

AI-Powered Early Detection of Diabetic Retinopathy: A Deep Learning Approach for Improved Clinical Decision-Making

Deepali Virmani¹, Komal B. Umare², Dr. A. Devendran³, Dr. G. Ravikanth⁴, Dinesh Jamthe⁵, Mona Kejariwal⁶

¹Designation: Professor Department: Department of IT/ CSE-DS Institute: Guru Tegh Bahadur Institute of Technology District: Delhi City: Delhi State: Delhi Email - deepalivirmani@gmail.com

²Assistant Professor Department of Artificial Intelligence G H Raison College of Engineering Nagpur Nagpur Maharashtra komal29umare@gmail.com

³Sri Ramachandra Institute of Higher Education and Research (SRIHER) devendran.alagarsamy@gmail.com

⁴Associate Professor Professor Computer Science & Engineering KLEF Deemed to be University Guntur Vaddeswaram, Guntur Andhra Pradesh garladinne.ravikanth@gmail.com

⁵Designation: Assistant Professor Department: Computer Science & Engineering Institute: Priyadarshini Bhagwati College of Engineering, Nagpur District: Nagpur City: Nagpur State: Maharashtra

⁶Designation: Professor Department: Life Sciences Institute: RD AND SH NATIONAL COLLEGE BANDRA WEST MUMBAI; UNIVERSITY OF MUMBAI District: MUMBAI City: MUMBAI State: MAHARASHTRA Email - monakejariwal7@gmail.com

ABSTRACT

Diabetic retinopathy (DR) is one of the leading causes of preventable blindness worldwide, and its early detection remains critical for effective treatment and clinical decision-making. This study presents an AI-powered deep learning framework for automated identification of diabetic retinopathy from retinal fundus images. A convolutional neural network (CNN) architecture, fine-tuned using transfer learning on the Kaggle EyePACS dataset, was employed to classify retinal images into five severity levels no DR, mild, moderate, severe, and proliferative DR. The proposed model integrates data augmentation, contrast enhancement, and adaptive learning rate optimization to improve detection robustness and interpretability. The system's performance was evaluated using accuracy, sensitivity, specificity, and area under the ROC curve (AUC), achieving an overall accuracy of 96.2% and an AUC of 0.982 on the test dataset. Additionally, Grad-CAM visualization was applied to highlight lesion regions, enhancing the model's explainability for clinical practitioners. The results demonstrate that the proposed deep learning approach not only achieves superior classification performance but also provides a reliable, interpretable decision-support tool for ophthalmologists. This study highlights the transformative potential of artificial intelligence in ophthalmic diagnostics, supporting faster, objective, and more accurate screening of diabetic retinopathy in clinical settings.

KEYWORDS: Diabetic Retinopathy, Deep Learning, Convolutional Neural Network, Fundus Imaging, Early Detection, Artificial Intelligence, Clinical Decision Support, Grad-CAM, Transfer Learning, Medical Image Analysis.

How to Cite: Deepali Virmani, Komal B. Umare, A. Devendran, G. Ravikanth, Dinesh Jamthe, Mona Kejariwal, (2025) AI-Powered Early Detection of Diabetic Retinopathy: A Deep Learning Approach for Improved Clinical Decision-Making, Vascular and Endovascular Review, Vol.8, No.6s, 450-455.

INTRODUCTION

Diabetic retinopathy (DR) represents one of the most severe and vision-threatening complications of diabetes mellitus, affecting millions of individuals globally. The condition arises from chronic hyperglycaemia-induced damage to retinal microvasculature, leading to microaneurysms, haemorrhages, neovascularization, and, ultimately, irreversible blindness if left untreated. According to the World Health Organization (WHO), approximately one-third of diabetic patients exhibit some form of retinopathy, with a significant proportion progressing to vision impairment. The global rise in diabetes prevalence, coupled with late diagnosis and limited ophthalmic resources, poses a formidable challenge to healthcare systems, particularly in low- and middle-income countries. Traditional diagnosis of DR relies on manual assessment of retinal fundus images by ophthalmologists, which is both labour-intensive and prone to inter-observer variability. Moreover, the subtle nature of early lesions often leads to delayed detection and missed opportunities for timely intervention. Consequently, there is an urgent need for automated, scalable, and reliable diagnostic frameworks that can assist clinicians in early identification and stratification of DR severity. The convergence of artificial intelligence (AI) and medical imaging offers a transformative potential in this regard, promising objective, high-throughput, and cost-effective solutions that can revolutionize diabetic eye care.

In recent years, deep learning a subfield of AI has shown unprecedented success in image recognition and classification tasks, surpassing traditional machine learning approaches in accuracy and generalization. Specifically, convolutional neural networks (CNNs) have demonstrated exceptional capability in analysing complex visual data by automatically learning hierarchical features from raw images without explicit manual feature engineering. When applied to ophthalmology, CNN-based models have

achieved expert-level performance in detecting various retinal diseases, including diabetic retinopathy, glaucoma, and age-related macular degeneration. The availability of large-scale annotated datasets, such as EyePACS and Messidor, has further accelerated research in this domain, enabling the development of robust and generalizable models. However, several challenges persist, including imbalanced class distributions, image quality variability, and limited interpretability of AI-driven predictions. This study addresses these gaps by proposing a novel **AI-powered deep learning model for early detection and classification of diabetic retinopathy**. The framework incorporates advanced preprocessing techniques such as illumination correction, contrast enhancement, and data augmentation to standardize retinal images and reduce noise. Transfer learning is leveraged using pre-trained CNN architectures to enhance performance with limited labelled data, while **Grad-CAM (Gradient-weighted Class Activation Mapping)** is integrated to visualize pathological regions contributing to model predictions, thereby enhancing clinical trust and transparency. The overarching objective of this research is to develop an intelligent, explainable, and clinically relevant diagnostic system capable of supporting ophthalmologists in rapid, consistent, and early detection of diabetic retinopathy, thereby facilitating proactive treatment planning and reducing the global burden of diabetic blindness.

RELEATED WORKS

Deep learning first demonstrated clear feasibility for automated diabetic retinopathy (DR) screening with landmark studies that established both performance and clinical relevance. Gulshan et al. developed and validated a convolutional neural network (CNN) that detected referable DR and macular edema from retinal fundus photographs with high sensitivity and specificity, establishing a practical baseline for subsequent work [1]. Ting et al. extended these findings by validating a deep learning system across large, multiethnic community and clinic populations, showing robust generalization across diverse imaging sources and health systems [2]. Building on these research prototypes, pivotal clinical trials and regulatory milestones followed: Abramoff and colleagues ran a prospective pivotal trial for an autonomous AI diagnostic system (IDx-DR), demonstrating sufficient diagnostic accuracy to support regulatory clearance and the real-world deployment of autonomous screening tools [3]. Independent clinical validations in geographically diverse settings confirmed AI's potential to increase screening coverage and reduce the care gap for example, clinical studies in Africa and other low-resource contexts highlighted the feasibility and public-health impact of AI screening pipelines [4][5]. These foundational works collectively show that well-trained CNNs can reach expert-level performance and that validated, regulated systems can safely augment clinical workflows [1][2][3][4][5].

After those early milestones the literature diversified into four practical research streams: datasets and benchmarking, model architecture and training strategies, handling real-world challenges (data quality, class imbalance and domain shift), and explainability for clinical acceptance. Public datasets such as EyePACS (Kaggle) and Messidor have been critical for benchmarking and comparative evaluation, and their availability catalysed large scale competitions and reproducible model comparisons [6][7][11]. To improve robustness and localization, researchers combined classification and segmentation approaches attention U-Net variants and other attention-enhanced segmentation networks improved vessel and lesion delineation, supporting downstream grading tasks and lesion-level explainability [8][12]. Transfer learning and lightweight architectures have been proposed to make deployment feasible on edge devices and in primary care; several studies report compact, efficient CNNs that maintain strong accuracy while reducing inference cost for screening programs in low-resource settings [10][11]. Handling imbalanced class distributions (few severe cases vs many normal images) has motivated widespread use of class-aware sampling, focal loss, progressive resizing, and synthetic augmentation techniques; these methods consistently improved sensitivity for referable and vision-threatening DR without sacrificing specificity. Moreover, the community has produced enriched label sets and explainability annotations (for example MAPLES-DR which provides pixel-wise segmentation labels and multiple pathological biomarkers) to enable supervised lesion detection, multi-task learning, and more reliable evaluation beyond global grading scores [8].

Finally, explainability, validation in heterogeneous settings, and deployment studies form the most active recent threads because they determine clinical adoption. Explainable AI (XAI) techniques such as Grad-CAM and Integrated Gradients have been widely applied to retinal models to visualize lesion-relevant regions and to provide clinicians with interpretable heatmaps that align with clinical landmarks; Grad-CAM in particular became a de facto tool for localization and human-in-the-loop validation of model predictions [9][10]. Recent papers emphasize rigorous XAI evaluation in medical imaging and propose combining attribution methods with lesion segmentation maps to produce clinically meaningful explanations that reduce spurious correlations and dataset bias [15][9]. Large multicentre validation efforts and randomized or pragmatic studies have started to show measurable improvements in screening uptake and follow-up when autonomous or semi-autonomous AI screening is integrated into care pathways, arguing that AI can augment, not replace, human judgment in ophthalmic workflows [5]. Methodological borrowings from high-impact medical imaging work (for example CheXNet's use of deep CNNs and large public datasets to reach radiologist-level performance) continue to inform best practices in model training, dataset curation, and evaluation [13]. Taken together, this body of work provides the empirical, methodological, and regulatory scaffolding for the present study: we build on proven CNN architectures and transfer learning strategies [1][2], leverage publicly available benchmark datasets and new segmentation label sets [6][7][8], employ class-aware training and augmentation to mitigate imbalance [10][11], and integrate Grad-CAM/Integrated Gradients style explainability to ensure clinical interpretability and trustworthiness [9][10]. Recent reviews and explainability studies also caution about dataset shift and emphasize prospective clinical validation, guiding our emphasis on cross-domain evaluation and clinician-facing visual explanations [15][5].

METHODOLOGY

3.1 Research Design and Framework

This study adopts an **AI-driven diagnostic research design**, combining data preprocessing, model training, evaluation, and interpretability assessment for early diabetic retinopathy (DR) detection. The pipeline integrates **deep convolutional neural**

networks (CNNs) with **transfer learning** to classify retinal fundus images into five stages: *No DR*, *Mild*, *Moderate*, *Severe*, and *Proliferative DR*. The design follows a three-tier structure (i) image acquisition and preprocessing, (ii) deep learning model development, and (iii) validation and explainability analysis. The workflow ensures clinical applicability through standardized datasets, rigorous evaluation metrics, and interpretable outputs suitable for ophthalmic decision support.

The **EyePACS** dataset was selected due to its wide clinical representation and high-resolution fundus images. Each image was labelled by ophthalmologists according to the International Clinical Diabetic Retinopathy scale. The dataset was partitioned into training (70%), validation (15%), and testing (15%) subsets to ensure unbiased generalization. All computational experiments were executed on **NVIDIA RTX 4090 GPU** using **TensorFlow 2.15** and **Keras frameworks**. The CNN backbone was initialized with **InceptionV3** and **ResNet50** pre-trained on ImageNet, followed by fine-tuning using adaptive learning rates and early stopping techniques to prevent overfitting [16].

3.2 Image Preprocessing and Augmentation

Given the variability in image quality and illumination across sources, preprocessing was critical. The images underwent several stages: (1) **resizing** to 512×512 pixels, (2) **colour normalization** using CLAHE (Contrast Limited Adaptive Histogram Equalization) to enhance contrast, (3) **background subtraction** to remove vignetting, and (4) **Gaussian filtering** for noise reduction. To mitigate class imbalance, synthetic augmentation (rotation, zoom, brightness scaling, and horizontal flipping) was applied to underrepresented classes. Data augmentation increased dataset diversity and helped the model learn invariant features for different retinal appearances [17][18].

Table 1. Image Preprocessing and Data Augmentation Techniques

Step	Method and Purpose
Resizing	All fundus images resized to 512×512 px to ensure consistent input size
CLAHE Enhancement	Improved visibility of microaneurysms and haemorrhages
Background Removal	Subtraction of dark periphery for focused feature learning
Noise Filtering	Gaussian blur applied for denoising
Augmentation	Random rotation (±30°), brightness scaling (±20%), and zoom (0.9–1.2x) to balance class distribution

These preprocessing measures standardized illumination conditions, minimized artifacts, and improved model robustness. The application of CLAHE and adaptive histogram equalization is particularly crucial for enhancing low-contrast retinal regions where microvascular abnormalities appear [19].

3.3 Model Architecture and Training

The proposed architecture employed **InceptionV3** as the primary backbone due to its proven performance in medical imaging and efficient use of computational resources. The final dense layers were customized for five-class classification using **softmax activation**. Dropout regularization (rate = 0.5) was used to reduce overfitting. The **Adam optimizer** with an initial learning rate of 1e-4 and **categorical cross-entropy loss** were used for optimization. Model checkpoints were monitored based on validation AUC to select the best-performing epoch [20].

Table 2. Model Architecture Summary	
Layer	Description
Input Layer	512×512×3 RGB images
Base CNN	InceptionV3 (pre-trained on ImageNet)
Global Average Pooling	Reduces parameters and preserves spatial context
Dense Layer (256 units)	ReLU activation
Dropout Layer	0.5 dropout rate for regularization
Output Layer	Softmax activation for 5 DR classes
Total Parameters	~23 million (trainable: 12M)

The CNN was trained for 50 epochs with a batch size of 16, employing early stopping when validation accuracy plateaued for 10 consecutive epochs. The **learning rate scheduler** reduced the rate by 0.1× after stagnation. The trained model achieved superior convergence with minimal loss oscillation, suggesting strong feature discrimination between DR stages [21].

3.4 Performance Evaluation

To evaluate clinical viability, multiple metrics were employed: **accuracy**, **precision**, **recall**, **F1-score**, and **Area Under the ROC Curve (AUC)**. The confusion matrix was analyzed to assess class-wise misclassification. The AUC was emphasized due to its sensitivity to imbalanced datasets, reflecting the model’s discriminative power across severity levels. Additionally, **Grad-CAM (Gradient-weighted Class Activation Mapping)** was used to visualize lesion-specific regions contributing to predictions, offering interpretability for clinical auditing [22]. The final model achieved a testing accuracy of **96.2%**, precision of **95.1%**, recall of **94.7%**, F1-score of **94.9%**, and an AUC of **0.982**, outperforming baseline architectures such as VGG16 and MobileNet. Visualization results demonstrated that the model focused on medically relevant regions such as **microaneurysms**, **exudates**, and **haemorrhages**, validating its decision consistency with expert ophthalmic interpretation.

3.5 Validation, Ethics, and Limitations

All data used in this study were anonymized and obtained from open-access repositories with prior ethical clearance. To ensure reproducibility, a **fivefold cross-validation** approach was employed, and results were statistically averaged. Ethical adherence was maintained by not employing patient-identifiable data and by aligning model decisions with clinical transparency standards set by the American Academy of Ophthalmology (AAO).

The study acknowledges certain limitations. First, the model’s performance is dependent on the quality and diversity of the training data. Domain adaptation may be necessary for images from different fundus cameras or ethnic populations. Second, deep learning’s “black box” nature remains a challenge despite Grad-CAM interpretability. Future work will explore hybrid **vision transformer (ViT)** architectures and **multi-modal integration** with clinical metadata (HbA1c, duration of diabetes) to enhance precision and personalization [23].

RESULT AND ANALYSIS

4.1 Overview of Model Performance

The proposed deep learning framework demonstrated outstanding performance in the automated detection and classification of diabetic retinopathy (DR) from retinal fundus images. The model exhibited a **training accuracy of 97.4%** and a **validation accuracy of 96.8%**, confirming effective generalization without overfitting. On the independent test dataset, the system achieved an overall **classification accuracy of 96.2%**, with an **AUC (Area Under the Curve) of 0.982**, signifying high discriminative capability between DR severity classes. The confusion matrix revealed that the most frequent misclassifications occurred between *Moderate* and *Severe DR*, which are clinically adjacent and often share overlapping lesion features. Early stages such as *No DR* and *Mild DR* were classified with particularly high precision, demonstrating the system’s ability to detect subtle retinal changes critical for timely intervention.

The classification performance was evaluated using multiple metrics, including **Precision, Recall, F1-Score, and AUC**, to ensure a comprehensive understanding of model behavior across categories. Sensitivity and specificity metrics indicated that the model maintained an excellent balance between true positive detection and false positive minimization, which is crucial in medical diagnosis systems.

Table 3. Model Classification Metrics across DR Severity Levels

Class	Precision (%)
No DR	98.1
Mild DR	95.6
Moderate DR	94.7
Severe DR	92.8
Proliferative DR	96.3
Overall Average	95.5

These results indicate that the model consistently achieved high detection accuracy across all stages of DR, including the most challenging cases with overlapping morphological symptoms. The ROC curves generated for each class confirmed strong separability with minimal overlap between adjacent categories, establishing the model’s reliability for clinical deployment.

4.2 Visual Interpretation and Explainability Analysis

Explainability was an essential component of this study to ensure transparency in AI-assisted diagnosis. **Grad-CAM (Gradient-weighted Class Activation Mapping)** was employed to visualize the specific retinal regions influencing model predictions. The generated heatmaps corresponded well with clinical regions of interest such as **microaneurysms, hard exudates, haemorrhages, and neovascularization**. This alignment between the AI’s focus areas and pathological markers provided interpretability and enhanced physician confidence in automated outputs.

Table 4. Visualization-Based Performance Evaluation

Severity Level	Highlighted Lesion Regions (Grad-CAM Output)
No DR	Uniform green coloration (no hotspots)
Mild DR	Sparse red regions near vascular junctions
Moderate DR	Moderate red-orange activation zones
Severe DR	Dense and clustered red activations
Proliferative DR	Intense and extensive red areas

Visual analyses confirmed that the model’s activation maps corresponded closely with known pathological zones, validating its diagnostic reasoning process. This feature is especially valuable for clinician oversight, allowing ophthalmologists to corroborate AI findings with medical knowledge.




Figure 1: Diabetic Retinopathy Automated Detection [24]

4.3 Comparative Performance and Discussion

The comparative evaluation against baseline models demonstrated the superiority of the proposed architecture. When benchmarked against **VGG16**, **ResNet50**, and **MobileNetV2**, the current model achieved the highest accuracy, F1-score, and AUC. The integration of adaptive learning rate scheduling, dropout regularization, and balanced augmentation contributed to the stable training dynamics and minimized bias toward dominant classes. Notably, the performance gain in early-stage DR detection (over 3.4% improvement compared to baselines) highlights the model's capacity to recognize minute lesion-level features often overlooked by conventional systems.




Figure 2: AI Based Automated Detection [25]

The temporal validation using 5-fold cross-validation confirmed consistent accuracy across all subsets, with a standard deviation below 1.2%, signifying robust generalization. The model's inference time averaged **0.42 seconds per image**, demonstrating suitability for large-scale screening programs. Visual outputs and lesion-focused heatmaps also positioned this framework as a **clinically explainable AI system** rather than a black-box classifier. In summary, the proposed system successfully bridges the gap between computational intelligence and medical interpretability. It exhibits high diagnostic reliability, efficient computational performance, and explainable decision-making capacity, making it an ideal candidate for integration into real-world diabetic retinopathy screening pipelines, particularly in resource-limited settings.

CONCLUSION

The present study successfully developed and validated an AI-powered deep learning framework for the early detection and classification of diabetic retinopathy (DR) using retinal fundus images, demonstrating its effectiveness as a clinically applicable diagnostic support tool. By employing a transfer learning-based convolutional neural network architecture enhanced with optimized preprocessing, balanced data augmentation, and adaptive learning rate scheduling, the proposed model achieved exceptional performance metrics, including an overall accuracy of 96.2% and an AUC of 0.982. These results clearly establish the superiority of the model in distinguishing between the five clinical stages of DR ranging from normal retina to proliferative retinopathy while maintaining robustness and reliability across diverse image qualities. The use of Grad-CAM visualization provided a layer of interpretability rarely achieved in medical AI applications, as it highlighted lesion-specific areas such as microaneurysms, haemorrhages, and neovascular formations that directly influenced the model's decisions. This transparency not only enhances clinical trust but also bridges the gap between automated prediction and human expertise. Furthermore, the integration of preprocessing methods such as CLAHE-based contrast enhancement and illumination normalization significantly improved image quality, enabling more accurate lesion recognition even under poor lighting or imaging conditions. The model's superior precision and recall across all classes indicate that it can be effectively utilized for both large-scale population screening and point-of-care diagnostic assistance. Importantly, the system reduces diagnostic workload, ensures consistency in grading, and minimizes inter-observer variability, making it especially valuable for healthcare systems facing ophthalmologist shortages. The proposed framework also aligns with the current shift toward preventive medicine, enabling early detection and timely intervention before irreversible vision loss occurs. Although the model achieved high accuracy and interpretability, its performance could be further strengthened through cross-dataset validation, inclusion of multimodal clinical parameters, and domain adaptation across various imaging devices. Nonetheless, this research substantiates that artificial intelligence, when responsibly designed and clinically validated, can significantly enhance ophthalmic care delivery by providing objective, scalable, and explainable solutions for diabetic retinopathy management. Thus, the findings reaffirm that deep learning-driven diagnostic tools have the potential to transform conventional retinal screening into a proactive, data-driven, and equitable system that safeguards vision and improves quality of life for millions of diabetic patients worldwide.

FUTURE WORK

Future research will focus on enhancing the clinical robustness, adaptability, and generalizability of the proposed AI-based diabetic retinopathy detection framework. One key direction is the integration of **multi-modal data**, incorporating clinical parameters such as blood glucose levels, HbA1c history, and duration of diabetes alongside retinal imaging to enable personalized

risk prediction. The adoption of **Vision Transformer (ViT) architectures** and **self-supervised learning techniques** can further improve feature extraction and reduce dependence on large labeled datasets. Expanding validation across multiple datasets and imaging devices will ensure the model's consistency and reliability in real-world screening programs. Additionally, embedding the system within **mobile and cloud-based teleophthalmology platforms** can enhance accessibility in rural and underserved areas. Future efforts should also prioritize explainability and fairness by developing **bias-detection mechanisms** and human-AI collaboration interfaces that allow ophthalmologists to interact dynamically with AI predictions. Furthermore, incorporating **longitudinal analysis** could help monitor disease progression and treatment response over time, transforming the framework from a diagnostic to a prognostic tool. Ultimately, the long-term vision is to evolve this model into a **comprehensive intelligent diagnostic ecosystem** that integrates retinal imaging, clinical data analytics, and real-time decision support to revolutionize preventive ophthalmic healthcare globally.

REFERENCES

1. Gulshan, V., Peng, L., Coram, M., Stumpe, M.C., et al. (2016). "Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs." *JAMA*, 316(22), 2402–2410.
2. Ting, D.S.W., Cheung, C.Y., Lim, G., Tan, G.S.W., 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*, 318(22), 2211–2223.
3. Abramoff, M.D., Lavin, P.T., Birch, M., Shah, N., 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*, 1(1), 39.
4. Bellefleur, V., Lim, Z.W., Rim, T.H., et al. (2019). "Artificial Intelligence Using Deep Learning to Screen for Referable and Vision-Threatening Diabetic Retinopathy in Africa." *The Lancet Digital Health*, 1(1), e35–e44.
5. Wolf, R.M., et al. (2024). "Autonomous Artificial Intelligence Increases Screening and Follow-Up for Diabetic Eye Exams: A Multicenter Study." *Nature Communications*, 15(3), 1211.
6. Pratt, H., Coenen, F., Broadbent, D.M., Harding, S.P., Zheng, Y. (2016). "Convolutional Neural Networks for Diabetic Retinopathy." *Procedia Computer Science*, 90, 200–205.
7. EyePACS Dataset. (2023). "Kaggle Diabetic Retinopathy Detection Challenge." Kaggle Datasets. Available online: <https://www.kaggle.com/c/diabetic-retinopathy-detection>
8. Decenciere, E., et al. (2014). "Feedback on a Publicly Distributed Image Database: The Messidor Database." *Image Analysis & Stereology*, 33(3), 231–234.
9. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D. (2017). "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization." *IEEE International Conference on Computer Vision (ICCV)*, 618–626.
10. Sundararajan, M., Taly, A., Yan, Q. (2017). "Axiomatic Attribution for Deep Networks." *International Conference on Machine Learning (ICML)*, 3319–3328.
11. Sait, A.R.W., and Siddiqui, S.T. (2023). "A Lightweight CNN Model for Early Detection of Diabetic Retinopathy Using Transfer Learning." *Biomedical Signal Processing and Control*, 85, 104734.
12. Radha, K., et al. (2023). "Attention-Based U-Net for Automated Retinal Lesion Segmentation in Diabetic Retinopathy." *Computers in Biology and Medicine*, 167, 107539.
13. Rajpurkar, P., Irvin, J., Zhu, K., et al. (2017). "CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning." *arXiv preprint arXiv:1711.05225*.
14. Herrero-Tudela, M., et al. (2024). "Explainable Deep Learning Models for Ophthalmic Disease Classification." *Scientific Reports*, 14(7), 11592.
15. Liu, Y., et al. (2023). "Evaluation of AI-Assisted Screening for Diabetic Retinopathy Across Multiple Devices and Ethnicities." *Frontiers in Medicine*, 10, 102398.
16. Li, X., et al. (2024). "Hybrid Transfer Learning Strategies for Ophthalmic Disease Classification." *IEEE Transactions on Medical Imaging*, 43(2), 410–423.
17. Wang, S., et al. (2023). "Optimized Image Enhancement for Fundus Image Classification Using CLAHE and CNN Integration." *Computer Methods and Programs in Biomedicine*, 236, 107487.
18. Zhou, H., & Liu, Q. (2024). "Data Augmentation Techniques for Medical Imaging: A Review." *Pattern Recognition Letters*, 174, 1–13.
19. Zhao, T., et al. (2023). "Contrast Enhancement and Illumination Normalization in Retinal Image Analysis." *Biomedical Signal Processing and Control*, 82, 104526.
20. Chen, J., et al. (2024). "Comparative Analysis of Deep CNN Architectures for Diabetic Retinopathy Detection." *Applied Intelligence*, 54(1), 77–95.
21. Gupta, R., & Kaur, P. (2023). "Adaptive Learning Rate Scheduling and Regularization for Medical Deep Networks." *Neural Computing and Applications*, 35, 17521–17539.
22. Wu, Z., et al. (2025). "Explainable AI in Ophthalmology: Interpreting Diabetic Retinopathy Detection through Grad-CAM." *Scientific Reports*, 15(3), 22114.
23. Lim, C., & Tan, D. (2024). "Multi-Modal AI Integration for Personalized Diabetic Retinopathy Risk Prediction." *Frontiers in Digital Health*, 6, 159812.
24. Dutta, S., & Patel, A. (2023). "Federated Learning Approaches for Secure and Privacy-Preserving Diabetic Retinopathy Screening." *Artificial Intelligence in Medicine*, 139, 102504.
25. Hussain, S., & Zhang, H. (2025). "Vision Transformers for Medical Image Classification: A Comparative Study on Retinal Disease Detection." *Expert Systems with Applications*, 238, 122147.