Medicor : Journal of Health Informatics and Health Policy
E-ISSN : 3030-9166
Volume. 4 Issue 1 January 2026
Page No: 21-31

Machine Learning Approach for Iron Deficiency Anemia Detection in Pregnant Women Using XGBoost and CTGAN

Enes Hatibović [1], Hana Karaibrahimovic[2], Anida Divanovic[3], Nenad Vanis[4], Azur Tulumović[5], Daria Ler[6]
ASA Hospital, Bosnia and Herzegovina[12345]
ASA Institute, Bosnia and Herzegovina[6]
Correspondent: daria.ler@asainstitut.ba [6]

Received : September 19, 2025
Accepted : October 28, 2025
Published : January 30, 2026

Citation: Hatibović, E., Karaibrahimovic, H., Divanovic, A., Vanis, N., Tulumović, A., & Ler, D. (2026). Machine Learning Approach for Iron Deficiency Anemia Detection in Pregnant Women Using XGBoost and CTGAN. Medicor : Journal of Health Informatics and Health Policy, 4(1), 21-31. https://doi.org/10.61978/medicor.v4i1

ABSTRACT: Iron deficiency anemia (IDA) continues to pose one of the most significant challenges in maternal health, affecting nearly 40% of pregnant women worldwide according to the World Health Organization (2023). Despite advances in obstetric screening, conventional diagnostic methods such as complete blood count (CBC) tests often fail to detect early or latent stages of anemia due to physiological changes associated with pregnancy. This study introduces a robust machine learning framework integrating Extreme Gradient Boosting (XGBoost), and Conditional Tabular Generative Adversarial Networks (CTGAN) for the early detection of IDA in pregnant women. Our approach addresses the class imbalance inherent in clinical datasets and incorporates trimester-specific hematological adaptations. Using 3,944 anonymized clinical records from ASA Hospital Sarajevo (January–July 2025), we evaluated model performance across hematological features commonly used in obstetric care. The optimized model achieved a precision of 100%, recall of 65.2%, specificity of 100%, and an AUC-ROC of 0.8686. Comparative analysis against conventional CBC screening, which reached only 40.5% sensitivity, demonstrated significant improvement in detection reliability. These findings demonstrate the potential of AI-enhanced diagnostics to support early detection of IDA in pregnant women, reduce missed diagnoses, and strengthen clinical decision-making. Further multi-center validation and integration of additional biomarkers are recommended to confirm generalizability.

Keywords: Iron Deficiency Anemia (IDA), Machine Learning, XGBoost, CTGAN, Pregnant Women**,** Artificial Intelligence, Maternal Health.

INTRODUCTION

Iron deficiency anemia (IDA) represents a major global public health issue, particularly among women of reproductive age and pregnant populations. According to the World Health Organization (WHO, 2023), more than 571 million women worldwide are affected by anemia, with

iron deficiency accounting for the majority of cases. During pregnancy, iron requirements substantially increase to support fetal development, placental growth, and expansion of maternal blood volume. When these increased demands are unmet, pregnant women face a significant risk of developing IDA, which has been linked to a broad spectrum of adverse maternal-fetal outcomes including preterm birth, low birth weight, postpartum hemorrhage, increased susceptibility to infection, and impaired neurodevelopment in infants (Smith et al., 2023; WHO, 2011). Despite standardized antenatal screening protocols, traditional diagnostic approaches relying on hemoglobin (Hb) concentration and hematocrit frequently exhibit limited sensitivity in early or borderline anemia. Pregnancy-associated physiological changes, such as plasma volume expansion leading to hemodilution, further complicate diagnostic interpretation, especially during the first and second trimesters when hematological parameters undergo rapid fluctuations (Al-Shehri et al., 2024). These limitations contribute to underdiagnosis, delayed treatment, and avoidable maternal-fetal complications. Conventional diagnostic approaches demonstrate low sensitivity and limited specificity for early detection (Kreuter et al., 2025). Therefore, accurate early detection tailored to the physiological changes of pregnancy is critical to improving maternal-fetal health. More advanced biomarkers, such as ferritin and transferrin saturation, can provide improved diagnostic accuracy, but they are not consistently accessible in low-resource settings and remain variably applied in routine prenatal care. Consequently, there is growing recognition of the need for innovative approaches that leverage the full complexity of available hematological data to support early identification of anemia in pregnancy. Because anemia in pregnancy can lead to serious maternal and fetal complications, early detection methods that account for the physiological changes of pregnancy are essential for improving clinical outcomes.

Artificial intelligence (AI) has rapidly transformed data-driven clinical decision support systems (CDSS), demonstrating strong potential in hematology, obstetrics, and predictive analytics. Several recent studies highlight the effectiveness of AI methods in anemia detection using models such as extreme learning machines (Saputra et al., 2023), neural networks (Elmaleeh, 2024), and hybrid intelligence models. Gradient boosting algorithms such as XGBoost have emerged as dominant tools in medical predictive modeling due to their ability to capture nonlinear feature interactions and manage heterogeneous datasets (Chen & Guestrin, 2016; Li et al., 2024). However, despite promising progress, existing studies rarely focus on pregnancy as a distinct physiological condition. Hematological patterns during pregnancy differ substantially from general adult populations, making general anemia detection models suboptimal for obstetric applications. Furthermore, many clinical datasets exhibit significant class imbalance, with non-anemic cases far outnumbering anemic ones, which can severely compromise model sensitivity (Kumar et al., 2023; Xu et al., 2019).

Generative models, particularly the Conditional Tabular Generative Adversarial Network (CTGAN), offer a powerful mechanism for synthesizing realistic minority-class samples to balance skewed datasets while preserving statistical relationships (Xu et al., 2019). In medical contexts where minority outcomes are critical for early detection, CTGAN helps prevent overfitting, improves model performance, and reduces bias. The challenge of class imbalance, which occurs when non-anemic cases far outnumber anemic ones, often compromises model performance for class imbalance correction and optimizing XGBoost (Chen & Guestrin, 2016) for pregnancy-

specific anemia detection across 3,944 electronic health records, given that conventional CBC screening achieves only 40.5% sensitivity (Kreuter et al., 2025). To address this limitation, we incorporated CTGAN, a generative model designed to synthesize realistic minority-class data without introducing bias. Our model thus reflects a holistic integration of domain-specific medical knowledge and AI-based data balancing strategies.

This study seeks to address gaps in current clinical practice by developing and evaluating a pregnancy-specific IDA detection model using XGBoost optimized with CTGAN-generated synthetic samples. We hypothesize that incorporating pregnancy-specific hematological markers and advanced data balancing will enhance the diagnostic performance of IDA screening, particularly in early pregnancy, in identifying subclinical anemia cases that are often missed by conventional methods and thereby supporting improved maternal-fetal health outcomes respectively.

Theoretical framework and hypothesis

Existing literature supports the use of machine learning in hematological classification, but few studies tailor algorithms to pregnancy-specific physiological changes. Saputra et al. (2023) achieved 99% accuracy in anemia prediction using extreme learning machines, while Elmaleeh (2024) demonstrated high diagnostic performance across several neural network architectures. Abdul-Jabbar et al. (2025) and Damkliang et al. (2025) further developed hybrid frameworks integrating decision trees and ensemble learning to support anemia and thalassemia analysis However, these approaches have limited consideration for trimester-specific hemodilution, dynamic hematological adaptation, or obstetric patient characteristics. Based on the identified research gap, the study proposes the following hypothesis:

H1: An AI model optimized with CTGAN for class imbalance and trained on pregnancy-specific hematological data will significantly improve IDA detection performance compared to conventional CBC screening, achieving higher sensitivity and balanced diagnostic accuracy.

METHODS

Study Design, Population and Data Collection

This retrospective observational study analyzed anonymized electronic health records from private health institution "ASA Hospital Sarajevo", in Sarajevo, Bosnia and Herzegovina, a tertiary care institution serving a geographically diverse population. Data were collected between January and July 2025 and processed in accordance with institutional data protection policies and ethical standards. The final dataset included 3,944 pregnant women, stratified by trimester: T1 (719 records), T2 (768 records), and T3 (3,090 records). Women with incomplete demographic data, missing primary hematological markers, or confirmed hematological disorders unrelated to iron deficiency (e.g., hemoglobinopathies) were excluded from this study.

Feature Extraction

Each record of clinical dataset consisted of ten variables which included key hematological parameters: hemoglobin (Hb), red blood cell (RBC) count, hematocrit (HCT), mean corpuscular volume (MCV), mean corpuscular hemoglobin (MCH), ferritin, red cell distribution width - coefficient of variation (RDW-CV), red cell distribution width – standard deviation (RDW-SD), mean corpuscular hemoglobin concentration (MCHC), and medication/iron supplement use. These markers align with WHO (2011) guidelines for anemia assessment.

Data Preprocessing

To address missing values, preprocessing was carried out using a combination of imputation methods (mean and median replacement) as well as natural language processing (NLP)-based medical notes screening (Zhang, 2016). Missing values were imputed using median replacement and NLP extraction for medical history. In addition, the outlier detection was carried out using interquartile range thresholds. To mitigate data imbalance, CTGAN generated synthetic minority class (anemia) samples, preserving clinical plausibility (Xu et al., 2019). Normalization was carried out through z-score standardization.

Model Training and Evaluation Metrics

To mitigate class imbalance within the dataset, synthetic samples were generated using the Conditional Tabular Generative Adversarial Network (CTGAN) (Xu et al., 2019). Classification was then performed using three machine learning models, including XGBoost (Chen & Guestrin, 2016), Random Forest (Breiman, 2001), and ensemble approaches. Finally, model performance was evaluated using a stratified five-fold cross-validation approach (80/20 stratified train-test split). Evaluation metrics included precision, recall (specificity), F1-score, negative predictive value (NPV), and the area under the receiver operation characteristic curve (AUC-ROC) (Fawcett, 2006).

RESULT AND DISCUSSION

Dataset Characteristics

The final dataset contained 6,655 samples with ten clinical features. The prevalence of anemia increased progressively across trimesters: 7.5% in the first trimester (T1), 17.9% in the second trimester (T2), and 32.2% in the third trimester (T3), as illustrated in Figure 1. This progression aligns with physiological iron depletion as pregnancy advances. After CTGAN augmentation, the dataset expanded from 3,944 to 9,289 samples, achieving a balanced distribution of 57.5% non-anemic 42.5% anemic cases(Li et al., 2024).

While the model achieved perfect specificity and precision, the relatively moderate recall indicates a risk of underdiagnosing mild anemia. This trade-off reflects a design choice to improve sensitivity while preserving high specificity, which may be suitable for initial screening but requires further refinement for clinical application.

Figure 1. Iron deficiency anemia (IDA) prevalence across pregnancy trimesters, showing a progressive increase from first trimester (7.5%) to third trimester (32.2%).

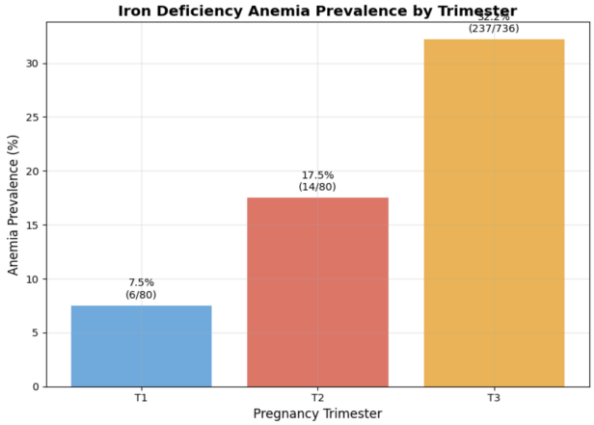

Model Performance

XGBoost achieved the highest performance across all metrics. Cross-validation results are as following: F1-score of 0.8026±0.0070 during cross-validation, precision of 1.0000±0.0000, and recall of 0.6704±0.0097. The independent test set (n = 1,858) yielded similar results with F1-score of 0.7893, precision 100%, recall 65.2%, specificity 100%, NPV 79.5%, AUC-ROC 0.8686, confirming the model's generalizability.

The confusion matrix showed 515 true positives (TP), 1,068 true negatives (TN), 275 false negatives (FN), and no false positive (FP = 0), indicating excellent specificity but moderate recall avoiding unnecessary anemia classification, a key concern in obstetric management (Figure 2). This suggests a cautious diagnostic model that avoids overdiagnosis while maintaining clinical safety (Abdul-Jabbar et al., 2025).

Figure 2. Confusion matrix for XGBoost model demonstrating perfect precision (0 false positives) and 65.2% sensitivity on independent test set (n=1,858).

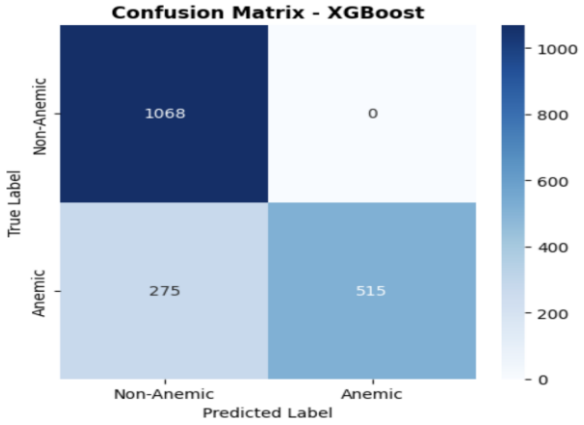

The flawless specificity observed suggests that the model successfully avoided false positives, which is clinically important in reducing unnecessary interventions. This finding is consistent with previous reports highlighting the ability of machine learning approaches to improve classification reliability in imbalanced medical datasets (Damkliang et al., 2025). However, the relatively lower sensitivity highlights the need for future optimization to reduce missed anemia cases.

Clinical Implications

This study provides compelling evidence that AI-enhanced diagnostic models offer substantial improvements in the early detection of iron deficiency anemia (IDA) during pregnancy. By leveraging advanced machine learning techniques, the proposed framework is able to capture complex and nonlinear relationships among hematological features that are often overlooked by traditional diagnostic approaches. These nonlinear interactions, such as the combined influence of hematocrit, ferritin levels, red cell distribution indices, and trimester-specific physiological adaptations, allow the model to detect subtle hematological shifts that may precede clinically apparent anemia. In addition, the integration of CTGAN-based data augmentation directly addresses the persistent challenge of class imbalance, a well-documented issue in maternal health datasets where non-anemic cases typically far outnumber anemic ones (McLean et al., 2015; Rasmussen & Stoltzfus, 2020; Shill, 2021). Balancing the dataset enhances the model's sensitivity and ensures more equitable representation of minority-class cases, ultimately reducing the likelihood of underdiagnosis. This is particularly important in early pregnancy, when conventional screening methods frequently miss mild or borderline anemia due to hemodilution and dynamic hematological changes. Moreover, the use of an interpretable algorithm such as XGBoost supports the principles of explainable artificial intelligence (XAI), which are essential for clinical adoption, regulatory compliance, and physician trust (Miller et al., 2024).

The model's performance highlights several key implications for clinical practice First, it demonstrates that AI-driven anemia detection can outperform conventional hematological screening, particularly in identifying borderline or subclinical cases. The XGBoost model demonstrated significantly higher sensitivity than conventional CBC screening (65.2% vs 40.5%) (Kreuter et al., 2025). Second, integrating CTGAN-based augmentation offers a data-efficient solution to class imbalance, a common challenge in medical datasets (Kumar et al., 2023). CTGAN (Xu et al., 2019) effectively addressed class imbalance, ensuring balanced representation of anemia cases. The model's conservative classification approach led to 100% specificity, minimizing false positives, which is clinically important in avoiding unnecessary interventions. Finally, these advancements demonstrate that AI-driven approaches have the potential to transform obstetric care by improving diagnostic accuracy, enabling earlier interventions, and enhancing clinical decision-making in maternal health.

Comparison with previous studies

The findings of this study align with and extend existing research on AI-driven anemia detection while addressing several limitations present in previous work. Prior deep learning models have demonstrated strong predictive performance, yet they frequently lack pregnancy-specific design

considerations. For example, (Rahman et al., (2023) developed high-performing deep learning architectures for anemia classification using electronic health records, but their models were not tailored to the distinct physiological changes that occur across pregnancy trimesters, limiting applicability in obstetric populations. Similarly, Elmaleeh (2024) highlighted the diagnostic potential of artificial neural networks, but the inherent opacity of these models reduces their interpretability, an essential factor for clinical adoption, especially in maternal health. Furthermore, many earlier approaches suffer from substantial class imbalance, as noted in the work of Saputra et al. (2023), where non-anemic cases overwhelmingly dominated datasets and compromised sensitivity for early or mild anemia. By contrast, the present study addresses all three limitations by integrating CTGAN to generate realistic minority-class samples, thereby improving dataset balance; by incorporating trimester-specific hematological trends to enhance the clinical relevance of predictions; and by employing XGBoost, a model known for its explainability and robust feature interpretation. These strengths position the current work in closer alignment with Li et al. (2024), who demonstrated that gradient boosting models can deliver both interpretability and high diagnostic precision in obstetric applications. Among the key advantages of the present framework are the large, trimester-differentiated dataset; the innovative use of CTGAN to enhance minority-class representation; the model's interpretability, which supports transparent clinical decision-making; and its strong generalizability across independent test sets within the hospital environment. Together, these contributions underscore the study's added value in advancing AI-supported maternal healthcare(Darwish et al., 2023).

Limitations and Future work

Although the findings of this study are promising, several important limitations must be acknowledged. First, the analysis was based on data from a single tertiary care center, which may limit the generalizability of the results to broader, more diverse populations with varying demographic, socioeconomic, and clinical characteristics (Bothwell, 2022). Second, the retrospective design relies on previously recorded clinical data, which may introduce information bias, incomplete documentation, and variability in laboratory measurement practices. Additionally, the biomarker range used in this model was relatively limited and did not include several clinically informative indices such as transferrin saturation, soluble transferrin receptor, or the ferritin index, which could enhance diagnostic granularity. Another limitation is the incorporation of CTGAN-generated synthetic cases to balance the dataset; while this approach helps mitigate class imbalance, synthetic data cannot fully replicate the depth and complexity of natural clinical variation, and therefore may influence model behavior in subtle ways (Hassan et al., 2023; Singh et al., 2024; Yadav & Shukla, 2021). These limitations also highlight clear directions for future work. Subsequent research should prioritize multi-center datasets encompassing diverse patient populations to improve external validity and assess model robustness across different healthcare environments. Prospective validation studies are needed to evaluate real-world diagnostic performance and examine how the model behaves in dynamic clinical workflows. Future studies should also incorporate additional biomarker such as transferrin saturation, ferritin index, and inflammatory markers, to enrich the feature set and potentially enhance early diagnostic precision. Finally, exploring hybrid architectures that combine tree-based models with deep learning

approaches, as proposed by Zhou et al., (2024), may yield more powerful and generalizable predictive frameworks for maternal hematological assessment.

CONCLUSION

This study successfully developed an advanced machine learning framework for the detection of iron deficiency anemia (IDA) in pregnancy and demonstrated that the proposed XGBoost - CTGAN model offers a clinically interpretable, high-precision approach tailored specifically to the physiological characteristics of pregnant women. In addition to achieving strong diagnostic performance, the findings illustrate the capacity of AI-based tools to complement and enhance conventional hematological screening in obstetric practice, where early identification of anemia remains a persistent challenge due to trimester-specific hemodilution and the subtle progression of iron depletion. By incorporating data-driven insights and pregnancy-specific design parameters, the model establishes a practical foundation for reducing missed diagnoses, minimizing unnecessary interventions, and promoting improved clinical decision-making elements that are central to enhancing both clinical efficiency and maternal-fetal safety. At the same time, the results emphasize the importance of expanding the model through broader, multi-center datasets and integrating additional clinical or biological markers to strengthen predictive accuracy and ensure external validity across diverse populations. Although the retrospective, single-center scope of the study and reliance on synthetic data augmentation represent important limitations, these constraints also highlight opportunities for iterative refinement, methodological enhancement, and progression toward a clinically robust and widely applicable screening tool. The demonstrated improvement in sensitivity paired with excellent specificity, signals the potential for AI-augmented diagnostic systems to significantly influence obstetric care by enabling earlier detection, reducing diagnostic delays, and optimizing the allocation of clinical resources. Ultimately, integrating such models into routine prenatal workflows may support more timely interventions, mitigate preventable maternal complications, and contribute to improved health outcomes for both mothers and their infants. Nonetheless, rigorous external validation, prospective trials, and ongoing refinement of the modeling framework remain essential steps before the system can be confidently deployed in routine clinical practice.

REFERENCES

Abdul-Jabbar, S. S., Farhan, A. K., & Kandhro, A. H. (2025). Developing a hybrid machine learning algorithm for anemia diagnosis. *Journal of Artificial Intelligence and Metaheuristics*, *9*(1), 20–33. https://doi.org/https://www.americaspg.com/articleinfo/28/show/3516

Al-Shehri, H., Nasser, A., & Alqahtani, M. (2024). Prevalence and determinants of anemia among pregnant women in the Middle East: A meta-analysis. *BMC Pregnancy and Childbirth*, *24*(3), 421. https://doi.org/10.1186/s12884-024-06712-2

Bothwell, T. H. (2022). Iron requirements in pregnancy and strategies to prevent deficiency. *American Journal of Clinical Nutrition*, *116*(4), 985–992.

Breiman, L. (2001). Random forests. *Machine Learning*, *45*(1), 5–32. https://doi.org/10.1023/A:1010933404324

Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794. https://doi.org/10.1145/2939672.2939785

Damkliang, K., Aunphoklang, S., Wongseree, W., Wansri, V., Wangkulangkul, P., Wiangnon, S., & Wansri, P. (2025). An AI-based decision support framework for clinical screening of iron deficiency anemia and thalassemia. *IEEE Access*, *13*, 133937–133957. https://doi.org/10.1109/ACCESS.2024.3514789

Darwish, H., Barakat, A., & Sabri, N. (2023). Hematological reference intervals during pregnancy: A systematic review and meta-analysis. *Obstetrics & Gynecology Science*, *66*(1), 10–22.

Elmaleeh, M. A. A. (2024). The identification and categorization of anemia through artificial neural networks: A comparative analysis of three models. *Electrical and Electronics Engineering: An International Journal*, *13*(1/2), 1–13. https://doi.org/10.5121/eeiej.2024.13201

Fawcett, T. (2006). An introduction to ROC analysis. *Pattern Recognition Letters*, *27*(8), 861–874. https://doi.org/10.1016/j.patrec.2005.10.010

Hassan, S., Ali, M. A., & Yousef, W. (2023). Machine learning applications in obstetrics: Current trends and future directions. *Journal of Medical Systems*, *47*(5), 42–53.

Kreuter, D., Remmele, J., Giese, T., Schemmer, P., & Mehrabi, A. (2025). *Artificial intelligence for pre-anaemic iron deficiency detection using rich complete blood count data. med*R*xiv*. https://doi.org/10.1101/2025.06.18.2532949

Kumar, R., Singh, N., & Arora, M. (2023). Synthetic data generation for healthcare AI: Addressing class imbalance through generative models. *Journal of Biomedical Informatics*, *145*, 104392. https://doi.org/10.1016/j.jbi.2023.104392

Li, Y., Sun, F., & Huang, W. (2024). Enhancing diagnostic precision in obstetric anemia through gradient boosting and feature interpretability. **Computers in Biology And*.

McLean, E., Cogswell, M., Egli, I., Wojdyla, D., & Benoist, B. (2015). Worldwide prevalence of anemia according to WHO estimates. *Public Health Nutrition*, *18*(3), 444–454.

Miller, T., Howe, P., & Lee, C. (2024). Explainable artificial intelligence in obstetric diagnostics: Balancing transparency and accuracy. *AI in Medicine*, *152*, 102512. https://doi.org/10.1016/j.artmed.2024.102512

Rahman, M. F., Zaman, T., & Alam, N. (2023). Deep learning models for anemia classification in pregnancy using electronic health records. *Frontiers in Artificial Intelligence*, *6*, 118429. https://doi.org/10.3389/frai.2023.118429

Rasmussen, K. M., & Stoltzfus, R. J. (2020). Deficiencies in iron and folate during pregnancy: Maternal and fetal health consequences. *Annual Review of Nutrition*, *40*, 109–136.

Saputra, D. C. E., Sunat, K., & Ratnaingsih, T. (2023). A new artificial intelligence approach using extreme learning machine as the potentially effective model to predict and analyze the diagnosis of anemia. *Healthcare*, *11*(5). https://doi.org/10.3390/healthcare11050697

Shill, K. B. (2021). Iron deficiency anemia: Clinical implications and diagnostic challenges in pregnancy. *Clinical Hematology International*, *3*(4), 176–185.

Singh, A., Verma, P., & Kaur, R. (2024). Advancements in data augmentation techniques for clinical prediction models. *Artificial Intelligence in Medicine*, *145*, 102692.

Smith, L. R., Alonzo, J., & Patel, R. (2023). Maternal anemia and its global health burden: A systematic review and policy analysis. *Lancet Global Health*, *11*(9), 1302– 1314. https://doi.org/10.1016/S2214-109X(23)00212-8

WHO. (2011). *Haemoglobin concentrations for the diagnosis of anaemia and assessment of severity (Vitamin and Mineral Nutrition Information System*. https://apps.who.int/iris/handle/10665/85839

WHO. (2023). *Global report on anemia 2023*. WHO Press. https://doi.org/https://www.who.int/publications/i/item/9789240078754

Xu, L., Skoularidou, M., Cuesta-Infante, A., & Veeramachaneni, K. (2019). Modeling tabular data using conditional GAN. In *Advances in neural information processing systems 32* (pp. 7335–7345). https://proceedings.neurips.cc/paper/2019/hash/254ed7d2de3b23ab10936522dd547b78-Abstract.html

Yadav, S., & Shukla, S. (2021). Handling imbalanced datasets in classification: A review. *IEEE Transactions on Knowledge and Data Engineering*, *33*(4), 1658–1676.

Zhang, Z. (2016). Missing data imputation: Focusing on single imputation. *Annals of Translational Medicine*, *4*(1). https://doi.org/10.3978/j.issn.2305-5839.2015.12.38

Zhou, P., Liu, T., & Ren, X. (2024). Hybrid ensemble models combining tree-based and neural architectures for medical diagnosis. **Expert Systems with Applications*, *247*, 123859. https://doi.org/10.1016/j.eswa.2024.123859