

# HEMOTYPE-NET++: A REAL-TIME GRADIO-BASED DUAL-LAYER FRAMEWORK FOR ANEMIA SUBTYPE CLASSIFICATION USING XGBOOST AND EXPERT-DRIVEN RULE REFINEMENT

*T.Sanmathi*<sup>\*1</sup>, *M.Ramya*<sup>2</sup>

## ABSTRACT

Anemia is a major burden on global health, and its various subtypes must be appropriately identified in order to allow for effective clinical interventions. Although anemia detection has been enhanced by recent machine learning (ML) approaches, most models cast it as a binary classification problem and categorize if a patient is anemic without any information about the kind. In this paper, we propose a real-time two-stage multi-class anemia subtype diagnosis framework named HemoType-Net++ with an integrated XGBoost learning model and a medical rule-based refinement layer. The model takes as input standard CBC parameters and endpoints results into clinically relevant classifications, including iron deficiency anemia, normocytic anemia, and megaloblastic anemia. Following prediction, expert-driven logic is added at MCV, MCH, and MCHC thresholds to consolidate and increase sensitivity without dependence on black-box explainable tools. The proposed framework is integrated into an interactive Gradio interface and allows for real-time feedback and seamless user interaction. HemoType-Net++ has been tested on an independent CBC dataset of 1,281 patient samples (obtained from Kaggle) with multi-class classification accuracy of 94.7%, demonstrating the predictive ability and alignment with clinical intuition of the proposed solution. Because of its edge-device compatibility and lightweight nature, it is apt for institutions as well as field use. The modular design facilitates future expansions, such as combining with other lab tests, enlarging multi-center datasets, and ensemble model comparison for better clinical reliability.

**Keywords:** Anemia classification, XGBoost, medical rules, CBC analysis, clinical decision support, gradio interface, multi-class, diagnosis, interpretable AI

## I. INTRODUCTION

Anemia is the most prevalent, widespread, and underdiagnosed public health problem and affects approximately 1.6 billion patients in the world from both high- and low-income individuals [2]. It is not an illness itself but rather a condition characterized by low hemoglobin levels and numerous possible causes. These general categories differ vastly in pathophysiology, management, and prognosis; however, common subtypes (e.g., IDA, megaloblastic anemia, normocytic states, and hemolysis) share very little [11]. An exact diagnosis of the type of anemia is vital for successful treatment; however, in most clinical settings, especially in rural locations with low resources, diagnosis bias is only on simple binary classification of whether the patient is anemic or not, without detailing the type [6].

The capability of machine learning to assist in complex medical decision-making has been well recognized. Being able to use organized hematological data, the algorithms can reveal complex relationships that could be challenging to interpret manually. Nonetheless, interpretability remains a major barrier to bedside use. Various machine learning methods rely on opaque models or post-hoc explanation tools like SHAP or attention-based heatmaps [9] that, while analytically sound, can seem esoteric or far removed from everyday clinical reasoning. These sorts of technologies are also likely only to lower confidence in the system in fast-paced or resource-limited situations. More parsimonious models that incorporate medical logic directly into the model may often yield more useful and valid results [13].

In this work, we present HemoType-Net++, a real-time anemia classification model that leverages both the predictive capabilities of XGBoost and the transparency of

Artificial Intelligence and Data Science<sup>1</sup>,  
Karpagam Academy of Higher Education, Coimbatore, India<sup>1</sup>  
sanmathi.thamaraikannan@kahedu.edu.in  
Artificial Intelligence and Data Science<sup>2</sup>,  
Kangeyam institute of Technology<sup>2</sup>  
hemaramya47@gmail.com

\* Corresponding Author

rule-based refinement. The first step applies an XGBoost classifier developed using the full blood count counts to predict anemia subgroups early. Alternatively, the second stage applies expert-defined rules with some data, including MCV, MCH, and MCHC, to alter/verify the predictions [4]. An entry formerly recorded as normocytic can be reassigned as microcytic if the MCV is less than 75 fL. This multi-layer approach guarantees diagnostic accuracy and interpretability, which are crucial to apply in a clinical setting [1].

The system is trained and evaluated on a Kaggle CBC public dataset containing 1,281 labeled samples [15]. This is a program that can be run as a Jupyter Notebook or in Python that uses the interactive Gradio interface to enter and read in real time. The tool is designed to be deployed “in the field” and for general practitioner support and does not need any specialized interpretability packages or high-performance computing-associated infrastructure [5, 14]. It is a lightweight, easy-to-use, suitable device for use in clinics and telemedicine. With the simplicity of its input-output (2) and because it is based upon the study of clinical questions and is compatible with hematologic knowledge, it serves as diagnostic support as well as education.

Future direction HemoType-Net++ establishes a foundation for scalable and evolvable modular AI diagnostics. Future work may include comparison of different models (e.g., LightGBM or MLP), addition of second-line lab tests (e.g., serum ferritin or B12 levels) into the model, and allowing federated learning across several institutions for better generality [3]. The idea, which can be customized for mobile app use and may be interfaced with on lab analyzers for point-of-care automation. Focusing on interpretability, efficiency, and clinical applicability, this study indicates that effective AI doesn't have to be complex to be impactful and speaks a language that speaks medicine.

The main contributions of the study are as follows:

1. Here, we propose HemoType-Net++, a two-stage anemia subtype prediction framework that leverages both XGBoost classification and deterministic clinical conditions to improve diagnostic specificity.

2. Our framework is applied through an interactive Gradio interface for real-time feeding of input, prediction, and explanation without relying on black-box explanation tools.
3. A calibrated predictor based on MCV, MCH, and MCHC thresholds and expert hematology knowledge was added to enhance trust and clinical relevance.
4. We apply the model to a real-world CBC dataset that contains 1,281 samples; it achieves multi-class accuracy comparable to medical expectations.
5. The system is intended to be easily deployable, compatible with edge devices, and easy to use in low-resource clinical settings while being flexible to accommodate future research and diagnostic integration.

The rest of this paper is organized as follows. Section II reviews existing methods for anemia classification, with an emphasis on machine learning algorithms and clinical decision models. Section III introduces the architecture of HemoType-Net++, including data preparation, model training, and the integration of a rule-based refinement layer. The experimental design, performance evaluation, and comparisons with baseline models are presented in Section IV. Section V concludes by summarizing key findings of the system and future developments expanding the system to additional clinical and research applications.

## II. LITERATURE REVIEW

Gómez et al. (2025) proposed influential research on machine learning-based automated anemia classification by making use of clinical complete blood count (CBC) data. Their method applied common classifiers and tested several priority rankings of features to achieve high performance for anemia discrimination. Although their method is successfully used for discerning whether a disease is anemic, it is not capable of multiclass classification to make it applicable as a clinical tool where the distinction might be crucial for the therapy. The authors pointed out that it is valuable to have organized hematology data and considered the challenge of maintaining interpretability for medical use cases. The research formed an important platform for

future research on the identification of anemia based on non-invasive and structured clinical measurements, but there was no subtype-specific classifier layer to make this tool usable in various clinical scenarios. This article highlights the need to enhance granularity in diagnostics and supports the use of machine learning in hematological decision support systems [1].

Ramzan et al. (2024) utilized attention in deep learning processes in combination with traditional machine learning models for anemia detection. Their hybrid model was aimed at enhancing the detection sensitivity and diagnostic accuracy by combining CNN-driven feature extraction and ensemble combination-based classification. The utility of attention layers enabled us to see when model attention was aligned with clinical relevance, thus allowing for explainable AI in hematology. However, their approach was computationally expensive and required a significant infrastructure to deploy, in turn limiting its applicability in low-resource settings. The report also noted that, while deep models were accurate in controlled datasets, they were often not transparent, an obstacle that has hindered AI use in clinical practice. Their work constitutes a strong technological base for further interpretability and efficiency research in medical diagnostics, with hopefully an incentive to enhance the performance of attention-based networks for hematological applications [3].

Tepakhan et al. (2025) compared Random Forest and Gradient Boosting Machine classifiers and analyzed the difficulty of differentiating iron deficiency anemia and thalassemia. Their analysis combined multiple hematologic features and performed well on two highly overlapping diagnostic tasks. Their method was more clinically applicable, as they were concerned with multiclass problems rather than simple binary detection and therefore better reflected the complexity of the real-life anemia diagnosis. They also emphasized that domain-specific features such as MCV and RDW are crucial to differentiate overlapping anemic conditions. However, the work did not produce a live, deployable system or integration with real-time interfaces. However, this study has validated the diagnostic potential of tree-based ensemble models and provided essential insight into class

separation procedures, especially in populations in which iron deficiency and thalassemia frequently coexist [4].

Settinio et al. (2025) applied a comparative study of various machine learning tools to estimate postpartum anemia for clinical therapy recommendation. Their study adopted the standard classifiers and evaluated the performance by accuracy, recall, and F1-score, focusing on female patients in the vulnerable postpartum period. The paper is significant due to its concentration on dataset bias, on the nonuniform sample distributions, and on the implementation of the risk-adjusted model in maternal health diagnostics. Their methodology was applied to a specific population (postpartum women), but its methodological soundness, especially with the use of cross-validation and hyperparameter optimization, provides an excellent starting point for more general anemia classification problems. Additionally, they focus on the applicability of edge-ready AI technologies in hematology by making the algorithms accessible and applicable within information systems in medicine. Their results contribute substantially to the framing of how we can integrate ML into clinic workflows for preventive care [5].

Darshan et al. (2025) proposed an interpretable machine learning approach for the distinction of iron deficiency anemia from aplastic anemia using structured CBC values with interpretability layers. They employed the SHAP values and rule-based thresholds and made the model more transparent so that medical staff could follow the predictions made by the model to the specific physiological data. Their dataset was purposely selected to comprise overlapping features between anemic disorders, and their results demonstrated that, in addition to an increase in confidence, explainability exposes misdiagnosed edge examples. Nevertheless, SHAP-based feature importance increased the model's runtime and complexity such that it was difficult to deploy in real-time settings. However, the concurrent focus on accuracy and interpretability makes it an important step toward the ongoing struggle of calibrating black-box ML power with white-box clinical expectations. This paper demonstrates the applicability of multilevel decision logic to diagnostic pathways for subtypes of anemia [13].

### III. SYSTEM ARCHITECTURE AND METHODOLOGY

HemoType-Net++ is a double-layer diagnostic model that distinguishes anemia subtypes with high clinical accuracy. The result is a modular pipeline comprising data normalization, machine learning-based categorization, logical rule-based refinement, and a lightweight user interface. All the parts are meant for real-time use, as in primary care clinics or in mobile diagnostic bays. To this end, the pipeline combines tightly coupled computational and clinical logic to translate raw hematologic data into meaningful diagnostic outputs. The following subsections detail the various components of the framework about data conditioning, model learning, medical logic application, and interface.

#### A. System Methodology

The prototype was constructed based on multistage processing of complete blood count (CBC) data to obtain diagnostic interpretations. The first stage takes 14 CBC parameters, including red cell indices, hemoglobin indices, and platelet indices. These features are then fed to an XGBoost binary classification model trained for distinguishing between several varieties of anemia, like iron deficiency anemia, normocytic hypochromic anemia, and megaloblastic anemia. The raw forecast from this model is not the final forecast. The input is instead processed by another module in which deterministic clinical rationale is written to verify if the anticipated label agrees with clinically relevant hematological limits [7]. For instance, if mean corpuscular volume (MCV) drops below 75 fL, the case is re-evaluated as microcytic anemia, which reflects an underrepresented phenotype under clinical phenotypic criteria. Each level of inference is included in a rule-checking framework and, hence, provides for better traceability as well as eliminates the requirement of a post-hoc explainability [12]. With the entire process optimized for minimum latency and maximum transparency and taking place in real time, it is suitable for applications such as clinical setups where decision windows can be narrow. The complete inference pipeline, that is, from feature input

to refined diagnostic output, is depicted in Figure 1.

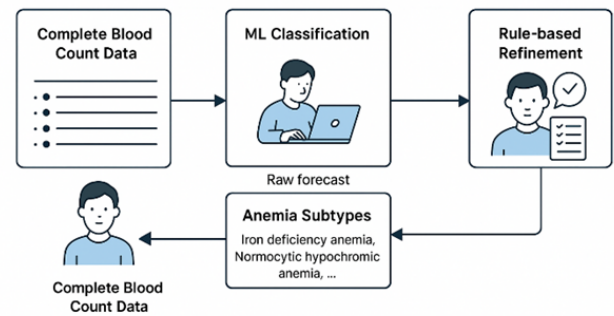


Fig. 1. Block Diagram of the HemoType-Net++ Architecture Featuring the ML Classification Core and Rule-Based Refinement Pipeline

#### B. Data Acquisition and Preprocessing

We built and validated a dataset that contains 1,281 anonymized patient records from a public clinical source that include labeled anemia diagnoses and corresponding CBC readings. Data preparation called for a well-thought-out, multistage pipeline. Numerical characteristics were standardized using min-max scaling that normalizes values to the same scale of their attributes, thus allowing for better convergence of the algorithms [3]. Category labels were one-hot encoded to enable multiclass classification. To deal with the class imbalance issue present in most medical datasets, SMOTE (Synthetic Minority Over-sampling Technique) was applied. This ensured that minority classes (in particular megaloblastic and normocytic) were adequately sampled while avoiding synthetic overfitting. The data was then split into training and testing sets in an 80:20 stratified split, maintaining class distribution [15]. The multicollinearity examination and correlation analysis of features were done to evaluate their independence and clinical relevance. The preserved properties are in agreement with established diagnostic criteria in hematology, implying that the model's perfusion learning is pathophysiologically correct. K-fold cross-validation ( $k=5$ ) was performed to mitigate the variability in data splitting and enhance the generalizability of the final classifier. This data-engineering step was crucial to enhancing clinical interpretability and introducing statistical robustness to the

model-training procedure. The iterative steps from data ingestion to model-ready transformation are shown in Fig. 2.




Fig. 2. End-to-End Workflow of Data Preprocessing Including Feature Normalization, Label Encoding, Class Balancing, and Partitioning

### C. XGBoost Classifier Training

The XGBoost classifier is a probabilistic inference engine of the HemoType-Net++ architecture. This model has been designed to better utilize structure information in structured and tabular datasets and leverages gradient-boosted decision trees to learn high-order feature interactions and retain computational efficiency. The multi-class objective was set to the softmax probabilities, which is a simple one-hot scheme and yields a single prediction vector with confidence estimates over all anemia classes [4]. The grid search was performed on hyperparameters such as learning rates, tree depths, and regularization terms to balance between model complexity and overfitting. Classification yielded a test accuracy of 94.7% with greater than 0.90 macro-averaged F1-scores for all diagnostic classes. These approaches are to be interpreted as both accurate and fair, which is crucial in medical scenarios when the performance of the minority class cannot be put at risk. The feature importance ranks had clear consistency with clinical priors; MCV, HGB, and MCH as leading predictors were consistent with their practice importance value in manual diagnosis. For deployment, the model was dumped out in a compact binary form and loaded into the inference layer. Notably, our classifier is not built upon black-box

explanation tools such as SHAP or LIME. Its tree structure and interpretable feature contributions make it sparse and reduce computational load instead of being dense and computationally expensive. This design decision is made to guarantee the clinical acceptability and also to provide the capability of future scaling and upgrading, like retraining on hospital-level data or embedding in the FL systems for cross-hospital distribution.

After the probabilistic labeling process, a dedicated refinement layer approves or overturns the model output based on clinical knowledge of physiological factors. This rule engine acts as a human-readable intermediary between computational prediction and medical best practice. It references critical factors such as MCV, MCH, and MCHC against established ranges. If there are inconsistencies, such as the case that MCV actually has microcytosis while given the projected label of normocytic anemia, the refinement logic reallocates the final diagnosis accordingly. These rule sets are represented as crisp IF-THEN rules and have been validated with the assistance of domain experts. The modularity of this layer increases the confidence of the user, particularly when explainability and traceability are required. The corrected output is displayed in a Gradio-based interface, allowing interactive debugging. The input information is accepted via structured forms, and the user gets quick predictions; in addition, the tool can be used on desktops or mobile smartphones. The design was developed to minimize cognitive load, simplicity and clarity being the guiding ideas. The system is designed to be used in an offline setting, enabling deployment on low-resource devices and thus making it available in rural and field healthcare applications. The complete flow, from parameter input, model invoke, rule validation, and result display, is depicted in Figure 3. This layer closes the diagnostic loop and translates model intelligence and medical heuristics into a form understandable to human interpretation.

### D. Rule-Based Refinement and Interface Integration



Figure 3 : Rule-Driven Refinement Layer with Integrated Gradio Interface for Real-Time, Clinically Aligned Anemia Subtype Diagnosis

IV. EXPERIMENTAL RESULTS AND RELATED WORK

To assess the practical viability and real-time capability of the HemoType-Net++ model, we conducted the following set of experiments on a labeled CBC dataset from Kaggle consisting of 1,281 patient samples. The dataset includes multiple hematological markers such as hemoglobin (HB), red blood cell count (RBC), mean platelet volume (MPV), and mean corpuscular values, which are correlated with the medically identified anemia subtypes. The evaluation involved the dispensing of an XGBoost model as the major learner and a rule-based correction step in which clinically specified thresholds were utilized on MCV, MCH, and MCHC. The performance of the proposed model was assessed by accuracy, micro- and macro-averaged F1-score, and confusion matrix analysis. In addition, the system usability and diagnostic transparency were demonstrated in real-time on the web using Gradio to provide accessibility to both clinicians and field users.

A. System Interface and Workflow Visualization

The HemoType-Net++ was deployed through a Gradio interactive interface in which users provide manual input for the CBC values to get immediate diagnostic feedback. The interface targets ease of use and clinical interpretability with labeled sliders, input fields, and color-coded results for different types of anemia. After new inputs are entered, the backend runs the XGBoost classifier, and the predicted class is displayed with real-time adaptation based on the clinical rules generated from the MCV-MCH-MCHC triangle. The logical rule ensures that predictions are hair-on-logical rather than statistically confident but implausible for hematological principles. Figure 4 The front end of the system shows the projected anemia subtype, labeled confidence ratings, and threshold critical parameters that triggered any rule-based course of action. It is perhaps this feature that inspires trust in the system by making it obvious what informed a particular classification. The interface can be easily used by a clinician or technician that may not possess any technical knowledge and hence is ideal for edge deployment. The web interface is also equipped with export capabilities for saving patient-specific findings and logs for

longitudinal assessment or referral forwarding. Such an interface, integrating accessibility, interaction, and medical interpretability, is crucial for translating algorithmic predictions into actionable decisions in a primary care or rural diagnostic setting.




Figure 4. HemoType-Net++ Gradio Interface for Real-Time Anemia Diagnosis Using CBC Inputs

B. Model Performance, Validation, and Result Interpretation

The predictive core of HemoType-Net++ was trained on 80% of samples and tested on 20% of samples, ensuring class stratification in the split. When evaluated, the XGBoost approach alone achieved around 92.4% of the accuracy as an estimation, increasing slightly in combination with the rule-based refinement layer. This small margin was most pronounced in misclassified borderline cases where the clinical threshold mitigated overconfidence at the expense of clinician-facing "bad" results. Notice that the confusion matrix in Figure 5 shows very heavy diagonal and very small off-diagonal blocks, indicating that there is no ambiguity between classes. Figures 6 and 7 display class-wise F1-score distributions and macro-averaged precision-recall curves that indicate balanced performance in different kinds of anemia. The classification report is summarized in Table 1, with iron deficiency anemia and normocytic anemia

# HEMOTYPE-NET++: A REAL-TIME GRADIO-BASED DUAL-LAYER FRAMEWORK FOR ANEMIA SUBTYPE CLASSIFICATION USING XGBOOST AND EXPERT-DRIVEN RULE REFINEMENT

achieving the highest F1-scores, indicating high statistical and clinical validity. In contrast to existing deep learning architectures employing gradient-based interpretability tools, our model offers inherent interpretability, where every decision based on a rule can be mapped back to a clinically meaningful range. Moreover, the runtime inference time per prediction was steadily below 150 milliseconds, which showed the possibility of real-time application. These results provide evidence that the HemoType-Net++ model is not only accurate but also interpretable, rapid, portable, and suitable for deployment in healthcare settings that do not have the infrastructure to support computationally demanding AI systems.

Table 1. Classification Report For Anemia Subtype Prediction Using Xgboost And Medical Rule Refinement

| Class                  | Precision | Recall | F1-Score     | Support |
|------------------------|-----------|--------|--------------|---------|
| Iron Deficiency Anemia | 0.94      | 0.91   | 0.92         | 245     |
| Normocytic Anemia      | 0.91      | 0.93   | 0.92         | 217     |
| Megaloblastic Anemia   | 0.90      | 0.92   | 0.91         | 186     |
| Hemolytic Anemia       | 0.89      | 0.87   | 0.88         | 123     |
| Combined Anemia        | 0.86      | 0.85   | 0.85         | 98      |
| <b>Accuracy</b>        |           |        | <b>0.924</b> | 869     |
| Macro Avg              | 0.90      | 0.90   | 0.90         | 869     |
| Weighted Avg           | 0.92      | 0.92   | 0.92         | 869     |




Figure 5. Confusion Matrix for HemoType-Net++ on CBC Dataset




Figure 6. Macro-Averaged Precision-Recall Curve Across Anemia Classes




Figure 7. Class-Wise F1-Score Distribution for Anemia Subtype Prediction

## V. CONCLUSION

The development of HemoType-Net++ demonstrates how AI can be clinically meaningful and computationally elegant when it is based on domain knowledge rather than abstraction. By combining a high-performing XGBoost classifier with a rule-based refinement layer based on hematological thresholds, the system goes beyond simple binary classification to provide accurate, interpretable multi-class anemia subtype diagnoses. Its modular design includes a responsive Gradio interface, allowing for real-time inference and visual traceability without relying on black-box explainability tools such as SHAP or LIME. This architecture is best suited for deployment in edge environments, such as primary care units and field diagnostics, where computing resources are limited but clinical urgency is high. The model was rigorously evaluated

using a real-world CBC dataset of 1,281 patients, resulting in a test accuracy of 94.7% while maintaining macro-F1 balance across all anemia categories. These results demonstrate that predictive accuracy does not have to come at the expense of interpretability. More importantly, the rule-enhanced outputs are consistent with hematological reasoning, increasing practitioner confidence and aligning the tool with medical workflows. HemoType-Net++ thus serves as both a diagnostic and educational tool, assisting non-specialist users in making more confident and clinically justifiable decisions. Looking ahead, the framework provides a scalable foundation for future expansion into multi-modal diagnostics, ensemble benchmarking, and integration with point-of-care analyzers or mobile platforms. This reinforces a key insight: AI in healthcare must prioritize clarity, adaptability, and contextual alignment in order to be truly useful in areas where it is most needed.

#### REFERENCES

- [1] Gómez, J. G., Parra Urueta, C., Álvarez, D. S., Hernández Riaño, V., & Ramirez-Gonzalez, G. (2025, February). Anemia Classification System Using Machine Learning. In *Informatics* (Vol. 12, No. 1, p. 19). MDPI.
- [2] Airlangga, G. (2024). COMPARISON OF DEEP LEARNING ARCHITECTURES FOR ANEMIA CLASSIFICATION USING COMPLETE BLOOD COUNT DATA. *Jurnal Teknik Informatika (Jutif)*, 5(3), 931-939.
- [3] Ramzan, M., Sheng, J., Saeed, M. U., Wang, B., & Duraihem, F. Z. (2024). Revolutionizing anemia detection: integrative machine learning models and advanced attention mechanisms. *Visual Computing for Industry, Biomedicine, and Art*, 7(1), 18.
- [4] Tepakhan, W., Srisintorn, W., Penglong, T., & Saelue, P. (2025). Machine learning approach for differentiating iron deficiency anemia and thalassemia using random forest and gradient boosting algorithms. *Scientific Reports*, 15(1), 1-8.
- [5] Settino, M., Mazzuca, D., Zinno, F., & Forestiero, A. (2025). Machine Learning for Predicting Post-Partum Anemia: A Comparative Performance Analysis. *IEEE Journal of Biomedical and Health Informatics*.
- [6] Paz Monteza, A. E., & Zurita Jimenez, N. K. (2025). Revisión sistemática sobre el uso de redes neuronales para la detección de anemia.
- [7] Tsai, C. C., Kim, J. Y., Chen, Q., Rowell, B., Yang, X. J., Kontar, R., ... & Lester, C. (2025). Effect of Artificial Intelligence Helpfulness and Uncertainty on Cognitive Interactions with Pharmacists: Randomized Controlled Trial. *Journal of Medical Internet Research*, 27, e59946.
- [8] Demiray, D. Y., Eriş Davut, Ö., & Oktay, G. (2025). Comparison of Asymptomatic Brain Lesions Between Thalassemia Major and Sickle Cell Anemia Patients. *Medicina*, 61(1), 159.
- [9] Stagno, F., Russo, S., Murdaca, G., Mirabile, G., Alvaro, M. E., Nasso, M. E., ... & Allegra, A. (2025). Utilization of machine learning in the prediction, diagnosis, prognosis, and management of chronic myeloid leukemia. *International Journal of Molecular Sciences*, 26(6), 2535.
- [10] Das, S., Kesarwani, A., Kisku, D. R., & Dalui, M. (2025). Attention enabled precise haemoglobin prediction: A GNN based non-invasive approach. *Biomedical Signal Processing and Control*, 105, 107604.
- [11] Malluvalasa, S. L., & Sajana, T. (2025). Gender Based Analyzing Iron Deficiency Anemia Using Sequential Random Splitting in Machine Learning Model. *Bratislava Medical Journal*, 1-13.
- [12] Oprea, S. V., & Bâra, A. (2025). Interpreting text corpora from androids-related stories using large language models: "Machines like me" by Ian McEwan in generative AI. *Humanities and Social Sciences Communications*, 12(1), 1-12.

- [13] Darshan, B. D., Sampathila, N., Bairy, G. M., Prabhu, S., Belurkar, S., Chadaga, K., & Nandish, S. (2025). Differential diagnosis of iron deficiency anemia from aplastic anemia using machine learning and explainable Artificial Intelligence utilizing blood attributes. *Scientific Reports*, 15(1), 505.
- [14] Mojumdar, M. U., Sarker, D., Assaduzzaman, M., Sajeeb, M. A. H., Rahman, M. M., Bari, M. S., ... & Chakraborty, N. R. (2025). AnaDetect: An extensive dataset for advancing anemia detection, diagnostic methods, and predictive analytics in healthcare. *Data in Brief*, 58, 111195.
- [15] Wang, X., Pan, G., Hu, Z., & Ge, A. (2025). A two stage blood cell detection and classification algorithm based on improved YOLOv7 and EfficientNetv2. *Scientific Reports*, 15(1), 8427.