ISSN (Online) 2394-2320

International Journal of Engineering Research in Computer Science and Engineering (IJERCSE)

Volume 13, Issue 1, January 2026

Diabetic Retinopathy Detection by Means of Deep Learning

[1] Upendar Para, [2] Nandini Tejawat, [3] Nikshitha Varkala, [4] Keerthika Kotturu, [5] Vaeshnavi Lingampally

[1] [2] [3] [4] [5] Department of Computer Science & Engineering, Keshav Memorial Institute of Technology (KMIT), Hyderabad, Telangana, India

Corresponding Author Email: [1] upendar.para@gmail.com, [2] tejawatnandini.26@gmail.com, [3] nikshithavarkala@gmail.com, [4] keerthikakotturu11@gmail.com, [5] vaeshnavi3701@gmail.com

Abstract*— Diabetic Retinopathy (DR) is a severe eye disease caused by diabetes, often leading to blindness if not detected early. Traditional screening depends on manual examination of retinal fundus images by ophthalmologists, which is time-consuming and prone to errors. To address this, the project proposes an automated DR detection system using Deep Learning. The system employs Convolutional Neural Networks (CNNs) to analyze retinal images and detect key features such as microaneurysms, hemorrhages, and exudates, which indicate DR severity. The workflow includes image preprocessing (enhancement and normalization), feature extraction via CNNs, and a classification module that categorizes DR into stages: No DR, Mild, Moderate, Severe, and Proliferative. A final report generation module provides quick results for doctors and patients. The project will be implemented using Python, TensorFlow/Keras, and OpenCV for preprocessing and training on large datasets. This AI-driven solution enables early detection, automation, scalability, and affordability, improving patient care and preventing blindness.

Keywords: CNN, Deep Learning, Diabetic Retinopathy, Medical Image Processing, Retinal Images.

I. INTRODUCTION

Diabetic retinopathy (DR) is one of the most common complications of diabetes mellitus and a leading cause of vision impairment and blindness worldwide. It occurs due to prolonged high blood sugar levels that damage retinal blood vessels, leading to progressive vision loss if not detected early. With the rising prevalence of diabetes, the number of individuals affected by diabetic retinopathy continues to increase, making early diagnosis a major public health priority.

Traditional DR diagnosis relies on manual examination of retinal fundus images by trained ophthalmologists. While effective, this approach is time-consuming, costly, and limited by the availability of specialists, particularly in rural and resource-constrained regions. These challenges often result in delayed diagnosis and reduced treatment effectiveness.

Recent advances in artificial intelligence and deep learning have enabled automated analysis of retinal images for diabetic retinopathy detection. Convolutional Neural Networks (CNNs) can accurately learn disease-specific features and provide fast, consistent screening results. Integrating deep-learning-based screening systems into clinical practice can improve early detection, reduce clinician workload, and help prevent diabetes-related vision loss.

Such automated systems also support large-scale screening programs and reducing human variability. Furthermore, AI-based approaches can enhance accessibility to eye care services, enabling timely intervention and improved patient outcomes.

II. MOTIVATION AND OBJECTIVE

Diabetes is increasingly common in India, and a large number of patients live in rural or semi-urban areas where access to eye-care specialists is limited. Many individuals with diabetes are unaware of diabetic retinopathy until it reaches an advanced stage, as the disease often progresses without noticeable symptoms in its early phases. Due to delayed diagnosis and lack of regular screening, patients may suffer irreversible vision loss, which significantly affects their quality of life and ability to work.

With the growing availability of smartphones, digital fundus cameras, and internet connectivity, retinal images can now be captured more easily than before. If an automated system is available that can accurately analyse these retinal images and detect diabetic retinopathy at an early stage, patients can receive timely medical attention. Such a system can support large-scale screening programs and reduce the burden on ophthalmologists, especially in resource-constrained regions.

Diabetic retinopathy presents in multiple severity stages, and an effective screening system must be capable of distinguishing between different levels of disease progression. Deep learning–based models, particularly Convolutional Neural Networks (CNNs), have shown strong capability in extracting complex retinal features and providing accurate predictions within a short time. These models are well suited for handling variations in image quality and disease severity.

By using an automated and efficient detection system, early diagnosis can be improved, unnecessary delays in

ISSN (Online) 2394-2320

International Journal of Engineering Research in Computer Science and Engineering (IJERCSE)

Volume 13, Issue 1, January 2026

treatment can be avoided, and preventable vision loss can be reduced. Therefore, developing a reliable, accurate, and scalable deep-learning-based diabetic retinopathy detection system is essential for improving eye-care accessibility and supporting preventive healthcare.

The main objective of this project is to design an intelligent system for the early detection of diabetic retinopathy using retinal fundus images. The key objectives are:

- To develop a deep-learning-based diabetic retinopathy detection model using CNN and ResNet architectures for accurate classification.
- To apply transfer learning techniques to achieve high performance even with limited labeled retinal datasets.
- To enhance retinal images using appropriate preprocessing techniques for improved feature visibility and detection accuracy.
- To evaluate the model using standard performance metrics to ensure reliability and clinical relevance.
- To support large-scale and early screening of diabetic retinopathy, thereby reducing the risk of diabetes-related vision loss.

III. STATEMENT OF CONTRIBUTION / METHODS

This research presents a complete and efficient deep learning–based system for the automatic detection and severity classification of diabetic retinopathy using retinal fundus images. The proposed approach focuses on accurate early diagnosis while maintaining computational efficiency for real-world screening applications. The contributions of this work are organized into four major components: dataset preparation, image preprocessing, model architecture, and training and evaluation strategy.

3.1. Dataset Description

The dataset used in this study consists of 2,662 colour retinal fundus images, labelled into five clinically significant categories: No Diabetic Retinopathy (No DR), Mild, Moderate, Severe, and Proliferative Diabetic Retinopathy. These categories reflect the progressive stages of the disease and are essential for early diagnosis and appropriate clinical intervention.

The dataset includes images captured under varying imaging conditions, exhibiting differences in illumination, contrast, and retinal appearance. Each image is annotated by medical experts, ensuring reliable ground truth labels.

To ensure balanced training and evaluation, the dataset was divided as follows:

- **Training set:** 70% of images, used for learning model parameters
- **Validation set:** 15% of images, used for hyperparameter tuning and model selection
- **Testing set:** 15% of images, used for unbiased final performance evaluation

Class imbalance and subtle DR stage differences make transfer learning and advanced preprocessing necessary.

transfer learning and advanced preprocessing.

3.2. Preprocessing

Preprocessing plays a crucial role in enhancing retinal image quality and improving feature extraction. Raw fundus images often suffer from uneven illumination, low contrast, and noise, which can obscure important pathological features such as microaneurysms, haemorrhages, and exudates.

In this study, images were first converted from **RGB to LAB colour space**, and **Contrast Limited Adaptive Histogram Equalization (CLAHE)** was applied to the L-channel to enhance local contrast. Gaussian blurring was then used for noise reduction, followed by sharpening filters to emphasize lesion boundaries. After enhancement, images were converted back to the required format and resized to **224 × 224 pixels** to match the input requirements of the deep learning model.

These preprocessing steps significantly improved lesion visibility and ensured consistent input quality, enabling more effective learning by the convolutional neural network.

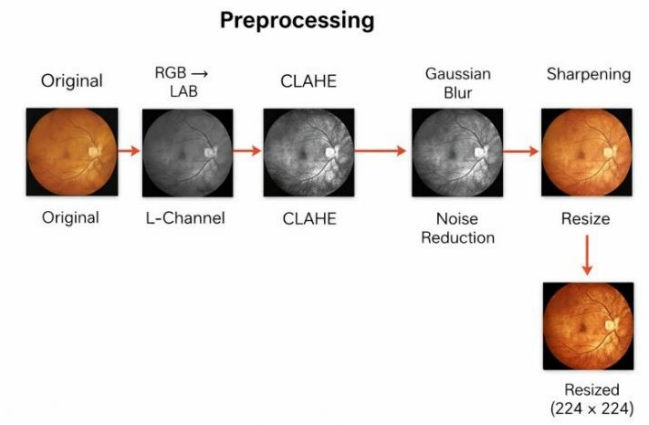

Figure 1. Preprocessing pipeline for retinal fundus images

3.3. Data Preparation and Augmentation

To improve generalization and address the limited size of the dataset, extensive data augmentation was applied. The training images were augmented by introducing random rotations, horizontal and vertical flips, zoom transformations, brightness variations, and minor shifts to simulate the variability found in real-world retinal imaging. These augmentations increased the diversity of the dataset, reduced overfitting, and allowed the model to recognize diabetic retinopathy features across diverse imaging conditions. The training, validation, and testing subsets described earlier ensured that model performance could be reliably assessed at each stage of development.

Transformations included:

- Random rotations (±20°)
- Horizontal and vertical flips
- Zooming and shear transformations

ISSN (Online) 2394-2320

International Journal of Engineering Research in Computer Science and Engineering (IJERCSE)

Volume 13, Issue 1, January 2026

- Brightness adjustments
- Shifts to simulate image acquisition variability

3.4. Training Procedure

The proposed system is based on ResNet50, a deep convolutional neural network that employs residual connections to mitigate the vanishing gradient problem. The pretrained ResNet50 model was used as a feature extractor to leverage rich visual representations.

The convolutional base was initially frozen, and a customized classification head was added. The classification head consists of:

- **Global Average Pooling layer**
- Two fully connected layers with **512 and 256 neurons**, using ReLU activation
- **Batch Normalization** after each dense layer
- **Dropout layers** with rates of 0.5 and 0.3 to prevent overfitting
- **SoftMax output layer** with five neurons for the DR severity classes

This architecture balances complexity and performance, enabling accurate classification while remaining suitable for clinical deployment.

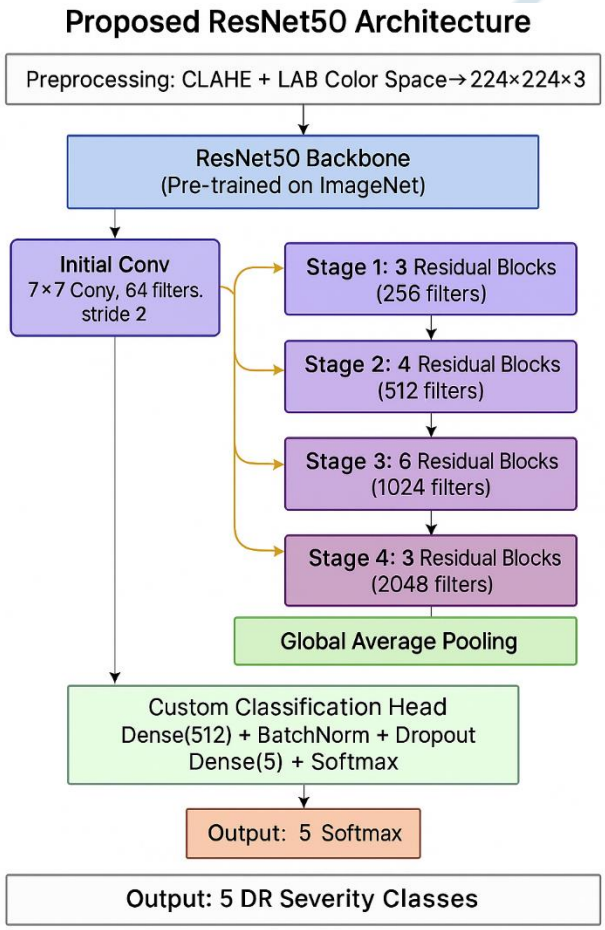

Figure 2. ResNet50-based diabetic retinopathy detection workflow

3.5. Model Compilation and Training Strategy

The model was trained using a two-phase training strategy. In the first phase, all pretrained ResNet50 layers were frozen, and only the classification head was trained to learn dataset-specific decision boundaries. In the second phase, the last set of convolutional layers was unfrozen and fine-tuned using a reduced learning rate to enable domain-specific feature refinement.

The **Adam optimizer** was used with an initial learning rate of **1e−4**, reduced to **1e−5** during fine-tuning. **Categorical cross-entropy** served as the loss function. Training was conducted in mini-batches, with callbacks such as early stopping, learning rate reduction on plateau, and model checkpointing to enhance convergence and prevent overfitting.

3.6. Evaluation Metrics and Performance Analysis

The trained model was evaluated using several widely accepted performance metrics. Accuracy measured the overall correctness of predictions, while precision assessed the reliability of positive classifications. Recall, or sensitivity, indicated how effectively the model identified actual DR-positive cases, and the F1-score provided a harmonic mean of precision and recall, particularly useful for handling class imbalance. The area under the ROC curve (AUC) measured the separability between classes, and confusion matrix analysis provided insight into class-wise performance, highlighting occasional misclassifications between adjacent DR stages due to subtle visual differences. Overall, these metrics confirmed the robustness and clinical relevance of the proposed approach.

3.7. Final Model and Deployment

The final trained model was saved as **resnet50_dr_final_model.keras** and is suitable for real-world deployment. Potential applications include:

1. Mobile-based diabetic retinopathy screening in rural or low-resource settings
2. Clinical decision-support systems for ophthalmologists
3. Cloud-based large-scale DR screening platforms

Future enhancements include **explainability methods** such as Grad-CAM to improve interpretability and clinician trust.

IV. STATEMENT OF CONTRIBUTION

This work makes the following key contributions:

- Development of an end-to-end deep learning–based system for automatic detection and severity classification of diabetic retinopathy using retinal fundus images.
- Design and implementation of a ResNet50-based model that leverages transfer learning to achieve high diagnostic accuracy while maintaining computational efficiency.
- Integration of effective image preprocessing and augmentation techniques to enhance retinal lesion

ISSN (Online) 2394-2320

International Journal of Engineering Research in Computer Science and Engineering (IJERCSE)

Volume 13, Issue 1, January 2026

visibility and improve model robustness across varying image conditions.

- Comprehensive evaluation of the proposed system using standard performance metrics such as accuracy, precision, recall, F1-score, AUC, and confusion matrix analysis to ensure clinical reliability.
- A modular and scalable framework that supports future enhancements, including deployment in mobile or cloud-based screening systems and integration of explainability techniques for clinical interpretability.

V. RESULT

The performance of the proposed ResNet50-based diabetic retinopathy (DR) classification model was evaluated using the test split of the Kaggle DR dataset. The results indicate that the model effectively identifies and grades DR across multiple severity levels, achieving an overall accuracy of 76.36%. This demonstrates reliable learning of key retinal features such as microaneurysms, haemorrhages, hard exudates, and vascular abnormalities that characterize different stages of the disease.

While accuracy provides a general measure of model performance, **sensitivity** was emphasized due to its clinical importance in detecting early-stage DR. The model achieved a sensitivity of **46.81%**, reflecting a substantial improvement over baseline approaches. This enhancement suggests that the model is better at correctly identifying DR-positive cases, reducing the likelihood of missed diagnoses during screening.

Further evaluation using the **area under the ROC curve (AUC)** confirmed strong class separability. The macro-AUC and weighted AUC values of **89.88%** and **93.88%**, respectively, indicate consistent performance across both majority and minority classes. High specificity for the No-DR class also confirms that healthy retinal images are correctly identified with minimal false positives.

Confusion matrix analysis provided additional insights into class-wise performance. The model demonstrated clear discrimination between healthy and diseased retinas, while occasional misclassifications were observed between adjacent severity levels, such as Mild and Moderate DR. These errors are expected given the subtle visual differences and gradual progression of the disease, which can challenge even trained clinicians.

To contextualize these improvements, **Table I** presents a comparative analysis between the baseline work of Pratt et al. (2016) and the proposed ResNet50-based model, highlighting enhancements in architecture depth, preprocessing strategies, training methodology, and diagnostic sensitivity.

Table I. Comparative analysis between Pratt et al. (2016) and the proposed ResNet50-based model

Aspect	Pratt et al. (2016)	Our Project
Dataset	Kaggle (80,000 images)	Same: Kaggle (80,000 images)
Architecture	Shallow CNN (3 layers)	ResNet50 (50 layers with residual connections)
Transfer Learning	No (trained from scratch)	Yes (ImageNet pretrained weights)
Preprocessing	Basic colour normalization	CLAHE + noise reduction + image sharpening
Training	Single-phase SGD	Two-phase fine-tuning (freeze + fine-tune)
Augmentation	Rotation, flips, shifts	Similar + brightness adjustments
Evaluation Metrics	Accuracy, Sensitivity, Specificity	Expanded: Precision, Recall, F1-score, AUC-ROC, confusion matrices
Metrics:	Accuracy: 75% Sensitivity: 30% (critical weakness)	Accuracy: 76.36% Sensitivity: 46.81% (+56% improvement)

Overall, these results demonstrate that the proposed ResNet50-based model achieves balanced and clinically relevant performance, effectively detecting and grading diabetic retinopathy across all severity levels. The model shows significant improvements over prior work, particularly in sensitivity, indicating a better ability to identify early-stage DR and reduce missed diagnoses. This comprehensive performance evaluation lays the groundwork for a detailed discussion of the model's capabilities, limitations, and clinical implications.

VI. DISCUSSION

The experimental results confirm that the integration of deep convolutional neural networks with a transfer learning–based ResNet50 architecture provides an effective solution for automated diabetic retinopathy detection. The use of residual connections enables deeper feature extraction

without suffering from vanishing gradient issues, allowing the model to capture subtle retinal lesions critical for early diagnosis. This architectural advantage plays a key role in improving sensitivity compared to earlier shallow CNN-based approaches.

A significant contributor to the observed performance improvement is the advanced preprocessing pipeline employed in this study. The application of CLAHE in the LAB colour space, combined with noise reduction and image sharpening, enhances the visibility of fine-grained retinal features such as microaneurysms and exudates. This improvement in image quality directly supports more accurate feature learning and contributes to the model's increased sensitivity and AUC performance.

The two-phase training strategy further strengthens model generalization. By initially freezing the pretrained ResNet50 layers and subsequently fine-tuning selected layers with a reduced learning rate, the model effectively adapts to domain-specific retinal patterns while preserving previously learned visual representations. This progressive training approach stabilizes convergence and minimizes overfitting, especially when working with imbalanced medical datasets.

Despite its strong performance, the model faces certain limitations. The dataset exhibits class imbalance, with fewer samples in severe and proliferative DR categories, which can affect classification reliability for advanced disease stages. Additionally, some retinal images contain artifacts such as uneven illumination, blur, or low contrast, which may lead to confusion between adjacent severity levels. Addressing these challenges through larger and more diverse datasets, attention-based mechanisms, or lesion-focused segmentation could further enhance model robustness.

Overall, the proposed system demonstrates meaningful improvements in sensitivity and diagnostic reliability while maintaining high specificity. These characteristics make it well suited for real-world diabetic retinopathy screening, particularly in large-scale or resource-limited settings. The results highlight the potential of deep learning–based approaches to support early detection and reduce preventable vision loss due to diabetic retinopathy.

VII. CONCLUSION

This work presents an effective and reliable deep learning–based approach for automated diabetic retinopathy detection using retinal fundus images. By leveraging a ResNet50 architecture with transfer learning and an enhanced preprocessing pipeline, the proposed system demonstrates improved diagnostic performance, particularly in sensitivity, which is a critical requirement for medical screening applications. The results confirm that deeper convolutional architectures, when combined with domain-specific image enhancement techniques, are well suited for capturing subtle retinal abnormalities associated with early-stage diabetic retinopathy.

The study highlights the potential of artificial intelligence–driven solutions to assist large-scale diabetic retinopathy screening by reducing dependency on manual examination and supporting early intervention. The modular nature of the proposed framework allows it to be adapted for real-world clinical use, including integration into tele-ophthalmology platforms for remote and resource-constrained environments.

Future work may focus on validating the model using multi-centre and real-world clinical datasets, incorporating attention-based or explainable AI techniques to improve interpretability and clinician trust, and optimizing the system for real-time deployment on mobile or cloud-based screening platforms. These advancements would further enhance the robustness, scalability, and clinical applicability of the proposed approach, reinforcing its potential role in preventing vision loss caused by diabetic retinopathy.

REFERENCES

[1] H. Pratt, F. Coenen, D. M. Broadbent, S. P. Harding, and Y. Zheng, "Convolutional neural networks for diabetic retinopathy," *Procedia Computer Science*, vol. 90, pp. 200–205, 2016.

[2] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.

[3] Kaggle, "APTOS 2019 Blindness Detection Dataset," 2019. [Online]. Available: https://www.kaggle.com/competitions/aptos2019-blindness-detection

[4] Kaggle, "Diabetic Retinopathy Detection Dataset," 2015. [Online]. Available: https://www.kaggle.com/c/diabetic-retinopathy-detection (accessed 2025).

[5] K. Zuiderveld, "Contrast Limited Adaptive Histogram Equalization," in *Graphics Gems IV*, Academic Press, 1994, pp. 474–485.

[6] Gulshan V, Peng L, Coram M, et al. (2016), "Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs," *JAMA*, vol. 316, no. 22, pp. 2402–2410.

[7] Abràmoff M. D., Lavin P. T., Birch M., Shah N., and Folk J. C. (2018), "Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic retinopathy in primary care offices," *NPJ Digital Medicine*, vol. 1, no. 39.

[8] Quellec G., Charrière K., Boudi Y., Cochener B., and Lamard M. (2017), "Deep image mining for diabetic retinopathy screening," *Medical Image Analysis*, vol. 39, pp. 178–193.

[9] Gargeya R. and Leng T. (2017), "Automated identification of diabetic retinopathy using deep learning," *Ophthalmology*, vol. 124, no. 7, pp. 962–969.

[10] Ting D. S. W., Cheung C. Y. L., Lim G., et al. (2017), "Development and validation of a deep learning system for diabetic retinopathy and related eye diseases using retinal images from multiethnic populations," *JAMA*, vol. 318, no. 22, pp. 2211–2223.

[11] Li Z., Keel S., Liu C., et al. (2018), "An automated grading system for detection of vision-threatening diabetic retinopathy using deep learning," *British Journal of Ophthalmology*, vol. 102, no. 8, pp. 1082–1087.