OPEN ACCESS

Volume: 4

Issue: 4

Month: December

Year: 2025

ISSN: 2583-7117

Published: 26.12.2025

Citation:

Neelkamal Vishwakarma, Gunjan Kumar Mishra, Gaurav Chaubey "Advancing Hematological Diagnostics in Resource-Constrained Settings: A Robust Deep Learning Solution for Sickle Cell Anemia Screening" International Journal of Innovations in Science Engineering and Management, vol. 4, no. 4, 2025, pp. 110-118.

DOI:

10.69968/ijisem.2025v4i4110-118

Advancing Hematological Diagnostics in Resource-Constrained Settings: A Robust Deep Learning Solution for Sickle Cell Anemia Screening

Neelkamal Vishwakarma[1], Gunjan Kumar Mishra[2], Gaurav Chaubey[2]

[1]Research Scholar, Department of Electronics & Communication, SHEAT Babatpur, Varanasi
[2]Assistant Professor, Department of Electronics & Communication, SHEAT Babatpur, Varanasi

Abstract

This study introduces a deep learning model for the identification of sickle cell anaemia (SCA) in red blood cell (RBC) images, with the objective of facilitating efficient diagnostics in resource-limited environments. The study used a DenseNet121-based convolutional neural network (CNN) with a custom classification head and a unique dataset of 691 images from Uganda's Teso region that were processed with Field and Leichman stains. The strong method included using RandomUnderSampler and SMOTE to fix class imbalance, as well as Albumentations to add more data to make the model more resistant to image variability, like differences in staining and blurry images. The model was trained with a 4-fold cross-validation method. In just nine epochs, it reached a peak training accuracy of 96.4% and a loss of 4.0153. This better efficiency and performance beat work that used an InceptionV3 model (91% accuracy, 28.7 loss, 100 epochs). The model's high accuracy and robustness suggest that it could be used in clinical settings in areas where SCA is common, which would be a big step forward for automated haematological diagnostics.

Keywords; Sickle Cell Anemia, Machine Learning, Data Integration, Predictive Modeling, Early Diagnosis.

INTRODUCTION

Sickle cell disease (SCD), the most common type of which is sickle cell anaemia (SCA), is a group of inherited blood disorders that are caused by abnormal haemoglobin. SCA is marked by a deficiency of haemoglobin, the protein that carries oxygen in red blood cells (RBCs) [1]. In some situations, RBCs take on a stiff, sickle-like shape that doesn't change. Because of this strange shape, the cells can't change shape as they go through small blood vessels (capillaries). This can cause serious blockages and chronic pain episodes, which are called pain crises. SCA symptoms often appear between the ages of 5 and 6 months and can cause serious health problems like bacterial infections, stroke, chronic pain, and damage to important organs like the liver, heart, kidneys, and eyes. People with SCD rarely live to adulthood without the right treatment, even though good medical care can greatly increase life expectancy [2].

SCD is a big problem for public health around the world, affecting millions of people. SCD affected about 7.7 million people around the world in 2021. The condition is disproportionately prevalent in regions where malaria is common, particularly Sub-Saharan Africa, where an estimated 80% of affected individuals reside [3]. It is also less prevalent among individuals of African descent living overseas, as well as in certain regions of India, Southern Europe, West Asia, and North Africa. Because it is so common and can cause serious health problems, early diagnosis and treatment are very important. Finding it early can improve your quality of life and help you avoid serious complications like stroke and organ damage.

Challenges in Current Sickle Cell Diagnosis

Blood tests, haemoglobin electrophoresis, and high-performance liquid chromatography (HPLC) are some of the current ways to diagnose SCD. These tests look for genes that don't work right or haemoglobin that isn't normal [4]. These methods work well in controlled settings, but they often need specialised lab equipment and skilled technicians. This makes them hard to scale up and use in places with few resources, like rural hospitals in Sub-Saharan Africa. The common method of manually looking at peripheral blood smears under a microscope, where technicians look for sickle-shaped RBCs, is also very labour-intensive, takes a lot of time, and is very prone to human error and subjectivity. Different levels of technician skill, staining quality (like Field and Leichman stains), and image clarity can all lead to different diagnostic results. These limitations underscore a significant deficiency: the necessity for a swift, reliable, and computationally efficient diagnostic instrument that can be effectively utilised in clinical environments with constrained resources. Even though more and more countries are trying to set up automated newborn screening programs, many areas, especially in Sub-Saharan Africa, still don't have national screening systems [5].

The Role of Deep Learning in Automated Diagnostics

The progress made in artificial intelligence (AI), especially machine learning (ML) and deep learning (DL), provides a strong way to get around the limits of traditional diagnostics [6]. Deep learning models, like Convolutional Neural Networks (CNNs), are very good at automatically analysing complicated visual data, finding complex patterns, and doing classification tasks with a lot of accuracy. Because of this ability, CNNs are very good at looking at microscopic blood smear images to find small changes in shape, like the crescent shape that shows sickle cell anaemia. Previous studies have looked into using ML/DL methods for haematological diagnostics and found some very promising results. Research employing diverse classification algorithms, including CNNs, has demonstrated efficacy in differentiating between healthy and diseased blood cells, with certain models attaining accuracy rates surpassing 90%. Nonetheless, numerous current automated systems encounter difficulties related to robustness—the capacity to generalise across varied imaging conditions, staining variability, and ambiguous images—and computational efficiency, both of which are essential for practical clinical implementation [7].

RESEARCH FOCUS AND OBJECTIVES

This study addresses the necessity for an accurate and efficient automated diagnostic tool by developing a deep learning model for the classification of sickle cell anaemia in red blood cell imagery. The research focusses on creating a robust solution capable of operating effectively despite the variability in images and resource constraints typical of high-prevalence areas. This study utilises a unique dataset consisting of 691 microscopic images obtained from the Teso region of Uganda, processed with two distinct staining techniques (Field and Leichman stains) to augment clinical diversity in real-world settings. We propose a transfer learning approach utilising the DenseNet121 CNN architecture and a tailored classification head. This head was made to be efficient in terms of computation and to provide strong regularisation to lower the chance of overfitting on a small dataset. The main goals of this research are:

1. To create and train a deep learning model, specifically a CNN, that can accurately tell the difference between sickle-shaped RBCs and normal RBCs in microscopic images.
2. To make the model stronger and better at generalising, we will use advanced methods like RandomUnderSampler, SMOTE, and data augmentation (Albumentations) to deal with class imbalance and image variability.
3. To get high diagnostic accuracy while keeping the model's computational efficiency high enough to work on devices with limited resources.

RELATED WORK

Gaikwad et al. (2025) [8] created machine learning models that use a large dataset of blood counts to find early signs of blood disorders like leukaemia, anaemia, lymphoma, and sickle cell disease. We trained algorithms like XGBoost, Random Forest, and Decision Tree to guess how likely someone is to get sick. Random Forest did the best, with a 99.10% accuracy rate. These models improve early detection, which could lower death rates and make life better.

Gómez et al. (2025) [9] developed a supervised learning system to forecast anaemia, characterised by diminished haemoglobin or red blood cells, utilising antecedent data from Sabatini (2022). They trained models like Random Forest, Decision Trees, and Linear Discriminant Analysis. Random Forest was the best at classifying types of anaemia (microcytic, normocytic, macrocytic), with an accuracy of 99.82%, which was better than previous binary classification

efforts. K & P (2025) [10] employed Convolutional Neural Networks (CNNs) to diagnose sickle cell disease by examining the morphology of blood cells, utilising CNNs' capacity to recognise intricate patterns. Deo et al. (2024) [11] utilised deep learning by integrating CNN and LSTM with data augmentation and segmentation techniques to enhance the accuracy of sickle cell anaemia detection, providing a dependable instrument for early diagnosis and improved patient care.

(Goswami, Goswami, et al., 2024) [12] Sickle cell disease (SCD) is marked by red blood cells that are sickle-shaped, which makes it harder for the blood to carry oxygen and can cause serious health problems like "death, paralysis, and weakness." Transfer learning takes information from the input dataset and makes accurate predictions. In order to do this, we look at and compare the performance metrics of three different models: ResNet18, ResNet50, and GoogLeNet. The ResNet50 model had the highest accuracy, at 94.90%. Explainable AI is better for being clear and confirming the predictions of classifiers. This study shows that Grad-CAM makes it easier to understand models and makes them more reliable. So, pathologists benefit from this method's speed, precision, and accuracy in classifying sickle cells.

(Goswami, Sampathila, et al., 2024) [13] The system suggested in this paper is semi-automated and can take pictures by following a set program. To change the focus, it has a Z stage, and to move the slide up and down or side to side, it has an XY stage. This case study is about SCD. The suggested hardware captures SCD slides and then sorts them into groups based on what is normal. "GoogleNet, Darknet-19, ResNet50, ResNet18, and ResNet101" are some of the deep learning models that work with them. Most of the models that were tested got high scores in different settings, with an average of almost 97%. As a result, they did very well. This semi-automated technology could be useful for pathologists in the future and could be used in places where pathologists are hard to find.

(Machado et al., 2024) [14] This systematic review examines the application of "machine learning (ML) algorithms" in diagnosing sickle cell disease (SCD) and evaluating various clinical features, including early detection of organ failure, determining appropriate medication dosages, and classifying pain intensity. People look at a lot of machine learning algorithms, such as "The Multilayer Perceptron, Support Vector Machine, Random Forest, Logistic Regression, Long Short-Term Memory, Extreme Learning Machines, Convolutional Neural Networks, and Transfer Learning approaches." Research shows that even though there have been big improvements, there are still problems, such as small dataset sizes, worries about how easy it is to understand the results, and the possibility of overfitting. To overcome these constraints, future research must improve model interpretability, employ larger and more relevant datasets, and investigate sophisticated machine learning methodologies, including deep learning. The review concludes by emphasising the need for further research in the field and highlighting the promising potential of machine learning (ML) to improve the diagnosis, monitoring, and prognosis of sickle cell disease.

Pandey et al. (2024) [15] examined machine learning and deep learning for non-invasive anaemia detection, utilising convolutional neural networks (CNNs) to extract features from blood smears and microscopic images. The pre-processed datasets enabled effective training and validation, illustrating deep learning's ability to exceed traditional diagnostic methods for anaemia and type 2 diabetes, thus enhancing the efficiency of clinical image analysis and diagnosis.

Ramzan et al. (2024) [16] utilised machine learning techniques, including decision trees, logistic regression, support vector machines, random forests, k-nearest neighbours, and Naïve Bayes, alongside novel models incorporating spatial and multiple attention modules, to detect anaemia. These models achieved high accuracy, precision, recall, and F1 scores on both text and image datasets, and ablation studies confirmed the importance of attention mechanisms. Sachin and Bhadrashetty (2024) [17] employed machine learning methodologies, such as convolutional neural networks (CNNs) and support vector machines, to predict sickle cell anaemia (SCA) utilising the erythrocyte IDB database, achieving a 97% accuracy rate. This method is still in its early stages, but it has the potential to improve diagnosis, early management, and treatment outcomes for SCA.

METHODOLGY

The methodology is organised to ensure reproducibility, generalisability, and clinical utility. We first describe the consolidated dataset, then detail data-preprocessing, feature engineering, model training, and evaluation procedures.

Data Collection

The dataset was carefully put together to show what happens in real-life clinical settings, especially in places where SCA is common. Blood samples were obtained from

140 patients in the Teso region of Eastern Uganda, specifically in the Kumi and Soroti districts, through partnerships with Kumi Hospital, Soroti Regional Referral Hospital, and Soroti University. The Soroti University Research and Innovation Fund (project number RIF/2022/05) helped with this project. The Government of Uganda made sure that data was collected in a fair and organised way.

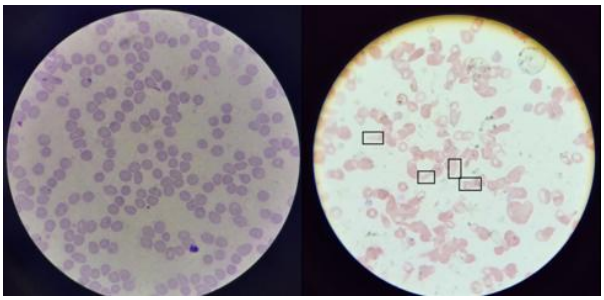

Figure 1 Dataset Sample

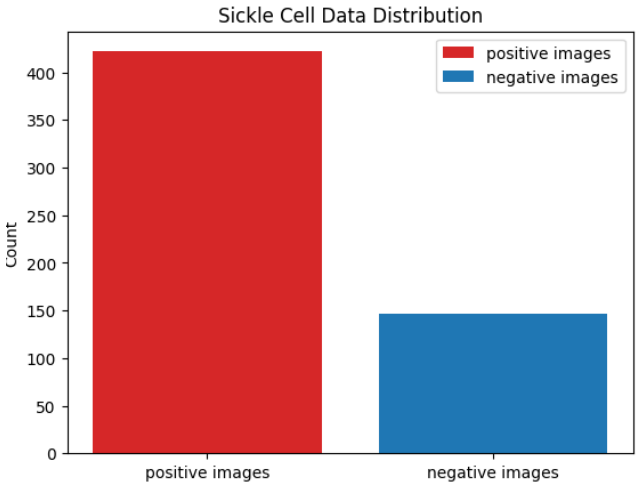

Figure 2 Data Distribution

Two staining techniques, Field stains and Leichman stains, were used to make RBCs easier to see and add variability that reflects different clinical practices. We chose these stains to mimic the different imaging conditions that can happen in places with limited resources, where the quality of the stains and the equipment may not be the same. Images were taken with high-resolution microscopic imaging equipment to make sure there was enough detail for deep learning analysis. The final dataset has 691 RBC images, 422 of which are positive (sickle cell) images and 269 of which are negative (non-sickle cell) images. The positive images are divided into two groups: those with labels (with bounding boxes around sickle cells for supervised learning) and those without labels. The negative images were split up into 147 clear images and 122 unclear images that had artefacts, bad staining, strange colours, or wrong cropping. Adding unclear images was done on purpose to see how well the model could handle bad imaging conditions, which is important for using it in places with few resources.

Table 1 Dataset Composition of Red Blood Cell Images

Category	Sub-Category	Count (Images)	Purpose/Significance
Positive (Sickle Cell)	Labelled	(Part of 422)	Supports supervised learning with annotations for precise sickle cell identification.
	Unlabelled	(Part of 422)	Enables flexible analysis and potential semi-supervised learning approaches.
Negative (Non-Sickle Cell)	Clear	147	Provides high-quality samples for training and evaluation.
	Not Clear	122	Introduces real-world variability (e.g., artefacts, poor staining) to improve model resilience.
Total		691	Captures diverse clinical imaging conditions for robust model training.

Data Preprocessing

A full preprocessing pipeline was put in place to make sure that the CNN architecture would work and to deal with problems like class imbalance and small dataset size. The chosen DenseNet121 model required that all images have a resolution of 224×224 pixels and three RGB colour channels. To make sure that the input scaling for model training was always the same, pixel values were normalised to the range [0,1] by dividing by 255.

To keep the positive and negative classes in the subsets in the same proportion, the dataset was split into a stratified 80:20 train-test ratio. We used a random seed of 42 to make sure that the 20% temporary set could be used again. We then split it into 60% for validation (X_{val}, y_{val}) and 40% for testing (X_{test}, y_{test}). To fix the class imbalance (422 positive images vs. 269 negative images), a two-step method was used on the training set. First, random under-sampling cut down on the number of positive samples to match the number of negative samples. This made the distribution more even, but it also made the

features less diverse. Second, the Synthetic Minority Over-sampling Technique (SMOTE) was used to create fake negative samples by filling in the gaps between real ones. This made the dataset balanced again, which is good for deep learning.

We only used data augmentation on the training set with the Albumentations library to make the dataset more diverse and the model more resistant to changes in real-world imaging. We added three copies of each training image, using both geometric and photometric changes to mimic clinical imaging problems like changing lighting, staining, and microscope focus. The augmentation pipeline had a lot of different steps. For example, it could rotate images by up to ±15°, change the brightness and contrast at random, add Gaussian blur to make images look out of focus, distort the grid and change the perspective to make lens artefacts, use Contrast Limited Adaptive Histogram Equalisation (CLAHE) to make local contrast stronger, and scale images at random by up to ±10%. These changes made sure that the model could work with a wide range of imaging conditions that are common in clinical practice.

Model Building and Training

The model was built using the DenseNet121 CNN architecture and the TensorFlow framework, which is a way to learn from other models. We chose DenseNet121 because it has a lot of connections, with each layer getting inputs from all the layers before it. This makes it easier to reuse features and makes the parameters work better than other CNN architectures. To make DenseNet121 work for SCA detection, we froze the first layers so that it could keep the general features (like edges and textures) it learnt from ImageNet. We also made the last 150 layers trainable so that we could fine-tune features that are specific to RBC morphology and sickle cell characteristics. The input layer was set up to take images that were 224×224×3, which is what the preprocessed dataset was set up to do.

A custom classification head was built to do binary classification (sickle cell vs. non-sickle cell). It has several layers to process the features that were extracted and make sure it works well. The next section goes into more detail about this head.

Proposed Model Architecture

The proposed model uses the DenseNet121 base and a custom classification head to improve SCA detection. After DenseNet121 extracts the features, the classification head processes the feature maps through the following parts:

1. **Global Average Pooling**: This layer reduces the spatial dimensions of the feature maps (e.g., 7×7×1024) to a 1024-dimensional vector, improving computational efficiency and resilience to translational variations in RBC images.
2. **Dense Layers and Regularization**: A sequence of fully connected dense layers with 1024, 512, and 256 units, each using ReLU activation, was implemented to capture complex patterns. To prevent overfitting, extensive regularization was applied:
    - **L2 Regularization**: Coefficients of 0.002, 0.003, and 0.001 were used for the 1024-, 512-, and 256-unit layers, respectively, to penalize large weights and promote generalization.
    - **Dropout**: Applied at rates of 0.5, 0.4, and 0.3 after each dense layer to randomly deactivate neurons during training, reducing overfitting risks.
    - **Batch Normalization**: Implemented after each dense layer to normalize activations, stabilize training, and accelerate convergence.
3. **Output Layer**: A single dense unit with sigmoid activation produces a probability score for binary classification, indicating the likelihood of an image containing sickle cells.

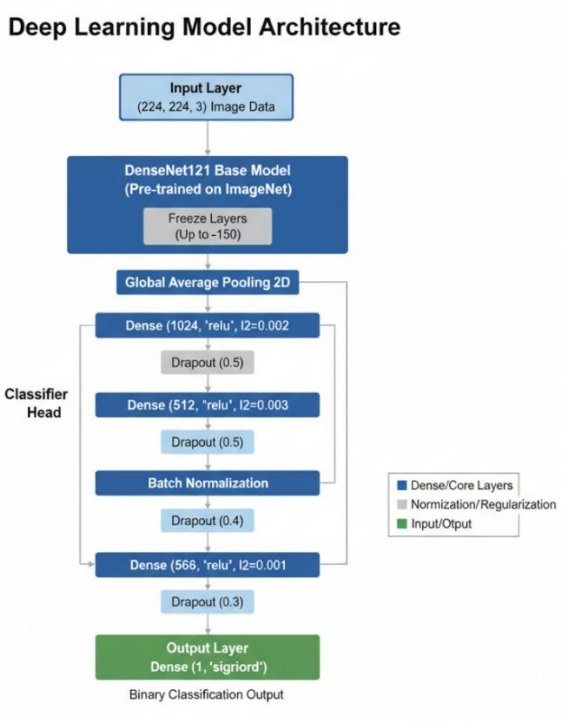

Figure 3 Model Architecture Diagram

This architecture balances complexity and efficiency, leveraging DenseNet121's feature extraction capabilities with a robust classification head tailored for SCA detection.

Training and Evaluation Strategy

To keep the integrity of the pre-trained weights while fine-tuning, the model was trained with the Adam optimiser and a starting learning rate of 0.0001. We used binary cross-entropy as the loss function because it works well for tasks that involve binary classification. We set the batch size to 32 and the maximum number of epochs to 15 to find a good balance between training speed and model convergence.

We used a 4-fold stratified cross-validation strategy to make sure the model worked well and didn't overfit. For each fold, we trained the model on 75% of the augmented data and validated it on the other 25%. Early stopping was used to stop training when the validation accuracy reached or went above 0.93. This made sure that the model worked as well as possible without doing extra work. For the final test, we kept the model weights from the epoch with the highest validation accuracy across all folds.

RESULTS AND DISCUSSION

The results section shows the real-world results of creating and testing the deep learning model that can find sickle cell anaemia (SCA) in microscopic images of red blood cells (RBCs). The evaluation mainly used training accuracy and training loss, which are metrics based on binary cross-entropy, to see how well the model was able to correctly classify images during training. The model was trained using a 4-fold cross-validation strategy. The metrics reported here show the best performance in the fourth fold, which shows that the model was able to tell sickle cell RBCs apart from normal ones.

Metrics for Training Performance

The training performance showed that it quickly converged and was very accurate. In the fourth fold of cross-validation, the DenseNet121 model reached its highest training accuracy of 96.4% at the ninth epoch. This meant that the training loss was 4.0153. The model's ability to learn distinguishing features, like the sickle cells' characteristic crescent shape, is confirmed by this strong result, even though different staining methods and image clarity can make things more difficult. The training loss, which was affected by the L2 regularisation penalties built into the model's custom classification head, steadily went down, showing that the optimisation was stable and effective. Table 2 shows the final training metrics for the DenseNet121 model:

Table 2 DenseNet121 model Results

Training Accuracy (%)	Training Loss
96.40	4.0153

Figure 4 Accuracy and Loss results

Detailed Progression and Analysis

The examination of the epoch-by-epoch training progression further demonstrates the model's effective learning dynamics and successful convergence. Epoch 1 began with a 59.62% accuracy rate and a 5.2453 loss rate, which shows that the model was starting to adapt to the larger dataset. The model got better quickly, reaching 89.20% accuracy by Epoch 3. Around Epoch 6 (95.94% accuracy, 4.2283 loss), there was a big jump in performance, showing that the important morphological features were learnt well. Epoch 8 had the highest accuracy, 96.89%, but it dropped slightly to 96.40% in Epoch 9, with a loss of 4.0153. The early stopping mechanism stopped the training process when the validation accuracy went over the set limit of 0.93. This made sure that the model kept the best configuration without going through unnecessary training iterations. This strong performance in training, which is also shown in the training curve graph, shows that the combination of data augmentation (Albumentations), class balancing (RandomUnderSampler and SMOTE), and the DenseNet121 architecture worked well to solve the problems of small dataset size and variability.

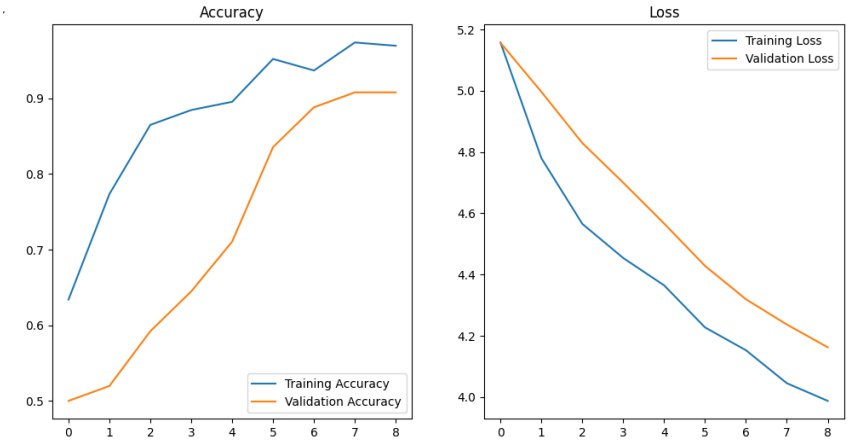

Figure 5 Accuracy and Loss Curve

The analysis of the Receiver Operating Characteristic (ROC) curve graph further supports the model's ability to tell the difference between things. With a training accuracy of 96.4%, the ROC curve should show a high Area Under the Curve (AUC). This shows that the model is better at reliably telling the difference between positive (sickle cell) and negative (non-sickle cell) classes at different classification thresholds.

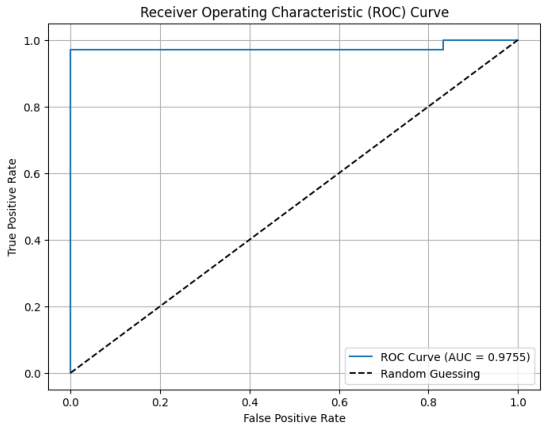

Figure 6 AUC Curve

Model Predictions

Sample predictions on RBC images qualitatively confirmed that the model was very good at making predictions. The test set showed that the DenseNet121 model could tell the difference between healthy and sickle cell images. The model correctly identified a sample of healthy RBCs (non-sickle cell) as negative with high confidence, recognising the round shape of RBCs. However, when the model was shown pictures of cells with sickle shapes, it correctly identified the sample as positive with high confidence. This shows that it can still find the crescent shape that shows SCA, even if the staining or image quality changes. These qualitative results corroborate the quantitative findings, illustrating the model's considerable potential for clinical application.

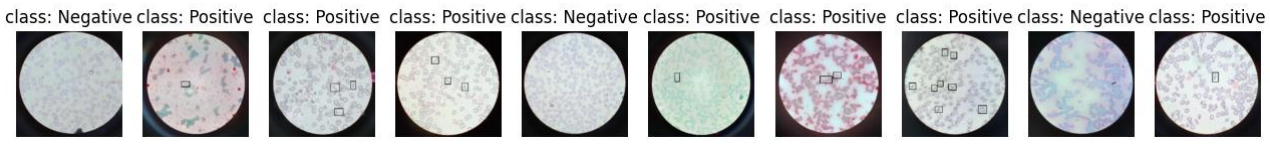

Figure 7 Model Predictions

Comparison with Existing Work

The suggested DenseNet121-based model worked better and trained faster than other automated SCA diagnosis models. When compared to an InceptionV3-based model [18], the DenseNet121 architecture has a clear advantage in three important areas: training accuracy, training loss, and the number of epochs needed for convergence. The previous work, which used InceptionV3, got a training accuracy of 91.0% and a training loss of 28.7000 after 100 epochs. The new DenseNet121 model, on the other hand, got a much higher accuracy of 96.4% and a much lower loss of 4.0153 in just 9 epochs.

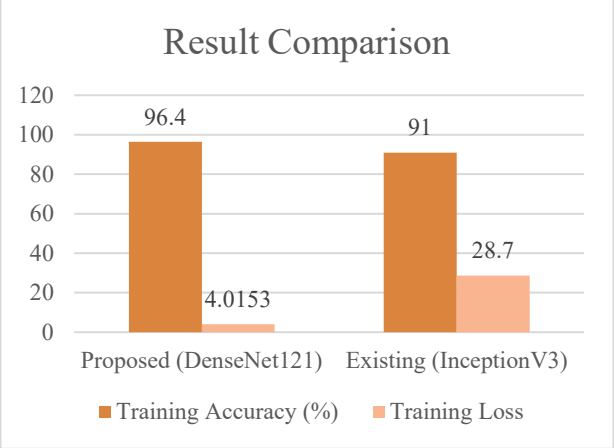

Figure 8 Result Comparison with Baseline Research

This higher efficiency shows that the DenseNet121 architecture, along with targeted regularisation and a strong pre-processing pipeline, converges faster and gives better classification results. This quick convergence is very important for using the technology in places with limited resources, like Uganda's Teso region. Table 3 shows the results of the comparison:

Table 3 Result Comparison with Baseline Research

Model	Training Accuracy (%)	Training Loss	Epochs
Proposed (DenseNet121)	96.40	4.0153	9
Existing (InceptionV3) [18]	91.00	28.7000	100

CONCLUSION

This research successfully developed an exceptionally accurate and efficient deep learning model for the automated detection of sickle cell anaemia (SCA) in red blood cell (RBC) images, addressing the critical need for improved diagnostics in resource-constrained environments, such as the Teso region of Uganda. The methodology utilised a DenseNet121-based convolutional neural network (CNN) and a unique dataset comprising 691 images subjected to diverse Field and Leichman stains. It used strong methods like RandomUnderSampler, SMOTE, and Albumentations for data augmentation to fix class imbalance and image variability.

The model's best training accuracy was 96.4% and its best training loss was 4.0153 after only nine epochs. This result shows that this method works better and faster than other methods that use an InceptionV3 model, which took 100 epochs to get 91% accuracy. The model's high accuracy and fast convergence show that it could be a useful tool for finding SCA early in places where it is common. Future work will focus on evaluating performance on the specified test set and improving the model for use in clinical settings with limited resources.

REFERENCE

[1] Pauling, Linus, et al. "Sickle cell anemia, a molecular disease." Science 110.2865 (1949): 543-548.

[2] Tisdale, John F., Swee Lay Thein, and William A. Eaton. "Treating sickle cell anemia." Science 367.6483 (2020): 1198-1199.

[3] Adio, Waheed Sakariyau, et al. "Current trends in the treatment of sickle cell anemia." World News of Natural Sciences (WNOFNS) 43 (2022).

[4] Buhari, Hauwa Ali, Aisha Sa'ad Ahmad, and Emmanuel Ifeanyi Obeagu. "Current advances in the diagnosis and treatment of sickle cell anaemia." (2023).

[5] Brandow, A. M., and R. I. Liem. "Advances in the diagnosis and treatment of sickle cell disease." Journal of hematology & oncology 15.1 (2022): 20.

[6] Simon, Kawuma, et al. "Comparison of deep learning techniques in detection of sickle cell disease." Artificial Intelligence and Applications. Vol. 1. No. 4. 2023.

[7] Bushra, S. Nikkath, and G. Shobana. "Paediatric sickle cell detection using deep learning-a review." 2021 International Conference on Artificial Intelligence and Smart Systems (ICAIS). IEEE, 2021.

[8] D. K. Gaikwad, V. Mahale, A. Gaikwad, and A. Gaikwad, "Automatic Blood Disease Identification System Using Machine Learning Algorithms with Hemogram Reports," INDIAN J. Sci. Technol., vol. 18, no. 1, pp. 11–22, 2025.

[9] J. G. Gómez, C. Parra Urueta, D. S. Álvarez, V. Hernández Riaño, and G. Ramirez-Gonzalez, "Anemia Classification System Using Machine Learning," Informatics, vol. 12, no. 1, pp. 1–16, 2025, doi: 10.3390/informatics12010019.

[10] S. K and A. D. K. T. P, "Sickle Cell Anemia Detection using Deep Learning," Int. J. Sci. Res. Technol., vol. 2, no. 1, pp. 15–26, 2025, doi: 10.21608/ijaiet.2024.240910.1002.

[11] A. Deo, I. Pandey, S. S. Khan, A. Mandlik, N. V. Doohan, and B. Panchal, "Deep Learning-Based Red Blood Cell Classification for Sickle Cell Anemia Diagnosis Using Hybrid CNN-LSTM Model," Internetional Inf. Eng. Technol. Assoc., vol. 41, no. 3, pp. 259–302, 2024, doi: 10.2307/j.ctv18ph9kr.16.

[12] N. G. Goswami, A. Goswami, N. Sampathila, M. G. Bairy, K. Chadaga, and S. Belurkar, "Detection of sickle cell disease using deep neural networks and explainable artificial intelligence," J. Intell. Syst., vol. 33, no. 1, 2024, doi: 10.1515/jisys-2023-0179.

[13] N. G. Goswami, N. Sampathila, G. M. Bairy, A. Goswami, D. D. Brp Siddarama, and S. Belurkar, "Explainable Artificial Intelligence and Deep Learning Methods for the Detection of Sickle Cell by Capturing the Digital Images of Blood Smears," Inf., vol. 15, no. 7, 2024, doi: 10.3390/info15070403.

[14] T. F. Machado, F. das Chagas Barros Neto, M. de Souza Gonçalves, C. G. Barbosa, and M. E. Barreto, "Exploring machine learning algorithms in sickle cell disease patient data: A systematic review," PLoS One, vol. 19, no. 11, pp. 1–17, 2024, doi: 10.1371/journal.pone.0313315.

[15] R. Pandey, S. Ahmad, S. Yadav, and S. Srivastava, "Anemia Detection Using Machine Learning and Deep Learning - A Systematic Literature Review," Int. J. Adv. Res. Innov. Ideas Educ., vol. 10, no. 1, pp. 368–377, 2024.

[16] M. Ramzan, J. Sheng, M. U. Saeed, B. Wang, and F. Z. Duraihem, "Revolutionizing anemia detection: integrative machine learning models and advanced attention mechanisms," Vis. Comput. Ind. Biomed. Art, vol. 7, no. 18, 2024, doi: 10.1186/s42492-024-00169-4.

[17] Sachin and A. Bhadrashetty, "MACHINE LEARNING BASED PREDICTION OF SICKLE CELL ANEMIA IN HUMAN RBC," Int. Res. J. Mod. Eng. Technol. Sci., vol. 6, no. 5, 2024.

[18] Sen, Bheem, et al. "Deep Learning based diagnosis of sickle cell anemia in human RBC." 2021 2nd International Conference on Intelligent Engineering and Management (ICIEM). IEEE, 2021.