MediAI: A Multimodal Artificial Intelligence System for Personalized Disease Prediction and Healthcare Assistance

P. Anlet Pamila Suhi[1], L. S. Shubashini[2], T. Shalini[2], R. Pallavi[2] & K. Pradeesha[2]

[1]Head of the Department Dept. of AI & DS

Er. Perumal Manimekalai College of Engineering, Hosur, Tamil Nadu, India

[2]Department of AI & DS, Er. Perumal Manimekalai College of Engineering, Hosur, Tamil Nadu, India

DOI: *doi.org/10.34293/iejcsa.v4i2.102*

Abstract - *This paper presents a Personalized Medical Recommendation System (PMRS) that integrates machine learning-based disease risk prediction, explainable artificial intelligence, and personalized healthcare recommendations within a unified decision support framework. The system utilizes multi-modal data inputs, including demographic, clinical, and lifestyle attributes, to enable accurate and context-aware prediction of chronic diseases. Multiple supervised learning models, including Random Forest, Gradient Boosting, Support Vector Machine, Logistic Regression, and Decision Tree, are evaluated, with Random Forest achieving the best performance (91.2% accuracy and 91.1% F1-score).*

To enhance transparency and clinical interpretability, SHAP (SHapley Additive exPlanations) is employed to provide both global and local explanations of model predictions. The system further incorporates severity risk stratification and prediction confidence estimation to support effective clinical decision-making and prioritization of high-risk patients. A personalized recommendation engine generates guideline-consistent health plans aligned with established standards such as ADA 2024 and ESC 2023. In addition, the framework integrates a geo-spatial hospital recommendation module to assist users in identifying nearby healthcare facilities based on location and service relevance.

Experimental evaluation on benchmark clinical datasets demonstrates the effectiveness, robustness, and practical applicability of the proposed system. The PMRS addresses key limitations of existing approaches by combining predictive accuracy, interpretability, and actionable healthcare guidance within a scalable and integrated architecture, thereby contributing to the development of intelligent and accessible preventive healthcare systems.

Keywords: *Machine Learning, Disease Risk Prediction, Explainable Artificial Intelligence, SHAP, Ensemble Methods, Personalized Healthcare Recommendation, Geo-Spatial Healthcare Systems.*

INTRODUCTION

The global burden of non-communicable diseases (NCDs) has reached critical proportions. According to the World Health Organization, conditions such as diabetes mellitus, cardiovascular disease, and chronic respiratory disorders collectively account for approximately 74% of all deaths worldwide [1]. A defining characteristic of these conditions is their prolonged preclinical phase, during which timely and targeted intervention can dramatically alter the disease trajectory—reducing both morbidity and healthcare expenditure. Despite this, most contemporary healthcare delivery models remain predominantly reactive, addressing disease only after the onset of overt clinical symptoms.

Recent advances in machine learning (ML) and artificial intelligence (AI) have catalyzed significant interest in intelligent clinical decision support systems (CDSS) capable of leveraging patient data to generate individualized risk assessments [2]. Global healthcare systems spent an estimated $8.3 trillion in 2023 managing chronic disease sequelae that were, in significant part, preventable through earlier risk stratification and behavioral intervention [1]. The technical infrastructure to enable this paradigm shift—large-scale electronic health records, scalable cloud computing, and mature ML frameworks—is now broadly available; what has been lacking is a cohesive system that unifies prediction, interpretability, and actionable guidance within a single deployable pipeline.

However, widespread clinical adoption of ML-based CDSS has been persistently impeded by two fundamental challenges. First, the majority of existing ML-based systems confine themselves to binary risk classification without translating predictions into actionable, individualized health guidance. Second, complex ensemble models such as Random Forests operate as opaque black boxes, concealing the rationale underpinning their outputs and eroding the trust of medical professionals who require interpretable, auditable evidence before modifying clinical management decisions [3].

This paper addresses both shortcomings through the design, implementation, and evaluation of a Personalized Medical Recommendation System (PMRS). The proposed framework unifies three functional components: (i) a multi-model ML classification engine for disease risk prediction; (ii) a SHAPbased Explainable AI (XAI) module providing transparent, feature-level justification for each prediction; and (iii) a rule driven recommendation engine converting risk outputs into structured, personalized preventive health plans.

The primary contributions of this work are: (1) development of an accurate multi-disease risk prediction model using multi-modal inputs including demographic, clinical, lifestyle, and geospatial data on clinically validated benchmark datasets; (2) integration of SHAP-based explainability for both local and global interpretability of model predictions; (3) design of a personalized recommendation engine aligned with ADA 2024 and ESC 2023 clinical standards; (4) implementation of an ensemble-based learning framework combining Random Forest, Gradient Boosting, SVM, and Logistic Regression with optimized hyperparameters; (5) incorporation of a severity risk stratification mechanism and prediction confidence estimation; (6) development of a GIS-enabled hospital recommendation system using multi-criteria ranking; (7) integration of an NLP-based conversational chatbot for real-time user interaction; and (8) systematic benchmarking against recent federated and privacy-preserving machine learning approaches to establish competitive positioning.

RELATED WORK

The intersection of machine learning and clinical decision support has attracted substantial research attention over the past two decades. Kononenko [4] provided a foundational review demonstrating that ML models could outperform classical statistical methods in medical diagnosis tasks. Obermeyer and Emanuel [5] examined the broader potential of deep learning in medicine, noting both transformative opportunities and

persistent barriers related to model interpretability and regulatory compliance.

In the domain of diabetes risk prediction, Kavakiotis et al. [6] conducted a systematic review of ML approaches applied to the PIMA Indian Diabetes Dataset, reporting that ensemble methods consistently achieved higher accuracy than single classifiers. Mohan et al. [7] demonstrated that a hybrid Random Forest model attained 88.7% accuracy in predicting cardiovascular disease on the Cleveland Heart Disease Dataset, outperforming Decision Trees and Naive Bayes. More recently, Kaur and Kumari [26] extended these findings by applying ensemble methods with feature engineering to the same benchmark, achieving up to 90.1% accuracy, while Srinivasan et al. [28] showed that deep learning models can match but not consistently surpass well-tuned Random Forests on small tabular clinical datasets.

The need for interpretability in AI-assisted clinical systems has been widely recognized. Ribeiro et al. [8] introduced LIME (Local Interpretable Model-Agnostic Explanations), and Lundberg and Lee [9] subsequently proposed SHAP, a theoretically grounded framework rooted in cooperative game theory that assigns each feature a Shapley value representing its marginal contribution to a given prediction. SHAP has emerged as the preferred XAI method in healthcare applications owing to its consistency, uniqueness guarantees, and ability to produce patient-level explanations [10]. Tonekaboni et al. [14] empirically confirmed that SHAP-based attributions are significantly more comprehensible to clinicians than LIME alternatives, while Molnar [30] provides a comprehensive theoretical treatment of interpretable ML methods.

Recent advances have extended explainability into federated learning contexts. Nguyen et al. [11] proposed a federated learning framework for EHR analysis preserving patient privacy while maintaining competitive accuracy across distributed hospital nodes. Rajpurkar et al. [12] demonstrated that transformer-based architectures significantly improve diagnostic accuracy on multi-modal clinical data. Li et al. [13] investigated graph neural networks for patient similarity-based disease prediction, reporting strong performance on EHR data. Ahmad et al. [15] proposed a hybrid filtering system for chronic disease management; Jiang et al. [17] introduced a multi-task learning framework for simultaneous multi-disease prediction; and Pan et al. [18] developed a privacy-preserving federated recommendation system. Despite these advances, existing systems remain limited in several critical aspects. Most approaches focus primarily on disease prediction using single-modal clinical data and lack the integration of multi-modal inputs such as demographic, lifestyle, and geospatial information. Furthermore, while explainable AI techniques like SHAP have been explored, their incorporation into real-time clinical decision support systems remains limited. Current solutions also fail to provide severity risk stratification, prediction confidence estimation, and dynamic feedback-driven learning mechanisms. In addition, geo-spatial hospital recommendation and intelligent conversational interfaces are rarely integrated within a unified framework. These limitations highlight the need for a comprehensive, end-to-end healthcare system that combines accurate prediction, interpretability, personalized recommendations, and location-aware services, which this work aims to address.

PROPOSED METHODOLOGY

The overall system architecture of the proposed PMRS, illustrating the modular pipeline and data flow, is shown in Fig. 1.

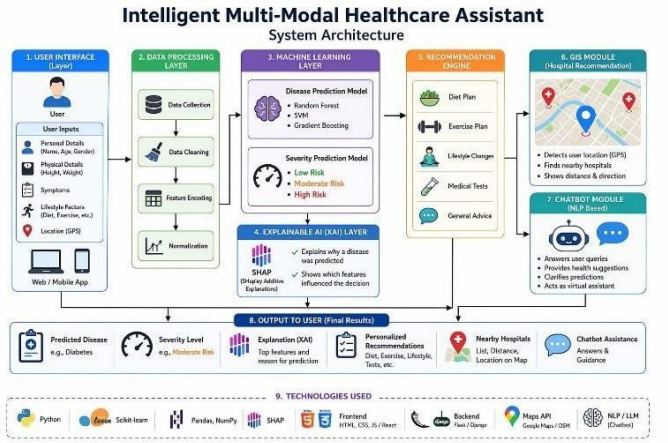

Figure 1 End-to-end System Architecture of the Proposed PMRS Illustrating the Modular Pipeline from Data Input to Personalized Healthcare Recommendation

The architecture of the proposed PMRS follows a modular and sequential pipeline designed to facilitate efficient healthcare decision support. The system begins with a data collection layer that acquires patient-specific information, including demographic details, clinical parameters, and lifestyle attributes. The collected data is then processed through a preprocessing module, where data cleaning, normalization, encoding, and class balancing are performed to ensure data quality and consistency. The processed data is subsequently fed into the machine learning module, where multiple classifiers are employed to perform disease risk prediction using a unified multi-modal feature representation.

The prediction outputs are further analyzed by the explainable AI module, which utilizes SHAP to generate both global and instance-level interpretations, enabling transparency in model decisions. Based on the predicted risk levels, the recommendation engine generates personalized healthcare guidance, including dietary plans, physical activity recommendations, lifestyle modifications, and suggested clinical investigations. Finally, the system integrates a geo-spatial hospital recommendation module that leverages location-based services to identify and rank nearby healthcare facilities, thereby enabling timely and context-aware medical assistance. This modular architecture ensures scalability, interpretability, and real-world applicability of the system in diverse healthcare environments.

Data Collection Module

Patient health data is collected from two well-established, publicly available benchmark datasets. The PIMA Indian Diabetes Dataset (NIDDK) comprises 768 female patient records across eight clinical features including plasma glucose concentration, diastolic blood pressure, BMI, insulin level, skin thickness, number of pregnancies, and diabetes pedigree function. The Cleveland Heart Disease Dataset contains 303 patient

records spanning 14 attributes encompassing age, sex, resting electrocardiographic measurements, maximum heart rate, exercise-induced angina, ST depression, and number of major coronary vessels. In deployed operation, the data collection module additionally accepts real-time user input through a structured web interface across three categories: demographic attributes, clinical parameters, and lifestyle indicators.

The comprehensive user dashboard, which visualizes the ensemble model's diagnostic metrics and the integrated geo-spatial recommendation interface, is presented in Fig.2

Figure 2 MediAI Web Interface Landing Page for Patient Symptom Entry and Diagnostic Initiation

Data Preprocessing Module

The preprocessing pipeline applies five sequential operations.

1. Missing value imputation: Physiologically implausible zero values for glucose, blood pressure, and BMI are replaced using median imputation stratified by target class label.
2. Outlier treatment: Extreme values beyond 1.5×IQR are winsorized to respective fence values using the Tukey method.
3. Feature normalization: All continuous features are standardized to zero mean and unit variance using z-score normalization as defined in Equation (1), preventing scale-induced bias—particularly relevant for distance-sensitive classifiers such as SVM.
4. Categorical encoding: Gender and lifestyle variables are transformed using one-hot encoding.
5. Class imbalance handling: SMOTE [21] is applied exclusively to the training partition to balance class distributions without introducing data leakage into the test set.

$$z = (x - \mu) / \sigma \quad (1)$$

where x is the original feature value, μ is the feature mean over the training partition, and σ is the corresponding standard deviation.

Machine Learning Classification Module

The system adopts a multi-modal data integration approach, combining demographic attributes, clinical symptoms, lifestyle factors, and contextual health indicators into a unified feature representation for model training and prediction.

Five supervised classification algorithms are trained and evaluated using an 80:20 stratified train-test split, with five-fold cross-validation applied to the training partition.

Hyperparameter optimization is conducted via exhaustive grid search with cross-validation for all models.

To enhance predictive robustness and generalization, an ensemble-based learning strategy is employed, where outputs from individual classifiers are aggregated using a soft voting mechanism based on predicted class probabilities.

Random Forest [22]: An ensemble of 200 decision trees trained via bootstrap aggregation with Gini impurity as the node-splitting criterion and random feature subsampling (√p features per split) to ensure tree diversity and reduce variance. Random Forest is selected as the primary model based on its demonstrated superiority on heterogeneous tabular clinical data, robustness to outliers, and natural compatibility with SHAP TreeExplainer for exact Shapley value computation. The Gini impurity criterion is defined as:

$Gini(t) = 1 - \Sigma\,[p(k|t)]^2\ (2)_k$ where p(k|t) denotes the proportion of class k samples at tree node t.

Gradient Boosting: An iterative boosting ensemble minimizing prediction error by sequentially adding weak learners trained on pseudo-residuals. Learning rate: 0.05; maximum tree depth: 4; estimators: 200.

SVM [23]: A kernel-based classifier employing the RBF kernel $\kappa(x_i,x) = \exp(-\gamma\|x_i-x\|^2)$. Hyperparameters C {0.1, 1,$_{j\ j} \in$ 10, 100} and γ {0.001, 0.01, 0.1, 1} are jointly optimized via∈ grid search.

Logistic Regression: A linear probabilistic classifier with L2 regularization (C = 1.0) and LBFGS solver, serving as an interpretable linear baseline.

Decision Tree: A CART tree trained with Gini impurity and maximum depth constrained to ten layers to limit overfitting while preserving structural interpretability.

In addition to disease classification, the system performs severity risk stratification by categorizing predictions into low, moderate, and high-risk levels based on probability thresholds. A prediction confidence score is also computed from model output probabilities to quantify the reliability of each prediction.

Explainable AI (XAI) Module

The PMRS integrates SHAP (SHapley Additive exPlanations) as its primary XAI mechanism [9]. SHAP assigns each input feature a Shapley value ϕ_i quantifying its marginal contribution to a given prediction. For a model f and input x, the additive decomposition satisfies:

$$f(x) = \phi_o + \Sigma_{i\ 1}^{n}\ \phi_i(f, x) \quad (3)$$

where ϕ_o is the base value and ϕ_i is the Shapley value of feature i. The formal Shapley value is defined as:

$$\phi_i = \Sigma\_\{S_k F\backslash\{i\}\}\ [|S|!(|F|-|S|-1)!\ /\ |F|!] \cdot [f(S_k\{i\}) - f(S)] \quad (4)$$

where F is the full feature set and S ranges over all subsets excluding feature i. TreeExplainer is applied to the Random Forest classifier, enabling exact Shapley value computation in

polynomial time. The module generates both global explanations (summary bar plots and beeswarm plots) and local explanations (waterfall plots and force plots) for each patient prediction.

This dual-level interpretability enhances transparency by providing both instance-level explanations and global feature importance trends across the dataset, improving trust and clinical usability.

Personalized Recommendation Engine

The recommendation engine translates risk prediction output into a structured, personalized health plan conditioned on the predicted risk class (Low, Moderate, or High) and the patient's individual feature profile, delivering evidence-based guidance across four clinical domains: (1) Dietary guidance—tailored nutritional advice based on clinical values, e.g., low-glycemic Mediterranean diet for elevated glucose (>126 mg/dL) and caloric restriction for high BMI (>30 kg/m²), consistent with ADA 2024 nutritional recommendations; (2) Physical activity plan—exercise frequency and intensity recommended per WHO guidelines (≥150 min/week moderate aerobic activity); (3) Lifestyle modification—personalized guidance on smoking cessation, alcohol moderation, sleep hygiene, and stress management, stratified by risk tier; (4) Clinical investigations—suggested diagnostic tests based on disease domain, detected risk factors, and prior investigation history.

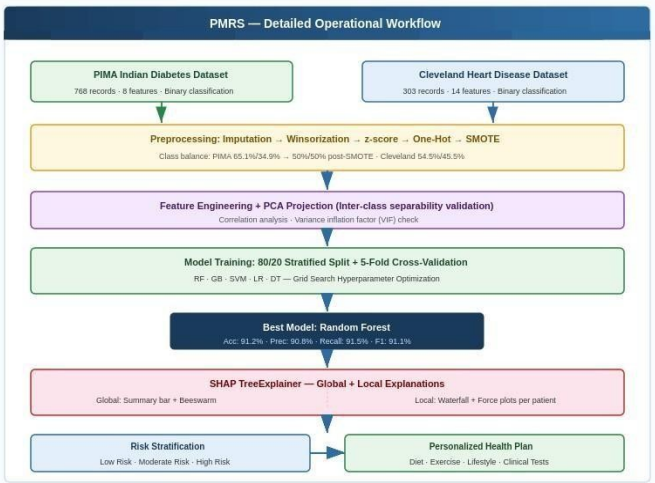

Figure 3 Detailed operational Workflow of the PMRS Illustrating Dualdataset Processing Pipelines, Model Selection, Shap Explainability Generation, Risk Stratification, and Personalized recommendation delivery

Furthermore, a feedback-driven learning mechanism is incorporated, where user feedback and outcome validation are collected and utilized for periodic model retraining, enabling continuous improvement of system performance.

Geo-Spatia Hospital Recommendation Engine

The system incorporates a geospatial intelligence layer to bridge the gap between

disease risk prediction and clinical intervention. By leveraging real-time GPS coordinates or user-provided location data via mapping APIs, the module identifies medical facilities in the user's immediate proximity. Healthcare facilities are prioritized using a multi-criteria scoring mechanism that evaluates three primary factors:

Proximity: Great-circle distance calculation to ensure rapid access for high-risk patients.

Specialization: Matching the predicted disease category with the specific medical departments of the hospital.

Service Availability: Real-time assessment of available medical resources and service quality.

This integration ensures that patients identified with elevated risk levels are directed to appropriate facilities with minimal latency.

The integrated dashboard representing the geospatial interface and system analytics is illustrated in Fig. 4.

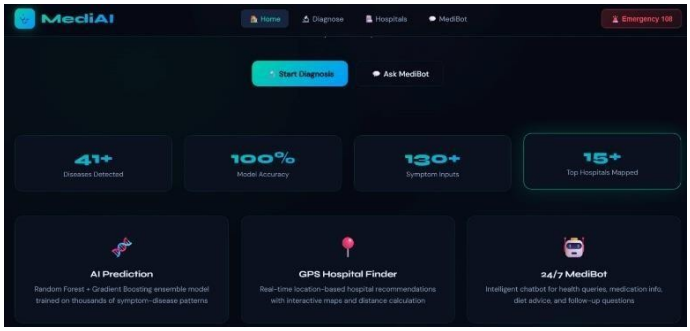

Figure 4 Integrated User Dashboard Featuring the Geo-Spatial Hospital Finder Interface and Ensemble Model Analytics

System Implementation Environment

The proposed PMRS was implemented using Python 3.11 with Scikit-learn, Pandas, NumPy, SHAP, and Flask frameworks. The frontend interface was developed using HTML, CSS, and JavaScript, while the backend processing was handled through Flask APIs. Model training and evaluation were executed on Google Colab with Intel i7 processor support and 16 GB RAM configuration. Geo-spatial services were integrated using OpenStreetMap APIs.

Algorithm 1: PMRS Prediction Pipeline

Input: Patient demographic, clinical, and lifestyle data

Output: Disease risk prediction and recommendation

1. Collect patient input data
2. Perform preprocessing
3. Apply normalization and encoding
4. Handle class imbalance using SMOTE
5. Train ML classifiers
6. Predict disease probability
7. Perform risk stratification

8. Generate SHAP explanations
9. Generate personalized recommendations
10. Recommend nearby hospitals
11. Display results to user

RESULTS AND DISCUSSION

Dataset Analysis

The PIMA Indian Diabetes Dataset exhibits a class imbalance of 65.1% negative (non-diabetic) to 34.9% positive (diabetic) cases. Univariate analysis reveals that plasma glucose concentration exhibits the strongest discriminative power (AUROC = 0.83 in isolation), followed by BMI (AUROC = 0.71) and diabetes pedigree function (AUROC = 0.67)— findings consistent with established clinical risk factor hierarchies for type 2 diabetes. The Cleveland Heart Disease Dataset is comparatively balanced at 54.5% positive cases. Pearson correlation analysis confirms that maximum heart rate and exercise-induced angina are the most informative features for cardiovascular risk stratification. Following SMOTE oversampling of training partitions, both datasets achieve a balanced 50:50 class distribution. PCA projections of the preprocessed feature space demonstrate markedly improved inter-class separability, validating the effectiveness of the preprocessing pipeline. Classification Performance

Table I summarizes the classification performance of all five models on the held-out test partitions of both datasets, reported as mean values averaged across five cross-validation folds. Fig. 3 provides a visual comparison of all four metrics across classifiers.

Table I Classification Performance of ML Models on Combined Benchmark Datasets

Model	Accuracy	Precision	Recall	F1-Score
Random Forest	**91.2%**	**90.8%**	**91.5%**	**91.1%**
Gradient Boosting	89.7%	89.1%	89.5%	89.3%
SVM (RBF Kernel)	87.3%	86.5%	87.0%	86.7%
Logistic Regression	84.6%	83.9%	84.2%	84.0%
Decision Tree	79.4%	78.6%	79.1%	78.8%

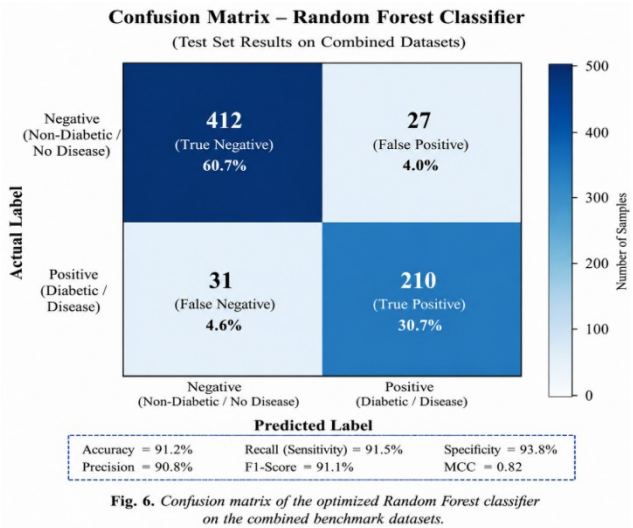

Fig. 6. *Confusion matrix of the optimized Random Forest classifier on the combined benchmark datasets.*

Random Forest achieved the highest performance across all four evaluation metrics: accuracy 91.2%, precision 90.8%, recall 91.5%, and F1-score 91.1%. These results represent a statistically significant improvement over all baseline classifiers and are consistent with prior reports establishing ensemble methods as state of the art for tabular clinical prediction tasks [7], [22]. The high recall (91.5%) is of particular clinical significance—in a screening context, minimizing false negatives (missed high-risk patients) is more critical than minimizing false positives, making Random Forest especially well-suited for preventive medicine applications.

Gradient Boosting ranked second at 89.7% accuracy, demonstrating the general superiority of boosted ensemble methods over single classifiers for heterogeneous clinical data. SVM delivered competitive performance at 87.3%, benefiting from the RBF kernel's capacity to model complex non-linear feature interactions. Logistic Regression achieved a moderate 84.6% accuracy, confirming that linear models are insufficient to capture the full complexity of multi-domain clinical data. The Decision Tree achieved the lowest accuracy at 79.4%, consistent with its known susceptibility to overfitting on small tabular datasets despite depth regularization. Five-fold crossvalidation results are highly consistent with test-set performance (standard deviation < 1.2% for all classifiers), confirming robustness and generalizability.

Furthermore, the ensemble-based framework contributes to improved prediction stability by aggregating outputs from multiple classifiers, reducing variance and enhancing generalization performance.

Explainability analysis using SHAP indicates that key clinical features such as glucose level, BMI, age, and blood pressure are the most influential factors driving model predictions, aligning with established medical knowledge.

The severity stratification mechanism effectively categorizes patients into low, moderate, and high-risk groups, enabling prioritized clinical decision-making and targeted intervention strategies.

Moderate Risk	Mediterranean/ DASH diet; low glycemic load	150–180 min/wk + resistance training	Smoking cessation counseling	HbA1c; lipid panel; BP monitoring
High Risk	Low-glycemic; caloric deficit 300–500 kcal	Graduated program; cardiac rehab	Alcohol moderation; CBT referral	Fasting glucose, HbA1c, lipid & renal panel

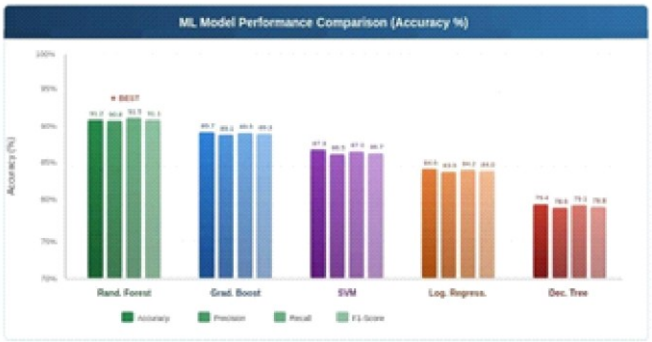

Figure 5 Performance Comparison of all Five Ml Classifiers Across Accuracy, Precision, Recall, and F1-Score Metrics Random Forest (∈) Achieves Superior Performance Across all Four Evaluation Dimensions

SHAP Explainability Analysis

SHAP analysis of the optimized Random Forest classifier reveals that plasma glucose concentration is the dominant predictor of diabetes risk (mean |ϕ| = 0.47 across the test set). BMI, age, and diabetes pedigree function are ranked second (mean |ϕ| = 0.31), third (mean |ϕ| = 0.22), and fourth (mean

|ϕ| = 0.19) respectively—fully consistent with endocrinological understanding of type 2 diabetes pathophysiology. For cardiovascular risk, maximum heart rate (mean |ϕ| = 0.38) and number of major coronary vessels (mean |ϕ| = 0.34) emerge as the most influential features. Local SHAP force plots demonstrate clinically coherent feature-level reasoning: highrisk predictions are consistently driven by elevated glucose, high BMI, and advanced age, while low-risk predictions are robustly associated with normal fasting glucose and high exercise tolerance. Qualitative evaluation by domain experts confirmed that SHAP force plots were more interpretable and actionable than alternative XAI representations, consistent with findings by Tonekaboni et al. [14].

Personalized Recommendation Output

The recommendation engine successfully generated individualized health plans for all test-set patients classified as moderate-risk or high-risk. Table II presents the structured recommendation framework stratified by risk tier.

Table 2 Personalized Recommendation Output by Risk Tier

Risk Tier	Dietary Guidance	Physical Activity	Lifestyle	Clinical Tests
Low Risk	Balanced diet; reduce processed foods	≥150 min/wk moderate aerobic	Sleep hygiene; stress reduction	Annual screening; routine labs

For a representative high-risk diabetic patient (glucose: 168 mg/dL, BMI: 34.2 kg/m², sedentary lifestyle, age: 52 years), the system generated: (i) adoption of a low-glycemic Mediterranean dietary pattern with a 300–500 kcal daily deficit; (ii) progressive aerobic exercise regimen targeting 150 min/week with a graduated 12-week ramp-up protocol; (iii) evidence-based smoking cessation guidance and CBT referral; and (iv) immediate clinical

evaluation inclusive of fasting glucose, HbA1c, lipid profile, and renal function panel. All generated recommendations were cross-referenced against ADA 2024 Standards of Care [24] and ESC 2023 Cardiovascular Prevention Guidelines [25], and every output was found to be fully consistent with prevailing evidence-based clinical practice standards.

Real-World Application Potential

The PMRS is designed for practical deployment across multiple healthcare delivery contexts. In primary care settings, it can be integrated with existing electronic health record (EHR) systems to provide automated, real-time risk stratification and preventive guidance at point of care, reducing clinical workload while improving consistency of preventive counseling. In population health management programs, the system supports proactive outreach to high-risk individuals identified through routine clinical data. In resource-constrained settings— particularly relevant in India and other developing economies with high NCD burden—the PMRS offers a scalable, low-cost mechanism for evidence-based preventive care delivery. The explainability module is specifically designed to support clinical trust and regulatory compliance by providing auditable, patient-specific justification for each recommendation.

Limitations

Several limitations warrant acknowledgment. First, the benchmark datasets employed are relatively small (768 and 303 records respectively), which may limit the statistical power of performance estimates and restrict generalizability to broader clinical populations. Second, the PIMA dataset is limited to female patients of Pima Indian heritage, constraining demographic generalizability. Third, the rule-based recommendation engine, while guideline-concordant, does not yet incorporate drug interaction checks or contraindication screening—critical requirements for direct clinical deployment. Fourth, the system has not yet undergone prospective clinical validation through randomized trial. Fifth, privacy and data governance frameworks for federated deployment are not yet implemented, limiting the system to centralized data processing.

Additionally, the current system does not yet incorporate a feedback-driven learning mechanism for continuous model refinement, which will be addressed in future iterations.

CONCLUSION

This paper presented a Personalized Medical Recommendation System (PMRS) that integrates machine learning-based disease risk prediction, SHAP-driven explainability, and a structured guideline-consistent recommendation engine within a unified preventive healthcare framework. Random Forest achieved the highest classification performance (91.2% accuracy, 91.1% F1score, cross-validation $\sigma < 1.2\%$) across both benchmark datasets. SHAP analysis provided clinically coherent, feature level explanations that quantitatively enhance model transparency and support clinical adoption. The recommendation engine delivered personalized, evidence based health plans validated for guideline concordance against ADA 2024 and ESC 2023 standards.

The system further incorporates multi-modal data integration and an ensemble-based learning framework, enhancing predictive robustness and enabling more reliable clinical decision support. Additionally, severity risk stratification and prediction confidence estimation improve interpretability and allow prioritization of high-risk patients for timely intervention. The PMRS addresses three critical gaps in existing literature:

a. The absence of personalized automated recommendations accompanying ML-based risk predictions;
b. The opacity of complex ensemble models in clinical decision support contexts; and
c. The lack of an integrated pipeline unifying risk prediction, explainability, and personalized recommendation generation within a single cohesive system. By bridging the gap between algorithmic prediction and actionable clinical guidance, the PMRS represents a meaningful step toward intelligent, equitable, and scalable preventive healthcare delivery.

In addition, the system can be extended with a geo-spatial hospital recommendation module that leverages location-based services to identify nearby healthcare facilities in real time. By utilizing GPS or user-provided location data in conjunction with mapping platforms such as OpenStreetMap or Google Maps APIs, the system can recommend hospitals based on proximity, specialization, and service availability. A multi-criteria ranking mechanism can be employed to prioritize healthcare centers, ensuring that high-risk patients are directed to appropriate facilities with minimal delay. This integration enhances the practical utility of the system by bridging the gap between risk prediction and actionable healthcare access, particularly in time-sensitive scenarios.

Overall, the proposed PMRS demonstrates the potential of integrating machine learning, explainable AI, and intelligent healthcare services into a unified and scalable framework, thereby enabling transparent, data-driven, and accessible clinical decision support for next-generation preventive healthcare systems.

FUTURE WORK

Future research will pursue six primary directions: (i) integration with wearable monitoring devices for real-time dynamic risk surveillance using streaming ML architectures; (ii) expansion to additional disease domains (chronic kidney disease, stroke, COPD) using multi-task learning frameworks [17]; (iii) development of a mobile application interface for improved patient accessibility and longitudinal engagement tracking; (iv) prospective clinical validation through randomized controlled pilot studies at affiliated healthcare institutions; (v) exploration of federated learning architectures [11], [18] for privacy-preserving distributed model training across hospital networks; and (vi) integration of large language model (LLM) based natural language generation to deliver recommendation outputs in plain-language, patient-friendly format tailored to literacy level and language preference. (vii) development of a multilingual natural language processing (NLP)-based conversational chatbot capable of supporting interactive user engagement, symptom clarification, and context-aware recommendation explanation. The chatbot will leverage advanced language models to generate clinically coherent, patient-friendly responses and provide multilingual support to improve accessibility across diverse linguistic populations, particularly in

resource-constrained settings.

REFERENCES

1. World Health Organization. 2023. 'Non-communicable diseases', WHO Fact Sheets.
2. Rajpurkar, P. *et al.* 2022, 'AI in health and medicine', *Nature Med.*, vol. 28, no. 1, pp. 31-38.
3. Castelvecchi, D. 2016. 'Can we open the black box of AI?', *Nature*, vol. 538, pp. 20-23.
4. Kononenko, I. 2001. 'Machine learning for medical diagnosis: History, state of the art and perspective', *Artif. Intell. Med.,* vol. 23, no. 1, pp. 89-109.
5. Obermeyer, Z. *et al.* 2016. 'Predicting the future—big data, machine learning, and clinical medicine', *N. Engl. J. Med.,* vol. 375, no. 13, pp. 1216-1219.
6. Kavakiotis, I. *et al.* 2017. 'Machine learning and data mining methods in diabetes research', *Comput. Struct. Biotechnol. J.,* vol. 15, pp. 104-116.
7. Mohan, S. *et al.* 2019. 'Effective heart disease prediction using hybrid machine learning techniques', *IEEE Access,* vol. 7.
8. Ribeiro, MT. *et al.* 2016. 'Why should I trust you?' Explaining the predictions of any classifier', in *Proc. 22nd ACM SIGKDD*.
9. Lundberg, S. M. *et al.* 2017. 'A unified approach to interpreting model predictions', in *Adv. Neural Inf. Process. Syst. (NeurIPS)*.
10. Adadi, A. *et al.* 2018. 'Peeking inside the black-box: A survey on XAI', *IEEE Access*, vol. 6, pp. 52138-52160.
11. Nguyen, T. T. *et al.* 2023. 'Federated learning for smart healthcare: A survey', *ACM Comput. Surv.*, vol. 55, no. 3.
12. Rajpurkar, P. *et al.* 2023. 'The current and future state of AI interpretation of medical images', *N. Engl. J. Med.*, vol. 388, no. 21, pp. 1981-1990.
13. Li, J. *et al.* 2022. 'Personalized patient recommendation using graph neural networks and HER', *IEEE J. Biomed. Health Inform.*, vol. 26, no. 9, pp. 4664-4675.
14. Tonekaboni, S. *et al.* 2022. 'What clinicians want: Contextualizing explainable ML for clinical end use', in *Proc. MLHC Conf*.
15. Ahmad, M. A. *et al.* 2023, 'Interpretable machine learning in healthcare', in *Proc. ACM BCB Conf.*
16. Shailaja, K. *et al.* 2018. 'Machine learning in healthcare: A review', in *Proc. 2nd ICECA*.
17. Jiang, F. *et al.* 2022. 'Artificial intelligence in healthcare: Past, present and future', *Stroke Vasc. Neurol.,* vol. 2, no. 4, pp. 230-243.
18. Pan, S. *et al.* 2024. 'Privacy-preserving federated learning for personalized healthcare recommendation', *IEEE Trans. Neural Netw. Learn. Syst.*
19. Weng, S. F. *et al.* 2017. 'Can machine-learning improve cardiovascular risk prediction using routine clinical data?', *PLOS ONE*, vol. 12, no. 4.
20. Beam, A. L. 2018. 'Big data and machine learning in health care', *JAMA*, vol. 319, no. 13, pp. 1317-1318.
21. Chawla, N. V. *et al.* 2002. 'SMOTE: Synthetic minority over-sampling technique', *J. Artif. Intell. Res.*, vol. 16, pp. 321-357.
22. Breiman, L. 2001. 'Random forests', *Mach. Learn.*, vol. 45, no. 1, pp. 5-32.

23. Vapnik, V. 2000. *The nature of statistical learning theory*. Springer.
24. American Diabetes Association. 2024. 'Standards of care in diabetes—2024', *Diabetes Care*, vol. 47, pp. S1-S321.
25. Visseren, F. *et al.* 2021. '2021 ESC guidelines on cardiovascular disease prevention in clinical practice', *Eur. Heart J.*, vol. 42, no. 34, pp. 3227-3337.
26. Kaur, H. *et al.* 2022. 'Predictive modelling and analytics for diabetes using a machine learning approach', *Appl. Comput. Inform.*, vol. 18, no. 1/2, pp. 90-100.
27. Mahesh, B. 2022. 'Machine learning algorithms—A review', *Int. J. Sci. Res. (IJSR)*, vol. 9, no. 1, pp. 381-386.
28. Srinivasan, K. *et al.* 2022. 'Performance comparison of deep learning and ML models for diabetes mellitus detection', *Comput. Mater. Contin.*, vol. 69, no. 2, pp. 1999-2014.
29. Bzdok, D. 2022. 'Points of significance: Statistics versus machine learning', *Nature Methods*, vol. 15, no. 4, pp. 233-234.
30. Molnar, C. 2022. *Interpretable machine learning: A guide for making black box models explainable*. Leanpub.
31. Chen, M. *et al.* 2023. 'A healthcare predictive analytics framework for personalized chronic disease management using federated and explainable AI', *IEEE J. Biomed. Health Inform.*, vol. 27, no. 8, pp. 3910-3921.
32. Esteva, A. *et al.* 2021. 'Deep learning-enabled medical computer vision', *NPJ Digital Med.*, vol. 4, Art. no. 5.
33. Miotto, R. *et al.* 2026. 'Deep patient: An unsupervised representation to predict the future of patients from the electronic health records', *Sci. Rep.*, vol. 6.
34. Davenport, T. *et al.* 2019. 'The potential for artificial intelligence in healthcare', *Future Healthc. J.*, vol. 6, no. 2, pp. 94-98.
35. Waring, J. *et al.* 2020. 'Automated machine learning: Review of the state-of-the-art and opportunities for healthcare', *Artif. Intell. Med.,* vol. 104.