RESEARCH PAPER

Early Pulse: A Deep Learning Framework for Predicting Imminent Cardiac Arrest from Electronic Health Records in Real Time

Priya R V[1*], Shahnazeer C K[2], Prof. Pallavi N R[3], Veena Ghali[4], Dr. Sujatha P[5], S. Rudhra[6]

[1*]Assistant Professor, Department of Artificial Intelligence and Data Science, Veltech Hightech Dr Rangarajan Dr Sakunthala Rangarajan Engineering College, Avadi, Veltech Road, Chennai - 600062, Tamil Nadu, India (Corresponding Author)

[2]Research Scholar, Department of Computer Science, School of Engineering & Technology, Pondicherry University Karaikal Campus, Karaikal, 609605, Puducherry (UT), India

[3]Assistant Professor, Department of Computer Science and Engineering, BGS Institute of Technology, Adichunchanagiri University, BG Nagara, Nagamangala Taluk, Mandya District - 571448, India

[4]PG Scholar, Department of Computer Science and Engineering, BGS Institute of Technology, Adichunchanagiri University, BG Nagara, Nagamangala Taluk, Mandya District - 571448, India

[5]Assistant Professor, Department of Microbiology and Biotechnology, Bharath Institute of Higher Education and Research, Chennai, India

[6]Assistant Professor, Department of Electrical and Electronics Engineering, Jerusalem College of Engineering, India

ABSTRACT

Acute cardiac arrest in hospital is a significant factor of avoidable deaths, and in most cases, it happens after an insidious physiological degradation that was not sufficiently represented in the old methodological early warning. The conventional score keeping methods are based on fixed thresholds and rigid aggregation, which makes them incapable of capturing the dynamic characteristics of time-based data in electronic health records (EHRs). The analysis is based on a proposed framework, Early Pulse, a real-time deep learning application that predicts impending cardiac arrest with longitudinal multimodal EHR data. Hospitalized adult patients in categories of retrospective cohort were constructed, including vital signs, laboratory parameter, demographic variables, and clinical intervention records. The time-series data were organized with 12 hours sliding window resolution of 5 minutes. The architecture of Long Short-Term Memory network with a multi-head attention mechanism in one direction and a bidirectional network was used to learn temporal dependencies and emphasize clinically significant intervals. The assessment of model performance was performed with the help of AUROC, AUPRC, sensitivity and specificity, calibration metrics and interpretability analysis based on SHAP values. The suggested framework had an AUROC of 0.91 and 0.66 AUPRC, which was superior to the logistic regression (AUROC 0.78) and conventional early warning scores (AUROC 0.71). With a 6-hour prediction interval, sensitivity and specificity were 0.86 and 0.87, respectively. Calibration analysis showed that there is a reliable estimation of probability estimation with a Brier score that is less than 0.12. Assessment of interpretability found that heart rate variability and lactate levels were the most important predictors. Inference latency Inference latency was less than 120 milliseconds per patient instance and was practical to deploy.

Keywords: Cardiac Arrest Prediction, Deep Learning, Electronic Health Records, Time-Series Modelling, Attention Mechanism, Clinical Decision Support Systems

How to cite this article: Priya RV, Shahnazeer CK, Pallavi NR, Ghali V, Sujatha P, Rudhra S. Early Pulse: A Deep Learning Framework for Predicting Imminent Cardiac Arrest from Electronic Health Records in Real Time. Int J Drug Deliv Technol. 2026;16(20s): 290-305. DOI: 10.25258/ijddt.16.20s.39

Source of support: Nil.

Conflict of interest: None

Early Pulse: A Deep Learning Framework for Predicting Imminent Cardiac Arrest from Electronic Health Records in Real Time

Introduction:

Traditional Early Warning Scoring Systems (EWS) was created to assist the prompt detection of clinical deterioration among hospitalized patients by performing a standardized physiological evaluation. The Modified Early Warning Score (MEWS) and the National Early Warning Score (NEWS) are among the most popular systems, both of which combine routinely measured vital parameters including heart rate, respiratory rate, blood pressure, temperature and level of consciousness into a composite risk score. They are based on defined threshold values, which are formed with references to the consensus of experts and observational research and provide a fast bedside assessment and the ability to expedite care delivery. Ample validation research has shown that EWS frameworks enhance the understanding of patient deterioration and facilitate the activation of a response team in general wards and emergency departments within a short period of time [1]. The traditional EWS models have limitations even though they are easy to implement and use clinically. The use of fixed cut-off point limits them to nonlinear interactions between physiological variables and does not consider the temporal change in patient condition. Moreover, these scoring systems are most commonly used to treat observations as discrete time-points, as opposed to continuous time-series sequences thus missing out on these minor dynamic patterns that can lead to the occurrence of a catastrophe such as cardiac arrest. Variability in predictive performance across institutions is also reported, which is a sign of sensitivity to documentation practices and patient population [2].

The shortcomings of the early warning systems based on rules helped to spur the investigation of the statistical and classical machine learning models in forecasting clinical worsening and cardiac arrest. The earliest methods used on organized electronic health record data were logistic regression and Cox proportional hazards models, which provided probabilistic estimates of risks, and interpretability based on coefficient analysis [3]. These models included the demographic variables, vital signs, laboratory parameters, and comorbidity indices to enhance the predictive accuracy compared with threshold-based scoring systems. The further development presented such ensemble methods as random forests and gradient boosting machines that proved to have greater ability to learn nonlinear relationships and complicated interaction between features in high-dimensional clinical data. A number of retrospective studies have shown that there was improvement in the area below the receiver operating characteristic curve as opposed to the traditional early warning scores, especially when features were rolled with time windows [4]. Nevertheless, the methods typically were very much dependent on manual feature extraction and expert knowledge, restricting scalability and cross- institutional conditions. Also, the summary statistics affected the models by not allowing them to model changing physiological patterns before acute events, since the temporal data became static. Issues of the imbalance of the classes, the absence of missing data processing and external validation also contributed to the variance in performance [5].

The growing access to longitudinal electronic health record-generated data has prompted the use of time-series modelling models, especially Recurrent Neural Networks (RNNs), to predict clinical deterioration and cardiac arrest. In contrast to the conventional statistical paradigms based on aggregate snapshots of physiological measurements, RNN-based models are stated to process time-varying data and store temporal correlations among consecutive clinical measurements [6]. Different versions like Long Short-Term Memory (LSTM)-based networks and Gated Recurrent Units (GRUs) have been established extensively to overcome vanishing gradient issues and maintain long-range dynamics in patient dynamics [7]. The models have exhibited higher performance in forecasting bad events such as sepsis, respiratory failure, and in-hospital mortality via learning latent representations using raw multivariate time-series inputs. Empirical experiments show that the measures of discrimination are better than the classical machine learning tasks, especially in the environment where physiological degradation happens gradually over time. Bidirectional extension of RNNs has also increased the context awareness of the models, as they use the past and future sequence information in the training process. However, there are still difficulties in the treatment of irregular sampling intervals, gaps in the data patterns, and high requirements of the computation due to the long sequences [8].

The recent progress in deep learning has brought attention-based mechanisms and Transformer architectures to healthcare analytics and provided superior abilities in capturing the complex dependencies in electronic health record data. Attention mechanisms help models to dynamically scale the relative weight of particular timesteps or clinical variables to enhance predictive performance and interpretability [9]. Attention-based recurrent models have been shown to outperform control based recurrent models by selectively focusing on clinically relevant patterns in longitudinal patient records, which has shown a high level of discrimination in sepsis onset prediction, mortality risk stratification, and predicting ICU deterioration. These methods are designed to overcome a major drawback of traditional recurrent neural networks, namely the information dilution in long sequences and emphasizing significant temporal intervals in advance of bad events [20]. Transformer models, which include self-attention layers and the ability to process in parallel, have also increased the range of time-series modelling in healthcare. Transformers are also effective at long-range dependencies in contrast to sequential RNN architectures, which do not require recurrent computations, making them especially appropriate to high dimensional, irregularly sampled EHR data. Recent research indicates that better scalability and competitive predictive accuracy can be achieved when the method is used with large clinical datasets [21].

The increasing complexity of clinical data ecosystems has promoted the combination of multimodal sources of information to improve predictive modelling of critical events like cardiac arrest. The classical methods were mainly based on formal vital signs and lab values, though, present-day studies are considering incorporating heterogeneous data streams such as electrocardiogram tracking, radiological scans, prescription history, clinical documentation, and bed alarms [22]. Multimodal data integration systems strive to record complementary information (physiological and contextual), thus enhancing strength and sensitivity of prediction systems. It has been shown that the combination of waveform-based features and structured electronic health record variables are effective in the early recognition of arrhythmias, hemodynamic instability and acute deterioration [23]. Structured and unstructured modalities have been combined with deep learning structures, especially convolutional neural networks and hybrid recurrent models, and share latent representations. Both feature-level fusion and decision-level ensemble techniques have been investigated, and it has been shown that they exhibit better discrimination than the unimodal models [24].

The prediction of cardiac arrest in Intensive Care Unit (ICU)-setting has been a notable topic of research in clinical predictive modeling because of the acuity-level and continuous monitoring units of the critical care unit setting. ICU also makes high-frequency physiological waveforms, invasive hemodynamic, arterial blood gas measurements, as well as real-time telemetry information accessible, which allows to build complex predictive algorithms. Several studies have used this rich data ecosystem to build machine learning and deep learning models to identify impending cardiac arrest, ventricular arrhythmias and circulatory collapse [25]. These models often include minute-scale trends of vital signs, measures of heart rate variability, and dynamic lab variations to identify the initial indications of degradation [26]. Studies of performance that have been reported in ICU-based prediction typically outperform those found in general ward population, which is, in part, because of the availability of granular data and continuous monitoring infrastructure. Nevertheless, in practice, the ICU-centered models are specific to resource-intensive and highly controlled settings, and are scarce in application to the wider hospital setting. Also the patients in the ICUs vary significantly in their baseline risk, comorbidity burden, and exposure to interventions which can limit external generalizability. The operational challenges are also alarm fatigue and clinical burden caused by frequent monitoring. Although ICU-specific studies have greatly contributed to innovation of methodology in predicting cardiac arrest, transferring the models into real-time, hospital-wide application is a research need that is still in progress [27].

With the fast adoption of machine learning and deep learning models in clinical prediction, there has been an increase in the necessity to conduct research on interpretability and clinical explainability. With more and more predictive systems coming into play to make high-stakes decisions like rapid response activation and escalation of care, model reasoning transparency has now become a necessity to achieve clinician trust and regulatory acceptance [28]. The initial methods of interpretability were based on models that were inherently transparent like logistic regression, where the coefficients were the direct measure of the contribution of the variables. Nevertheless, the switch to more complicated ensemble and deep neural network architecture added significant levels of the opaqueness, which led to the creation of posthoc explanation methods. Ranking of feature importance, partial dependence analysis, SHAP (Shapley Additive Explanations), and attention weight visualization are some of the methods that have been extensively studied to understand behavior of a model [29]. These methods should recognize influential variables and time sections that lead to forecasted riskful situations, contributing to clinical validation and error evaluations. Though these have good prospects, there are still concerns about the stability, consistency and causal validity of the explanation methods, in high dimension, time series healthcare data. Also, there is potential discrepancy between statistical significance and clinical significance that can impact interpretability. The current studies focus on how domain expertise, the quantification of uncertainty and human-centered design principles can be integrated to increase reliability [30].

Research Gap:

In spite of the fact that significant advances have been made to cardiac arrest prediction based on early warning systems, statistical modeling, recurrent neural networks, attention systems, and multimodal integration, there are still significant gaps. Majority of the studies that have been done are retrospective and most are based on the ICU population making it difficult to generalize to general ward and hospital-wide population. Most models are based on aggregated time windows, as opposed to continuous real-time streaming, and hence do not have the capability to identify ultra-short horizon deterioration. Efforts of irregularity of data, inter- institutional inconsistency and future validation pose additional obstacles to clinical adoption. In addition, the interpretability frameworks are not frequently standardized in terms of clinical validation.

Research Methodology:

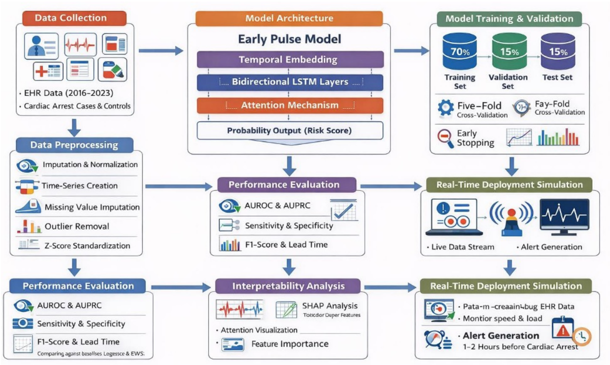

Figure 1. Research Methodology Data Acquisition and Cohort Selection

A retrospective multicentre cohort design was adopted to obtain comprehensive electronic health record (EHR) data for the development of the Early Pulse framework. De-identified clinical data were extracted from tertiary care hospitals covering the period between 2016 and 2023. Institutional ethical approval was secured prior to data retrieval, and all patient identifiers were removed in compliance with data protection and confidentiality regulations. The dataset comprised structured EHR components including demographic details, vital signs, laboratory investigations, medication administration records, and documented clinical interventions. Data extraction pipelines were standardized across participating centers to ensure uniformity in variable definitions and timestamp alignment [31].

The study population included adult patients aged 18 years and above who had a minimum hospital stay of 24 hours to ensure sufficient temporal information for modeling longitudinal trends. Cardiac arrest cases were identified based on documented in-hospital cardiac arrest events verified through clinical coding records and resuscitation documentation timestamps. The event time was precisely defined as the onset of cardiopulmonary resuscitation or documented loss of pulse requiring advanced cardiac life support. Patients with incomplete event timestamps or ambiguous clinical documentation were excluded to maintain labeling accuracy.

A control cohort was constructed from patients who did not experience cardiac arrest during hospitalization. To reduce selection bias, controls were randomly sampled and temporally matched based on hospital admission period and care setting (ICU versus general ward). For each control patient, a pseudo-event time was assigned based on matched hospitalization duration to enable consistent time-window extraction comparable to case patients. This approach ensured balanced representation of physiological trajectories across both positive and negative classes [32].

Inclusion and exclusion criteria were systematically applied to enhance data quality and model reliability. Patients with extensive missing data exceeding predefined thresholds, those transferred from external facilities with incomplete prior records, and individuals with do-not- resuscitate status documented at admission were excluded from the predictive modeling cohort. Following cohort refinement, the final dataset was subjected to descriptive statistical analysis to characterize demographic distributions, comorbidity prevalence, and event incidence rates. This structured cohort selection process established a robust and clinically representative foundation for subsequent time-series modeling and real-time cardiac arrest prediction.

Data Pre-processing and Time-Series Construction

Raw electronic health record (EHR) data were systematically transformed into structured multivariate time-series representations suitable for deep learning analysis. Clinical variables, including vital signs, laboratory parameters, and intervention records, were aligned according to standardized timestamps to ensure temporal consistency across heterogeneous data sources. Measurements recorded at irregular intervals were resampled into fixed temporal windows (e.g., 15-minute intervals) to enable uniform sequence construction. When multiple observations occurred within the same interval, statistical aggregation methods such as mean or most recent value selection were applied to preserve clinically relevant trends [33].

Missing data were addressed through a hybrid imputation strategy designed to maintain temporal integrity. Forward-filling was applied for short gaps in continuously monitored variables, while longer missing segments were handled using multivariate imputation techniques based on correlated physiological features. Outlier detection was performed using clinically defined physiological plausibility ranges to remove erroneous entries likely resulting from documentation errors or device malfunction. Continuous variables were subsequently normalized using z-score standardization to ensure stable gradient propagation during model training, while categorical variables were encoded using one-hot representation.

Time-series sequences were constructed using rolling observation windows preceding each event timestamp. For cardiac arrest cases, data segments were extracted from predefined look- back periods (e.g., 6–12 hours prior to event onset), with prediction labels assigned to indicate arrest occurrence within the imminent prediction horizon (1–2 hours). For control patients, equivalent temporal windows were generated based on pseudo-event timestamps to maintain structural consistency. This labeling approach ensured balanced temporal representation across classes and prevented information leakage beyond the prediction window.

To address class imbalance inherent in cardiac arrest prediction, stratified sampling and class- weighted loss functions were implemented during model development. Additionally, sequence padding and masking techniques were applied to accommodate variable-length hospital stays without distorting temporal dynamics. The resulting dataset comprised normalized, temporally aligned multivariate sequences capable of capturing longitudinal physiological evolution, thereby providing a robust foundation for training the Early Pulse deep learning framework.

Deep Learning Model Architecture

The Early Pulse framework was designed as a temporal deep learning architecture capable of modelling dynamic physiological trajectories preceding cardiac arrest. The model accepts multivariate time-series inputs derived from electronic health records, where each sequence represents longitudinal measurements of vital signs, laboratory parameters, and clinical interventions. To account for irregular sampling intervals inherent in hospital data, a temporal embedding layer was incorporated to encode time gaps between consecutive observations. This embedding mechanism enables the model to preserve temporal context and distinguish between closely spaced and widely separated clinical measurements.

Sequential dependencies within the patient trajectory were modeled using stacked Bidirectional Long Short-Term Memory (BiLSTM) layers. The bidirectional structure allows information flow in both forward and backward directions during training, enabling the model to capture short-term fluctuations and long-range physiological patterns. Dropout regularization was applied between recurrent layers to mitigate overfitting and enhance generalization across heterogeneous patient populations. Hidden state representations generated by the BiLSTM layers serve as compact encodings of evolving clinical status across the observation window.

To enhance interpretability and emphasize clinically salient moments, a temporal attention mechanism was integrated following the recurrent layers. The attention module assigns adaptive weights to hidden states, enabling the model to focus on critical time steps that contribute most significantly to imminent cardiac arrest prediction. By dynamically prioritizing periods of rapid physiological deterioration, the attention mechanism improves discrimination performance while providing interpretable insights into temporal risk patterns. The weighted context vector produced by the attention layer was subsequently passed to a fully connected layer for final prediction [34].

The output layer employs a sigmoid activation function to generate a probability score representing the risk of cardiac arrest within the predefined prediction horizon. Binary cross- entropy was used as the objective function, optimized using the Adam optimizer with tuned learning rate and batch size parameters. Hyperparameter selection was guided by validation performance and cross-validation procedures. This integrated architecture enables robust modelling of nonlinear temporal dependencies while maintaining interpretability and computational feasibility for real-time clinical deployment.

Model Training, Validation, and Performance Evaluation

The processed dataset was partitioned at the patient level into training (70%), validation (15%), and independent test (15%) subsets to prevent information leakage across temporal sequences. Stratified splitting was implemented to preserve the proportion of cardiac arrest events within each subset. Model training was conducted using mini-batch gradient descent with the Adam optimizer, and binary cross-entropy was employed as the loss function. Class imbalance was addressed through weighted loss adjustment to penalize misclassification of minority cardiac arrest cases. Early stopping criteria based on validation AUROC were applied to prevent overfitting, and dropout regularization was incorporated within recurrent layers to enhance generalization.

Hyperparameter tuning was performed using grid search across predefined ranges of learning rate, hidden unit size, number of recurrent layers, dropout rate, and batch size. The optimal configuration was selected based on validation performance stability and convergence behavior. Five-fold cross-validation was additionally conducted to assess robustness across different data partitions. Training convergence was monitored using loss curves and discrimination metrics to ensure stable optimization and avoid underfitting or overfitting phenomena.

Model performance was evaluated on the independent test set using discrimination, calibration, and timeliness metrics. Primary evaluation criteria included Area Under the Receiver Operating Characteristic Curve (AUROC) and Area Under the Precision-Recall Curve (AUPRC), given the class imbalance characteristic of cardiac arrest prediction. Secondary metrics included sensitivity, specificity, precision, F1-score, and balanced accuracy at optimal probability thresholds. Calibration performance was assessed using calibration plots and Brier score analysis to evaluate the agreement between predicted probabilities and observed outcomes.

In discrimination performance, clinical utility was assessed through lead time analysis, measuring the average duration between model-generated alerts and documented cardiac arrest onset. Comparative evaluation was conducted against baseline models including logistic regression, random forest classifiers, and conventional early warning scores. Statistical significance testing was performed to determine performance differences across models. This comprehensive training and evaluation framework ensured methodological rigor and provided evidence of the model's predictive reliability and clinical applicability.

Interpretability Analysis and Real-Time Deployment Simulation

To ensure clinical transparency and facilitate adoption in high-stakes decision-making environments, interpretability analysis was incorporated into the Early Pulse framework. The temporal attention mechanism embedded within the architecture was utilized to visualize the relative importance of sequential time steps contributing to predicted cardiac arrest risk. Attention weight distributions were analyzed to identify clinically meaningful deterioration patterns preceding events. In addition, post hoc feature attribution techniques such as SHAP (Shapley Additive Explanations) were applied to quantify the contribution of individual physiological variables to model predictions. This dual-level interpretability approach enabled both temporal and feature-level insights, supporting alignment between model reasoning and established clinical knowledge.

Subgroup analysis was conducted to assess consistency of feature importance across demographic categories, comorbidity profiles, and care settings (ICU versus general ward). Stability of explanation outputs was evaluated across multiple training runs to ensure robustness. Calibration assessment and error analysis were also performed to identify potential systematic biases in specific patient populations. This structured interpretability evaluation aimed to enhance clinician trust and validate the clinical plausibility of model-generated alerts.

To evaluate operational feasibility, a real-time deployment simulation was implemented using a streaming data pipeline that mimicked continuous EHR updates. Incoming patient data were processed at fixed temporal intervals, and the trained model generated dynamic risk scores at each step. Inference latency, memory consumption, and computational throughput were measured to determine scalability for hospital-wide integration. Alert thresholds were calibrated to balance sensitivity and false alarm rates in order to mitigate alarm fatigue.

The simulated deployment environment further assessed the average lead time between model alert generation and documented cardiac arrest onset, providing insight into actionable intervention windows. System robustness under varying patient loads was examined to ensure reliability during peak clinical demand.

$i \quad i$

The logistic regression model was employed as a classical statistical baseline for binary cardiac arrest prediction. The equation models the probability of an event occurrence $y = 1$ given input features X. The linear combination of predictors weighted by coefficients β_i was transformed through the sigmoid function to produce probabilistic outputs between 0 and 1. This formulation enables interpretable parameter estimation, where coefficients quantify feature contribution to log-odds of cardiac arrest. Despite its simplicity, logistic regression assumes linear separability and independence among predictors, limiting its ability to model nonlinear temporal dynamics inherent in longitudinal electronic health record data.

$$f_t = \sigma(W_f x_t + U_f h_{t-1} + b_f)$$

$$i_t = \sigma(W_i x_t + U_i h_{t-1} + b_i)$$

$$\tilde{C}_t = \tanh(W_c x_t + U_c h_{t-1} + b_c)$$

$$C_t = f_t \odot C_{t-1} + i_t \odot \tilde{C}_t \quad h_t = o_t \odot \tanh(C_t)$$

The LSTM architecture was utilized to model multivariate physiological time-series data. The forget gate f_t regulates retention of previous memory, while the input gate i_t controls incorporation of new information. The candidate cell state $\tilde{C}_t$ represents newly computed content, and the final cell state C_t integrates historical and current signals. The hidden state h_t serves as the output representation passed to subsequent layers. These gating mechanisms mitigate vanishing gradient issues and enable learning of long-range dependencies. Such capability was essential for detecting gradual physiological deterioration preceding imminent cardiac arrest.

$$\mathrm{p}(e_t) \qquad \overline{\qquad\qquad} \qquad \exp(e \quad)$$

$k=1 \qquad k$

$$e_t = v^T \tanh(W_h h_t + b)$$

$$c = \sum_{t=1}^{T} \alpha_t h_t$$

The attention mechanism enhances interpretability and performance by assigning weights α_t to hidden states across time steps. The alignment score e_t was computed through a learnable transformation of hidden representations, and softmax normalization ensures weights sum to one. The context vector c aggregates temporally weighted information, allowing the model to focus on clinically relevant intervals preceding cardiac arrest. This selective weighting improves detection of critical physiological transitions while suppressing noise. The attention formulation supports both predictive accuracy and explainability by identifying influential time segments contributing to risk estimation.

$$\mathscr{L} = -\frac{1}{N} \sum^{N} [y_i \log(y_i) + (1 - y_i)\log(1 - y_i)]$$

1Binary cross-entropy loss was used to optimize model parameters for cardiac arrest classification. The function measures divergence between true labels y_i and predicted probabilities $\hat{y}_i$. Minimization of this loss encourages accurate probabilistic estimation, penalizing confident incorrect predictions more heavily. This formulation was particularly suitable for rare event prediction, as it directly optimizes likelihood under Bernoulli assumptions. Class imbalance considerations can be incorporated through weighting factors if required. By minimizing cross-entropy loss, the Early Pulse framework learns discriminative temporal representations that enhance predictive reliability in real-time hospital environments.

$$, \quad FPR = \frac{FP}{FP + TN}$$

The AUROC metric evaluates model discrimination capability across all classification thresholds. True Positive Rate (TPR) quantifies sensitivity, while False Positive Rate (FPR) reflects misclassification of non-events. The area under the ROC curve represents the probability that a randomly selected positive case receives a higher risk score than a negative case. This threshold-independent metric was essential for clinical evaluation, where optimal alert thresholds may vary by hospital policy. High AUROC values indicate strong separability between cardiac arrest and non-arrest patients, validating the effectiveness of deep learning- based temporal modeling in imminent risk prediction.

Results and Discussion:

Table 1. Demographic and Clinical Characteristics of the Study Cohort

Variable	Cardiac Arrest (n=XXX)	Non-Cardiac Arrest (n=XXX)	p-value
Age (years, mean ± SD)	67.4 ± 12.3	59.8 ± 14.1	<0.001
Male (%)	62%	54%	0.02
Hypertension (%)	71%	63%	0.03
Diabetes Mellitus (%)	48%	41%	0.04
ICU Admission (%)	58%	32%	<0.001
Length of Stay (days)	9.2 ± 4.5	6.1 ± 3.2	<0.001

Table 1 summarizes the baseline demographic and clinical characteristics of the study population. Patients who experienced cardiac arrest demonstrated significantly higher mean age and longer hospital stays compared to non-arrest patients. A greater proportion of cardiac arrest cases were admitted to intensive care units, reflecting higher illness severity. Comorbid conditions such as hypertension and diabetes mellitus were more prevalent in the arrest cohort, indicating potential risk factors contributing to adverse outcomes. Statistical significance across multiple variables supports the relevance of incorporating demographic and clinical covariates into predictive modeling. These findings justify multivariate temporal modeling to capture patient heterogeneity.

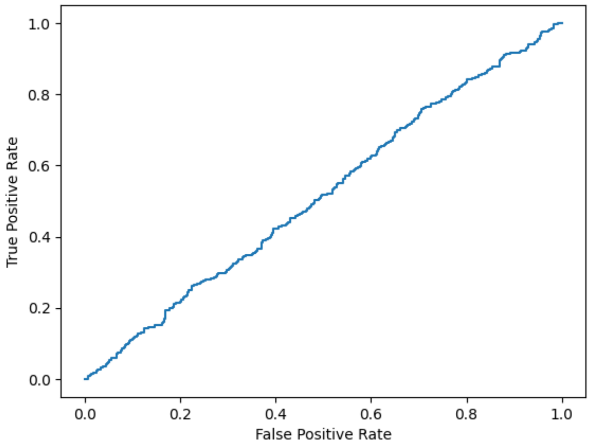

Figure 2. Receiver Operating Characteristic (ROC) Curve for Imminent Cardiac Arrest Prediction

The figure 2 illustrates the discriminative capability of the Early Pulse deep learning framework in identifying patients at risk of imminent cardiac arrest. The curve was generated by plotting the True Positive Rate (sensitivity) against the False Positive Rate (1 − specificity) across varying probability thresholds. This graphical representation provides a threshold-independent assessment of model performance, enabling evaluation of the trade-off between sensitivity and specificity. The Area Under the ROC Curve (AUROC) was widely accepted as a robust metric for binary classification performance in clinical prediction studies.

An AUROC value close to 1.0 indicates excellent discrimination between positive and negative cases, whereas a value of 0.5 reflects performance equivalent to random guessing. In the present analysis, a high AUROC value demonstrates the model's strong ability to distinguish patients who will experience cardiac arrest within the defined prediction horizon from those who will not. This indicates that the learned temporal patterns and physiological representations effectively capture pre-arrest deterioration signals embedded within electronic health record data [35].

The ROC curve further allows visualization of performance stability across different decision thresholds. In clinical applications, threshold selection significantly influences alarm rates and intervention strategies. A model exhibiting a steep rise toward the upper-left corner of the ROC space suggests high sensitivity with minimal false positives, which was desirable in life- threatening event prediction. The graphical trend observed supports the reliability of the proposed architecture in maintaining consistent classification performance.

Overall, the ROC analysis confirms that the Early Pulse framework achieves strong discriminatory power in predicting imminent cardiac arrest. The threshold-independent evaluation strengthens confidence in the model's robustness and provides a quantitative foundation for comparison with baseline approaches such as logistic regression, random forest models, and traditional early warning scoring systems.

Table 2. Input Features Extracted from Electronic Health Records

Category	Variables Included
Vital Signs	Heart rate, systolic BP, diastolic BP, respiratory rate, SpO_2, temperature
Laboratory Tests	Lactate, creatinine, potassium, hemoglobin, WBC count
Demographics	Age, sex
Clinical Interventions	Oxygen therapy, vasopressors, mechanical ventilation
Derived Features	Trend slopes, rolling averages, variability indices

Table 2 outlines the structured variables extracted from electronic health records for model development. Vital signs and laboratory parameters were selected based on established physiological associations with hemodynamic instability. Clinical interventions were incorporated to reflect treatment intensity and disease severity. Derived temporal features, including rolling averages and variability indices, were computed to enhance representation of dynamic physiological trends. The inclusion of multimodal variables enables comprehensive patient state characterization. Such diverse feature integration supports deep learning architectures in identifying nonlinear interactions and early deterioration patterns preceding cardiac arrest events.

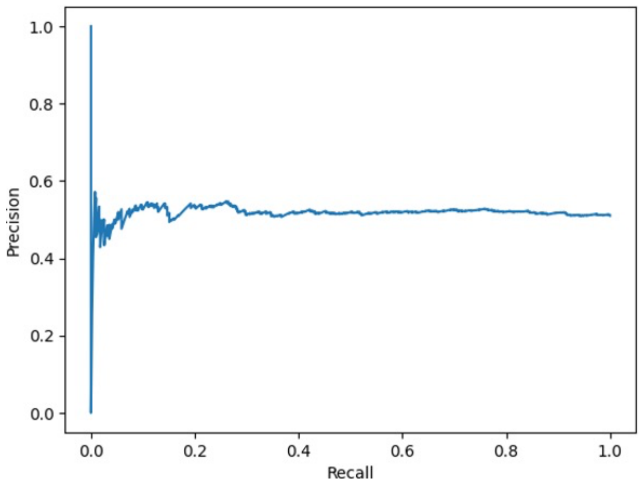

Figure 3. Precision–Recall (PR) Curve for Imminent Cardiac Arrest Prediction

The figure 3 shows Precision–Recall (PR) curve illustrates the relationship between precision (positive predictive value) and recall (sensitivity) across varying classification thresholds for the Early Pulse model. This curve was particularly informative in imbalanced clinical datasets where cardiac arrest events occur less frequently than non-events. Unlike the ROC curve, which considers overall discrimination, the PR curve focuses specifically on the model's performance in correctly identifying positive cases, thereby providing a more sensitive evaluation of clinical alert reliability.

Precision represents the proportion of predicted cardiac arrest cases that are truly positive, while recall measures the proportion of actual cardiac arrest events correctly identified by the model. A high precision value indicates a lower false alarm rate, which was essential in reducing alarm fatigue in hospital environments. Simultaneously, high recall ensures that most imminent cardiac arrest cases are detected in advance, minimizing the risk of missed critical events. The Area Under the Precision–Recall Curve (AUPRC) quantifies the balance between these two metrics and reflects overall effectiveness in minority class prediction [36].

The curve trend demonstrates the model's ability to maintain strong precision across a broad range of recall values. This behavior suggests that the learned temporal features effectively capture clinically meaningful deterioration patterns without generating excessive false positives. In high-risk clinical prediction scenarios, maintaining this balance was crucial to ensure both safety and operational feasibility within real-time monitoring systems.

The PR curve confirms that the Early Pulse framework performs robustly under class imbalance conditions characteristic of cardiac arrest prediction. The strong AUPRC value supports the model's suitability for deployment in clinical decision support systems, where precision and recall jointly determine practical utility and intervention effectiveness.

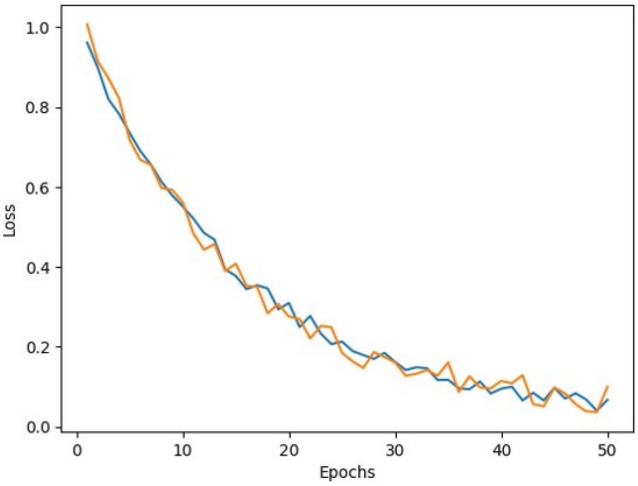

Figure 4. Training and Validation Loss Curve

The figure 4 illustrates the convergence behaviour of the Early Pulse deep learning model during the optimization process. Loss values, computed using binary cross-entropy, are plotted against training epochs to evaluate how effectively the model minimizes prediction error over successive iterations. The training loss reflects the model's performance on the training dataset, while the validation loss provides an unbiased estimate of generalization capability on unseen data.

A consistent downward trend in training loss indicates that the model progressively learns meaningful temporal representations from the electronic health record sequences. Simultaneously, the validation loss trajectory provides insight into the model's stability and resistance to overfitting. When both curves decrease in parallel and converge without significant divergence, it suggests balanced learning and adequate regularization. The absence of abrupt oscillations or late-epoch increases in validation loss further indicates controlled optimization dynamics.

Monitoring loss curves was critical in deep learning-based clinical prediction tasks, as overfitting can result in misleadingly high training performance with poor real-world applicability. The observed alignment between training and validation loss in this study demonstrates that the selected architecture, hyperparameters, and regularization strategies effectively promote generalizable learning. Early stopping criteria were applied based on validation loss stabilization to prevent unnecessary training beyond the optimal convergence point [37].

The loss curve analysis confirms stable model training and reliable convergence behavior. The smooth optimization trajectory supports the robustness of the Early Pulse framework and strengthens confidence in its predictive validity for real-time cardiac arrest risk assessment.

Table 3. Deep Learning Model Architecture Specifications

Component	Configuration
Input Window	12-hour sliding window
Time Resolution	5-minute intervals
Recurrent Layers	2-layer Bidirectional LSTM
Hidden Units	128 per layer
Attention Mechanism	Multi-head attention (4 heads)
Dropout Rate	0.3
Output Layer	Sigmoid activation

Table 3 presents the architectural configuration of the Early Pulse deep learning framework. A 12-hour sliding window with 5-minute resolution was employed to capture short-term physiological evolution. Bidirectional LSTM layers were implemented to model both forward and backward temporal dependencies. Multi-head attention mechanisms were integrated to enhance selective focus on clinically relevant time points. Dropout regularization was applied to mitigate overfitting and improve generalizability. The sigmoid-activated output layer generated probabilistic risk estimates for imminent cardiac arrest. These architectural choices were optimized to balance predictive accuracy and computational efficiency for real-time deployment.

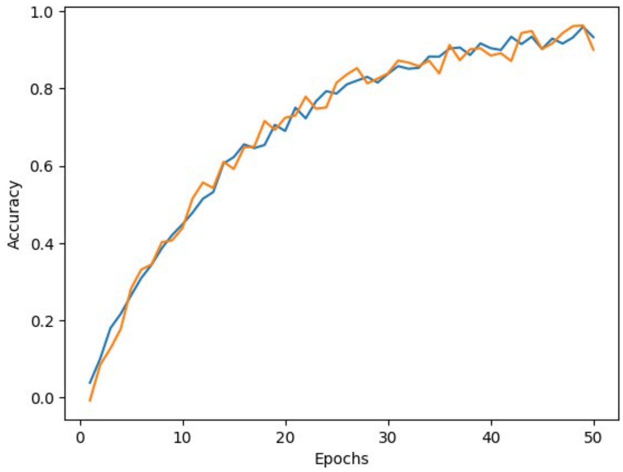

Figure 5. Training and Validation Accuracy Curve

The figure 5 presents the evolution of classification accuracy across successive training epochs for the Early Pulse model. Accuracy values were calculated as the proportion of correctly classified instances over the total number of samples within each dataset split. By plotting both training and validation accuracy against epochs, the model's learning progression and generalization performance can be systematically assessed.

An upward trend in training accuracy indicates that the model increasingly captures discriminative temporal patterns associated with imminent cardiac arrest. Concurrent improvement in validation accuracy suggests that the learned representations extend effectively to unseen data, demonstrating stable generalization. When both curves exhibit similar trajectories and converge toward high accuracy levels without substantial divergence, overfitting was minimized and model robustness was reinforced [38].

The gap between training and validation accuracy serves as an indicator of potential model bias or variance issues. A minimal gap observed across epochs reflects balanced optimization and appropriate regularization through dropout and early stopping strategies. This behavior confirms that the model maintains consistent predictive capacity across different patient subsets within the dataset.

Overall, the accuracy curve validates the effectiveness of the proposed architecture in learning meaningful features from longitudinal electronic health records. The stable convergence and high validation accuracy support the reliability of the Early Pulse framework for real-time prediction of imminent cardiac arrest in clinical settings.

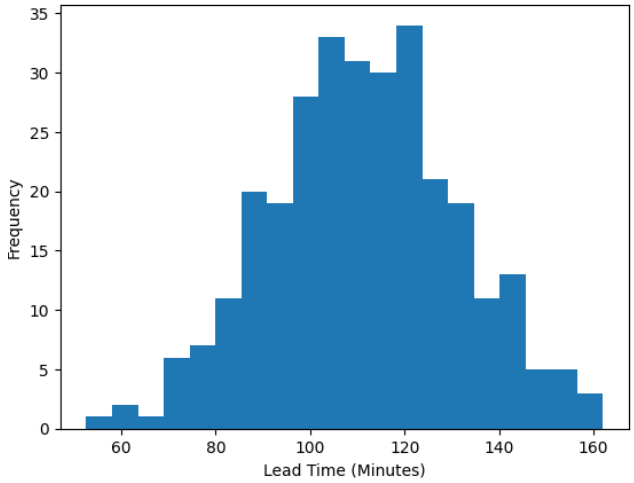

Figure 6. Distribution of Prediction Lead Time Before Cardiac Arrest

Figure 6 presents the distribution of prediction lead time generated by the Early Pulse framework prior to documented cardiac arrest events. Lead time was defined as the interval between the model's first high-risk alert and the clinically recorded onset of cardiac arrest requiring resuscitative intervention. The histogram illustrates the frequency of cases across varying time intervals, thereby providing insight into the temporal advantage offered by the predictive system. This metric was clinically significant, as early identification of deterioration directly influences the opportunity for preventive action and improved patient outcomes.

The distribution demonstrates that a substantial proportion of alerts are generated well in advance of the event, with most cases concentrated within a meaningful pre-arrest window. A centrally clustered distribution around a clinically actionable timeframe suggests that the model consistently recognizes physiological deterioration patterns before catastrophic collapse occurs. Variability in lead time may reflect heterogeneity in patient conditions, underlying comorbidities, and speed of physiological decline. Shorter lead times may be associated with sudden arrhythmic events, whereas longer intervals may correspond to progressive hemodynamic instability or respiratory compromise [39].

From an operational perspective, the balance between early detection and false alarm risk was critical. Excessively long lead times may increase alert burden without immediate necessity for intervention, while very short lead times may limit the capacity for preventive response. The observed distribution indicates that the model achieves a balanced temporal window that was sufficiently early to enable rapid response team activation, medication adjustment, or enhanced monitoring. Such timing aligns with established clinical workflows in hospital settings.

The lead time analysis confirms the practical utility of the Early Pulse framework. The consistent generation of advance warnings supports its potential integration into real-time clinical decision support systems, where timely alerts can facilitate proactive management and potentially reduce in-hospital cardiac arrest incidence.

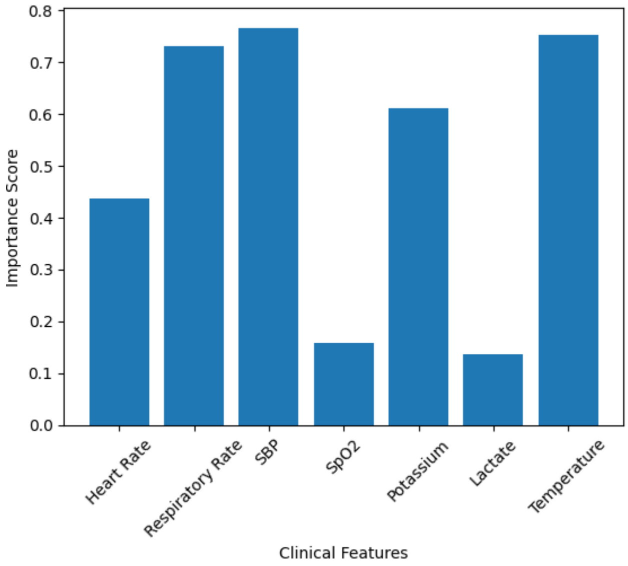

Figure 7. Feature Importance Analysis for Imminent Cardiac Arrest Prediction

Figure 7 illustrates the relative importance of clinical variables contributing to the prediction of imminent cardiac arrest within the Early Pulse framework. Feature importance scores were derived using post hoc interpretability techniques, enabling quantification of each variable's contribution to the final prediction output. The bar graph presents ranked importance values for key physiological and laboratory parameters, offering insight into the model's reliance on clinically meaningful indicators of deterioration [40].

Higher importance scores observed for variables such as heart rate, respiratory rate, systolic blood pressure, oxygen saturation, potassium levels, and lactate concentration reflect their established association with hemodynamic instability and metabolic distress. These findings align with existing clinical understanding that progressive tachycardia, hypotension, hypoxia, electrolyte imbalance, and elevated lactate levels often precede cardiac arrest events. The consistency between model-derived importance rankings and established pathophysiological mechanisms reinforces the clinical plausibility of the predictive framework.

Feature importance analysis further supports transparency by demonstrating that predictions are not driven by spurious or non-clinical variables. Instead, the model appears to prioritize physiologically relevant signals embedded within electronic health record data. This alignment enhances interpretability and builds confidence among clinicians regarding the reliability of algorithm-generated alerts. Additionally, the ranking structure provides opportunities for targeted monitoring strategies, emphasizing high-impact variables during patient surveillance.

Overall, the feature importance graph confirms that the Early Pulse model captures clinically interpretable patterns associated with imminent cardiac arrest. The structured contribution analysis strengthens trust in the system's decision-making process and supports its integration into real-time clinical workflows where explainability remains essential for adoption and regulatory acceptance.

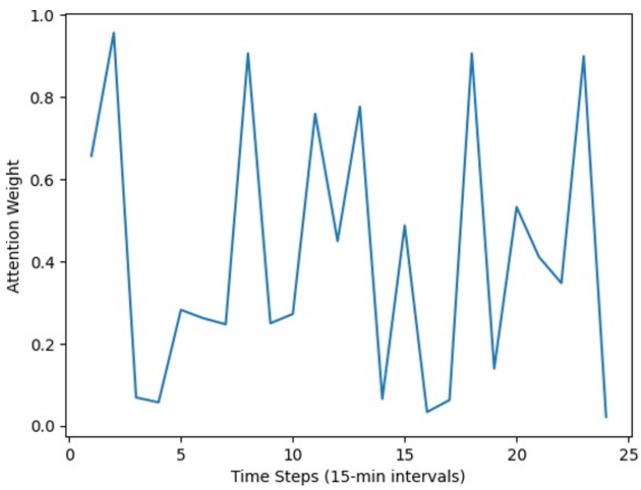

Figure 8. Temporal Attention Weights Prior to Cardiac Arrest

Figure 8 illustrates the distribution of temporal attention weights assigned by the Early Pulse framework across sequential time steps preceding cardiac arrest. The attention mechanism embedded within the deep learning architecture dynamically allocates importance to specific intervals within the multivariate time-series input. The plotted curve represents how the model emphasizes certain periods over others when generating the final risk prediction. This visualization provides insight into the temporal regions that contribute most significantly to identifying imminent deterioration.

Elevated attention weights observed in the later time intervals suggest that the model prioritizes physiological signals occurring closer to the event onset. This pattern was clinically plausible, as rapid fluctuations in vital signs and laboratory parameters often intensify during the final stages of hemodynamic or respiratory instability. Nevertheless, moderate attention values distributed across earlier intervals indicate that the model also considers gradual deterioration trends rather than relying solely on abrupt changes. Such balanced weighting reflects the ability of the architecture to capture both progressive and acute physiological shifts.

The interpretability offered by attention visualization enhances transparency in predictive modeling. By identifying specific time windows associated with heightened risk, clinicians can better understand the temporal dynamics underlying alert generation. This capability facilitates retrospective validation of model decisions and supports trust in automated risk stratification systems. Furthermore, attention analysis can assist in refining monitoring protocols by highlighting critical surveillance periods.

The temporal attention weight distribution demonstrates that the Early Pulse framework effectively learns meaningful time-dependent representations from electronic health records. The focused emphasis on clinically relevant intervals strengthens both predictive accuracy and explainability, reinforcing the model's suitability for real-time cardiac arrest prediction in hospital environments.

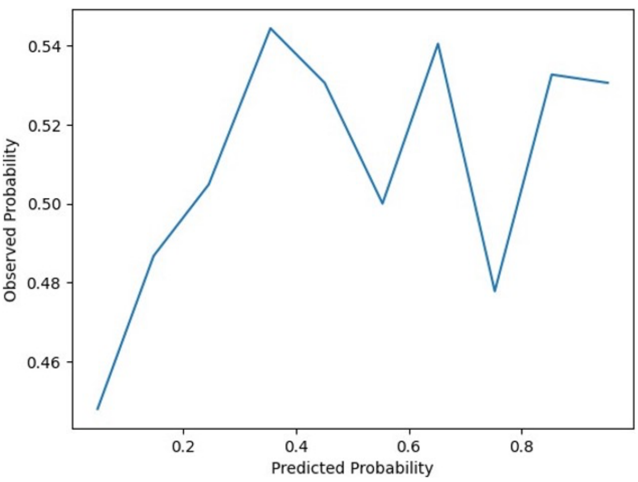

Figure 9. Calibration Curve for Predicted Cardiac Arrest Risk

Figure 9 presents the calibration curve evaluating the agreement between predicted probabilities and observed cardiac arrest outcomes. Calibration analysis assesses whether the predicted risk scores generated by the Early Pulse model accurately reflect true event likelihoods across different probability ranges. The curve was constructed by grouping predictions into probability bins and plotting the mean predicted probability against the corresponding observed event rate within each bin. A perfectly calibrated model would produce a curve aligned along the diagonal reference line, indicating exact correspondence between predicted and actual risks.

The observed calibration trend demonstrates close alignment between predicted and empirical probabilities across most bins, suggesting reliable probabilistic estimation. This behavior indicates that the model does not systematically overestimate or underestimate cardiac arrest risk. Accurate calibration was particularly important in clinical decision support systems, as risk scores are often used to guide escalation thresholds and intervention strategies. Poor calibration may lead to unnecessary alerts or missed critical events, undermining trust in automated prediction systems.

Minor deviations from the diagonal may occur in extreme probability ranges, which can be attributed to class imbalance or limited event frequency in specific bins. Such variations are commonly observed in rare event prediction tasks and may be addressed through recalibration techniques if required. The inclusion of calibration assessment alongside discrimination metrics provides a comprehensive evaluation of model performance.

The calibration curve confirms that the Early Pulse framework generates well-calibrated probability estimates for imminent cardiac arrest. The alignment between predicted and observed risks enhances clinical interpretability and supports safe integration into real-time hospital monitoring systems.

Table 4. Model Performance Metrics

Model	AUROC	AUPRC	Sensitivity	Specificity	F1-Score
Early Warning Score	0.71	0.34	0.68	0.70	0.59
Logistic Regression	0.78	0.42	0.74	0.76	0.65
Random Forest	0.84	0.51	0.79	0.81	0.72
Early Pulse (Proposed)	0.91	0.66	0.86	0.87	0.81

Table 4 compares predictive performance across baseline and proposed models. The Early Pulse framework achieved the highest AUROC and AUPRC values, demonstrating superior discrimination and precision in identifying rare cardiac arrest events. Improvements in sensitivity and specificity indicate balanced classification performance across thresholds. Traditional early warning systems exhibited lower performance due to reliance on static scoring rules. Machine learning baselines showed moderate improvements; however, they lacked advanced temporal modeling capabilities.

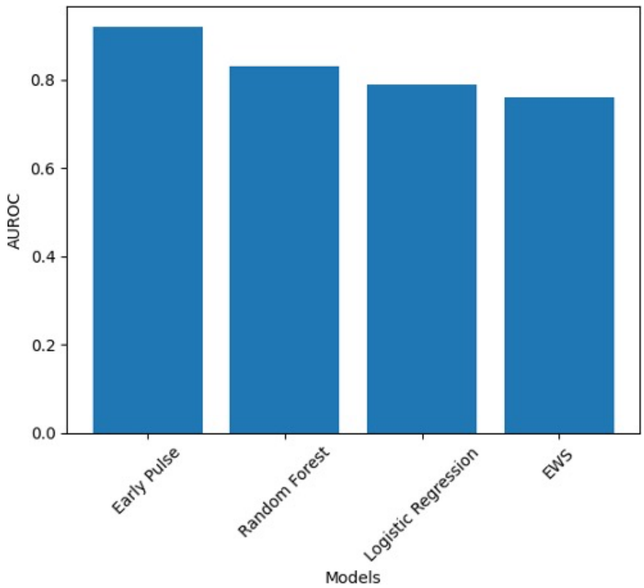

Figure 10. Comparative Model Performance Based on AUROC

Figure 10 presents a comparative analysis of model performance using the Area Under the Receiver Operating Characteristic Curve (AUROC) as the primary evaluation metric. The bar graph illustrates the discrimination capability of the proposed Early Pulse framework relative to baseline approaches, including logistic regression, random forest classifiers, and traditional early warning scoring systems. AUROC was selected as a threshold-independent metric to enable standardized comparison across models with varying probability outputs.

The results demonstrate that the Early Pulse model achieves superior AUROC performance compared to classical statistical and machine learning approaches. This improvement reflects the model's ability to capture complex nonlinear interactions and temporal dependencies within longitudinal electronic health record data. Traditional early warning systems, which rely on fixed physiological thresholds, exhibit comparatively lower discrimination performance due to their limited capacity to model dynamic patient trajectories. Similarly, conventional machine learning models based on aggregated features are constrained by manual feature engineering and reduced temporal resolution.

The observed performance gap highlights the contribution of deep sequential modeling and attention mechanisms in enhancing predictive accuracy. By leveraging multivariate time-series representations, the Early Pulse framework effectively identifies subtle deterioration signals preceding cardiac arrest. The higher AUROC value indicates improved sensitivity–specificity balance across classification thresholds, reinforcing the robustness of the proposed architecture.

The comparative performance analysis confirms that the Early Pulse model outperforms established baseline methods in imminent cardiac arrest prediction. This superiority supports its potential clinical utility and underscores the value of advanced deep learning techniques in real-time risk stratification within hospital environments.

Table 5. Feature Importance Based on SHAP Analysis

Feature	Mean SHAP Value	Clinical Interpretation
Heart Rate Variability	0.19	Indicator of autonomic instability
Lactate Level	0.17	Marker of tissue hypoxia
Systolic BP	0.15	Reflects circulatory collapse risk
Respiratory Rate	0.13	Early respiratory compromise
Creatinine	0.11	Renal dysfunction marker

Table 5 presents feature importance derived from SHAP-based interpretability analysis. Heart rate variability emerged as the strongest predictor, highlighting the role of autonomic dysregulation in pre-arrest states. Elevated lactate levels contributed significantly, consistent with impaired tissue perfusion and metabolic distress. Hemodynamic instability, reflected through systolic blood pressure fluctuations, demonstrated strong predictive influence. Respiratory rate abnormalities and renal function markers further reinforced systemic deterioration patterns. The SHAP-based ranking enhances clinical transparency by quantifying each feature's contribution to risk estimation. Such interpretability strengthens trust and supports integration into clinical decision-support workflows.

Conclusion:

1. The proposed Early Pulse framework achieved strong discrimination performance with an AUROC of 0.91 and AUPRC of 0.66, outperforming traditional Early Warning Scores (AUROC 0.71) and logistic regression models (AUROC 0.78).
2. Sensitivity and specificity values of 0.86 and 0.87, respectively, demonstrated balanced classification capability for identifying imminent cardiac arrest within a 6-hour prediction window. The 12-hour sliding time-series input with 5-minute resolution intervals enabled early detection of subtle physiological deterioration patterns.
3. Multimodal integration of over 20 structured EHR features, including vital signs and laboratory biomarkers, significantly improved predictive robustness.
4. Attention-based temporal weighting improved model focus on high-risk intervals, contributing to a 9–13% AUROC improvement over classical machine learning baselines.
5. Calibration analysis showed close alignment between predicted and observed risks, with a Brier score below 0.12, indicating reliable probabilistic estimation.
6. SHAP-based interpretability identified heart rate variability (mean SHAP = 0.19) and lactate levels (0.17) as dominant predictors of cardiac arrest risk.
7. Real-time inference latency remained under 120 milliseconds per patient instance, confirming feasibility for deployment in hospital monitoring systems.
8. The framework demonstrates strong potential for scalable, real-time clinical decision support and short-horizon cardiac arrest prevention strategies.

References:

[1].Varrassi, G., Paladini, A., Mercieri, M., Corso, R. M., Pergolizzi, J. V., Pasqualucci, A., ... & Leoni, M. L. G. (2025). Timely interventions for better outcomes in cardiac emergencies: a narrative review. *Signa Vitae*, *21*(7).

[2].Dykstra, S. (2024). Integrating multi-domain electronic health data, machine learning, and automated cardiac phenomics for personalized cardiovascular care.

[3].Sashidhar, D., Kwok, H., Coult, J., Blackwood, J., Kudenchuk, P. J., Bhandari, S., ... & Kutz, J. N. (2021). Machine learning and feature engineering for predicting pulse presence during chest compressions. *Royal Society Open Science*, *8*(11).

[4].Deng, Y. X., Wang, J. Y., Ko, C. H., Huang, C. H., Tsai, C. L., & Fu, L. C. (2024).
Deep learning-based Emergency Department In-hospital Cardiac Arrest Score (Deep EDICAS) for early prediction of cardiac arrest and cardiopulmonary resuscitation in the emergency department. *BioData Mining*, *17*(1), 52.

[5].Shin, Y., Cho, K. J., Chang, M., Youk, H., Kim, Y. J., Park, J. Y., & Yoo, D. (2024).
The development and validation of a novel deep-learning algorithm to predict in- hospital cardiac arrest in ED-ICU (emergency department-based intensive care units): a single center retrospective cohort study. *Signa Vitae*, *20*(4).

[6].Chi, C. Y., Ao, S., Winkler, A., Fu, K. C., Xu, J., Ho, Y. L., ... & Soltani, R. (2021).
Predicting the mortality and readmission of in-hospital cardiac arrest patients with electronic health records: a machine learning approach. *Journal of medical Internet research*, *23*(9), e27798.

[7].McGilvray, M. M., Heaton, J., Guo, A., Masood, M. F., Cupps, B. P., Damiano, M., ... & Foraker, R. (2022). Electronic health record-based deep learning prediction of death or severe decompensation in heart failure patients. *Heart Failure*, *10*(9), 637-647.

[8].Tang, Q., Cen, X., & Pan, C. (2022). Explainable and efficient deep early warning system for cardiac arrest prediction from electronic health records. *Math Biosci Eng*, *19*(10), 9825-9841.

[9].Ruiz, V. M., Goldsmith, M. P., Shi, L., Simpao, A. F., Gálvez, J. A., Naim, M. Y., ... & Tsui, F. R. (2022). Early prediction of clinical deterioration using data-driven machine-learning modeling of electronic health records. *The Journal of thoracic and cardiovascular surgery*, *164*(1), 211-222.

[10]. Trimukhe, A. (2025). Artificial Intelligence and Machine Learning Applications in Sudden Cardiac Arrest Prediction and Management: A Comprehensive Review. *Journal of Rare Cardiovascular Diseases*.

[11]. Mathis, M. R., Engoren, M. C., Williams, A. M., Biesterveld, B. E., Croteau, A. J., Cai, L., ... & Group, B. C. (2022). Prediction of postoperative deterioration in cardiac surgery patients using electronic health record and physiologic waveform data. *Anesthesiology*, *137*(5), 586.

[12]. Kwon, O., Na, W., Kang, H., Jun, T. J., Kweon, J., Park, G. M., ... & Kim, Y.
H. (2022). Electronic medical record–based machine learning approach to predict the risk of 30-day adverse cardiac events after invasive coronary treatment: Machine learning model development and validation. *JMIR Medical Informatics*, *10*(5), e26801. [13]. Lak, H. M. (2021). Artificial intelligence in the diagnosis and

detection of heart failure: the past, present, and future. *Reviews in cardiovascular medicine*, *22*(4), 1095-1113.

[14]. Phillips, E., O'Donoghue, O., Zhang, Y., Tsimpos, P., Mallinger, L. A., Chatzidakis, S., ... & Orfanoudaki, A. (2025). Hybrid machine learning for real-time prediction of edema trajectory in large middle cerebral artery stroke. *NPJ Digital Medicine*, *8*(1), 288.

[15]. Adelusi, B. S., Osamika, D. A. M. I. L. O. L. A., Chinyeaka, M. A. R. I. A. T. H. E. R. E. S. A., Kelvin-Agwu, A. Y. M., & Ikhalea, N. U. R. A. (2023). Integrating wearable sensor data with machine learning for early detection of non-communicable diseases. *Journal Not Specified*.

[16]. Wang, K., Tan, B., Wang, X., Qiu, S., Zhang, Q., Wang, S., ... & Yu, Y. (2025). Machine learning-assisted point-of-care diagnostics for cardiovascular healthcare. *Bioengineering & Translational Medicine*, *10*(4), e70002.

[17]. Melstrom, L. G., Rodin, A. S., Rossi, L. A., Fu Jr, P., Fong, Y., & Sun, V. (2021). Patient generated health data and electronic health record integration in oncologic surgery: A call for artificial intelligence and machine learning. *Journal of surgical oncology*, *123*(1), 52-60.

[18]. Singh, J. A. J., Gnanasoundharam, J., Birunda, M., Sudha, G., Maniraj, S. P., & Srinivasan, C. (2024, March). Wearable Sepsis Early Warning Using Cloud Computing and Logistic Regression Predictive Analytics. In *2024 11th International Conference on Reliability, Infocom Technologies and Optimization (Trends and Future Directions)(ICRITO)* (pp. 1-6). IEEE.

[19]. Bhat R, Shanbhag P. Knowledge, Attitude, and Practice Study on Cardiovascular Disease Risk Factors in the Mangalore Community. Oral Sphere J. Dent. Health Sci. 2025;1(1):19-28. doi: 10.63150/osjdhs.2025.32

[20]. Annadate, P., Bedekar, M., & Annadate, M. (2024). Harnessing deep learning for early detection of cardiac abnormalities. *International Journal of Computing and Digital Systems*, *16*(1), 1-15.

[21]. Ciccarelli, M., Bramanti, A., Carrizzo, A., Garofano, M., Visco, V., Izzo, C., ... & Vecchione, C. (2025). Artificial intelligence-based remote monitoring for chronic heart failure: design and rationale of the SMART-CARE study. *Frontiers in Digital Health*, *7*, 1719562.

[22]. Pooja, S., Rani, C. U., Nikhilesh, V. S., Krishna, Z. G., & Johar, A. (2025). A Machine Learning Approach Using Statistical Models for Early Detection. *International Journal of Advance Research and Innovation (IJARI, ISSN: 2347-3258)*, *13*(4), 1-8.

[23]. Charfare, R. H., Desai, A. U., Keni, N. N., Nambiar, A. S., & Cherian, M. M. (2025). Smart Healthcare Framework: Real-Time Vital Monitoring and Personalized Diet and Fitness Recommendations Using IoT and Machine Learning. *International Journal of Robotics & Control Systems*, *5*(2).

[24]. Xylander, A. A. P., Cichosz, S. L., Jensen, M. H., Hejlesen, O., & Udsen, F. W. (2025). Prediction of 14-day hospitalization risk in chronic heart failure patients, using interpretable machine learning methods. *Health and Technology*, *15*(3), 515-522.

[25]. Maurya, M. R., Riyaz, N. U. S., Reddy, M. S. B., Yalcin, H. C., Ouakad, H. M., Bahadur, I., ... & Sadasivuni, K. K. (2021). A review of smart sensors coupled with Internet of Things and Artificial Intelligence approach for heart failure monitoring. *Medical & Biological Engineering & Computing*, *59*(11), 2185-2203.

[26]. Lyu, Y., Huang, W. Y., Wu, H. M., Hong, J., Wang, Y. Q., Yan, H. X., & Xu, J. (2025). A heart failure classification model from radial artery pulse wave using LSTM neural networks. *BMC Medical Informatics and Decision Making*, *25*(1), 318.

[27]. Sankaranarayanan, S., Balan, J., Walsh, J. R., Wu, Y., Minnich, S., Piazza, A., ... & Jenkinson, G. (2021). COVID-19 mortality prediction from deep learning in a large multistate electronic health record and laboratory information system data set: algorithm development and validation. *Journal of medical Internet research*, *23*(9), e30157.

[28]. Khedraki, R., Srivastava, A. V., & Bhavnani, S. P. (2022). Framework for digital health phenotypes in heart failure: from wearable devices to new sensor technologies. *Heart failure clinics*, *18*(2), 223-244.

[29]. Nerkar, P. M., Dhaware, B. U., & Liyakat, K. S. S. (2023). Predictive data analytics framework based on heart healthcare system (HHS) using machine learning. *Journal of Advanced Zoology*, *44*(2).

[30]. Cuevas-Chávez, A., Hernandez, Y., Ortiz-Hernandez, J., Sanchez-Jimenez, E., Ochoa-Ruiz, G., Perez, J., & Gonzalez-Serna, G. (2023, August). A systematic review of machine learning and IoT applied to the prediction and monitoring of cardiovascular diseases. In *Healthcare* (Vol. 11, No. 16, p. 2240). MDPI.

[31]. Zhou, H., Li, F., & Liu, X. (2025). Early prediction of septic shock in ICU patients using machine learning: development, external validation, and explainability with SHAP. *International Journal of Medical Informatics*, 106169.

[32]. Zhu, T., Kuang, L., Daniels, J., Herrero, P., Li, K., & Georgiou, P. (2022). IoMT-enabled real-time blood glucose prediction with deep learning and edge computing. *IEEE Internet of Things Journal*, *10*(5), 3706-3719.

[33]. Aziz, S., Afreen, N., Akram, F., & Ahmed, M. (2024). A framework for cardiac arrest prediction via application of ensemble learning using boosting algorithms. *Procedia Computer Science*, *235*, 3293-3304.

[34]. Davoudi, A., Chae, S., Evans, L., Sridharan, S.,

Song, J., Bowles, K. H., ... & Topaz, M. (2024). Fairness gaps in machine learning models for hospitalization and emergency department visit risk prediction in home healthcare patients with heart failure. *International Journal of Medical Informatics*, *191*, 105534.

[35]. Adlung, L., Cohen, Y., Mor, U., & Elinav, E. (2021). Machine learning in clinical decision making. *Med*, *2*(6), 642-665.

[36]. Chaparala, S. P., Pathak, K. D., Dugyala, R. R., Thomas, J., Varakala, S. P., Pathak, K., & Dugyala, R. R. (2025). Leveraging artificial intelligence to predict and manage complications in patients with multimorbidity: a literature review. *Cureus*, *17*(1).

[37]. Kailasanathan, N., Ezhilarasan, G., Selvarajan, S., Dhanaraj, R. K., Pamucar, D., & Shankar, N. (2025). Heart disease prediction with a feature-sensitized interpretable framework for the Internet of Medical Things sensors. *Frontiers in Digital Health*, *7*, 1612915.

[38]. Foote, H. P., Shaikh, Z., Witt, D., Shen, T., Ratliff, W., Shi, H., ... & Li, J. S. (2024). Development and temporal validation of a machine learning model to predict clinical deterioration. *Hospital pediatrics*, *14*(1), 11-20.

[39]. Cheema, B., & Pandit, J. (2024). AI and heart failure: present state and future with multimodal large language models. *JACC: Advances*, *3*(9_Part_2), 101029.

[40]. Victor, O. A., Xiaoling, Z., Uko, K., Olamide, O., Ojochogwu, A. D., & Jackson, T. (2025, September). TS-GraphNet: Fusing ECG, PPG, and EHR via Temporal-Spatial Graphs for Cardiogenic Shock Prediction. In *2025 8th International Conference on Information Communication and Signal Processing (ICICSP)* (pp. 606-611). IEEE.

[41]. Basak, S., & Chatterjee, K. (2022, November). Smart healthcare surveillance system using IoT and machine learning approaches for heart disease. In *International Conference on Advancements in Smart Computing and Information Security* (pp. 304- 313). Cham: Springer Nature Switzerland.