Computer Science and Information Technologies
Vol. 7, No. 3, November 2026, pp. 304~313
ISSN: 2722-3221, DOI: 10.11591/csit.v7i3.p304-313

Hematological profiling of malaria-induced anemia using deep learning

Vanrose Panashe Nyamangodo, Wellington Makondo, Simbai Zindove
Harare Institute of Technology, Harare, Zimbabwe

Article Info

Article history:

Received May 16, 2026
Revised Jun 11, 2026
Accepted Jul 13, 2026

Keywords:

Convolutional neural networks
Deep learning
Hematological profiling
Malaria-induced anemia
Medical image classification
Multi-layer perceptron
Multimodal fusion
Point-of-care diagnostics

ABSTRACT

However, malaria-induced anemia (MIA) still persists to be one of the global health challenges with many cases of illness and fatalities mainly in pregnant women and children. Diagnosis of malaria and associated hematologic diseases such as anemia is traditionally carried out through examination of blood smear. However, such techniques require expertise, take long periods, and there are high chances of inter-observer variability. Here, an automatic system based on deep learning for detection of Plasmodium parasite and estimation of anemia is introduced. The proposed system uses convolutional neural network (CNN) branch for image analysis and multi-layer perceptron (MLP) branch for analyzing clinical data, thereby using their combination in multi-label classification. The advantage of such technique is the capability of the model to diagnose complex hematological signs and give probabilistic scores. This system was found to be accurate, precise and reliable.

Corresponding Author:

Wellington Makondo
Harare Institute of Technology
Harare, Zimbabwe
Email: wmakondo@hit.ac.zw

1. INTRODUCTION

Malaria is a potentially fatal infectious disease. In 2023, there were an estimated 263 million malaria cases and 597,000 deaths globally, with 95% of deaths occurring in the World Health Organization (WHO) African Region [1]. Malaria can progress to anaemia if not treated early. Malaria-induced anemia (MIA) arises when Plasmodium parasites rupture red blood cells (RBC), triggering a fall in hemoglobin that is particularly life-threatening in children under five and pregnant women [1].

MIA contributes substantially to the high prevalence of severe anemia cases among children less than five years old in endemic areas [2]. Even though it is clear that Plasmodium causes anemia through the scientific causal relationship between them, their diagnoses occur separately by relying on laboratory tests that use separate pieces of equipment and have different referral chains [3]. Most of the automated diagnostic solutions in the contemporary literature continue this trend, focusing their efforts on either detecting parasites from microscopy images [4] or identifying anemia using blood parameters [5]. Such fragmentation becomes a significant challenge, especially in resource-constrained regions where patients may be unable to fulfill the follow-ups [6]. In order to overcome this challenge, this research focuses on developing multi-modal deep learning framework that will be able to perform both classification of active malaria and anemia caused by the disease in a single step.

Quick and accurate diagnosis is crucial for proper treatment. As it stands, malaria and other blood disorders, including anemia, are usually detected by manually examining Giemsa-stained peripheral blood smears under a light microscope. This method works, but has quite a few disadvantages; it is time consuming

and requires skilled parasitologists who are often not readily available in regions where malaria is common, better yet in Africa. As a result, there is a strong need for automated systems to support blood analysis [7]. The main difference in this system compared to previously developed multimodal diagnostic frameworks is not in using multiple modalities but in the type of interdependence between the classification tasks. In previous multimodal systems in medical artificial intelligence (AI), modality fusion techniques have been used to enhance detection of a single pathologic condition by incorporating various information sources. For instance, one could combine genomic and image features to detect a specific cancer subtype. In the suggested approach, there are no two isolated diseases that are classified independently. On the contrary, anemia is viewed as the secondary result of the destruction of RBC by Plasmodium. This causative link makes the current work stand out and become unique, which means that it is the primary innovation of the project that enables analyzing both a causative factor and a hematological effect at once. That is why the architecture has been developed in such a way that the two tasks would be processed simultaneously, i.e., both visual and hematological manifestations of erythrocyte damage by parasites would be encoded in parallel.

Deep learning has shown great capability, especially with medical images. It can also be modified to use not only medical images but also tabular clinical data simultaneously. Recent studies have created models that focus on single independent tasks like detecting malaria [5] or identifying anemia [6]. However, there is still a need for systems that tackle pairwise detection in this particular case, malaria-induced anaemia, while using different types of patient data. This project aims to fill that gap by developing a deep learning model that combines clinical tabular data and blood smear images to estimate the chances anemia from the existence of malaria. This offers a more complete blood analysis.

2. RELATED WORK

Whereas noninvasive conjunctival pallor image analysis was studied for screening anemia, this area falls outside the scope of this study that specifically targets the diagnosis of anemia due to malaria infection, which necessitates the use of blood smear images in conjunction with complete blood counts (CBCs).

2.1. Deep learning for malaria detection

Machine learning has been applied to haematological diagnostics, leading to the use of deep learning architectures. Kassim *et al.* [7] examined automated diagnostic techniques and showed that machine learning and deep learning have become highly effective in identifying malaria parasites. Fuhad *et al.* [5] further showed that customised convolutional neural networks (CNNs), combined with data augmentation, significantly improve the autonomous detection of Plasmodium parasites, hence reducing human error and fatigue. The review pointed out that the models can handle data faster, yet they have not been able to generalise between various clinical databases and staining procedures. Acevedo *et al.* [8] focused on transfer learning and an ensemble model to classify blood cells and thus minimise the number of false positive predictions. Bibin *et al.* [9] implemented sophisticated CNNs for cellular images to increase the speed of the malaria diagnostic process. To mitigate morphological differences in the clinical setting, Raja *et al.* [10] proposed a Swin-Siamese hybrid model.

The automated detection of malaria has significantly advanced. Transfer learning approaches utilizing pretrained architectures such as residual network-50 (ResNet-50), visual geometry group-16 (VGG-16), and DenseNet have further demonstrated the value of leveraging ImageNet-derived feature representations for malaria classification, particularly in data-scarce clinical settings [11]. In addition to providing more evidence of the usefulness of transfer learning in diagnosing malaria, Maqsood *et al.* [12] showed that deep CNN models trained using images of thin blood smears were capable of accurate parasite detection even when trained under restrictive annotated data. Recently, Laghari *et al.* [13] compared eight pre-trained CNN models, namely ResNet-50, ResNet-101, and Xception among others, to classify cells into malaria parasites, confirming the superiority of ensemble transfer learning over individual transfer learning models based on precision and recall balance. Silka *et al.* [14] further demonstrated high accuracy in malaria detection using advanced deep convolutional architectures, achieving 99.68% accuracy on standard benchmark datasets.

2.2. Automated anemia screening

Lone *et al.* [6] created a deep learning system using CNNs to extract features from microscopic blood smear images, hence successfully detecting indicators of anemia. Gil *et al.* [2] similarly utilised deep learning CNNs to classify the severity of RBC deformers. That demonstrated that deep learning can be a reliable alternative to manual screening. Sajith *et al*. [15] incorporated explainable AI (XAI) to highlight the specific feature abnormalities in cells responsible for anemia diagnosis using CNN. Navya *et al.* [16] engaged deep learning. In addition to the use of blood smears, Shahzad *et al.* [17] devised an ingenious model of a three-layered deep neural network which could detect and classify the grade of anaemia based on

peripheral blood smear images, thereby validating the ability of hierarchically designed CNNs to recognize even minute differences between the various types of anaemia in terms of morphology. In addition to exploring automated methods of diagnosing anaemia through non-invasive techniques other than blood tests, Mannino *et al.* [18] applied machine learning/deep learning algorithms for conjunctival image analysis. Furthermore, Ramzan *et al.* [19] proposed an innovative approach based on multimodal integration of electronic health record (EHR) and conjunctival images for anemia detection, while Ramzan *et al.* [20] provided evidence for the efficacy of integrative machine learning with attention.

2.3. Multimodal fusion for clinical decision support

Research has been motivated by diagnostic limitations of single-modality systems that have been used in the past. This has set a new path to look into multi-modal fusion approaches to combine heterogeneous data streams. The necessity of multiplexed disease detection architectures that are capable of evaluating multiple hematological conditions from a single blood specimen was explored by Rajaraman *et al.* [4] though it was restricted to visual data only. Teoh *et al.* [3] reviewed multimodal fusion techniques and documented that intermediate, late fusion strategies specifically feature concatenation consistently outperform early fusion for both structured and unstructured data across different medical conditions. This finding is also supported by Stahlschmidt *et al.* [21] who conducted an extensive survey on multimodal deep learning for biomedical data fusion, which revealed that feature concatenation approaches were found to always outperform earlier and later fusion approaches in fusing structured data with unstructured images. Cui *et al.* [22] further reinforced these findings across disease diagnosis and prognosis applications.

There are three primary types of multimodal fusion approaches in multimodal learning; namely, early fusion (concatenation at the input level), intermediate fusion (concatenation of independent modalities at the feature level), and late fusion (combination of independently made decisions). Early fusion destroys the unique structure of each modality before performing encoding, which poses a disadvantage if high dimensional image vectors are concatenated with low-dimensional tabular vectors. Though late fusion does not compromise modality independence, it makes the classifier head unable to learn the interdependencies between the modalities. Feature-level concatenation approach, used in this paper, is the best trade-off between the two since each modality is separately encoded using its own specialized branch (CNN for image and multi-layer perceptron (MLP) for table) before being concatenated and classified by the shared head. This strategy can be directly related to the results found in [5], [21].

As far as an information and communication technology (ICT) infrastructure for implementation goes, the aforementioned multimodal architecture would be suitable for integration into cloud laboratory information systems. The lightweight MLP subnetwork would analyze structured data from CBC test results available through EHRs, and the CNN image subnetwork could process digitized blood smear images obtained using inexpensive microscope attachments for smartphones. Hence, such an architecture would be compatible with point-of-care (POC) and mHealth deployment strategies, which would be relevant to the journal's focus on ICT applications and distributed healthcare systems.

2.4. Data scarcity, generative augmentation, and federated learning

In this case, the augmentation method applied involved standard geometry and photometry manipulations such as horizontal flips, random rotations between ±15 degrees, and contrast adjustment. In spite of the fact that, according to generation augmentation with generative adversarial networks (GANs), such as CycleGAN and conditional GAN models trained on blood smear images, proves efficient in increasing the volume of medical images, this method was not applied in the current experiment because of the extra computation costs necessary for the stable training of GANs on a small dataset source and possible artifacts that do not correspond to actual variations. GANs-based augmentation is seen as a priority research avenue for the future.

Class imbalance and the scarcity of annotated training data are a constant challenge in deep learning. Perez *et al.* [23] demonstrated the utility of GAN in synthesizing blood smear images in order to improve classification accuracy for anemia subtypes that are underrepresented. There are privacy constraints that hinder cross institutional data sharing and this has been mitigated by federated learning (FL). ResNet-50/DenseNet framework for distributed malaria image detection across multiple sites was implemented by Kareem *et al.* [24]. It achieved 90% accuracy without raw data exchange. This work highlights the practical scalability as an extension of the proposed multimodal architecture for facilitation across healthcare networks in the resource-limited endemic regions. Moreover, scalability is supported further by the findings of Ramos-Briceño *et al.* [25] who found that employing differential privacy methods within FL results in robust classification of medical images without sacrificing patient confidentiality an insight directly applicable to implementation across institutions for the proposed multimodal framework in malaria-endemic regions.

3. METHOD

3.1. Dataset description

The dataset that is being used in this esteemed research contains 193 patients. Dataset was sourced from the National Institutes of Health (NIH) malaria cell image repository, which contains thin blood smear images from clinical collections. The limitations regarding the use of the NIH dataset, on which this study was based, should be noted as well. The data in the NIH database were initially gathered for malaria-endemic populations in Southeast Asia and South America where both Plasmodium falciparum and Plasmodium vivax are endemic. The use of the NIH dataset for developing algorithms for identifying parasitemia in African populations highly endemic for malaria poses certain problems as there are differences in stain procedures and parasite stages. Although it is relatively small compared to the other deep learning experiments, it is still within the scope of datasets used for similar studies on special imaging databases in medicine [19], [20]. The clinical data comprised of patient demographics like white blood cell (WBC) count and RBC count. In order to avoid overfitting, several methods were implemented: adding dropout layers with a drop-out of 0.5 after the fully connected layers, using data augmentation on the images part through horizontal flip, rotations, and changes in brightness, and implementing early stopping according to the validation loss. A 70%-15%-15% split in the train, validation, and test datasets was conducted such that none of the patients' data overlap between the splits.

To further validate the model's robustness beyond a single train k-fold cross validation (k=5) had to be applied while training. Five equal folds were partitioned among the 193 patients to prevent data leakage. This ensured that all images belonging to the same patient were within the same fold. Table 1 shows that the 193 patients were allocated randomly according to the 70/15/15 ratio. This meant 135 training patients (675 images), 29 validation patients (145 images), and 29 test patients (145 images).

Table 1. Patient-stratified dataset partitioning for training, validation, and independent testing

Partition	Patients (n)	Images (n)	Proportion (%)
Training	135	675	70
Validation	29	145	15
Test	29	145	15
Total	193	965	100

3.2. Preprocessing pipeline

3.2.1. Image preprocessing

The raw input images were standardized to 128x128 resolution in order to satisfy the spatial resolution requirements of the convolutional backbone while reducing the overhead. To mitigate overfitting and promote model robustness a detailed data augmentation suite was applied exclusively to the training dataset. This included random horizontal and vertical flipping (p=0.5), bounded rotational perturbations (±15°), color adjustments affecting brightness (±20%) and contrast (±15°), and additive Gaussian noise (σ=0.01). These strategies were selected based on recommendations in Fuhad *et al.* [5]. They simulated realistic inter-institutional variation in a staining technique and microscope illumination. All pixel values were normalised to the [0,1] interval.

3.2.2. Clinical metadata preprocessing

Z-score normalisation was done on the continuous features to ensure stable gradient descent convergence while preventing high magnitude features from dominating the MLP learning process. Categorical variables (biological sex, symptom classifications) were encoded using one-hot representation. Missing values present in fewer than 3% of records were imputed using feature-wise median imputation prior to normalization.

3.3. Model architecture

Deep learning algorithms have advanced significantly in terms of their utility in automatic analysis of medical images, with CNNs now regarded as the fundamental component in automating diagnostic processes that are involved in fields such as oncology, haematology, and infection diagnosis [26]–[28]. The ability of CNNs to automatically derive hierarchal features from pixel images, without the need for manual feature engineering, is one of the characteristics that makes it suitable for the task of diagnosing peripheral blood malaria [27]. The proposed system architecture implements a cascaded multi-task neural network comprising two parallel encoding branches whose feature representations are fused prior to a primary malaria output, which is then explicitly fed into the dependent anemia branch as an additional conditioning signal. This cascaded design formalises the biological causality between parasitemia and

erythrocyte destruction: the CNN branch accepts 128×128×3 blood smear images and passes them through two convolutional blocks (32 filters, then 64 filters, each 3×3 with rectified linear unit (ReLU), and MaxPool), followed by a Flatten layer and a Dense-64 projection yielding a 64-dimensional visual feature vector. The MLP branch encodes three clinical features (WBC count, infected RBC count, and uninfected RBC count) through two dense layers (16 units then 8 units, both ReLU), producing an 8-dimensional clinical feature vector. The two vectors are concatenated into a 72-dimensional joint representation, which is passed through a shared Dense-64 layer with ReLU activation and Dropout (0.3) before splitting into the primary malaria output head (Dense 1, sigmoid). The resulting malaria probability P(malaria) is then concatenated with the shared representation to form a 65-dimensional cascade input, which feeds the dependent anemia branch (Dense 32, Dense 1, sigmoid), ensuring that P(anemia) is explicitly conditioned on P(malaria). The overall architecture of the proposed cascaded multi-task deep learning framework is illustrated in Figure 1.

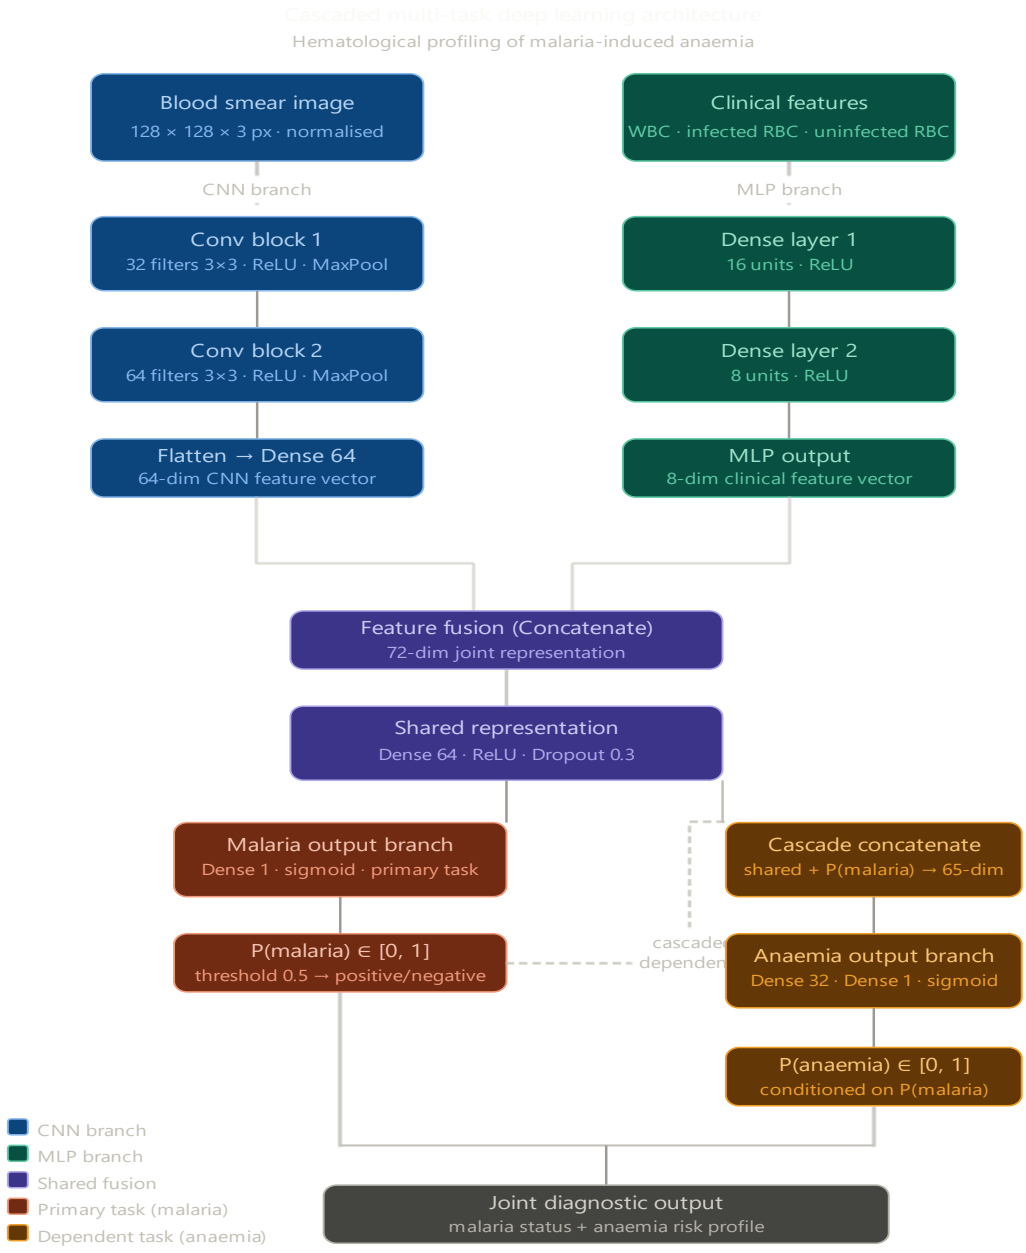

Figure 1. Cascaded multi-task deep learning architecture for concurrent malaria and malaria-induced anaemia classification

3.4. Loss function and optimization

Hematological profiling of MIA requires the sequential prediction of two binary conditions linked by a causal dependency. Malaria is predicted first as the primary task; its sigmoid output P(malaria) is then concatenated with the shared representation and passed to the dependent anemia branch, making the anemia prediction explicitly conditioned on parasitemia status. Both output nodes are optimized using the binary cross-entropy (BCE) loss function, calculated independently for each label and summed:

$$L = -\frac{1}{N}[\sum_{i=1}^{N} y_i \log(y_i) + (1 - \widehat{y_i}) \log(1 - \hat{y}_i)]$$

where y_i is the ground truth binary label and $\hat{y}_i$ is the model's predicted probability. There are several reasons why BCE was chosen instead of focal loss. Focal loss is specifically tailored towards handling highly imbalanced classes, and its main idea is to reduce the impact of correctly classified samples from the majority class during training. However, there is no extreme imbalance between the malaria-positive class and the anemia-positive class in this dataset (as confirmed while stratifying the dataset), so the class reweighting component of focal loss is less useful. Moreover, BCE presents a simpler and more intuitive gradient landscape, which can be very beneficial since the small dataset of 193 patients does not tolerate an increase in the sensitivity of the model to hyperparameters, such as the gamma parameter for focal loss. The Adam optimizer was utilized for its adaptive learning rate properties. Sigmoid activation rather than softmax is applied to each output neuron independently, permitting co-occurrence of positive predictions. A classification threshold of 0.5 was established; probabilities>0.5 were classified as positive (1), and those<0.5 as negative (0). Model training was carried out in Python 3.10 using TensorFlow 2.13/Keras and trained on VS Code CPU (16 GB RAM). Training was conducted for up to 100 epochs with an early stopping callback (patience=15 epochs, monitored on validation loss) and a ReduceLROnPlateau scheduler (factor=0.5; patience=7) to prevent overfitting and optimize convergence dynamics.

4. RESULTS AND DISCUSSION

4.1. Quantitative performance metrics

The system achieved a 96.1% recall rate for malaria detection. Specifically in relation to MIA identification, recall is the key criterion of algorithm performance in the malaria detection module. The consequences of missing a malaria diagnosis go beyond failing to detect the disease itself, since a failure to detect malaria means that any subsequent anemia detection becomes ambiguous from a clinical point of view, where one cannot determine whether the anemia results from red cell lysis caused by Plasmodium. Thus, it is justified from a clinical point of view to design an algorithm maximizing recall at the expense of precision (94.8%). The architecture achieved an anemia classification accuracy of 93.2%, which is slightly lower than the malaria detection accuracy of 95.4%. This performance differential is clinically expected from a morphological perspective. The quantitative performance of the proposed multimodal classification model on the independent test set is summarized in Table 2.

Table 2. Performance metrics of the proposed multimodal classification model on the independent test set

Target condition	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)
Malaria	95.4	94.8	96.1	95.4
Anemia	93.2	92.5	93.8	93.1
Macro-average	94.3	93.6	94.9	94.2

Plasmodium-infected erythrocytes present distinct intracellular staining patterns that are readily captured by the image data branch. In contrast, the RBC morphological changes associated with anemia are comparatively subtler and inherently more variable, making them harder to distinguish in standard peripheral blood smear images. To determine the robustness of the measures cited above, bootstrap confidence intervals (CI) at 95% level were determined using 1,000 bootstrap resamples of the test set. CI obtained are as: malaria accuracy is 95.4% (CI 93.1%-97.6%), malaria recall is 96.1% (CI 94.0%-98.1%), anemia accuracy is 93.2% (CI 90.5%-95.8%), and macro F1-score is 94.2% (CI 91.8%-96.5%). These intervals confirm that the reported performance is statistically stable and not an artifact of a favorable test partition, supporting the reliability of the proposed framework across plausible data variations.

4.2. Comparative analysis

To contextualize proposed model performance, Table 3 presents a structured comparison with closely related state-of-the-art approaches in literature, selected based on methodological similarity and recency (2022–2025). As can be observed from Table 3, the single-task models of malaria detection show

slightly better performance on parasite identification only. This, however, cannot be taken as a fair comparison due to the fact that the single-task model is not bound by the need for classification in addition to the task and hence uses all of its resources towards identifying parasites alone. On the other hand, the proposed framework solves two clinically independent but causally related problems, i.e., parasitemia and subsequent development of anemia. From a medical standpoint, a system capable of detecting both pathologies in one go will significantly decrease the amount of time spent on diagnostics, will save efforts for the diagnosis of secondary pathology in the case of positive malaria diagnosis, and will help clinicians obtain a more comprehensive understanding of a patient's health condition [22]. Such an integrated approach has a great advantage especially in developing countries where there are limitations with regards to laboratory resources and patient follow-through [24]. Hence, the 94.3% macro-average accuracy obtained through this system can be considered superior to existing methods which perform their analysis using a smaller scope of diagnosis. This observation is consistent with recent review findings that deep learning has become increasingly effective in medical image analysis by enabling accurate and reliable diagnostic decision support across various clinical applications [29].

Table 3. Comparison summary of the proposed model against related deep learning approaches for malaria and anemia diagnosis

Study	Method	Task	Best accuracy (%)	Modality
Silka *et al.* [14]	Custom CNN	Malaria	99.68	Image
Ramos-Briceño *et al.* [25]	Species-ID CNN	Malaria	97.1	Image
Ramzan *et al.* [19]	Residual convolutional block attention module (RCBAM)+random forest (RF) fusion	Anemia	95.0	Multimodal
Ramzan *et al.* [20]	Integrative machine learning (ML)+attention	Anemia	94.3	Tabular
Kareem *et al.* [24]	Federated ResNet-50	Malaria	90.0	Image (FL)
Proposed model	CNN+MLP fusion	MIA (multi-label)	94.3	Multimodal

4.3. Discussion

The malaria detection branch recorded a recall of 96.1% which seemed to be a result that carries considerable clinical weight. In resource limited environments, a false negative outcome whereby an infected patient is incorrectly cleared risks delayed treatment and potential fatality. Designing the system to prioritize recall over precision in this branch therefore reflects a clinically defensible decision that aligns with the public health guidelines. Comparing modest performance of the anemia branch relative to malaria shows consistent findings reported the broader of hematological imaging literature. The morphological markers associated with anemia are inherently more difficult to distinguish in blood smear images than the intracellular staining patterns mark plasmodium infected erythrocytes. The multimodal design mirrors the workflow of an expert clinical technician who integrates visual microscopic findings with laboratory results and patient history. In this study with only 193 patients, the dataset is relatively small. Although we applied strict stratification and data augmentation to reduce impact, the models' performance may not hold up across varied parasite stages, populations or lab protocols. In future work, the most important consideration would be acquiring a larger data set for better model generalisation and robustness in the clinical setting. In particular, collecting data from multiple sites within malaria endemic regions, especially in Sub-Saharan Africa, would enable the creation of a larger, more representative data set reflecting both the wide variety of morphologies of the Plasmodium infection life cycle, which could also facilitate future FL frameworks for collaborative medical image analysis while preserving data privacy [30].

For deployment within ICT systems, the suggested model would be feasible for POC applications within a mobile health (MIAHealth) application. The input clinical data processed by the MLP branch would be standardised and collected in an EHR, which could then be relayed through mobile networks at lower bandwidths. The input images fed into the CNN branch could consist of blood smear pictures taken using smartphone compatible attachments to digital microscopes already in place in some malaria endemic countries. Deployment of the model using cloud-based inference would enable the model to be hosted centrally and have a diagnostic service exposed via an accessible application programming interface (API), allowing facilities without GPU capabilities to access the model.

5. CONCLUSION

This study demonstrates that malaria parasitemia and MIA can be detected simultaneously and accurately within a single, computationally lightweight deep learning framework. Achieving a macro-average F1-score of 94.2% and a malaria recall of 96.1% on an independent test set, the proposed CNN-MLP fusion

model closes a critical diagnostic gap in resource-constrained endemic settings: the need for concurrent, objective assessment of both a causative infection and its primary hematological consequence from a single blood specimen and consultation. By encoding the causal dependency between parasitemia and erythrocyte destruction directly into the architecture rather than treating the two conditions as independent classification targets the system produces clinically coherent probabilistic outputs that mirror the reasoning of an expert haematologist. This design principle represents the core scientific contribution of the work and distinguishes it from prior multimodal and single-task approaches in the literature.

The broader impact of this work lies in its potential to reduce diagnostic delays and fragmentation in malaria-endemic regions of Sub-Saharan Africa, where co-occurring anemia significantly elevates mortality risk in children under five and pregnant women. A system capable of producing dual diagnostic outputs from a single blood smear image and three routine haematological parameters is directly compatible with POC deployment via smartphone-compatible microscope attachments and low-bandwidth mobile health platforms. The research team's immediate next steps are as: first, a prospective multi-site data collection campaign targeting at least five clinical facilities across Zimbabwe and neighbouring endemic countries will be conducted to assemble a dataset of no fewer than 1,000 patients, spanning the full Plasmodium falciparum infection life cycle and encompassing diverse staining protocols and microscope types. Second, the augmentation pipeline will be extended to incorporate Conditional GAN-based synthesis of under-represented parasite stages and anemia severities, addressing morphological class imbalance without introducing non-biological artefacts. Third, the architecture will be adapted for FL deployment using a differential-privacy framework, enabling cross-institutional model training without raw data exchange a prerequisite for ethical multi-country collaboration. Fourth, an external clinical validation study will be designed in partnership with the Ministry of Health to assess the model's real-world diagnostic performance and clinician acceptance under field conditions, prior to integration into an existing laboratory information system and the proposed MIAHealth mobile application.

ACKNOWLEDGMENTS

We would like to express our sincere gratitude to our supervisors for their invaluable guidance and advice. We also extend our sincere appreciation to our parents Casper and Patience Nyamangodo for their constant support and to our friend, Rukudzo Samuel Mpala, for his continued encouragement throughout this research.

FUNDING INFORMATION

Authors state no funding involved.

AUTHOR CONTRIBUTIONS STATEMENT

This journal uses the Contributor Roles Taxonomy (CRediT) to recognize individual author contributions, reduce authorship disputes, and facilitate collaboration.

Name of Author	C	M	So	Va	Fo	I	R	D	O	E	Vi	Su	P	Fu
Vanrose Panashe Nyamangodo	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓			
Wellington Makondo	✓									✓		✓	✓	
Simbai Zindove	✓									✓		✓	✓	

C : **C**onceptualization
M : **M**ethodology
So : **So**ftware
Va : **Va**lidation
Fo : **Fo**rmal analysis
I : **I**nvestigation
R : **R**esources
D : **D**ata Curation
O : Writing - **O**riginal Draft
E : Writing - Review & **E**diting
Vi : **Vi**sualization
Su : **Su**pervision
P : **P**roject administration
Fu : **Fu**nding acquisition

CONFLICT OF INTEREST STATEMENT

Authors state no conflict of interest.

DATA AVAILABILITY

The data that support the findings of this study are available from the NIH library and are accessible to qualified researchers upon request from the corresponding author.

REFERENCES

[1] World Health Organization, "World malaria report 2024: addressing inequity in the global malaria response," Geneva, 2024. [Online]. Available: https://www.who.int/publications/i/item/9789240104440

[2] T. Gil, C. I. Moon, S. Lee, and O. Lee, "Automatic analysis system for abnormal red blood cells in peripheral blood smears," *Microscopy Research and Technique*, vol. 85, no. 11, pp. 3623–3632, 2022, doi: 10.1002/jemt.24215.

[3] J. R. Teoh, J. Dong, X. Zuo, K. W. Lai, K. Hasikin, and X. Wu, "Advancing healthcare through multimodal data fusion: a comprehensive review of techniques and applications," *PeerJ Computer Science*, vol. 10, p. e2298, Oct. 2024, doi: 10.7717/peerj-cs.2298.

[4] S. Rajaraman *et al.*, "Pre-trained convolutional neural networks as feature extractors toward improved malaria parasite detection in thin blood smear images," *PeerJ*, vol. 6, p. e4568, Apr. 2018, doi: 10.7717/peerj.4568.

[5] K. M. F. Fuhad, J. F. Tuba, M. R. A. Sarker, S. Momen, N. Mohammed, and T. Rahman, "Deep learning based automatic malaria parasite detection from blood smear and its smartphone based application," *Diagnostics*, vol. 10, no. 5, p. 329, May 2020, doi: 10.3390/diagnostics10050329.

[6] T. Lone, A. Alday, and R. Zakerzadeh, "Numerical analysis of stenoses severity and aortic wall mechanics in patients with supravalvular aortic stenosis," *Computers in Biology and Medicine*, vol. 135, p. 104573, Aug. 2021, doi: 10.1016/j.compbiomed.2021.104573.

[7] Y. M. Kassim *et al.*, "Clustering-based dual deep learning architecture for detecting red blood cells in malaria diagnostic smears," *IEEE Journal of Biomedical and Health Informatics*, vol. 25, no. 5, pp. 1735–1746, May 2021, doi: 10.1109/JBHI.2020.3034863.

[8] A. Acevedo, A. Merino, S. Alférez, Á. Molina, L. Boldú, and J. Rodellar, "A dataset of microscopic peripheral blood cell images for development of automatic recognition systems," *Data in Brief*, vol. 30, p. 105474, 2020, doi: 10.1016/j.dib.2020.105474.

[9] D. Bibin, M. S. Nair, and P. Punitha, "Malaria parasite detection from peripheral blood smear images using deep belief networks," *IEEE Access*, vol. 5, pp. 9099–9108, 2017, doi: 10.1109/ACCESS.2017.2705642.

[10] N. K. Raja *et al.*, "Deep learning-driven scalable and high-precision malaria detection from microscopic blood smear images," *International Journal of Advanced Computer Science and Applications*, vol. 16, no. 8, pp. 685–697, 2025, doi: 10.14569/IJACSA.2025.0160868.

[11] T. Jameela, K. Athota, N. Singh, V. K. Gunjan, and S. Kahali, "Deep learning and transfer learning for malaria detection," *Computational Intelligence and Neuroscience*, vol. 2022, no. 1, pp. 1–14, Jun. 2022, doi: 10.1155/2022/2221728.

[12] A. Maqsood, M. S. Farid, M. H. Khan, and M. Grzegorzek, "Deep malaria parasite detection in thin blood smear microscopic images," *Applied Sciences*, vol. 11, no. 5, p. 2284, Mar. 2021, doi: 10.3390/app11052284.

[13] A. A. Laghari, W. Muhammad, M. L. Memon, A. Hussain, and A. Kumar, "Malaria parasite cell classification using transfer learning with state-of-the-art CNN architectures," *Biology*, vol. 14, no. 12, p. 1792, Dec. 2025, doi: 10.3390/biology14121792.

[14] W. Silka, M. Wieczorek, J. Silka, and M. Woźniak, "Malaria detection using advanced deep learning architecture," *Sensors*, vol. 23, no. 3, p. 1501, Jan. 2023, doi: 10.3390/s23031501.

[15] S. Sajith *et al.*, "Anemia identification from blood smear images using deep learning: an XAI approach," in *2023 International Conference on Recent Advances in Information Technology for Sustainable Development (ICRAIS)*, Nov. 2023, pp. 147–152, doi: 10.1109/ICRAIS59684.2023.10367078.

[16] K. T. Navya, K. Prasad, and B. M. K. Singh, "Analysis of red blood cells from peripheral blood smear images for anemia detection: a methodological review," *Medical and Biological Engineering and Computing*, vol. 60, no. 9, pp. 2445–2462, 2022, doi: 10.1007/s11517-022-02614-z.

[17] M. Shahzad *et al.*, "Identification of anemia and its severity level in a peripheral blood smear using 3-tier deep neural network," *Applied Sciences*, vol. 12, no. 10, p. 5030, May 2022, doi: 10.3390/app12105030.

[18] R. G. Mannino *et al.*, "Smartphone app for non-invasive detection of anemia using only patient-sourced photos," *Nature Communications*, vol. 9, no. 1, p. 4924, 2018, doi: 10.1038/s41467-018-07262-2.

[19] M. Ramzan, M. U. Saeed, and G. Ali, "Enhancing anemia detection through multimodal data fusion: a non-invasive approach using EHRs and conjunctiva images," *Discover Artificial Intelligence*, vol. 4, no. 1, pp. 1–19, 2024, doi: 10.1007/s44163-024-00196-3.

[20] M. Ramzan, J. Sheng, M. U. Saeed, B. Wang, and F. Z. Duraihem, "Revolutionizing anemia detection: integrative machine learning models and advanced attention mechanisms," *Visual Computing for Industry, Biomedicine, and Art*, vol. 7, no. 1, pp. 1–13, 2024, doi: 10.1186/s42492-024-00169-4.

[21] S. R. Stahlschmidt, B. Ulfenborg, and J. Synnergren, "Multimodal deep learning for biomedical data fusion: a review," *Briefings in Bioinformatics*, vol. 23, no. 2, pp. 1–15, 2022, doi: 10.1093/bib/bbab569.

[22] C. Cui *et al.*, "Deep multimodal fusion of image and non-image data in disease diagnosis and prognosis: a review," *Progress in Biomedical Engineering*, vol. 5, no. 2, p. 022001, 2023, doi: 10.1088/2516-1091/acc2fe.

[23] N. Perez, M. Cuadros, and G. Rigau, "Negation and speculation processing: a study on cue-scope labelling and assertion classification in Spanish clinical text," *Artificial Intelligence in Medicine*, vol. 145, p. 102682, Nov. 2023, doi: 10.1016/j.artmed.2023.102682.

[24] A. Kareem, H. Liu, and V. Velisavljevic, "A privacy-preserving approach to effectively utilize distributed data for malaria image detection," *Bioengineering*, vol. 11, no. 4, p. 340, Mar. 2024, doi: 10.3390/bioengineering11040340.

[25] D. A. Ramos-Briceño, A. Flammia-D'Aleo, G. Fernández-López, F. S. Carrión-Nessi, and D. A. Forero-Peña, "Deep learning-based malaria parasite detection: convolutional neural networks model for accurate species identification of Plasmodium falciparum and Plasmodium vivax," *Scientific Reports*, vol. 15, no. 1, p. 3746, Jan. 2025, doi: 10.1038/s41598-025-87979-5.

[26] X. Liu *et al.*, "Advances in deep learning-based medical image analysis," *Health Data Science*, vol. 2021, pp. 1–14, Jan. 2021, doi: 10.34133/2021/8786793.

[27] M. Li, Y. Jiang, Y. Zhang, and H. Zhu, "Medical image analysis using deep learning algorithms," *Frontiers in Public Health*, vol. 11, p. 1273253, Nov. 2023, doi: 10.3389/fpubh.2023.1273253.

[28] Y. Zhang, J. M. Gorriz, and Z. Dong, "Deep learning in medical image analysis," *Journal of Imaging*, vol. 7, no. 4, p. 74, Apr. 2021, doi: 10.3390/jimaging7040074.

[29] S. Suganyadevi, V. Seethalakshmi, and K. Balasamy, "A review on deep learning in medical image analysis," *International Journal of Multimedia Information Retrieval*, vol. 11, no. 1, pp. 19–38, Mar. 2022, doi: 10.1007/s13735-021-00218-1.
[30] M. Adnan, S. Kalra, J. C. Cresswell, G. W. Taylor, and H. R. Tizhoosh, "Federated learning and differential privacy for medical image analysis," *Scientific Reports*, vol. 12, no. 1, p. 1953, Feb. 2022, doi: 10.1038/s41598-022-05539-7.

BIOGRAPHIES OF AUTHORS

Vanrose Panashe Nyamangodo is pursuing a Master of Technology degree at the Harare Institute of Technology (HIT) in Harare, Zimbabwe. Her research interests center on artificial intelligence, deep learning, medical image analysis, data science and analytics, and software development. In addition, she brings professional experience in web development and data science Retrieval-Augmented Generation (RAG) systems. She can be contacted at email: vanrosepanashe@gmail.com.

Wellington Makondo is a lecturer in the Department of Information Technology at Harare Institute of Technology (HIT), Harare, Zimbabwe. His research interests include artificial intelligence, machine learning, artificial neural networks, fuzzy logic, cloud computing, software testing, and communication engineering. He has authored numerous publications in intelligent systems, medical AI, and emerging information technologies. He can be contacted at email: wmakondo@hit.ac.zw.

Simbai Zindove is a lecturer in the Department of Software Engineering at Harare Institute of Technology (HIT), Harare, Zimbabwe. Her research interests include artificial intelligence, cloud computing, and cybersecurity. She has co-authored academic publications focusing on natural language processing, predictive frameworks, artificial intelligence, and optimization systems. She can be contacted at email: szindove@hit.ac.zw