Leukemia Detection with Instance Segmentation and LeukemiaNet Research Study

Posted on June 8, 2026

Abstract

Leukemia, a malignant blood cancer, poses a significant health threat, particularly when not diagnosed promptly and accurately. This paper developed a comprehensive approach to leukemia cancer detection leveraging advanced techniques and technology. To enhance the model's robustness and ability to generalize, a Generative Adversarial Network (GAN) is employed for image augmentation. The proposed research model features an instance segmentation-based Mask Region-based Convolutional Neural Network (Mask-RCNN). Experimental findings and discussion demonstrate the efficacy of the proposed instance segmentation model (LeukemiaSeg) in detecting normal blood cells and Leukemia cells. Hence, the LeukemiaSeg model achieved a precision of 98.4%, a sensitivity of 97.3%, and an F1 score of 97.9%. The accuracy analysis shows a remarkable accuracy of 98.2%. The blood cell count is performed by a Fully Convolutional Neural Network (FCN), and the Leukemia net classification model performs the classification process. The experiment's findings revealed a remarkable accuracy of 99.85% in the classification process. A comparative evaluation highlights the model's superiority over existing classification approaches. By combining these advanced techniques, the proposed model presents a comprehensive solution for accurate leukemia cancer detection, facilitating timely diagnosis and improved patient outcomes.

Keywords: Leukemia Detection, Medical Imaging, Image Processing, Deep Learning, Instance Segmentation, Generative Adversarial Networks (GANs).

Introduction

1.1 Background of the Study

Blood disorders include a broad range of illnesses that impact the body's production, operation, or destruction of blood cells. Among them, leukemia is a particularly difficult and dangerous kind of blood cancer that is typified by the aberrant growth of white blood cells in the bone marrow. Bone marrow, which generates red blood cells, white blood cells, and platelets, is where blood cell synthesis primarily occurs. Though less severe than leukemia, conditions including anemia, iron deficiency, vitamin B12 deficiency, and folate deficiency are nevertheless quite common and

frequently treatable. The typical cause of anemia is a deficiency in red blood cells or a decrease in the capacity of the blood to carry oxygen. Iron deficiency can cause symptoms such as weariness, weakness, and shortness of breath. It is one of the most prevalent nutritional deficits. Because vitamin B12 and folate are necessary for DNA synthesis and red blood cell development, deficiencies in these nutrients can lead to megaloblastic anemia. Treatment options for these illnesses include dietary modifications, oral or injectable supplements, and sometimes, medicines. However, leukemia is characterized by the bone marrow producing immature or aberrant white blood cells at an uncontrolled rate, which impairs the marrow's capacity to produce healthy cells. Acute or chronic, lymphoid or myeloid leukemia might have different progression patterns and clinical manifestations. While more recent techniques, such as digital pathology and artificial intelligence-assisted diagnosis, are increasingly being investigated to increase accuracy and efficiency, traditional diagnostic procedures depend on peripheral blood smears, bone marrow biopsies, and immunophenotyping. Early and accurate diagnosis is critical because leukemia may progress rapidly and can be fatal without timely treatment.

Leukemia has several forms, the most common of which include Acute Lymphoblastic Leukemia (ALL), Acute Myeloid Leukemia (AML), Chronic Lymphocytic Leukemia (CLL), and Chronic Myeloid Leukemia (CML). ALL is most common among children, whereas AML can affect both children and adults. Chronic leukemias generally progress more slowly but can still cause serious health complications over time. The disease arises from mutations or abnormalities in blood-forming cells, leading to excessive proliferation and impaired differentiation. Genetic factors, radiation exposure, certain chemicals, previous chemotherapy, and some inherited syndromes are associated with increased risk. Nevertheless, in many patients, the exact cause remains unknown.

The diagnosis of leukemia traditionally involves complete blood count, microscopic examination of blood smears, bone marrow aspiration or biopsy, flow cytometry, cytogenetic analysis, and molecular testing. Microscopic analysis is essential for evaluating cell morphology, identifying blasts, and distinguishing abnormal cells from normal leukocytes. However, manual examination is time-consuming, labor-intensive, and dependent on the experience of pathologists. In settings with limited medical resources, shortages of trained specialists may delay diagnosis. Moreover, subtle morphological differences among leukemic cells can create diagnostic challenges and inter-observer variability.

Advancements in medical image processing and machine learning have created opportunities to automate portions of leukemia detection. Deep learning, particularly convolutional neural networks, can learn discriminative features directly from images and classify blood cells with high accuracy. Object detection and segmentation approaches can localize individual cells, identify cell boundaries, and distinguish abnormal cells from background tissue. Instance segmentation is especially valuable because it can detect and segment each cell separately, enabling cell counting and detailed morphological assessment.

The availability of high-quality, diverse datasets remains an important challenge. Medical datasets are

often small because collecting and annotating images requires expert knowledge, time, and resources. Class imbalance may also occur when normal samples greatly outnumber malignant samples or vice versa. Data augmentation can reduce these limitations by generating transformed or synthetic images. Generative Adversarial Networks have shown promise in producing realistic medical images that expand training datasets and improve model generalization.

This research proposes a comprehensive framework for leukemia detection that combines image preprocessing, GAN-based augmentation, instance segmentation, cell counting, and classification. The proposed LeukemiaSeg model is based on Mask R-CNN and is designed to identify and segment normal and leukemic blood cells. A fully convolutional network supports the cell-counting process, while LeukemiaNet performs final classification. The integrated approach aims to improve diagnostic accuracy, reduce manual workload, and provide useful support to medical professionals.

1.2 Problem Statement

Leukemia diagnosis is heavily dependent on manual microscopic examination and specialist interpretation. These processes can be slow, subjective, and difficult to access in resource-constrained settings. Variations in staining, illumination, cell overlap, image quality, and morphology further complicate automated analysis. Existing models may achieve high classification accuracy but often lack precise localization, segmentation, or cell-counting capabilities. Many studies also rely on limited datasets and do not sufficiently address class imbalance or generalization.

Therefore, there is a need for an integrated deep-learning model that can detect, segment, count, and classify blood cells accurately. Such a model should be robust to variations in images and capable of distinguishing normal cells from leukemia cells. The use of GAN-generated images may improve data diversity, while instance segmentation may provide detailed cell-level information. This study addresses these needs through the development of LeukemiaSeg and LeukemiaNet.

1.3 Research Objectives

The major objectives of the study are:

1. To preprocess microscopic blood-cell images for effective feature learning.
2. To use a Generative Adversarial Network for image augmentation and dataset expansion.
3. To develop an instance segmentation model for identifying and separating normal and leukemia cells.
4. To perform automated blood-cell counting using a fully convolutional neural network.
5. To develop a deep-learning classification model for distinguishing leukemia from normal samples.
6. To evaluate the proposed framework using accuracy, precision, sensitivity, specificity, and F1 score.

7. To compare the proposed approach with existing leukemia-detection methods.

1.4 Research Contributions

The primary contribution of this research is the creation of an integrated architecture that combines multiple stages of blood-cell analysis. Rather than limiting the system to image-level classification, the proposed framework performs instance-level segmentation, localization, counting, and classification. The GAN augmentation stage increases the diversity of the training data. The LeukemiaSeg model supplies detailed masks for individual cells, while the FCN and LeukemiaNet support quantitative and diagnostic analysis. The framework may be used as a decision-support tool for medical experts and could be adapted for other blood disorders.

1.5 Organization of the Study

The remainder of the study is structured as follows. The literature review discusses prior research on leukemia diagnosis, medical image processing, deep learning, segmentation, and GAN-based augmentation. The methodology describes the datasets, preprocessing, proposed network architecture, training procedure, and evaluation metrics. The results section presents the experimental outcomes and comparisons. The final sections discuss the findings, limitations, future work, and conclusion.

Literature Review

2.1 Leukemia Diagnosis and Blood-Cell Morphology

Leukemia diagnosis commonly begins with a complete blood count and microscopic evaluation of peripheral blood or bone marrow samples. Morphological characteristics such as cell size, nuclear shape, chromatin pattern, cytoplasmic appearance, and nucleus-to-cytoplasm ratio assist in distinguishing different cell types. Blast cells usually have large nuclei, fine chromatin, visible nucleoli, and limited cytoplasm. Nevertheless, variations among leukemia subtypes and similarities with reactive or immature normal cells can complicate interpretation.

Manual microscopy is widely used because it is relatively accessible and provides important cellular information. However, it depends strongly on the skill and experience of the observer. Fatigue, workload, and subjective judgment can influence results. Automated image-analysis systems have therefore been developed to improve consistency and support specialists.

2.2 Conventional Image-Processing Approaches

Early automated leukemia-detection methods relied on handcrafted features and conventional classifiers. Common preprocessing methods included color normalization, contrast enhancement, noise reduction, and conversion between color spaces. Segmentation was performed using thresholding, watershed transforms, clustering, edge detection, and morphological operations. Features such as area, perimeter, circularity, texture, color statistics, and nucleus-to-cytoplasm ratio were extracted and supplied to support vector machines, k-nearest-neighbor classifiers, decision trees, or artificial neural networks.

These methods provided useful results but were sensitive to image variations, parameter selection, and segmentation errors. Handcrafted features may also fail to capture complex or subtle patterns. Deep learning reduces the need for manual feature design by learning hierarchical representations from data.

2.3 Convolutional Neural Networks

Convolutional neural networks have achieved strong performance in medical-image classification. CNNs apply convolutional filters to learn spatial patterns and increasingly abstract features. Pooling layers reduce spatial dimensions, while fully connected layers or global pooling produce classification outcomes. Transfer learning allows models pretrained on large natural-image datasets to be fine-tuned for medical tasks, which is beneficial when medical datasets are limited.

Architectures such as AlexNet, VGG, ResNet, DenseNet, Inception, and EfficientNet have been applied to leukemia classification. ResNet uses residual connections to facilitate the training of deep networks. DenseNet promotes feature reuse by connecting each layer to subsequent layers. EfficientNet scales network dimensions systematically. Although classification models can identify whether an image contains leukemia, they may not show the exact location or boundaries of abnormal cells.

2.4 Object Detection and Segmentation

Object-detection models identify the location and category of objects using bounding boxes. Common architectures include Faster R-CNN, SSD, and YOLO. Faster R-CNN uses a region proposal network followed by classification and box regression. YOLO performs detection in a single stage and is known for speed. Object detection can locate blood cells but does not provide detailed cell boundaries.

Semantic segmentation labels every pixel according to class but does not necessarily separate neighboring objects of the same class. Instance segmentation combines detection and segmentation by assigning a separate mask to each object. This makes instance segmentation appropriate for

blood-cell images containing multiple touching or overlapping cells.

Mask R-CNN extends Faster R-CNN by adding a mask-prediction branch. The model first extracts features using a backbone network and feature pyramid. A region proposal network generates candidate object regions. ROI Align preserves spatial alignment when extracting features for each proposal. Separate branches predict class labels, bounding boxes, and masks. This architecture has been widely applied in medical imaging due to its ability to provide accurate object-level segmentation.

2.5 Generative Adversarial Networks

GANs consist of a generator and a discriminator trained competitively. The generator attempts to create synthetic images resembling real data, while the discriminator distinguishes real samples from generated ones. Through adversarial learning, the generator gradually produces more realistic outputs. Variants such as DCGAN, conditional GAN, Wasserstein GAN, CycleGAN, and StyleGAN have been used for medical-image synthesis and augmentation.

GAN-based augmentation can increase dataset size, improve class balance, and expose models to a wider range of visual patterns. However, synthetic images must be evaluated carefully because unrealistic artifacts or reduced diversity can affect training. Combining conventional augmentation with GAN-generated samples may provide a balanced strategy.

2.6 Cell Counting

Cell counting can be performed through detection, segmentation, density estimation, or regression. Counting segmented instances is direct but depends on accurate separation of overlapping cells. Fully convolutional networks can produce density maps or pixel-level predictions that support automated counting. Accurate counts may provide clinically relevant information and help quantify abnormal-cell burden.

2.7 Research Gap

A review of prior literature shows that many studies focus on a single task, such as image classification or nucleus segmentation. Fewer approaches integrate augmentation, instance segmentation, counting, and classification within one framework. Dataset limitations, staining differences, overlapping cells, and generalization remain unresolved challenges. The proposed study addresses this gap through a multi-stage model designed for comprehensive leukemia analysis.

Methodology

3.1 Overview of the Proposed Framework

The proposed system contains the following major stages:

1. Collection and organization of microscopic blood-cell images.
2. Image preprocessing and normalization.
3. Data augmentation using conventional transformations and a GAN.
4. Instance segmentation using the LeukemiaSeg model.
5. Blood-cell counting with a fully convolutional network.
6. Image classification using LeukemiaNet.
7. Evaluation and comparison with existing methods.

3.2 Dataset

The dataset contains microscopic images representing normal blood cells and leukemia cells. Images are organized into training, validation, and testing subsets. Expert annotations supply class labels, bounding boxes, and masks for cells used in the segmentation stage. The division of data is performed carefully to avoid leakage between training and testing samples.

The dataset may include variations in illumination, magnification, staining, image resolution, cell density, and background. These differences increase the complexity of the problem but also help evaluate model robustness.

3.3 Preprocessing

Preprocessing standardizes the images and improves their suitability for learning. The major operations include:

- Resizing images to a consistent input dimension.
- Removing noise with suitable filters.
- Applying contrast enhancement where necessary.
- Normalizing pixel intensity values.
- Performing color normalization to reduce staining variations.
- Encoding segmentation masks and class labels.

Image quality is inspected before training. Images with severe corruption, missing labels, or unusable artifacts are excluded or corrected.

3.4 Data Augmentation

Conventional augmentation includes rotation, horizontal and vertical flipping, translation, scaling, cropping, brightness adjustment, and contrast modification. These transformations preserve the underlying biological class while increasing visual diversity.

A GAN is additionally trained to generate synthetic blood-cell images. The generator takes random noise or conditional class information and produces images. The discriminator evaluates whether images are real or synthetic. The adversarial objective encourages realistic generation. Synthetic samples are reviewed for quality and then added to the training set. This method helps address limited sample size and class imbalance.

3.5 Proposed LeukemiaSeg Architecture

LeukemiaSeg is based on Mask R-CNN. Its major components are:

- Backbone network: A deep convolutional feature extractor learns visual features from blood-cell images.
- Feature Pyramid Network: Multi-scale feature maps improve the detection of cells with different sizes.
- Region Proposal Network: Candidate cell regions are generated using anchors and objectness scores.
- ROI Align: Features are extracted without the quantization errors associated with conventional ROI pooling.
- Classification branch: Each region is classified as normal, leukemia, or background.
- Bounding-box branch: The position and dimensions of each detected cell are refined.
- Mask branch: A pixel-level mask is generated for each identified instance.

The total loss combines classification loss, bounding-box regression loss, mask loss, and region-proposal losses. Optimization minimizes the combined objective over the training dataset.

3.6 Fully Convolutional Network for Counting

The FCN receives processed or segmented images and produces pixel-level predictions or density maps. Cell counts are derived from the predicted output. The model consists of convolutional and upsampling layers that preserve spatial information. Counting performance is evaluated by comparing predicted counts with expert-annotated counts.

3.7 LeukemiaNet Classification Model

LeukemiaNet is designed to classify images or extracted cell regions as normal or leukemic. The network contains convolutional blocks, normalization, nonlinear activation, pooling, and regularization layers. Global average pooling reduces the number of parameters, and a final classification layer produces class probabilities.

Transfer learning and fine-tuning may be used to initialize the backbone. Class weights or balanced sampling can reduce the effect of class imbalance. Dropout and augmentation help prevent overfitting.

3.8 Training Procedure

The models are trained using the training subset and evaluated periodically on the validation subset. Hyperparameters include learning rate, batch size, optimizer, number of epochs, confidence threshold, and non-maximum suppression threshold. Early stopping and learning-rate scheduling are used when validation performance stops improving.

The optimizer updates network weights based on gradients of the loss. Checkpoints preserve the best-performing model. The independent testing subset is used only after training and model selection are complete.

3.9 Evaluation Metrics

The following measures evaluate classification and detection performance:

Accuracy measures the proportion of correct predictions.

Accuracy = (TP + TN) / (TP + TN + FP + FN)

Precision measures how many predicted positive samples are truly positive.

Precision = TP / (TP + FP)

Sensitivity or recall measures how many actual positive samples are detected.

Sensitivity = TP / (TP + FN)

Specificity measures how many actual negative samples are correctly identified.

Specificity = TN / (TN + FP)

F1 score is the harmonic mean of precision and recall.

F1 = 2 × Precision × Recall / (Precision + Recall)

Segmentation performance can also be evaluated with Intersection over Union and Dice coefficient. Average precision across IoU thresholds measures instance-segmentation quality.

3.10 Experimental Environment

The experiments are carried out using a deep-learning framework and GPU-enabled computing environment. Libraries for image processing, numerical computation, visualization, and model evaluation support implementation. Random seeds and consistent data splits improve reproducibility.

Results and Discussion

4.1 Instance Segmentation Results

The proposed LeukemiaSeg model successfully detects and segments normal and leukemia cells. The model achieved a precision of 98.4%, sensitivity of 97.3%, F1 score of 97.9%, and accuracy of 98.2%. These results indicate that the model can identify abnormal cells accurately while maintaining a low false-positive rate.

Instance masks provide more detailed information than image-level labels or bounding boxes alone. They allow examination of cell boundaries, shape, and spatial distribution. Accurate separation of individual cells also supports counting and reduces errors caused by touching cells.

4.2 Classification Results

LeukemiaNet achieved an accuracy of 99.85% during the classification experiment. The result demonstrates strong discrimination between normal and leukemia images. GAN augmentation and preprocessing contributed to improved generalization, while the network architecture captured important morphological patterns.

The high accuracy should nevertheless be interpreted alongside precision, recall, specificity, and testing conditions. Medical diagnostic systems must avoid both false negatives and false positives. Independent clinical validation is needed before real-world deployment.

4.3 Effect of GAN Augmentation

GAN-generated images increased the variability of training data and reduced overfitting. Models trained with augmented data performed better on unseen samples than those trained only on the original dataset. The synthetic images represented a range of cell appearances and backgrounds. Quality inspection remained necessary to prevent unrealistic artifacts from influencing the model.

4.4 Cell Counting

The FCN provided reliable estimates of the number of cells in microscopic images. Counting results were consistent with segmentation outputs and expert annotations. The combination of segmentation and FCN-based counting provides complementary information. Segmentation supplies object-level masks, while the FCN can support counting in dense or complex regions.

4.5 Comparative Analysis

The proposed framework outperformed several existing leukemia-classification and segmentation approaches. The improvement is attributed to the combination of GAN augmentation, multi-scale instance segmentation, automated counting, and specialized classification. Existing methods based solely on handcrafted features or image-level classification provide less detailed output.

The performance also shows the value of integrating multiple deep-learning tasks. A comprehensive pipeline can supply diagnosis, localization, morphology, and quantitative information rather than a single class prediction.

4.6 Clinical Relevance

An automated leukemia-analysis system can assist pathologists by screening images, highlighting suspicious cells, and providing preliminary counts. It may reduce workload and support rapid evaluation, particularly in regions with limited access to specialists. The model is intended as a decision-support system rather than a replacement for medical professionals. Clinical diagnosis should incorporate patient history, laboratory tests, molecular findings, and expert judgment.

4.7 Challenges and Limitations

Several limitations should be considered. The performance of deep-learning models depends on the quality and representativeness of training data. Images from different laboratories may vary in staining protocols, microscopes, cameras, and preparation methods. A model trained on one dataset

may not generalize fully to another institution.

Expert annotation of masks is expensive and time-consuming. Errors or inconsistencies in annotations can affect training. Synthetic images may introduce artifacts. The system may also face difficulty with severely overlapping cells, poor image quality, rare subtypes, or cells with ambiguous morphology.

The experimental results are strong, but prospective clinical studies and multi-center validation are necessary. Model explainability, calibration, data privacy, computational requirements, and integration with laboratory workflows also require further investigation.

Conclusion

This study presents an integrated deep-learning framework for leukemia detection from microscopic blood-cell images. The approach combines preprocessing, GAN-based augmentation, Mask R-CNN instance segmentation, FCN-based cell counting, and LeukemiaNet classification. LeukemiaSeg achieved 98.4% precision, 97.3% sensitivity, 97.9% F1 score, and 98.2% accuracy. LeukemiaNet achieved a classification accuracy of 99.85%.

The findings demonstrate that deep learning can provide accurate and detailed analysis of normal and leukemic blood cells. Instance segmentation identifies individual cells and boundaries, while automated counting and classification add clinically useful information. GAN augmentation helps overcome limited data and improves model robustness.

The proposed system has the potential to support timely and consistent leukemia diagnosis. However, it should undergo broader validation with diverse datasets and clinical environments before deployment. Future research should examine multi-center data, leukemia subtype classification, explainable artificial intelligence, real-time implementation, and integration with clinical decision-support systems.

Future Work

Future work may focus on:

- Increasing the size and diversity of annotated datasets.
- Evaluating the model across multiple hospitals and laboratories.
- Classifying different leukemia subtypes and stages.
- Improving segmentation of overlapping and clustered cells.
- Applying self-supervised and semi-supervised learning to reduce annotation requirements.
- Developing explainable visualizations for medical experts.
- Optimizing the architecture for mobile or low-resource devices.
- Combining image findings with clinical, genetic, and laboratory data.

References

Abbas, N., Mohamad, N., Mansor, M., & Kadir, M. F. A. (2021). A review of deep learning methods in leukemia diagnosis. Journal of Medical Systems, 45, 1–18.

Ahmed, N., Yigit, A., Isik, Z., & Alpkocak, A. (2019). Identification of leukemia subtypes from microscopic images using convolutional neural networks. Diagnostics, 9(3), 104.

Al-jaboriy, S. S., Sjarif, N. N. A., Chuprat, S., & Abduallah, W. M. (2019). Acute lymphoblastic leukemia segmentation using local pixel information. Pattern Recognition Letters, 125, 85–90.

Anilkumar, K. K., Manoj, V. J., & Sagi, T. M. (2017). A survey on image segmentation of blood and bone marrow smear images for automated detection of leukemia. International Journal of Computer Applications, 167(5), 22–27.

Badrinarayanan, V., Kendall, A., & Cipolla, R. (2017). SegNet: A deep convolutional encoder-decoder architecture for image segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence, 39(12), 2481–2495.

Çinar, A., & Tuncer, S. A. (2021). Classification of lymphocytes, monocytes, eosinophils, and neutrophils on white blood cell images using hybrid AlexNet-GoogleNet-SVM. SN Applied Sciences, 3, 1–11.

Faroq, S., & Raza, B. (2019). Automatic leukemia detection and classification using microscopic images. Biomedical Signal Processing and Control, 52, 142–150.

Girshick, R. (2015). Fast R-CNN. Proceedings of the IEEE International Conference on Computer Vision, 1440–1448.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. Advances in Neural Information Processing Systems, 27, 2672–2680.

He, K., Gkioxari, G., Dollár, P., & Girshick, R. (2017). Mask R-CNN. Proceedings of the IEEE International Conference on Computer Vision, 2961–2969.

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 770–778.

Huang, G., Liu, Z., Van Der Maaten, L., & Weinberger, K. Q. (2017). Densely connected convolutional networks. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 4700–4708.

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional

neural networks. Advances in Neural Information Processing Systems, 25, 1097–1105.

Labati, R. D., Piuri, V., & Scotti, F. (2011). All-IDB: The acute lymphoblastic leukemia image database for image processing. Proceedings of the IEEE International Conference on Image Processing, 2045–2048.

Long, J., Shelhamer, E., & Darrell, T. (2015). Fully convolutional networks for semantic segmentation. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 3431–3440.

Matek, C., Schwarz, S., Spiekermann, K., & Marr, C. (2019). Human-level recognition of blast cells in acute myeloid leukemia with convolutional neural networks. Nature Machine Intelligence, 1(11), 538–544.

Mohapatra, S., Patra, D., & Satpathy, S. (2014). An ensemble classifier system for early diagnosis of acute lymphoblastic leukemia in blood microscopic images. Neural Computing and Applications, 24, 1887–1904.

Rajaraman, S., Antani, S. K., Poostchi, M., Silamut, K., Hossain, M. A., Maude, R. J., Jaeger, S., & Thoma, G. R. (2018). Pre-trained convolutional neural networks as feature extractors toward improved malaria parasite detection in thin blood smear images. PeerJ, 6, e4568.

Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You only look once: Unified, real-time object detection. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 779–788.

Ren, S., He, K., Girshick, R., & Sun, J. (2015). Faster R-CNN: Towards real-time object detection with region proposal networks. Advances in Neural Information Processing Systems, 28, 91–99.

Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. Medical Image Computing and Computer-Assisted Intervention, 234–241.

Shafique, S., & Tehsin, S. (2018). Acute lymphoblastic leukemia detection and classification of its subtypes using pretrained deep convolutional neural networks. Technology in Cancer Research & Treatment, 17, 1–7.

Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. International Conference on Learning Representations.

Tan, M., & Le, Q. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. International Conference on Machine Learning, 6105–6114.

Togacar, M., Ergen, B., & Comert, Z. (2020). Classification of white blood cells using deep features obtained from convolutional neural network models based on the combination of feature selection methods. Applied Soft Computing, 97, 106810.