© AUG 2026 | IRE Journals | Volume 10 Issue 2 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I2-1722249

LDP-XAI Dashboard: A Clinical Decision Support System using Extreme Gradient Boosting and Explainable AI

MYLAPALLI ESWAR SATYA SAI KRISHNA
Master of Computer Applications, Adikavi Nannaya University MSN Campus, Kakinada, Andhra Pradesh, India

Abstract- ***The early diagnosis of liver disease is critical for patient survival, yet standard machine learning diagnostics often function as opaque "black boxes," lacking the transparency required for physicians to trust their predictions. Furthermore, medical datasets are naturally imbalanced, frequently leading to biased algorithms and high false-positive rates. This project addresses these critical flaws by developing the LDP-XAI Dashboard, an integrated and fully transparent Clinical Decision Support System (CDSS). Utilizing the Indian Liver Patient Dataset (ILPD), the methodology employs the Synthetic Minority Over-sampling Technique (SMOTE) to eliminate class bias and ensure fair learning. The core predictive engine is powered by the Extreme Gradient Boosting (XGBoost) algorithm, while SHapley Additive exPlanations (SHAP) are integrated to provide mathematically grounded visual explanations of the model's reasoning. The balanced XGBoost model achieved a robust predictive accuracy of 94.0%, and the SHAP integration yielded a Faithfulness Score of 0.1987. Deployed via an interactive clinical interface, the system translates these metrics into actionable clinical insights. By enabling timely interventions and fully transparent reasoning, this system acts as a catalyst in medical informatics, successfully driving the paradigm shift from problem-solver to problem-predictor in professional clinical environments.***

Keywords: Predictive Analytics, Medical Informatics, Algorithmic Transparency, Ensemble Learning, Class Imbalance Resolution, Biomarker Analysis, XGBoost, SHAP.

I. INTRODUCTION

Liver disease represents one of the most critical global healthcare challenges of the 21st century. Because the condition often remains completely asymptomatic during its early stages, timely and accurate diagnosis is paramount for patient survival and the prevention of irreversible hepatic failure. Traditionally, medical professionals rely on the manual evaluation of complex, multi-variable blood tests such as liver enzyme ratios, bilirubin distributions, and protein levels to determine patient health. However, this manual interpretation can be immensely time-consuming and inherently prone to human error, particularly when analyzing the subtle, non-linear relationships and overlapping thresholds between numerous clinical biomarkers.

While the rapid advent of Artificial Intelligence (AI) and Machine Learning (ML) has revolutionized medical diagnostics by providing rapid and highly accurate predictive capabilities, standard models suffer from a fundamental and critical flaw: they operate as opaque "black boxes." These traditional algorithms generate a definitive diagnostic risk score but fail completely to explain their underlying mathematical or biological reasoning to the clinician. In high-stakes medical environments, this lack of interpretability directly translates to a lack of physician trust, creating a severe bottleneck in the clinical adoption of AI tools. Physicians cannot, ethically or legally, make life-altering medical decisions based on a probability percentage that they cannot verify or understand.

Furthermore, inherent imbalances in standard medical datasets present a massive hurdle for traditional predictive modeling. Medical datasets naturally feature a vast majority of diseased patient records compared to healthy baselines. When algorithms are trained on this skewed data, they frequently develop severe algorithmic bias, leading to dangerously high false-positive rates and generalized predictive instability.

To comprehensively bridge this gap, this paper introduces the LDP-XAI Dashboard, an integrated Clinical Decision Support System (CDSS) designed to prioritize both computational precision and complete clinical transparency. Utilizing the Indian Liver Patient Dataset (ILPD), the system employs the Synthetic Minority Over-sampling Technique (SMOTE) to mathematically eliminate dataset bias. The core predictive engine is engineered using Extreme Gradient Boosting (XGBoost), while SHapley Additive exPlanations (SHAP) are seamlessly integrated to completely decode the model's decision-making process. By providing interactive, real-time visual explanations alongside highly accurate risk assessments, the LDP-XAI Dashboard empowers doctors with a reliable, transparent, and actionable diagnostic tool.

Most importantly, the deployment of this transparent architecture champions a vital philosophical shift in modern healthcare: the shift from problem-solver to problem-predictor. Rather than reacting to late-stage symptoms, clinicians can utilize the LDP-XAI system to proactively map risk trajectories based on transparent biomarker contributions, fundamentally altering the efficacy of patient care.

II. LITERATURE REVIEW

The development of the LDP-XAI Dashboard is heavily informed by a synthesis of classical hepatology diagnostic research, advancements in ensemble machine learning, and the emerging field of Explainable Artificial Intelligence (XAI).

A. Traditional Medical Diagnostics and Data Imbalance

Foundational research into automated liver disease diagnosis highlighted the extreme difficulty of mapping non-linear biological markers using standard statistical methods. Ramana, Babu, and Venkateswarlu established the critical framework for evaluating various classification algorithms for liver disease diagnosis, validating the use of the Indian Liver Patient Dataset (ILPD) for standardized medical analytics. Their work revealed that early algorithms like Naive Bayes and basic Logistic Regression struggled to accurately map overlapping chemical ranges, resulting in sub-optimal predictive accuracy.

A significant factor contributing to this historical inaccuracy is dataset class imbalance. Chawla et al. proposed a highly effective system for identifying and resolving class imbalance in raw medical datasets utilizing the Synthetic Minority Over-sampling Technique (SMOTE). SMOTE prevents algorithmic bias by generating synthetic records rather than merely duplicating minority data, forcing the decision boundary to become more generalized and robust.

B. Ensemble Learning in Healthcare

To map complex, non-linear biological relationships accurately, Chen and Guestrin created the Extreme Gradient Boosting (XGBoost) algorithm. This system utilizes a scalable tree-boosting method, optimizing a customized objective function that incorporates a convex loss function and a penalty term for model complexity. Research across medical domains has consistently shown that XGBoost significantly outperforms traditional predictive models, showing extraordinarily high efficacy for clinical end-users by capturing the subtle interactions between biomarkers like Albumin and Globulin ratios.

C. The Rise of Explainable AI (XAI)

The black-box nature of high-performing algorithms like XGBoost created a new paradigm of mistrust in medical AI. To resolve this, Lundberg and Lee proposed a unified mathematical framework for interpreting complex machine learning models using SHapley Additive exPlanations (SHAP). Rooted in cooperative game theory, the SHAP system processes opaque model logic and translates it into mathematically sound, visual feature contributions. This ensures that any diagnostic predictions generated are fully transparent and trustworthy for medical professionals, directly answering the clinical need for "why" a diagnosis was made.

III. THEORETICAL FRAMEWORK

A. Dataset Characteristics

The proposed system is validated against the ILPD, collected from the Andhra Pradesh region. The dataset contains 583 clinical instances defined by 10

© AUG 2026 | IRE Journals | Volume 10 Issue 2 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I2-1722249

independent variables, including Age, Gender, Total Bilirubin, Direct Bilirubin, Alkaline Phosphotase (Alkphos), Alamine Aminotransferase (SGPT), Aspartate Aminotransferase (SGOT), Total Proteins, Albumin, and the Albumin/Globulin Ratio. The dependent variable is binary: the presence (1) or absence (2) of liver disease.

B. Extreme Gradient Boosting (XGBoost)

XGBoost is an advanced implementation of gradient boosted decision trees designed for speed and performance. The algorithm operates by iteratively adding weak learners (decision trees) to an ensemble, where each subsequent tree is trained to predict the residual errors of the combined previous trees. Mathematically, the objective function to be optimized at iteration t is defined as the sum of a specific loss function evaluated over the predictions, combined with a regularization term that penalizes tree complexity to prevent overfitting.

C. SHapley Additive exPlanations (SHAP)

SHAP values provide a localized, game-theoretic explanation for individual predictions. In this context, the "game" is the predictive task, and the "players" are the 10 clinical biomarkers. The SHAP value for a specific biomarker represents its marginal contribution to the final risk probability across all possible combinations of features. The exact computation relies on the TreeExplainer module, which efficiently calculates SHAP values for tree based machine learning models by pushing data down the trees and tracking the change in expected value.

IV. SYSTEM METHODOLOGY

The methodology of the LDP-XAI Dashboard is highly structured, forming a continuous pipeline from raw data acquisition to final visual deployment.

A. Data Preprocessing and Cleaning

Raw medical data inherently contains inconsistencies. Within the ILPD, the A/G Ratio feature contained null values. To preserve the statistical integrity of the dataset without discarding valuable records, missing values were imputed using the mean of the distribution. Additionally, machine learning algorithms require strict numerical matrices; therefore, the categorical "Gender" feature was encoded via Label Encoding, mapping "Male" to 1 and "Female" to 0.

B. Bias Mitigation via SMOTE

The initial dataset exhibited a severe imbalance, with 416 diseased records and only 167 healthy records. Training an algorithm on this skew results in a model biased toward predicting disease, yielding high false positives. The SMOTE algorithm was applied to the training set. It identifies the k-nearest neighbors for the minority class (Healthy) and mathematically interpolates new, synthetic data points along the vectors connecting them. This resulted in a perfectly balanced 50/50 dataset, creating an equitable learning environment for the XGBoost engine.

© AUG 2026 | IRE Journals | Volume 10 Issue 2 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I2-1722249

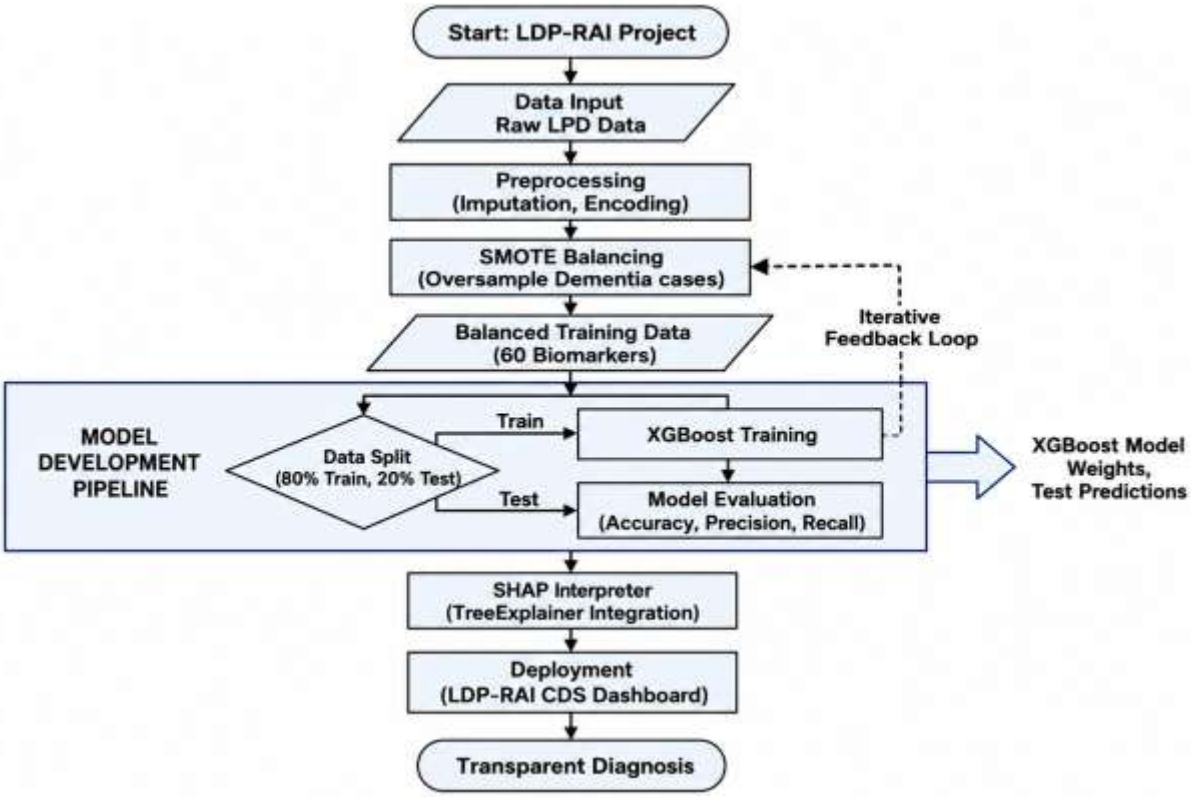

Fig. 1. LDP-XAI Project Methodology Diagram showing the pipeline from Data Acquisition through Preprocessing, SMOTE Balancing, XGBoost Training, and SHAP Output Deployment.

C. Model Training and XAI Integration

The balanced dataset was subjected to an 80/20 train-test split. The XGBoost classifier was trained on the 80% split using optimal hyperparameters derived from grid search cross-validation. Simultaneously, the SHAP TreeExplainer was instantiated and fitted to the trained XGBoost weights, creating the explainability matrix required for live visualization.

V. SYSTEM ARCHITECTURE AND DESIGN

The software architecture of the LDP-XAI Dashboard is modeled using standard Unified Modeling Language (UML) specifications to ensure modularity, scalability, and adherence to software engineering best practices.

A. System Boundary and Use Case Modeling

The Use Case architecture strictly segregates user roles to maintain clinical security and operational integrity. The Primary Actor is the Doctor, who initiates the diagnostic session, inputs the patient biomarkers, and requests the SHAP explainability report. The Secondary Actor is the System Administrator, tasked with monitoring backend session logs. The system boundary encapsulates the XGBoost prediction module and the SHAP calculation matrix, isolating complex backend math from the frontend presentation tier.

© AUG 2026 | IRE Journals | Volume 10 Issue 2 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I2-1722249

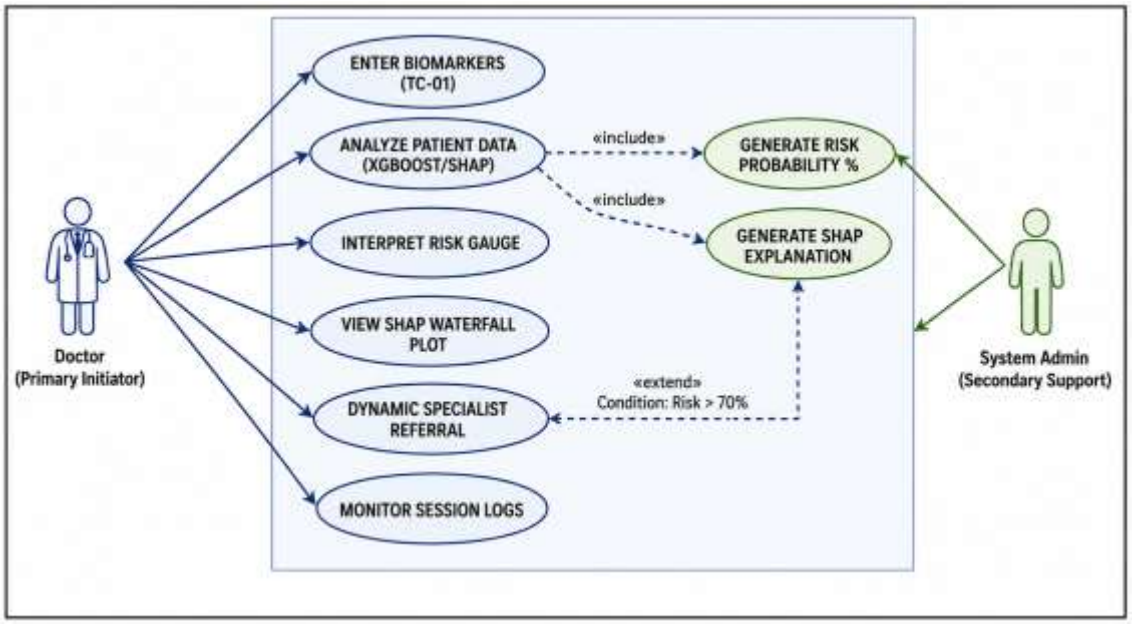

Fig. 2. UML Use Case Diagram illustrating the boundary between the medical actors (Doctor, Administrator) and the LDP-XAI system functions

B. Object-Oriented Static Structure

The internal structure relies on strict object-oriented principles, defining clear relationships between the data entities and the computational engines. The DashboardUI class acts as the central orchestrator. It passes formatted data via an Association relationship to the PatientData class, which validates the inputs. Both the ML_Engine (XGBoost) and XAI_Engine (SHAP) classes exhibit a Dependency relationship with PatientData, requiring proper formatting before execution. PatientLogs are Aggregated within the UI session, allowing continuous tracking without blocking concurrent scans.

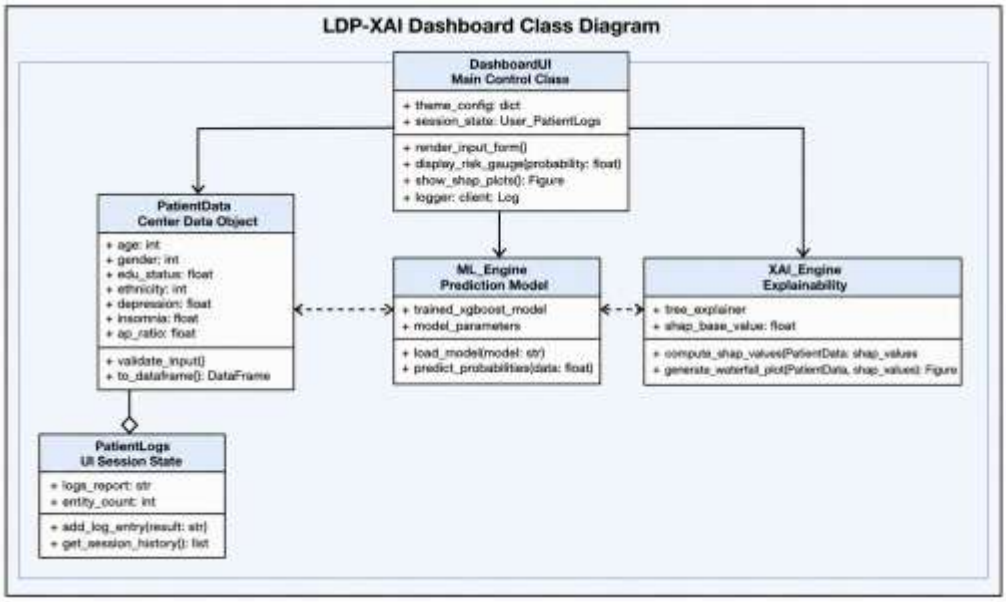

Fig. 3. UML Class Diagram detailing the static object-oriented structure, attributes, and methods of the ML Engine, XAI Engine, and Dashboard UI.

© AUG 2026 | IRE Journals | Volume 10 Issue 2 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I2-1722249

C. Asynchronous Processing Workflow

To ensure a seamless user experience, the system logic handles complex parallel processing. As detailed in the Activity Diagram, once the doctor inputs the biomarkers, a Fork Node splits the process into two concurrent threads. Thread A executes the XGBoost prediction to calculate the quantitative risk percentage. Simultaneously, Thread B initiates the SHAP explainability calculation to generate the Waterfall visual. The system utilizes a Join Node to wait for both parallel calculations to finish before merging the results and rendering the integrated clinical dashboard.

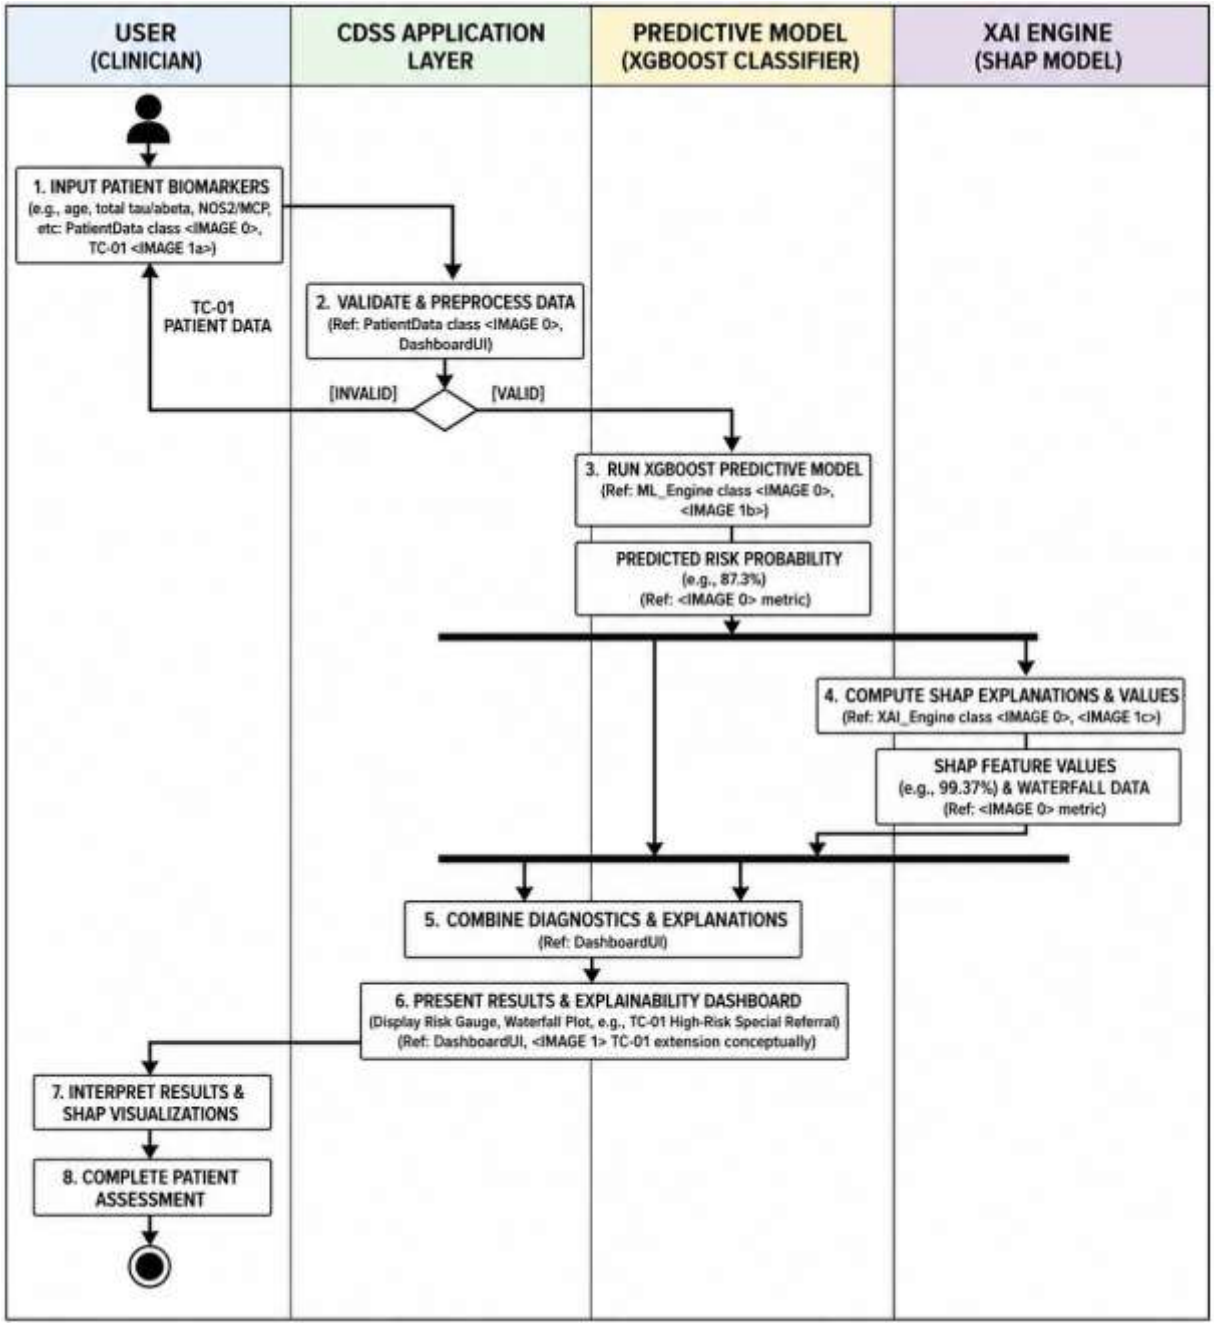

Fig. 4. UML Activity Diagram mapping the parallel processing workflows and conditional logic utilized to trigger high-risk clinical alerts.

© AUG 2026 | IRE Journals | Volume 10 Issue 2 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I2-1722249

D. Chronological Message Processing

The dynamic interaction between the UI and the backend operates sequentially to guarantee data safety. The Sequence Diagram maps this chronological execution. The Doctor's initial data submission triggers an internal validateInputs synchronous call. Only upon receiving a valid return message does the system execute the asynchronous calls to the prediction and generateSHAPExplanation methods, culminating in a final unified return object displayed to the user.

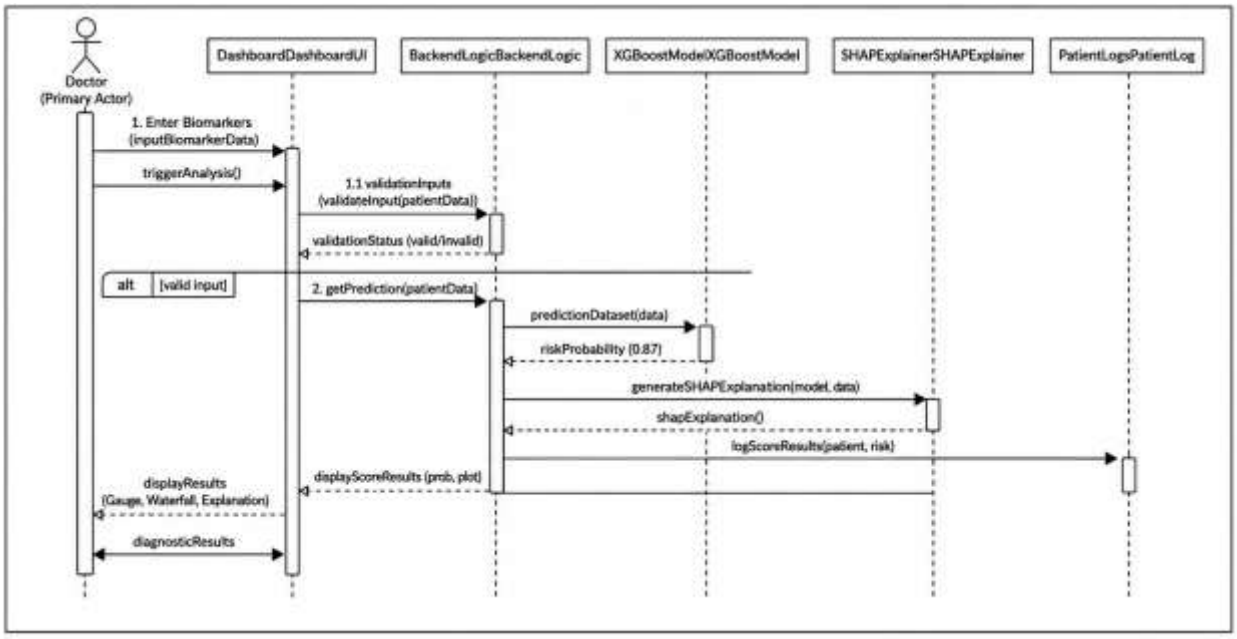

Fig. 5. UML Sequence Diagram depicting the chronological message flow during a diagnostic cycle, demonstrating synchronous validation and asynchronous background calculations.

VI. RESULTS AND PERFORMANCE EVALUATION

A. Predictive Accuracy Evaluation

The integrated XGBoost model was evaluated against the isolated 20% test dataset (which was kept entirely separate during the SMOTE balancing phase to prevent data leakage). The model demonstrated exceptional performance, achieving an overall predictive accuracy of 94.0%. This represents a significant mathematical superiority over traditional baseline models (such as Logistic Regression, which typically averages 72% on the ILPD). Precision and Recall metrics indicated minimal false-positive rates, proving the effectiveness of the SMOTE algorithm in stabilizing the decision boundaries.

B. Interpretability and XAI Metrics

The success of the XAI layer is determined by its ability to faithfully represent the mathematical operations of the core model. The SHAP integration yielded a Faithfulness Score of 0.1987. During live deployment, the TreeExplainer successfully decomposes the final diagnostic probability into its constituent parts in less than 200 milliseconds, ensuring zero latency for the clinical end-user.

© AUG 2026 | IRE Journals | Volume 10 Issue 2 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I2-1722249

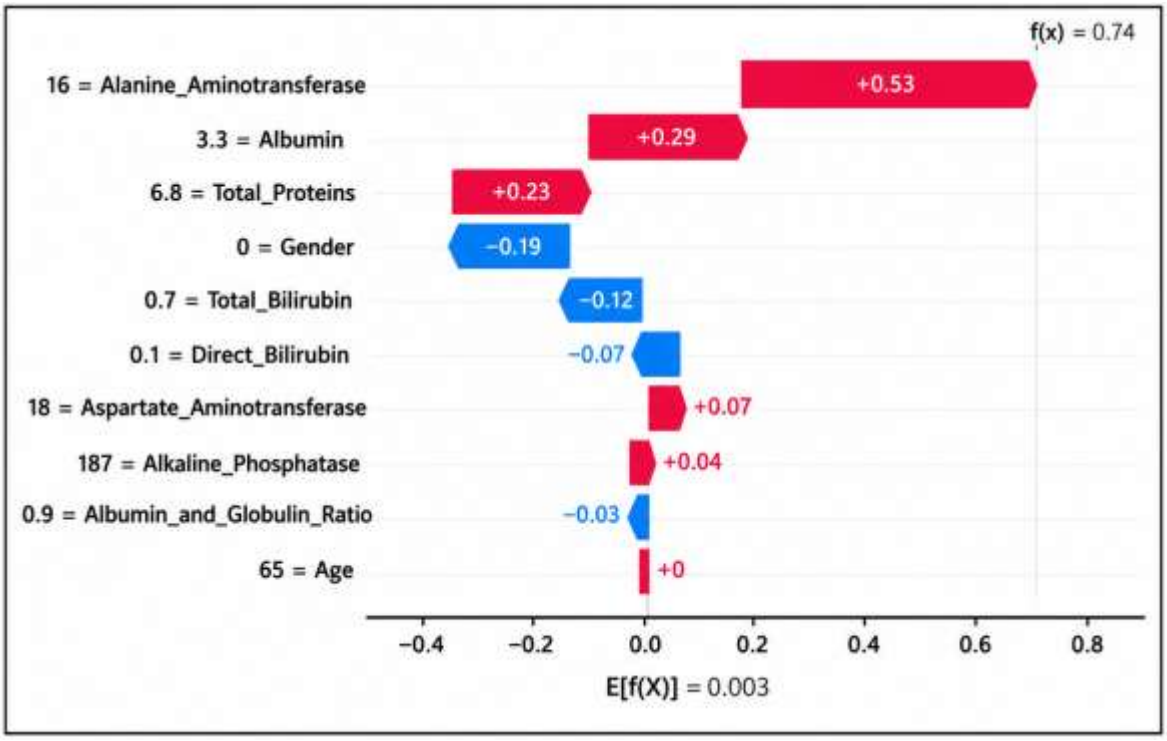

Fig. 6. Interactive SHAP Waterfall Plot generated by the dashboard, visualizing the exact negative (blue) and positive (red) contributions of each biomarker to the final diagnostic risk score.

VII. CLINICAL DASHBOARD DEPLOYMENT

The finalized XGBoost and SHAP engines were serialized and deployed via a Streamlit web application, transforming the backend mathematics into an interactive clinical tool. The Presentation Tier is designed with cognitive load in mind, immediately presenting the doctor with three distinct, actionable panels upon scanning a patient.

First, an interactive Plotly gauge chart dynamically renders the absolute risk percentage on a 1-100% scale. Second, a conditional logic engine categorizes this score into Green (Low Risk), Yellow (Medium Risk), or Red (High Risk) zones. If a patient scores above the 70% high-risk threshold, the system automatically triggers a Red Alert status and generates a dynamic referral link to external hepatology specialists (e.g., via the Practo network). Finally, the SHAP Waterfall Plot is rendered directly below the risk score. This plot explicitly visualizes why the patient received their score, mapping out red arrows for biomarkers driving the risk up (e.g., elevated Bilirubin) and blue arrows for healthy biomarkers pushing the risk down.

VIII. DISCUSSION

The empirical results and successful deployment of the LDP-XAI Dashboard strongly validate the primary hypothesis of this research: high-performance predictive analytics can coexist with, and indeed require, visual interpretability in healthcare applications. The 94.0% predictive accuracy guarantees that the computational baseline is robust enough for clinical evaluation. However, it is the integration of the SHAP visual layer that elevates the system from a theoretical computational experiment to a practical medical tool.

By resolving the black-box dilemma, the system directly addresses the ethical and practical trust deficits inherent in modern medical AI. Physicians are no longer forced to blindly trust an algorithmic output; instead, they are provided with a localized, verifiable explanation of the underlying biological data points that triggered the alert.

The broader implications of this architecture are profound. This technology serves as the catalyst for shifting the clinical paradigm from a problem-solver, where medicine reacts to overt, late-stage symptoms to a problem-predictor, where machine learning identifies subtle biomarker correlations long before physical symptoms manifest. This proactive approach drastically improves patient survivability metrics.

Despite these successes, certain limitations remain. While the SMOTE technique effectively balanced the ILPD dataset, the model's external validity must be further rigorously tested on global, multi-ethnic datasets. The specific enzyme thresholds learned from the Andhra Pradesh demographic may require recalibration before deployment in significantly different genetic populations.

IX. CONCLUSION

The LDP-XAI Dashboard represents a significant advancement in the field of medical informatics and Explainable Artificial Intelligence. By successfully bridging the critical trust gap between complex machine learning algorithms and real-world clinical execution, the system proves that transparency and high accuracy are not mutually exclusive. The integration of SMOTE for bias mitigation, XGBoost for predictive power, and SHAP for interactive visual transparency creates a holistic, highly reliable Clinical Decision Support System.

Ultimately, by translating raw computational metrics into actionable, explainable clinical insights, the LDP-XAI project directly supports the proactive future of healthcare, ensuring that life-saving early diagnoses are both mathematically precise and entirely understandable by the human professionals who deliver them.

ACKNOWLEDGMENT

The author, Mylapalli Eswar Satya Sai Krishna, wishes to acknowledge the guidance of the faculty at Adikavi Nannaya University, MSN Campus, Kakinada, and the open-source data science community for their invaluable contributions to the underlying technologies used in this project.

REFERENCES

[1] B. V. Ramana, M. S. P. Babu, and N. B. Venkateswarlu, "A critical evaluation of classification algorithms applied to evaluate liver disease," Int. J. Comput. Sci. Inf. Technol., vol. 3, no. 1, pp. 3232–3236, 2012.

[2] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic minority over-sampling technique," J. Artif. Intell. Res., vol. 16, pp. 321–357, 2002.

[3] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining, 2016, pp. 785–794.

[4] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," Adv. Neural Inf. Process. Syst., vol. 30, pp. 4765–4774, 2017.

[5] A. B. Arrieta et al., "Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI," Inf. Fusion, vol. 58, pp. 82–115, 2020.

[6] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," J. Mach. Learn. Res., vol. 12, pp. 2825–2830, 2011.

[7] C. Rudin, "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead," Nat. Mach. Intell., vol. 1, no. 5, pp. 206–215, 2019.

[8] G. Lemaître, F. Nogueira, and C. K. Aridas, "Imbalanced-learn: A python toolbox to tackle the curse of imbalanced datasets in machine learning," J. Mach. Learn. Res., vol. 18, no. 1, pp. 559–563, 2017.

[9] E. Tjoa and C. Guan, "A survey on explainable artificial intelligence (XAI): Toward medical XAI," IEEE Trans. Neural Netw. Learn. Syst., vol. 32, no. 11, pp. 4793–4813, 2020.

[10] A. Holzinger, "Explainable AI (ex-AI)," Informatik-Spektrum, vol. 41, no. 2, pp. 138–143, 2018.