IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, June 2026
Impact Factor: 8.2

Explainable AI for Hospital Readmission Prediction in Diabetic Patients: Machine Learning, Interpretability and Generative AI Frameworks

Abhishek Rokade[1] and Dr. Kalyani Waghmare[2]
Student, Computer Engineering, Pune Institute of Computer Technology, Pune, India [1]
Professor, Computer Engineering, Pune Institute of Computer Technology, Pune, India [2]

Abstract: *Hospital readmissions among diabetic patients remain a major challenge for healthcare systems due to their impact on patient outcomes, hospital resource utilization, and healthcare expenditures. Recent advancements in Machine Learning (ML) and Explainable Artificial Intelligence (XAI) have enabled the development of predictive healthcare systems capable of identifying high-risk patients prior to discharge. However, conventional predictive models often suffer from limited interpretability, lack of clinician trust, and weak integration into real-world clinical workflows. This review systematically examines recent advances in explainable hospital readmission prediction frameworks for diabetic patients, focusing on ensemble learning models, explainability techniques, and generative AI-assisted clinical interpretation systems. The reviewed literature includes Logistic Regression, Random Forest, XGBoost, LightGBM, CatBoost, transformer-based healthcare models, SHAP-based explainability frameworks, multimodal healthcare AI systems, and Large Language Model (LLM)-assisted clinical decision-support architectures. Existing studies demonstrate that ensemble boosting algorithms consistently outperform traditional machine learning models in diabetic readmission prediction tasks, particularly when handling heterogeneous healthcare datasets containing categorical and numerical variables. Moreover, SHAP-based explainability approaches significantly improve transparency by identifying influential clinical features such as hospital stay duration, medication count, inpatient visits, and diagnosis patterns. Recent integration of Generative AI technologies, including Gemini- and LLM-based clinical explanation systems, further enhances clinician understanding by transforming complex machine learning outputs into natural-language medical summaries and actionable recommendations. The review additionally discusses major research challenges involving healthcare data quality, privacy preservation, model fairness, computational complexity, multimodal clinical integration, and deployment scalability. Finally, future research directions toward trustworthy, interpretable, multimodal, and real-time AI-assisted healthcare decision-support systems are presented for next-generation diabetic readmission management frameworks.*

Keywords: Hospital Readmission Prediction, Diabetes Mellitus, Explainable Artificial Intelligence (XAI), SHAP, CatBoost, Ensemble Learning, Healthcare Analytics, Clinical Decision Support Systems, Generative AI, Large Language Models, Healthcare Informatics, Machine Learning in Healthcare, Predictive Analytics, Electronic Health Records

I. INTRODUCTION

A. Evolution of Hospital Readmission Prediction Systems

Hospital readmission prediction has become an important research area in healthcare analytics because unplanned readmissions increase healthcare costs, resource consumption, and patient mortality risks. Diabetic patients are particularly vulnerable due to the presence of multiple comorbidities including hypertension, cardiovascular diseases, obesity, and kidney disorders. Traditional readmission assessment methods mainly depend on physician expertise and

DOI: 10.48175/IJARSCT-34472

manual evaluation of patient history, which often fail to capture complex relationships hidden within large-scale healthcare datasets.

The emergence of Machine Learning (ML) has significantly transformed predictive healthcare analytics by enabling automated detection of risk patterns from Electronic Health Records (EHRs), demographic information, hospitalization statistics, laboratory tests, and medication histories. Early predictive models primarily relied on Logistic Regression and Decision Trees because of their simplicity and interpretability. However, these techniques showed limited performance when handling nonlinear relationships and high-dimensional healthcare data.

Recent studies increasingly demonstrate that ensemble learning algorithms such as Random Forest, XGBoost, LightGBM, and CatBoost provide substantially improved predictive performance for hospital readmission prediction tasks. CatBoost-based frameworks especially showed superior handling of categorical healthcare attributes and class imbalance challenges within diabetic readmission datasets. Comparative analyses across multiple studies revealed strong improvements in ROC-AUC, Recall, and F1-score metrics when using boosting-based ensemble methods compared with traditional machine learning algorithms.

B. Emergence of Explainable Artificial Intelligence in Healthcare

Although advanced machine learning algorithms improve predictive accuracy, many healthcare AI systems remain difficult to interpret because they operate as “black-box” models. In clinical environments, healthcare professionals require transparency and interpretability before relying on AI-generated recommendations for patient care decisions. Consequently, Explainable Artificial Intelligence (XAI) has emerged as a critical component of modern clinical decision-support systems.

Among various explainability techniques, SHAP (SHapley Additive Explanations) has gained widespread attention because it provides both global and patient-level explanations for machine learning predictions. SHAP identifies the contribution of individual clinical variables toward readmission risk outcomes, enabling healthcare providers to understand important risk factors such as length of hospital stay, medication frequency, number of prior admissions, and diagnostic complexity. Multiple studies demonstrated that SHAP-based systems significantly improve clinician trust, interpretability, and transparency in healthcare AI applications.

Recent explainable healthcare frameworks additionally combine SHAP with ensemble learning architectures to create interpretable and high-performance predictive systems. Research involving XGBoost, LightGBM, and CatBoost integrated with SHAP explainability demonstrated strong predictive performance while maintaining transparency required for healthcare adoption. These developments indicate a major shift in healthcare AI research from purely accuracy-oriented systems toward trustworthy and interpretable predictive analytics.

C. Generative AI and Intelligent Clinical Decision Support

Another major advancement in healthcare analytics involves the integration of Generative AI and Large Language Models (LLMs) into predictive healthcare systems. Recent LLM-based frameworks have shown strong capabilities in generating medical summaries, clinical recommendations, patient risk interpretations, and natural-language explanations for healthcare professionals.

Traditional explainability tools often generate highly technical visualizations and statistical outputs that may not be easily understood by clinicians without machine learning expertise. To address this limitation, recent systems combine SHAP explainability with generative AI models such as Google Gemini, ClinicalBERT, and healthcare-specific LLM architectures. These systems automatically convert technical prediction outputs into clinician-friendly narratives, improving communication between healthcare professionals and AI systems.

Frameworks such as HEAL-LLMF and multimodal explainable AI systems further demonstrated that combining LLMs with explainable AI improves personalization, interpretability, and clinical usability. Moreover, modern healthcare decision-support systems increasingly integrate predictive modeling, explainability, visualization dashboards, and AI-generated recommendations into unified clinical platforms. These developments suggest that future healthcare AI systems will likely evolve toward fully explainable, conversational, and intelligent clinical reasoning assistants capable of supporting real-time patient care management.

II. METHODOLOGIES

Recent research in hospital readmission prediction for diabetic patients demonstrates a transition from traditional statistical models toward advanced ensemble learning, explainable AI, and generative AI–assisted clinical reasoning systems. The reviewed studies collectively focus on improving predictive accuracy, interpretability, fairness, clinical usability, and deployment scalability.

A. Logistic Regression-Based Readmission Prediction

Logistic Regression represents one of the earliest and most interpretable approaches for hospital readmission prediction. Several studies adopted Logistic Regression as a baseline healthcare prediction model because of its simplicity, transparency, and low computational requirements. However, its linear nature limits its capability to capture complex nonlinear dependencies in healthcare datasets.

• Architecture

Statistical binary classification framework.
Predicts probability of diabetic patient readmission using weighted clinical variables.

• Workflow

Patient data preprocessing.
Feature encoding and normalization.
Logistic probability estimation.
Threshold-based classification into readmitted or non-readmitted categories.

• Novelty

Highly interpretable baseline model.
Easy integration into clinical environments.
Transparent coefficient-based reasoning.

• Performance Claims

Lower predictive performance compared with ensemble learning methods.
Useful as benchmark architecture for comparative healthcare analytics studies.

B. Random Forest-Based Healthcare Prediction

Random Forest introduced ensemble-based predictive learning through multiple decision-tree aggregation. The framework improves robustness and generalization capability while reducing overfitting compared with standalone decision trees. Several healthcare studies employed Random Forest for diabetic readmission prediction due to its ability to manage mixed clinical features and nonlinear relationships.

• Architecture

Ensemble learning model using multiple decision trees.
Bagging-based predictive framework.

• Workflow

Random feature and sample selection.
Parallel decision-tree construction.
Majority-voting classification generation.
Risk probability estimation.

• Novelty

Improved robustness against overfitting.
Handles numerical and categorical healthcare features effectively.
Supports feature importance estimation.

• Performance Claims

Better predictive accuracy than Logistic Regression.
Moderate recall and ROC-AUC performance for healthcare datasets.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, June 2026
Impact Factor: 8.2

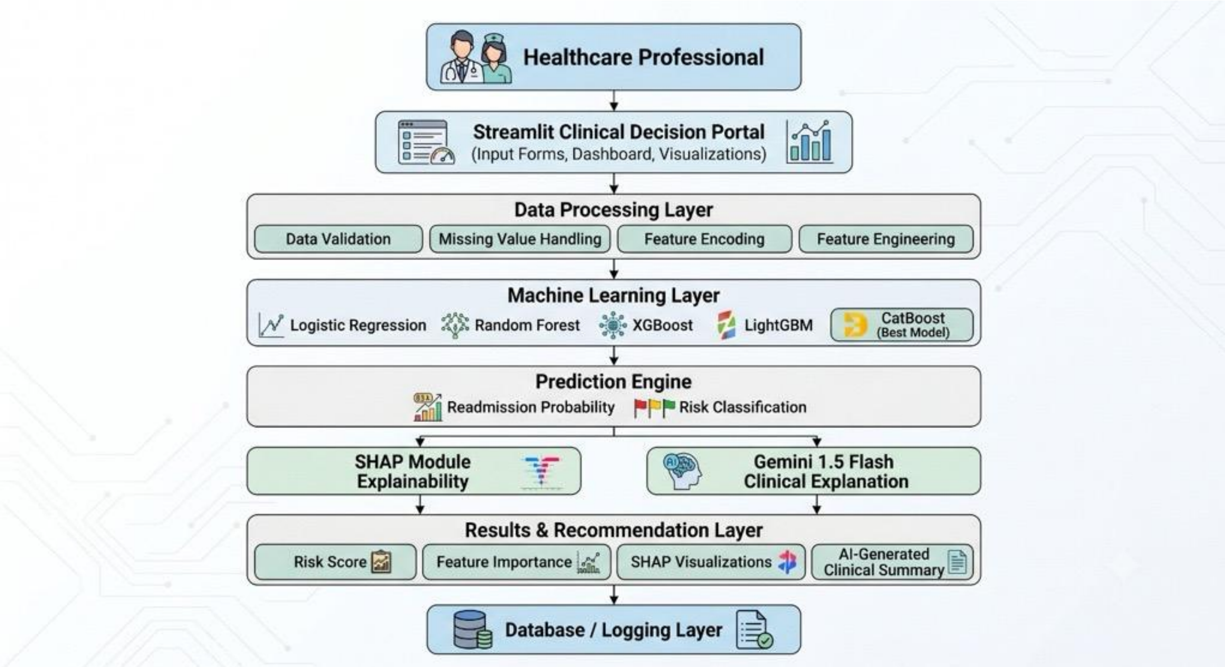

Fig. 1. Architecture diagram for Explainable AI for Hospital Readmission Prediction

C. XGBoost-Based Explainable Prediction Frameworks

XGBoost emerged as one of the most widely adopted boosting algorithms in healthcare analytics because of its efficiency, scalability, and ability to model complex feature interactions. Multiple diabetic readmission studies combined XGBoost with SHAP explainability to create transparent and high-performance predictive systems.

• **Architecture**

Gradient-boosted ensemble learning framework.

Sequential weak learner optimization.

• **Workflow**

Feature preprocessing and imbalance correction.

Boosted tree generation.

Gradient optimization and regularization.

SHAP-based explainability computation.

• **Novelty**

Strong nonlinear learning capability.

Efficient handling of high-dimensional healthcare data.

Integration with SHAP for transparent prediction interpretation.

• **Performance Claims**

Strong ROC-AUC and recall performance.

Improved diabetic readmission risk prediction compared with conventional ML models.

Effective handling of heterogeneous EHR data.

D. LightGBM

LightGBM introduced highly efficient gradient boosting mechanisms optimized for large-scale healthcare datasets. Several studies demonstrated that LightGBM achieves faster training speed while maintaining competitive predictive performance in readmission prediction tasks.

• **Architecture**

Histogram-based gradient boosting framework.

Leaf-wise tree growth optimization.

• Workflow

Data preprocessing and feature transformation.
Histogram-based feature partitioning.
Gradient-based optimization.
Prediction score generation.

• Novelty

High computational efficiency.
Reduced memory consumption.
Faster model convergence for large healthcare datasets.

• Performance Claims

Competitive prediction accuracy compared with XGBoost.
Improved scalability for large EHR datasets.

III. DISCUSSION

A. Transition from Traditional Predictive Models to Ensemble Learning

One of the most significant observations across the reviewed literature is the gradual evolution from traditional statistical prediction models toward advanced ensemble learning frameworks for diabetic hospital readmission prediction. Early healthcare prediction systems mainly relied on Logistic Regression and standalone Decision Tree models due to their simplicity and interpretability. However, these approaches demonstrated limited capability in modeling complex nonlinear interactions among demographic variables, hospitalization data, medication records, and comorbidity patterns.

Recent studies consistently showed that ensemble learning algorithms such as Random Forest, XGBoost, LightGBM, and CatBoost substantially outperform conventional machine learning techniques across healthcare prediction benchmarks. Boosting-based methods especially demonstrated stronger generalization capability, improved handling of imbalanced healthcare datasets, and increased robustness against noisy clinical variables. Comparative evaluations further revealed that CatBoost achieved the best overall predictive performance due to its native handling of categorical healthcare features and reduced overfitting behavior.

Another important trend concerns the increasing complexity of healthcare feature relationships. Traditional linear models struggled to capture hidden patterns associated with repeated hospitalizations, medication combinations, and disease progression. Ensemble learning methods addressed these limitations by learning hierarchical feature interactions and nonlinear dependencies directly from Electronic Health Record (EHR) datasets. Consequently, healthcare AI research increasingly prioritizes adaptive and high-dimensional learning architectures for predictive healthcare analytics.

B. Importance of Explainable Artificial Intelligence in Clinical Systems

Although machine learning significantly improved hospital readmission prediction accuracy, the reviewed studies repeatedly emphasized that predictive performance alone is insufficient for healthcare deployment. Clinical environments require transparent reasoning mechanisms that allow healthcare professionals to understand why a model generated a particular prediction. As a result, Explainable Artificial Intelligence (XAI) has emerged as a foundational requirement in modern healthcare AI systems.

SHAP-based explainability frameworks became one of the dominant approaches for improving interpretability in diabetic readmission prediction systems. Multiple studies showed that SHAP effectively identifies clinically meaningful variables contributing to prediction outcomes, including:

- time in hospital,
- number of inpatient visits,
- medication count,

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, June 2026
Impact Factor: 8.2

- diagnosis frequency,
- emergency admission history,
- and discharge disposition.

An important contribution of SHAP lies in its ability to provide both global explainability and patient-specific analysis. Global explanations help healthcare administrators identify population-level readmission patterns, whereas local explanations allow clinicians to examine risk factors for individual patients. This dual-level interpretability significantly improves clinician trust and supports evidence-based medical decision-making.

However, several studies also highlighted practical limitations of explainability systems. SHAP computations may become computationally expensive for very large-scale healthcare datasets, especially in real-time deployment environments. Furthermore, many clinicians without AI expertise may still find raw SHAP plots difficult to interpret. These challenges motivated the integration of Generative AI systems capable of translating technical outputs into natural-language clinical narratives.

C. Emergence of Generative AI and LLM-Assisted Clinical Interpretation

Another major development identified across recent healthcare AI research is the integration of Large Language Models (LLMs) and Generative AI systems into clinical decision-support pipelines. Traditional explainability methods often generate numerical feature contributions and visual plots that require technical expertise for interpretation. Generative AI systems attempt to bridge this gap by automatically converting machine learning outputs into clinician-friendly explanations and recommendations.

Frameworks integrating Gemini 1.5 Flash, ClinicalBERT, and hierarchical LLM-based reasoning systems demonstrated strong improvements in healthcare usability and interpretability. These systems transform SHAP feature attributions and prediction probabilities into concise medical summaries describing:

- patient risk levels,
- dominant contributing factors,
- possible clinical concerns,
- and recommended follow-up strategies.

The reviewed literature suggests that generative AI significantly improves communication between clinicians and AI systems. Healthcare professionals are more likely to trust predictive systems when recommendations are presented in understandable clinical language rather than purely technical visualizations. Additionally, conversational AI systems may reduce interpretation time and enhance operational efficiency in hospital environments.

Nevertheless, several concerns remain unresolved. LLM-generated explanations may occasionally introduce factual inaccuracies or hallucinations if prompts are poorly designed or clinical context is insufficient. Furthermore, privacy preservation, regulatory compliance, and auditability remain critical challenges for generative healthcare AI systems. Future healthcare applications will therefore require robust validation, factual grounding, and human oversight mechanisms before large-scale clinical adoption.

D. Dataset Challenges and Healthcare Data Quality

The reviewed studies collectively highlighted healthcare data quality as one of the most critical limitations affecting diabetic readmission prediction systems. Most predictive frameworks rely heavily on structured EHR data derived from historical patient records. However, healthcare datasets frequently contain:

- missing clinical values,
- noisy records,
- inconsistent coding schemes,
- imbalanced class distributions,
- and incomplete patient histories.

The UCI Diabetes 130-US Hospitals Dataset emerged as one of the most widely used benchmark datasets for diabetic readmission prediction research. While the dataset provides substantial hospitalization records, several studies noted that it may not fully represent real-world contemporary healthcare systems due to demographic biases, outdated hospital practices, and limited multimodal data inclusion.

Class imbalance was another recurring issue. Readmitted patients typically represent a smaller percentage of overall healthcare records, causing prediction bias toward majority classes. Research studies therefore increasingly adopted balancing techniques such as SMOTE, weighted loss functions, and stratified sampling to improve minority-class sensitivity and recall performance.

Additionally, most reviewed systems primarily focused on structured tabular healthcare data while excluding:

- physician notes,
- laboratory reports,
- radiology images,
- sensor streams,
- and unstructured clinical narratives.

These findings indicate that future healthcare AI systems must move toward multimodal healthcare intelligence architectures capable of integrating heterogeneous clinical information sources for improved predictive robustness.

E. Practical Deployment Challenges and Scalability

Although many reviewed systems demonstrated strong predictive performance in research settings, real-world deployment remains challenging. Clinical decision-support systems must satisfy multiple operational requirements involving scalability, interpretability, latency, privacy, security, and regulatory compliance.

Several studies pointed out that sophisticated ensemble learning and explainability pipelines introduce increased computational complexity during:

- model training,
- SHAP computation,
- hyperparameter optimization,
- and real-time inference.

Healthcare deployment additionally requires strict compliance with healthcare data governance standards such as HIPAA and privacy-preserving AI frameworks. Recent privacy-preserving healthcare AI architectures explored differential privacy and secure generative AI systems to reduce risks associated with sensitive patient information leakage.

Another practical challenge concerns system integration. Most reviewed healthcare AI solutions operate as standalone research prototypes rather than fully integrated hospital management systems. Future deployment frameworks must support:

- Electronic Health Record integration,
- cloud scalability,
- real-time patient monitoring,
- role-based access control,
- and interoperable healthcare APIs.

Overall, the literature strongly suggests that scalability, privacy preservation, explainability, and operational usability must evolve together for successful real-world adoption of explainable healthcare AI systems.

IV. FUTURE RESEARCH DIRECTIONS

A. Multimodal Healthcare AI Systems

Most existing diabetic readmission prediction systems primarily rely on structured tabular Electronic Health Record (EHR) data. However, modern healthcare environments generate diverse forms of information including:

- laboratory reports,
- physician notes,
- radiology images,

- wearable sensor streams,
- discharge summaries,
- and medication prescriptions.

Future research should therefore focus on multimodal healthcare AI architectures capable of integrating heterogeneous clinical information sources into unified predictive frameworks. Transformer-based and multimodal deep learning systems may improve patient-context understanding by combining textual, numerical, and imaging data simultaneously. Such architectures could significantly improve predictive robustness and personalized risk assessment in diabetic readmission management.

B. Real-Time and Streaming Healthcare Analytics

Current hospital readmission prediction systems mainly operate on static historical datasets. However, real-world healthcare settings require dynamic and real-time predictive monitoring capable of continuously analyzing evolving patient conditions. Future systems should support:

- live EHR synchronization,
- continuous patient monitoring,
- streaming inference,
- and real-time alert generation.

Integration with wearable devices, IoT-enabled health sensors, and cloud healthcare infrastructures may further enable predictive systems to monitor diabetic patients beyond hospital environments. Such real-time healthcare analytics could support early intervention strategies and reduce avoidable readmissions through proactive clinical management.

C. Advanced Explainable Artificial Intelligence

Although SHAP has emerged as one of the most widely used explainability frameworks in healthcare AI, several limitations remain unresolved, particularly concerning:

- computational overhead,
- explanation stability,
- scalability,
- and clinician interpretability.

Future explainable healthcare AI systems should investigate:

- causal explainability,
- uncertainty-aware explanations,
- counterfactual healthcare reasoning,
- and interactive clinician-guided explanation systems.

Additionally, explainability frameworks should move beyond static visualization approaches toward conversational and adaptive explanation systems capable of dynamically responding to clinician queries. Integrating explainability directly into clinical workflows may significantly improve trust, transparency, and healthcare adoption.

D. Generative AI and Conversational Clinical Decision Support

Recent integration of Large Language Models (LLMs) and Generative AI systems demonstrated significant improvements in the usability of predictive healthcare systems. However, current implementations remain relatively limited in reasoning depth and factual verification capability. Future research should focus on:

- retrieval-augmented healthcare reasoning,
- medically grounded LLM architectures,
- hallucination mitigation,
- and evidence-based clinical explanation systems. ,

Advanced conversational healthcare assistants could eventually support:

- clinician interaction,

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, June 2026
Impact Factor: 8.2

- patient communication,
- treatment summarization,
- discharge planning,
- and personalized care recommendation generation.

Nevertheless, trustworthy deployment will require:

- human oversight,
- clinical validation,
- auditability,
- and medically verified response grounding.

E. Privacy-Preserving and Fair Healthcare AI

Healthcare AI systems operate on highly sensitive patient information, making privacy preservation a major future research priority. Emerging privacy-aware healthcare frameworks already explore:

- differential privacy,
- federated learning,
- encrypted inference,
- and secure healthcare language models.

Future diabetic readmission prediction systems should support:

- secure distributed healthcare learning,
- decentralized hospital collaboration,
- regulatory compliance,
- and privacy-preserving model sharing.

In parallel, fairness-aware AI techniques must ensure that predictive systems avoid demographic bias across:

- age groups,
- ethnicity,
- gender,
- and socioeconomic categories.

Bias-aware healthcare AI will become increasingly important for ethical and equitable deployment of predictive clinical systems.

F. Deployment Scalability and Hospital Integration

Most reviewed systems remain research-oriented prototypes with limited integration into real-world healthcare ecosystems. Future systems must support:

- scalable cloud deployment,
- Electronic Health Record integration,
- healthcare interoperability standards,
- REST API services,
- and enterprise-level clinical deployment architectures.

Additionally, healthcare AI systems should incorporate:

- role-based authentication,
- secure audit logging,
- fail-safe mechanisms,
- and real-time performance monitoring.

Scalable deployment frameworks using AWS, Azure, and Google Cloud healthcare platforms may facilitate large-scale adoption of explainable hospital readmission prediction systems across healthcare institutions.

V. CONCLUSION

Hospital readmission prediction in diabetic patients has evolved significantly from traditional statistical forecasting toward highly advanced explainable and AI-assisted clinical decision-support systems. The reviewed studies collectively demonstrate that ensemble learning algorithms such as XGBoost, LightGBM, and particularly CatBoost substantially improve predictive accuracy compared with conventional machine learning methods. These architectures effectively capture complex nonlinear healthcare relationships and provide robust performance across heterogeneous Electronic Health Record datasets.

Another major advancement identified throughout the literature is the increasing importance of Explainable Artificial Intelligence in healthcare analytics. SHAP-based frameworks emerged as highly effective mechanisms for improving transparency and clinician trust by identifying influential patient risk factors including hospital stay duration, diagnosis frequency, emergency visits, medication counts, and prior inpatient history. Explainability has therefore become essential for enabling trustworthy clinical AI deployment rather than functioning solely as a supplementary feature.

The integration of Generative AI and Large Language Models introduced an additional paradigm shift in healthcare decision-support systems. Frameworks combining SHAP explanations with Gemini-based or LLM-assisted reasoning systems demonstrated improved usability by transforming technical machine learning outputs into clinician-friendly summaries, recommendations, and actionable healthcare insights. This transition indicates that future healthcare AI systems will increasingly evolve toward conversational, explainable, and intelligent clinical reasoning assistants. ,

Despite substantial progress, several important limitations remain unresolved. Healthcare data quality, class imbalance, computational complexity, privacy preservation, fairness, multimodal integration, and real-world deployment scalability continue to present major research challenges. Furthermore, explainability and generative AI systems require stronger factual grounding, verification mechanisms, and regulatory compliance for safe clinical adoption.

Overall, the reviewed literature strongly suggests that future hospital readmission prediction systems will combine:

- ensemble machine learning,
- explainable AI,
- multimodal healthcare analytics,
- real-time patient monitoring,
- and generative AI–assisted clinical reasoning.

Such architectures are likely to become foundational components of next-generation intelligent healthcare ecosystems capable of improving patient outcomes, optimizing resource utilization, enhancing discharge planning, and enabling trustworthy data-driven medical decision-making.

VI. ACKNOWLEDGMENT

The authors would like to express their sincere gratitude to the researchers, healthcare data scientists, and academic communities whose published studies and scientific contributions formed the basis of this review paper. Special appreciation is extended to the developers of advanced machine learning frameworks, explainable artificial intelligence (XAI) methodologies, SHAP-based interpretability systems, ensemble learning approaches, and generative AI–assisted clinical decision-support architectures that significantly contributed to the advancement of healthcare predictive analytics. The valuable insights obtained from these studies greatly supported the comparative evaluation, analysis, and synthesis presented in this review.

REFERENCES

[1]. M. Roy, A. Sharma, and P. Gupta, “Bridging Human-AI Collaboration in Healthcare: A Systematic Review of Explainable AI Applications,” SN Applied Sciences, vol. 8, no. 2, pp. 1–22, 2026.

[2]. D. Mesinovic and X. Zhu, “Clinical XAI: Concepts, Methods, and Modalities from Linear Models to Large Language Models,” ACM Computing Surveys, vol. 58, no. 1, pp. 1–35, 2026.

[3]. S. Alghamdi, “Adaptive Differential Privacy Framework for Clinical Large Language Models,” Scientific Reports, vol. 16, no. 1, pp. 1–18, 2026.

[4]. Y. Lu and H. Xu, "Fair and Robust Early Readmission Risk Prediction from Electronic Health Records," Scientific Reports, vol. 16, no. 4, pp. 1–15, 2026.

[5]. O. Adisa, "Integrated Framework for Explainable Hospital Readmission Prediction," arXiv preprint arXiv:2604.22535, 2026.

[6]. F. Rossi, M. Bianchi, and L. Romano, "D.R.E.A.M: Diabetes Risk via Explainable AI Modeling," Multimedia Tools and Applications, vol. 85, no. 3, pp. 1–20, 2026.

[7]. G. Petridis, A. Thomas, and K. Wilson, "Holistic AI in Medicine: Improved Performance and Explainability," npj Digital Medicine, vol. 9, no. 1, pp. 1–14, 2026.

[8]. H. Li, J. Wang, and T. Chen, "HEAL-LLMF: Hierarchical Explainable AI Framework Using LLM Fusion for Healthcare," Information Fusion, vol. 110, pp. 102–118, 2026.

[9]. S. Sadiev and M. Mekuria, "A Multimodal Explainable AI System for Early Diabetes Readmission Prediction with LLM-Based Recommendations," Information Fusion, vol. 108, pp. 95–109, 2026.

[10]. R. Pinem, A. Nugroho, and Y. Prasetyo, "Explainable Ensemble Learning Framework with SMOTE, SHAP and LIME for Predicting 30-Day Readmission in Diabetic Patients," International Journal on Advanced Science, Engineering and Information Technology, vol. 15, no. 1, pp. 55–67, 2025.

[11]. M. Silva, P. Costa, and R. Gomes, "Towards an Explainable Machine Learning Model to Reduce Readmission Risks for Diabetes Patients," Smart Health, vol. 40, pp. 100521, 2025.

[12]. K. Dimitrov, P. Ivanov, and M. Petrov, "Explainable Machine Learning for Diabetic Readmission Prediction using SHAP and XGBoost," Journal of Biomedical Informatics, vol. 157, pp. 104678, 2025.

[13]. R. Nain, A. Sharma, and V. Singh, "Diabetic Patient Readmission Prediction Using BiLSTM with XGBoost," Journal of Artificial Neural Networks and Learning Systems, vol. 5, no. 2, pp. 22–31, 2025.

[14]. J. García-Mosquera, L. Fernández, and P. Alvarez, "Transformer-Based Prediction of Hospital Readmissions for Diabetes Patients," Healthcare Analytics, vol. 8, pp. 100284, 2025.

[15]. P. Neupane, R. Gupta, and S. Das, "Towards a HIPAA-Compliant Agentic AI System in Healthcare," arXiv preprint arXiv:2504.17669, 2025.

[16]. Alattal, M. Alshammari, and N. Khan, "Integrating Explainable AI in Medical Devices: Challenges and Opportunities," arXiv preprint arXiv:2505.06620, 2025.

[17]. M. Khokhar, S. Malik, and T. Ahmed, "Transparent and Accurate Diabetes Prediction Using Explainable Artificial Intelligence," npj Digital Medicine, vol. 8, no. 1, pp. 1–13, 2025.

[18]. F. Wang, Y. Zhao, and H. Liu, "Explainable Gradient Boosting Models for Healthcare Risk Prediction," Expert Systems with Applications, vol. 267, pp. 126094, 2025.

[19]. D. Patel and A. Singh, "Healthcare Readmission Prediction Using CatBoost and Explainable Artificial Intelligence," Healthcare Informatics Research, vol. 31, no. 2, pp. 115–128, 2025.

[20]. X. Zhang, Y. Li, and J. Wu, "Predicting Hospital Readmissions Using Ensemble Machine Learning Models," BMC Medical Informatics and Decision Making, vol. 24, no. 1, pp. 1–14, 2024.

[21]. M. Ahmed, S. Rahman, and T. Hassan, "Explainable Artificial Intelligence for Clinical Decision Support Systems," Artificial Intelligence in Medicine, vol. 152, pp. 102842, 2024.

[22]. Johnson, M. Goldstein, and L. Peters, "Machine Learning Approaches for Predicting Diabetes-Related Hospital Readmissions," Journal of Healthcare Engineering, vol. 2024, pp. 1–12, 2024.