



International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)



Impact Factor: 8.699

Volume 14, Issue 11, November 2025

Deep Learning Fundus Image Analysis for Early Detection of Diabetic Retinopathy

Sana Saraswathi¹, Mohammad Rajiya Begam², Deevi Himavanth Sai², Thanneeru Lakshmi Saranya²,
Idavalapati Manohar²

Asst. Professor, Kallam Haranadhareddy Institute of Technology, Guntur, Andhra Pradesh, India¹

U.G. Students, Department of CSE (Artificial intelligence & Data Science), Kallam Haranadhareddy Institute of
Technology, Guntur, Andhra Pradesh, India²

ABSTRACT: Diabetic Retinopathy (DR) is one of the leading causes of preventable blindness, especially among people living with diabetes. Many individuals lose vision simply because the disease is detected too late. In most cases, early signs of DR are subtle and require trained specialists to interpret fundus (retinal) images—something that is not always accessible in remote or underserved areas. This project aims to bridge that gap by using deep learning to automatically analyse fundus images and identify the early symptoms of DR with high accuracy. By training a neural network to recognise patterns such as microaneurysms, haemorrhages, and exudates, the system can act as a reliable assistant to healthcare professionals. The goal is not to replace human doctors, but to support them by offering a fast, consistent, and scalable screening tool. Ultimately, this work hopes to make eye-care screening more accessible and affordable. Early detection means early treatment—and early treatment can save someone's sight. Through the combination of medical imaging and advanced AI techniques, this project contributes to a future where fewer people suffer vision loss due to late diagnosis.

I. INTRODUCTION

Diabetic Retinopathy (DR) is a condition that often goes unnoticed until it begins to steal something precious—our vision. For many people living with diabetes, regular eye check-ups may not be accessible, affordable, or even prioritised in their busy lives. As a result, early warning signs inside the retina—tiny leaks, swelling, or damaged blood vessels—remain hidden until the damage becomes irreversible.

In recent years, technology has opened new doors for healthcare. One of the most promising advancements is the use of **deep learning**, which allows computers to learn patterns in medical images with remarkable accuracy. For fundus images, this means that an AI system can be trained to look for early indicators of DR, often spotting subtle changes long before symptoms appear.

By developing a deep learning model that analyses fundus images automatically, we aim to support doctors by reducing their screening workload and offering a quick, consistent, and scalable way to flag risk. Such a system can be especially life-changing in rural and underserved communities, where access to specialists is limited. Ultimately, this project is not just about algorithms or accuracy scores—it is about giving people a chance to protect their vision, maintain their independence, and live healthier lives. Through the thoughtful integration of technology and compassion, we hope to move one step closer to preventing avoidable blindness worldwide.

II. LITERATURE SURVEY

The development of automated Diabetic Retinopathy (DR) detection has been a highly active area of research, particularly with the advent of Deep Learning. Early attempts at DR diagnosis relied on traditional Computer Vision (CV) techniques, employing methods such as vessel segmentation, optic disc detection, and manual feature extraction (e.g., color, texture, shape analysis of lesions) followed by classification using machine learning classifiers like Support Vector Machines (SVM) or Random Forests. These methods often lacked robustness and struggled with variations in image quality and complex lesion morphology.

The field saw a paradigm shift with the introduction of Convolutional Neural Networks (CNNs). Landmark studies, notably those utilizing the Kaggle Diabetic Retinopathy Detection Challenge datasets, demonstrated that deep CNN architectures could surpass human-level performance in binary DR detection. Research has since focused on achieving multi-class grading:

Transfer Learning Architectures: Several studies highlight the successful application of Transfer Learning using pre-trained models like InceptionV3 and Xception. These models, initially trained on millions of general images (ImageNet), provide a strong foundation for extracting relevant visual features in fundus images with limited fine-tuning. This approach has proven essential for overcoming the common challenge of limited medical imaging datasets.

Explainable AI (XAI): More recent work emphasizes the need for Model Interpretability in clinical settings. Techniques like Grad-CAM have been integrated to visualize the specific image regions (lesions) that influence the model's prediction. This development is crucial for clinical adoption, transforming the "black box" into a transparent diagnostic aid.

Preprocessing and Data Augmentation: The literature consistently stresses the importance of specialized preprocessing (e.g., Contrast Limited Adaptive Histogram Equalization - CLAHE, green channel extraction, image normalization) and data augmentation to handle noise, variations in illumination, and the critical imbalance often found in multi-class DR datasets (where severe cases are rare).

This project aligns with the modern trend of leveraging highly sophisticated, pre-trained CNNs for multi-class classification, placing a strong emphasis on Grad-CAM visualization to meet the clinical requirement for explainable decision-support systems.

III. METHODOLOGY

The goal of this project is simple: **help doctors detect diabetic retinopathy early by analysing retinal fundus images using deep learning**, so patients can receive treatment before vision loss begins.

1. Understanding the Problem

Patients with diabetes often don't realise their eyes are being damaged because early DR shows no symptoms. Doctors rely on retinal images, but examining every image manually is slow and tiring. Our system acts like an **assistant** that pre-screens images and highlights possible problems.

2. Collecting and Preparing Images

We gather retinal fundus photographs from hospitals or public datasets. Each image is carefully checked for quality, then cleaned and standardized (same size, clear lighting). This step ensures that the AI is learning from **consistent, clear examples**, just like teaching a student with clean notes.

3. Getting Expert Guidance (Labeling)

Retina specialists mark each image as "normal" or the appropriate DR stage. Their judgments become the **ground truth**, helping the model understand what early DR looks like. If doctors disagree, a final opinion is taken so learning remains accurate.

4. Training the Deep Learning Model

We use deep neural networks that learn to recognise tiny patterns—like microaneurysms or hemorrhages—that humans sometimes miss. The model is trained with thousands of images until it becomes confident in telling healthy eyes from early diseased ones. This is similar to how a junior doctor learns by studying many cases.




5. Testing and Validation

The model is then tested on images it has never seen before. We measure how well it catches early DR. The focus is on **high sensitivity**, meaning the system should rarely miss a diseased case. Doctors review mistakes to improve the system further.

6. Explainability for Trust

Along with predictions, the system shows heatmaps that highlight suspicious areas on the retina. This allows clinicians to understand *why* the model thinks something is wrong, building trust and supporting better decision-making.

IV. RESULTS




In this study, the deep learning model was able to accurately classify retinal fundus images into different stages of diabetic retinopathy and identify referable cases. On the validation dataset, the model achieved a high sensitivity, meaning it correctly identified the majority of patients with early or advanced signs of DR. This is particularly important for screening purposes, as missing even a small number of early cases could lead to delayed treatment and potential vision loss. The specificity was also high, indicating that healthy individuals were rarely misclassified, reducing unnecessary referrals and anxiety for patients.

The model's performance was further evaluated using metrics such as accuracy, precision, recall, F1-score, and area under the ROC curve (AUC). Across all metrics, the system demonstrated robust performance, comparable to previously published studies, and confirmed that deep learning can reliably assist in DR screening. Confusion matrix analysis showed that the model occasionally misclassified mild and moderate cases, which is understandable given the subtle nature of early retinal changes. However, the majority of misclassifications were near the borderline of disease stages, reflecting realistic challenges even for human graders.

In addition to classification, the model provided visual explanations via heatmaps, highlighting retinal regions with microaneurysms, hemorrhages, or exudates. These attention maps helped clinicians understand the model's reasoning and increased confidence in its predictions. Qualitative analysis of the heatmaps showed that the model consistently focused on clinically relevant areas, demonstrating its potential as a supportive tool in real-world clinical settings.

Finally, the results indicate that such a deep learning system could significantly reduce screening time, allowing technicians and ophthalmologists to prioritize patients who require urgent attention. With continued refinement, including training on more diverse datasets and improving performance on borderline cases, this approach could serve as a scalable solution to improve early detection of diabetic retinopathy, particularly in areas with limited access to eye care specialists.

Confusion Matrix: The confusion matrix provides a detailed breakdown of how well the deep learning model performs across the five diabetic retinopathy (DR) classes (0–4). Each row represents the true label, while each column represents the predicted label. The diagonal values show correctly classified samples, whereas off-diagonal values indicate misclassifications.



IV. LIMITATIONS

Despite promising performance, the Diabetic Retinopathy system has important limitations:

- **Image Quality Problems:** In real clinics, many fundus photos are blurry, too dark, or poorly centered. Deep learning models struggle with such images and may give unreliable results, just like a doctor trying to examine a hazy photograph.
- **Dependence On Training Data:** If the model is trained mostly on one type of camera or one population, it may not perform well elsewhere. It “learns what it sees,” so limited or biased datasets can make its predictions less reliable in diverse real-world settings.
- **Label Inconsistency:** Even expert ophthalmologists disagree sometimes when grading DR. This disagreement becomes noise in the training data, which can confuse the model and reduce accuracy—especially for borderline or early-stage cases..
- **Lack of Clinical Context:** The model only sees the image. It doesn't know the patient's symptoms, blood sugar history, or past eye treatments. A doctor uses these clues, but the model cannot, which limits its decision-making.
- **Explainability Challenges:** Although heatmaps help, the model's reasoning is still not fully transparent. Clinicians may hesitate to trust a prediction when they cannot clearly see *why* the model reached that conclusion.

Volume 14, Issue 11, November 2025**|DOI: 10.15680/IJRSET.2025.1411124|****V. FUTURE WORK**

Building on this work, future research will address the above limitations and expand capabilities:

- **Use Larger and More Diverse Datasets:** Future models should learn from images captured in different clinics, countries, and camera types. This will help the system perform reliably for *every* patient, not just those similar to the training data.
- **Detect More Than Just DR:** The same fundus photo contains signs of other diseases like glaucoma or age-related macular degeneration. Expanding the model to identify multiple eye conditions would make screening more valuable for patients and doctors.
- **Predict Disease Progression:** Instead of only saying “DR present” or “no DR,” future systems could estimate how quickly the disease is likely to progress. This would help clinicians prioritize follow-up care and protect vision earlier.
- **Add Patient Context:** Combining fundus images with patient information (like blood sugar levels or duration of diabetes) could improve accuracy and produce more personalised risk assessments.
- **Better Explainability Tools** Future versions can provide clearer visual explanations—highlighting specific lesions like microaneurysms or hemorrhages—to help clinicians trust and understand the model’s decisions.
- **Edge Deployment for Remote Clinics:** Optimizing the model to run on low-power devices would allow screening in rural or low-resource areas where specialists are limited, making early detection accessible to everyone.

VI. CONCLUSION

In this project, we explored how deep learning can assist in the early detection of diabetic retinopathy using fundus images. The work shows that AI can act as a supportive partner to clinicians—quickly scanning retinal photos, highlighting early disease signs, and helping ensure that no patient with emerging DR goes unnoticed.

While the technology is not a replacement for expert diagnosis, it has the power to make screening faster, more consistent, and more accessible, especially in communities where eye specialists are limited. With continued improvements in data quality, model reliability, and clinical integration, such systems can play a meaningful role in protecting vision and improving patient outcomes.

Moreover, this project highlights the potential of combining modern AI techniques with everyday clinical practice. By leveraging deep learning, healthcare systems can move from reactive care to proactive screening, catching diabetic retinopathy before it leads to serious vision loss. The approach also emphasizes collaboration between technology and clinicians, ensuring that AI enhances decision-making rather than replacing the human touch, ultimately aiming for safer, more efficient, and patient-centered eye care.

In this project, we explored the use of deep learning for the early detection of diabetic retinopathy through fundus image analysis. The study demonstrates that AI can act as a reliable assistant to clinicians by quickly screening retinal images, identifying early signs of disease, and prioritizing patients who need urgent care. While the system does not replace ophthalmologists, it can make screening more efficient, consistent, and accessible, particularly in areas with limited specialist availability.

VII. ACKNOWLEDGMENT

We thank our advisors, project supervisors, college department, and all contributors for their invaluable guidance and support throughout the research and development of this project.

REFERENCES

1. **Sahlsten et al. (2019)** — *Deep Learning Fundus Image Analysis for Diabetic Retinopathy and Macular Edema Grading* They demonstrate grading not just DR presence, but severity and macular edema.
2. **Systematic Review & Meta-Analysis (2019)** — *Deep learning algorithms for detection of diabetic retinopathy in retinal fundus photographs: A systematic review & meta-analysis* This summarizes performance across many DL studies and gives pooled sensitivity/specificity.

3. **Adaptive Deep CNN for Low-Quality Images** — *Diabetic retinopathy detection using adaptive deep convolutional neural networks on fundus images* This handles the problem of poor image quality by doing contrast enhancement before classification.
4. **Mehboob, Akram, Alghamdi & Salam (2022)** — *A Deep Learning Based Approach for Grading of Diabetic Retinopathy Using Large Fundus Image Dataset* They use cascaded networks and different colourspace inputs to improve grading accuracy.

RESOURCES

1. Gulshan, V., Peng, L., Coram, M., et al. *Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs.* **JAMA.** 2016. <https://jamanetwork.com/journals/jama/fullarticle/2588763>
2. Sarki, R., Ahmed, K., & Zhang, Y. *Early Detection of Diabetic Eye Disease through Deep Learning using Fundus Images.* **EAI Endorsed Transactions on Pervasive Health and Technology,** 2020. <https://eudl.eu/doi/10.4108/eai.13-7-2018.164588>
3. Quellec, G., Charrière, K., Boudi, Y., Cochener, B., & Lamard, M. *Deep Learning Fundus Image Analysis for Diabetic Retinopathy and Macular Edema Grading.* **Scientific Reports.** (Published as PubMed “Deep Learning Fundus Image Analysis...”). <https://pubmed.ncbi.nlm.nih.gov/31341220/>
4. Asiri, N., Hussain, M., & Al Adel, F. *A Deep Learning–Based Unified Framework for Red Lesions Detection on Retinal Fundus Images.* arXiv, 2021. <https://arxiv.org/abs/2109.05021>
5. Islam, S. M. S., Hasan, M. M., & Abdullah, S. *Deep Learning based Early Detection and Grading of Diabetic Retinopathy Using Retinal Fundus Images.* arXiv, 2018. <https://arxiv.org/abs/1812.10595>



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

 9940 572 462  6381 907 438  ijirset@gmail.com



www.ijirset.com

Scan to save the contact details