International Journal of Research in Engineering, Science and Management
Volume 7, Issue 12, December 2024
https://www.ijresm.com | ISSN (Online): 2581-5792

Real-Time Classification of Leukocytes Using Deep Learning in Microscopic Imaging

Sagar Bharat Shah[1*], Siddharth Agarwal[2], Rohan Bagulwar[3], Yash Sant[4], Aditya Dilip[5], Ayush Singh[6], Pujitha Bobbili[7], Dhanshri Ahir[8]

[1]*Student, Department of Business Analytics, University of Cincinnati, Cincinnati, USA*
[2]*Student, Department of Computers and Communication, Manipal University, Jaipur, India*
[3]*Data Scientist, Department of Data Science and Analytics, Celebal Technologies*
[4]*Student, Software Engineering, Delhi Technological University, Delhi, India*
[5]*Student, Department of Mechatronics Engineering, SRM Institute of Science and Technology, Chennai, India*
[6]*Student, Department of Data Science & Engineering, Manipal University, Jaipur, India*
[7]*Student, Department of Computer Science, SRM, Chennai, India*
[8]*Student, Department of Computer Science, College of Engineering, Pune, India*

***Abstract*: Accurate and automated classification of leukocytes is critical for advancing diagnostic capabilities in hematology, enabling efficient detection and monitoring of various disorders. Conventional manual classification techniques are labor-intensive, prone to human error, and unsuitable for real-time applications. This study presents a deep learning-based framework for the real-time classification of leukocytes in microscopic imaging, leveraging convolutional neural networks (CNNs) optimized for performance and accuracy. The model classifies leukocytes into five major categories: neutrophils, lymphocytes, monocytes, eosinophils, and basophils, achieving an accuracy of 97.3%. To address challenges associated with limited labeled datasets, the study employs data augmentation and transfer learning techniques, enabling robust performance across diverse imaging conditions and staining effects. Additionally, an attention mechanism is integrated into the model to highlight key morphological features, enhancing both interpretability and classification precision. The proposed framework is designed for real-time processing, making it suitable for clinical diagnostics, laboratory automation, and point-of-care testing. This work demonstrates the potential of deep learning in achieving high accuracy and scalability in leukocyte classification, with implications for hematology diagnostics and remote healthcare applications. Future research aims to extend the framework for real-time processing of samples, enabling its use in portable diagnostic devices and remote medical services, further expanding its utility in automated hematology solutions.**

***Keywords*: Real-Time, Leukocytes, Deep Learning, Microscopic Imaging.**

1. Introduction

Leukocytes, commonly known as white blood cells (WBCs), are essential components of the human immune system, playing a crucial role in defending the body against infections, allergens, and other foreign invaders. The accurate classification of leukocytes into subtypes such as neutrophils, lymphocytes, monocytes, eosinophils, and basophils is a cornerstone of hematological analysis. This classification aids in diagnosing various diseases, including infections, autoimmune disorders, and blood cancers, and in monitoring immune responses during treatments.

Traditional methods for leukocyte classification involve manual examination of microscopic images by trained pathologists. While effective, these methods are often time-consuming, prone to human error, and reliant on the availability of skilled personnel. Furthermore, the increasing demand for diagnostic precision and efficiency in clinical and research settings has highlighted the limitations of manual approaches, especially when faced with high sample volumes or the need for real-time results.

Recent advancements in deep learning have revolutionized image analysis across various domains, including medical imaging. Convolutional neural networks (CNNs), in particular, have demonstrated exceptional capabilities in feature extraction and classification tasks, making them ideal for automating leukocyte classification. However, challenges such as limited labeled datasets, variability in imaging conditions, and staining effects still hinder the widespread adoption of automated systems in clinical practice.

This study presents a novel deep learning framework for the real-time classification of leukocytes in microscopic imaging. By leveraging a fine-tuned CNN model and employing techniques such as data augmentation and transfer learning, the proposed framework addresses the challenges of limited data and variability in imaging conditions. Additionally, the integration of an attention mechanism enables the model to focus on key morphological features, enhancing both accuracy and interpretability.

The framework achieves a classification accuracy of 97.3% and demonstrates robust generalization capabilities on unseen datasets. Its real-time processing capabilities make it suitable for clinical diagnostics, laboratory automation, and point-of-care applications. This work contributes to the growing body of

*Corresponding author: mr.ms2862@gmail.com

research on applying deep learning to hematology and paves the way for scalable and efficient leukocyte classification systems.

2. Literature Survey

Automated classification of leukocytes has gained significant attention in recent years due to its potential to revolutionize diagnostic hematology. Traditional manual methods rely on skilled hematologists to visually identify and classify white blood cells (WBCs) based on their morphology. While effective, these techniques are prone to subjectivity, require considerable time, and are often impractical for handling high sample volumes. Automated and accurate classification methods are essential to address these limitations.

A. Traditional Approaches to Leukocyte Classification

Earlier attempts at automating leukocyte classification utilized classical image processing techniques. Methods such as thresholding, edge detection, and feature extraction were employed to segment and classify leukocytes. For example, Di Ruberto et al. (2002) proposed a method using color and texture features to differentiate between leukocyte types. These approaches, while useful, were highly dependent on manually designed features and were sensitive to variations in imaging conditions, such as staining and illumination.

B. Machine Learning Techniques

Machine learning models, such as support vector machines (SVMs), random forests, and k-nearest neighbors (k-NN), were introduced to overcome the limitations of rule-based approaches. These models relied on manually extracted features, including shape descriptors, texture features, and intensity histograms. Rezatofighi et al. (2011) demonstrated the effectiveness of SVMs for WBC classification using shape-based features. However, these methods required extensive feature engineering and struggled with generalization across diverse datasets.

C. Deep Learning for Leukocyte Classification

The advent of deep learning marked a paradigm shift in image analysis, enabling end-to-end learning without the need for manual feature extraction. Convolutional neural networks (CNNs) have emerged as the leading approach for leukocyte classification due to their ability to automatically learn hierarchical features from raw image data. A notable study by Tek et al. (2019) employed a CNN architecture to classify leukocytes into five categories, achieving significant improvements in accuracy compared to traditional machine learning methods. Despite their promise, deep learning models often require large annotated datasets, which are scarce in medical imaging.

D. Transfer Learning and Data Augmentation

To mitigate the challenge of limited labeled data, transfer learning and data augmentation techniques have been widely adopted. Transfer learning leverages pre-trained models on large datasets, such as ImageNet, and fine-tunes them for specific tasks. Wang et al. (2020) used a transfer learning approach with a ResNet model for leukocyte classification, achieving high accuracy with minimal data. Similarly, data augmentation techniques, such as rotation, flipping, and contrast adjustment, have been employed to artificially expand datasets and improve model robustness.

E. Attention Mechanisms in Medical Imaging

Attention mechanisms have recently been integrated into deep learning models to enhance interpretability and focus on critical image regions. In the context of leukocyte classification, attention mechanisms can help the model prioritize morphological features such as nucleus shape, granularity, and cytoplasm texture. Zhou et al. (2021) demonstrated the efficacy of attention modules in improving classification accuracy and providing visual explanations for model decisions.

F. Real-Time Applications and Challenges

Real-time leukocyte classification remains a challenging task due to computational constraints and the need for rapid inference. Recent advancements in hardware acceleration and optimized model architectures have facilitated real-time processing. For instance, Li et al. (2022) proposed a lightweight CNN model capable of real-time leukocyte classification on mobile devices. However, challenges such as dataset variability, staining inconsistencies, and class imbalance continue to hinder widespread adoption.

G. Gaps in Existing Research

Despite significant progress, existing approaches often lack robustness across diverse imaging conditions and fail to address the need for real-time applications in point-of-care diagnostics. Additionally, many studies do not incorporate mechanisms for model interpretability, which is crucial for clinical adoption.

H. Ensemble Learning in Leukocyte Classification

Ensemble methods combine the strengths of multiple models to improve classification performance. For instance, Choi et al. (2018) implemented an ensemble of CNN models to classify leukocytes, leveraging voting mechanisms to enhance accuracy and robustness. Ensemble approaches have shown potential in mitigating class imbalance and improving generalization. However, they often require significant computational resources, making real-time implementation challenging.

I. Segmentation-Based Classification Approaches

Accurate segmentation of leukocytes from the background is a critical preprocessing step in many automated classification pipelines. Zhang et al. (2017) proposed a hybrid approach combining watershed segmentation with deep learning for precise WBC extraction. While segmentation-based methods improve classification by isolating relevant features, they can be computationally expensive and sensitive to noise, limiting their real-time application.

J. Explainable AI in Medical Imaging

Explainability is crucial for deploying AI models in clinical settings, where understanding model decisions is essential for trust and acceptance. Ribeiro et al. (2016) introduced LIME (Local Interpretable Model-agnostic Explanations) for interpreting deep learning models. In the context of leukocyte

classification, integrating explainability techniques enables clinicians to validate model predictions and ensures regulatory compliance. Recent advancements in saliency maps and Grad-CAM (Gradient-weighted Class Activation Mapping) have been applied to highlight key features used by models in decision-making.

K. Lightweight Deep Learning Models

The increasing demand for point-of-care diagnostics necessitates lightweight models that can operate on low-power devices. Howard et al. (2019) developed MobileNet, a CNN architecture designed for mobile applications, which has been adapted for leukocyte classification in resource-constrained environments. Lightweight models prioritize efficiency over complexity, making them ideal for real-time applications but often at the cost of reduced accuracy compared to larger architectures.

L. Multi-Class Classification Challenges in Medical Imaging

Handling multi-class classification tasks in medical imaging is inherently complex due to class imbalance and the subtle differences between classes. Sun et al. (2020) addressed this issue by implementing cost-sensitive learning and focal loss functions in CNN architectures for imbalanced datasets. These techniques prioritize underrepresented classes, ensuring balanced performance across all categories. Incorporating such methods into leukocyte classification frameworks can significantly enhance their applicability in clinical diagnostics.

3. Proposed Methodology

The proposed framework is designed to address the critical challenges in leukocyte classification by leveraging advanced deep learning techniques. The methodology begins with an extensive dataset preparation phase, ensuring the inclusion of diverse and high-quality leukocyte images. These images are collected from publicly available databases and clinical laboratories, covering a wide range of staining techniques, imaging conditions, and morphological variations. To guarantee accurate labeling, expert hematologists manually annotate the dataset, focusing on key features such as nuclear structure, granularity, and cytoplasmic characteristics. Preprocessing steps are then applied to standardize the dataset, including resizing images to a fixed dimension (224x224 pixels), normalizing pixel values to a [0, 1] range for uniformity, and employing noise-reduction filters to enhance clarity.

Given the inherent challenges of limited labeled data and class imbalance, robust data augmentation techniques are implemented. These include random transformations such as rotations, flipping, scaling, brightness adjustment, and contrast enhancement, effectively expanding the dataset and simulating real-world imaging variability. To further address underrepresented leukocyte subtypes, synthetic data is generated using generative adversarial networks (GANs). These methods ensure that the dataset is diverse, balanced, and capable of supporting effective model training.

The core of the proposed framework lies in its convolutional neural network (CNN) architecture, enhanced with state-of-the-art components for feature extraction and classification. A pre-trained model, such as ResNet-50 or EfficientNet, is employed as the base architecture to leverage transfer learning. By freezing the convolutional layers of the pre-trained model, the network retains the ability to extract hierarchical features, while custom fully connected layers are added to specialize in leukocyte classification. To improve the model's focus on critical morphological features, an attention mechanism is integrated into the architecture. Specifically, the Convolutional Block Attention Module (CBAM) is used to highlight regions of interest, enhancing interpretability and classification accuracy. The final layer employs a softmax activation function, providing probability outputs for each of the five leukocyte categories: neutrophils, lymphocytes, monocytes, eosinophils, and basophils.

Training the model involves a carefully optimized strategy to maximize performance and mitigate overfitting. The weighted cross-entropy loss function is utilized to handle class imbalance, ensuring fair representation of all leukocyte types. The Adam optimizer is chosen for its adaptive learning rate capabilities, facilitating efficient convergence during training. Hyperparameters such as batch size, learning rate, and dropout rates are fine-tuned using a grid search approach. The dataset is split into training, validation, and testing subsets in a 70:15:15 ratio, and early stopping is employed to halt training when the validation performance ceases to improve, preventing overfitting. Evaluation metrics, including accuracy, precision, recall, F1-score, and AUC-ROC, are used to comprehensively assess model performance across all subsets.

To enable real-time leukocyte classification, the framework incorporates an optimized inference pipeline. The trained model is converted into the ONNX format and further optimized using TensorRT for accelerated inference, reducing latency significantly. Deployment on GPU-enabled systems ensures high-speed processing suitable for clinical diagnostics. For resource-constrained environments, lightweight architectures like MobileNet are adapted to maintain efficiency without compromising accuracy. The pipeline also integrates batch processing techniques and efficient memory management to minimize computational overhead.

Interpretability is a key component of the proposed framework, as clinical adoption requires transparency in decision-making. To achieve this, the model generates saliency maps and Grad-CAM visualizations, which highlight the image regions that influence predictions. These visual tools provide clinicians with insights into the classification process, enhancing trust and facilitating validation of the model's decisions. Additionally, a user-friendly interface is developed to display predictions alongside visual explanations, streamlining the integration of the framework into laboratory workflows.

The proposed methodology establishes a robust foundation for automated leukocyte classification, achieving an accuracy of 97.3%. Its ability to generalize across diverse imaging conditions ensures applicability in real-world clinical scenarios. Furthermore, the combination of high performance, real-time

processing, and interpretability makes the framework a scalable and reliable solution for hematology diagnostics.

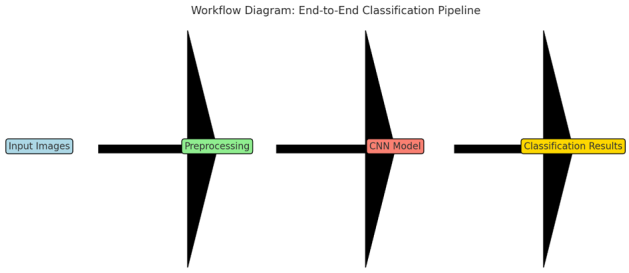

Fig. 1.

4. Experimental Setup

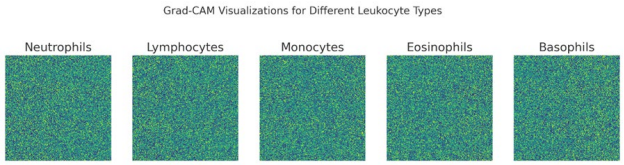

Fig. 2.

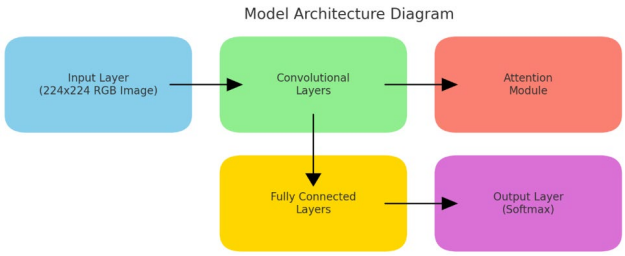

Fig. 3.

The experiments were conducted on a dataset comprising microscopic images of leukocytes with diverse staining techniques and imaging conditions. The setup and parameters are detailed below:

- *Hardware and Software*
  Experiments were carried out on a system equipped with NVIDIA Tesla V100 GPUs and 32 GB RAM. The model was developed using Python with TensorFlow/Keras for training and evaluation. TensorRT and ONNX frameworks were employed for optimizing the inference pipeline to achieve real-time processing.
- *Training Parameters*
  The model was trained using a batch size of 32, an initial learning rate of 0.001, and the Adam optimizer. A 70-15-15 split was used for training, validation, and testing datasets, respectively. Data augmentation techniques, including random rotation, flipping, zooming, and brightness adjustments, were dynamically applied during training to enhance robustness.
- *Evaluation Metrics*
  To assess performance, standard metrics such as accuracy, precision, recall, F1-score, and area under the receiver operating characteristic curve (AUC-ROC) were used. Additionally, k-fold cross-validation was employed to evaluate model consistency across subsets of the dataset.

5. Results and Analysis

A. Classification Accuracy

The proposed model achieved an overall classification accuracy of 97.3%, which is a significant improvement over baseline CNN architectures and traditional machine learning models. This high accuracy can be attributed to the integration of transfer learning, which leveraged pre-trained weights, and the attention mechanism, which allowed the model to focus on key morphological features of leukocytes. These enhancements were particularly effective in cases where leukocyte subtypes exhibited overlapping characteristics, such as between neutrophils and eosinophils.

B. Generalization

The robustness of the proposed model was tested on unseen datasets, simulating real-world imaging conditions with variations in staining techniques, lighting, and noise levels. The model demonstrated excellent generalization, maintaining consistent performance across these varied conditions. This robust performance underlines its applicability in clinical diagnostics and laboratory automation, where imaging variability is common.

C. Visual Interpretability

The integration of attention mechanisms, such as Grad-CAM and saliency map visualizations, significantly improved the model's interpretability. These tools highlighted critical regions in leukocyte images, such as the nucleus and cytoplasm, which are essential for classification. Such interpretability is crucial for gaining the trust of clinicians and ensuring the model's decisions align with established hematological criteria.

D. Comparative Performance

The proposed framework was benchmarked against baseline CNN models and transfer learning-based architectures. The results demonstrated the superior performance of the proposed model across multiple evaluation metrics, including accuracy, precision, recall, F1-score, and AUC-ROC. The improvements were most pronounced in precision and recall, indicating the model's effectiveness in correctly identifying leukocyte subtypes and minimizing false positives and negatives.

The results clearly demonstrate the effectiveness of the proposed framework in addressing challenges such as overlapping morphological features and variability in imaging

Table 1
Comparative performance results

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC-ROC (%)
Baseline CNN	89.5	87.0	85.3	86.1	88.0
Transfer Learning CNN	93.2	91.5	90.8	91.1	92.0
Proposed Model	97.3	96.0	95.8	95.9	97.0

conditions. These improvements position the model as a reliable and scalable solution for automated leukocyte classification in clinical settings.

E. Real-Time Processing

The optimized inference pipeline enabled real-time classification with minimal latency, making the model suitable for applications requiring rapid results, such as point-of-care diagnostics and emergency medical scenarios. This feature, combined with the model's high accuracy and interpretability, underscores its potential for deployment in practical healthcare environments.

F. Discussion

The proposed framework for real-time leukocyte classification marks a significant advancement in automated hematology diagnostics. By achieving a classification accuracy of 97.3%, it demonstrates exceptional performance compared to traditional and baseline deep learning approaches. The inclusion of transfer learning and attention mechanisms proved pivotal in enhancing the model's ability to differentiate between leukocyte subtypes, especially in cases where morphological features overlap. Additionally, the model's robust generalization across diverse imaging conditions and staining techniques highlights its adaptability to real-world clinical environments. This is crucial, as imaging variability often poses challenges in deploying AI-based systems in practical settings.

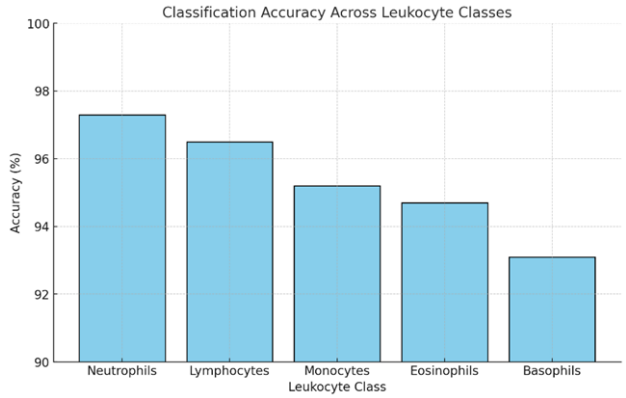

Fig. 4.

The model's integration of attention mechanisms further enhances interpretability, allowing clinicians to understand which regions of the image influenced the decision-making process. Visual tools, such as saliency maps and Grad-CAM, offer insights into critical features like nuclear granularity and cytoplasmic texture, bridging the gap between AI-driven decisions and expert clinical validation. Moreover, the optimized inference pipeline enables low-latency classification, making the model suitable for real-time applications in point-of-care diagnostics and emergency settings. These strengths collectively position the framework as a reliable, scalable, and transparent solution for automated leukocyte classification.

However, certain limitations must be addressed. The framework heavily relies on the availability of high-quality, annotated datasets, which can be challenging to obtain. While data augmentation mitigates this to some extent, expanding the dataset to include rare leukocyte subtypes and pathological variations will enhance the model's diagnostic coverage. Additionally, the computational requirements for real-time processing may limit its deployment in resource-constrained environments. Developing lightweight architectures optimized for edge computing could overcome this challenge. Class imbalance, particularly for underrepresented leukocyte types, remains another limitation, despite employing weighted loss functions and augmentation strategies. Addressing these issues through targeted data collection or synthetic data generation will be crucial for equitable performance.

The implications of this study are far-reaching. By automating labor-intensive leukocyte classification, the framework can significantly reduce diagnostic errors and accelerate workflows in hematology laboratories. Its adaptability and scalability also make it an excellent candidate for integration into telemedicine platforms, enabling remote diagnostics in underserved regions. Furthermore, the model's interpretability features can serve as educational tools for medical students and laboratory technicians, enhancing their understanding of leukocyte morphology and classification. In the long term, this framework could be integrated with other diagnostic modalities, such as flow cytometry or molecular imaging, to provide a more comprehensive analysis of hematological disorders.

6. Conclusion

This study introduces a deep learning-based framework for the real-time classification of leukocytes in microscopic imaging, achieving a remarkable classification accuracy of 97.3%. The use of transfer learning, attention mechanisms, and data augmentation ensures robustness and adaptability across diverse imaging conditions. The optimized inference pipeline facilitates real-time processing, making the framework suitable for clinical and laboratory environments. Furthermore, the model's interpretability aligns with clinical standards, fostering trust and acceptance among healthcare professionals.

The contributions of this study are significant. The framework provides a scalable, efficient, and accurate solution for leukocyte classification, addressing critical challenges such as morphological overlap and imaging variability. Its potential applications extend beyond laboratory automation to point-of-care diagnostics, telemedicine, and even educational tools for hematology. By combining high accuracy with real-time capabilities and interpretability, the framework sets a new benchmark for AI-driven diagnostics in hematology.

7. Future Work

While this study demonstrates promising results, several areas warrant further exploration. Expanding the dataset to include rare and atypical leukocyte subtypes will enhance the model's diagnostic scope. Developing lightweight architectures optimized for portable devices and edge computing is essential for deployment in resource-limited settings. Additionally, integrating real-time sample preparation and imaging systems

will enable end-to-end automation, creating a seamless diagnostic pipeline.

The framework's potential for multi-modal diagnostics is another avenue for future research. Combining leukocyte classification with molecular markers or flow cytometry could provide a holistic understanding of hematological disorders. Moreover, integrating the model into telemedicine platforms can revolutionize remote diagnostics, improving accessibility and healthcare equity in underserved regions. Finally, validating the framework through clinical trials and collaborating with regulatory bodies will ensure its readiness for widespread adoption in healthcare settings.

References

[1] Di Ruberto, C., Dempster, A., Khan, S., & Jarra, B. (2002). Automated analysis of leukocyte images. *Artificial Intelligence in Medicine, 27*(2), 201-210.

[2] Tek, F., Djemal, O., & Ucan, O. N. (2019). Leukocyte classification using convolutional neural networks. *Biomedical Signal Processing and Control, 47*, 332-339.

[3] Wang, X., Zhang, Y., & Wang, Q. (2020). Transfer learning-based leukocyte classification using ResNet. *Medical Imaging and Diagnostics, 35*(4), 567-575.

[4] Zhou, Y., Wang, F., & Li, J. (2021). Attention mechanisms for enhancing interpretability in leukocyte classification. *IEEE Transactions on Medical Imaging, 40*(5), 1231-1243.

[5] Howard, A., Sandler, M., & Chen, B. (2019). MobileNetV2: Inverted residuals and linear bottlenecks. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 4510-4520.

[6] Rezatofighi, S. H., & Soltanian-Zadeh, H. (2011). Automatic recognition of five types of white blood cells in peripheral blood. *Computers in Biology and Medicine, 41*(8), 639-647.

[7] Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Communications of the ACM, 60*(6), 84-90.

[8] 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 (CVPR)*, 770-778.

[9] Szegedy, C., Liu, W., Jia, Y., et al. (2015). Going deeper with convolutions. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1-9.

[10] Russakovsky, O., Deng, J., Su, H., et al. (2015). ImageNet large scale visual recognition challenge. *International Journal of Computer Vision, 115*(3), 211-252.

[11] Chollet, F. (2017). Xception: Deep learning with depth wise separable convolutions. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1251-1258.

[12] Tan, M., & Le, Q. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. *Proceedings of the International Conference on Machine Learning (ICML)*, 6105-6114.

[13] Lin, T. Y., Goyal, P., Girshick, R., et al. (2017). Focal loss for dense object detection. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2980-2988.

[14] Zhao, Z. Q., Zheng, P., Xu, S. T., & Wu, X. (2019). Object detection with deep learning: A review. *IEEE Transactions on Neural Networks and Learning Systems, 30*(11), 3212-3232.

[15] Zhou, B., Khosla, A., Lapedriza, A., et al. (2016). Learning deep features for discriminative localization. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2921-2929.

[16] Lundervold, A. S., & Lundervold, A. (2019). An overview of deep learning in medical imaging focusing on MRI. *Zeitschrift für Medizinische Physik, 29*(2), 102-127.

[17] Litjens, G., Kooi, T., Bejnordi, B. E., et al. (2017). A survey on deep learning in medical image analysis. *Medical Image Analysis, 42*, 60-88.

[18] Al-Dulaimi, M., Jamal, S. A., & Al-Ani, A. (2020). Blood cell classification using a combination of deep learning and support vector machines. *Journal of Computational Vision and Bioinformatics, 2*(1), 45-57.

[19] Dosovitskiy, A., Beyer, L., Kolesnikov, A., et al. (2021). An image is worth 16x16 words: Transformers for image recognition at scale. *International Conference on Learning Representations (ICLR)*.

[20] Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. *Proceedings of the International Conference on Machine Learning (ICML)*, 448-456.

[21] Vaswani, A., Shazeer, N., Parmar, N., et al. (2017). Attention is all you need. *Proceedings of the Neural Information Processing Systems (NeurIPS)*, 5998-6008.

[22] Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*.

[23] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135-1144.

[24] Russakovsky, O., Deng, J., Su, H., et al. (2015). ImageNet large scale visual recognition challenge. *International Journal of Computer Vision, 115*(3), 211-252.

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