

A HYBRID CNN-GNN FRAMEWORK FOR FINE-GRAINED PERIPHERAL BLOOD CELL MORPHOLOGY CLASSIFICATION

Toufique Ahmed¹, Sadia Bhutto^{*2}, Iftikhar A. Koondhar³, Zulfiqar Ali Bhutto⁴, Kirtika⁵

^{1,*2,3,4,5}Department of Computer Systems Engineering, Quaid-e-Awam University of Engineering, Science and Technology (QUEST), Nawabshah, Pakistan

¹toufiqueahmed204@gmail.com, ²sadiabhutto999@gmail.com, ³iftikhar.ahmed@quest.edu.pk
⁴mirzulfiqarb@gmail.com, ⁵kirtika1428@gmail.com

DOI: <https://doi.org/10.5281/zenodo.23275291>

Keywords

Peripheral blood cells; CNN; GNN; cell morphology; ResNet18; medical image classification; OI-PBC

Article History

Received: 25 June 2026

Accepted: 16 September 2026

Published: 10 October 2026

Copyright © Authors

Corresponding Author: *
Sadia Bhutto

Abstract

Identifying blood cell types correctly under a microscope is a task doctors rely on constantly, since red cells, white cells, and platelets each carry information about a patient's health, and several conditions leave visible traces in how these cells look. Doing this by hand takes time, varies from one observer to another, and needs specialists who are not always available, especially in smaller clinics. In this paper we build a Hybrid CNN-GNN model that tries to combine the best of two worlds: a pretrained ResNet18 branch that picks up on color and texture patterns, and a graph neural network branch that represents each cell as three connected parts, the cell itself, the nucleus, and the cytoplasm, so the model can also reason about how these parts relate to one another. We train and test the approach on the OI-PBC dataset, which has close to 300,000 labeled peripheral blood cell images spread across twelve cell types, and we balance the training data to 3,000 images per class so that common cell types do not drown out the rarer ones. On a test set that is also balanced this way, our hybrid model reaches 91.15% accuracy and a macro F1 of 91.17%, ahead of a GNN-only version (48.42%) and a CNN-only version (90.38%). On the original, unbalanced test set it hits 99.35% accuracy, a Kappa of 0.966, and an AUC of 0.990. The task here is twelve-class cell-type classification, not disease diagnosis, since our labels describe cell morphology rather than confirmed patient outcomes. We also wrapped the pipeline in a small Gradio research prototype that returns a predicted cell type, a confidence score, and a downloadable PDF summary, and we deliberately keep this framed as an exploratory screening aid rather than a validated diagnostic tool, since that would require patient-level labels and clinical validation we do not yet have.

INTRODUCTION

Looking at blood under a microscope is still one of the cheapest and most useful ways doctors have to catch early signs of disease. Anemia,

leukemia, and several other conditions all leave visible traces in how red cells, white cells, and platelets look, so a hematologist trained to spot these patterns can often flag a problem long

before more expensive tests are ordered. The catch is that someone has to sit down and do this manually, slide after slide, and that person gets tired, sees things a little differently on different days, and is not always around when needed, particularly outside of major hospitals.

Over the last few years, convolutional neural networks have become a common tool for teaching a computer to recognize cell types from images (Zhu, Wang, and Zhang 2023; Abou Ali, Dornaika, and Arganda-Carreras 2023). They are good at this, but they mostly just look at pixels, they do not really know that a cell has a nucleus sitting inside a cytoplasm and that the relationship between those two parts matters. Hematologists lean on exactly that relationship when they are trying to tell apart cell types that look almost identical at first glance, immature myeloid cells being a good example. Graph neural networks are one way to give a model this kind of structural awareness, since they let you describe an image as a set of connected parts rather than a flat grid (Zhang et al. 2023; Jiang et al. 2024). The problem is that a GNN is only as good as the features you hand it, and most work in this space still relies on features that were designed by hand rather than learned.

That is the gap we are trying to close here. We keep the CNN around, but instead of only using it to classify the whole image, we also use it to generate the features that feed into a small graph built from the cell, the nucleus, and the cytoplasm. Section 2 lays out the three main contributions of this paper, and the rest of the paper is organized around them: Related Work compares each contribution to the closest prior work, Methodology explains how each one was built, and Evaluation reports how each one was validated.

Main Contributions

This work makes three main contributions, and we use the same three-way split in every section that follows.

2.1 An Intra-Cell Morphology Graph

We build a lightweight three-node graph for every blood cell, one node for the whole cell,

one for the nucleus, and one for the cytoplasm, connected by edges that carry position and size information. Unlike graphs built across multiple cells in an image, this graph captures structure inside a single cell, which is the level at which hematologists actually reason when telling apart visually similar cell types.

A Hybrid CNN-GNN Model with a Fair Head-to-Head Comparison

We feed CNN-derived visual features into the intra-cell graph instead of handcrafted measurements, and we train a GNN-only model, a CNN-only model, and this hybrid model under identical data splits, class balancing, and training schedules, so the comparison actually isolates what each component contributes.

A Research Prototype for Interactive Screening

We package the trained hybrid model into a Gradio-based research prototype that takes a blood cell image and returns a predicted class, a confidence score, and an automatically generated PDF report, so the underlying classifier can be tried interactively rather than only evaluated offline. We treat this as a demonstration of how the model's output could eventually be surfaced to a clinician, not as a validated diagnostic instrument.

Related Work

This section is organized around the same three contributions above, so each subsection compares our work to the literature closest to that specific contribution.

Structural and Graph-Based Representations for Cell Images

Graph-based representations of cell images are not new on their own. Altarawneh et al. (Altarawneh et al. 2023) build a hybrid CNN-GNN for leukocyte detection where the graph models relationships between different leukocytes in a blood sample, with each whole leukocyte image treated as a graph node and edges

capturing inter-cell context. Zhang et al. (Zhang et al. 2023) review a broader set of graph-based methods that similarly model relationships between anatomical regions or between cells for disease diagnosis. What is common across this line of work is that the graph operates *between* cells or regions of a whole slide. Our contribution in Section 2.1 instead builds the graph *inside* a single cell, treating the cell, nucleus, and cytoplasm as the nodes, which is a different and, to our knowledge, less explored level of structure for this task.

Hybrid CNN-GNN and CNN-Transformer Models for Blood Cell Classification

Several recent systems combine CNN features with a second modeling stage to capture relationships beyond raw pixels. Alshehri et al. (Alshehri et al. 2025) pair pretrained CNN backbones with Vision Transformer layers to combine local and global feature modeling for normal blood cell classification. Jiang et al. (Jiang et al. 2024) and Duc, Bilik, and Martinek (Duc, Bilik, and Martinek 2025) propose GNN-based white blood cell recognizers built on super pixel graphs and multi-hop attention, respectively. Zhu, Wang, and Zhang (Zhu, Wang, and Zhang 2023) and a follow-up study (Zhu et al. 2023) instead use ensembles of CNN and randomized-network classifiers without any graph component, while Abou Ali, Dornaika, and Arganda-Carreras (Abou Ali, Dornaika, and Arganda-Carreras 2023) and Asghar et al. (Asghar et al. 2023) compare a range of CNN and transformer backbones on blood cell datasets. None of these studies runs a controlled, side-by-side comparison of a GNN-only model, a CNN-only model, and a CNN-fed hybrid under matched training conditions on the same class-balanced data, which is exactly what our second contribution (Section 2.2) sets out to do.

Deployable Tools for Blood Cell Screening

On the applied side, Sazak and Kotan (Sazak and Kotan 2024) apply YOLOv10 and YOLOv11 to detect and classify RBCs, WBCs, and platelets as broad categories

directly in microscopic images, and Mondal et al. (Mondal et al. 2024) present Blood Cell-Net, a lightweight CNN pipeline with a segmentation stage ahead of classification for nine blood cell types. Both are aimed at practical deployment, but neither couples a fine-grained twelve-class classifier with a structural reasoning branch or produces a clinician-facing report. Our third contribution (Section 2.3) builds on the classification model from Section 2.2 and wraps it in an interactive tool that a user can actually run, complete with a downloadable report.

Methodology

This section is again organized by contribution. It is primarily an image processing pipeline: raw microscope images are segmented, converted into graph and feature representations, and passed through the models described below.

Building the Intra-Cell Morphology Graph

We work with the OI-PBC dataset (Abouali 2023), which contains 298,850 labeled peripheral blood cell images across twelve classes: Basophil (BA), Band Neutrophil (BNE), Eosinophil (EO), Erythroblast (ERB), Lymphocyte (LY), Metamyelocyte (MMY), Monocyte (MO), Myelocyte (MY), Platelet (PLT), Promyelocyte (PMY), Red Blood Cell (RBC), and Segmented Neutrophil (SNE). Annotations come in YOLO format, with each image paired with a text file giving the class and bounding box of every cell present, and individual cell images are cropped from these boxes with a small 2-pixel padding on each side; crops smaller than 20 pixels in either dimension are discarded as too small to carry meaningful morphological information. RBC images alone make up around 90% of the raw dataset, which would badly skew training if left as is, so we capped the training set at 3,000 images per class, 36,000 images in total, using under sampling for classes with more than 3,000 available images and

oversampling for classes with fewer. We report results on both the original imbalanced test set and a class-balanced version we call the fair test set. Table 1 shows the class-wise composition of the original test

set, which preserves the dataset's natural distribution and is what the balanced-test numbers in Section 5.2 are compared against. Table 1: Blood cell classes and their distribution in the original test set.

Code	Cell Type	Test Samples	Category
BA	Basophil	94	WBC
BNE	Band Neutrophil	170	WBC
EO	Eosinophil	321	WBC
ERB	Erythroblast	169	Precursor
LY	Lymphocyte	125	WBC
MMY	Metamyelocyte	107	Precursor
MO	Monocyte	148	WBC
MY	Myelocyte	120	Precursor
PLT	Platelet	1371	Platelet
PMY	Promyelocyte	65	Precursor
RBC	Red Blood Cell	25483	RBC
SNE	Segmented Neutrophil	177	WBC
Total		28350	

Figure 1: Class Distribution — Original vs Balanced Training Set




Figure 1: Class distribution of the training data before and after balancing to 3,000 samples per class.

For each cell image we segment out three regions: the whole cell, the nucleus, and the cytoplasm. The whole-cell boundary is separated from the background using Otsu's thresholding

on the grayscale image, followed by morphological opening and closing to smooth the boundary. The nucleus is then segmented from within that boundary by clustering pixels into two

groups with k-means in the LAB color space; the cluster with the lower mean L-channel intensity, corresponding to the darker, more heavily stained pixels, is taken as the nucleus, which is appropriate since the nucleus stains darker than the cytoplasm under standard blood smear stains. Morphological cleanup is applied to remove small artifacts from the nucleus mask, and the cytoplasm region is what is left after subtracting the nucleus mask from the cell mask. When a nucleus cannot be reliably segmented this way, we fall back to a single whole-cell node rather than forcing a bad split. Each region becomes a node in a small, fully connected graph, and the edges carry information about how far apart the regions are and how their sizes compare. We kept this graph deliberately simple; three nodes are enough to capture the cell-nucleus-cytoplasm relationship without slowing down inference.

Designing the Hybrid CNN-GNN Model

We trained three models under identical conditions so the comparison would actually mean something. **GNN baseline.** Each node gets twelve handcrafted features: normalized area, centroid position, mean and standard deviation of the blue, green, and red channels, normalized perimeter, and circularity. These are passed through three Graph Convolutional Network (GCN) layers, each followed by Batch Normalization and a ReLU activation, with a hidden dimension of 128 throughout. A residual connection links the input and output of the second GCN layer so the model can preserve information from the original node features alongside what it aggregates from neighbors. After the

three GCN layers, global mean pooling collapses all node representations into a single graph-level vector, which is passed through a two-layer fully connected classification head, with a ReLU activation between the two layers, to produce the final twelve-way prediction.

CNN baseline. A pretrained ResNet18 backbone, fine-tuned end to end with a new 12-way classification layer on top, replacing the original 1000-way ImageNet head. This model only ever sees the full, cropped cell image, resized to 224×224 .

Hybrid CNN-GNN (ours). Here we reuse the same pre-trained ResNet18, but instead of classifying the whole image directly, we run it separately over each of the three region crops from Section 4.1 and pull out a 512-dimensional feature vector from the global average pooling layer for each one. These CNN-derived vectors replace the twelve handcrafted features from the GNN baseline as node inputs to the same three-layer GCN architecture described above. The graph-level output of the GNN branch is concatenated with the CNN branch's own 512-dimensional whole-image feature vector, and this combined vector is passed through a two-layer fully connected classification head to produce the final prediction. Fusing by concatenation is a simple choice, but it lets the classification head learn how much weight to give the visual and structural signals for each cell type, rather than fixing that balance by hand

Figure 2: Hybrid CNN-GNN Model Architecture





Figure 2: Hybrid CNN-GNN architecture, combining a ResNet18 feature extractor with a graph convolutional net- work for blood cell classification.

All three models were trained under identical conditions on the balanced training set, using a batch size of 64 for training and 128 for validation and testing, with a Weighted Random Sampler so that individual batches stayed approximately class-balanced even within an epoch. The loss function was Focal Loss with $\gamma = 2.0$, combined with per-class weights set by inverse class frequency in the training set, chosen specifically to keep the rare classes contributing meaningfully to the gradient. All three models were trained with the Adam W optimizer (learning rate $3e-4$, weight decay $2e-4$), a cosine-

annealed learning rate schedule down to a minimum of $1e-6$ over 100 epochs, gradient clipping at a max norm of 1.0, and early stopping with a patience of 20 epochs tracked on validation macro F1, on Google Co- lab with GPU acceleration. We tracked six evaluation metrics: overall accuracy, balanced accuracy, macro F1, Cohen's Kappa, Matthews Correlation Coefficient (MCC), and ROC AUC, with balanced accuracy and macro F1 weighted most heavily since neither is inflated just because a model is good at recognizing the majority class.



The prototype wraps the trained hybrid model from Section 4.2 in a Gradio interface. A user uploads a blood cell image, which is passed through the same segmentation and graph construction steps from Section 4.1 and then through the hybrid model. The predicted class and confidence score are rendered into a downloadable PDF report so the classification output is easy to share and review. We do not attach a disease-level interpretation to this output

in the current version, since doing so responsibly would need patient-level disease labels and clinical input we have not yet collected; the prototype currently reports cell-type predictions only. Figure 3 shows how this pipeline connects to the two branches described above.




Figure 3: Overall system architecture of the proposed Hybrid CNN-GNN framework, from input image through the CNN and graph branches to final classification and disease-level reporting.

Evaluation

This section follows the same three-way structure used in Related Work and Methodology, so that each contribution is validated on its own terms before we discuss the results as a whole.

Validating the Intra-Cell Morphology Graph

To see whether the intra-cell graph is actually contributing anything, we look at where the Hybrid CNN-GNN model gains the most over the CNN-only baseline: the myeloid precursor classes MY, MMY, and PMY, which are consecutive, visually similar stages of granulocyte maturation and are exactly the case where the nucleus-cytoplasm relationship should matter most. The Hybrid CNN-GNN confusion matrix in Figure 4 shows a clean diagonal overall, with what confusion remains concentrated almost entirely among these three classes, meaning the graph branch is doing real work precisely where cell-internal structure is diagnostically relevant, even though it does not fully resolve visual

similarity between successive maturation stages.

Figure 4: Normalized confusion matrix for the Hybrid CNN- GNN model on the balanced fair test set. Residual confusion is concentrated among the myeloid precursor classes, the same classes where the intra-cell graph is expected to help most.

Validating the Hybrid CNN-GNN Model

Table 2 reports performance on the original, imbalanced test set, and Table 3 reports performance on the balanced fair test set. In both cases the Hybrid CNN-GNN model outperforms both baselines on every metric. The gap is widest against the GNN-only baseline: balanced accuracy on the fair test set is just 48.42% for the GNN baseline, even though it looked reasonable on the original, RBC-heavy test set. That gap shows that twelve handcrafted numbers per region are not enough to separate twelve cell types once the majority-class advantage is removed, and that CNN-derived features are what fix most of the problem.

Table 2: Overall performance on the original (imbalanced) test set.

Model	Acc.	Bal. Acc.	Macro F1	Kappa	AUC
GNN baseline	92.07%	50.51%	44.05%	0.646	0.901
CNN baseline	99.34%	90.64%	90.45%	0.965	0.989
Hybrid CNN-GNN	99.35%	90.92%	90.96%	0.966	0.990

Table 4 gives the full per-class precision, recall, and F1 on the balanced fair test set for both the GNN baseline and the Hybrid CNN-GNN model, side by side. Every one of the twelve classes gets an F1 score above 0.85 from the Hybrid CNN-GNN model, with RBC (0.96), PLT (0.95), and LY (0.94) strongest, and MY (0.85) and MMY (0.86) weak- est but still solid. The GNN baseline, by

contrast, falls apart on exactly the classes that matter most clinically: EO gets an F1 of 0.00, meaning the model essentially never identifies an eosinophil correctly, and PMY gets an F1 of just 0.04 despite a recall of 0.71, meaning it flags promyelocytes constantly but almost always flags the wrong cell as one. Figure 5 shows this same comparison as a bar chart across all three models.

Table 3: Overall performance on the class-balanced fair test set.

Model	Acc.	Bal. Acc.	Macro F1	M C C	AUC
GNN baseline	48.42%	48.42%	47.68%	0.436	0.882
CNN baseline	90.38%	90.38%	90.22%	0.904	0.988
Hybrid CNN-GNN	91.15%		91.15%	0.911	0.990

Table 4: Per-class precision, recall, and F1 on the balanced fair test set, GNN baseline vs. Hybrid CNN-GNN.

	GNN baseline			Hybrid CNN-GNN		
Class	Prec.	Rec.	F1	Prec.	Rec.	F1
BA	0.41	0.40	0.31	0.89	0.88	0.88
BNE	0.49	0.41	0.45	0.88	0.90	0.89
EO	0.00	0.00	0.00	0.86	0.88	0.87
ERB	0.53	0.37	0.44	0.91	0.90	0.90
LY	0.05	0.76	0.09	0.95	0.94	0.94
MMY	0.23	0.36	0.28	0.87	0.86	0.86
MO	0.21	0.42	0.28	0.88	0.89	0.88
MY	0.26	0.25	0.25	0.85	0.84	0.85
PLT	0.63	0.63	0.63	0.96	0.95	0.95
PMY	0.02	0.71	0.04	0.85	0.87	0.86
RBC	1.00	0.83	0.91	0.97	0.96	0.96
SNE	0.33	0.64	0.44	0.92	0.93	0.92
Macro avg.	0.35	0.48	0.31	0.90	0.90	0.90




Figure 5: Precision, recall, and F1-score for all twelve classes, compared across the GNN baseline, CNN baseline, and Hybrid CNN-GNN model.

Validating the Screening Tool

For the prototype contribution, we evaluated it qualitatively rather than with a benchmark score, since there is no established metric for report usability at this stage. Running the pipeline described in Section 4.3 end to end, from image

upload to classification, confidence score, and PDF report, took a few seconds per image on standard hardware, confirming the prototype is fast enough for interactive, single-image use. We are not aware of a comparable blood cell tool in the reviewed literature. that pairs a fine-grained twelve-class hybrid CNN-GNN classifier with an interactive, report-generating interface, which is the gap our third contribution targets. We see this as a starting point for later clinician-facing work rather than a finished product.

Discussion and Limitations

If we had to point to one takeaway from these numbers, it is that the CNN branch is doing most of the heavy lifting. Balanced accuracy jumps from 48.42% to 90.38% just by switching from the GNN baseline to the CNN baseline, a much bigger move than the further step up to 91.15% once the graph reasoning is added on

top. So, for a task like this, at this scale, learned visual features are clearly the dominant signal, and structural reasoning adds a real but smaller benefit on top, mostly in exactly the cases where it should matter, the myeloid precursor classes where the nucleus-cytoplasm relationship carries diagnostic weight. That said, the gain from adding the graph branch is modest relative to the CNN baseline, and we have not yet tested whether it is statistically reliable across repeated training runs, so we treat it here as a promising but unconfirmed trend rather than a settled result. There are several things this work does not do, and we want to be upfront about them rather than let the framing overstate what has been shown. First, all of our results come from a single dataset, OI-PBC, so we do not yet know how the model behaves on images from a different lab, scanner, or staining protocol, and domain shift of that kind is a well-known problem for microscopy models. Second, our labels describe cell morphology type, not confirmed patient-level disease status, so nothing here has been validated against actual diagnoses, and we do not treat the myeloid-precursor pattern the model is more accurate on as evidence of diagnostic capability. Third, the

model classifies single, already- cropped cells rather than scanning and counting cells across a whole slide, which is a different and harder problem that a real screening pipeline would eventually need to solve. Fourth, our segmentation step relies on color-based thresholding tuned to this dataset's staining, and we have not measured how often it fails or how failures affect downstream classification, so segmentation error is a real but currently unquantified source of noise in our pipeline. Finally, none of this has been reviewed by a practicing hematologist or tested in a clinical setting, and every output from the current prototype should be treated as exploratory rather than something a clinician could rely on.

Conclusion and Future Work

We set out to build a Hybrid CNN-GNN model for fine- grained peripheral blood cell classification around three contributions: an intra-cell morphology graph, a fairly compared hybrid model, and an interactive research prototype. Across both the original and class-balanced test sets, the hybrid model consistently beat both a GNN-only and a CNN-only baseline, landing at 91.15% accuracy and a 91.17% macro F1 on the harder, balanced test set, with the largest single- baseline jump coming from adding CNN features at all, and a further, smaller gain from the graph branch on top. We see this as evidence that morphology-aware graph representation is a promising direction for this task, though not yet a fully confirmed one.

Several next steps are needed before the claims here can be strengthened further. On the modeling side, we plan to run a proper ablation across CNN-only, GNN-only, CNN features without graph aggregation, and the full hybrid model, and to repeat training across multiple random seeds so we can report the hybrid model's gain with a confidence interval rather than a single run. We also want to com- pare against at least

one stronger contemporary CNN back- bone beyond ResNet18. On the structural side, swapping in attention-based graph layers such as Graph Attention Net- works instead of the current GCN, and running node and edge ablations, would help isolate exactly what the graph is contributing. On the applied side, we want to quantify seg- mentation failure rate, test robustness to realistic microscopy variation such as brightness, contrast, and blur, and, if an independent blood-cell dataset becomes available, evaluate external validity outside OI-PBC. Longer term, extending the pipeline to whole-slide images and automatic cell counting, aggregating cell-level predictions toward a patient level read, adding explainability tools such as class activation maps, and running the prototype past practicing hematologists for clinical feedback are the steps that would need to happen before any disease-screening framing would be justified.

Acknowledgments

We are grateful to our co-supervisor, Arif Raza khuhawar, for their guidance throughout this project, and to the Department of Computer Systems Engineering, QUEST Nawab- shah, for its support.

REFERENCES

- Abou Ali, M.; Dornaika, F.; and Arganda-Carreras, I. 2023. White Blood Cell Classification: Convolutional Neural Net- work (CNN) and Vision Transformer (ViT) under Medical Microscope. *Algorithms*, 16(11): 525.
- Abouali, M. 2023. OI-PBC Dataset: 298,850 Blood Cells and 12 Classes. Kaggle. <https://www.kaggle.com/datasets/mohamadabouali1/oi-pbc-dataset-298850-blood-cells-and-12-classes>.

- Alshehri, O. M.; Shaf, A.; Irfan, M.; Jalal, M. M.; Alta- yar, M. A.; Abu-Alghayth, M. H.; Al Shmrany, H.; Ali, T.; Soomro, T. A.; and Alkhathami, A. G. 2025. A Hy- brid CNN-Transformer Framework for Normal Blood Cell Classification: Towards Automated Hematological Analy- sis. *CMES-Computer Modeling in Engineering & Sciences*, 144(1): 1165–1196.
- Altarawneh, K.; et al. 2023. Hybrid Convolution Neural Net- work and Graph Neural Network for Detection of Leuko- cyte. *International Journal on Recent and Innovation Trends in Computing and Communication*, 11(9).
- Asghar, R.; Kumar, S.; Mahfooz, A.; et al. 2023. Classification of All Blood Cell Images using ML and DL Models. *arXiv preprint arXiv:2308.06300*.
- Duc, M. L.; Bilik, P.; and Martinek, R. 2025. White Blood Cell Classification using Multi-hop Attention Graph Neural Networks. *Expert Systems with Applications*, 272: 126725.
- Jiang, Y.; Shen, Y.; Wang, Y.; and Ding, Q. 2024. Automatic Recognition of White Blood Cell Images with Memory Efficient Superpixel Metric GNN: SMGNN. *Mathematical Biosciences and Engineering*, 21(2): 2163–2188.
- Mondal, S. K.; Md. Talukder, M. S. H.; Aljaidi, M.; Sulaiman, R. B.; Md Tushar, M. M. S.; and Alsuwaylimi, A. A. 2024. BloodCell-Net: A Lightweight Convolutional Neural Network for the Classification of All Microscopic Blood Cell Images of the Human Body. *arXiv preprint arXiv:2405.14875*.
- Sazak, H.; and Kotan, M. 2024. Automated Blood Cell Detection and Classification in Microscopic Images Using YOLOv11 and Optimized Weights. *Diagnostics*, 15(1): 22.
- Zhang, L.; Zhao, Y.; Che, T.; Li, S.; and Wang, X. 2023. Graph Neural Networks for Image-Guided Disease Diagnosis: A Review. *iRADIOLOGY*, 1(2): 151–166.
- Zhu, Z.; Ren, Z.; Lu, S.; Wang, S.; and Zhang, Y. 2023. DLBCNet: A Deep Learning Network for Classifying Blood Cells. *Big Data and Cognitive Computing*, 7(2): 75.
- Zhu, Z.; Wang, S.-H.; and Zhang, Y.-D. 2023. ReRNet: A Deep Learning Network for Classifying Blood Cells. *Technology in Cancer Research & Treatment*, 22: 1–14.