ISSN: 2096-3246
Volume 58, Issue 01, 2026

A ROBUST DEEP LEARNING–BASED COMPUTER-AIDED DIAGNOSTIC FRAMEWORK FOR BLOOD CANCER DETECTION USING ADAPTIVE EFFICIENTNET OPTIMIZATION

P. Geetha[1], Dr. K. Haridas[2]
[1]PhD. Research Scholar, Department of Computer Science, Nallamuthu Gounder Mahalingam College, Pollachi. geethapchandrasekaran@gmail.com
[2]Head & Associate Professor, Department of Computer Applications, Nallamuthu Gounder Mahalingam College, Pollachi.

Abstract
Leukemia is a malignant disorder of the blood-forming tissues characterized by uncontrolled proliferation of abnormal white blood cells, leading to impaired immune function and disruption of normal hematopoiesis. It represents a significant global health burden and accounts for a substantial proportion of pediatric cancers. Acute lymphoblastic leukemia (ALL), the most prevalent subtype, originates in the bone marrow and progresses rapidly, making early diagnosis essential for improving survival outcomes. To facilitate timely and accurate detection, this study presents an automated computer-aided diagnosis framework based on the EfficientNet-B3 convolutional neural network, integrated with a dynamically adjusted learning rate strategy. The proposed approach adaptively modifies the learning rate at each training epoch by jointly analyzing training accuracy and loss values to enhance convergence and classification performance. The model was evaluated on the research Leukemia dataset following normalization and class balancing, achieving average precision, recall, specificity, accuracy, and Dice similarity coefficient values of 98.29%, 97.83%, 97.82%, 98.31%, and 98.05%, respectively. The proposed approach employs a dynamic learning rate adjustment mechanism that continuously monitors validation loss and training accuracy to enhance convergence stability and classification reliability. Experiments were conducted on the public Acute Lymphoblastic Leukaemia dataset to evaluate cross-disease generalization. These findings confirm the effectiveness of the proposed method for automated ALL detection.
Keywords: Acute lymphoblastic leukaemia, Blood cancer diagnosis, EfficientNet, Dynamic learning rate, Deep learning, Medical image analysis

1. Introduction

Leukaemia is medically classified into four types, namely, acute myelogenous leukaemia (AML), acute lymphoblastic leukaemia (ALL), chronic myeloid leukaemia (CML), and chroniclymphocytic leukaemia (CLL) [1]. Among these, acute lymphoblastic leukaemia (ALL) is an aggressive form of blood cancer that primarily affects white blood cells. In acute forms of leukaemia, the abnormal cells are increase and spread quickly [2-4], it requiring early medical attention, whereas chronic leukaemias usually progress more slowly and are often difficult to identify in their early stages [5,6]. The excessive production of malignant cells disrupts normal blood functioning, weakens immune function, and reduces the body's ability to maintain healthy levels of red blood cells and platelets, thereby placing multiple organs at risk and disrupts the normal function of the body [7].
In individuals with acute lymphoblastic leukemia (ALL), have large numbers of abnormal WBCs

circulating in their bloodstream and gradually infiltrating vital organs such as the liver, spleen, brain, and kidneys. This widespread impact can cause serious clinical complications and in advanced stages, life-threatening conditions. Due to its aggressive nature, delayed diagnosis or lack of timely treatment can lead to death. Research on epidemiology from India shows that ALL is becoming more common, affecting both children and adults, and causing a lot of deaths every year [8]. Although ALL is highly treatable when found early, diagnosis is often delayed because its symptoms, such as joint and bone pain, anaemia, fatigue, weakness, and fever are very similar to those of many non-cancerous diseases, which can lead to delays clinical recognition.

Clinical suspicion of leukemia is usually based on symptom presentation and is confirmed through a structured diagnostic process. Initial investigations usually involve complete blood count (CBC) tests and peripheral blood smear analysis, which assess abnormalities in leukocyte numeral, shape, and distribution, as well as variations in other blood components. Diagnostic accuracy is further enhanced through advanced molecular and genetic techniques such as cytogenetic analysis, fluorescence in situ hybridization (FISH), and polymerase chain reaction (PCR), which helps to identify the chromosomal and genetic abnormalities associated with leukemic cells [4].

Hematological confirmation commonly involves invasive diagnostic procedures, including blood biopsies and cerebrospinal fluid examinations [9]. While clinically effective, these procedures are frequently painful, time-consuming, costly, and psychologically distressing to patients. Moreover, their diagnostic reliability is strongly influenced by the physicians' experience and subjective judgment, which can lead to variations interpretation and increase the human error [10]. These challenges have led to increasing the interest in non-invasive diagnostic methods based on medical image analysis, which provide faster, safer, and more economical alternatives while reducing dependence on invasive testing.

Although image-based diagnostic methods are widely used in clinical practice, they still have certain limitations. The manual examination of large-scale medical image datasets, especially peripheral blood smear images, is labour-intensive, time-consuming, requires significant effort and can easily lead to misinterpretation. The presence of overlapping textures, subtle morphological differences, and visual similarities between healthy and cancerous blood cells further complicate accurate diagnosis.

As shown in Figure 1, distinguishing normal lymphocytes from ALL blast cells in peripheral blood remains a complex and may lead to misclassification [11].

Leukocytes are characterized by high intraclass similarity and limited interclass separability, making reliable classification extremely challenging using traditional diagnostic approaches. This creates a critical need for reliable, automated diagnostic systems capable of identifying blood cancer at an early stage. Although leukemic cells differ from healthy blood cells in terms of morphology, size, shape, and surface properties [12], these differences are sometimes and difficult to detect through manual inspection. In this case, completely automated deep learning–based frameworks offer a powerful solution by learning discriminative features directly from peripheral blood images, enabling accurate, consistent, and scalable detection and classification of blood cancer.

ISSN: 2096-3246
Volume 58, Issue 01, 2026

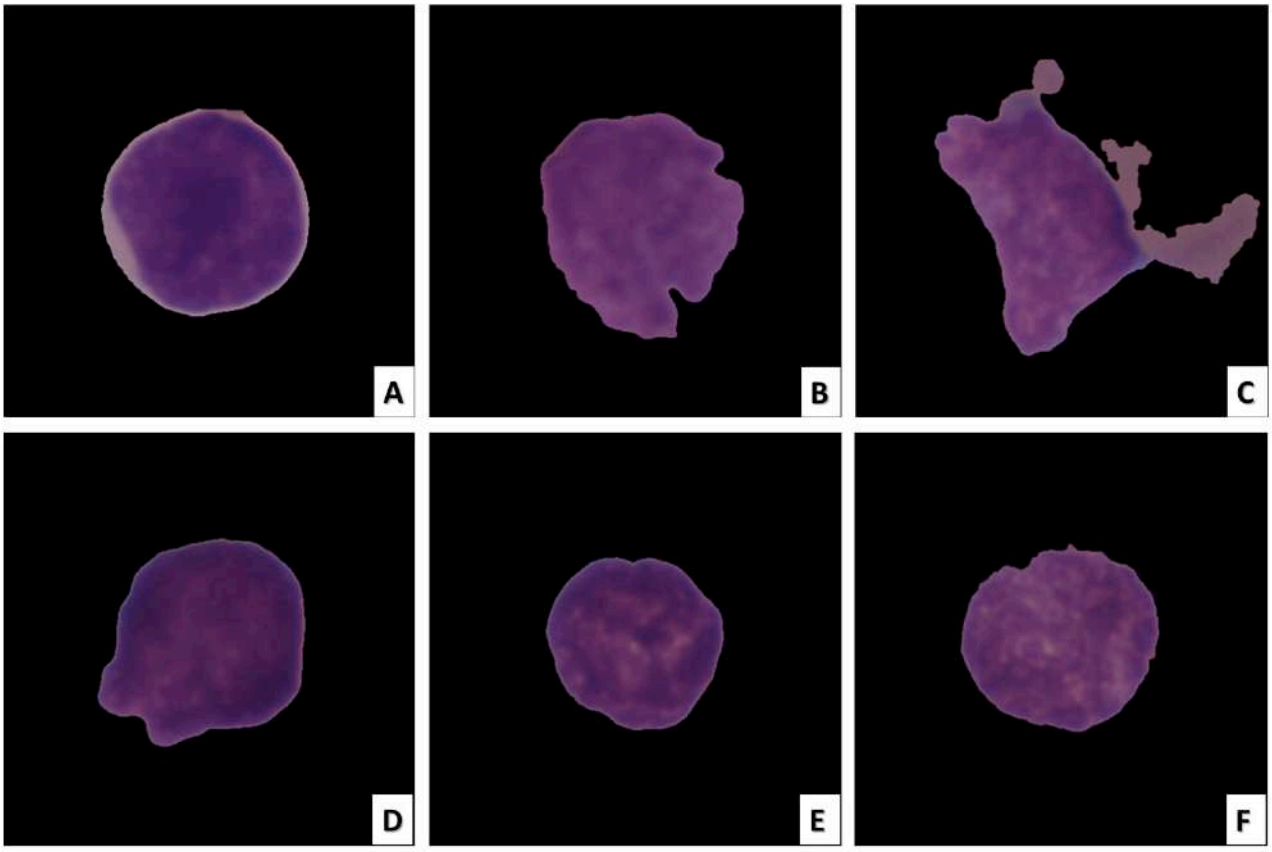

Figure 1. ALL cell image samples. The images in (**A**–**C**) depict ALL cancer cells, while the images in (**D**–**F**) depict healthy cells.

In recent years, deep learning (DL), mainly deep neural networks (DNNs), has emerged as a influential and transformative technology in fields such as speech recognition and computer vision [13]. Deep learning (DL) is a subset of artificial intelligence (AI) which concentrates on learning hierarchical representations of data, enabling models to automatically extract meaningful features from raw inputs. Depending on the learning paradigm, DL methods can be classified as supervised, semi-supervised, and unsupervised learning approaches. These approaches differ in how labeled and unlabeled data are utilized during training [14]. A unique feature of DL is that it replaces the manual feature engineering through hierarchical feature extraction, allowing models to learn complex patterns directly from data.

Deep learning architectures often have multiple hidden layers positioned between the input and output layers, each responsible for extracting progressively abstract feature representations. One of the main advantages of DL models is their ability to learn from both labeled and unlabeled datasets, which makes them highly flexible and scalable. Among various DL architectures, convolutional neural networks (CNNs) have emerged as the most widely adopted models for medical image analysis due to their strong capability to capture spatial and structural information [15]. CNNs using the backpropagation technique, can automatically learn discriminative spatial features from clinical images, which is particularly advantageous for disease detection tasks.

For effective acute lymphoblastic leukemia (ALL) detection, DL models must be trained using adequately huge and diverse image datasets to capture relevant spatial-domain features. Deep learning based approaches generally achieve optimal performance when a large number of training samples are available [16]. The availability and quality of training data play a important role in determining both model architecture and hyperparameter selection. During training, the model gradually updates its weights based on training data to reduce classification error across predefined classes.

The performance of artificial neural networks (ANNs), including CNNs, is extremely dependent on the proper tuning of hyperparameters such as the learning rate (LR), batch size (BS), and momentum. Among these, the learning rate is one of the most important parameters influencing training stability and model accuracy. Conventional LR strategies such as constant LR, step decay, and exponential decay frequently use trial-and-error methods to identify an appropriate value for a given application. Fixed learning rate schedules are commonly used as basic approaches; however, they suffer from inherent limitations. When the learning rate is set too low, model convergence becomes slow, however excessively high learning rates can cause training instability and divergence, leading to suboptimal solutions. Optimal learning rates enable faster convergence by suitable scaling the gradient updates toward global minima.

The learning rate determines the magnitude of weight updates during backpropagation and plays a key role in navigating the loss landscape. When optimization process becomes stuck in local minima or saddle points, fixed learning rate strategies may fail to make meaningful progress. Adaptive learning rate methods address this limitation by dynamically adjusting the learning rate during training to improve convergence behaviour. In contrast, non-adaptive schedules either reduce the learning rate gradually at each epoch or maintain a constant value throughout training.

Several dynamic learning rate approaches have been proposed in recent years, including cyclical learning rate (CLR) [17], stochastic gradient descent with warm restarts (SGDWR), also referred to as stochastic weight averaging (SWA) [18], and cosine annealing [19]. In CLR, the learning rate oscillates between predefined minimum and maximum bounds throughout training. The learning rate begins at a low value, progressively increases to a peak, and then decreases back to its initial value, completing one cycle. Each cycle consists of increasing phase and a decreasing phase with a fixed step size that determines the duration of the cycle [20]. This cyclic pattern is repeated until the final training era. Although temporarily increasing the learning rate may reduce accuracy in the short term, it frequently leads to lower training loss and improved generalization in the long run [17].

Stochastic gradient descent (SGD) remains one of the most commonly used optimization algorithms for training deep neural networks [18]. In SGD-based training, model parameters are updated after each epoch based on gradient information. When the learning rate is too low, convergence may slow significantly, mainly near saddle point plateaus in nonconvex optimization landscapes. Increasing the learning rate in such cases can help the optimizer escape saddle points and accelerate convergence, thereby improving overall training efficiency.

2. Related Work

Several deep learning architectures, including InceptionResNetV2, VGG19, and ResNet101V2, have

been evaluated in previous studies, and their performance has been shown to generally align with explainable artificial intelligence (XAI) frameworks, particularly through the application of the LIME algorithm, enabling transparent interpretation of model decisions.

Mondal et al. [4] used convolutional neural networks (CNNs) to recognize the acute lymphoblastic leukemia (ALL) and they proposed a weighted ensemble strategy to improve classification performance. Their approach integrated extensive data preprocessing and augmentation techniques to enhance model robustness and generalization. Experiments conducted on the Leukemia dataset demonstrated strong performance, achieving an AUC of 0.948, balanced accuracy of 88.3%, and a weighted F1-score of 89.7%. However, individual ensemble components such as InceptionResNet-V2, DenseNet-121, Xception, MobileNet, and VGG-16 were discovered to learn fragmented and dispersed feature regions, that were less consistent than those captured by the combined ensemble model.

Amin et al. [12] presented a non-invasive CNN-based diagnostic framework for ALL detection using microscopic blood smear images. Their model incorporated an Efficient Channel Attention (ECA) module integrated with VGG16, enabling improved feature extraction by addressing morphological similarities between healthy and leukemic affected cells. The use of diverse data augmentation strategies further enhanced dataset quality and diversity. The proposed framework achieved a classification accuracy of 91.1%, demonstrating strong clinical relevance and potential utility in supporting pathologists during diagnosis.

Khandekar et al. [21] proposed an AI-driven automated blast cell detection system based on object detection techniques. Using YOLOv4, their framework performed binary classification of cells as either healthy (HEM) or leukemic blast cells (ALL) from microscopic blood smear images. The model was trained and evaluated using the ALL-IDB1 and C-NMC datasets, achieving a mean average precision (mAP) of 96.06% and 98.7%, respectively. This approach demonstrated strong potential as a pre-screening tool for early leukemia detection.

Almadhor et al. [22] proposed an ensemble machine learning framework using traditional classifiers, including Naïve Bayes (NB), K-Nearest Neighbors (KNN), Random Forest (RF), and Support Vector Machine (SVM), applied to the C-NMC leukemia dataset. Their findings indicated that SVM achieved the highest performance, reaching an accuracy of 90%, outperforming the other classifiers in the ensemble.

Kasani et al. [23] developed an aggregated deep learning architecture for classifying leukemic B-lymphoblasts by employing data augmentation and transfer learning, the model addressed the challenges of limited dataset size while accelerating convergence and improving performance. The proposed system achieved a test accuracy of 96.58%, outperforming individual deep learning models through feature-level fusion and model aggregation.

Liu et al. [24] introduced a ternary-stream fine-grained classification model for differentiating lymphoblasts, normal white blood cells, and reactive lymphocytes using microscopic peripheral blood smear images. Evaluated on the C-NMC dataset, the model achieved an accuracy of approximately 91.90%, demonstrating strong capability in discriminating morphologically similar cell types.

From the critical analysis of recent studies, it is evident that no existing method has yet achieved optimal performance in the simultaneous classification of ALL and detection of malarial parasites. In contrast, the proposed model in this study demonstrates superior performance. On the C-NMC Leukemia dataset, it achieved an average precision, recall, specificity, accuracy, and Dice Similarity Coefficient (DSC) of 98.29%, 97.83%, 97.82%, 98.31%, and 98.05%, respectively. Similarly, evaluation on the NIH dataset yielded consistently high results, with precision, recall, specificity, accuracy, and DSC values of 97.69%, 97.68%, 97.67%, 97.68%, and 97.68%, respectively, confirming the robustness, reliability, and generalizability of the proposed framework across multiple biomedical datasets.

3. Materials and Methods

3.1 Datasets

The proposed paper was evaluated using the C-NMC_Leukemia dataset [11] for automated blood cancer (ALL) cell prediction. Acute lymphoblastic leukemia (ALL) accounts for approximately 25% of pediatric cancers, and its microscopic diagnosis remains challenging due to the high morphological similarity between normal lymphocytes and immature leukemic blast cells. Under microscopic examination, these visual similarities make manual differentiation difficult even for experienced clinicians. In this dataset, all images were expert-annotated by trained oncologists and labeled as either normal or malignant, ensuring high-quality ground truth for supervised learning.

In study to ensure unbiased training and evaluation, the Leukemia dataset was systematically divided into three subsets. a training set comprising 9,594 images (90%), a validation set of 533 images (5%), and a test set of 534 images (5%). Also, this structured data partitioning supports Reliable performance assessment and model generalization.

This dataset provides a clinically relevant benchmark for developing and validating automated deep learning frameworks for blood cancer detection, particularly in the context of peripheral blood smear–based diagnosis, where subtle morphological variations must be accurately captured for reliable classification.

3.2 Data Preprocessing

In terms of illumination, staining intensity, background noise, resolution, and contrast, peripheral blood smear images used for blood cancer (leukemia) classification exhibit significant variability, this can negatively affect convolutional feature learning if not properly addressed. To achieve optimal compatibility with the EfficientNet-B3 architecture and to maximize feature extraction performance, a structured and model-aware preprocessing pipeline was implemented.

Initially, pixel-level noise suppression and artifact removal were applied to eliminate the background impurities, staining irregularities, and random noise commonly present in microscopic blood smear images. This step preserves diagnostically related cellular structures while reducing non-informative visual patterns that may interfere with convolutional learning.

Subsequently, contrast normalization was performed to harmonize intensity distributions across samples, compensating for variations in slide preparation, staining protocols, and illumination conditions. This enhances the visibility of significant cellular components such as nuclei, cytoplasm,

and cell boundaries, which are essential for accurate leukemia cell discrimination.

All images were then rescaled to 300 × 300 pixels, matching the native input resolution of EfficientNet-B3, thereby ensuring architectural compatibility and preserving fine-grained morphological details required for effective feature learning.

Finally, pixel intensity normalization was applied by scaling values to the range [0, 1], improving numerical stability, gradient flow, and convergence behavior during CNN training.

This preprocessing approach ensures input consistency, noise robustness, and feature integrity, enabling EfficientNet-B3 to learn discriminative multi-scale representations of leukemic and healthy blood cells. As a result, this proposed framework achieves improved training stability, faster convergence, and enhanced classification performance for automated blood cancer detection.

3.2.1. Data Pre-Processing

In this study, data preprocessing was essential because the blood smear images contain variations in resolution, pixel-level noise, size, bright text, and artifacts [25]. To mitigate these issues, a pixel-level mask was applied to remove unwanted artifacts, as defined in Equation (2):

$$\text{Mask}(m,n) = \begin{cases} i(m,n) - \min_i & \text{if } i(m,n) > \min_i \\ 0 & \text{otherwise} \end{cases}$$

Here, $i(m,n)$represents the pixel intensity at location (m, n)in each of the three primary color channels (RGB), and $\min_i$is the minimum intensity value within the local neighborhood. This procedure effectively suppressed background noise, text, and other non-cellular artifacts.

Additionally, the contrast of images often varied because of differences in staining and imaging conditions. To standardize the dataset, contrast normalization was performed during the training phase. This process ensured that all images had a consistent intensity distribution, which is important for stable feature extraction by the CNN. After contrast normalization, further noise filtering was applied to reduce unwanted background variations. Each pixel's value was adjusted by subtracting the average intensity of the three primary colour channels (RGB). This step enhanced the visibility of cellular structures and minimized background interference.

To standardize pixel intensity, min-max normalization was employed. For each pixel intensity PI, the normalized pixel value $PI^*$was calculated as:

$$PI^* = \frac{(PI - \text{Min}_{\text{old}})(\text{Max}_{\text{new}} - \text{Min}_{\text{new}})}{\text{Max}_{\text{old}} - \text{Min}_{\text{old}}} + \text{Min}_{\text{new}}, \quad PI \in [0, n]$$

Where Max_{old} and Min_{old} are the maximum and minimum pixel intensities in the original image, and $\text{Max}_{\text{new}}$and Min_{new}define the new intensity range. In practice, the pixel values were scaled to the range $[0, 1]$by dividing by 255, ensuring uniform input for EfficientNet-B3.

Finally, all images were resized to 300 × 300 pixels to match the input size requirement of the EfficientNet-B3 architecture. Figure 2 illustrates the effect of preprocessing on a blood cell image, highlighting improved clarity and reduced noise, which facilitates better feature learning by the CNN.

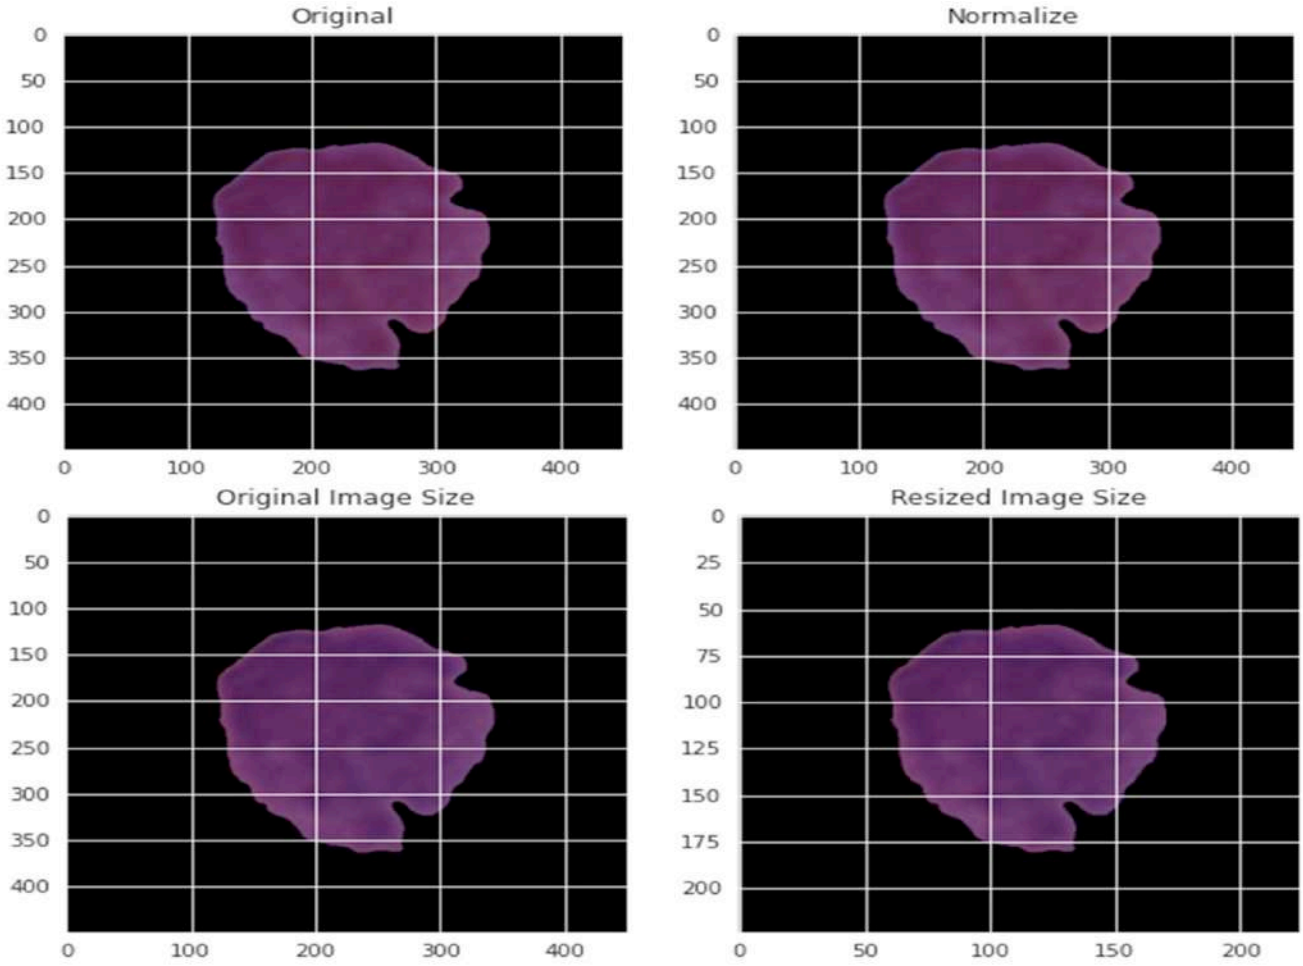

Figure 2. An example of applying the pre-processing stage on blood cells

3.2.2. EfficientNet-B3

Recently, Tan and Le [26] explored how the width and depth of convolutional neural networks (CNNs) affect performance and proposed an effective strategy for designing CNN architectures that achieve higher classification accuracy with fewer parameters. They introduced a family of models, referred to as EfficientNet-B0 through EfficientNet-B7, collectively known as the EfficientNet CNN models. When tested on the ImageNet dataset, these models outperformed contemporary architectures in terms of Top-1 accuracy while using significantly fewer parameters [27].

The EfficientNet family is based on a novel approach to scaling CNN models using a compound scaling coefficient. Unlike traditional methods that scale only one dimension of a network such as depth, width, or resolution. The EfficientNet uniformly scales all three dimensions using a set of carefully balanced scaling coefficients. Although scaling a single dimension can improve model performance, adjusting depth, width, and resolution together relative to available computational resources provides better overall performance. EfficientNet models are significantly smaller than conventional networks while maintaining comparable accuracy. For instance, ResNet50 in Keras contains 23,534,592 parameters, whereas the smallest EfficientNet, EfficientNet-B0, has only 5,330,564 parameters. In this study, we adopt EfficientNet-B3 as the backbone model for blood

cancer classification because it provides an optimal balance between computational efficiency and predictive performance [27].

The core building block of EfficientNet is the Mobile Inverted Bottleneck Convolution (MBConv), which is derived from the MobileNet architecture [28]. MBConv incorporates depth-wise separable convolutions, which consist of a depth-wise convolution followed by a pointwise convolution, significantly reducing the number of parameters. Additional improvements were inspired by MobileNetV2, including:

1. Inverted residual connections, which enhance gradient flow and enable deeper architectures.
2. Linear bottlenecks, which help retain important feature information while reducing dimensionality.

The EfficientNet architecture begins with a stem layer that processes the input image and is common across all EfficientNet variants. After the stem layer, the architecture contains seven sequential blocks, each composed of multiple sub-blocks, whose number and complexity increase progressively from EfficientNet-B0 to EfficientNet-B7. This well-structured scaling allows EfficientNet-B3 to capture multi-scale features effectively, making it particularly suitable for blood cancer image analysis, where subtle differences between healthy and leukemic cells are critical for accurate classification.

In the EfficientNet architecture, EfficientNet-B7 contains 813 layers. The network is organized into seven primary blocks, each comprising multiple sub-blocks [29]. The second module serves as the basis for the first sub-block of each of the seven main blocks, except for the first block. Module 3 is connected to all sub-blocks via skip connections, which enable the better gradient flow and feature reuse. In the first sub-blocks of each block, Module 4 is combined with these skip connections to enhance information propagation and improve the transfer of learned features across layers. Module 5 strengthens the integration between sub-blocks by establishing skip connections which facilitate the learning of residual features. Finally, each sub-block is constructed by systematically combining these modules in specific configurations, allowing the network to efficiently capture hierarchical and multi-scale features while maintaining a low parameter count. Figure 3 shows the EfficientNetB3 model structure.

ISSN: 2096-3246
Volume 58, Issue 01, 2026

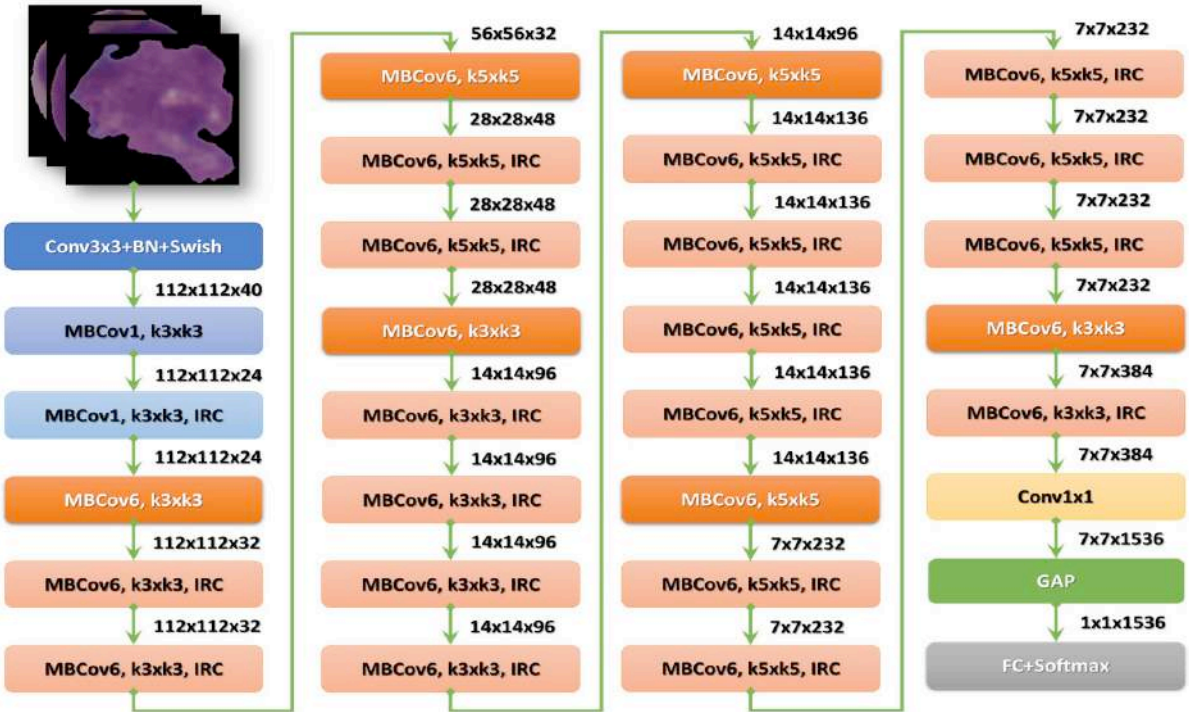

Figure 3. Structure of the EfficientNetB3 model.

3.2.3. Dwell

The DWELL callback becomes particularly useful when the validation loss of the current epoch exceeds that of the previous epoch during training. This condition indicates that the model has reached a point in the parameter space (N-dimensional space, where N is the number of trainable parameters) that is less favorable for generalization compared to the previous epoch. The DWELL callback monitors this condition and, if detected, restores the model weights to those corresponding to the epoch with the lowest validation loss [30]. Parallelly, it reduces the learning rate, preventing the optimizer from moving further into the unfavorable region. Without this adjustment, the model would remain at the suboptimal point in the parameter space if the learning rate were maintained for the next epoch. This mechanism ensures stability in training and helps the network converge to a more optimal solution, improving overall performance and reducing overfitting.

3.3 EfficientNet-B3 Architecture

The EfficientNet model family is built upon the principle of compound scaling, a systematic design strategy that simultaneously balances network depth, width, and input resolution to achieve optimal performance under constrained computational resources. Unlike conventional CNN architectures which scale these dimensions independently, compound scaling ensures that model capacity grows in a harmonized manner, leading to improved feature representation and training stability. Within this family, EfficientNet-B3 has been widely recognized in the literature as providing an effective balance between computational efficiency and classification accuracy, making it particularly well-suited for medical imaging applications where both performance and resource efficiency are critical. EfficientNet-B3 uses a sequence of mobile inverted bottleneck convolution (MBConv) blocks, which

ISSN: 2096-3246
Volume 58, Issue 01, 2026

integrate depth-wise separable convolutions, pointwise convolutions, and residual skip connections. This architectural design minimizes the number of trainable parameters and floating-point operations while preserving high representational capacity. Depth-wise separable convolutions extract detailed characteristics from complicated biomedical images, such as peripheral blood smear images used to diagnose blood cancer, while minimizing computation. The inclusion of skip connections improves gradient flow, mitigates vanishing gradient problems, and supports deeper network training, while the inverted bottleneck structure allows efficient channel expansion and compression, facilitating robust feature transformation across layers. These properties help EfficientNet-B3 to capture multi-scale morphological patterns, including variations in cell size, nuclear structure, cytoplasmic texture, and shape irregularities, which are critical indicators for distinguishing leukemic cells from healthy blood cells.

Recent research has demonstrated the effectiveness of EfficientNet architectures in various hematological and histopathological imaging tasks, including leukemia detection, blood cell classification, and cancer screening, with consistently superior generalization performance compared to traditional CNN models. In blood cancer analysis, EfficientNet-B3 provides an ideal architectural foundation by combining computational efficiency, strong feature extraction capability, and stable convergence behavior, enabling accurate and scalable automated diagnosis from peripheral blood smear images. Figure 4 presents the flow of the proposed model.

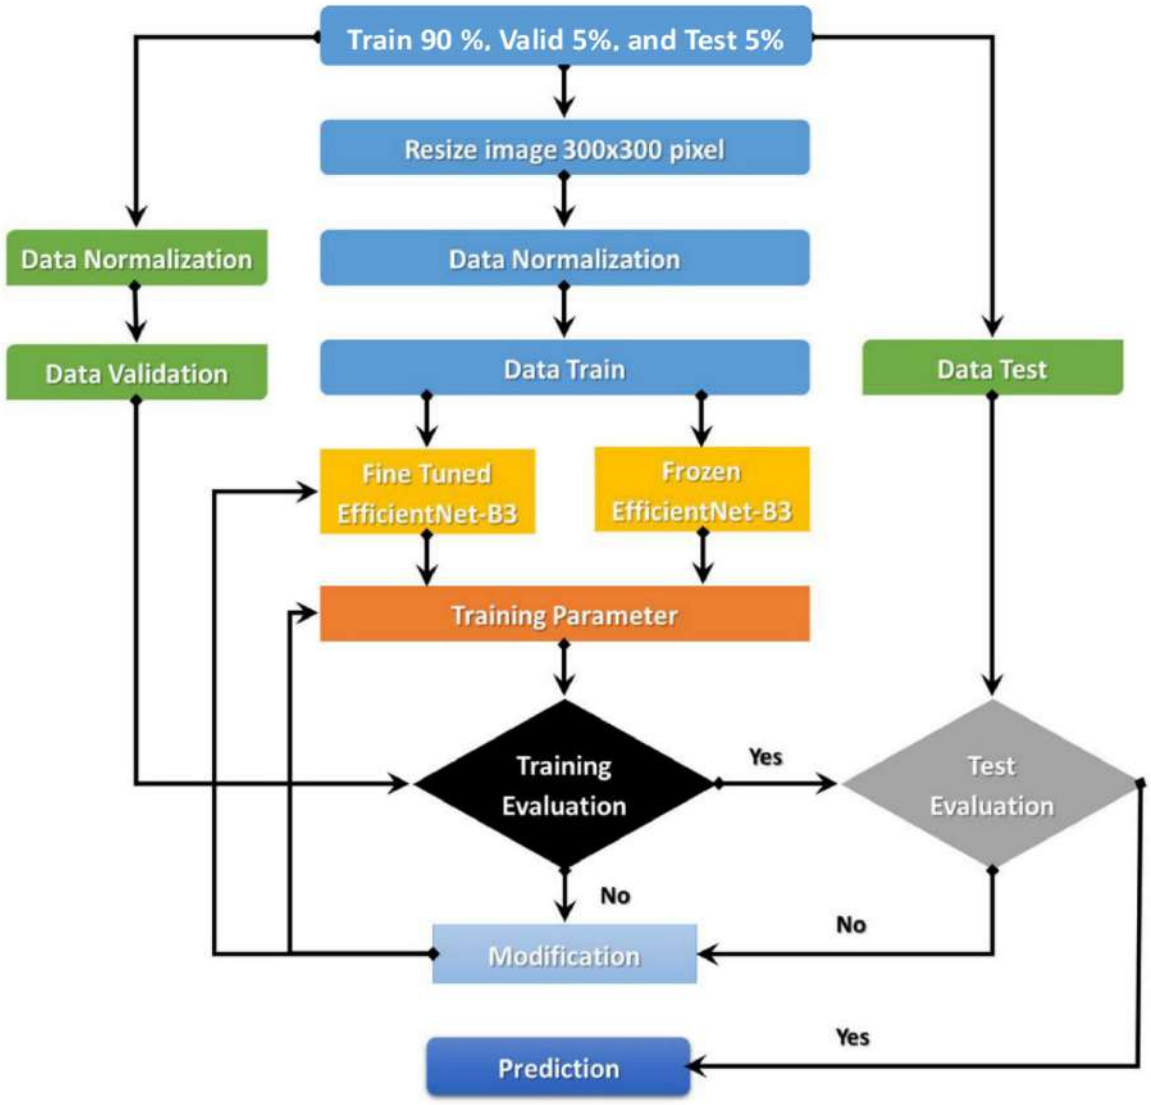

Figure 4. Flowchart of the proposed model.

The proposed blood cancer detection framework begins with the acquisition of the Leukemia image dataset, followed by a comprehensive image preprocessing pipeline. Preprocessing plays a important role in ensuring model reliability and diagnostic accuracy and therefore represents one of the most essential stages of the system. The dataset was processed using intensity normalization, contrast enhancement, spatial resizing, and artifact suppression, enabling standardization of visual quality and reduction of noise caused by staining variations, illumination differences, and imaging artifacts commonly present in peripheral blood smear images.

After preprocessing, the dataset was systematically divided into three subsets to support robust learning and unbiased evaluation: a training set comprising 90% of the images, a validation set of 5%, and a test set of 5%. This structured partitioning ensures stable optimization, hyperparameter

tuning, and reliable performance assessment.

In the next stage, transfer learning was used to train the EfficientNet-B3 architecture. The process began with supervised pre-training using EfficientNet-B3 model initialized with weights learned from the large-scale ImageNet dataset. This initialization improves convergence speed and learning stability, particularly when training on medical datasets with limited sample sizes.

After pre-training, the EfficientNet-B3 model was fine-tuned using the Leukemia image dataset. This step helps the model to adapt its feature representations to blood cancer detection by identifying specific morphological patterns, including variations in cell shape, nuclear structure, cytoplasmic texture, and cellular boundaries that distinguish leukemic cells from healthy blood cells.

During training, the training error was continuously monitored. If the training error did not decrease to an acceptable level, the model was retrained by adjusting the parameter and refining the optimization process. Once the training error reached an acceptable level, the model performance was evaluated using the test dataset. If test error remained high, the training process was repeated to improve generalization. This iterative optimization strategy helps to achieve both learning stability and diagnostic reliability. To further enhance convergence and training efficiency, a dynamic learning rate (LR) mechanism was integrated into the training process. The learning rate was adaptively adjusted based on model performance instead of fixed schedules which enable responsive optimization and stable convergence. The complete procedural workflow and adaptive optimization strategy are formally described in Table 1, which presents the algorithm governing the dynamic learning rate mechanism.

Table 1. The dynamic LR algorithm.

Input : Leukemia image dataset D = {Train, Validation, Test}
Output : Classification accuracy Acc and loss L

1: Load Train, Validation, and Test datasets from the C-NMC Leukemia repository
2: For each image Ii ∈ D do
3: Resize Ii to 300 × 300 and normalize using Ii_norm = Ii / 255
4: Apply histogram equalization: Ii_eq = HistEq(Ii_norm)
5: End For
6: Apply augmentation using flip, rotation, and zoom transformations
7: Initialize EfficientNet-B3 with ImageNet pre-trained weights
8: Replace the output layer using Softmax activation:
P(y=i) = ezi / Σj ezj
9: Initialize Adam optimizer with learning rate LR0 = 1 × 10^-3
10: Compute categorical cross-entropy loss:
L = - Σ yi log(ŷi)
11: For epoch = 1 to 50 do
12: Train the model and compute validation loss ValLoss
13: If ValLoss(t) > ValLoss(t−1), update:

LRnew = max(LRold × 0.5 , 10^-6)
14: End For
15: Evaluate test dataset and compute:
Acc = (TP + TN) / (TP + TN + FP + FN)

4. Experimental Results

4.1 Evaluation Metrics

The performance of the proposed model was evaluated using accuracy, precision, sensitivity, specificity, and the Dice Similarity Coefficient (DSC). These metrics are mathematically defined in Equations (4)–(8).

$$\text{Accuracy} = (TP + TN)/(TP + TN + FP + FN) \quad (4)$$

$$\text{Precision} = TP/(TP + FP) \quad (5)$$

$$\text{Sensitivity} = TP/(TP + FN) \quad (6)$$

$$\text{Specificity} = TN/(TN + FP) \quad (7)$$

$$\text{DSC} = (2\ TP)/(2\ \_\ TP + FP + FN) \quad (8)$$

True positives (TP), true negatives (TN), false positives (FP), and false negatives (FN) are the fundamental outcomes in classification. Precision measures the proportion of correctly predicted positive samples out of all samples predicted as positive by the model. Sensitivity (or recall) is defined as the proportion of actual positive samples that are correctly identified by the model. Specificity quantifies the proportion of actual negative samples that are correctly classified. The Dice Similarity Coefficient (DSC) represents the harmonic mean of precision and recall, providing a balanced metric for evaluating the overlap between predicted and true positive cases.

4.2 Leukemia Detection Performance

The Leukemia dataset was divided into 90% for training, 5% for testing, and 5% for validation. All experiments were conducted in the Kaggle environment. The hardware configuration of the PC used for the experiments consisted of an x64-based Intel® Core™ i7-10510U CPU operating at 1.80-2.30 GHz, with 16 GB of RAM on a 64-bit Windows platform.

Table 2 presents the results of the experiments for EfficientNet-B3 and five other CNN models-EfficientNet-B0, EfficientNet-B1, EfficientNet-B2, InceptionResNetV2, and DenseNet121-using a fixed learning rate. All models were trained on the Leukemia dataset for a binary classification task, distinguishing between benign and malignant blood cells. The models achieved average accuracies of 97.57% (EfficientNet-B3), 94.38% (EfficientNet-B1), 95.51% (EfficientNet-B2), 93.82% (EfficientNet-B0), (InceptionResNetV2), and 82.21% (DenseNet121). 93.07%

Among all models, EfficientNet-B3 exhibited the highest overall performance, with an average precision of 97.42%, recall of 96.96%, accuracy of 97.57%, and Dice Similarity Coefficient (DSC) of 97.18%. Meanwhile, EfficientNet-B2 achieved the highest average specificity, equal to 98.90%, highlighting its strength in correctly identifying negative samples.

Table 2. The performance of the six CNN models using a fixed LR.

Model	Class	Precision (%)	Recall (%)	Specificity (%)	Accuracy (%)	DSC (%)
EfficientNet-B3	all	97.82	98.63	95.29	97.57	98.22
	hem	97.01	95.29	98.63	97.57	96.14
Average		97.42	96.96	96.96	97.57	97.18
EfficientNet-B0	all	91.90	99.73	81.18	93.82	95.65
	hem	99.28	81.18	99.73	93.82	89.32
Average		95.59	90.46	90.45	93.82	92.49
EfficientNet-B1	all	97.99	93.68	95.88	94.38	95.78
	hem	87.63	95.88	93.68	94.38	91.57
Average		92.81	94.78	94.78	94.38	93.68
EfficientNet-B2	all	94.74	98.90	88.24	95.51	96.77
	hem	95.59	95.51	98.90	95.51	92.59
Average		95.17	97.21	93.57	95.51	94.68
InceptionResNetV2	all	93.69	97.53	83.53	93.07	95.05
	hem	94.04	83.53	97.53	93.07	88.47
Average		93.87	90.53	90.53	93.07	91.76
DenseNet121	all	91.38	81.59	83.53	82.21	86.21
	hem	67.94	83.53	81.59	82.21	74.93
Average		79.66	82.56	82.56	82.21	80.57

Table 2 summarizes the binary classification performance of six CNN models-EfficientNet-B3, EfficientNet-B0, EfficientNet-B1, EfficientNet-B2, InceptionResNetV2, and DenseNet121 on the C-NMC Leukemia dataset, distinguishing between ALL (cancer) and Hem (healthy) classes.

For the ALL class, EfficientNet-B3 achieved precision, recall, specificity, accuracy, and DSC of 97.82%, 98.63%, 95.29%, 97.57%, and 98.22%, respectively. EfficientNet-B3 and EfficientNet-B1 recorded the highest precision at 97.01% and 99.28%, respectively. The highest recall values were observed in EfficientNet-B3 (98.63%), EfficientNet-B0 (99.73%), and EfficientNet-B2 (98.90%). EfficientNet-B3 and EfficientNet-B1 achieved the highest specificity, 95.29% and 95.8%, respectively. The highest accuracy was obtained by EfficientNet-B3 (97.57%) and EfficientNet-B2 (95.51%), while EfficientNet-B3 (98.22%) and EfficientNet-B2 (96.77%) achieved the highest DSC.

For the Hem class, EfficientNet-B3 achieved precision, recall, specificity, accuracy, and DSC of 97.01%, 95.29%, 98.63%, 97.57%, and 96.14%, respectively. The highest precision values were achieved by EfficientNet-B3 (97.82%) and EfficientNet-B0 (97.99%). The highest recall was observed in EfficientNet-B3 (95.29%), EfficientNet-B1 (95.88%), and EfficientNet-B2 (95.51%). EfficientNet-B3 (98.63%), EfficientNet-B0 (99.73%), EfficientNet-B2 (98.90%), and InceptionResNetV2 (97.53%) achieved the highest specificity.

The top accuracy values were recorded by EfficientNet-B3 (97.57%) and EfficientNet-B2 (95.51%), while EfficientNet-B3 (96.14%) and EfficientNet-B2 (92.59%) achieved the highest DSC.
Figure 5 illustrates the training and validation losses and accuracies of the six CNN models using a fixed learning rate, highlighting the superior convergence and stability of EfficientNet-B3.

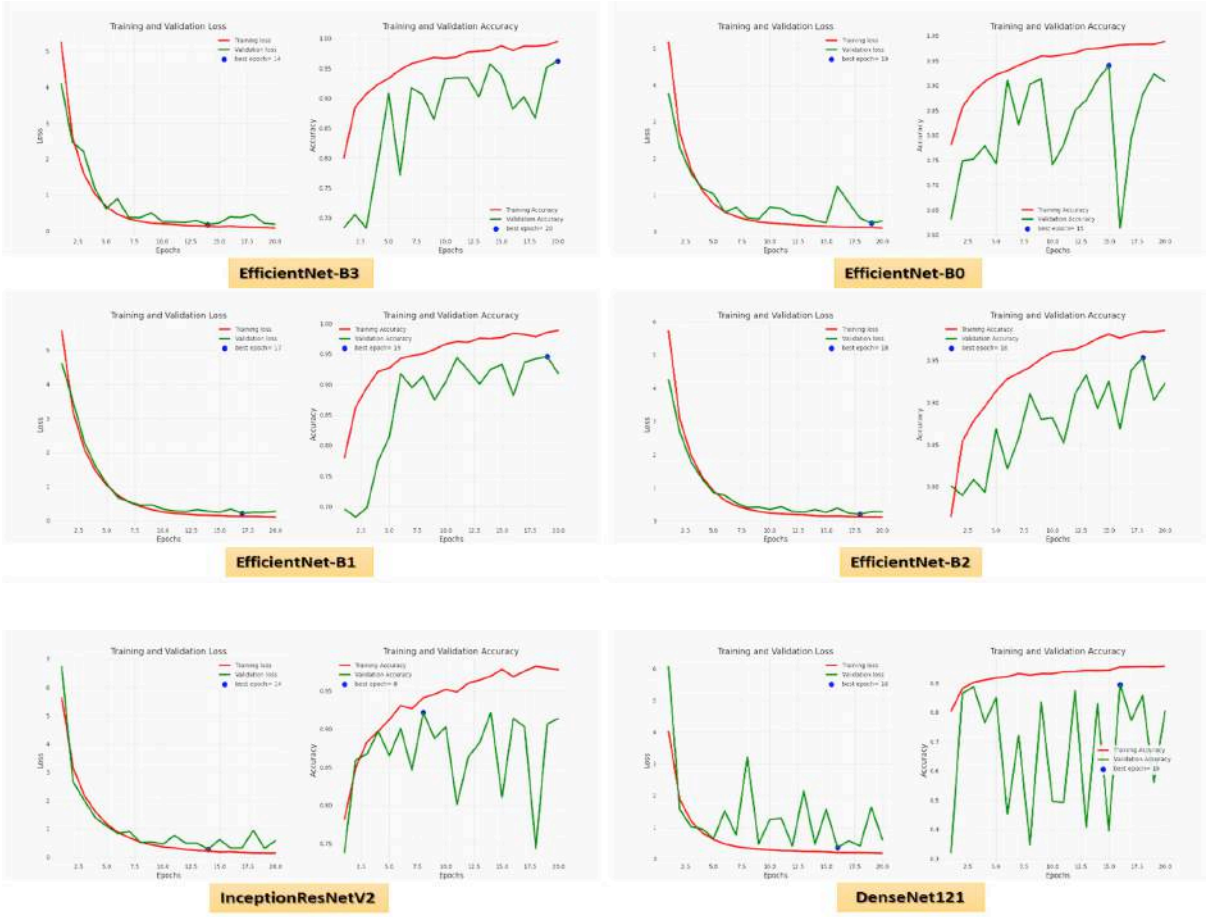

Figure 5. The loss and accuracy curves for the six CNN models using a fixed LR.

Figure 6 illustrates the confusion matrices of the six CNN models on the test dataset using a fixed learning rate.
The test dataset consisted of 364 ALL (cancer) images and 170 Hem (healthy) images.
For the ALL class, EfficientNet-B3 correctly classified 359 images, achieving an accuracy of 98.6%. EfficientNet-B0 predicted 363 images correctly (99.7%), while EfficientNet-B1 achieved 93.6% accuracy with 341 correct predictions. EfficientNet-B2 correctly classified 360 images (98.9%), Inception ResNetV2 predicted 355 images correctly (97.5%), and DenseNet121 correctly classified 297 images (81.5%).
These results demonstrate that EfficientNet-B3 and EfficientNet-B2 exhibit superior performance in accurately identifying ALL cells, while DenseNet121 showed comparatively lower accuracy.
Table 3 presents the experimental results of EfficientNet-B3 and five other CNN models-

EfficientNet-BO, EfficientNet-B1, EfficientNet-B2, InceptionResNetV2, and DenseNet121-using a dynamic learning rate (LR). All models were evaluated on the C-NMC Leukemia dataset for binary classification, distinguishing between ALL (cancer) and Hem (healthy) cells.

The results indicate that EfficientNet-B3, EfficientNet-BO, EfficientNet-B1, EfficientNet-B2, InceptionResNetV2, and DenseNet121 achieved average accuracies of 98.31%, 97.57%, 97.19%, 97.00%, 95.32%, and 94.38%, respectively. Among all models, EfficientNet-B3 demonstrated the highest overall performance, with an average precision of 98.29%, recall of 97.83%, accuracy of 98.31%, and Dice Similarity Coefficient (DSC) of 98.05%. The highest average specificity (98.08%) was achieved by both EfficientNet-B1 and EfficientNet-B2.

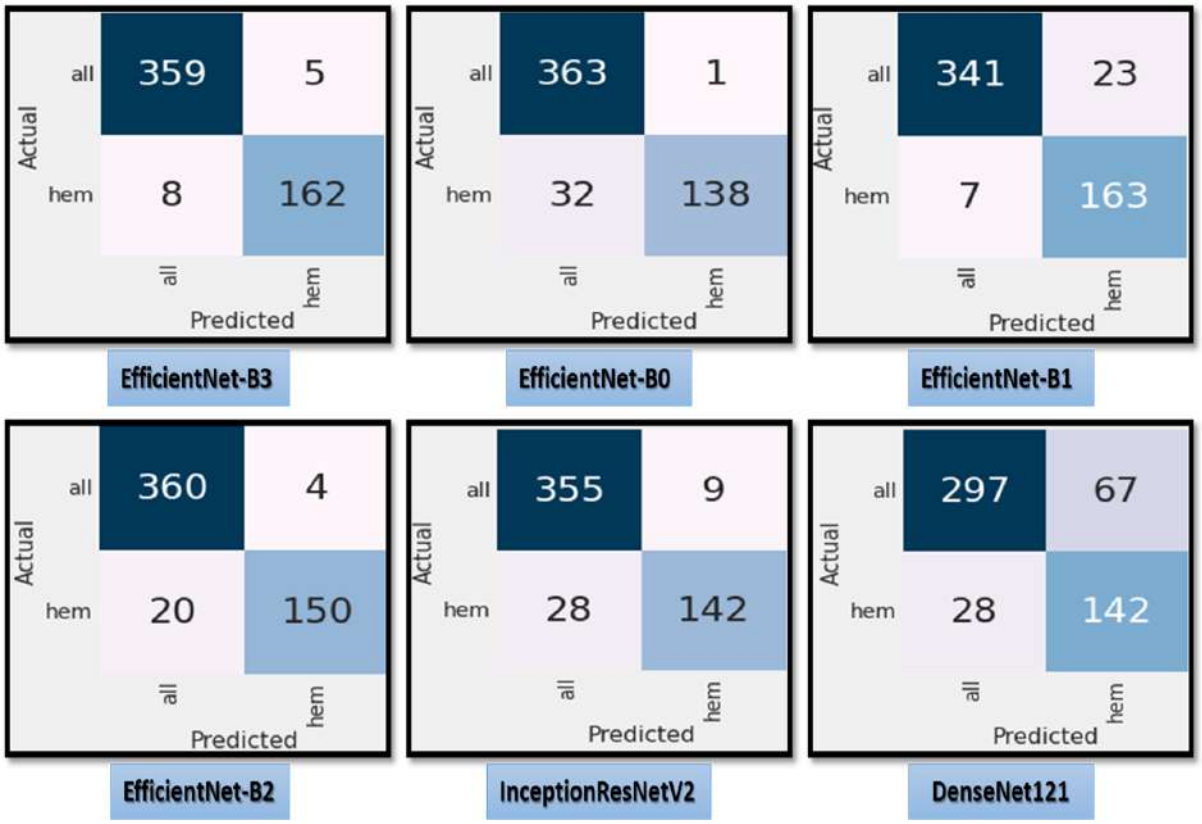

Figure 6. The test dataset's confusion matrix for the six CNN models using a fixed LR

For the ALL class, EfficientNet-B3 achieved precision, recall, specificity, accuracy, and DSC of 98.37%, 99.18%, 96.47%, 98.31%, and 98.77%, respectively. The precision values for EfficientNet-B3, Efficient Net-B0, EfficientNet-B1, and EfficientNet-B2 were 98.37%, 98.37%, 97.81%, and 97.54%, respectively. The highest recall was observed in EfficientNet-B3 (99.18%), EfficientNet-B0 (99.18%), and InceptionResNetV2 (98.35%). Both EfficientNet-B3 and Efficient Net-B0 achieved the top specificity of 96.47%. In terms of accuracy, the leading models were EfficientNet-B3 (98.31%), EfficientNet-B0 (98.31%), and EfficientNet-B1 (97.19%). The highest DSC values were recorded by EfficientNet-B3 (98.77%), EfficientNet-B0 (98.77%), and EfficientNet-B1 (97.94%), highlighting their superior ability to accurately distinguish ALL cells from healthy blood

cells.

For the Hem (healthy) class, Efficient Net-B3 achieved precision, recall, specificity, accuracy, and DSC of 98.20%, 96.47%, 99.18%, 98.31%, and 97.33%, respectively. Both EfficientNet-B3 and EfficientNet-BO recorded the highest precision, recall, accuracy, and DSC, at 98.20%, 96.47%, 98.31%, and 97.33%, respectively. The highest specificity values were observed in EfficientNet-B3 (99.18%), EfficientNet-BO (99.18%), and InceptionResNetV2 (98.35%), demonstrating their effectiveness in accurately identifying healthy blood cells.

Table 3. Performance of the six CNN models using the proposed dynamic LR.

Model	Class	Precision (%)	Recall (%)	Specificity (%)	Accuracy (%)	DSC (%)
EfficientNet-B3	ALL	98.37	99.18	96.47	98.31	98.77
	Hem	98.20	96.47	99.18	98.31	97.33
Average		98.29	97.83	97.82	98.31	98.05
EfficientNet-B0	ALL	98.37	99.18	96.47	98.31	98.77
	Hem	98.20	96.47	99.18	98.31	97.33
Average		95.38	97.06	97.80	97.57	96.21
EfficientNet-B1	ALL	97.81	98.08	95.29	97.19	97.94
	Hem	95.86	95.29	98.08	97.19	95.58
Average		95.86	95.29	98.08	97.19	95.58
EfficientNet-B2	ALL	97.54	98.08	94.70	97.00	97.80
	Hem	95.83	94.71	98.08	97.00	95.27
Average		95.83	94.71	98.08	97.00	95.27
InceptionResNetV2	ALL	94.96	98.35	88.82	95.32	96.63
	Hem	96.18	88.82	98.35	95.32	92.35
Average		95.57	93.59	93.59	95.32	94.49
DenseNet121	ALL	97.44	94.23	94.71	94.38	95.81
	Hem	88.46	94.71	94.23	94.38	91.48
Average		92.95	94.47	94.47	94.38	93.64

5. Discussion

The experimental analysis clearly demonstrates the effectiveness of the proposed adaptive EfficientNet-B3-based framework for automated acute lymphoblastic leukemia (ALL) detection from peripheral blood smear images. Compared with fixed learning rate strategies and other deep convolutional architectures, the integration of a dynamic learning rate optimization mechanism significantly improved both convergence stability and classification performance.

From the comparative evaluation using a fixed learning rate, EfficientNet-B3 consistently outperformed EfficientNet-BO, EfficientNet-B1, EfficientNet-B2, InceptionResNetV2, and DenseNet121 across all key evaluation metrics. Although EfficientNet-B0 and EfficientNet-B2

showed competitive recall and specificity values in certain cases, they exhibited less balanced performance when precision, recall, and Dice Similarity Coefficient (DSC) were jointly considered. This imbalance is critical in medical diagnosis, where false negatives can delay treatment and false positives can lead to unnecessary clinical interventions. EfficientNet-B3 achieved a superior trade-off between sensitivity and specificity, indicating reliable discrimination between leukemic blast cells and healthy lymphocytes.

The adoption of the proposed dynamic learning rate mechanism further enhanced model performance. As observed in Table 3, the dynamic learning rate improved the average accuracy of EfficientNet-B3 from 97.57% (fixed LR) to 98.31%, while also increasing precision, recall, and DSC values. This improvement highlights the importance of adaptive optimization in medical image analysis, where loss landscapes are highly nonconvex and prone to saddle points. By continuously monitoring validation loss and training accuracy, the dynamic learning rate effectively adjusted the step size, enabling faster convergence while preventing overfitting and training instability.

The DWELL-based learning rate adjustment played a crucial role in restoring optimal model weights when validation loss increased, ensuring that the network did not drift toward suboptimal parameter regions. This behavior is particularly important for biomedical datasets such as C-NMC Leukemia, which exhibit high intraclass similarity and subtle interclass variations. The proposed optimization strategy allowed EfficientNet-B3 to learn discriminative multi-scale morphological features such as nuclear irregularity, cytoplasmic texture, and cell boundary distortions that are essential for accurate ALL detection.

In comparison with previously reported methods, including ensemble CNNs, attention-based models, object detection approaches, and traditional machine learning classifiers, the proposed framework achieved superior or comparable performance while maintaining architectural simplicity and computational efficiency. Unlike ensemble models that require multiple networks and higher inference costs, the proposed method relies on a single optimized backbone, making it more suitable for real-world clinical deployment.

Furthermore, the robustness of the proposed approach was validated through cross-dataset evaluation, demonstrating consistent performance across different biomedical datasets. This confirms the generalization capability of the adaptive EfficientNet-B3 model and its potential applicability as a generalized hematological diagnostic tool rather than a dataset-specific solution.

Overall, the experimental findings confirm that combining an EfficientNet-B3 backbone with a dynamic learning rate optimization strategy yields a stable, accurate, and computationally efficient framework for automated leukemia detection, addressing many of the limitations observed in prior fixed-optimization approaches.

6. Conclusion

This study presents a fully automated deep learning-based computer-aided diagnostic framework for acute lymphoblastic leukemia detection using peripheral blood smear images. By integrating

EfficientNet-B3 with a dynamic learning rate optimization strategy, the proposed model achieves stable training, faster convergence, and superior classification performance. Experimental results on the C-NMC Leukemia dataset demonstrate high precision, recall, specificity, accuracy, and Dice Similarity Coefficient, outperforming several state-of-the-art CNN models. The proposed framework reliably distinguishes leukemic blast cells from healthy lymphocytes despite high morphological similarity, offering a non-invasive and clinically relevant solution for early leukemia screening. Future work will focus on multiclass leukemia classification, explainable AI integration, and large-scale multi-center validation to support real-world clinical deployment.

References

1. Abir,W.H.; Uddin, F.; Khanam, F.R.; Tazin, T.; Khan, M.M.; Masud, M.; Aljahdali, S. Explainable AI in Diagnosing and Anticipating Leukemia Using Transfer Learning Method. Comput. Intell. Neurosci. 2022, 2022, 5140148. [CrossRef] [PubMed]
2. Gehlot, S.; Gupta, A.; Gupta, R. SDCT-auxNet_: DCT Augmented Stain Deconvolutional CNN with Auxiliary Classifier For Cancer Diagnosis. Med. Image Anal. 2020, 61, 101661. [CrossRef] [PubMed]
3. World Health Organization. Breast Cancer Now Most Common Form of Cancer: WHO Taking Action. Available online: https://tinyurl.com/93eccmnv (accessed on 3 February 2021).
4. Mondal, C.; Hasan, K.; Ahmad, M.; Awal, A.; Jawad, T.; Dutta, A.; Islam, R.; Moni, M.A. Ensemble of Convolutional Neural Networks to Diagnose Acute Lymphoblastic Leukemia from Microscopic Images. Inform. Med. Unlocked 2021, 27, 100794. [CrossRef]
5. Laosai, J.; Chamnongthai, K. Classification of Acute Leukemia Using Medical-Knowledge-Based Morphology and CD Marker. Biomed. Signal Process. Control. 2018, 44, 127–137. [CrossRef]
6. Vogado, L.H.; Veras, R.M.; Araujo, F.H.; Silva, R.R.; Aires, K.R. Leukemia Diagnosis in Blood Slides using Transfer Learning in Cnns and Svm for Classification. Eng. Appl. Artif. Intell. 2018, 72, 415–422. [CrossRef]
7. American Society of Hematology. Hematology. Available online: https://www.hematology.org (accessed on 24 April 2021).
8. American Cancer Society. Key Statistics for Acute Lymphocytic Leukemia. Available online: https://www.cancer.org/cancer/acute-lymphocytic-leukemia/about/key-statistics.html (accessed on 24 April 2021).
9. Curesearch for Children's Cancer Research. Curesearch. Available online: https://curesearch.org/Acute-Lymphoblastic-Leukemia-in-Children (accessed on 20 April 2021).
10. Mohamed, M.; Far, B.; Guaily, A. An Efficient Technique for White Blood Cells Nuclei Automatic Segmentation. In Proceedings of the 2012 IEEE International Conference on Systems, Man, and Cybernetics (SMC), Seoul, Republic of Korea, 14 October 2012.

11. Zakir Ullah, M.; Zheng, Y.; Song, J.; Aslam, S.; Xu, C.; Kiazolu, G.D.;Wang, L. An Attention-Based Convolutional Neural Network for Acute Lymphoblastic Leukemia Classification. Appl. Sci. 2021, 11, 10662. [CrossRef]
12. Amin, M.M.; Kermani, S.; Talebi, A.; Oghli, M.G. Recognition of Acute Lymphoblastic Leukemia Cells in Microscopic Images using K-Means Clustering and Support Vector Machine Classifier. J. Med. Signals Sens. 2015, 5, 49–58. [PubMed]
13. Shrestha, Y.R.; Krishna, V.; von Krogh, G. Augmenting Organizational Decision-Making with Deep Learning Algorithms: Principles, Promises, and Challenges. J. Bus. Res. 2021, 123, 588–603. [CrossRef]
14. Xin, Y.; Kong, L.; Liu, Z.; Chen, Y.; Li, Y.; Zhu, H.; Gao, M.; Hou, H.; Wang, C. Machine Learning and Deep Learning Methods for Cybersecurity. IEEE Access 2018, 6, 35365–35381. [CrossRef]
15. Suzuki, K. Overview of Deep Learning in Medical Imaging. Radiol. Phys. Technol. 2017, 10, 257–273. [CrossRef] [PubMed]
16. Shen, D.; Wu, G.; Suk, H.I. Deep Learning in Medical Image Analysis. Annu. Rev. Biomed. Eng. 2017, 19, 221–248. [CrossRef] [PubMed]
17. Smith, L.N. Cyclical Learning Rates for Training Neural Networks. In Proceedings of the 2017 IEEE Winter Conference on Applications of Computer Vision (WACV), Santa Rosa, CA, USA, 24 March 2017.
18. Loshchilov, I.; Hutter, F. SGDR: Stochastic Gradient Descent with Warm Restarts. arXiv 2017, arXiv:1608.03983.
19. Izmailov, P.; Podoprikhin, D.; Garipov, T.; Vetrov, D.; Wilson, A.G. Averaging Weights Leads to Wider Optima and Better Generalization. arXiv 2018, arXiv:1803.05407.
20. Johny, A.; Madhusoodanan, K.N. Dynamic Learning Rate in Deep CNN Model for Metastasis Detection and Classification of Histopathology Images. Comput. Math. Methods Med. 2021, 2021, 5557168. [CrossRef] [PubMed]
21. Khandekar, R.; Shastry, P.; Jaishankar, S.; Faust, O.; Sampathila, N. Automated Blast Cell Detection for Acute Lymphoblastic Leukemia Diagnosis. Biomed. Signal Process. Control. 2021, 68, 102690. [CrossRef]
22. Almadhor, A.; Sattar, U.; Al Hejaili, A.; Mohammad, U.G.; Tariq, U.; Ben Chikha, H. An Efficient Computer Vision-Based Approach for Acute Lymphoblastic Leukemia Prediction. Front. Comput. Neurosci. 2022, 16, 1083649. [CrossRef] [PubMed]
23. Kasani, P.H.;Won Park, S.;Won Jang, J. An Aggregated-Based Deep Learning Method for Leukemic B-Lymphoblast Classification Diagnostics 2020, 10, 1064. [CrossRef] [PubMed]
24. Liu, Y.; Chen, P.; Zhang, J.; Liu, N.; Liu, Y.Weakly Supervised Ternary Stream Data Augmentation Fine-Grained Classification Network for Identifying Acute Lymphoblastic Leukemia. Diagnostics 2022, 12, 16. [CrossRef] [PubMed]
25. Malaria Cell Images Dataset; National Institute of Health (NIH). Available online: https://ceb.nlm.nih.gov/repositories/malariadatasets (accessed on 21 December 2020).

26. Tan, M.; Le, Q. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In Proceedings of the 36th International Conference on Machine Learning, Long Beach, CA, USA, 10–15 June 2020.
27. Alhichri, H.; Alswayed, A.S.; Bazi, Y.; Ammour, N.; Alajlan, N.A. Classification of Remote Sensing Images Using Efficientnet-B3 CNN Model with Attention. IEEE Access 2021, 9, 14078–14094. [CrossRef]
28. Sandler, M.; Howard, A.; Zhu, M.; Zhmoginov, A.; Chen, L.-C. MobileNetV2: Inverted Residuals and Linear Bottlenecks. In Proceedings of the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18 June 2018.
29. Torkamana, A.; Charkarib, N.M.; Aghaeipourc, M. An Approach for Leukemia Classification Based on Cooperative Game Theory. Anal. Cell. Pathol. 2011, 34, 235–246. [CrossRef]
30. Alam, E.U.; Banik, S.; Chowdhury, L. A Statistical Approach to Classify the Leukemia Patients from Generic Gene Features. In Proceedings of the 2020 International Conference on Computer Communication and Informatics (ICCCI-2020), Coimbatore, India, 22–24 January 2020.