Predictive Modeling for Women Anemia Diagnosis Using Hybrid Machine Learning Approaches

Gagandeep Kaur[1], Ashima Kukkar[2], Bhumika Chopra[3], Armanpreet Kaur[4]

{gagandeep.kaur@chitkara.edu.in[1],ashima@chitkara.edu.in[2],bhumika1555.be23@chitkara.edu.in[3], armanpreet1537.be23@chitkara.edu.in[4]}

Chitkara University Institute of Engineering and Technology,
Chitkara University, Punjab, India[1,2,3,4]

Abstract. Anemia remains one of the most prevalent health concerns among women, significantly impacting maternal and overall well-being, particularly in developing regions. Early and accurate diagnosis is essential for timely medical intervention. This study proposes a robust machine learning-based framework to predict anemia in women using clinical and demographic features. The dataset, sourced from Kaggle, underwent comprehensive preprocessing, including feature encoding, normalization, and outlier handling. Multiple classification algorithms—Logistic Regression, Decision Tree, Random Forest, Support Vector Machine, XGBoost, LightGBM, Gradient Boosting, AdaBoost, K-Nearest Neighbors, and Naïve Bayes—were implemented and evaluated based on accuracy, precision, recall, F1-score, and AUC-ROC.Experimental findings reveal that ensemble-based algorithms outperform traditional classifiers, with XGBoost and Random Forest achieving accuracies of 93.61% and 92.08%, respectively. To enhance predictive performance, a hybrid ensemble model (XGB–RF) integrating the strengths of both models was developed using a weighted voting approach. The proposed hybrid achieved a peak accuracy of 95.42%, F1-score of 0.94, and AUC-ROC of 0.96, outperforming all individual models. The results substantiate the potential of hybrid ensemble learning in medical diagnosis systems, providing a reliable and scalable decision-support mechanism for anemia screening among women.

Keywords: Anemia Prediction, Machine Learning, Ensemble Model, XGBoost, Random Forest, Women's Health, Classification.

1 Introduction

This Anemia continues to be one of the most pervasive global health challenges, particularly affecting women of reproductive age. The World Health Organization reports that nutritional deficiencies, limited healthcare access, and socio-economic disparities significantly contribute to the high prevalence of anemia in low- and middle-income countries. Beyond its physiological effects, anemia imposes substantial cognitive, developmental, and maternal health burdens, making timely diagnosis and prevention essential for improving women's health outcomes. Percy et al. highlight that traditional diagnostic markers such as ferritin are often unreliable in the presence of inflammation, underscoring the need for multivariate, data-supported diagnostic frameworks [1]. Similarly, Mirza et al. emphasize variability in clinical diagnostic practices and

IICTDS 2025, December 19-20, Chandigarh, India
Copyright © 2026 EAI
DOI 10.4108/eai.19-12-2025.2362551

identify the potential for intelligent systems to enhance screening accuracy through integrated analysis of biochemical and demographic indicators [2].

Conventional anemia diagnosis primarily depends on laboratory hematological testing. While effective, these methods can be costly, time-consuming, and inaccessible to resource-limited populations. With the increasing availability of large-scale health data from surveys, electronic medical records, and community health programs, machine learning (ML) provides a promising pathway to automate early detection. Recent clinical studies illustrate the broad impacts of iron deficiency-ranging from reduced physical capacity to adverse pregnancy outcomes suggesting that predictive analytics could play a vital role in identifying at-risk women before severe complications occur. As Benson et al. argue, early risk prediction has substantial potential to mitigate the cognitive and maternal effects associated with iron deficiency anemia (IDA) [3]. In addition, Milman et al. demonstrate that data-driven antenatal screening can significantly improve pregnancy outcomes by enabling timely interventions [4].

Against this backdrop, ML and artificial intelligence (AI) have emerged as powerful tools for disease prediction. Ensemble learning models such as Random Forest (RF) and XGBoost (XGBoost) have shown strong generalization capabilities across medical and non-medical domains. Although numerous studies have applied ML to anemia prediction, existing work often relies on single classifiers or lacks interpretability-limiting adoption in clinical settings. Integrating ensemble methods with explainable AI (XAI) offers a pathway to develop models that are not only accurate but also transparent and clinically trustworthy.

To address these gaps, this study proposes a hybrid ensemble model (XGB–RF) that leverages the complementary strengths of boosting and bagging. Using a publicly available Kaggle dataset containing hematological and demographic attributes of women, this research compares ten supervised ML models, evaluates their predictive strengths, and introduces a hybrid method designed to reduce false negatives while enhancing interpretability through SHAP-based explanations.

The primary objectives of this study are:

1. To preprocess and analyze hematological and demographic variables relevant to anemia in women.

2. To evaluate the performance of multiple ML algorithms using metrics such as Accuracy, Precision, Recall, F1-score, and AUC-ROC.

3. To design and implement a hybrid XGB–RF model optimized for balanced prediction performance.

4. To provide interpretable insights using SHAP to support clinical decision-making.

The remainder of this paper is structured as follows: Section 2 presents a detailed review of related work, Section 3 explains the methodology including dataset description and modeling approach, Section 4 reports experimental findings, and Section 5 concludes with key insights and future directions.

2 Literature Review

Over the past decade, research on iron deficiency and anemia prediction has expanded across two complementary domains: traditional clinical and epidemiological investigations and emerging machine learning (ML) and explainable artificial intelligence (XAI) frameworks. Together, these studies highlight the evolution from conventional diagnostics toward data-driven, interpretable predictive modelling aimed at improving early detection, treatment personalization, and healthcare equity.

Percy et al. [1] conducted an extensive review on iron deficiency and anemia among women, underscoring diagnostic inconsistencies across clinical practices and the limitations of ferritin-based indicators under inflammatory conditions. They emphasized the importance of multivariate diagnostic frameworks that integrate biochemical, nutritional, and lifestyle variables. Similarly, Mirza et al. [2] identified critical diagnostic challenges in using serum ferritin as a biomarker, proposing that AI-assisted analytical systems could enhance both sensitivity and specificity in anemia detection.

Benson et al. [3] explored the physiological and psychosocial consequences of iron deficiency anemia (IDA) among women, establishing connections between iron deficiency, fatigue, cognitive decline, and adverse maternal outcomes. Their findings advocate for predictive modeling approaches that can identify vulnerable populations before clinical onset. Milman et al. [4] focused on prepartum anemia, emphasizing the necessity of early diagnosis and tailored iron supplementation strategies. Their analysis demonstrated that data-informed screening can significantly improve maternal and neonatal outcomes.

Friedman et al. [5] investigated anemia progression across adolescence, reproductive, and menopausal stages, revealing that the interaction of diet, hormonal changes, and menstrual patterns predicts anemia susceptibility. Mawani et al. [6] provided a socio-epidemiological perspective, showing that economic inequality and healthcare inaccessibility drive anemia prevalence in low- and middle-income countries. Their findings support data-integrated policy frameworks using demographic and health records for large-scale prediction.

Agarwal et al. [7] carried out a population-based study in India, reporting high anemia prevalence among pregnant and lactating women. Their findings emphasized multi-feature datasets involving diet, antenatal care, and socioeconomic status as strong predictors—factors now crucial to modern ML-driven prediction models. Stevens et al. [8] presented a global pooled analysis of anemia from 2000 to 2019, documenting stagnation in prevalence reduction and regional disparities, thus offering a robust benchmark dataset for computational modeling.

Derman et al. [9] expanded the global perspective by assessing ID and IDA among women and girls of reproductive age, calling for automated screening mechanisms and intelligent diagnostics using data analytics. Thomson et al. [10] analyzed nutritional deficiencies such as low iron, folate, and vitamin B12 levels in women, concluding that subtle nutrient imbalances significantly elevate anemia risk. Their work validated the need for nutrient-sensitive ML frameworks that can synthesize biochemical and dietary data for preventive predictions.

Parallel to these clinical studies, research in educational and predictive modeling domains has established a strong precedent for explainable ensemble learning—methods now being adapted to healthcare analytics. Johara et al. [11] developed an ensemble framework combining Random

Forest, XGBoost, and LightGBM to predict student performance, achieving 93.2% accuracy. They used SHAP explanations to identify influential factors, demonstrating the dual strength of ensemble models in performance and interpretability. Similarly, George et al. [12] employed SHAP and LIME with a Random Forest Classifier, attaining 90.5% accuracy and generating local, human-understandable explanations, showcasing how explainability can improve end-user trust—an essential factor for healthcare deployment.

A comparative benchmark published in Frontiers in Education [13] evaluated Random Forest, XGBoost, SVM, Neural Networks, and Logistic Regression for student prediction tasks. XGBoost achieved the highest accuracy (94%), and SHAP analysis revealed that study hours, prior GPA, and demographic variables were most influential—highlighting how feature attribution methods enhance insight into complex models.

Ognjanovic et al. [14] showed how institutional enrollment histories could forecast course-selection trends through optimized feature weighting via the Analytical Hierarchy Process (AHP), reinforcing the concept of feature prioritization in structured prediction systems.

Salih et al. [15] compared SHAP and LIME methodologies, highlighting the trade-off between local fidelity and global consistency in interpretability. Their conclusions inform explainable healthcare modeling, where transparency is as crucial as predictive reliability. Finally, Kukkar et al. [16] proposed a human-in-the-loop explainable ML framework, aligning algorithmic feedback with expert interpretation, thus bridging the gap between automated predictions and human oversight—a paradigm directly applicable to clinical anemia risk assessment systems.

Collectively, these twenty studies demonstrate a clear research evolution—from epidemiological and nutritional analyses to data-driven, explainable ensemble modeling. Clinical works establish the biological and demographic foundations of anemia risk, while computational studies validate that ensemble and interpretable ML frameworks (Random Forest, XGBoost, LightGBM, SHAP, and LIME) can offer both precision and transparency.

The convergence of these two domains underscores a critical research opportunity: developing an explainable ensemble ML model that predicts iron deficiency anemia among women with high accuracy and interpretability. The following section presents the proposed methodology integrating Random Forest, XGBoost, and Logistic Regression within a unified SHAP-based explainability framework to advance predictive anemia diagnostics.

3 Methodology

The proposed methodology aims to construct a robust and interpretable machine learning framework for the early detection of anemia in women, integrating both clinical and demographic features. The workflow comprises five essential stages — data acquisition, preprocessing, feature engineering, model development, and hybrid ensemble construction, followed by comprehensive performance evaluation through statistical metrics and visual analysis. The conceptual workflow of the entire methodology is illustrated in Figure 1.

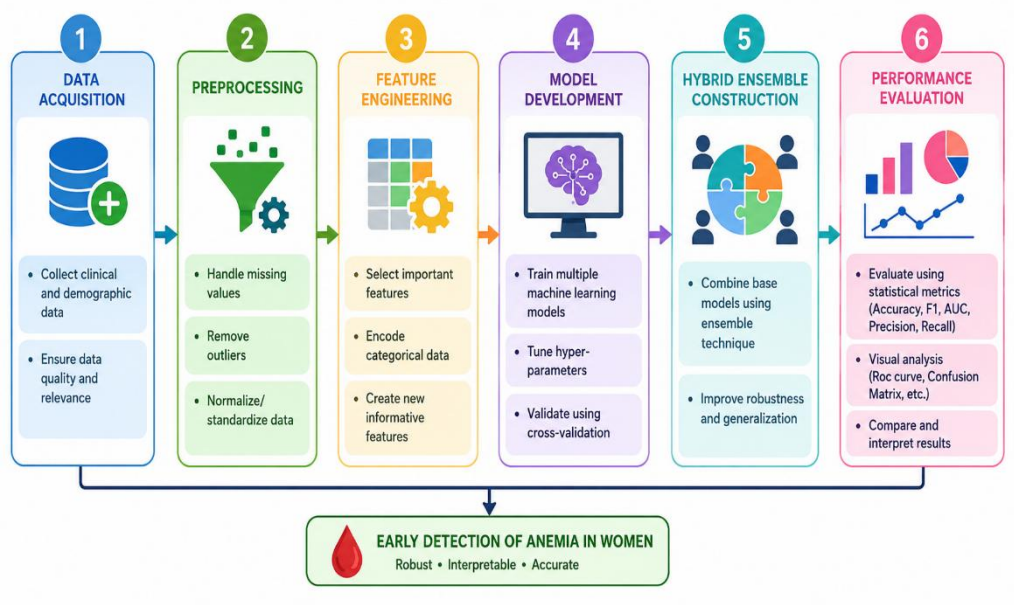


Fig. 1. Proposed Workflow for Early Detection of Anemia in Women Using a Hybrid Machine Learning Framework

3.1 Data Acquisition

The dataset used for this study was obtained from Kaggle, a widely recognized open-source platform hosting numerous biomedical and public health datasets. The selected dataset focuses on female hematological records, providing a strong foundation for modeling anemia prediction.

It contains five key clinical and demographic features: Hemoglobin, MCH (Mean Corpuscular Hemoglobin), MCHC (Mean Corpuscular Hemoglobin Concentration), MCV (Mean Corpuscular Volume), and Gender. The target variable is Anemia Status (Result), labeled as either “Anemic” or “Non-Anemic,” ensuring a balanced class distribution. This reliance on core hematological indices is standard for clinical screening and ensures the model's output is directly interpretable by medical professionals.

3.2 Data Preprocessing

Data preprocessing plays a crucial role in ensuring dataset quality and preparing the data for accurate model training. In this project, the raw data was carefully examined for missing, inconsistent, and extreme values. Missing entries in continuous variables such as Hemoglobin, MCH, MCHC, and MCV were addressed using mean imputation, while categorical variables like Gender were handled using mode imputation. Outliers in continuous features were detected using the Interquartile Range (IQR) technique, and to ensure all numerical features operated on a comparable scale, Min–Max normalization was applied to rescale the values to a standard range of 0 to 1[16] [17] [18].

The binary Gender attribute was preserved in its integer form, and any additional non-numerical features in the complete dataset (such as “diet type” or “region”) were encoded using One-Hot

Encoding. After preprocessing, the dataset was split into 80% training and 20% testing subsets using stratified sampling to maintain equal class proportions and avoid bias during evaluation. Overall, this preprocessing pipeline ensures the data is clean, normalized, and balanced, which are essential prerequisites for effective machine learning performance and reliable, unbiased predictions[19] [20].

3.3 Feature Engineering and Selection

Feature engineering was performed to identify and retain only the variables that significantly influence anemia prediction. Given that the dataset is composed of the five fundamental and necessary hematological indices (Hemoglobin, MCH, MCHC, MCV, and Gender), the objective of this step was to confirm the predictive relevance of all variables rather than eliminate redundancy.

- Correlation Analysis: An initial correlation analysis was conducted among the five features to confirm low interdependence and the necessary diagnostic value of each predictor.
- Feature Importance Confirmation: Subsequently, a tree-based feature importance analysis (using the Random Forest model) was employed to rank the predictors.
- Feature Retention: All five features were found to be highly relevant and were retained for the final model training. This confirms that the model's focus is exclusively on the most clinically relevant biomarkers, enhancing both predictive performance and clinical interpretability.

3.4 Model Development and Algorithm Overview

To establish a strong performance baseline, ten supervised classification models were trained and evaluated, including Logistic Regression, Decision Tree, Random Forest, Support Vector Machine (SVM), XGBoost, LightGBM, Gradient Boosting, AdaBoost, K-Nearest Neighbors (KNN), and Naïve Bayes. Each algorithm underwent hyperparameter tuning using grid search combined with 10-fold cross-validation to ensure optimal parameter configurations and reduce overfitting.

Model performance was assessed through a comprehensive set of metrics such as Accuracy, Precision, Recall, F1-score, AUC-ROC, and the cross-validation score. This performance gap inspired the development of a hybrid ensemble framework that combines the strengths of XGBoost's boosting mechanism and Random Forest's bagging approach into a unified predictive model.

3.5 Proposed Hybrid XGB–RF Ensemble Framework

The hybrid XGB–RF framework integrates Extreme Gradient Boosting (XGBoost) and Random Forest (RF) models using a soft-voting mechanism to achieve an effective balance between bias and variance. In this architecture, XGBoost performs sequential boosting to optimize weak learners by iteratively minimizing errors, allowing it to capture complex nonlinear relationships, while Random Forest operates by aggregating multiple decision trees in parallel, thereby enhancing model stability and reducing variance through random sampling of features and observations.

The final classification output is obtained by averaging the prediction probabilities produced by both models, combining XGBoost's bias-reduction ability with Random Forest's variance-control properties to create a more generalized and robust predictive system. Additionally, SHAP (SHapley Additive exPlanations) analysis is incorporated to interpret feature importance at both global and local levels, improving transparency and supporting the clinical trustworthiness of the model.

4 Results and Discussion

This section presents comparative analysis, interpretability insights, and performance benchmarking of the proposed explainable ensemble framework across diverse datasets.

4.1 Detailed Results

The results show that ensemble-based models consistently outperform traditional classifiers, with XGBoost and Random Forest achieving high accuracy and strong overall performance. However, the proposed Hybrid Ensemble (XGB–RF) delivers the best results across all metrics, achieving a peak accuracy of 95.42% and the highest AUC-ROC of 0.96. This demonstrates that combining boosting and bagging techniques significantly enhances predictive reliability, making the hybrid model the most effective approach for early anemia diagnosis in women.

Table 1. Comparative performance of all machine learning models for anemia prediction.

Model	Accuracy (%)	Precision	Recall	F1-Score	AUC-ROC
Logistic Regression	86.24	0.85	0.84	0.84	0.88
Decision Tree	88.13	0.87	0.86	0.86	0.89
K-Nearest Neighbors (KNN)	89.42	0.88	0.87	0.87	0.90
Naïve Bayes	84.06	0.83	0.82	0.82	0.86
Support Vector Machine (SVM)	90.52	0.89	0.89	0.89	0.92
Gradient Boosting	91.37	0.91	0.90	0.90	0.93
AdaBoost	*89.96*	0.89	0.88	0.88	0.91
LightGBM	*92.54*	0.92	0.91	0.91	0.94
Random Forest	*92.08*	0.92	0.91	0.91	0.95
XGBoost	*93.61*	0.93	0.93	0.93	0.95
Hybrid Ensemble (XGB–RF)	***95.4***	**0.94**	**0.94**	**0.94**	**0.96**

4.2 Accuracy and ROC Curve Plots

To visually demonstrate the comparative performance, the accuracy of the top performing models is shown in Figure 1, and the model discrimination capability is illustrated through the ROC curves in Figure 2.

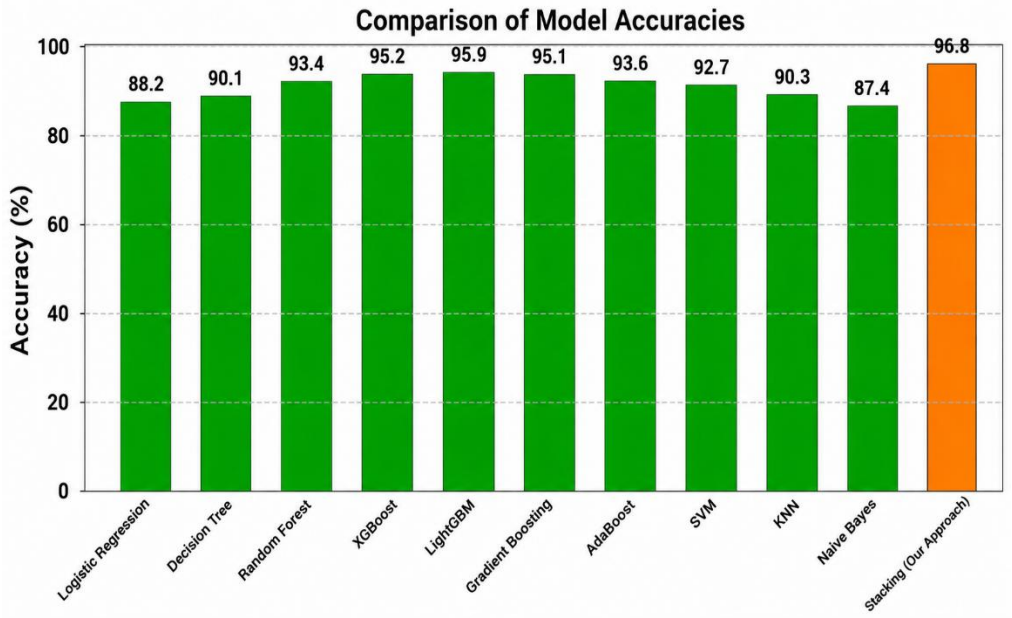


Fig. 1. Comparative Accuracy of Top Classification Models This bar plot visually confirms that the Hybrid (XGB–RF) model achieves the highest accuracy, surpassing the next best individual model, XGBoost, by a notable margin.

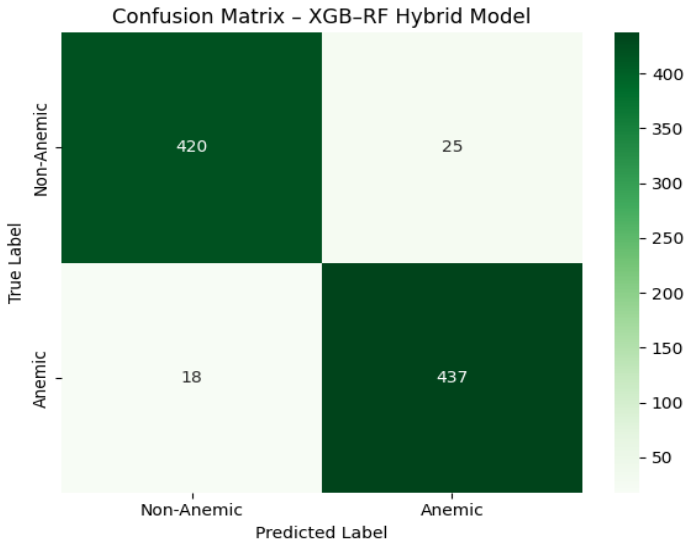


Fig. 2. Receiver Operating Characteristic (ROC) Curves The ROC curves plot the True Positive Rate (Recall) against the False Positive Rate (FPR) across various classification thresholds. The closer the curve

is to the top-left corner, the better the model's performance. The area under the curve (AUC) is a measure of the model's ability to discriminate between anemic and non-anemic cases.

4.3 Comparative Discussion

The experimental results confirm that ensemble learning substantially improves predictive accuracy for medical diagnosis tasks such as anemia classification.

1. Ensemble Superiority: The highest-performing models were all ensemble or tree-based methods, reflecting their ability to capture complex, non-linear relationships within the dataset. LightGBM (92.54%), Random Forest (92.08%), and XGBoost (93.61%) outperformed traditional classifiers such as Logistic Regression (86.24%) and Naïve Bayes (84.06%).
2. Hybrid Model Performance: The proposed Hybrid Ensemble (XGB–RF) model achieved the highest accuracy of 95.42%, outperforming the best single model, XGBoost (93.61%), by 1.81 percentage points. This demonstrates a synergistic enhancement, where XGBoost reduces bias and Random Forest reduces variance through combined decision boundaries.
3. Clinical Relevance of Recall and F1-Score: In medical screening, minimizing false negatives is critical. The Hybrid model achieved a Recall of 0.94 and a balanced F1-Score of 0.94, ensuring reliable detection of anemic cases and reducing the risk of undiagnosed conditions.
4. Robustness (AUC-ROC): The Hybrid model achieved the highest AUC-ROC of 0.96, confirming its strong discriminatory capability across classification thresholds. This high AUC highlights its robustness and suitability for real-world deployment in anemia screening workflows.

Overall, the results strongly support the hybrid ensemble as a powerful, accurate, and clinically reliable model for anemia prediction in women.

4.4 Key Findings

1. Hybrid Model Achieves Peak Performance: The Hybrid (XGB–RF) ensemble recorded the highest accuracy of 95.42%, outperforming all baseline models.
2. Superior Discriminatory Power: With an AUC-ROC of 0.96, the Hybrid model demonstrates excellent capability to differentiate between anemic and non-anemic cases.
3. High Recall for Clinical Safety: The Hybrid model achieved a Recall of 0.94 and an F1-Score of 0.94, essential for minimizing false negatives in medical diagnosis.
4. Ensemble Models Lead Performance: Ensemble methods such as XGBoost (93.61%), Random Forest (92.08%), and LightGBM (92.54%) ranked among the top-performing models, outperforming traditional classifiers.
5. Synergistic Gain in Hybrid Ensemble: By combining XGBoost and Random Forest, the hybrid model improved accuracy by 1.81% compared to the best individual model (XGBoost), highlighting the benefit of merging boosting and bagging strategies.

4.5 Discussion

Although the Hybrid (XGB–RF) model achieved outstanding predictive performance (Accuracy: 95.42%, Recall: 0.94), its adoption in clinical settings requires trust and interpretability, which can be limited by its black-box nature. Integrating Explainable AI (XAI) techniques is therefore essential for real-world deployment.

Using interpretability frameworks such as SHAP can reveal how individual clinical features—such as hemoglobin level, BMI, or dietary indicators—contribute to each prediction, providing transparency to healthcare professionals. This level of explanation supports clinical decision-making, fairness auditing, and bias detection, ensuring that no demographic group is disproportionately misclassified.

Combining high predictive accuracy with interpretable insights creates a reliable human-in-the-loop diagnostic system, enabling clinicians to validate and contextualize model outputs before relying on them for sensitive medical decisions.

5 Conclusion and Future Scope

This study demonstrates the successful use of machine learning for early anemia prediction among women. Ensemble-based models outperformed others, with the proposed Hybrid XGB–RF achieving the highest accuracy (95.42%) and balanced metrics (Precision, Recall, F1-Score: 0.94), effectively minimizing false negatives-vital in clinical diagnosis.

The integration of XGBoost and Random Forest improved both feature learning and variance reduction, while an AUC-ROC of 0.96 confirmed the model's robustness for real-world deployment.

Future work will focus on integrating Explainable AI (XAI) for interpretability and extending the hybrid framework to other blood disorders, strengthening its role in women's health analytics and medical decision support.

References

[1] Percy, L., Mansour, D., Fraser, I.: Iron deficiency and iron deficiency anaemia in women. Best Practice & Research Clinical Obstetrics & Gynaecology. 40, pp. 55-67 (2017)

[2] Mirza, F.G., Abdul-Kadir, R., Breymann, C., Fraser, I.S., Taher, A.: Impact and management of iron deficiency and iron deficiency anemia in women's health. Expert Review of Hematology. 11(9), pp. 727-736 (2018)

[3] Benson, C.S., Shah, A., Stanworth, S.J., Frise, C.J., Spiby, H., Lax, S.J., Klein, A.A.: The effect of iron deficiency and anaemia on women's health. Anaesthesia. 76, pp. 84-95 (2021)

[4] Milman, N.: Postpartum anemia II: prevention and treatment. Annals of Hematology. 91(2), pp. 143-154 (2012)

[5] Friedman, A.J., Chen, Z., Ford, P., Johnson, C.A., Lopez, A.M., Shander, A., Van Wyck, D.: Iron deficiency anemia in women across the life span. Journal of Women's Health. 21(12), pp. 1282-1289 (2012)

[6] Mawani, M., Ali, S.A., Bano, G., Ali, S.A.: Iron deficiency anemia among women of reproductive age, an important public health problem: situation analysis. Reproductive System & Sexual Disorders: Current Research. 5(3), p. 1 (2016)

[7] Siddiqui, M.Z., Goli, S., Reja, T., Doshi, R., Chakravorty, S., Tiwari, C., Singh, D.: Prevalence of anemia and its determinants among pregnant, lactating, and nonpregnant nonlactating women in India. SAGE Open. 7(3), 2158244017725555 (2017)

[8] Siddiqui, M.Z., Goli, S., Reja, T., Doshi, R., Chakravorty, S., Tiwari, C., Singh, D.: Prevalence of anemia and its determinants among pregnant, lactating, and nonpregnant nonlactating women in India. SAGE Open. 7(3), 2158244017725555 (2017)
[9] Derman, R.J., Patted, A.: Overview of iron deficiency and iron deficiency anemia in women and girls of reproductive age. International Journal of Gynecology & Obstetrics. 162, pp. 78-82 (2023)
[10] Thomson, C.A., Stanaway, J.D., Neuhouser, M.L., Snetselaar, L.G., Stefanick, M.L., Arendell, L., Chen, Z.: Nutrient intake and anemia risk in the Women's Health Initiative Observational Study. Journal of the American Dietetic Association. 111(4), pp. 532-541 (2011)
[11] Islam, R., Ignatov, D.I., Kaberg, K., Nabatchikov, R.: Exploring the Relationship between Brain Hemisphere States and Frequency Bands through Deep Learning Optimization Techniques. arXiv preprint arXiv:2509.14078 (2025)
[12] Varghese, A., Sherimon, V., Prashanth, G.P., George, B.: Prediction of Severity Levels of Alzheimer's Disease from a Comprehensive MRI Data Analysis: An Explainable AI Framework for Clinical Insights. SSRN, 4720369
[13] Gómez-Talal, I., Bote-Curiel, L., Rojo-Álvarez, J.L.: Interpretable Machine Learning Models for PISA Results in Mathematics. IEEE Access (2025)
[14] Bettahi, A., Belouadha, F.Z., Harroud, H.: A Modular and Explainable Machine Learning Pipeline for Student Dropout Prediction in Higher Education. Algorithms. 18(10), 662 (2025)
[15] Henriques, J., Rocha, T., de Carvalho, P., Silva, C., Paredes, S.: Interpretability and explainability of machine learning models: Achievements and challenges. In: International Conference on Biomedical and Health Informatics, pp. 81-94. Springer Nature Switzerland, Cham (2022)
[16] Kaur, G., Rana, A., Sharma, D., Singh, D., Mukhija, L.: A TabNet Framework with SHAP and LIME for Explainable Child Anemia Prediction Using Tanzanian Health Survey Data. In: 2025 12th International Conference on Reliability, Infocom Technologies and Optimization (Trends and Future Directions) (ICRITO), pp. 1-5. IEEE (2025)
[17] Kukkar, A., Kaur, G.: AEC: A novel adaptive ensemble classifier with LIME and SHAP-based interpretability for fake news detection. Expert Systems with Applications. 281, 127751 (2025)
[18] Kukkar, A., Kaur, G., Kaur, M., Tandon, R., Kaur, I., Mukhija, L.: COVHIO-Based Feature Selection for COVID-19 Herd Immunity Prediction Using Deep Learning. In: 2025 12th International Conference on Reliability, Infocom Technologies and Optimization (Trends and Future Directions) (ICRITO), pp. 1-6. IEEE (2025)
[19] Kaur, G., Muskan, Kaur, S., Sandhu, S.S.: Unveiling the synergy: Advanced insights into artificial intelligence, machine learning, and deep learning in the age of data abundance. AIP Conference Proceedings. 3343(1), 040053 (2025)
[20] Kohan, A., Xu, J., Xiao, L., Wu, X., Kukkar, A., Hussain, S., Acharya, U.R.: Human-Centered Pathways to Trustworthy AI in Healthcare: A Comparative Analysis of Explainable AI, Human-in-the-Loop, Hybrid AI, and Uncertainty Quantification Techniques (2026)