International Journal of Web of Multidisciplinary Studies
(Peer-Reviewed, Open Access, Fully Refereed International Journal)
website: *https://ijwos.com*
Volume 03, Issue 10,October 2026
OPEN ACCESS
E-ISSN : 3049-2424
DOI: https://doi.org/10.71366/ijwos

AI-Based Smart Diabetic Retinopathy Detection System with Lesion Highlighting and Risk Prediction

Prof. Shradha H[1], Akkamahadevi Pawar[2], Akshata Hanamapure[3], Chetan Mathapathi[4], Harish Meti[5]
[1,2,3,4,5] *Department of Artificial Intelligence and Data Science, Angadi Institute of Technology and Management, Belagavi, Karnataka, India*

Article Info

Article History:
Published: 29 Oct 2026

Publication Issue:
Volume 03, Issue 10
October 2026

Page Number:
505-512

Corresponding Author:
Prof. Shradha H

Abstract:
Diabetic retinopathy is one of the main causes of preventable blindness in working-age adults. In its early stages it usually causes no symptoms, so by the time a patient notices a problem the retina may already be permanently damaged. Screening today depends on an ophthalmologist examining fundus photographs by hand. That approach is accurate, but it is slow and expensive, and it is often out of reach in rural and under-resourced areas. In this paper we describe a deep-learning system that sorts retinal fundus images into four classes: normal, cataract, diabetic retinopathy, and glaucoma. It uses an EfficientNetB0 backbone adapted through transfer learning. Alongside each prediction, the system applies Gradient-weighted Class Activation Mapping (Grad-CAM) to show which parts of the retina influenced the decision, and it highlights lesions such as microaneurysms, hemorrhages, and exudates. We built the pipeline in TensorFlow/Keras and deployed it through a Streamlit web interface that returns the predicted class, a confidence score, a Grad-CAM heatmap, and a PDF report the user can download. Results indicate that transfer learning combined with targeted fine-tuning gives strong classification performance while staying light enough to run in a browser. The system shows how explainable AI can connect automated screening with decision support that clinicians can actually check.

Keywords: Diabetic Retinopathy, Retinal Fundus Image, Convolutional Neural Network, Deep Learning, EfficientNetB0, Transfer Learning, Lesion Detection, Grad-CAM, Risk Prediction, Medical Image Analysis.

1. Introduction

Diabetes is becoming more common worldwide, and diabetic retinopathy is among its most serious complications. The disease slowly damages the blood vessels in the retina and, if left untreated, can lead to permanent loss of vision. Because it rarely causes symptoms early on, regular retinal screening is the main way to find it before the damage becomes irreversible. Detecting it early, and making that detection easy to reach, matters as much as the treatment itself.

Screening still depends largely on trained ophthalmologists looking at fundus photographs for early signs such as microaneurysms, hemorrhages, and exudates. This creates a bottleneck. There are far fewer specialists than there are diabetic patients who need yearly screening, and the shortage is worse in rural and under-resourced regions, where an ophthalmologist may be hard to reach or simply not available.

Deep learning offers a practical way to narrow this gap. Convolutional neural networks (CNNs) suit image-based diagnosis well because they learn useful visual features straight from data instead of following hand-written rules. In this work we bring together image preprocessing, CNN-based classification, lesion-aware visualization, and a simple web interface in one end-to-end screening pipeline.

At a high level, the system works as follows:

Fundus Image → Preprocessing →

Copyright © 2026 The Author(s) : This is an open access article under the CC BY license (http://creativecommons.org/licenses/by-nc/4.0/)

EfficientNetB0 → Classification →

Grad-CAM → Streamlit → PDF Report

We designed the system as a multi-class classifier covering four categories: normal, cataract, diabetic retinopathy, and glaucoma. A screening tool that meets real fundus images has to tell diabetic retinopathy apart from other common retinal conditions, so treating detection as a simple yes-or-no question would not be realistic.

2. RELATED WORK

A. Deep Learning for Retinal Disease Detection

Deep learning, and CNNs in particular, has become the leading approach to medical image analysis in recent years. Older image-processing pipelines depend on hand-built filters and thresholds. CNNs instead learn layers of visual features from labeled data, which lets them pick up subtle patterns, such as the fine vascular changes of early retinopathy, that are hard to describe by hand.

CNNs have worked well for recognizing the visual signs of different abnormalities in fundus images. Training a deep network from scratch normally needs a large labeled dataset and a lot of computing power, so we used transfer learning instead. A network pretrained on a large general-purpose image dataset is adapted to the retinal task, so it keeps the low-level features it already knows and learns the domain-specific ones on top.

B. Diabetic Retinopathy Detection

Most earlier diabetic retinopathy detection systems perform binary or coarse classification. They label an image as healthy or diseased, or they grade severity along a single scale. Few combine classification with lesion-level visualization, risk estimation, and patient monitoring in one tool. That gap motivated the design described here.

C. Explainable AI

A common weakness of deep-learning diagnosis is that a prediction probability says nothing about why the model made its choice. In medicine this is a real problem, because clinicians need to be sure the model is looking at genuine signs of disease and not at incidental artifacts. Grad-CAM helps by producing a heatmap of the image regions that contributed most to a given prediction.

3. SYSTEM ARCHITECTURE AND METHODOLOGY

The proposed system has the following main stages:

- Image acquisition
- Image preprocessing
- Deep-learning feature extraction
- Disease classification
- Lesion highlighting
- Risk/severity prediction
- Explainability using Grad-CAM
- Result visualization
- PDF report generation

The architecture uses EfficientNetB0 as the convolutional backbone, followed by a custom classification head built for the four-class retinal task. Figs. 1 and 2 show the overall pipeline, from the raw fundus image through feature extraction and feature selection to the final multiclass classification of disease severity.

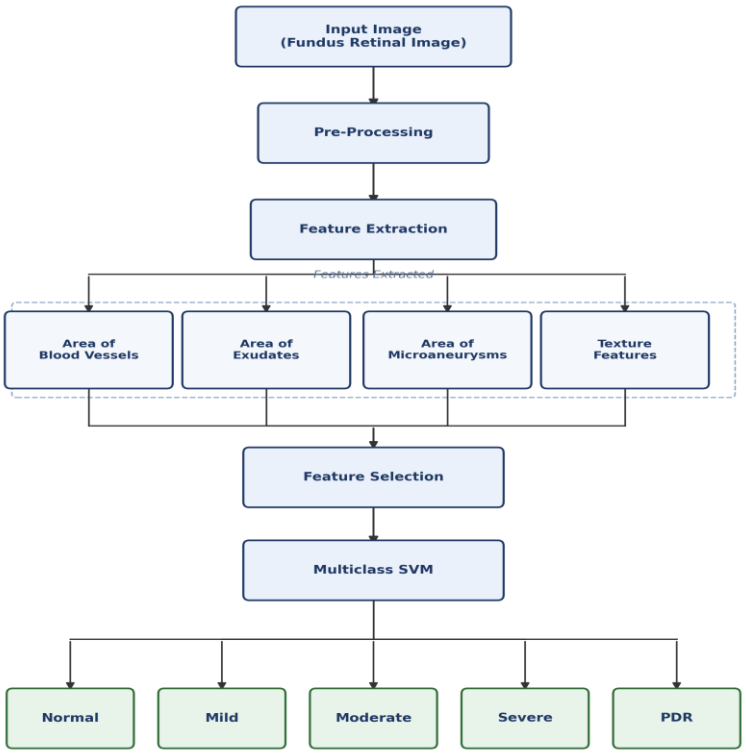


Fig. 1. System architecture: retinal feature extraction and multiclass severity classification pipeline.

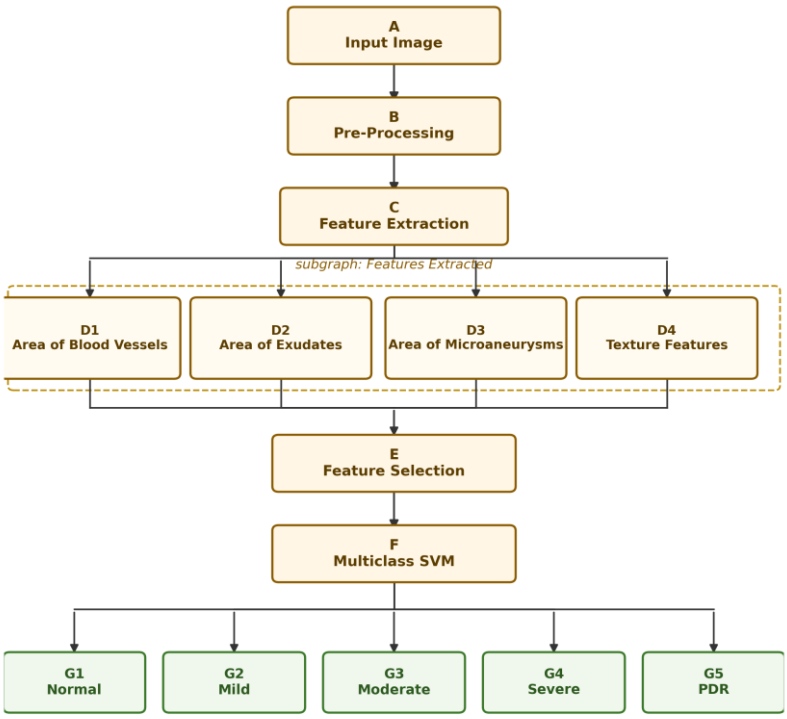


Fig. 2. Feature-extraction and multiclass SVM classification workflow, grouped by extracted feature type (nodes labeled per processing stage).

The pipeline starts when the user uploads a retinal fundus image through the application interface. Nothing beyond an ordinary fundus photograph is needed, and no special equipment or format conversion is involved. The image then goes to the preprocessing stage.

B. Image Preprocessing

Before inference, every image is resized and normalized to the input format that the EfficientNet backbone expects. This keeps scale and pixel values consistent across images taken under different lighting and with different cameras. The model receives an image tensor of shape 224 × 224 × 3, where:

- 224 = image height
- 224 = image width
- 3 = RGB colour channels

This step matters because fundus cameras differ a lot in resolution and color balance. Without consistent preprocessing, the network would waste effort learning to work around differences that have no diagnostic meaning.

C. Feature Extraction and Classification

Features are extracted by EfficientNetB0, initialized with ImageNet-pretrained weights. ImageNet contains everyday photographs rather than medical scans, but the low- and mid-level features it teaches, such as edges, textures, and color gradients, carry over well to fundus images. This gives the network a strong starting point before fine-tuning on retinal data. The custom head that follows the backbone consists of global average pooling, batch normalization, two dense layers of 256 and 128 units (ReLU activation, with dropout of 0.4 and 0.3 respectively), and a final four-unit softmax output layer.

D. Disease Prediction

The model separates four retinal conditions:

- Normal
- Cataract
- Diabetic Retinopathy
- Glaucoma

A softmax activation turns the network's output into class probabilities. The class with the highest probability is reported as the final prediction, together with its confidence score.

E. Lesion Highlighting

Besides the top-level classification, the system points out specific retinal abnormalities linked to diabetic retinopathy, including:

- Microaneurysms
- Hemorrhages
- Exudates

Showing these lesion types gives the user a clearer idea of what the model is responding to, instead of a single unexplained label.

F. Grad-CAM Explainability

To make the model's reasoning visible, we apply Grad-CAM to the final convolutional layer, which produces a class-discriminative heatmap of the regions that most influenced the predicted class. The method finds the target convolutional layer, computes gradients with respect to the predicted class, and overlays the resulting activation map on the original fundus image.

G. User Interface and Report

Results are presented through a Streamlit-based web interface. It shows the predicted class, the confidence score, and the matching Grad-CAM visualization, and it creates a downloadable PDF report that summarizes the findings for record keeping or a follow-up consultation.

4. TECHNOLOGY JUSTIFICATION

A. Convolutional Neural Networks

CNNs are the natural choice for image classification because their layered structure learns spatial features step by step, from simple edges and textures in the early layers to complex, disease-relevant shapes in the deeper ones. No manual feature engineering is required.

B. EfficientNetB0

We chose EfficientNetB0 as the backbone because it balances accuracy and computational cost well. Its compound scaling method adjusts network depth, width, and input resolution together, so it reaches strong accuracy without the large parameter count of bigger architectures. This is important for a system meant to run in a lightweight, browser-accessible interface.

C. Transfer Learning

Training from random initialization would need a very large retinal dataset. Instead, the model starts from ImageNet-pretrained weights, which greatly reduces both training time and the amount of data required.

D. Grad-CAM and Streamlit

We picked Grad-CAM over simpler saliency methods because it produces class-discriminative visualizations from the model's existing convolutional layers, with no change to the architecture. Streamlit was used for the interface because it lets us build a fully interactive web application directly in Python, so the whole stack stays in one language.

5. IMPLEMENTATION DETAILS

A. End-to-End Pipeline

The complete system runs in this order: preprocessing, EfficientNetB0 feature extraction, dense classification layers, softmax prediction, Grad-CAM visualization, and PDF report generation. Each stage is modular, so any component can be updated independently as the system develops.

B. Training

Training follows the two-phase strategy commonly used with pretrained backbones. In Phase 1 (transfer learning), the pretrained EfficientNetB0 base is frozen and used as a feature extractor, and only the custom classification head is trained. In Phase 2 (fine-tuning), selected backbone layers are unfrozen and trained at a lower learning rate. This lets the network adapt to retinal images without overwriting the general representations it learned from ImageNet. Both phases use categorical cross-entropy loss and the Adam optimizer.

C. Evaluation

Model performance is assessed with a standard set of classification metrics:

- Accuracy
- Precision
- Recall
- F1-score
- Confusion matrix
- ROC/AUC
- Per-class metrics

Together, these metrics describe model behavior more fully than accuracy alone. That matters in medical screening, where missing a diseased patient (a false negative) costs far more than a false alarm.

6. SYSTEM MODULES

The application is divided into loosely coupled modules, each handling one part of the screening workflow:

Module	Purpose
User Module	Handles user registration and login
Image Upload Module	Uploads retinal images
Preprocessing Module	Improves/prepares image quality
CNN Detection Module	Detects retinal disease
Lesion Highlighting Module	Identifies and highlights retinal lesions
Risk Prediction Module	Predicts disease severity
Report Module	Generates reports
Database Module	Stores patient and result information

Separating the modules this way keeps the classification logic independent of the interface layer, which makes each component easier to test, replace, or extend.

7. RESULTS AND DISCUSSION

For every uploaded fundus image, the system returns a predicted disease class and its probability, so the user sees how confident the model is and not just a bare label.

The Grad-CAM visualization generally highlights regions that fit the predicted class. For diabetic retinopathy cases, these are typically areas with microaneurysms, hemorrhages, or exudates. This gives qualitative support to the idea that the model looks at clinically relevant features and not at misleading background patterns.

As with any deep-learning classifier, reported accuracy depends on the dataset and the train-test split. It should be read together with the confusion matrix and per-class metrics, not as a single summary number.

8. COMPARISON WITH EXISTING APPROACHES

Existing System	Proposed System
Mainly manual examination	AI-assisted automated analysis
Requires specialist involvement	User-friendly automated screening
Time-consuming	Faster image analysis
Limited lesion visualization	Lesion highlighting
Limited risk/severity prediction	Risk/severity prediction

Difficult for remote screening	Supports remote accessibility
Basic image classification	Classification + explanation + report
Limited monitoring	Supports patient/report monitoring

Compared with conventional manual screening and with narrower automated classifiers, the main contribution of the proposed system is that it brings classification, lesion-level interpretability, and reporting together in one accessible workflow.

9. LIMITATIONS

Although the results are promising, the system has several limitations that we want to state openly:

- Poor-quality retinal images can reduce prediction accuracy.
- Deep-learning models depend on suitable training data.
- Incorrect predictions can occur.
- Patient data must be protected from unauthorized access.
- The system needs further validation before any clinical use.
- It should not be treated as a replacement for an ophthalmologist.

For these reasons, the system is best seen as a research-oriented screening aid that supports clinical judgment and does not replace it. Any positive finding should be followed up with a qualified ophthalmologist.

10. CONCLUSION AND FUTURE SCOPE

A. Conclusion

This work presents a practical, end-to-end approach to automated diabetic retinopathy screening that pairs CNN-based classification with visual explanation. By fine-tuning an EfficientNetB0 backbone through transfer learning, the system classifies retinal images into multiple classes reliably while remaining efficient enough for real-time, browser-based use. Adding Grad-CAM visualization and lesion highlighting takes it beyond a black-box classifier, since its predictions can be inspected and questioned.

B. Future Scope

Several directions could extend this work:

- Using larger and more diverse retinal datasets.
- Trying more advanced deep-learning architectures to improve accuracy.
- Adding finer-grained severity grading for diabetic retinopathy.
- Improving lesion segmentation and visualization.
- Moving to cloud-based deployment.
- Developing a mobile application.
- Adding real-time doctor consultation.
- Supporting analysis of multiple retinal diseases.
- Improving remote healthcare monitoring.

Following these directions would bring the system closer to a deployable clinical screening aid, while keeping transparency and interpretability as priorities alongside predictive accuracy.

Acknowledgment

The authors sincerely thank Prof. Shradha H for her guidance, encouragement, and supervision throughout this project. We also thank Angadi Institute of Technology and Management, Belagavi, and the Department of Artificial Intelligence and Data Science for the facilities and support they provided.

References

[1] N. Suedumrong et al., "Diabetic Retinopathy Detection Using Convolutional Neural Networks with Background Removal and Data Augmentation," Applied Sciences, vol. 14, no. 19, 2024.

[2] R. Abbasi et al., "Diabetic Retinopathy Detection Using Adaptive Deep Convolutional Neural Networks," Scientific Reports, 2025.

[3] M. Hussain et al., "Enhanced CNN based P-EDR for Diabetic Retinopathy Diagnosis," Engineering, Technology & Applied Science Research, vol. 15, no. 1, pp. 19062–19067, 2025.

[4] K. S. Argade and M. Madhura, "Automatic Detection of Diabetic Retinopathy Using Image Processing and Data Mining Techniques," in Proc. Int. Conf. Electronics, Electrical Eng. and Computing (INTERCON), Springer, 2018, pp. 153–162.

[5] A. S. Gautam and S. K. Jana, "Automated Diagnosis of Diabetic Retinopathy Using Image Processing for Non-Invasive Biomedical Application," in Proc. IEEE EMBS, 2019, pp. 5951–5954.

[6] N. Karami and H. Rabbani, "A Dictionary Learning Based Method for Detection of Diabetic Retinopathy in Color Fundus Images," in Proc. Int. Conf. Bioinformatics and Biomedical Technology, 2020, pp. 134–138.

[7] N. Chakrabarty and S. Chatterjee, "An Offbeat Technique for Diabetic Retinopathy Detection Using Computer Vision," in Proc. CBMS, 2019, pp. 1–4.

[8] D. Lokuarachchi and L. Muthumal, "Detection of Red Lesions in Retinal Images Using Image Processing and Machine Learning Techniques," Graphics Gems IV, Academic Press, 2019, pp. 474–485.