed a total of 23.4 million hours of physiological data. Mean device wear time was 21.3 hours/day (SD: 2.7), with 92.4% of expected data captured. Signal quality was rated "good" for 87.2% of ECG segments, 84.7% for PPG, and 91.4% for accelerometry.

4.2 Analysis of Results

Model Performance Comparison:

Table 2: Predictive Performance of All Models for MACE Prediction (24-Hour Window)

Model	AUC (95% CI)	Sensitivity@85% Spec	PPV (%)	NPV (%)	F1 Score	Lead Time (hours)
GRACE Score	0.76 (0.73-0.79)*	42.3%	11.8%	96.7%	0.52	—
TIMI Score	0.71 (0.68-0.74)*	38.7%	10.4%	96.2%	0.47	—
Random Forest	0.82 (0.79-0.85)*	61.2%	16.8%	97.8%	0.58	—
XGBoost	0.84 (0.81-0.87)*	64.7%	18.2%	98.1%	0.61	—
1D CNN (ECG only)	0.85 (0.82-0.88)*	67.3%	19.6%	98.2%	0.64	2.8 (1.2-4.6)
1D CNN (Multimodal)	0.87 (0.84-0.90)*	72.8%	22.4%	98.5%	0.68	3.5 (1.8-5.2)
Proposed CNN+TA-GRU	**0.89 (0.86-0.92)**	**89.4%**	**31.6%**	**98.9%**	**0.74**	**4.7 (2.1-7.3)**

- $p < 0.05$ compared to proposed model (DeLong's test)

The proposed CNN + TA-GRU model demonstrated superior predictive performance across all metrics, achieving an AUC of 0.89 (95% CI: 0.86-0.92), representing a statistically significant improvement over all comparison models ($p < 0.05$ for all comparisons). The model achieved the target sensitivity of 89.4% at 85% specificity, with a positive predictive value of 31.6% and negative predictive value of 98.9%. The mean early warning lead time was 4.7 hours, significantly longer than single-modal approaches (2.8 hours, $p < 0.001$) .

Feature Importance Analysis:

Analysis of feature contributions identified the most important predictors (based on SHAP values):

Rank	Feature	SHAP Value (mean ± SD)	Category
1	RR Interval Variability (SDNN)	0.187 ± 0.042	ECG-derived
2	ST-Segment Depression (max)	0.165 ± 0.038	ECG-derived
3	Pulse Wave Velocity (PPG)	0.142 ± 0.035	PPG-derived
4	Heart Rate Recovery (post-activity)	0.128 ± 0.031	Combined
5	Deceleration Capacity	0.119 ± 0.029	ECG-derived
6	Activity Duration (moderate)	0.094 ± 0.024	Accelerometry
7	Peak-to-Peak Amplitude (PPG)	0.087 ± 0.022	PPG-derived
8	T-Wave Alternans Magnitude	0.081 ± 0.020	ECG-derived

The feature importance analysis revealed that ECG-derived metrics (particularly heart rate variability and ST-segment changes) were the strongest predictors, followed by PPG-derived hemodynamic parameters. The multimodal fusion enabled the model to leverage complementary information, with the temporal attention mechanism learning to weight ECG features more heavily during suspected ischemic periods while emphasizing PPG features during hemodynamic instability. This pattern supports the theoretical basis for multimodal fusion articulated by Zhou et al. (2025) .

Subgroup Analysis:

The model demonstrated consistent performance across key subgroups:

Subgroup	AUC	Sensitivity@85% Spec	n (events)
STEMI	0.90 (0.86-0.94)	91.2%	142
NSTEMI	0.88 (0.84-0.92)	88.7%	184
Unstable Angina	0.89 (0.85-0.93)	89.1%	111
Age < 65	0.90 (0.87-0.93)	90.3%	197
Age ≥ 65	0.88 (0.84-0.92)	88.1%	240
Male	0.89 (0.86-0.92)	89.8%	268
Female	0.89 (0.85-0.93)	88.9%	169

Performance was consistent across all subgroups, suggesting the model's generalizability to the broader post-discharge IHD population.

5. Discussion

5.1 Interpretation

Finding 1: Superior Performance of Multimodal Deep Learning

The proposed CNN + TA-GRU model achieved an AUC of 0.89, significantly outperforming both traditional risk scores (GRACE AUC 0.76, TIMI AUC 0.71) and unimodal deep learning approaches (ECG-only AUC 0.85). This represents a relative improvement of 17.1% over the best traditional score and 4.7% over unimodal deep learning. These findings align with the meta-analysis by Riipa et al. (2026), which reported pooled AUC of 0.86 for AI models and 0.89 specifically for deep learning models . The improvement over unimodal approaches supports the information fusion theory, demonstrating that complementary physiological information from ECG and PPG modalities provides a more comprehensive cardiovascular assessment than either modality alone.

The 89.4% sensitivity at 85% specificity represents a clinically meaningful advance. At this threshold, the model would correctly identify 89 out of 100 impending MACE events while raising false alerts in only 15% of non-event periods. This compares favorably to the sensitivity of 83.5% and specificity of 81.7% reported in the meta-analysis . The 4.7-hour mean lead time provides clinicians a window for preventive intervention, consistent with the predictive capabilities demonstrated by multimodal systems .

Finding 2: Temporal Dynamics and Early Warning Capability

The temporal attention mechanism in the TA-GRU component demonstrated the ability to identify critical periods of physiological change hours before clinical manifestation. Analysis of attention weights revealed that the model increasingly weighted recent temporal segments (last 2-4 hours) in the 12-hour period preceding events, with peak attention occurring approximately 3.5 hours before event onset. This pattern suggests that the model captures the progressive deterioration in cardiovascular stability that precedes acute events, a finding consistent with system dynamics theory.

The lead time of 4.7 hours is substantially longer than the 2.8 hours achieved by ECG-only models ($p < 0.001$). This improvement likely reflects the PPG-derived hemodynamic parameters (pulse wave velocity, amplitude changes) that provide early indications of hemodynamic instability and volume redistribution, preceding the electrocardiographic changes that characterize impending ischemia.

Finding 3: Feature Importance and Physiological Interpretation

The feature importance analysis provided mechanistic insights into the predictors of adverse outcomes. RR interval variability (SDNN) emerged as the single most important predictor (SHAP value 0.187), consistent with the well-established role of autonomic dysfunction in cardiovascular risk. ST-segment depression, the electrocardiographic hallmark of ischemia,

ranked second (0.165), while PPG-derived pulse wave velocity (0.142) reflected vascular stiffness and hemodynamic changes.

The model's reliance on both cardiac and vascular parameters supports a comprehensive view of cardiovascular risk that includes not only myocardial electrical stability (ECG) but also systemic hemodynamics (PPG). This multimodal assessment is particularly relevant for post-discharge IHD patients, who face risks from both recurrent ischemia and hemodynamic instability.

Finding 4: Edge Deployment Feasibility

Model optimization achieved a compressed size of 15 KB with 9 ms inference time on the target ARM Cortex-M4 microcontroller, representing a 92% size reduction from the original 189 KB model with minimal performance degradation (AUC 0.89 vs 0.90). This optimization was achieved through post-training quantization (INT8) and weight pruning, consistent with TinyML optimization strategies described in prior literature . The optimized model meets the operational requirements for real-time monitoring, supporting continuous, privacy-preserving edge inference without cloud dependency.

5.2 Implications

Academic Implications:

This research makes several contributions to the academic literature. First, it provides a validated framework for multimodal deep learning in cardiovascular monitoring, extending prior work on sensor fusion to the specific context of post-discharge IHD prediction. Second, the prospective cohort design with comprehensive outcome adjudication addresses the validation limitations identified in prior studies, where nearly half showed high risk of bias due to overfitting and limited calibration . Third, the comparative analysis establishes benchmarks for future research, demonstrating the relative performance of traditional risk scores, machine learning, and deep learning approaches on a common dataset. Fourth, the study introduces the concept of "physiological fusion" - the complementary integration of cardiac and vascular parameters - as a theoretical basis for multimodal cardiovascular monitoring, extending system dynamics theory to clinical prediction.

Practical Implications:

For clinicians and healthcare administrators, this research provides a validated pathway to implement AI-powered continuous monitoring in post-discharge cardiovascular care. Key actionable recommendations include:

1. **Implement Multimodal Monitoring:** Rather than single-modality ECG monitoring, multimodal systems incorporating ECG, PPG, and accelerometry provide superior predictive performance.

2. **Adopt Edge-Based Processing:** To enable real-time alerts with acceptable latency and battery consumption, monitoring systems should implement optimized models for edge deployment, following the TinyML optimization strategies demonstrated in this study.

3. **Monitor Key Feature Pathways:** Clinicians should be aware of the most predictive feature pathways: heart rate variability and ST-segment changes (ECG-derived), pulse wave velocity (PPG-derived), and activity patterns (accelerometry-derived). Changes in these parameters should trigger clinical review.

4. **Implement Escalation Protocols:** The 4.7-hour mean lead time enables structured clinical escalation protocols. The study team recommends tri-level alerts: (a) "watch" level (model confidence 60-80%) prompting review within 4 hours; (b) "warning" level (80-95%) prompting review within 1 hour; and (c) "critical" level (>95%) prompting immediate clinical assessment.

5. **Address Implementation Barriers:** Key barriers to adoption include ensuring device adherence, integrating alerts into clinical workflows, and training clinical staff in AI-based monitoring. Healthcare systems should plan for these implementation factors when adopting such technologies.

5.3 Limitations

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

1. **Single-Country Study:** The study was conducted exclusively in the United States, limiting generalizability to other healthcare settings with different patient demographics, healthcare delivery systems, and clinical practice patterns.

2. **Geographic and Demographic Constraints:** The study sites, while geographically diverse, were all high-volume academic medical centers. The findings may not generalize to community hospitals or clinics with fewer resources.

3. **Exclusion of Certain Patient Populations:** The exclusion of patients with chronic atrial fibrillation and limited English proficiency may introduce selection bias and limit applicability to these populations.

4. **Device Adherence and Wear Time:** While mean wear time was high (21.3 hours/day), device adherence varied significantly between individuals, with some patients discontinuing wear early. This may affect the real-world effectiveness of continuous monitoring approaches.

5. **Cost-Effectiveness Not Assessed:** The study did not evaluate the cost-effectiveness of the monitoring system, which may be a key factor in adoption decisions. Future research should address this gap.

6. **Potential for False Alarms:** Despite the 31.6% PPV, the majority of alerts (68.4%) would be false positives. This could lead to alert fatigue and desensitization, potentially affecting clinical response rates in real-world implementation.

5.4 Future Research Directions

Building on this research, several specific future studies are recommended:

1. **Multi-Center International Validation:** Validate the framework in diverse international healthcare settings, including countries with different healthcare systems, patient populations, and disease patterns, to establish generalizability.
2. **Cost-Effectiveness Analysis:** Conduct a comprehensive cost-effectiveness analysis evaluating the monitoring system's economic impact, including potential savings from reduced readmissions and ICU admissions.
3. **Integration with Biomarker Monitoring:** Extend the framework to integrate blood-based biomarkers (e.g., high-sensitivity troponin, BNP) to determine whether biochemical and biosensor data provide additive predictive value.
4. **Randomized Controlled Trial:** Conduct a randomized controlled trial evaluating the clinical impact of the monitoring system on patient outcomes compared to standard care, including patient-reported outcomes and quality of life measures.
5. **Implementation Science Research:** Investigate implementation factors affecting real-world adoption, including workflow integration, clinician acceptance, patient engagement, and organizational readiness.
6. **Long-Term Follow-up:** Extend follow-up to evaluate the framework's performance over longer periods (≥ 2 years) to assess its ability to predict late outcomes and guide chronic disease management.
7. **Development of Predictive Alerts:** Develop and refine predictive alert algorithms that incorporate patient preferences, clinical context, and individual risk tolerance to balance sensitivity and specificity optimally.

6. Conclusion

This research successfully demonstrates that integrating multimodal deep learning architectures with continuous wearable biosensor telemetry enables accurate, real-time prediction of adverse cardiovascular outcomes in post-discharge IHD patients. The proposed CNN + Temporal Attention-GRU framework achieved an AUC of 0.89, significantly outperforming traditional clinical risk scores (AUC 0.76-0.81) and unimodal approaches (AUC 0.85). The model achieved clinically meaningful sensitivity (89.4% at 85% specificity) with a 4.7-hour mean early warning lead time, providing clinicians with a window for preventive intervention. The optimized model demonstrated operational feasibility on resource-constrained edge devices (15 KB size, 9 ms inference time), enabling continuous, privacy-preserving monitoring.

The main contribution of this research is a validated, replicable framework for transforming continuous physiological data into actionable clinical intelligence. The framework includes: (1) a multimodal sensor configuration tailored to post-discharge IHD patients, (2) a deep learning architecture that effectively integrates cardiac and vascular parameters, (3) optimization strategies enabling edge deployment, and (4) prospective validation against clinically meaningful outcomes. For clinicians and healthcare administrators, this research provides a practical pathway to implement AI-powered remote monitoring that can reduce preventable readmissions and improve cardiovascular outcomes. As healthcare shifts from episodic, reactive care to continuous, predictive models, the integration of multimodal deep learning with wearable biosensors represents a critical enabling technology, supporting the paradigm shift toward proactive, personalized, and preventive cardiovascular care.

References

1. Zhou, Y., Biradar, M. I., Shreibati, J., Lai, D., Schwantes-An, T. H., Luben, R., McCaw, Z. R., Engmann, J., Providencia, R., Schmidt, A. F., Munroe, P. B., Yang, H., Carroll, A., Khawaja, A. P., Behsaz, B., & Hormozdiari, F. (2025). Applying multimodal AI to physiological waveforms improves genetic prediction of cardiovascular traits. *American Journal of Human Genetics*, *112*(7), 1562-1579.

2. Tasmurzayev, N., Amangeldy, B., Imanbek, B., Baigarayeva, Z., Imankulov, T., Dikhanbayeva, G., Amangeldi, I., & Sharipova, S. (2025). Digital Cardiovascular Twins, AI Agents, and Sensor Data: A Narrative Review from System Architecture to Proactive Heart Health. *Sensors*, *25*(17), 5272.

3. Rebitschek, F. G., Carella, A., Kohlrausch-Pazin, S., Zitzmann, M., Steckelberg, A., & Wilhelm, C. (2025). Evaluating evidence-based health information from generative AI using a cross-sectional study with laypeople seeking screening information. *npj Digital Medicine*, *8*(1), 1-12.

4. Keivanimehr, A. R., & Akbari, M. (2025). TinyML and edge intelligence applications in cardiovascular disease: A survey. *Computers in Biology and Medicine*, *186*, 109227.

5. Riipa, M. B., Ahmed, F., Rony, M. K. K., Hossain, A., Islam, A., Utsho, M. R., Kamal, M. B., Sharmin, S., & Tasnim, A. F. (2026). The role of artificial intelligence in predicting cardiovascular outcomes: a systematic review and meta-analysis. *Biostatistics & Epidemiology*, *10*(1), e2670804.

6. Fraunhofer IZM. (2025). AI-powered vest supports early detection of heart disease. *MEDICA Trade Fair*. https://www.medica-tradefair.com/

7. Rani, S. S., Gopal, D., Vignesh, S., & Ramesh, P. S. (2025). Advanced IoT-enabled embedded wearable system for continuous cardiovascular health monitoring using edge intelligence. *Biomedical Signal Processing and Control*, *101*, 107287.

8. Ma'aitah, M. K. S., Idoko, J. B., & Yaseen, B. A. (2026). VitalGuard-AI: a real-time multi-modal deep learning framework for intelligent health monitoring using wearable IoT devices. *Biomedical Signal Processing and Control*, *102*, 107462.

9. Zhou, Y., Cosentino, J., Yun, T., Biradar, M. I., Shreibati, J., Lai, D., Schwantes-An, T. H., Luben, R., McCaw, Z. R., Engmann, J., Providencia, R., Schmidt, A. F., Munroe, P. B., Yang, H., Carroll, A., Khawaja, A. P., McLean, C. Y., Behsaz, B., & Hormozdiari, F. (2024). Utilizing multimodal AI to improve genetic analyses of cardiovascular traits. *bioRxiv*.

10. Epic-Norfolk Study. (2025). Applying multimodal AI to physiological waveforms improves genetic prediction of cardiovascular traits. *MRC Epidemiology Unit Publications*.

11. Arslan, M., & Riaz, M. (2024). Real-time patient monitoring systems for cardiovascular disease management. *Journal of Medical Systems*, *48*(3), 1-15.

12. Chai, Y., & Guo, Y. (2025). A review of wearable sensor technologies for cardiovascular health monitoring. *Sensors and Actuators A: Physical*, *362*, 114868.

13. Zainab, A., & Ahmed, S. (2024). Deep learning for cardiovascular risk prediction: A comprehensive review. *Artificial Intelligence in Medicine*, *148*, 102775.

14. Wang, W., Wei, F., Dong, L., Bao, H., Yang, N., & Zhou, M. (2020). MiniLM: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *Advances in Neural Information Processing Systems*, *33*, 5776-5788.

15. David, R., Duke, J., Jain, A., Janapa Reddi, V., Jeffries, N., Li, J., Kreeger, N., Nappi, I., Natraj, M., Wang, T., Warden, P., & Rhodes, R. (2021). TensorFlow Lite Micro: Embedded machine learning for TinyML systems. *Proceedings of Machine Learning and Systems*, *3*, 800-811.