

EfficientNet-B0 Convolutional Network for Early Detection and Severity Classification of Diabetic Retinopathy

Yuli Astuti¹, Siti Nurhaisah^{*2}

¹Diploma Program of Information Management, Faculty of Computer Science, Universitas Amikom Yogyakarta, Indonesia

²Department of Information System, Faculty of Computer Science, Universitas Amikom Purwokerto, Indonesia
e-mail: yuli@amikom.ac.id¹, sitnurhaisyah05@gmail.com²

Article History	Abstract
Received : 06 January 2026 Revised : 22 February 2026 Accepted : 02 March 2026 Published : 28 March 2026 Keywords: <i>Diabetic Retinopathy, EfficientNet-B0, Deep Learning, Medical Imaging, Transfer Learning.</i>	Diabetic retinopathy (DR) is one of the leading causes of preventable blindness worldwide. Early detection is essential to reduce the risk of vision loss; however, screening programs are often constrained by a shortage of qualified specialists. This study implemented a lightweight EfficientNet-B0 architecture for the automated classification of five diabetic retinopathy severity levels using a dataset of 3,662 color fundus images. The proposed methodology incorporated transfer learning, an 80:20 stratified data split, and a WeightedRandomSampler to address severe class imbalance (9.4:1 ratio between the majority and minority classes). The model was trained using the Adam optimizer with a learning rate of 1×10^{-4} and an early stopping mechanism. Experimental results showed an overall accuracy of 84% and a macro-average AUC of 0.94, demonstrating strong discriminative capability. The highest performance was achieved in the No_DR class (F1-score = 0.98) and the Moderate class (F1-score = 0.79), whereas the Severe and Proliferative_DR classes remained more challenging because of their limited number of training samples.

^{*}Corresponding author : sitnurhaisyah05@gmail.com

1. INTRODUCTION

Diabetes mellitus is one of the most pressing global health challenges, and its prevalence has risen steadily over the past decades. The International Diabetes Federation reported approximately 537 million adults living with diabetes in 2021, a number projected to reach 783 million by 2045 [1]. Diabetic retinopathy (DR) is a major microvascular complication of chronic hyperglycemia that causes progressive damage to the retinal blood vessels. It is one of the leading causes of preventable blindness, particularly among the working-age population. DR progresses through five stages, from no DR to mild, moderate, and severe non-proliferative DR, and finally to proliferative DR (PDR), which carries a high risk of permanent vision loss if left untreated [2], [3]. Early diagnosis and timely intervention substantially reduce the risk of severe vision loss and improve long-term patient outcomes [4].

Despite this clinical urgency, manual DR screening still faces several challenges. Evaluating fundus images requires experienced ophthalmologists and specialized clinical facilities, both of which are often limited in rural and resource-constrained settings. Manual interpretation is also time-consuming and susceptible to inter-observer variability, which can affect diagnostic consistency and the overall effectiveness of screening programs [2], [5]. Consequently, many DR cases are detected only at advanced stages, when visual damage has already become irreversible.

Recent advances in artificial intelligence (AI), particularly deep learning, have opened significant opportunities to automate medical image analysis, including the detection and classification of DR from fundus images. Convolutional Neural Networks (CNNs) extract discriminative features hierarchically and have consistently outperformed traditional machine learning methods in fundus image classification tasks [6]. Regulator-approved AI systems for DR detection have already demonstrated high diagnostic performance, validating the reliability of deep learning in real-world clinical settings [7], [8]. Among CNN architectures, EfficientNet-B0 uses a compound scaling strategy that uniformly

balances network depth, width, and input resolution, enabling an optimal trade-off between accuracy and computational efficiency compared with conventional architectures such as ResNet-50 [9], [10]. With only about 5.3 million parameters, EfficientNet-B0 can be trained and deployed on resource-constrained devices without significantly compromising classification performance [9].

Several studies have applied deep learning to diabetic retinopathy classification using different architectures. Bodapati et al. [11] proposed a composite deep neural network with a gated-attention mechanism that combined features from multiple pretrained CNNs to classify five DR severity levels. Majumder and Kehtarnavaz [12] combined a classification model with a regression model to exploit the ordinal relationship between DR stages, reporting weighted Kappa scores of 0.90 and 0.88 on the APTOS and EyePACS datasets, respectively. Mehboob et al. [13] introduced cascaded and parallel CNN frameworks that progressively classify fundus images into DR stages using a large Kaggle fundus dataset. Ahmed [14] compared ResNet and EfficientNet variants with extensive data augmentation on the APTOS 2019 dataset, reaching 84.6% accuracy and a 94.1% AUC for five-class grading, although the approach required generating tens of thousands of synthetic training samples per class to balance the dataset.

These findings indicate that although complex and hybrid architectures can improve classification performance, they often introduce substantial computational overhead that limits their practical deployment in large-scale screening and resource-constrained clinical settings. Therefore, this study develops and evaluates a Computer-Aided Diagnosis (CAD) system based on EfficientNet-B0 for multi-class diabetic retinopathy classification, emphasizing both diagnostic accuracy and computational efficiency.

2. METHOD

In the implementation method section, describe clearly and concisely the methods used to achieve the goals set out in the service activities. The results of the dedication must be measurable and the writer was asked to explain the measuring instrument used, both descriptively and qualitatively. Explain how to measure the level of achievement of the success of community service activities. The level of achievement can be seen from the attitude, socio-cultural and economic changes of the target community. The overall research workflow is illustrated in Figure 1, covering data collection, preprocessing, model development, training and optimization, and evaluation.



2.1. Dataset

This study utilizes a publicly available diabetic retinopathy dataset from Mendeley Data [15]. The dataset comprises 3,662 color fundus images categorized into five DR severity levels. Table I presents the class distribution in the dataset.

Table 1. Diabetic Retinopathy Dataset Distribution

Class	Number of Images	Percentage	Characteristics
No_DR	1,805	49.3%	Healthy retina without DR signs
Mild	370	10.1%	Early microaneurysms
Moderate	999	27.3%	Hemorrhages and exudates
Severe	193	5.3%	Venous beading
Proliferative_DR	295	8.0%	Neovascularization

As shown in Table 1, the dataset exhibits significant class imbalance, with the ratio between the majority and minority classes reaching 9.4:1 (No_DR to Severe).

2.2. Preprocessing and Data Augmentation

Fundus images typically present variability in resolution, illumination, and acquisition conditions. Preprocessing was therefore carried out following established medical imaging practices [3], [16]. All images were resized to 224×224 pixels to match the input requirement of EfficientNet-B0. Pixel values were normalized using the ImageNet mean (0.485, 0.456, 0.406) and standard deviation (0.229, 0.224, 0.225), consistent with transfer learning protocols for models pretrained on ImageNet [9], [17].

The dataset was divided into 80% training (2,930 images) and 20% validation (732 images) using stratified sampling to preserve class distribution. Stratified splitting is recommended for imbalanced medical datasets to ensure fair class representation during training and evaluation [9].

To improve generalization and reduce overfitting, the following augmentation techniques were applied to the training set:

- Random horizontal flip with a probability of 0.5
- Random rotation up to $\pm 15^\circ$
- Color jitter with brightness (0.2), contrast (0.2), and saturation (0.2) adjustments

Such augmentation techniques have been shown to improve robustness against real-world clinical variability in retinal imaging [18].

2.3. Handling Class Imbalance

Medical image datasets often suffer from class imbalance, where severe DR cases are far less common than normal cases. To mitigate this issue, a WeightedRandomSampler was implemented during training, in which the sampling weight for each class is calculated as shown in Equation (1):

$$W_i = N / (C \times n_i) \quad (1)$$

Where N is the total number of samples, C is the number of classes, and n_i is the number of samples in class i . This inverse-frequency weighting strategy increases the sampling probability of minority classes and reduces bias toward majority classes, an approach widely validated for mitigating class imbalance in deep learning models [9], [17].

2.4. Model Architecture

This study employs EfficientNet-B0 as the backbone architecture [9], selected for its:

- Parameter efficiency (approximately 5.3 million parameters)
- Compound scaling mechanism balancing depth, width, and resolution
- High performance with low computational cost

EfficientNet models have demonstrated competitive performance in DR classification tasks while maintaining computational efficiency [7], [19]. The network was initialized with pretrained ImageNet weights to leverage transfer learning. Transfer learning is a widely used strategy for medical imaging tasks with limited training data, although the resulting performance gain depends on the network architecture and the degree of feature reuse from the source domain [17].

The final fully connected layer was replaced with a new linear layer containing five output neurons corresponding to the five DR classes. The resulting model formulation is expressed in Equation (2):

$$y = \text{Linear}(\text{EfficientNet-B0}(x; \theta_{\text{pretrained}}); \theta_{\text{new}}) \quad (2)$$

where x is the input image, $\theta_{\text{pretrained}}$ denotes the pretrained parameters transferred from ImageNet, and θ_{new} denotes the classifier parameters trained from scratch.

2.5. Training Configuration

The model was trained using the PyTorch framework with GPU CUDA acceleration; the training hyperparameters are summarized in Table 2.

Table 2. Training Hyperparameters

Parameter	Value
Optimizer	Adam
Learning Rate	1×10^{-4}
Batch Size	32
Maximum Epochs	25
Loss Function	Cross Entropy Loss
Early Stopping Patience	5 epochs

The model was trained using the PyTorch framework with GPU CUDA acceleration; the training hyperparameters are summarized in Table 2. The Adam optimizer was selected for its adaptive learning rate mechanism and stable convergence behavior in deep neural networks, consistent with training configurations used in similar EfficientNetB0-based DR grading frameworks [20]. Cross-entropy loss was used as the objective function for the multi-class classification task, and early stopping monitored the validation loss to prevent overfitting; training was terminated if the validation loss did not improve for five consecutive epochs. All experiments were conducted on a system equipped with an NVIDIA Tesla T4 GPU, 16 GB RAM, CUDA 11.8, Python 3.10, PyTorch 2.0, and the timm library.

2.6. Model Evaluation

Model performance was evaluated using comprehensive quantitative metrics and analytical visualizations, computed from the number of true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN) obtained for each class using a one-vs-all approach.

2.6.1. Quantitative Metrics

Model performance was evaluated using accuracy, precision, recall, and F1-score. For imbalanced multi-class datasets, class-wise metrics provide more informative evaluation than accuracy alone [9]. Overall accuracy is defined in Equation (3):

$$\text{Accuracy} = (\text{TP} + \text{TN}) / (\text{TP} + \text{TN} + \text{FP} + \text{FN}) \quad (3)$$

For each class, precision and recall are computed as shown in Equations (4) and (5):

$$Precision = TP / (TP + FP) \quad (4)$$

$$Recall = TP / (TP + FN) \quad (5)$$

Precision and recall are then combined into the F1-score, as shown in Equation (6):

$$F1-Score = 2 \times (Precision \times Recall) / (Precision + Recall) \quad (6)$$

2.6.2. Visual Analysis

A confusion matrix was used to analyze misclassification patterns between severity levels [3]. Receiver Operating Characteristic (ROC) curves and the Area Under the Curve (AUC) were computed for each class and combined using macro-averaging to give equal weight to each class regardless of its size [9], [16], as shown in Equation (7):

$$AUC_{macro} = (1/C) \sum_{i=1}^C AUC_i \quad (7)$$

Where C is the number of classes.

3. RESULTS AND DISCUSSION

Table 3 summarizes the classification performance of the proposed EfficientNet-B0 model on the validation set.

Table 3. Classification Performance Metrics

Class Name	Precision	Recall	F1-Score	Support
Mild	0,60	0,70	0,65	74
Moderate	0,78	0,79	0,79	200
No_DR	0,98	0,98	0,98	361
Proliferate_DR	0,77	0,51	0,61	59
Severe	0,53	0,62	0,57	39
Accuracy			0,84	733
Macro Avg	0,73	0,72	0,72	733
Weighted Avg	0,85	0,84	0,84	733

Based on the quantitative results in Table 3, the EfficientNet-B0 model achieved an overall accuracy of 84%. The model performed particularly well on the No_DR class, with an F1-score of 0.98, indicating a near error-free ability to identify normal retinal images. In the Moderate category, the model also achieved stable performance, with an F1-score of 0.79.

Nevertheless, the Severe and Proliferative_DR classes were more challenging, with F1-scores of 0.57 and 0.61, respectively. The relatively low recall for the Proliferative_DR class (0.51) suggests that the model still has difficulty distinguishing the complex visual features associated with advanced disease stages. The WeightedRandomSampler nonetheless maintained competitive recall for the Mild (0.70) and Severe (0.62) classes despite their limited sample sizes. Overall, the weighted-average F1-score of 0.84 confirms the robustness of the proposed model for multi-class DR classification.

Figure 2 presents the confusion matrix illustrating the model's classification performance across the five severity levels.




Figure 2. Confusion Matrix

Figure 2 presents the confusion matrix, indicating that the model classified the No_DR class exceptionally well, with 352 out of 361 samples (97.5%) correctly identified. The model also achieved strong results on the Moderate class, correctly predicting 158 out of 200 samples (79.0%). Several misclassification patterns were nonetheless observed the Mild class tends to be confused with No_DR, likely due to the subtle visual characteristics of microaneurysms, while the Severe class is frequently confused with Moderate because of visual similarity in exudate accumulation and hemorrhages at the threshold level. The Proliferative_DR class shows the lowest recall (0.51), as the complex features of neovascularization are often misidentified as lower disease stages.

Figure 3 shows the ROC curves and the corresponding AUC values for each class.




Figure 3. ROC-AUC Curve

As shown in Figure 3, the ROC analysis demonstrates that the EfficientNet-B0 model has strong discriminative capability, achieving a macro-average AUC of 0.94. This high AUC value indicates that the model is effective as a medical decision-support system, providing strong sensitivity and specificity for prioritizing patients according to their risk of vision loss.

Figure 4 illustrates the web-based implementation of the trained model for practical use.




Figure 4. Website Diabetic Retinopathy Prediction

As shown in Figure 4, users can upload retinal fundus images directly through a web browser. The system then performs real-time inference using the stored best-performing model weights to predict the diabetic retinopathy severity level.

4. CONCLUSION

This study demonstrates the implementation of EfficientNet-B0 for automated diabetic retinopathy classification, achieving 84% overall accuracy and a 0.94 macro-average AUC across five severity levels. With only 5.3 million parameters, the model offers greater computational efficiency than previous approaches based on ensemble methods or larger architectures, making it well suited for deployment in resource-constrained clinical settings. The model performed particularly well in detecting healthy retinas (F1-score: 0.98) and moderate-stage DR (F1-score: 0.79), while the WeightedRandomSampler effectively mitigated the 9.4:1 class imbalance ratio, as reflected in the balanced macro-average metrics.

The model's performance on minority classes nonetheless remains a limitation, with the Proliferative_DR and Severe classes achieving F1-scores of only 0.61 and 0.57, respectively. The frequent misclassification of Proliferative_DR cases as Moderate (20 of 59 cases) represents a clinical concern that warrants further refinement. In addition, evaluation on a single dataset limits generalizability across different imaging protocols and demographic populations. These findings suggest that although the model shows promise as a screening tool, it should serve as a supportive diagnostic aid rather than a replacement for expert clinical examination, particularly for advanced disease stages.

Future work should explore advanced loss functions, such as focal loss, to improve minority-class performance; validate the model across multiple public datasets to assess generalizability; and integrate explainability techniques, such as Grad-CAM [20], to enhance clinical interpretability. Investigating probability calibration and optimal threshold-selection strategies could also help bridge the gap between high AUC values and practical classification performance. As the global diabetes epidemic continues to grow, efficient deep learning architectures such as EfficientNet-B0 have significant potential to expand access to diabetic retinopathy screening in underserved populations, ultimately helping to prevent avoidable vision loss through timely detection and intervention.

BIBLIOGRAPHY

- [1] D. J. Magliano and E. J. Boyko, *IDF DIABETES ATLAS, 10th edition*. National Center for Biotechnology Information, 2021. [Online]. Available: <https://www.ncbi.nlm.nih.gov/books/NBK581934/>
- [2] F. M. Dejene, T. G. Debelee, F. Schwenker, Y. M. Ayano, and D. W. Feyisa, "Diabetic retinopathy screening using machine learning : a systematic review," *BMC Biomed. Eng.*, vol. 7, p. 12, 2025, doi: 10.1186/s42490-025-00098-0.
- [3] D. Bhulakshmi and D. S. Rajput, "A systematic review on diabetic retinopathy detection and classification based on deep learning techniques using fundus images," *PeerJ Comput. Sci.*, p. 14, 2024, doi: 10.7717/peerj-cs.1947.
- [4] D. Dave, S. R. Steinhubl, and R. B. McQueen, "Deep Learning – Enabled Diabetic Retinopathy Screening : A Techno-Clinical Revolution or Just More Artificial Intelligence Hype ?," *Diabetes Care*, vol. 49, no. March, pp. 381–383, 2026.
- [5] T. Wang, Wei-TingLuo, Y.-K. Tu, Y.-B. Chou, and Y.-T. Wu, "Systematic review and meta-analysis of regulator-approved deep learning systems for fundus diabetic retinopathy detections," *Digit. Med.*, vol. 9, p. 110, 2026, doi: 10.1038/s41746-025-02223-8.
- [6] H. S. Alghamdi, "applied sciences Towards Explainable Deep Neural Networks for the Automatic Detection of Diabetic Retinopathy," *Appl. Sci.*, vol. 12, p. 9435, 2022, doi: 10.3390/app12199435.
- [7] L. Arora *et al.*, "Ensemble deep learning and EfficientNet for accurate diagnosis of diabetic retinopathy," *Sci. Rep.*, vol. 14, p. 30554, 2024, doi: 10.1038/s41598-024-81132-4.
- [8] R. Alsohemi and S. Dardouri, "Fundus Image-Based Eye Disease Detection Using EfficientNetB3 Architecture," *J. Imaging*, vol. 11, p. 279, 2025, doi: 10.3390/jimaging11080279.
- [9] S. Akhtar *et al.*, "A deep learning based model for diabetic retinopathy grading," *Sci. Rep.*, vol. 15, pp. 1–20, 2025, doi: 10.1038/s41598-025-87171-9 2.
- [10] Y. Huang, L. Lin, P. Cheng, J. Lyu, R. Tam, and X. Tang, "Identifying the Key Components in ResNet-50 for Diabetic Retinopathy Grading from Fundus Images : A Systematic Investigation," *Diagnostics*, vol. 13, p. 1664, 2023, doi: 10.3390/diagnostics13101664.
- [11] J. D. Bodapati, N. S. Shaik, and V. Naralasetti, "Composite deep neural network with gated-attention mechanism for diabetic retinopathy severity classification," *J. Ambient Intell. Humaniz. Comput.*, vol. 12, pp. 9825–9839, 2021, [Online]. Available: <https://link.springer.com/article/10.1007/s12652-020-02727-z#Abs1>
- [12] S. Majumder and N. Kehtarnavaz, "Multitasking Deep Learning Model for Detection of Five Stages of Diabetic Retinopathy," *Digit. Object Identifier*, vol. 9, pp. 123220–123230, 2021, doi: 10.1109/ACCESS.2021.3109240.
- [13] A. Mehboob, M. U. Akram, and N. S. Alghamdi, "A Deep Learning Based Approach for Grading of Diabetic Retinopathy Using Large Fundus Image Dataset," *Diagnostics*, vol. 12, p. 3084, 2022, doi: 10.3390/diagnostics12123084.
- [14] F. Ahmed, "Addressing High Class Imbalance in Multi-Class Diabetic Retinopathy Severity Grading with Augmentation and Transfer Learning," 2025, doi: 10.48550/arXiv.2507.17121.
- [15] P. Tuna, "Diabetic Retinopathy," Mendeley Data. [Online]. Available: <https://data.mendeley.com/datasets/nxcd8krdhg/1>
- [16] S. Tummala, V. Sainath, G. Thadikemalla, S. Kadry, and M. Sharaf, "EfficientNetV2 Based Ensemble Model for Quality Estimation of Diabetic Retinopathy Images from DeepDRiD," *Diagnostics*, vol. 13, p. 622, 2023, doi: 10.3390/diagnostics13040622.
- [17] S. Yi, X. Yang, T. Wang, F. She, X. Xiong, and J. He, "applied sciences Diabetic Retinopathy Diagnosis Based on RA-EfficientNet," *Appl. Sci.*, vol. 11, p. 11035, 2021, doi: 10.3390/app112211035.
- [18] P. Macsik, J. Pavlovicova, S. Kajan, J. Goga, and V. Kurilova, "Image preprocessing-based

-
- ensemble deep learning classification of diabetic retinopathy,” *IET Image Process.*, vol. 18, no. November 2023, pp. 807–828, 2024, doi: 10.1049/ipr2.12987.
- [19] S. Batool *et al.*, “Deploying efficient net batch normalizations (BNs) for grading diabetic retinopathy severity levels from fundus images,” *Sci. Rep.*, vol. 13, p. 14462, 2023, doi: 10.1038/s41598-023-41797-9.
- [20] V. Anand, D. Koundal, W. Y. Alghamdi, and B. M. Alsharbi, “Smart grading of diabetic retinopathy : an intelligent fine-tuned EfficientNetB framework,” *Front. Artif. Intell.*, vol. 7, p. 1396160, 2024, doi: 10.3389/frai.2024.1396160.