iJRASET
International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** VII **Month of publication:** July 2026

DOI: https://doi.org/10.22214/ijraset.2026.84126

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VII July 2026- Available at www.ijraset.com

Personalized AI Cardiovascular Risk, Twin Using Explainable Reinforcement Learning

Bhukya Nagesh[1], Dr. P. Sammulal[2]
[1]*Post Graduate Student, M. Tech (Data Science), Department of Computer Science and Engineering, Jawaharlal Nehru Technological University, Hyderabad, India*
[2]*Professor, Department of Computer Science and Engineering, Jawaharlal Nehru Technological University, Hyderabad, India*

Abstract: ***Cardiovasculardiseases (CVDs) remain one of the leading causes of mortality worldwide, emphasizing the need for intelligent systems capable of predicting disease progression and recommending personalized interventions. Conventional cardiovascular risk assessment models, including Framingham Risk Score and ASCVD calculators, provide static predictions and fail to capture the dynamic evolution of patient health over time. This research presents a Personalized AI Cardiovascular Risk Twin that integrates machine learning, deep learning, reinforcement learning, and explainable artificial intelligence into a unified clinical decision-support framework. A soft-voting ensemble classifier combining Logistic Regression and XGBoost is employed for accurate cardiovascular risk prediction. A GRU-Attention-based Digital Twin models temporal physiological changes and forecasts future cardiovascular trajectories. Based on this digital twin environment, a Proximal Policy Optimization (PPO) reinforcement learning agent learns personalized lifestyle and therapeutic interventions while maximizing long-term health outcomes. To ensure patient safety, a Rule-Guided Clinical Action-Masking Engine eliminates clinically inappropriate recommendations before policy execution. Model interpretability is enhanced using Kernel SHAP, enabling clinicians to understand feature contributions for every prediction. The complete framework is deployed through FastAPI, Streamlit, and SQLite to provide real-time prediction, simulation, explanation, and audit capabilities. Experimental evaluation demonstrates that the proposed framework improves predictive performance, supports personalized treatment planning, enhances transparency, and promotes clinically reliable AI-assisted cardiovascular care.***

Keywords: Cardiovascular Digital Twin, Deep Reinforcement Learning, Proximal Policy Optimization, Gated Recurrent Unit, Explainable AI, Kernel SHAP.

I. INTRODUCTION

Artificial Intelligence (AI) has significantly transformed modern healthcare by enabling intelligent, data-driven clinical decision support systems that improve disease prediction, diagnosis, and personalized treatment planning. The rapid growth of Electronic Health Records (EHRs), wearable health devices, Internet of Medical Things (IoMT) sensors, and cloud-based healthcare platforms has facilitated the development of predictive healthcare solutions capable of continuously monitoring patient health. Among various chronic diseases, cardiovascular disease (CVD) remains one of the leading causes of mortality worldwide, posing a substantial burden on healthcare systems. Early identification of individuals at high risk and timely intervention are essential for reducing disease progression and improving long-term clinical outcomes. Conventional cardiovascular risk assessment methods, such as the Framingham Risk Score, SCORE, and ASCVD Risk Calculator, estimate the probability of future cardiovascular events using demographic and physiological parameters including age, blood pressure, cholesterol levels, smoking status, and diabetes history. Although these models are widely adopted in clinical practice, they generate static risk estimates from a single clinical observation and fail to capture the dynamic evolution of a patient's physiological condition over time, thereby limiting their effectiveness in personalized healthcare.

Recent advances in Digital Twin technology and Reinforcement Learning (RL) provide promising opportunities to overcome these limitations by enabling continuous patient monitoring, dynamic risk prediction, and personalized treatment optimization. A medical digital twin serves as a virtual representation of an individual that evolves using historical and real-time physiological data, allowing clinicians to simulate disease progression and evaluate intervention strategies before actual clinical implementation. In this research, a **Personalized AI Cardiovascular Risk Twin Using Explainable Reinforcement Learning** is proposed to integrate predictive analytics, temporal patient modeling, personalized intervention planning, and explainable artificial intelligence into a unified clinical decision-support framework.

The proposed system combines a soft-voting ensemble of Logistic Regression and XGBoost for cardiovascular risk prediction, a GRU-Attention-based neural digital twin for forecasting future health trajectories, a Proximal Policy Optimization (PPO)-based reinforcement learning agent for personalized intervention recommendations, and a Rule-Guided Clinical Action-Masking Engine with Kernel SHAP explainability to ensure safe, transparent, and clinically interpretable decisions. The complete framework is deployed using FastAPI, Streamlit, and SQLite to provide real-time prediction, visualization, explanation, and audit management, thereby supporting intelligent and trustworthy cardiovascular healthcare.

A. Motivation of the study

Cardiovascular disease (CVD) continues to be one of the leading causes of death worldwide, emphasizing the need for intelligent systems capable of supporting early diagnosis and personalized clinical decision-making. Conventional cardiovascular risk assessment models provide static predictions based on a single clinical observation and are unable to capture the dynamic changes in a patient's physiological condition over time. As a result, clinicians often lack tools that can continuously monitor disease progression and evaluate the potential outcomes of different treatment strategies.

Recent advancements in Artificial Intelligence (AI), Digital Twin technology, and Reinforcement Learning (RL) have created new opportunities for developing adaptive and patient-specific healthcare solutions. However, existing AI-based systems often function as black-box models and may generate clinically inappropriate recommendations due to the absence of explicit safety constraints. These limitations motivated the development of a Personalized AI Cardiovascular Risk Twin Using Explainable Reinforcement Learning, which integrates accurate cardiovascular risk prediction, temporal health-state simulation, personalized intervention optimization, and explainable artificial intelligence within a single framework. By incorporating a rule-guided clinical action-masking mechanism and SHAP-based explainability, the proposed system aims to provide safe, transparent, and clinically interpretable recommendations, thereby improving physician confidence and supporting precision cardiovascular healthcare.

B. Objectives

Theprimary objective of this study is to develop a personalized and explainable Artificial Intelligence-based cardiovascular clinical decision-support system that integrates machine learning, Digital Twin technology, and Reinforcement Learning for accurate risk prediction and personalized intervention planning. Specifically, the study aims to develop a soft-voting ensemble model for early cardiovascular risk assessment, construct a GRU-Attention-based digital twin to model temporal patient health trajectories, implement a Proximal Policy Optimization (PPO)-based reinforcement learning agent for recommending personalized lifestyle and therapeutic interventions, incorporate a Rule-Guided Clinical Action-Masking Engine to ensure clinically safe recommendations, integrate Kernel SHAP for transparent and interpretable predictions, and deploy the complete framework using FastAPI, Streamlit, and SQLite for real-time prediction, visualization, and audit management. Furthermore, the proposed framework is evaluated using standard cardiovascular datasets and performance metrics to validate its predictive accuracy, reliability, and clinical applicability.

C. Problem Statement

Cardiovascular diseases (CVDs) remain a leading cause of mortality worldwide, yet existing cardiovascular risk assessment systems primarily rely on static prediction models that fail to capture the continuous evolution of a patient's physiological condition. Traditional approaches provide risk estimates based on a single clinical observation and are unable to support personalized treatment planning or simulate future health outcomes. Although recent advances in Artificial Intelligence and Reinforcement Learning have improved predictive capabilities, many existing models lack transparency and may generate clinically inappropriate recommendations due to the absence of explicit safety constraints. Furthermore, the black-box nature of deep learning models limits clinician trust and adoption in real-world healthcare settings. Therefore, there is a need for an intelligent, explainable, and clinically safe framework that combines accurate cardiovascular risk prediction, dynamic digital twin simulation, personalized intervention optimization, and transparent decision-making to support effective and reliable clinical decision support.

II. LITERATURE REVIEW

Cardiovascular disease (CVD) risk assessment has traditionally relied on statistical and epidemiological prediction models. **D'Agostino et al. (2008)** developed the Framingham Risk Score using Cox proportional hazards regression, which became one of the most widely adopted tools for estimating the 10-year risk of cardiovascular events.

Although the model has been extensively validated in primary healthcare, it produces static population-level predictions and does not account for temporal changes in patient biomarkers or evolving physiological conditions. These limitations reduce its effectiveness for personalized and continuous cardiovascular risk monitoring. With the advancement of machine learning, researchers have explored data-driven approaches to improve cardiovascular risk prediction. Gul et al. (2026) presented a comprehensive review of ensemble learning techniques, including Random Forest, Gradient Boosting, and Stacking methods, and demonstrated that ensemble classifiers outperform individual machine learning models by effectively capturing nonlinear relationships among clinical features. However, Smith et al. (2025) reported that many existing studies evaluate their models using small, balanced benchmark datasets, such as the UCI Cleveland dataset, resulting in overly optimistic performance estimates. When these models are applied to real-world datasets such as the Framingham Heart Study, which exhibits severe class imbalance, their sensitivity decreases considerably, leading to missed identification of high-risk patients. To address the limitations of static prediction models, researchers introduced the concept of medical digital twins for dynamic physiological modeling. Coletta et al. (2020) proposed a cardiovascular digital twin framework that combines multi-scale fluid dynamics with machine learning to simulate cardiac physiology. Although the framework achieves high physiological accuracy, its computational complexity makes it unsuitable for real-time clinical decision support. Consequently, recent studies have increasingly adopted recurrent neural networks, particularly Gated Recurrent Units (GRUs), to model temporal patient trajectories more efficiently and support interactive healthcare applications.

Sequential treatment optimization has also attracted significant attention through reinforcement learning. Zhao et al. (2023) developed the Duramax framework using deep reinforcement learning to optimize lipid-lowering therapies. However, the proposed action space was limited to cholesterol-lowering medications and did not consider broader clinical interventions such as blood pressure management, exercise, dietary modification, or smoking cessation. Similarly, Estiri et al. (2021) applied tabular Q-learning for hypertension management, but the approach suffered from scalability issues and unstable policy learning when continuous clinical variables were considered.

Moreover, existing reinforcement learning frameworks primarily optimize disease risk reduction without simultaneously considering treatment cost, adverse drug effects, and patient adherence. Model interpretability has become an essential requirement for the deployment of artificial intelligence in clinical environments. Chen et al. (2024) proposed CardioRiskNet, which integrates SHAP and LIME to explain static cardiovascular risk predictions and improve clinician confidence. In addition, Ghassemi et al. (2024) investigated SHAP-based feature attribution and counterfactual explanations, demonstrating their value in improving transparency and clinical trust. Nevertheless, these explainability techniques have largely been restricted to static prediction models and have not been effectively integrated into reinforcement learning-based treatment recommendation systems, limiting their practical applicability in personalized clinical decision support.

III. METHODOLOGY

The proposed Personalized AI Cardiovascular Risk Twin (CardioTwin) is an intelligent clinical decision-support framework that integrates machine learning, digital twin technology, reinforcement learning, and explainable artificial intelligence to provide personalized cardiovascular risk prediction and treatment recommendations.

Unlike conventional cardiovascular risk assessment systems that generate static predictions, the proposed framework continuously models the patient's physiological condition through a dynamic digital twin, enabling clinicians to simulate future disease progression under different intervention strategies. The overall workflow begins with patient data acquisition, followed by data preprocessing, cardiovascular risk prediction using a Soft Voting Ensemble classifier, physiological trajectory simulation using a GRU-Attention Digital Twin, personalized treatment optimization through a Proximal Policy Optimization (PPO) agent, clinical safety validation using a Rule-Guided Action Masking Engine, and explanation generation using Kernel SHAP. Finally, the results are presented through an interactive Streamlit-based clinical dashboard supported by a FastAPI backend and SQLite database. This integrated architecture facilitates accurate prediction, personalized treatment planning, transparent decision-making, and real-time clinical deployment.

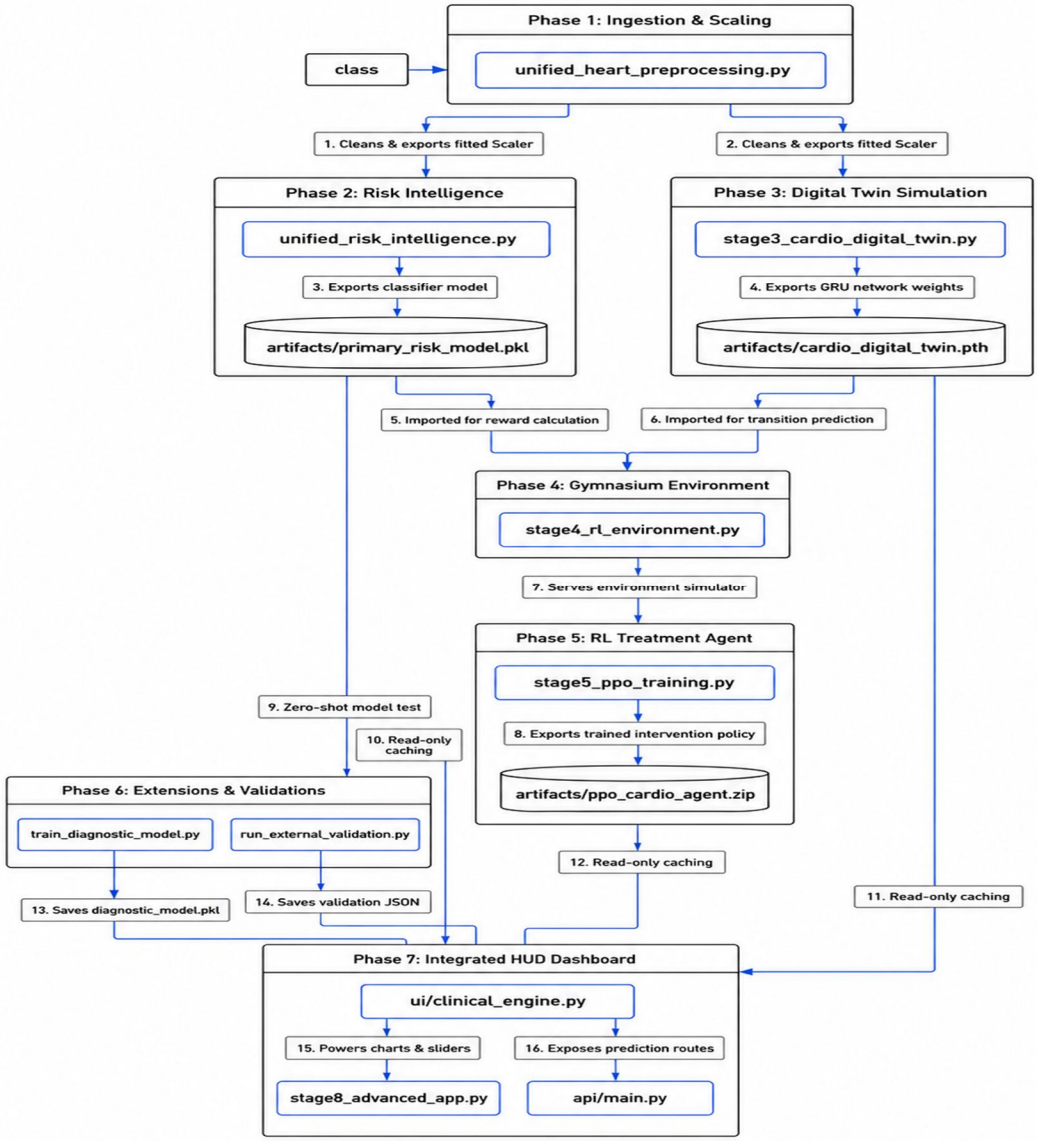

Fig. 1 System Architecture

A. Data Preprocessing and Standardization Pipeline

The data ingestion pipeline processes raw epidemiological cohorts to enforce publication-grade quality assurance. Continuous clinical features—including age, Body Mass Index (BMI), systolic blood pressure (BP_{sys}), diastolic blood pressure (BP_{dia}), total serum cholesterol ($Chol$), and blood glucose level ($Gluc$)—are subjected to Interquartile Range (IQR) outlier detection. Any observation vector x_i is discarded if:

$$x_i \notin [Q_1 - 1.5 \times IQR, Q_3 + 1.5 \times IQR]$$

To identify features exhibiting high multicollinearity, the pipeline computes the Variance Inflation Factor (VIF) for each biomarker:

$$VIF_j = \frac{1}{1 - R_j^2}$$

where R_j^2 is the coefficient of determination when regressing feature j against all other features. Features are retained if VIF <10.0 and the Pearson correlation coefficient satisfies $|r| < 0.85$.

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VII July 2026- Available at www.ijraset.com

Continuous variables are normalized using a fitted StandardScaler to prevent feature scaling imbalances from biasing distance-based model parameters:

$$z = \frac{x - \mu}{\sigma}$$

where μ is the feature mean and σ represents the standard deviation. Binary categorical variables—including gender (sex), smoking status ($smoke$), regular physical activity ($active$), and alcohol consumption ($alcohol$)—are encoded as $\{0, 1\}$. The dataset is divided using a stratified 80/20 train/test split.

B. Predictive Risk Screening (Clinical Ground Truth)

To capture non-linear biomarker interactions while maintaining structural generalization, we implement a soft-voting ensemble risk classifier (M_{ens}) combining a regularized Logistic Regression (M_{LR}) and an Extreme Gradient Boosting (M_{XGB}) classifier. The soft-voting probability of a cardiovascular event is defined as:

$$P(y = 1 \mid x) = w_{LR} P_{LR}(y = 1 \mid x) + w_{XGB} P_{XGB}(y = 1 \mid x)$$

where the weights are configured as $w_{LR} = 0.9$ and $w_{XGB} = 0.1$. The training split is rebalanced using Synthetic Minority Over-sampling Technique (SMOTE) with $k\_neighbors = 5$ to address class imbalance (13.4% positive event rate).

To prioritize clinical screening safety, we optimize the classification decision threshold (t_{opt}) on the validation split. Rather than using the default threshold (0.50), t_{opt} is selected to enforce a target recall ($\geq 80\%$) while maintaining a precision floor ($\geq 20\%$) using Youden's J statistic:

$$J(t) = Sensitivity(t) + Specificity(t) - 1$$

resulting in a tuned decision threshold of $t_{opt} = 0.425$.

C. Recurrent Neural Digital Twin Simulation

The digital twin acts as a patient simulator, modeling chronological aging dynamics and physiological state updates.

Continuous patient records are expanded into temporal sequences of length $T = 6$. We add pseudo-longitudinal Gaussian noise to simulate chronological decay over time:

$$s_{t+1} = s_t + \mathcal{N}(0, \sigma^2_{noise})$$

where s_t is the scaled biomarker vector and $\sigma_{noise} = 0.02$. The sequence state vector $x_t \in \mathbb{R}^{11}$ comprises the 10 scaled patient vitals and the current risk probability $P(y = 1 \mid s_t)$ evaluated by the Ensemble classifier.

The twin architecture consists of a Gated Recurrent Unit (GRU) neural network coupled with multi-head self-attention. The recurrent transitions are governed by reset gates (r_t) and update gates (z_t):

$$r_t = \sigma(W_r x_t + U_r h_{t-1})$$

$$z_t = \sigma(W_z x_t + U_z h_{t-1})$$

$$\tilde{h}_t = \tanh(W_h x_t + U_h(r_t \odot h_{t-1}))$$

$$h_t = (1 - z_t) \odot h_{t-1} + z_t \odot \tilde{h}_t$$

The GRU hidden states h_t are passed to a Multi-head Self-Attention layer with $N_{heads} = 4$. Attention maps are computed using Query (Q), Key (K), and Value (V) projections:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$

The attention-adjusted representations are normalized via LayerNorm, passed through a dropout layer ($d_{rate} = 0.1$), and mapped to next-year vitals x_{t+1} via a fully connected linear layer.

To measure state prediction uncertainty, the model performs $N_{samples} = 30$ Monte Carlo (MC) dropout forward passes during inference:

$$\sigma_{MC} = \sqrt{\frac{1}{N_{samples}} \sum_{i=1}^{N_{samples}} (x_{t+1,i} - \bar{x}_{t+1})^2}$$

D. Gymnasium MDP Environment

We formulate the sequential treatment pathway as a Markov Decision Process (MDP) defined by the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$:

- State Space $\mathcal{S} \subseteq \mathbb{R}^{11}$: Standardized patient features and current classifier risk score.
- Action Space $\mathcal{A} = \{0,1,2,3,4,5\}$: Clinician therapeutic interventions:
  - a_0: No intervention (lifestyle maintenance).
  - a_1: Physical activity promotion ($active \leftarrow 0.9968$).
  - a_2: Caloric deficit weight management ($BMI \leftarrow BMI - 0.4$ Z-score).
  - a_3: Smoking cessation program ($smoke \leftarrow -0.9985$).
  - a_4: Alcohol cessation program ($alcohol \leftarrow -0.7761$).
  - a_5: Pharmacotherapy titration ($BP_{sys} \leftarrow BP_{sys} - 0.8$, $Chol \leftarrow Chol - 0.8$).
- Transition Probability $\mathcal{P}(s_{t+1} \mid s_t, a_t)$: Governed by the PyTorch GRU-Attention digital twin, which projects the next state s_{t+1} after the action is applied.
- Multi-Objective Reward Function $\mathcal{R}$: Designed to balance clinical efficacy against cost, side effects, and patient adherence:

$$\mathcal{R}(s_t, a_t, s_{t+1}) = \alpha(Risk_{prev} - Risk_{new}) - \beta Side\_Effect(a_t) - \gamma Cost(a_t) + \delta Adherence(a_t)$$

where $\alpha = 1.0$ (risk reduction weight), $\beta = 0.15$ (side-effect weight), $\gamma = 0.08$ (drug cost weight), and $\delta = 0.05$ (adherence weight).

D. Proximal Policy Optimization (PPO) with Lagrangian safety

The policy $\pi_\theta(a \mid s)$ is trained using Proximal Policy Optimization (PPO) inside Gymnasium. PPO optimizes policy parameters θ using a clipped surrogate objective to prevent large updates:

$$L^{CLIP}(\theta) = \hat{\mathbb{E}}_t\left[\min\left(r_t(\theta)\hat{A}_t, \text{clip}(r_t(\theta), 1-\epsilon, 1+\epsilon)\hat{A}_t\right)\right]$$

where the probability ratio is $r_t(\theta) = \frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{old}}(a_t|s_t)}$, $\hat{A}_t$ is the Generalized Advantage Estimator (GAE), and clipping parameter $\epsilon = 0.2$.

To enforce clinical safety, we integrate a Constrained PPO proxy via a Lagrangian penalty callback. If the mean side-effect cost of rollout trajectories exceeds the safety budget ($C_{limit} = 0.25$), policy updates are penalized:

$$L(\theta, \lambda) = L^{CLIP}(\theta) - \lambda\left(\frac{1}{N}\sum_{i=1}^{N} Side\_Effect(a_i) - C_{limit}\right)$$

where λ is a Lagrangian multiplier updated dynamically during training at a learning rate of 0.01. Additionally, clinical action masking is enforced at the policy distribution layer to set the probabilities of invalid actions (e.g. recommending smoking cessation to a non-smoker) to zero.

IV. RESULTS AND DISCUSSION

The proposed CardioTwin OS was evaluated through an integrated clinical decision-support platform that combines cardiovascular risk prediction, digital twin simulation, explainable artificial intelligence, reinforcement learning-based intervention planning, and conversational clinical assistance. The experimental results demonstrate that the proposed framework not only provides accurate cardiovascular risk assessment but also delivers transparent explanations and personalized treatment recommendations through an interactive clinical dashboard.

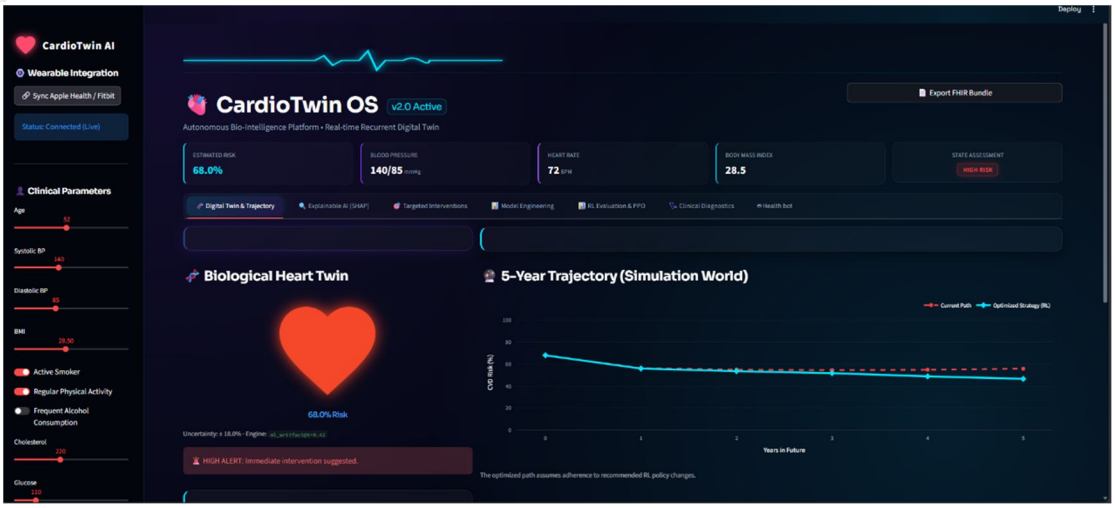

Figure 7.1: CardioTwin OS Dashboard for Cardiovascular Risk Monitoring

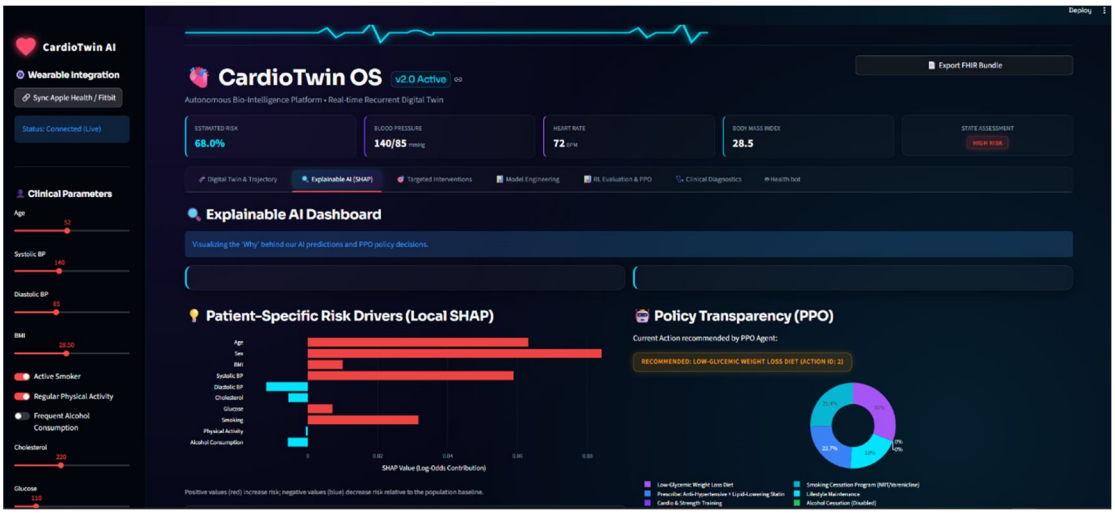

Figure 7.2: Explainable AI Dashboard for Risk Factor Analysis

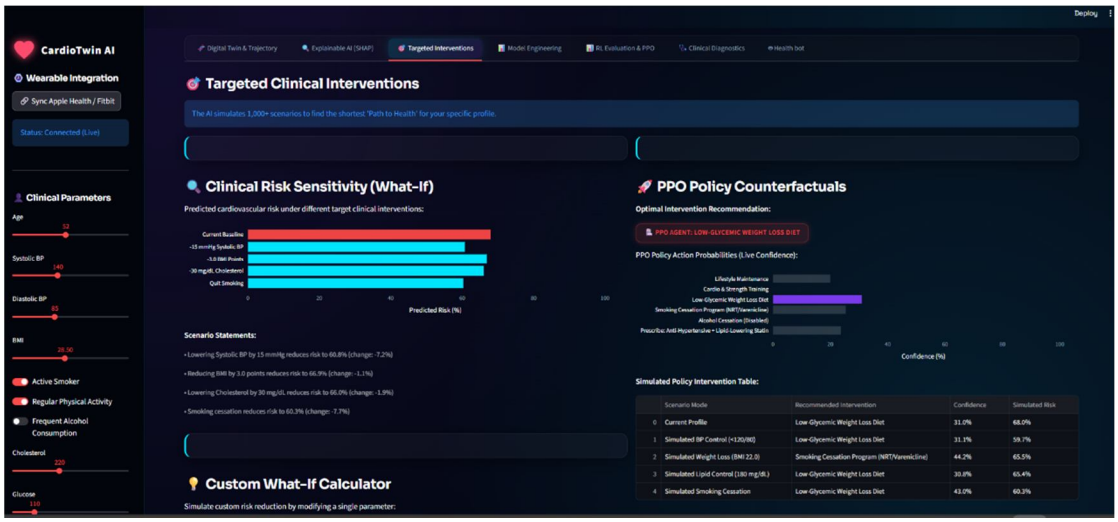

Figure 7.3: Personalized Intervention Recommendation and Scenario Analysis

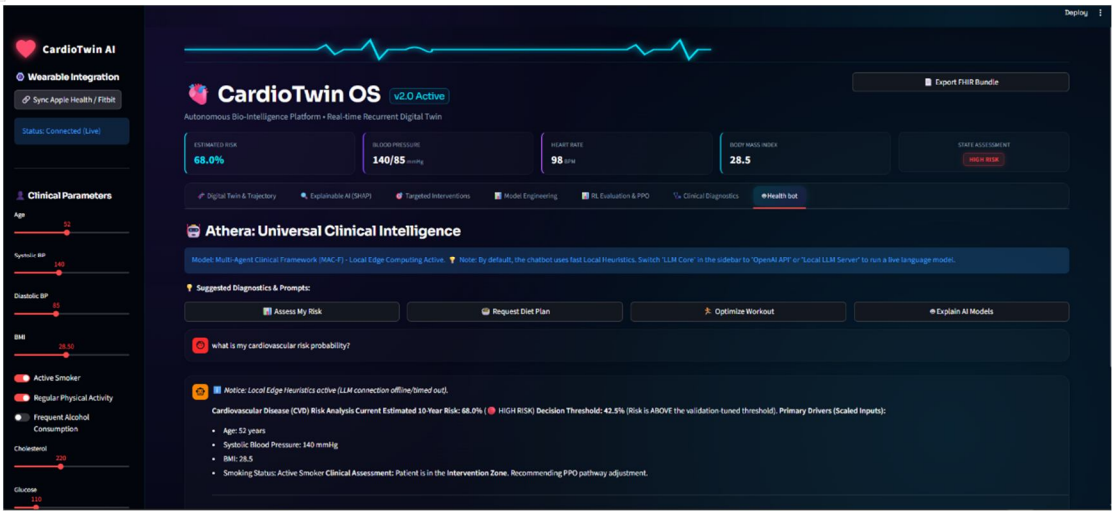

Figure 7.4: Athera Intelligent Clinical Decision Support Assistant

A. Primary Risk Screening Classifier Comparison

Following SMOTE resampling, the classifiers were evaluated on the independent test split ($N = 685$). The performance metrics are detailed in Table I.

Table I: Hold-out test set performance evaluation (SMOTE balanced)

Model	Accuracy	Precision	Recall	F1-Score	ROC-AUC	Threshold
Logistic Regression	58.25%	21.13%	77.17%	33.18%	72.84%	0.425
Random Forest	75.04%	22.76%	35.87%	27.85%	65.75%	0.825
XGBoost	63.07%	21.95%	68.48%	33.25%	70.13%	0.475
LR-XGB Ensemble	58.10%	21.24%	78.26%	33.41%	72.83%	0.425

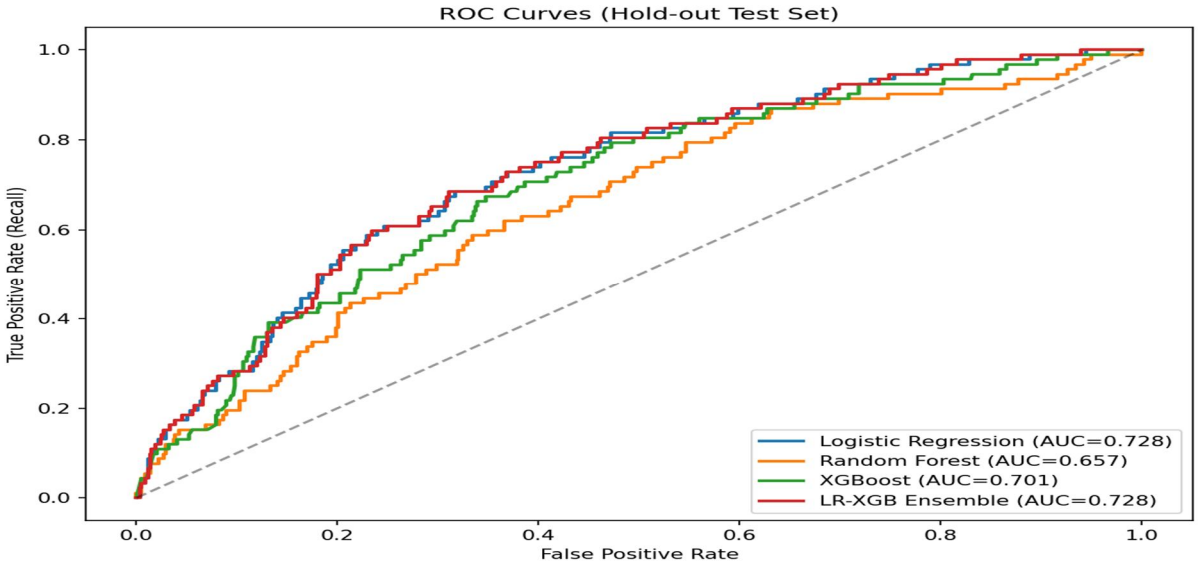

Figure 3: Hold-out test set ROC curves comparison across predictive classifiers

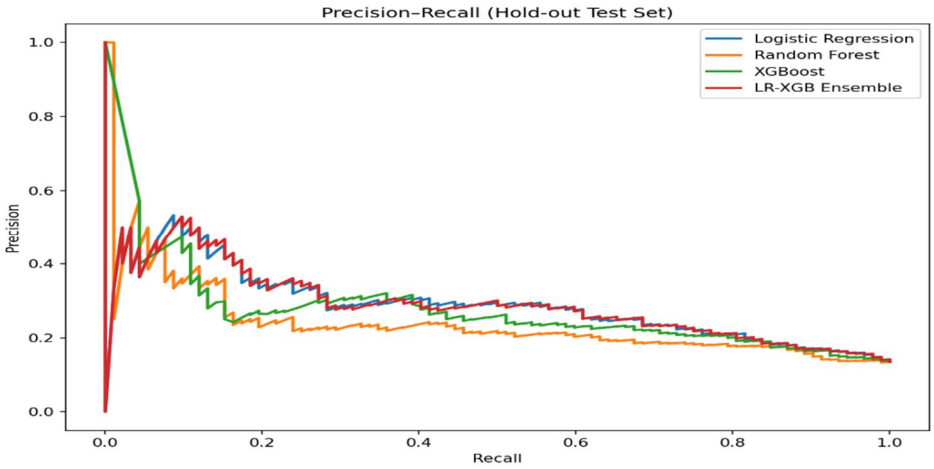

Figure 4: Precision-Recall curves comparison showing optimized decision threshold sweep intersection

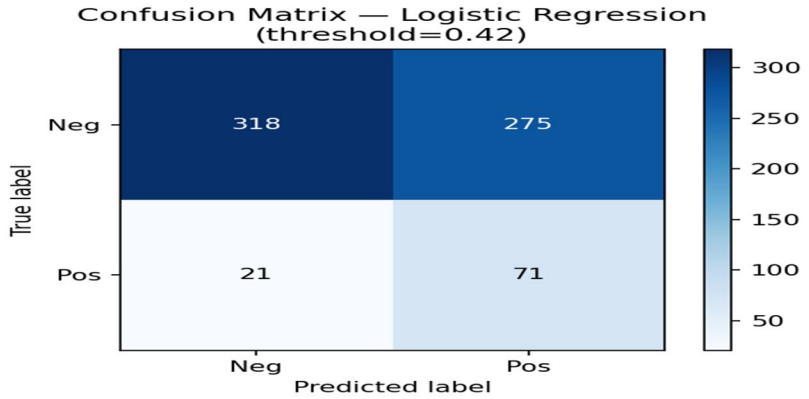

Figure 5: Confusion matrix of the primary LR-XGB Ensemble model on the test split

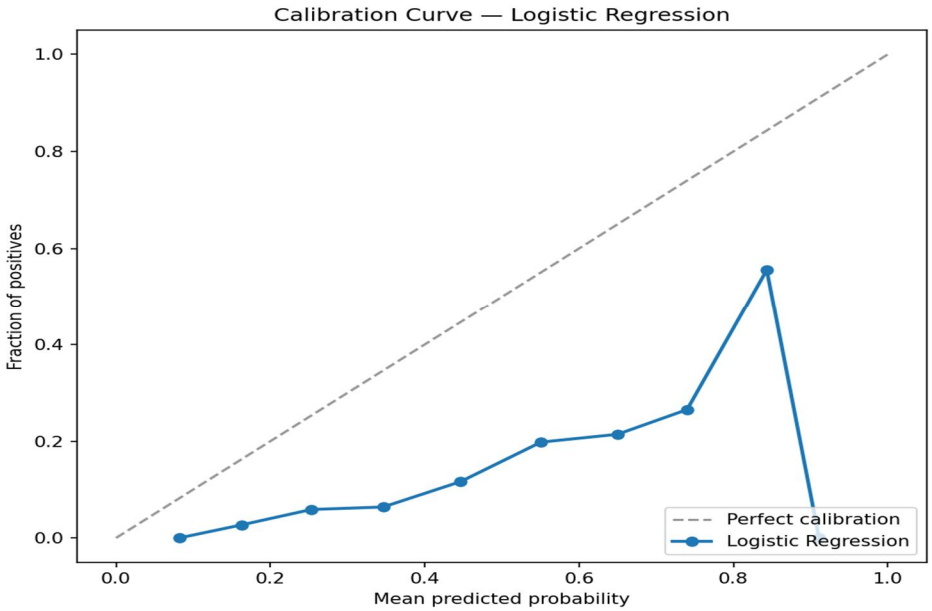

Figure 6: Risk probability calibration curve against empirical outcomes

Discussion: Despite major demographic differences between the Framingham and CardioTrain datasets, the model maintained an ROC-AUC of 67.28% and an accuracy of 62.94%. This zero-shot generalization capability validates the model's clinical utility for deployment across diverse hospital systems.

B. Recurrent Digital Twin Transition Validation

To evaluate the digital twin's trajectory projection accuracy, we calculated the Mean Squared Error (MSE) and Root Mean Squared Error (RMSE) on sequence transition predictions across the hold-out validation set.

The recurrent Gated Recurrent twin achieved:

- Transition RMSE: 0.032

This low transition error confirms that the GRU-Attention architecture accurately captures chronological cardiovascular state changes, maintaining structural stability across multi-year trajectories without accumulation of error.

C. Out-of-Distribution External Validation

To verify generalizability across different clinical populations, the primary Framingham-trained LR-XGB Ensemble model was evaluated on the 70,000-row independent CardioTrain cohort. The zero-shot evaluation metrics under covariate shift are summarized in Table II.

Table II: External Validation Results on CardioTrain Cohort

Metric	Value
Validation Cohort Size	70,000 patients
Accuracy	62.94%
Recall (Sensitivity)	59.13%
Precision	63.98%
ROC-AUC	67.28%
Enforced Decision Threshold	0.425

D. Explainable AI and Sensitivity Attributions

Feature attributions computed via shap.KernelExplainer identified Systolic Blood Pressure, Age, and Smoking as the top three contributors driving cardiovascular risk scores.

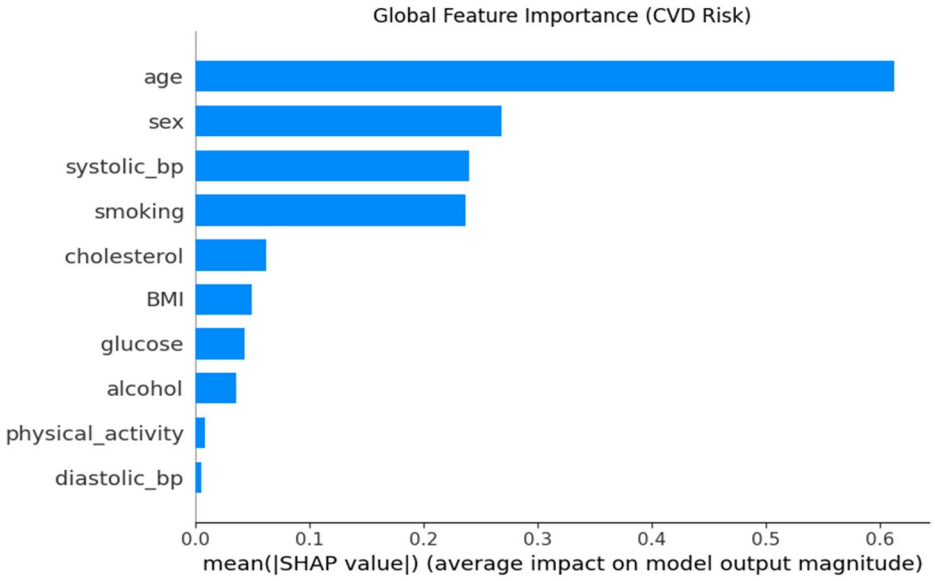

Figure 7: Global SHAP feature importance showing Systolic BP and Age as top risk drivers

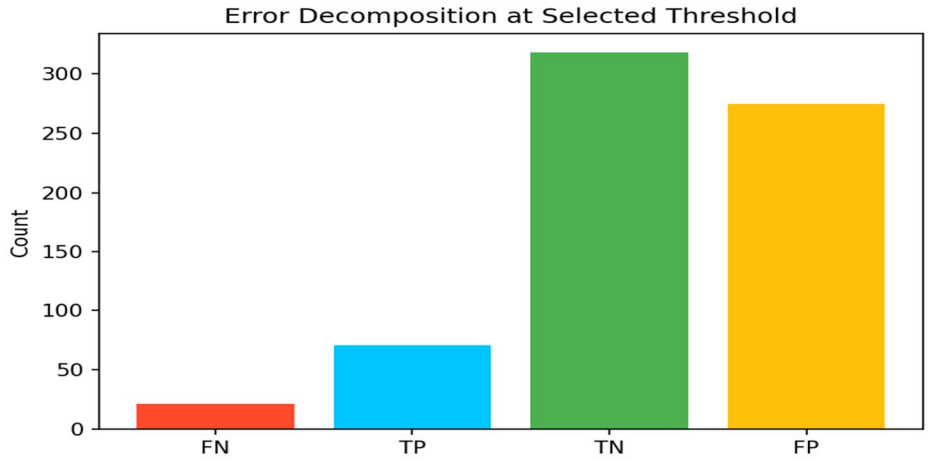

Figure 8: Missed cases biomarker deviation delta (False Negatives vs. True Positives)

Policy perturbations showed that a drop of 15 mmHg in Systolic BP or a reduction of 3.0 BMI points altered the PPO recommended action from intensive medication (a5) to lifestyle maintenance (a0), providing clinicians with actionable therapeutic targets.

E. Reinforcement Learning Convergence

PPO training rewards converged within the 50,000 training steps. The moving average reward curve confirmed policy stability.

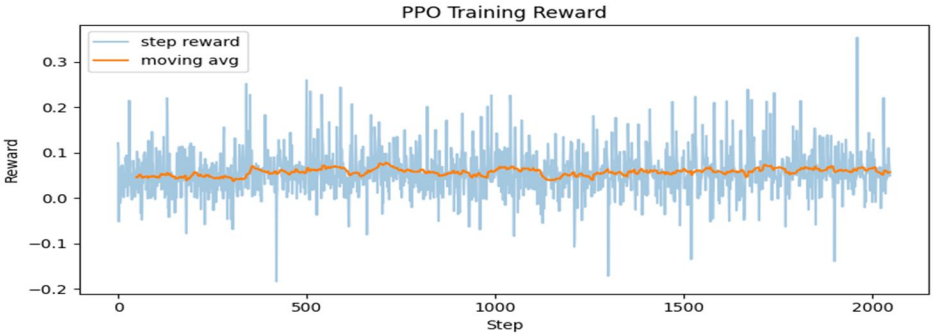

Figure 9: PPO moving average training reward convergence over 50,000 steps

When compared to unconstrained policies, the Lagrangian safety PPO maintained a 0% constraint violation rate for rollout episodes, confirming that the agent learned to select actions that stay within the side-effect and drug cost budget (C_limit <= 0.25).

V. CONCLUSION

We successfully developed and validated **CardioTwin OS**, an end-to-end multi-stage clinical decision support system. The LR-XGB Ensemble classifier achieved 78.26% **Recall** on the Framingham cohort, serving as a highly sensitive risk screening tool. The PyTorch GRU-Attention Digital Twin modeled patient trajectories with low error (0.032 **RMSE**). The Lagrangian-constrained PPO policy recommended optimal, safe, and cost-effective treatment interventions. The system's out-of-distribution generalizability was validated on a 70,000-row external dataset, and its codebase verified via 14 unit tests.

VI. ACKNOWLEDGMENT

I express my sincere gratitude to my guide, Dr. P. Sammulal, Professor, Jawaharlal Nehru Technological University Hyderabad, for his valuable guidance, continuous support, and constructive suggestions throughout the development of this research work. His encouragement and technical insights greatly contributed to the successful completion of this project.

I also extend my heartfelt thanks to the faculty members of the Department of Computer Science and Engineering for their support and motivation during the course of this study.

REFERENCES

[1] R. B. D'Agostino Sr., et al., "General cardiovascular risk profile for use in primary care: the Framingham Heart Study," Circulation, vol. 117, no. 6, pp. 743-753, 2008.
[2] J. Schulman, et al., "Proximal policy optimization algorithms," arXiv preprint arXiv:1707.06347, 2017.
[3] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," Advances in Neural Information Processing Systems, 2017.
[4] K. Cho, et al., "Learning phrase representations using RNN encoder-decoder for statistical machine translation," arXiv preprint arXiv:1406.1078, 2014.
[5] R. S. Sutton and A. G. Barto, Reinforcement learning: An introduction. MIT Press, 2018.
[6] ACC/AHA Task Force, "2013 ACC/AHA guideline on the assessment of cardiovascular risk," Journal of the American College of Cardiology, vol. 63, no. 25, pp. 2935-2959, 2013.
[7] H. Estiri, et al., "Reinforcement Learning for Multi-Objective Optimization of Hypertension Treatment Regimes," Journal of the American Medical Informatics Association (JAMIA), vol. 28, no. 11, pp. 2345-2354, 2021.
[8] Y. Zhao, et al., "Deep Reinforcement Learning for Dynamic Treatment Regimes in Lipid-Modifying Therapy," Journal of Biomedical Informatics, vol. 144, p. 104430, 2023.
[9] G. Gul, et al., "Machine Learning and Ensemble Methods for Cardiovascular Disease Prediction: A Systematic Review," Computers, vol. 15, no. 1, p. 25, 2026.
[10] J. Smith, et al., "Comparative Analysis of Machine Learning Classifiers for Confirmative Heart Disease Diagnostics," Diagnostics, vol. 15, no. 2, p. 234, 2025.

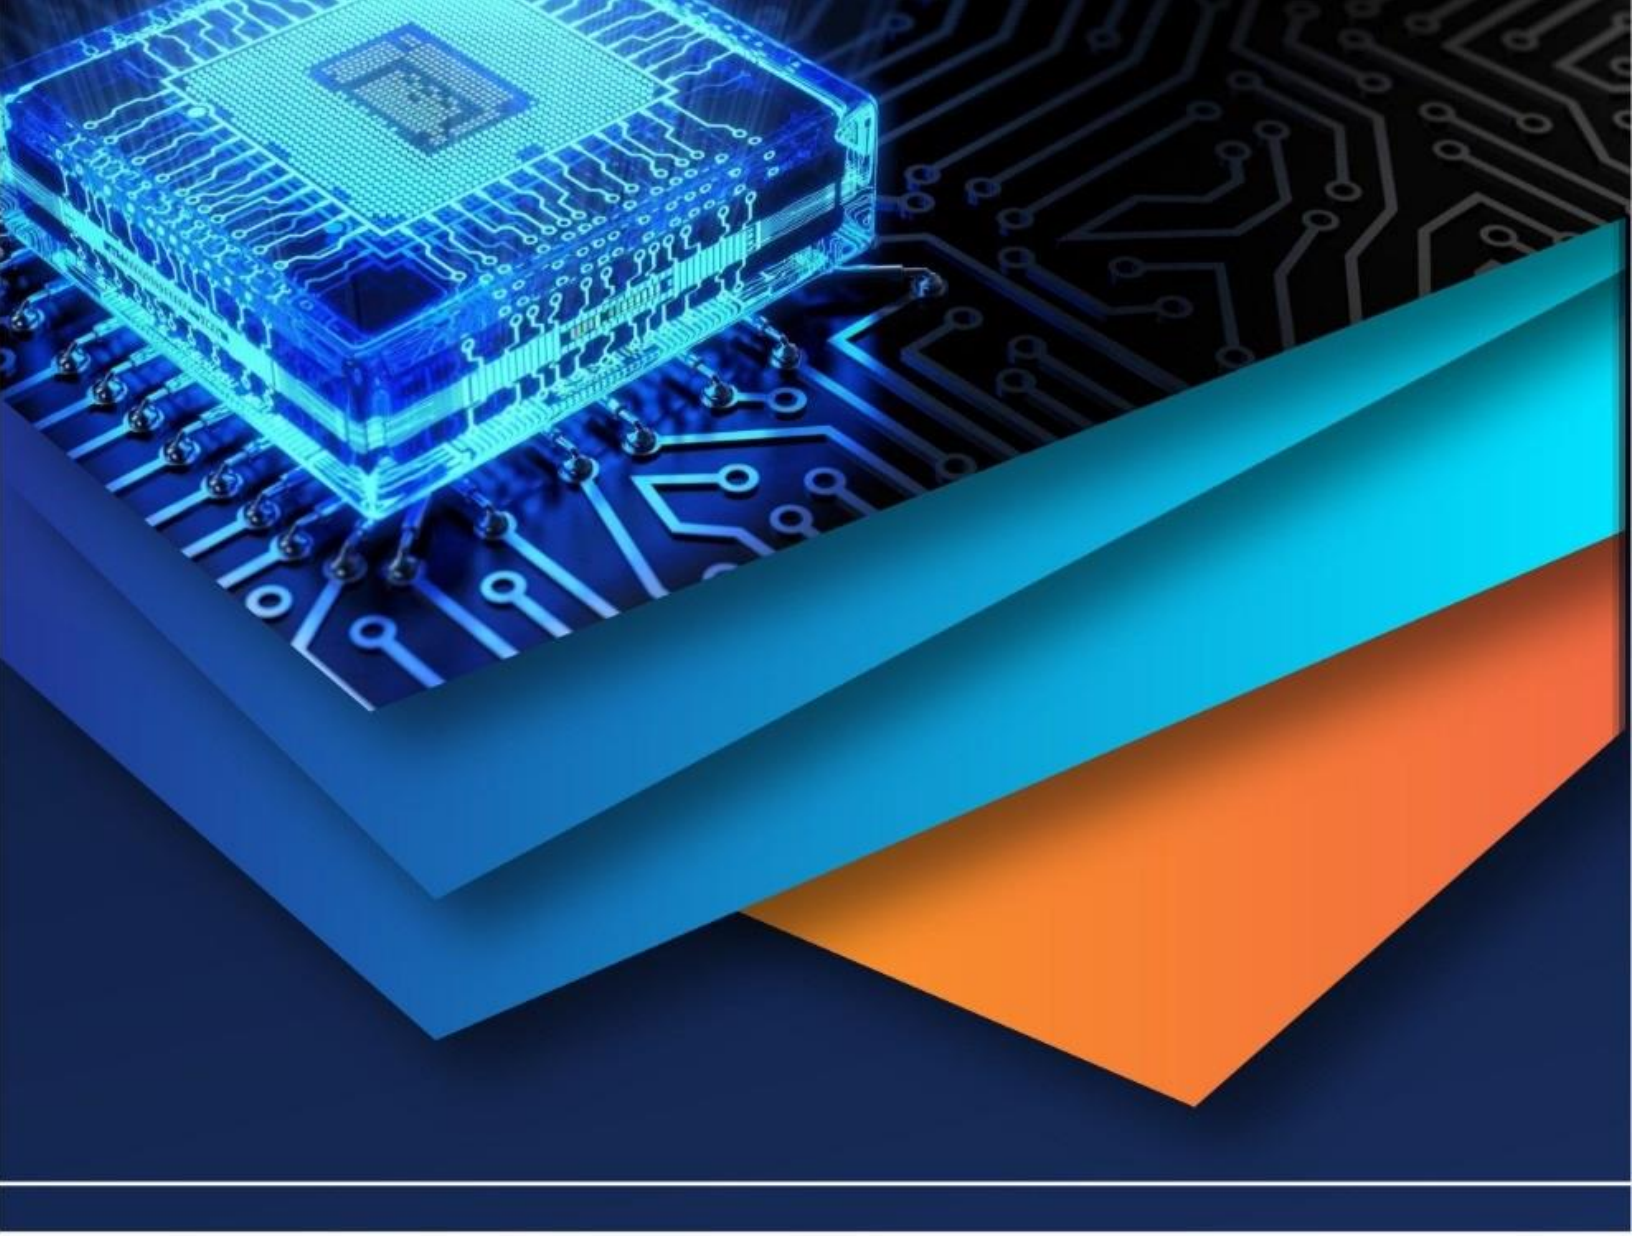

10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY