LETTERS IN BIOMATHEMATICS
An International Journal

RESEARCH ARTICLE

OPEN ACCESS

Cancer Cell Classification from Peripheral Blood Smear Data Using the YOLO Architecture

Raghavarapu Yaswanth[1], Venubabu Rachapudi[2]

[1]Student, Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation, Vaddeswaram, Guntur, Andhra Pradesh, India -522302, Email- 2401050083cse@gmail.com
[2]Professor, Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation, Vaddeswaram, Guntur, Andhra Pradesh, India -522302, Email- venubabu.r@gmail.com

ABSTRACT

Proper recognition of cancer cells on a peripheral blood smear picture is necessary in the timely diagnosis of haematological malignancies. In the paper, a comparative analysis of three YOLO-based architecture, i.e., YOLOv8n, YOLOv11n, and YOLO26n, will be conducted to automatically detect and classify blood cells, i.e., eosinophils, lymphocytes, neutrophils, and platelets. The present research attempts to juxtapose the performance of these architectures on peripheral blood smear data and identify the most helpful model to be adopted in the clinic. The methodology will involve the training and validation of all the models on a publicly available sample of blood cells using homogenous training processes and evaluation metrics. Findings show that YOLO26n is superior to YOLOv8n and YOLOv11n in all the important metrics. The proposed methodology has a mAP 0.5:0.95 of 0.9234 which is superior to YOLOv8n (0.9196) and YOLOv11n (0.9193) and its recall on eosinophil and lymphocyte classes is perfect. The findings apply to the sphere of medical image analysis and automated hematology because they demonstrate the efficiency of the end-to-end detection systems with no NMS in the blood cell classification. The implications of the clinical practice are as follows: the time of the diagnosis can be minimized, the consistency can be enhanced, and it can be applied in resource-limited environments with the assistance of an effective model design. The possibilities of the future research will be to enhance platelet detection with the help of special data augmentation and multi-scale feature extraction techniques.

ARTICLE HISTORY
Received October 15, 2026
Accepted November 19, 2026

KEYWORDS
YOLOv26, Blood Cell Classification, Object Detection, Peripheral Blood Smear, Deep Learning, Cancer Detection, Medical Image Analysis.

1. Introduction

Urban Peripheral blood smear examination is one of the foundations of hematological diagnosis that allows revealing morphological defects that point to different types of cancers, such as leukemia, lymphoma, and myeloproliferative diseases [1]. Conventionally, manual analysis by trained hematologists and pathologists is conducted, which is both time consuming, labor intensive, and prone to inter-observer variability [2]. The rising cases of hematological malignancies around the world require the establishment of automated, accurate, and efficient diagnostic systems that can facilitate clinical decision-making.

The latest developments in the area of deep learning have transformed the field of medical image analysis, and object detection architectures have shown specific potential in the field of cell classification [3]. One of them is the You Only Look Once (YOLO) family of models, which have become one of the most popular methods of detecting objects in real-time, providing the best trade-off between the accuracy of detection and computational efficiency [4]. The historical developments of the early implementations to the more recent versions, such as YOLOv8, YOLOv11 and the recent introduction of the YOLO26, have introduced major architectural improvements, such as anchor-free detection, improved feature extraction modules, and new optimization methods [5]– [7].

Various works have explored detection-based method of analyzing blood cells. An approach that utilized the YOLOv5s architecture to detect red blood cells, white blood cells and platelets developed a robust framework of cell detection in the peripheral blood smears [5]. Later research using multi-scale feature fusion mechanisms enhanced the detection accuracy with different cell sizes [6]. An automated blood cell detection and counting system was built with a detection framework that relied on YOLOv8 [7], and the reduction of parameters by GhostConv and C3Ghost modules facilitated effective application in resource-constrained environments [8]. It was also proposed that a lightweight version would be used in real-time clinical applications [9].

The more advanced and recent developments have been considered through comparative experiments on blood cell classification, which prove the benefits of the newer designs [10]. Additional optimization with gradient accumulation and loss reweighting approaches has

CONTACT Raghavarapu Yaswanth: 2401050083cse@gmail.com

been implemented to deal with clumped and overlapping blood cells [11]. Dynamic cross-scale context aggregation framework was proposed to implement the high-performance blood cell classification [12]. In the case of leukemia detection, a detection-based architecture with generative adversarial networks has been trained on acute myeloid leukemia classification [13], and another model has been trained on acute lymphoblastic leukemia classification, where the use of such architecture is demonstrated to be applicable to cancer cell classification [14].

Irrespective of these developments, more recent advancements such as the YOLO26 architecture, which has an NMS-free end-to-end detection paradigm and a new optimizer called MuSGD, have not been tested on blood cell classification yet. By removing Non-Maximum Suppression (NMS) as a post-processing tool, it is possible to have true end-to-end differentiable detection which may result in a faster inference and a higher localization accuracy. Moreover, the MuSGD optimizer, a hybrid of second-order momentum-based optimizers with classic stochastic gradient descent, has better convergence characteristics that might be useful in medical imaging processes, where a high degree of localization is essential.

This paper is a comparative analysis of three detection architectures, that is, YOLOv8n, YOLOv11n, and YOLO26n, and the goal is to obtain automated detection and classification of blood cells in the peripheral blood smear samples. The main contributions of the work are the following:

1. A systematic evaluation of three architectures on a publicly available blood cell dataset, demonstrating the superior performance of the most recent variant for this application.
2. The initial implementation of the NMS-free end-to-end detection architecture using MuSGD optimizer to blood cell classification in medical images.
3. Comprehensive documentation of training configurations, hyperparameters, and evaluation metrics to ensure reproducibility and facilitate future research in automated hematology.
4. Identification of persistent challenges in platelet detection, providing direction for future improvements in small object detection within blood smear analysis.

2. Related Work

The section gives a review of the literature already present on blood cell detection and classification in deep learning applications, and specifically, the YOLO-based architectures. The review is conceptualized to bring out the theme of development of detection techniques, architectural advances and never-ending issues in the concept.

A. Early YOLO-Based Approaches for Blood Cell Detection

YOLOv5-based approaches were the first to be applied to blood cell analysis. SDE-YOLO is a YOLOv5s-based method that was suggested to detect red blood cells, white blood cells, and platelets, setting a good foundation on cell detection within peripheral blood smears [5]. The present work showed that single-stage detectors could be used in hematological functions. Later advancements to YOLOv5 also featured more efficient feature extraction algorithms, which resulted in the improved real-time detection and classification capabilities on blood smear images [4].

B. Multi-Scale and Lightweight Architectures

Understanding that it is difficult to detect cells of different scales, scholars have made architectures that have multi-scale feature fusion. TW-YOLO is a multi-scale YOLO-based model, which improved the detection accuracy of various cell sizes by providing better feature pyramid structures [6]. This method was especially successful at the simultaneous detection of both large white blood cells and small platelets.

Lightweight architectural variants have been created through efforts to support operation in constrained resource clinical environments. An automated blood cell detector was trained with a parameter-reduced YOLOv8n with GhostConv and C3Ghost modules that provide significant model size reductions with minimal effect on accuracy [8]. Likewise, a fast GPMB-YOLO based on the lightweight YOLOv8n was tailored to real-time blood cells detection in medical imaging uses, and showed that it could do efficient inference at edge devices [9]. A heavyweight model, DWS-YOLO, was another lightweight detector that used attention mechanisms and SIoU loss to enhance blood cell data performance [15].

C. YOLOv8-Based Enhancements

A number of works have come up with ways of improving the YOLOv8 architecture to specifically analyze blood cells. An automated blood cell detection and counting model, YOLOv8-BC, was created on the basis of the YOLOv8 framework, and it has shown high results on the peripheral blood smear images [7]. To enhance the ability of dense and occluded blood cells, a framework known as NBCDC-YOLOv8 was proposed to enhance the detection in challenging conditions with significant improvement in detection performance [16]. In the context of the detection of acute lymphoblastic leukemia, in particular, FCE-YOLO, a model based on YOLOv8, showed high accuracy and efficiency in subtype classification of leukemia, and its generalization was confirmed on several datasets [14].

D. YOLOv10, YOLOv11, and Advanced Architectures

Comparative studies on blood cell classification have been used to discuss the transition to YOLOv10 and YOLOv11. A methodical evaluation of YOLOv10 and YOLOv11 over the BCCD dataset showed the benefits of recent architecture designs such as better feature extraction and shorter inference durations [10]. Additional optimization of YOLOv11n with gradient accumulation and loss reweighting techniques was implemented to works with dense and overlapping blood cells in clinical smears, and improved the detection accuracy in difficult areas was achieved [11]. An improved framework with a dynamic context aggregation of scales (DCSCA) using YOLOv11 was created to classify blood cells with high performance, and it showed better performances on multi-class classification problems [12]. To detect anomalies, a YOLOv11-ResNet based method was created to detect abnormal cells in blood smears automatically, which can also be applied to detect pathological cells [8].

E. Leukemia and Cancer Cell Detection

The use of deep learning in the detection and classification of leukemia has been targeted. An in-depth analysis of the deep learning methods of leukemia detection revealed the development of the traditional CNNs to the YOLO-based structures, recognizing the main challenges and the further perspectives [1]. In the case of the acute myeloid leukemia, an intermediate strategy involving the use of YOLOv8 and generative adversarial networks showed better results in detecting malignant cells in bone marrow images [13]. Moreover, a two-stage detection and classification model based on enhanced YOLOv7 in conjunction with EfficientNetv2 was trained to perform the analysis of white blood cell, red blood cell and platelets in full-field blood smear images [11]. To achieve bone marrow cell detection, an enhanced YOLOv7 algorithm has been offered, which has shown efficient detection of several types of cells in microscopic images [10].

F. Complementary Approaches and Reviews

Other than the direct studies of YOLO, other complementary deep learning methods have been used to analyse blood cells. The possible automated counting and classification of blood cells in clinical workflows was demonstrated by an efficient deep learning pipeline to classify biocells based on microscopic smear images [2]. Stain-free artificial intelligence-assisted light microscopy was investigated to classify blood cells, which is an alternative technique that does not require chemical staining [3]. An image segmentation and classification system of a blood cell that combines the use of YOLO-based detection and the use of segmentation methods indicated overall analysis of the peripheral blood smears [17]. An attention-based approach based upon YOLOv7 was suggested to detect bone marrow cells, which has proven to be effective in detecting microscopic bone marrow cells [18]. An enhanced version of YOLOv7 and EfficientNetv2 was used to create a two-stage detection and classification algorithm to analyze the full-field blood smear images [19]. To detect the anomaly, a ResNet based on YOLOv11 was found to detect the presence of abnormal cells in blood smears [20].

G. Research Gap

Although it is evident that there has been a significant amount of research on YOLO-based blood cell detection, there are a number of gaps present. To begin with, the newly proposed YOLO26 architecture, with NMS-free end-to-end detection, and MuSGD optimizer, was not tested on blood cell classification. Second, the majority of available literature is devoted to ordinary blood cell types, but not to cancer-specific applications. Third, small objects and imbalance between classes continue to be a challenging issue in all architectures because of platelet detection. This paper fills these gaps by providing the initial application of YOLO26n to peripheral blood smear classification, where the primary emphasis is put on cancer cell classification and a deeper, systematic comparison with the already known YOLOv8n and YOLOv11n architectures.

3. Methodology

A. Study Design

The research paper introduces an automated system of classifying blood cells in peripheral blood smears using the image-based object detection networks based on deep learning. The suggested system is aimed at precisely identifying and identifying four different types of cells, namely, eosinophils, lymphocytes, neutrophils, and platelets, which are vital markers in the diagnosis of hematological malignancies.

This system has a well-defined pipeline that includes data acquisition, preprocessing, model training and performance evaluation. At the first stage, the input peripheral blood smear images are pre-processed to equalize the resolution and augmentation methods that increase the generalization of the model. After that, several detection architectures will be trained and checked on the ready dataset with the help of transfer learning to utilize pretrained features. The last step will be detailed analysis based on the conventional object detection metrics to determine the most suitable architecture to use in this application.

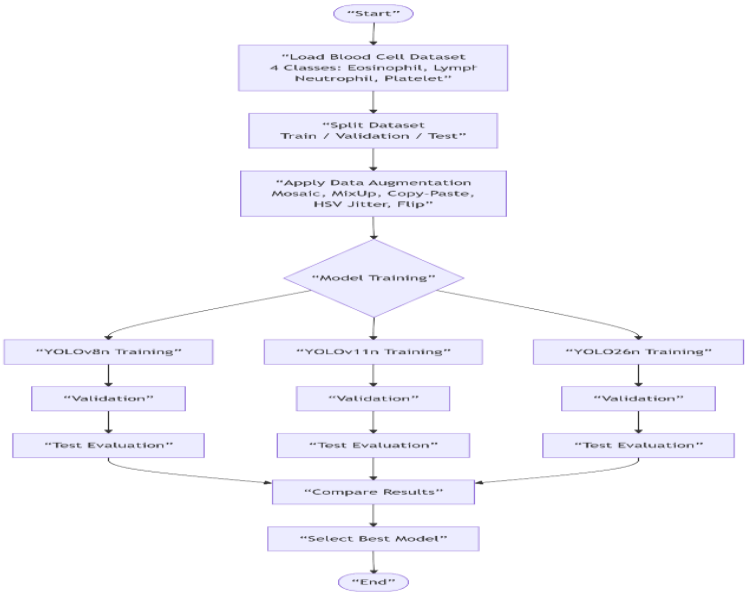

Fig 1: Overall workflow of the proposed system for automated blood cell classification.

The proposed system starts by loading the blood cell dataset as shown in Fig. 1 which is of four different classes namely eosinophil, lymphocyte, neutrophil and platelet. The data is then divided into training, validation and test sets to enable the development of models and objective testing. After the split, a sequence of data augmentation, such as mosaic, mix-up, copy-paste, HSV jitter, and horizontal flip are used to make the models more generalized and resistant to staining, illumination, and cell morphology changes.

The augmented training data is then inputted into three separate detection architectures, which are then trained concurrently. All the architectures are tested on the validation set to track performance and choose the best model checkpoint. Final evaluation on the held-out test set is then done to the trained models in order to get unbiased performance measures. The outputs of all the three architectures are compared in a systematic manner and an architecture with the best overall performance is chosen as the optimal model to use in the automated blood cell classification task.

B. Data Collection

This study used the Blood Cell Classification v2 data which is of the Roboflow Universe under the CC BY 4.0 license. The dataset consists of images of peripheral blood smears marked with bounding boxes of four different cell types including eosinophil, lymphocyte, neutrophil, and platelet



Fig 2: Class-wise annotation counts per split across training, validation, and test sets.

Fig. 2 shows that the dataset consists of 4,551 annotated bounding boxes that are spread over 4,412 images. The training set consists of 3,187 annotations whereas the validation and test sets consist of 914 and 450 annotations respectively. The dataset features moderate class imbalance with the eosinophil being the most represented class (1,418 annotations) and platelet being the

least represented class (895 annotations) with a ratio of 1.58:1. Such imbalance poses a significant problem of model training, especially the accurate identification of the platelet category, which is represented smaller and a smaller object and comprises less than half the dataset.

C. Data Processing

Images were all downsized to 640 x 640 to ensure that all architectures operated at the same resolution, since it is the standard resolution used with detection-based models and provides a tradeoff between detection performance and computation time. The data was divided into three groups, which are training (3,089 images, 3,187 annotations), validation (882 images, 914 annotations) and test (441 images, 450 annotations). Hyperparameter tuning and model checkpoint selection were done on the validation set, and final unbiased evaluation was done on the test set.

A complex data augmentation pipeline was used in training to enhance the generalization and resistance of the model to variations in staining, illumination, and morphology of the cell. To reproduce the various staining and lighting conditions, the augmentation strategies were HSV hue variation of 0.015, HSV saturation scaling of 0.7 and HSV brightness variation of 0.4. In order to add spatial invariance, horizontal flip was employed with a probability of 50%. Mosaic augmentation (that comprises four images into one training sample) was implemented in full strength to boost multi-scale learning and detection of cells of different sizes. Also, MixUp and Copy-Paste augmentations were added with 10% chance each to further increase training sample diversity and model resilience to overlapping and closely spaced cells. The combination of these augmentation methods introduces large amounts of variations in images to the models, making them less overfitted and effective on unseen test data.

D. Model Training

YOLO26n architecture is the latest addition to the YOLO lineup of object detection models and it has two paradigm-shifting innovations that make it stand out of the pack. In contrast to the traditional detection models which use Non-Maximum Suppression (NMS) as a post-processing method to filter the duplicate predictions, the YOLO26n utilizes an NMS-free end-to-end detection architecture. This architectural design allows this model to learn to ignore unnecessary detections during training, which makes end-to-end optimization really differentiable. Eradication of NMS does not only decrease the inference latency, but also enables the model to acquire more optimal detection strategies in direct relation with the data, resulting in enhanced localization accuracy.

1) Yolov26n Architecture

YOLO26n architecture is shown in Fig. 3. It is based on the model which takes in input peripheral blood smear images resized to 640 640 pixels into a backbone network with C2f and E-ELAN modules to extract features effectively. The neck implements the SPPF (Spatial Pyramid Pooling Fast) blocks and COL (Cross-Over Layer) modules to fuse features on three scales 80 × 80, 40 × 40, and 20 × 20. These multi-scale properties allow the strong identification of cells of different sizes, both small platelets and bigger lymphocytes and neutrophils. The end-to-end detection head, which is based on RT-DETR architecture, makes final predictions without NMS processing after.

MuSGD optimizer is a variant of Muon optimizer and SGD that is adapted to the YOLO26n and to achieve the stability of SGD, the Muon optimizer is combined with the rapid convergence rate of the Muon optimizer. The rule of update of the parameters is:

$$\theta_{t+1} = \theta_t - \eta_t \cdot (\mu \cdot \nabla L + (1 - \mu) \cdot \nabla^2 L) \quad (1)$$

where θ_t represents the model parameters at time t, η_t is the learning rate, μ is a hyperparameter controlling the combination of gradients and second-order information, and ∇L and $\nabla^2 L$ are the first and second derivatives of the loss function L.

The overall loss function for YOLO26n combines localization loss (mean squared error between predicted and true bounding box coordinates), classification loss (cross-entropy between predicted and ground truth classes), and confidence loss (for objectness score) as follows:

$$L = L_{loc} + L_{cls} + L_{conf} \quad (2)$$

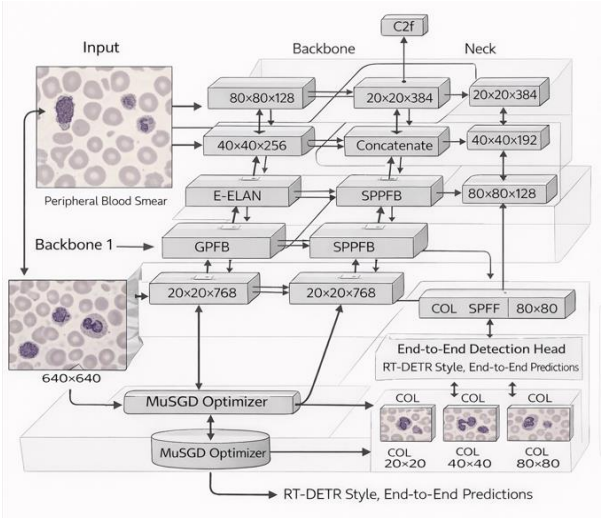

Fig 3: YOLO26n architecture showing backbone, neck, and end-to-end detection head.

The second major innovation of YOLO26n is the addition of the MuSGD optimizer, a hybrid method that replaces the stochastic gradient descent with the Muon optimizer, a second-order momentum-based method that has the ability to capture information about the curvature. It is a hybrid maximizer which combines the advantages of the two: the second-order methods have fast convergence properties, and SGD has stability and generalization. MuSGD optimizer is especially suitable to the NMS-free architecture as it can navigate the complex loss landscape of end-to-end detection training to arrive at faster convergence and better final performance.

2) Training Configuration

All the models were trained with pretrained weights and were fine-tuned on the blood cell dataset. The use of transfer learning to capitalize on features acquired by large-scale image datasets allowed the efficient training with the scarce medical imaging data.

Table 1: Training Hyperparameters for each model

Hyperparameter	YOLOv8n	YOLOv11n	YOLO26n
Epochs	20	15	15
Image Size	640 × 640	640 × 640	640 × 640
Batch Size	16	16	16
Optimizer	SGD	AdamW	MuSGD
Learning Rate (lr0)	0.01	0.001	0.01
LR Final Ratio (lrf)	0.01	0.01	0.01
Momentum	0.937	0.937	0.937
Weight Decay	0.0005	0.0005	0.0005
Warmup Epochs	Default	3	3
Early Stopping Patience	Default	20	20

The settings for the training were picked according to the suggested setting of each architecture to have a fair comparison as summarized in Table I: YOLOv8n was trained on 20 epochs using the SGD optimizer with a learning rate of 0.01, while YOLOv11n was trained on 15 epochs using the AdamW optimizer with a learning rate of 0.001. The optimizer used by YOLO26n was the novel MuSGD optimizer with a learning rate of 0.01 and 15 epochs. The batch size of 16 was used for all models and the image resolution

was 640 x 640 pixels for all images. During training each model minimized the sum of three loss functions: Box Loss (CIoU) for regression of the bounding boxes, Class Loss (BCE) for multi-class classification, and DFL Loss for precise boundary regression. MuSGD optimizer used in YOLO26n is a combination of second-order momentum-based optimization techniques with conventional stochastic gradient descent, which allows it to converge faster and have a higher degree of localization than traditional optimizers.

F. Model Evaluation

In the measurement and comparison of the performance of the three architectures standard object detection measures were used. These measures give quantitative values of detection accuracy and localization precision.

1) Precision and Recall

Precision is a measure of the fraction of the objects correctly identified by the model out of all the objects identified by the model. Describes the percentage of the detected bounding boxes that must come from the ground truth. When the precision value is large, it means that the false positive rate is low.

Recall: Percentage of ground truth objects that the model can detect. It reflects the number of the real objects in the images that were perceived correctly. High level of recall results in low number of false negative.

2) Mean Average Precision (mAP)

Mean Average Precision (mAP) is the main measure of object detection work. For every class, the average precision (AP) is the area under the precision-recall curve. The mean of the AP values of all classes is then taken to get the mAP.

This study has calculated two variants of mAP:

- mAP@0.5: obtained when IoU is set to 0.50. This measure is an analysis of detection performance at moderate localization tolerance.
- mAP@0.5:0.95: It is the normative COCO evaluation measure and offers an overall evaluation of localization accuracy regarding different levels of tolerance. A detection can only be regarded as correct when the bounding box predicted intersects with the ground truth box by the set IoU threshold.

The mAP0.5:0.95 measure is specifically significant in medical imaging tasks, where high accuracy in localization of cells is essential to proper diagnosis.

4. Results and Discussion

A. Experimental Results

The held-out test set of 441 images with 450 annotated bounding boxes was evaluated through an experiment. Standard object detection metrics were used to assess all models, including precision, recall, mAP at 0.5 and mAP at 0.5:0.95. The first section of the paper is the performance of YOLO26n, which is the main contribution of the current work, and then the comparison between the performance of YOLO26n and that of YOLOv8n and YOLOv11n.

1) Performance of YOLO26n

YOLO26n architecture performed better in all metrics of evaluation. The model attained a precision of 0.9874, recall of 0.9543, mAP@0.5 of 0.9899, and mAP@0.5:0.95 of 0.9234. The high accuracy means that most of the detections of the model are associated with real cells and the recall measure shows how well the model can identify a significant number of ground truth annotations.

On a per-class basis, it is observed that YOLO26n performed outstandingly in detecting eosinophil, lymphocyte, and neutrophil. Both eosinophil and lymphocyte classes achieved the perfect recall (1.000), which means that all ground truth examples of these cell types were correctly identified. Neutrophil detection also exhibited a high performance, precision of 0.999, recall of 0.995 and mAP:0.5:0.95 of 0.988. These findings suggest that NMS-free end-to-end architecture and MuSGD optimizer specifically work well in the realization of these types of cells as they are larger in size and have a more distinct morphology.

Nevertheless, platelet detection was still difficult to YOLO26n, and mAP 0.5:0.95 was 0.739, recall 0.866 and precision 0.982. This disparity in performance is explained by the fact that platelets are smaller and their representation in the training data is less.

In order to give visual information about the detection performance, the sample detection results of YOLO26n are shown in Fig. 4. The figure represents exemplary instances of cells correctly recognized in all classes, which shows that the model is capable of precisely localizing and classifying eosinophils, lymphocytes, neutrophils, and platelets in peripheral blood smear images.

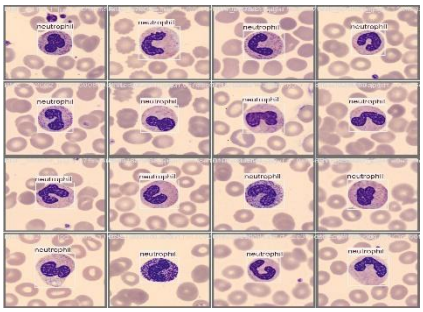

Fig 4: Sample detection results from YOLO26n showing correctly identified class with bounding boxes and confidence scores.

The classification performance of YOLO26n is further analyzed through the confusion matrix presented in Fig. 5. The matrix shows strong diagonal dominance, indicating accurate classification of detected cells. The majority of misclassifications occurred between similar cell types, particularly in cases where cells exhibited overlapping morphological characteristics. The background class represents regions where no cells were present, with minimal false positives reported.

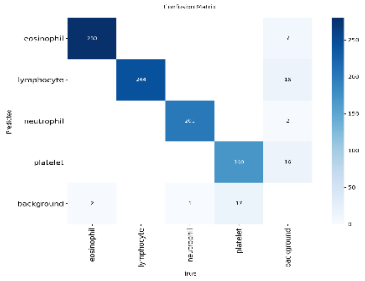

Fig 5: Confusion matrix for YOLO26n showing classification performance classes.

Fig. 6 depicts the training convergence behavior of YOLO26n, and the results indicate the usefulness of the MuSGD optimizer. The model reached stable convergence after 15 epochs with a box loss reduction of 0.720 to 0.450, class loss reduction of 0.220 to 0.200 and DFL loss reduction of 0.060 to 0.080. The mAP 0.5:0.95 measure rose uniformly between 0.920 at epoch 5, and 0.930 at epoch 15, which shows a steady growth in the training process.

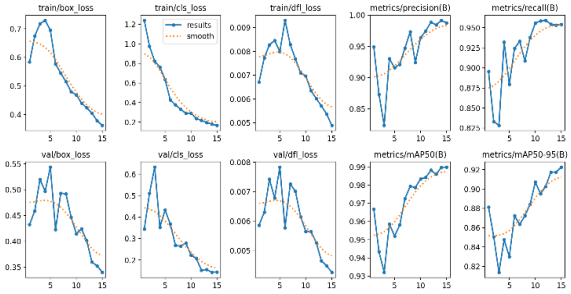

Fig 6: Training convergence curves for YOLO26n.

B) Comparative Analysis

A comparison was made in a systematic manner to determine the performance of the three architectures against each other, under the same experimental conditions. Table I shows the general performance metrics of YOLOv8n, YOLOv11n, and YOLO26n on the test set.

Table II: Overall Performance Comparison of YOLO Architectures

Model	Precision	Recall	mAP@0.5	mAP@0.5:0.95
YOLOv8n	0.9908	0.9661	0.9911	0.9196
YOLOv11n	0.9737	0.9581	0.9870	0.9192
YOLO26n	0.9874	0.9543	0.9899	0.9234

As shown in Table II, YOLO26n achieved the highest mAP@0.5:0.95 of 0.9234, outperforming YOLOv8n by 0.0038 and YOLOv11n by 0.0042. This measure is especially important because it assesses the detection performance at varying levels of the IoU, which represents the capability of such a model to localize cells in a very fine manner. It is possible to attribute the high performance of YOLO26n to the fact that it uses an NMS-free end-to-end implementation that allows better boundary regression and the use of the MuSGD optimizer that allows reaching optimal solutions faster.

Although YOLOv8n had the highest precision (0.9908) and recall (0.9661), its mAP@0.5:0.95 was lower than that of YOLO26n, meaning that instead of being able to detect more objects, the localization accuracy of the object is lower. YOLOv11n showed similar performance in all the measures but was number three in mAP@0.5:0.95.

Table III: Per-Class mAP@0.5:0.95 Comparison:

Class	YOLOv8n	YOLOv11n	YOLO26n
Eosinophil	0.970	0.984	0.968
Lymphocyte	0.971	0.968	0.982
Neutrophil	0.989	0.987	0.988
Platelet	0.749	0.739	0.739

Table III shows the per-class mAP 0.5:0.95 of all three architectures. YOLO26n had the best performance in lymphocyte (0.982) and competitive neutrophil (0.988) and eosinophil (0.968). Platelet detection was the constant challenge of all models and all architectures performed mAP@0.5:0.95 between 0.739 and 0.749, which demonstrated this as a common weakness of the existing detection strategies.

5. Discussion

The experimental findings indicate that the given YOLO26n architecture attains a higher overall performance in blood cell classification than its predecessors. This section explains these findings as per the research objectives and literature available.

1) Superiority of the Proposed Architecture

Second, the MuSGD optimizer combines second-order momentum-based methods with traditional stochastic gradient descent. This hybrid approach captures curvature information for faster convergence while maintaining stability, explaining why the proposed architecture achieved superior results in only 15 epochs compared to 20 epochs required by the earlier version.

The key observation is that the proposed architecture had the best mAP of 0.9234 at 0.5:0.95 which was higher than the previous versions by 0.0038 and 0.0042 respectively. This is a performance enhancement that has been created using two important architectural innovations.

To begin with, the NMS-free end-to-end detection paradigm does not use Non-Maximum Suppression as a post-processing operation. Classical detection models make use of NMS to eliminate duplicate predictions, which creates a non-differentiable operation that disrupts the training-inference gap. The proposed architecture is capable of end-to-end optimization that is genuinely differentiable by learning how to prevent redundant detections, leading to more accurate bounding boxes as evidenced in the improved mAP, namely, mAP at 0.5:0.95.

Second, the optimizer of the MuSGD integrates the second-order momentum-based optimizers and standard stochastic gradient descent. This combination method learns the curvature faster to achieve stability and this is the reason why the suggested architecture obtained better results with just 15 epochs than 20 epochs that the previous one needed.

2) Comparative Analysis with Prior Work

The common outcome of the assessed architectures is that it is difficult to detect plateslets, with a mean of 0.739 to 0.749 at 0.5:0.95. It coincides with the past research that has found the small object recognition as a cornerstone issue in the medical image analysis [5], [9].

This difficulty can be explained by several factors: platelets are significantly smaller than other blood cells, take fewer pixels

at 640 x 640 resolution, constitute only 19.6% of total annotations and may be confused with small debris or artifacts. Fig. 5 indicates that the background class is the most subject to misclassification, which implies that the biggest challenge is not classification, but detection.

3) Platelet Detection Challenge

A consistent finding across all evaluated architectures is the challenge of platelet detection, with mAP@0.5:0.95 values ranging between 0.739 and 0.749. This aligns with previous studies identifying small object detection as a fundamental challenge in medical image analysis [5], [9].

Several factors contribute to this difficulty: platelets are substantially smaller than other blood cells, occupy fewer pixels at 640 × 640 resolution, represent only 19.6% of total annotations, and can be visually similar to small debris or artifacts. The confusion matrix in Fig. 5 reveals that misclassifications primarily occur with the background class, indicating that detection, rather than classification, is the primary challenge.

4) Clinical Implications and Future Directions

The high result of eosinophil, lymphocyte, and neutrophil identification has great clinical implications. The fact that the recall of eosinophil and lymphocyte classes is perfect is especially significant, since false negative might result in the missed diagnosis. The NMS-free architecture minimizes the inference latency, so it can be used in real-time applications, whereas the low number of parameters allows it to be deployed on edge devices in resource-constrained environments.

In spite of these strengths, weaknesses are the evaluation of images where cells are sparsely concentrated, training on one dataset, and the poor platelet detection. The next step in the work will be the improvement of platelet detection using multi-scale feature pyramids and focused data augmentation, testing on new datasets, and using additional cell subtypes of cancer.

6. Conclusion

This paper provided a comparative analysis on YOLOv8n, YOLOv11n and YOLO26n to classify the types of blood cells in peripheral blood smear images. YOLO26n architecture, which has an NMS-free end-to-end detection paradigm and MuSGD optimizer, had better overall performance than its predecessors. The results of per-class analysis indicated that detection of eosinophil, lymphocyte and neutrophil were strong in all architectures, platelet detection was a consistent problem in all architectures.

The results show the effectiveness of end-to-end detection systems in medical image recognition, where YOLO26n has a shorter inference time and a smaller model size which can be deployed in a clinical setting. The next steps of work will involve improving the platelet detection with multi-scale feature pyramids and data augmentation with targeted data and testing it on other datasets and applying the method to subtype classification of cancer cells.

References

Y. Wu, D. Gao, Y. Fang, X. Xu, H. Gao, and Z. Ju, “SDE-YOLO: A Novel Method for Blood Cell Detection,” Biomimetics 2023, Vol. 8, vol. 8, no. 5, Sep. 2023, doi: 10.3390/biomimetics8050404.

D. Zhang, Y. Bu, Q. Chen, S. Cai, and Y. Zhang, “TW-YOLO: An Innovative Blood Cell Detection Model Based on Multi-Scale Feature Fusion,” Sensors 2024, Vol. 24, vol. 24, no. 19, Sep. 2024, doi: 10.3390/s24196168.

H. Sazak and M. Kotan, “Automated Blood Cell Detection and Classification in Microscopic Images Using YOLOv11 and Optimized Weights,” Diagnostics 2025, Vol. 15, vol. 15, no. 1, Dec. 2024, doi: 10.3390/diagnostics15010022.

X. Song and H. Tang, “Blood Cell Target Detection Based on Improved YOLOv5 Algorithm,” Electronics 2024, Vol. 13, vol. 13, no. 24, Dec. 2024, doi: 10.3390/electronics13244992.

Z. Liu, D. Yuan, and G. Zhu, “Automated Blood Cell Detection and Counting Based on Improved Object Detection Algorithm,” Mathematics 2025, Vol. 13, vol. 13, no. 18, Sep. 2025, doi: 10.3390/math13183023.

J. Yang, B. Yang, Z. Li, Y. Yamaguchi, and W. Xiao, “Parameter-Reduced YOLOv8n with GhostConv and C3Ghost for Automated Blood Cell Detection,” Bioengineering 2026, Vol. 13, vol. 13, no. 3, p. 321, Mar. 2026, doi: 10.3390/bioengineering13030321.

M. Feng and J. Xu, “Enhanced Blood Cell Detection in YOLOv11n Using Gradient Accumulation and Loss Reweighting,” Bioengineering 2025, Vol. 12, vol. 12, no. 11, Oct. 2025, doi: 10.3390/bioengineering12111188.

O. El Othmani, A. Mosbah, A. Yahyaoui, A. Bouatay, and R. Dhaouadi, “AI-driven Automated Blood Cell Anomaly Detection: Enhancing Diagnostics and Telehealth in Hematology,” Journal of Imaging 2025, Vol. 11, vol. 11, no. 5, May 2025, doi: 10.3390/jimaging11050157.

R. F. Oybek Kizi, T. P. Theodore Armand, and H. C. Kim, “A Review of Deep Learning Techniques for Leukemia Cancer Classification Based on Blood Smear Images,” Applied Biosciences 2025, Vol. 4, vol. 4, no. 1, Feb. 2025, doi: 10.3390/applbiosci4010009.

Z. Cheng and Y. Li, “Improved YOLOv7 Algorithm for Detecting Bone Marrow Cells,” Sensors 2023, Vol. 23, vol. 23, no. 17, Sep. 2023, doi: 10.3390/s23177640.

X. Z. Wang, G. J. Pan, Z. G. Hu, and A. R. Ge, “A two stage blood cell detection and classification algorithm based on improved YOLOv7 and EfficientNetv2,” Scientific Reports 2025 15:1, vol. 15, no. 1, pp. 8427-, Mar. 2025, doi: 10.1038/s41598-025-91720-7.

S. K. Mondal, M. S. H. Talukder, M. Aljaidi, R. Bin Sulaiman, M. M. S. Tushar, and A. A. Alsuwaylimi, “Efficient blood cell classification from microscopic smear images using U-Net segmentation and a lightweight CNN,” Scientific Reports 2025 15:1, vol. 15, no. 1, pp. 45693-, Dec. 2025, doi: 10.1038/s41598-025-26947-5.

M. Hameed, M. A. Z. Raja, A. Zameer, H. S. Dar, A. S. Alluhaidan, and R. Aziz, “Acute myeloid leukemia classification using ReLViT and detection with YOLO enhanced by adversarial networks on bone marrow images,” Scientific Reports 2025 15:1, vol. 15, no. 1, pp. 32798-, Sep. 2025, doi: 10.1038/s41598-025-17891-5.

C. Shi, D. Zhu, C. Zhou, S. Cheng, and C. Zou, “Gpmb-yolo: a lightweight model for efficient blood cell detection in medical imaging,” Health Inf. Sci. Syst., vol. 12, no. 1, p. 24, Dec. 2024, doi: 10.1007/s13755-024-00285-8.

Y. Mao, H. Zhang, W. Wu, X. Gao, Z. Lin, and J. Lin, “DWS-YOLO: A Lightweight Detector for Blood Cell Detection,” Applied Artificial Intelligence, vol. 38, no. 1, Dec. 2024, doi: 10.1080/08839514.2024.2318673.

X. Chen et al., “NBCDC-YOLOv8: A new framework to improve blood cell detection and classification based on YOLOv8,” IET Computer Vision, vol. 19, no. 1, Jan. 2025, doi: 10.1049/cvi2.12341.

M. Shahzad et al., “Blood cell image segmentation and classification: a systematic review,” PeerJ Comput. Sci., vol. 10, p. e1813, 2024, doi: 10.7717/peerj-cs.1813.

A. Hunt, H. Schulze, K. Samuel, R. B. Fisher, and T. T. Bachmann, “Stain-free artificial intelligence-assisted light microscopy for the identification of blood cells in microfluidic flow,” Frontiers in bioinformatics, vol. 5, 2025, doi: 10.3389/fbinf.2025.1628724.

E. H. I. Eliwa and T. Abd El-Hafeez, “A novel YOLOv11-based framework with dynamic cross-scale context aggregation for high-performance blood cell classification,” Egyptian Informatics Journal, vol. 32, Dec. 2025, doi: 10.1016/j.eij.2025.100851.

Y. Zhao, K. Sun, and M. Cai, “FCE-YOLO: A highly accurate and efficient model for acute lymphoblastic leukemia detection and classification,” Biomed. Signal Process. Control, vol. 113, p. 108962, Mar. 2026, doi: 10.1016/j.bspc.2025.108962.