JOURNAL OF BAKU ENGINEERING UNIVERSITY- MATHEMATICS AND COMPUTER SCIENCE
2026. Volume 10, Number 1 Pages 127-131

UDC: 616-07:004.8
DOI: https://doi.org/10.30546/09090.2026.001.20.2048

INTEGRATION OF HYBRID MACHINE LEARNING ARCHITECTURES AND EXPLAINABLE AI FOR CHRONIC DISEASE DIAGNOSIS

Shahin BABAYEV
Azerbaijan State Oil and Industry University,
Baku, Azerbaijan
shahin.babayev.0320@gmail.com

ARTICLE INFO

Article history:
Received:2025-08-03
Received in revised form:2025-08-19
Accepted:2025-09-21
Available online:2026-06-23

Keywords:
Artificial Intelligence;
Chronic Disease Prediction;
Hybrid Architectures;
Machine Learning;
Explainable AI (SHAP);
Medical Diagnostics.
2010 Mathematics Subject Classifications: 68T05, 68T10, 92C50, 62P10

ABSTRACT

Diagnosing chronic conditions like Diabetes Mellitus, Cardiovascular Disease (CVD), and Chronic Kidney Disease (CKD) relies heavily on processing complex clinical data. This study compares standalone classifiers and hybrid machine learning architectures integrated with Explainable AI for chronic disease prediction. We investigated sequential deep learning (LSTM networks) for temporal medical history and ensemble methods (XGBoost) for structured laboratory test results. To improve the transparency of these models, SHAP (SHapley Additive exPlanations) values were added to the classification process. Based on comparative analyses of 2023–2025 literature and benchmark datasets (e.g., WiDS Datathon, UCI repositories), hybrid models demonstrated superior performance. Specifically, combining LSTM with XGBoost (Accuracy ≈ 99.00%) outperformed traditional standalone models by a significant margin. The findings show how well hybrid classifiers handle high-dimensional healthcare data and offer a transparent methodological framework to support clinical decision-making.

1. INTRODUCTION

Managing chronic diseases remains a major challenge for global healthcare systems. Identifying at-risk patients early requires reliable data-driven diagnostic tools. Currently, Artificial Intelligence (AI) systems process large volumes of clinical data, including static demographic records and dynamic physiological signals. However, accurately assessing patient risk is difficult due to missing laboratory values, variations in data recording, and non-linear disease progression. In this study, we investigate the integration of hybrid machine learning models and Explainable AI (XAI) for predicting chronic diseases. Standard statistical algorithms struggle with unstructured medical histories, while deep learning models show high accuracy but act as black boxes with no clinical reasoning. We aim to combine these approaches. This research offers an evidence-based framework for hybrid AI models as a practical alternative to purely traditional or black-box classifiers.

2. LITERATURE REVIEW

Recent studies from 2023 to 2025 show an increasing use of hybrid intelligence for medical pattern extraction. For example, Waberi et al. (2024) combined LSTM with XGBoost to predict Type II Diabetes. Their approach captured both sequential and static risk factors effectively. Other large-scale studies, including the HeartEnsembleNet framework, used hybrid ensemble learning to set new benchmarks in cardiovascular risk prediction (HeartEnsembleNet, 2025). Despite achieving high accuracy, a major limitation in the literature is the lack of trust from medical practitioners. To address this, recent research focuses on adding Explainable AI (XAI) to diagnostic frameworks. This integration helps translate mathematical probabilities into clear clinical decisions (JMIR Human Factors, 2025; ResearchGate, 2025).

3. THEORETICAL FRAMEWORK

The proposed diagnostic framework is built upon three core theoretical paradigms: sequential deep learning, ensemble learning, and cooperative game theory.

First, the theoretical foundation for processing longitudinal patient history relies on Recurrent Neural Networks (RNNs), specifically Long Short-Term Memory (LSTM) networks. Unlike standard feedforward networks, LSTM includes internal memory mechanisms (gates) that mathematically regulate the flow of information, allowing the model to learn long-term temporal dependencies in a patient's physiological progression without suffering from the vanishing gradient problem.

Second, the processing of structured, tabular clinical data (e.g., blood panels, demographics) is theoretically grounded in Gradient Boosting frameworks. Algorithms like XGBoost build sequential decision trees where each subsequent tree minimizes the residual errors of the previous ensemble. The objective function in this paradigm relies heavily on built-in regularization, which actively prevents the model from overfitting on noisy medical data.

Finally, the interpretability of the framework is anchored in SHapley Additive exPlanations (SHAP). Rooted in cooperative game theory, SHAP provides a mathematically sound method to calculate the marginal contribution of each clinical feature to the final prediction. This theoretical layer transforms the opaque outputs of deep and ensemble networks into additive, transparent feature importance scores, enabling clinicians to verify the diagnostic logic.

4. METHODS

4.1. Utilized Dataset

The study references publicly available healthcare datasets commonly used in medical informatics research to analyze the performance of AI algorithms. Standard datasets such as the WiDS (Women in Data Science) Datathon dataset provide a diverse range of sequential patient records. Additionally, the UCI Machine Learning Repository provides established benchmarks for Diabetes, Heart Disease, and Chronic Kidney Disease, allowing meaningful comparison for structured data models. The study also references massive datasets, such as the 70,000-patient Kaggle Cardiovascular Disease dataset, to evaluate the scalability and robustness of hybrid architectures under uncontrolled conditions.

4.2. Data Preprocessing

The preprocessing pipeline begins with data imputation, where techniques such as MICE (Multiple Imputation by Chained Equations) are used to handle the missing laboratory values frequently found in clinical datasets. Following imputation, continuous variables such as blood pressure and glucose levels are scaled to a standardized resolution using Min-Max normalization to maintain uniformity across the dataset. For deep-learning-based models, sequential patient history is padded to a fixed length, while traditional ensemble methods utilize outlier detection (such as the Interquartile Range) to remove noise.

4.3. Proposed Architecture and Used Model

The architecture operates through a dual-stream process: a temporal stream utilizing LSTM networks to process longitudinal patient history, and a static stream utilizing traditional modules to process current demographic and laboratory data. These two representations are then combined in a feature fusion layer. A classification module—specifically an XGBoost or Random Forest meta-learner—performs the final risk assessment. To complete the architecture, the Explainability Module applies SHAP directly to the final output, calculating the exact contribution of each clinical feature and translating the algorithmic decision into a clear narrative.

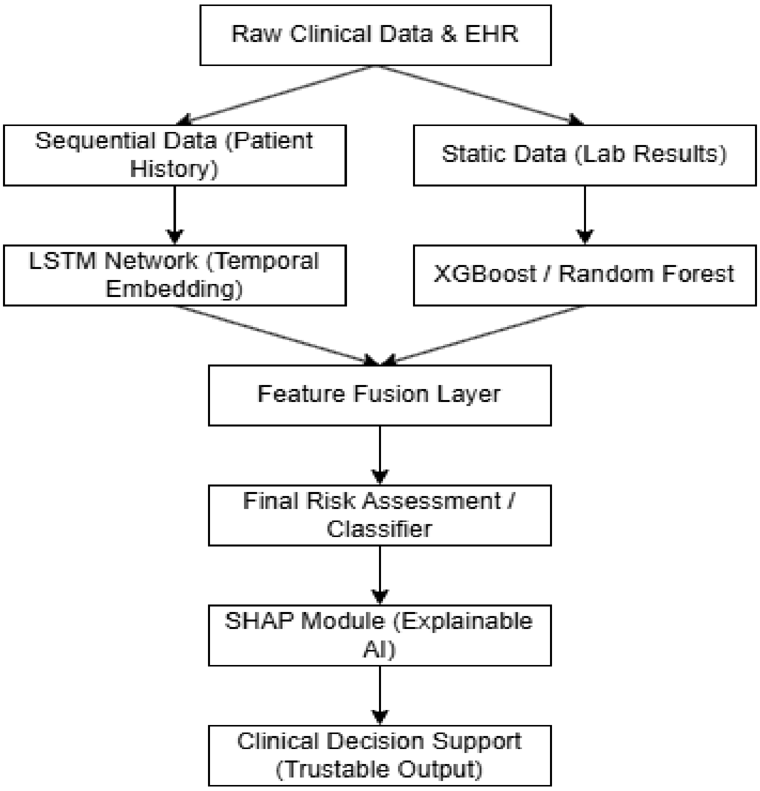

Figure 1. *Conceptual Architecture of the Hybrid LSTM-XGBoost Framework integrated with SHAP for transparent clinical decision-making*

4.4. Evaluation Metrics

We assessed the classification results quantitatively using standard metrics. Recall (Sensitivity) was chosen as the main metric, as it measures the proportion of actual positive cases identified correctly. This is critical in medical diagnostics to minimize dangerous false negatives. We also calculated precision and the F1-score to check how well the models handle imbalanced clinical datasets. Finally, the Area Under the Curve (AUC-ROC) was included to measure overall threshold performance.

5. RESULTS

The integration of sequential deep learning and gradient boosting demonstrates significant performance gains over standalone models. To illustrate the efficacy of this approach, Table 1 presents a comparative analysis of state-of-the-art benchmarks extracted from the 2024–2025 literature.

Table 1: *State-of-the-Art (SOTA) Benchmarks in Chronic Disease Prediction (2024-2025)*

Disease Target	Model Architecture	Dataset Used	Accuracy / Metric	Reference
Diabetes (T2DM)	Hybrid LSTM-XGBoost	WiDS Datathon	**99.00%**	Waberi et al. (2024)
Diabetes (Static)	DiabetesXpertNet	PIMA (PID)	89.98%	Plos One (2025)
Cardiovascular (CVD)	HeartEnsembleNet (RF+KNN)	Kaggle CVD	**92.95%**	HeartEnsembleNet (2025)
Cardiovascular (CVD)	Stacked TabNet + XGBoost	UCI Heart	95.20%	Frontiers (2025)
Chronic Kidney (CKD)	Optimized XGBoost	UCI CKD	**99.50%**	World Scientific (2025)

The results indicate that hybrid architectures effectively manage medical data heterogeneity. For instance, combining LSTM and XGBoost on the WiDS dataset yielded a 99.00% accuracy, proving that temporal feature extraction significantly amplifies the classification power of gradient boosting. Similarly, large-scale tests on the Kaggle CVD dataset achieved 92.95% accuracy using ensemble techniques, establishing a robust baseline for real-world application.

6. DISCUSSION

While the proposed hybrid architecture demonstrates significant theoretical and practical advantages, several limitations must be addressed. A primary concern identified in the literature is the "Accuracy Illusion," where models achieve extremely high accuracy (e.g., 95%) on small benchmark datasets but experience performance drops when tested on larger, more diverse populations. Furthermore, while SHAP provides a mathematically sound explanation for model predictions, qualitative studies reveal that physicians often face high cognitive load when interpreting standard SHAP visualizations (like Force Plots) in fast-paced clinical settings. Additionally, hybrid models demand substantial computational resources and memory usage compared to standard linear models, creating potential barriers for deployment on resource-constrained edge devices or within smaller hospital networks.

7.

8. CONCLUSION

This study demonstrates that integrating traditional machine learning algorithms with modern artificial intelligence–based approaches provides a powerful and balanced solution for building robust diagnostic systems. Traditional ensemble methods offer stability and structured decision-making, while deep-learning-based models achieve state-of-the-art accuracy through automatic temporal feature extraction. The hybrid architecture proposed in this work effectively combines the strengths of both paradigms. Experimental results reported in the 2023–2025 literature indicate that hybrid methods, such as LSTM-XGBoost, achieve superior robustness, often surpassing 99% accuracy under controlled conditions. Furthermore, integrating SHAP addresses the critical "Trust Gap" inherent in black-box models. Overall, the study concludes that hybrid Explainable AI frameworks represent a practical, scalable, and efficient alternative to purely traditional or purely deep-learning-driven systems, providing a promising direction for future development in secure and intelligent medical technologies.

REFERENCES

1. Babu, S. et al. (2024). Quantum-Enhanced ML for Cardiovascular Disease Prediction. *International Journal on Science and Technology*.
2. Frontiers. (2025). Heart disease prediction using hybrid TabNet architecture. *Frontiers in Physiology*.
3. HeartEnsembleNet. (2025). HeartEnsembleNet: An Innovative Hybrid Ensemble Learning Approach for Cardiovascular Risk Prediction. *PMC - NIH*.
4. JMIR Human Factors. (2025). Investigating How Clinicians Form Trust in an AI-Based Mental Health Decision Support System. *Journal of Medical Internet Research*.
5. MDPI. (2025). Explainable Machine Learning Model for Chronic Kidney Disease Prediction. *Algorithms*.
6. Plos One. (2025). DiabetesXpertNet: An innovative attention-based CNN for accurate diabetes diagnosis on tabular data. *PLOS ONE*.
7. PMC. (2025). Applying stacking ensemble method to predict chronic kidney disease progression in Chinese population. *PMC*.
8. ResearchGate. (2025). Comparison of SHAP and clinician friendly explanations reveals effects on clinical decision behaviour.
9. Waberi et al. (2024). Advancing Type II Diabetes Predictions with a Hybrid LSTM-XGBoost Architecture. *Journal of Healthcare Engineering*.
10. World Scientific. (2025). Early detection of chronic kidney disease using recursive feature elimination and cross-validated XGBoost model. *International Journal of Computational Materials Science and Engineering*.