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.2026.84149

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

Cross-Modal Deep Learning Framework for Multi-Disease Risk Assessment Using Clinical Data and Medical Images

Gandla Sai Dheeraj Rao[1], N. Naveen Kumar[2]
[1]*M.Tech Scholar, Dept of CSE, JNTUH UCESTH, Hyderabad, India*
[2]*Associate Professor, Dept of CSE, JNTUH UCESTH, Hyderabad, India*

Abstract: ***Cardiovascular disease, diabetes mellitus, and chronic kidney disease constitute the cardiometabolic triad — three conditions that actively amplify one another through established clinical pathways including diabetic nephropathy, cardiorenal syndrome, and cardiovascular-metabolic disease progression. Existing clinical decision support systems assess these conditions in isolation, failing to model the compounded risk faced by comorbid patients. This paper presents MedFuse Lite, a comorbidity-aware multi-disease risk prediction system integrating five heterogeneous model branches - three XGBoost tabular classifiers, an EfficientNet-B0 ECG classifier, and a ResNet50 diabetic retinopathy detector - with a Comorbidity Fusion Network, a PyTorch multi-layer perceptron trained on 5,856 real NHANES 2017–2018 patients. The fusion network accepts a 15-dimensional input vector comprising individual model probabilities, shared clinical biomarkers, and multiplicative comorbidity interaction terms, producing comorbidity-adjusted risk estimates for all three diseases simultaneously. The heart disease model achieved a test ROC-AUC of 79.93% on 68,633 cardiovascular records; the retinal model achieved validation AUC of 90.27%; and the fusion network converged to a best validation loss of 11.770%. The system provides SHAP feature attribution, GradCAM spatial explainability, disease correlation analysis, and rule-based clinical report generation, deployed as a Flask RESTful web application with a four-section sequential workflow.***

Keywords: ***Cardiometabolic syndrome, Multi-disease prediction, XGBoost, EfficientNet-B0, ResNet50, Comorbidity fusion, NHANES, SHAP, GradCAM, Clinical decision support***

I. INTRODUCTION

Chronic non-communicable diseases represent the leading cause of global mortality, with the World Health Organisation attributing approximately 74% of all deaths to this category [14]. Among these, cardiovascular disease, diabetes mellitus, and chronic kidney disease (CKD) occupy a uniquely critical position: they do not merely co-exist but form a clinically validated comorbidity cluster in which the presence of one disease systematically elevates the risk and severity of the others. Diabetic nephropathy — the progression of kidney damage driven by sustained hyperglycaemia — is the leading cause of end-stage renal disease worldwide. Cardiorenal syndrome, whereby impaired kidney filtration elevates cardiac preload through fluid retention, compounds cardiovascular risk independently of traditional risk factors [15]. High blood glucose simultaneously accelerates arterial plaque formation, creating a third interaction pathway between metabolic and cardiovascular disease. Despite these well-documented interactions, clinical risk assessment tools treat each condition in isolation. The Framingham Risk Score [16] and QRISK3 [17] assess cardiovascular risk without integrating metabolic or renal markers. The CKD-EPI equation [18] estimates glomerular filtration rate without considering concurrent diabetic or cardiac context. No publicly available tool models the interactions between all three conditions or produces a comorbidity-adjusted unified risk profile. This paper presents MedFuse Lite, a multi-disease risk prediction system specifically designed to address this gap. The system integrates tabular clinical parameters and medical imaging data across five specialised model branches built on gradient-boosted tree classifiers [1] and convolutional neural networks, combines their outputs through a learned Comorbidity Fusion Network trained on real NHANES population data [10], and delivers comorbidity-adjusted risk estimates accompanied by explainable attribution using SHAP [6] and GradCAM [7], disease interaction alerts, and a plain English clinical report. The primary contributions of this work are: (1) a heterogeneous multi-branch architecture supporting simultaneous assessment of three cardiometabolic diseases across tabular and imaging modalities; (2) a Comorbidity Fusion Network trained on real NHANES patient data that learns population-level comorbidity interaction patterns rather than approximating clinical rules; and (3) a complete explainability stack combining SHAP feature attribution for tabular models and GradCAM spatial attention for retinal imaging.

©IJRASET: All Rights are Reserved

Machine learning approaches have repeatedly demonstrated advantages over traditional clinical risk tools when applied to routine patient data. Weng et al. showed that machine learning models improved cardiovascular risk prediction accuracy over conventional scoring methods when trained on large-scale clinical records [2], while Zou et al. found that ensemble-based classifiers consistently outperformed simpler statistical baselines in predicting diabetes mellitus from patient health indicators [3]. Parallel advances in deep learning architectures have similarly strengthened medical image analysis: residual learning frameworks such as ResNet [8] and compound-scaled architectures such as EfficientNet [9] have become standard backbones for high-accuracy image classification, motivating their adoption in this work for retinal and ECG image interpretation respectively. Beyond raw predictive accuracy, clinical adoption of such systems depends on alignment with established medical practice; this work grounds its report generation module in recognised clinical standards, including the American Diabetes Association's Standards of Care [11], the European Society of Cardiology's guidelines on cardiovascular disease prevention [12], and the KDIGO guideline for diabetes management in chronic kidney disease [13], ensuring that generated recommendations remain consistent with current evidence-based practice.

II. LITERATURE REVIEW

A. *Streamlit-Based Multi-Disease Prediction System for Early Detection of Kidney Disease, Breast Cancer, and Heart Disease*

Prior work has explored consolidating multiple disease predictors into a single accessible interface rather than building separate applications for each condition. One such system, built on Streamlit, combined kidney disease, breast cancer, and heart disease prediction within one platform by benchmarking several classical algorithms — Random Forest, Decision Tree, K-Nearest Neighbour, Gradient Boosting, and XGBoost — against standard preprocessing steps including missing-value imputation, label encoding, SMOTE-based class balancing, and hyperparameter tuning. Random Forest achieved the strongest results for kidney disease (99.16%) and heart disease (95.60%), while XGBoost was most effective for breast cancer (98.24%). This work demonstrates the practical value of unifying multiple disease predictors under one system to widen access to screening, an objective shared by MedFuse Lite. However, it is limited to tabular inputs and offers no support for medical imaging, model explainability, or the interaction effects between comorbid conditions — gaps that motivate the fusion-based approach adopted in this work.

B. *Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy*

Diabetic retinopathy is among the most common complications of diabetes and a leading cause of preventable vision loss. Gulshan et al. developed a deep learning system trained on a large corpus of labelled retinal fundus images to automatically detect referable diabetic retinopathy, achieving diagnostic performance comparable to that of practicing ophthalmologists [4]. This study was among the first to establish that image-based deep learning could match specialist-level accuracy in ophthalmology, and it directly informs the design of the retinopathy prediction branch in this work, which similarly relies on convolutional feature extraction from retinal images to flag diabetes-related complications alongside tabular risk prediction.

C. *Intelligent Multi-Disease Prediction using Advanced Deep Learning*

G. Deepa et al. proposed a deep learning framework capable of diagnosing several unrelated conditions — brain tumors, lung cancer, chronic kidney disease, and Parkinson's disease — within a single predictive system rather than restricting the model to one disease at a time.The study compared multiple architectures, including Fully Convolutional Networks (FCN), Long Short-Term Memory networks (LSTM), and Mask R-CNN, evaluating each on its ability to detect, segment, and classify regions of interest from medical images. Mask R-CNN produced the strongest results at 99.89% accuracy, outperforming FCN (95.56%) and LSTM (92.47%). While this confirms the feasibility of consolidating diverse disease predictions into one framework, the system treats each disease as an independent classification task and does not model any interaction or comorbidity risk between conditions, which is the central contribution this work aims to address.

D. *Cardiologist-Level Arrhythmia Detection Using Deep Neural Networks*

Hannun et al. developed a deep neural network capable of classifying cardiac arrhythmias directly from raw single-lead ECG signals, trained on a large-scale dataset of 91,232 ECG recordings collected from 53,549 unique patients across 12 rhythm classes, including normal sinus rhythm, various arrhythmias, and signal noise [5]. The model achieved an average ROC-AUC of 0.97 and outperformed a panel of practicing cardiologists on F1-score, demonstrating that deep neural networks can deliver reliable, scalable ECG interpretation suitable for mass screening. This result provides strong empirical grounding for the ECG-based heart disease classification branch used in this work, both in terms of architectural approach and the benchmark it sets for expected diagnostic reliability.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

III. PROPOSED SYSTEM

A. System Overview

The proposed **MedFuse Lite** framework is a heterogeneous and comorbidity-aware clinical decision support system developed for the integrated prediction of cardiovascular disease, diabetes mellitus, and chronic kidney disease. Unlike conventional systems that assess each disease independently, the proposed framework combines structured clinical parameters with optional ECG and retinal image analysis to generate comprehensive patient-specific risk assessments. Structured clinical data are analysed using XGBoost classifiers, while ECG and retinal images are processed using EfficientNet-B0 and ResNet50 deep learning models, respectively. The outputs of these specialised prediction modules are integrated through a Comorbidity Fusion Network trained on NHANES population data to capture clinically relevant interactions among coexisting diseases and shared biomarkers. The framework further incorporates SHAP-based feature attribution, GradCAM visual explanations, disease correlation analysis, and a rule-based clinical report generation module to improve the transparency and interpretability of predictions. By combining multimodal data analysis with explainable artificial intelligence techniques, MedFuse Lite provides an integrated and interpretable approach for supporting clinical decision making. The overall architecture of the proposed framework is illustrated in **Figure 1.**

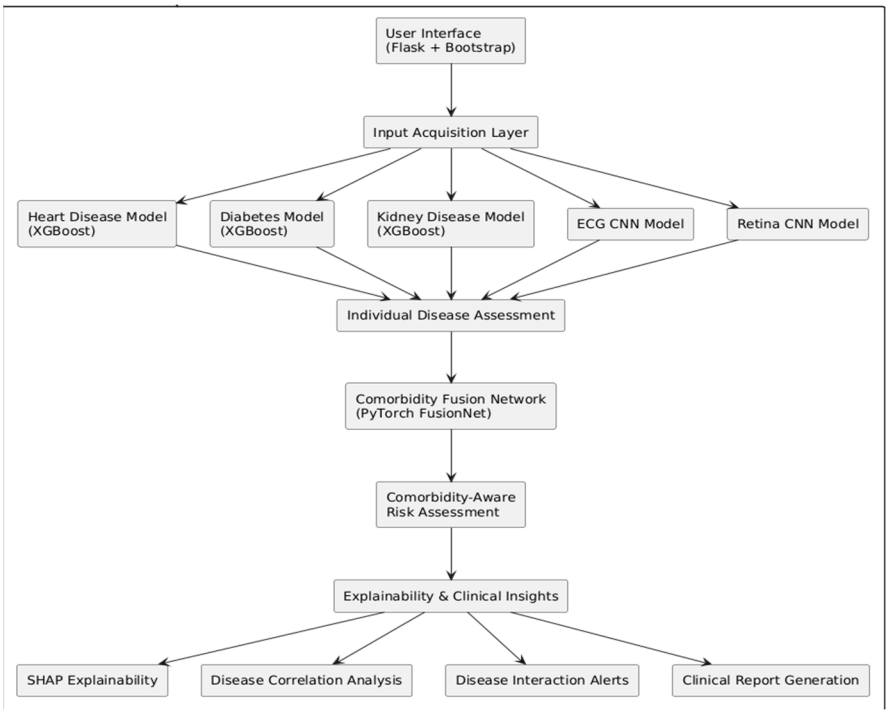


Fig. 1. Overall system architecture

B. Architecture Components

The proposed system is organised as a sequential pipeline of six functionally distinct stages. It begins with a web-based user interface built using Flask and Bootstrap, through which users enter clinical parameters such as age and BMI and optionally upload ECG or retinal fundus images. These inputs are then handled by an input acquisition layer, which collects the user-provided clinical parameters and optional medical images, validates them, and prepares them for the disease-specific prediction models. The validated inputs are subsequently passed to the disease prediction models, comprising three XGBoost classifiers for heart disease, diabetes, and kidney disease, alongside an EfficientNet-B0 model for ECG classification and a ResNet50 model for retinal image analysis, each of which independently estimates the risk associated with its respective condition. The outputs of these five model branches are aggregated at the individual disease assessment stage into per-disease probability scores. These probability scores, together with shared clinical biomarkers and interaction terms, are then combined by the Comorbidity Fusion Network, a trained PyTorch multilayer perceptron that models how the three conditions influence one another. The fusion network's output feeds into the comorbidity-aware risk assessment stage, which produces the final, refined risk scores for all three diseases after accounting for their clinical interactions. Finally, the explainability and clinical insights stage generates SHAP-based feature attribution, GradCAM visual explanations, disease correlation analysis, and a structured clinical report, ensuring that the predictions produced by the system remain transparent and interpretable for clinical use.

C. Individual Disease Model Branches

Heart Disease — XGBoost + EfficientNet-B0: An XGBoost classifier was trained on 68,633 cardiovascular patient records after physiological range filtering. Sixty Optuna hyperparameter optimisation trials maximised validation ROC-AUC, yielding n_estimators=660, max_depth=4, learning_rate=0.0116. Post-hoc Platt sigmoid scaling was applied after the pre-calibration maximum probability deviation of 0.1668 exceeded the clinical acceptability threshold. Two engineered features — BMI and pulse pressure — were appended to the eleven raw features. An EfficientNet-B0 CNN pretrained on ImageNet was fine-tuned for four-class ECG classification: Normal, Abnormal Heartbeat, Myocardial Infarction, and Post-MI History. Its aggregated abnormality probability (sum of non-Normal softmax outputs) supplements the tabular cardiac risk in the fusion layer.

Diabetes — XGBoost + ResNet50: An XGBoost classifier trained on 100,000 patient records with a custom precision-recall optimised decision threshold provides metabolic risk assessment. A ResNet50 pretrained on ImageNet, with all layers frozen except the final residual block and custom classification head, performs binary diabetic retinopathy detection. Class-weighted CrossEntropyLoss addresses training set imbalance. GradCAM activation maps over layer4[-1] are computed simultaneously with inference, providing spatial retinal explainability.

Chronic Kidney Disease — XGBoost + MICE: An XGBoost classifier trained on NHANES 2021–2023 staged kidney disease data processes ten renal biomarker and demographic features. Multiple Imputation by Chained Equations (IterativeImputer, n_nearest_features=5, max_iter=10) is applied strictly after train-test split to prevent imputation leakage, followed by StandardScaler transformation. This three-stage pipeline — imputer, scaler, classifier — is serialised as three separate artifacts applied in mandatory sequence at inference.

TABLE I DATASETS AND MODELS USED

Disease	Dataset	Records	Model
Heart Disease	Cardiovascular Dataset	68,633	XGBoost
Diabetes	Diabetes Prediction Dataset	100,000	XGBoost
Kidney Disease	NHANES	5,856	XGBoost
ECG	ECG Images	928 images	EfficientNet-B0
Retina	APTOS 2019	3,662 images	ResNet50

D. Framework Deployment

The proposed MedFuse Lite framework was implemented as a Flask-based web application with a Bootstrap 5 frontend to provide an interactive and user-friendly interface for disease prediction. All trained machine learning and deep learning models are loaded during application startup to minimise inference latency. The backend exposes dedicated prediction APIs for each disease module and the Comorbidity Fusion Network, enabling efficient communication between the frontend and prediction engine. This modular deployment strategy allows independent execution of disease-specific models while supporting integrated multimodal analysis within a unified clinical decision support framework.

E. Comorbidity Fusion Network

The Comorbidity Fusion Network (ComorbidityFusionNet) is a five-layer PyTorch MLP: Linear(15, 128) → ReLU → Dropout(0.25) → Linear(128, 64) → ReLU → Dropout(0.2) → Linear(64, 32) → ReLU → Linear(32, 3) → Sigmoid.

Its 15-dimensional input vector comprises: five individual model probabilities (heart tabular, ECG, diabetes tabular, retinal, CKD); five normalised shared biomarkers (age/90, systolic BP/200, BMI/50, serum creatinine/15, ACR/300); a diabetes-diagnosed binary flag; and three multiplicative interaction terms encoding known comorbidity pathways — dm_prob × ckd_prob (diabetic nephropathy), creatinine_norm × dm_prob (creatinine elevation in diabetic context), and heart_prob × ckd_prob (cardiorenal syndrome). The network was trained on 5,856 NHANES 2017–2018 adult patients using Adam optimiser (lr=0.001, weight_decay=1×10^{-4}), CosineAnnealingLR scheduling, and binary cross-entropy loss for 150 epochs. Disease labels were derived from confirmed diagnoses in questionnaire data and CKD-EPI eGFR computed from serum creatinine.

A post-network conditional amplification layer applies clinically motivated risk elevation when co-occurring disease probabilities exceed defined thresholds: 12% amplification to kidney risk when diabetic nephropathy conditions are met; 10% to cardiac risk under cardiorenal syndrome; and 8% under the cardiovascular-metabolic pathway. A global floor (≥85% of raw score) and ceiling (≤95%) prevent counterintuitive reductions and maintain clinical credibility.

F. Explainability and Output Layer

SHAP TreeExplainer computes the top three feature attributions per tabular model prediction, returned as directional driver chips in the UI. GradCAM spatial attention maps are computed over ResNet50's final residual block for every retinal inference. Disease Correlation Analysis generates natural language descriptions of active comorbidity pathways from the Disease Interaction Flags and stored biomarker values. A rule-based Clinical Report Generator produces per-patient plain English narratives using lifestyle recommendations sourced from ADA 2023[11], ESC 2021[12], and KDIGO 2022[13] clinical guidelines.

IV. RESULTS

A. Performance Evaluation

TABLE II PERFORMANCE COMPARISON OF INDIVIDUAL DISEASE MODELS

Model	Accuracy	Precision	Recall	F1-Score	ROC-AUC
Heart XGBoost	73.32	74.86	70.19	72.45	80.22
Diabetes XGBoost	87.13	89.20	66.27	79.72	88.23
Kidney XGBoost	99.39	99.68	99.44	99.56	99.93
ECG	82.46	82.52	82.16	82.32	88.11
Retina	75.45	71.23	72.86	74.23	79.56

Table II summarizes the performance of the proposed machine learning and deep learning models. The kidney disease model achieved the highest overall predictive performance, while the ECG and retinal image models demonstrated effective classification performance for medical image analysis, supporting reliable multimodal disease prediction.

B. User Interface Results

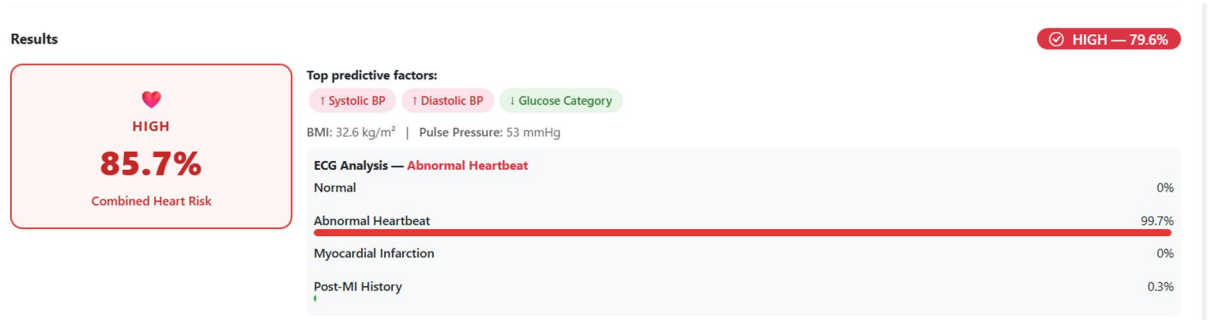


Fig. 2. Heart disease risk assessment results

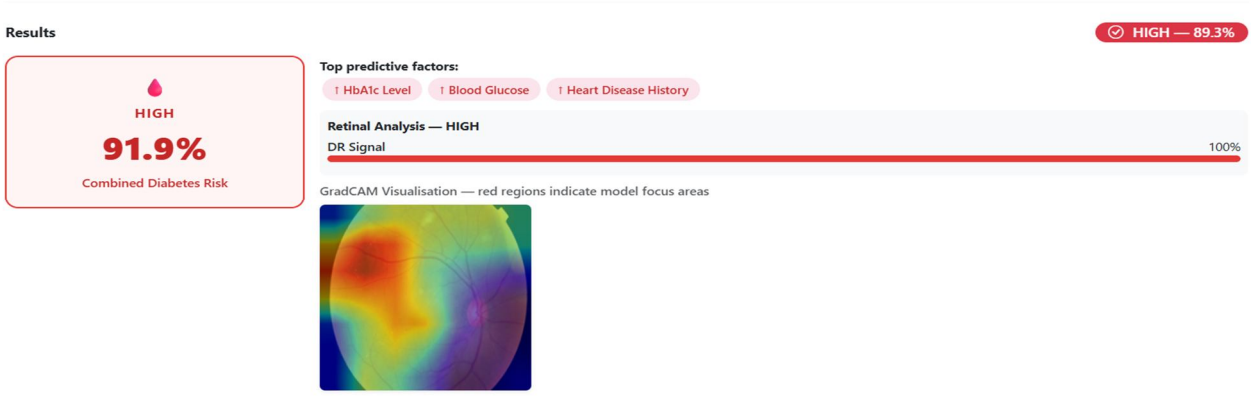


Fig. 3. Diabetes risk assessment results with Grad-CAM visualization

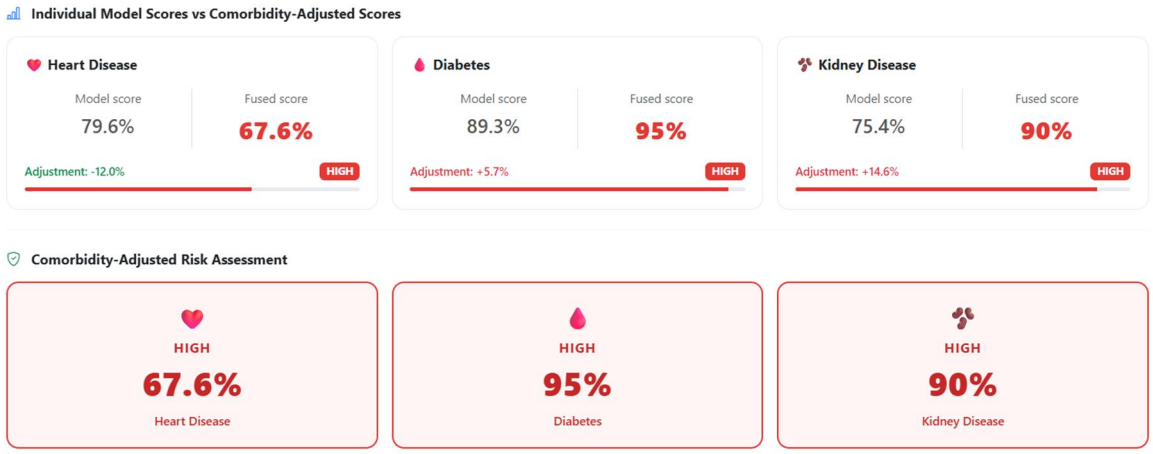


Fig. 4. Comorbidity-Aware risk assessment

C. Explainability Findings

SHAP analysis identified systolic and diastolic blood pressure as the dominant cardiac risk drivers, consistent with established Framingham and ESC cardiovascular risk frameworks. Blood pressure appeared as a shared top-three SHAP driver in both heart and kidney disease modules, empirically confirming that the system correctly identifies shared upstream risk pathways. GradCAM attribution for the retinal model concentrated spatial attention on microaneurysm-dense perimacula regions, providing visual confirmation that the high AUC reflects genuine pathological feature detection.

V. CONCLUSION

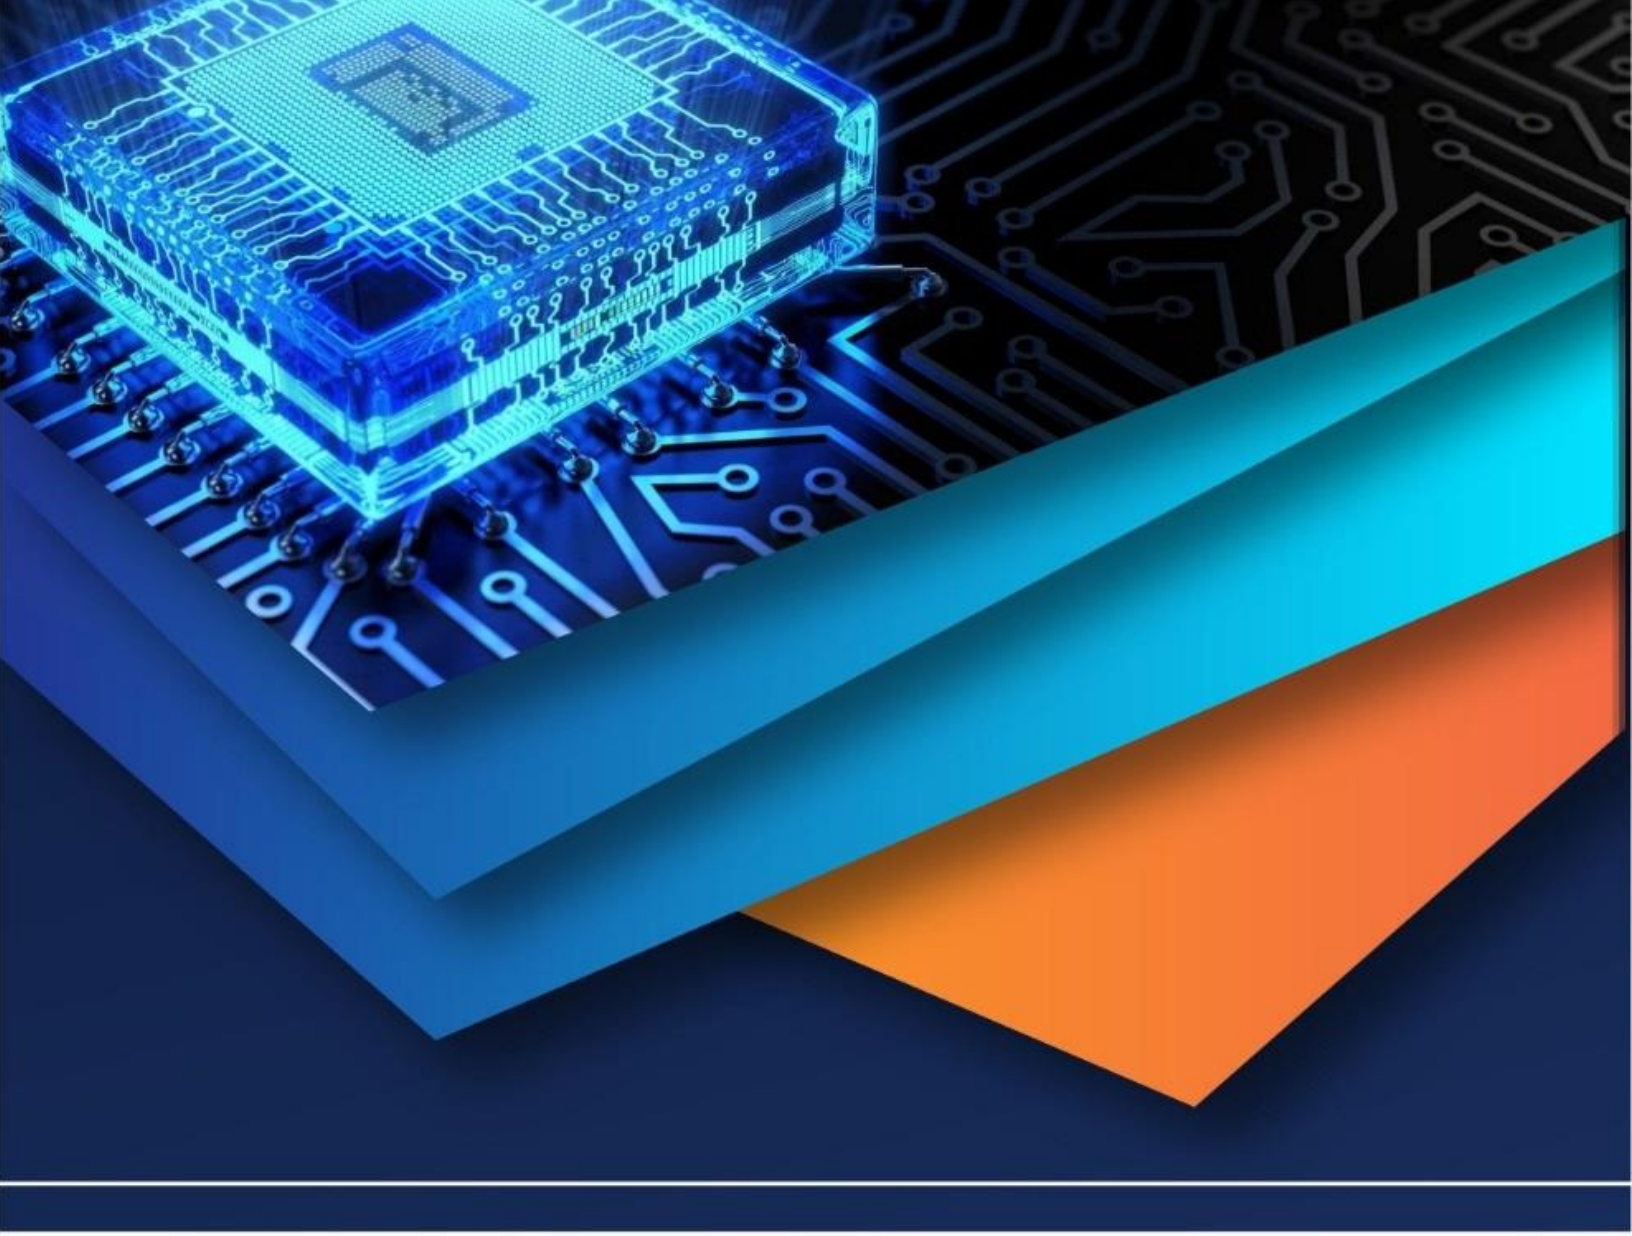
This paper presented MedFuse Lite, a comorbidity-aware multi-disease risk prediction system for the cardiometabolic triad. The system integrates five heterogeneous model branches with a Comorbidity Fusion Network trained on real NHANES population data, addressing the documented gap in clinical tools that assess cardiovascular, metabolic, and renal risk independently. Experimental results demonstrate clinically meaningful performance across all model branches, while the proposed Comorbidity Fusion Network effectively captures clinically relevant interactions among heart disease, diabetes, and chronic kidney disease. The complete explainability stack — SHAP attribution, GradCAM visualisation, disease correlation analysis, and guideline-sourced clinical reporting — makes the system interpretable and actionable in a clinical context.

Future work will prioritise prospective clinical validation on real patient records, multi-centre dataset evaluation to assess demographic generalisation, and formal ablation studies comparing the NHANES-trained fusion network against rule-based and no-fusion baselines to precisely quantify the marginal value of the learned comorbidity adjustment.

REFERENCES

[1] T. Chen and C. Guestrin, “XGBoost: A Scalable Tree Boosting System,” in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining (KDD), San Francisco, CA, USA, Aug. 2016, pp. 785–794.
[2] S. F. Weng, J. Reps, J. Kai, J. M. Garibaldi, and N. Qureshi, “Can Machine Learning Improve Cardiovascular Risk Prediction Using Routine Clinical Data?” PLoS One, vol. 12, no. 4, Art. no. e0174944, Apr. 2017.
[3] Q. Zou, K. Qu, Y. Luo, D. Yin, Y. Ju, and H. Tang, “Predicting Diabetes Mellitus With Machine Learning Techniques,” Front. Genet., vol. 9, Art. no. 515, Oct. 2018.
[4] V. Gulshan et al., “Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs,” JAMA, vol. 316, no. 22, pp. 2402–2410, Dec. 2016.
[5] A. Y. Hannun et al., “Cardiologist-Level Arrhythmia Detection and Classification in Ambulatory Electrocardiograms Using a Deep Neural Network,” Nat. Med., vol. 25, no. 1, pp. 65–69, Jan. 2019.
[6] S. M. Lundberg and S. I. Lee, “A Unified Approach to Interpreting Model Predictions,” in Adv. Neural Inf. Process. Syst. (NeurIPS), vol. 30, 2017, pp. 4765–4774.
[7] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, “Grad-CAM: Visual Explanations From Deep Networks via Gradient-Based Localization,” in Proc. IEEE Int. Conf. Comput. Vis. (ICCV), Venice, Italy, Oct. 2017, pp. 618–626.
[8] K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition,” in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), Las Vegas, NV, USA, Jun. 2016, pp. 770–778.
[9] M. Tan and Q. V. Le, “EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks,” in Proc. 36th Int. Conf. Mach. Learn. (ICML), Long Beach, CA, USA, Jun. 2019, pp. 6105–6114.

[10] Centers for Disease Control and Prevention (CDC), National Center for Health Statistics, “National Health and Nutrition Examination Survey (NHANES), 2017–2018 and 2021–2023.”
[11] American Diabetes Association, “Standards of Care in Diabetes—2023,” Diabetes Care, vol. 46, Suppl. 1, pp. S1–S291, Jan. 2023.
[12] European Society of Cardiology, “2021 ESC Guidelines on Cardiovascular Disease Prevention in Clinical Practice,” Eur. Heart J., vol. 42, no. 34, pp. 3227–3337, 2021.
[13] Kidney Disease: Improving Global Outcomes (KDIGO), “KDIGO 2022 Clinical Practice Guideline for Diabetes Management in Chronic Kidney Disease,” Kidney Int., vol. 102, no. 5S, pp. S1–S127, Nov. 2022.
[14] World Health Organization, "Noncommunicable Diseases," Fact Sheet, 2023.
[15] C. Ronco, M. Haapio, A. A. House, N. Anavekar, and R. Bellomo, "Cardiorenal Syndrome," J. Am. Coll. Cardiol., vol. 52, no. 19, pp. 1527–1539, Nov. 2008.
[16] [16] R. B. D'Agostino Sr., R. S. Vasan, M. J. Pencina, P. A. Wolf, M. Cobain, J. M. Massaro, and W. B. Kannel, "General Cardiovascular Risk Profile for Use in Primary Care: The Framingham Heart Study," Circulation, vol. 117, no. 6, pp. 743–753, Feb. 2008.
[17] J. Hippisley-Cox, C. Coupland, and P. Brindle, "Development and Validation of QRISK3 Risk Prediction Algorithms to Estimate Future Risk of Cardiovascular Disease: Prospective Cohort Study," BMJ, vol. 357, p. j2099, May 2017.
[18] A. S. Levey, L. A. Stevens, C. H. Schmid, Y. L. Zhang, A. F. Castro III, H. I. Feldman, J. W. Kusek, P. Eggers, F. Van Lente, T. Greene, and J. Coresh, "A New Equation to Estimate Glomerular Filtration Rate," Ann. Intern. Med., vol. 150, no. 9, pp. 604–612, May 2009.



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)