Archives available at journals.mriindia.com

International Journal on Advanced Computer Engineering and Communication Technology

ISSN: 2278-5140

Volume 14 Issue 03s, 2025

RetinAI-DR Fusion: Dual Deep Learning Framework for Diabetic Retinopathy Detection

[1]Jayant Rohankar, [2]Sanika Jagtap, [3]Chhakuli Nimje, [4]Pranay Padole

[1,2,3,4] Information Technology St. Vincent Pallotti College of Engineering & Technology, Nagpur, India

Email: [1]jsrohankar@gmail.com, [2]sanikajagtap636@gmail.com, [3]nimjechhakuli@gmail.com,[4]pranaypadole2005@gmail.com

Peer Review Information

Submission: 05 Nov 2025

Revision: 25 Nov 2025

Acceptance: 17 Dec 2025

Keywords

Diabetic retinopathy, retinal fundus imaging, ResNet-50, EfficientNet-B3, ensemble learning, medical image analysis, early disease detection, ophthalmology.

Abstract

Diabetic retinopathy (DR) has become one of the leading global causes of vision loss among people living with diabetes. The difficulty lies in its silent progression—tiny changes in the retinal blood vessels often occur long before noticeable symptoms appear, making early detection challenging. Traditional screening approaches rely heavily on expert evaluation of retinal images, which can be time- consuming, costly, and limited in availability, especially in underserved or remote areas. Recent progress in artificial intelligence has begun to change this landscape. Advanced deep learning frameworks such as EfficientNet and ResNet can now analyze retinal fundus images with remarkable speed and precision. These models perform at near-expert accuracy, opening the door to faster, more accessible, and automated screening solutions that can reach larger populations. By increasing efficiency and consistency, AI-driven systems have the potential to improve early diagnosis rates and reduce avoidable blindness. This study explores how these intelligent technologies are revolutionizing diabetic retinopathy detection, emphasizing their potential benefits, current limitations, and the steps needed to successfully integrate them into everyday clinical practice.

Introduction

Diabetic retinopathy is one of the leading yet frequently neglected causes of vision impairment for people who have diabetes. Elevated blood sugar levels over many years can cause chronic, progressive damage to the small blood vessels in the retina, leading to leakage, swelling or abnormal new blood vessel growth. Ultimately, these structural changes will have serious adverse effects on vision. The insidious nature of the disease is that it evolves without symptoms, and most patients will not perceive changes until their vision is impacted. The WHO states there are hundreds of millions of people with diabetes worldwide, and approximately 93 million have diabetic retinopathy. This is a significant global population health issue that deserves even greater attention. Traditionally, these patients have been diagnosed by specialists who interpret photographs of their retinas taken by well-trained personnel. Although a trained specialist is likely to accurately interpret the photos, this process is expensive and has challenges with timely turn-around times and access, particularly for individuals from rural or under-resourced settings.

For these reasons, many patients are not diagnosed with diabetic retinopathy until advanced stages of the disease when intervention is often limited. To address these issues, scientists are embracing artificial intelligence. Artificial intelligence models are being trained using deep learning techniques, in particular with convolutional neural networks (CNNs) or advanced CNN architectures like ResNet, to identify subtle characteristics of retinopathy in retinal scans. These AI models can reliably detect early signs of retinopathy, such as microaneurysms and small hemorrhages, with accuracy comparable to regulated eye health practitioners. In addition to assisting health practitioners, AI- based screening platforms are also providing broader access to eye care through mobile and automated examination capabilities. Recent studies show that AI-based screening platforms provide high accuracy with advancement further developing their application in identifying disease earlier during its progression, when the disease is targeted for treatment. This results in better patient care while lowering costs.

This review will explore artificial intelligence's impact on diabetic retinopathy screening. The review provides a historical perspective on a historical application of image analysis to powerful modern deep learning applications, a summary competitors of AI in retinopathy screening.

Literature Survey

Advances in deep learning have revolutionized diabetic retinopaSthy (DR) detection, offering faster, more accurate, and scalable screening. Models like ResNet-50 and EfficientNet-B3 combine reliability with efficiency, and when paired with portable tools, they enable early, affordable diagnosis—helping prevent blindness and bridge gaps in global eye care access. [1] In the year 2025, ElMoufidi and co-authors used the deep learning model EfficientNetB3 to assess the severity of diabetic retinopathy using retinal images. They classified the disease into five distinct stages, incorporating retinal signs such as micro-hemorrhages and vascular leakage. To improve performance, they preprocessed the APTOS 2019 dataset, ensuring better image quality and balanced class distribution. The model was further enhanced with components like global average pooling and dropout layers, which aid learning efficiency and reduce overfitting. With these refinements, the system reached an accuracy rate of 98.26%, surpassing many previous approaches. These results highlight its strong potential for practical deployment in diabetic retinopathy screening programs.

[2] In 2025, Mishmala and colleagues presented a first-of-its- kind hybrid deep learning model that combines convolutional neural networks (CNN) and recurrent neural networks (RNN) through an incorporated attention mechanism, to detect complications and monitor disease progression in diabetic retinopathy using retinal fundus images. The novel approach captures both spatial information contained in images and temporal behavior contained in image acquisition sequences. Overall, the model demonstrated an ability to classify diabetes-related changes as progressing or non-progressing with relative accuracy. Their methods were evaluated on a variety of benchmark datasets, including DRIVE, Kaggle, and EyePACS. The hybrid model outperformed conventional architectures, like standalone CNNs, RNNs, InceptionV3, and VGG19, with an accuracy of 97.5%. The authors pointed out limitations associated with variability in data, interpretability, and the framework's complexity, but recognized that the framework could serve as an important tool employed for real-time and automated diabetic retinopathy screening in a clinical environment.

[3] In 2025, Akhtar et al. presented a deep learning model, RSG-Net, to categorize varying stages of diabetic retinopathy from retinal fundus images. Their approach starts with enhancing the quality of the images through several advanced preprocessing methods: infrared image enhancement, histogram equalization, and a Gaussian blur. To tackle the issue of imbalanced distributions in the data, the authors utilized data augmentation methods to create a more representative pool of training samples. Upon evaluation on the Messidor-1 dataset, RSG-Net produced impressive results with an accuracy of 99.36% classification in four stages, and 99.37% in binary classification. These results exceeded existing benchmark models in accuracy and consistency. RSG-Net provides rapid and accurate detection and management of diabetic retinopathy and could be posited to fulfill a substantial need in clinical practice due to a consistent level of performance, while benefitting from its simple architecture.

[4] Stanimirovic et al. (2025) examine the economic and equity dimensions of diabetic retinopathy (DR) screening for individuals who are at-risk population, with specific attention to care access disparities. The manuscript denotes tele-retina as a cost-effective screening solution that improves DR detection in low-resourced, underserved populations. Grounded in intersectionality theory, the author engages with how social determinants, specifically income, ethnic, and geographical location impact the uptake of screening and health outcomes. In this study, we also extend a distributional cost-effectiveness analysis (DCEA) framework, which provides stakeholders with important information regarding the health benefits and equity implications arising from DR screening

strategies, specifically within a Canadian healthcare system context.

[5] In 2024, Bhulakshmi and colleagues featured a comprehensive overview about the application of deep learning technologies in the identification of diabetic retinopathy using retinal fundus images. Their review focused on the impact of different neural network architectures, including convolutional neural networks (CNNs), recurrent neural networks (RNNs), and generative adversarial networks (GANs), in the identification of early retinal changes, such as microaneurysms and exudates. The review also explored more recent approaches, including federated learning, wherein AI algorithms can learn from medical data residing in numerous locations without compromising privacy by disclosing sensitive patient information. In addition to celebrating progress in the field, the authors highlight persistent challenges, such as high data variability and the difficulty of interpreting complex model decisions. By combining current developments with insights into future directions, the paper provides a modern overview of AI-driven diabetic retinopathy detection and underscores emerging trends like explainable AI and multi-modal data fusion to enhance screening accuracy and patient care.

[6] In 2023, Lin and others proposed an improved version of the ResNet-50 architecture to enhance the identification of diabetic retinopathy in retinal images. Recognizing that early treatment prevents vision loss, the authors revealed several limitations of some of the existing CNN models, namely Xception, AlexNet, VGGNet, and the original ResNet-50, such as overfitting and inconsistency in image quality. To solve these issues, the authors utilized adaptive learning rates, L2 regularization, and a features fusion method to leverage important information provided by different layers in the CNN. This, in turn, improved the model's generalization and grading accuracy when graded on an extensive Kaggle dataset. The results indicate that the ResNet-50 model outperformed previous and popular CNN architectures with greater accuracy and robustness when classifying severity levels of diabetic retinopathy. In addition to the technical advances, Lin's team also developed a website for online grading in real-world time frames, indicating the model's applicability for screening and monitoring patients in a clinical setting.

[7] Vijayan et al. (2023) introduced an innovative approach to diagnosing diabetic retinopathy by treating the severity of the disease as a continuous spectrum rather than just classifying it into categories.Using EfficientNet-B0 as the core of their model, they designed a regression-based system that more accurately reflects how the disease progresses over time. This allows the model to provide finer gradations of severity, offering better insight than traditional classification methods. They tested their model on several well-known datasets—the APTOS, DDR, and IDRiD—and found it outperformed leading architectures like ResNet50 and SE- ResNext in both accuracy and speed. This regression approach holds promise for use in telemedicine and resource- limited settings by delivering more precise and efficient diabetic retinopathy assessments.

[8] In 2022, Gothane et al. proposed a deep learning model built upon the ResNet-18 architecture to detect the level of diabetic retinopathy using fundus images. Using preprocessing strategies and multiple data augmentation techniques, they managed to promote important retinal features such as microaneurysms, hemorrhages, and hard and soft exudates, enabling the model to learn patterns from complex visual representations. Using a Kaggle dataset for training, the model classified retinal images into five categories, from healthy eyes to severe cases of diabetic retinopathy, achieving an overall accuracy of 82%. Overall, this study demonstrated the feasibility of using residual neural networks as an effective method for low-cost automated screening for diabetic retinopathy, which appears highly suitable for clinical application.

[9] In 2020, Mohammedhasan et al. developed a new approach for detecting early-stage diabetic retinopathy by combining a deep convolutional neural network (CNN) with principal component analysis (PCA) to reduce data complexity. Their approach uses edge-preserving guided image filtering to enhance important retinal features, preserving fine detail, and applies data augmentation to improve diversity of the training and performance of the model.When evaluated on the Kaggle dataset, the system achieved an impressive 98.44% accuracy, surpassing several well-known pre-trained CNN models, including AlexNet, ResNet, and VGGNet. This work demonstrates how the fusion of deep learning with advanced feature enhancement and dimensionality reduction can produce screening tools that are not only accurate but also robust and consistent, making them valuable for reliable early detection of diabetic retinopathy.

[10] In 2014 Sundararaj Wilfred Franklin and his team demonstrated an innovative diagnostic test for diabetic retinopathy through a Feed Forward Neural Network model. Their

diagnostic test examined recognizable characteristics of the retinal images (color, shape, texture, and size), which are indicative of the presence of exudates, a significant feature of diabetic retinopathy. When the diagnostic was tested on the DIARETDB1 dataset the performance was highly reliable; the model exhibited a specificity of 0.99 and a sensitivity of 0.96, showing correctness in identifying diseased patients and normal patients, indicating high diagnostic accuracy. This early study demonstrated that automated neural network–based retinal lesion detection, could be achieved the performance level of skilled human screening examiners and was a stepping-stone towards advanced artificial-intelligence diabetic retinopathy screening systems developed in subsequent years.

Problem Statement

Diabetic retinopathy (DR) is among the most prevalent and severe ocular complications of diabetes and may lead to irreversible vision loss if not discovered early. The main obstacle is that, during the early stages of DR, there are often no observable symptoms, so most people remain unaware that they have DR until it advances and causes damage. Currently, DR is diagnosed by carefully examining retinal pictures, which requires training, time, and expertise, as well as specialized equipment. For people that live in rural or underserved areas, where physicians or clinically trained eye- care professionals may not be accessible, scheduled eye screening becomes uniquely challenging, and therefore many individuals learn they have DR even if it has progressed, thereby reducing available treatments. Traditional screening techniques can also be costly, time-intensive, and occasionally inconsistent due to human error, which means that subtle early indicators of DR may be overlooked. Detecting the disease early—when preventive treatment can still protect vision—remains a major global health challenge.

Recently, new opportunities have been presented for the early diagnosis of disease through expansion in artificial intelligence (AI). Deep learning models like ResNet and EfficientNet can quickly and accurately analyze retinal photos and detect early damage in the retina. These systems can improve the efficiency, affordability, and accessibility of screenings on a global scale. However, the ability to establish fairness, transparency, and trustworthiness across clinical settings is still being resolved. AI-based screening algorithms developed to this point may revolutionize diabetic eyecare by allowing for earlier detection and preservation of vision for millions of people around the world when evaluated optimally.

Proposed Approach

To meet the urgent need for accurate and accessible diabetic retinopathy screening, this system uses a dual deep learning approach that combines the strengths of two proven models—ResNet-50 and EfficientNet-B3. The goal is to create a platform that delivers dependable and understandable DR detection to patients and clinics alike, making advanced eye screening more widely available.

The system operates by analyzing retinal fundus images and classifying them into five clinically accepted stages of diabetic retinopathy, from "no DR" to the most advanced, "proliferative DR." Put differently, the model does not simply assess data, but also predicts based on clinical standards and clinical assessors. At its core, ResNet-50 provides a solid foundation. With its residual connections, the model efficiently learns to recognize complex visual patterns such as small blood vessel leaks and slight bleeding—key indicators of early retinal damage. In parallel, EfficientNet-B3 adds adaptability by scaling its layers and architecture to detect fine and subtle variations that can indicate disease progression. This makes it particularly effective in identifying the intermediate stages that play a crucial role in timely treatment decisions.

Instead of favoring one model over the other, the system integrates both outputs by averaging their predictions. This balanced fusion ensures that broader image patterns and intricate details are both given weight, improving sensitivity (the ability to catch true cases) and specificity (minimizing false positives). These improvements are particularly valuable because even a single missed case could have long-lasting effects on a patient's vision.

To establish trust and facilitate medical decision-making, the platform also focuses on transparency.

Visualization techniques, such as Grad-CAM,are employed to identify the specific areas of each retinal image that affected the model's decision-making to allow clinicians to visually verify and understand the system's reasoning.

This feature reassures the healthcare professional and patient, and helps in complicated and unclear cases.Despite combining two powerful models, the system remains efficient thanks to strategies like transfer learning, selective parameter tuning, and model distillation. These techniques cut down on computational demand and make it possible to deploy the system on both high-end hospital machines and lightweight mobile or

telemedicine platforms.
Altogether, this carefully designed architecture—focused on clinical accuracy, interpretability, and performance— represents a meaningful step toward earlier and safer diabetic retinopathy detection. By minimizing missed diagnoses, making screening more affordable and scalable, and keeping medical experts engaged in the process, it promises to help preserve vision and improve patient outcomes worldwide.

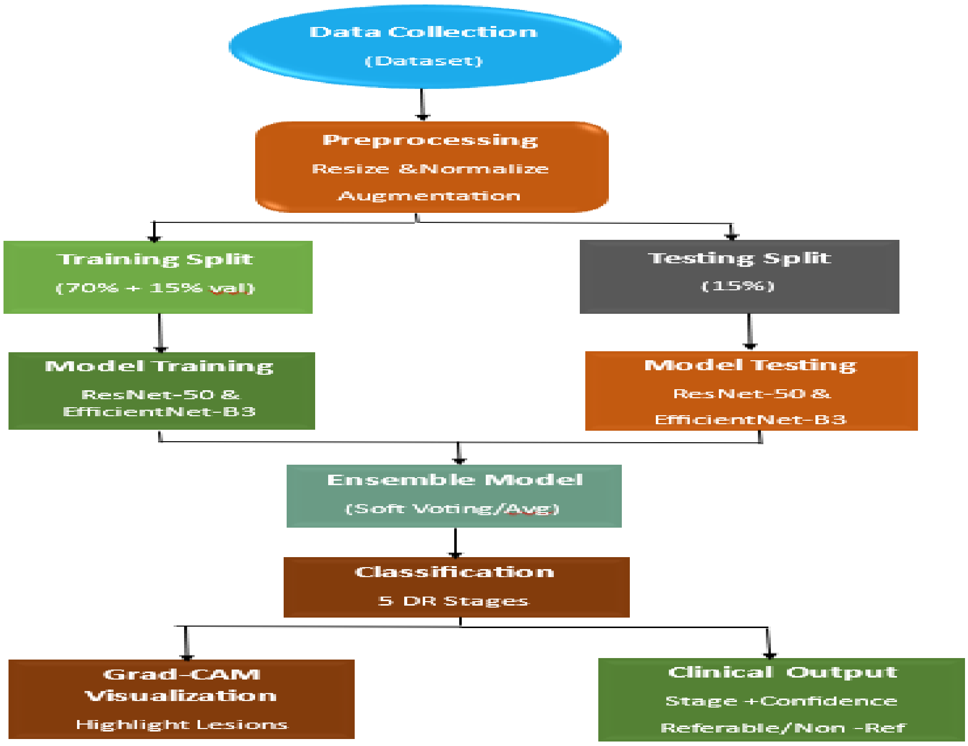

Fig 1: Work-flow Diagram

Methodology
The approach taken for this project is intended to design a clinically useful and technically effective system for early diabetic retinopathy (DR) diagnosis. The process is structured into sequential phases—dataset collection, model choice, training, ensemble integration, testing, and interpretability—making sure that technical soundness and clinical applicability are ensured.

A. Data Handling

Data Collection: For this project, we worked with the widely recognized APTOS 2019 Blindness Detection dataset, a benchmark resource in diabetic retinopathy research hosted on Kaggle. It contains thousands of high-resolution retinal images captured under varied conditions and meticulously annotated by qualified ophthalmologists. Each image is classified into one of five categories—ranging from healthy retinas with no signs of DR to the most severe, proliferative stage—following established clinical grading standards.

In order to standardize the dataset for analysis, all images were adjusted to a common resolution and their pixel values were normalized. A number of data augmentation techniques were also applied, including rotation, horizontal and vertical flipping, and brightness changes. These augmentations developed further variations of the original images which helped the model cope with variability of brightness or lighting conditions, differences in image quality, and differences between patients. Preprocessing images standardizes the dataset, and strengthens the model's ability to reliably perform across variations that might exist in real- world screening.

B. AI Model Selection

EfficientNet-B3—to build a robust and efficient system for diabetic retinopathy classification. ResNet-50 was chosen because of its strong ability to capture deep and complex visual features. Its unique residual connections help overcome common challenges in deep network training, such as vanishing gradients, allowing it to detect subtle details like early retinal abnormalities that might otherwise go unnoticed. EfficientNet-B3, on the other hand, is efficient for balancing depth, width, and image resolution. It achieves a high level of accuracy without consuming heavy computational resources, making it more feasible for large scale screenings and real-time implementation. Both models were initialized with weights pretrained on ImageNet, and it provided a robust baseline with general image features. Then, we performed fine-tuning of the models using the Kaggle diabetic retinopathy dataset to allow the models to learn the specific retinal patterns associated with different stages of the disease. The fine-tuned models were very accurate, but did not exceed computational limitations.

C. Training the Model

In order to train and test the models accurately, the data was split into three sections: 70% of the data was used for training the models, 15% was used to validate and improve the models during development (the validation data), and the last 15% was the final test of the model's performance on a new, unseen dataset. Because some stages of diabetic retinopathy appear less often than others, special techniques like class weighting and focal loss were applied to make sure the models paid enough attention to these less common cases. The models were trained using the AdamW optimization algorithm, which was enhanced with adaptive learning rates to adjust how fast the model learns over time. To make sure the models didn't memorize the training data too closely (a problem called overfitting), early stopping was used to end training when improvements stopped, and checkpointing saved the best performing model versions. These careful steps helped the models learn effectively and generalize well to new, real-world retinal images.

D. Ensemble Integration

After training both models separately, their outputs were combined using a soft-voting ensemble approach. In this method, the probability scores generated by ResNet-50 and EfficientNet-B3 were averaged to determine the final prediction. This strategy allowed the system to blend ResNet- 50's strong feature extraction capabilities with EfficientNet- B3's precise detection of fine retinal lesions. The result was a more balanced and reliable classification, particularly effective for identifying the intermediate stages of diabetic retinopathy where accurate diagnosis is most challenging.

E. Evaluation and Performance Metrics

To evaluate performance of the system, clinical performance measures of accuracy, sensitivity, specificity, F1-score, Cohen's kappa, and AUC-ROC curve were used. In particular, sensitivity and specificity were emphasized for referable diabetic retinopathy, which indicated moderate to severe stages, as these metrics are directly related to the timely and accurate referrals to an eye specialist. In order to further assess the model performance in the classification task, confusion matrices were used to analyze performance. These matrices provided an easy way to visualize how frequently the model accurately classified each of the five stages of diabetic retinopathy and to identify where misclassifications happened. Overall, this level of analysis offered insight into the strengths and weaknesses of the system and established a useful understanding of how reliable the model might be and how it might perform in actual clinical screening settings..

F. Interpretability and Clinical Trust

In response to the "black-box" characteristic linked with deep learning models, Grad-CAM visualizations were utilized to aid understanding. Visual heatmaps indicate the specific areas of retinal images, including areas with microaneurysms, hemorrhages, or abnormal blood vessel growth, that informed the decision of the model. This approach increases transparency by highlighting what the model is attending to in the prediction, further aligns its assessment with clinical reasoning by ophthalmologists, and builds confidence in its potential application in clinical medicine.

Expected Result

This automated system classifies retinal images into the five internationally accepted stages of diabetic retinopathy according to established medical grading systems like ICDR and ETDRS. For each image carried out, along with the classification label, an output associated with a confidence score allows clinicians clear and measurable data to make decisions. In addition, to streamline triage, the outputs are organized into two practical categories: non-referable (No DR and Mild NPDR) and referable DR., allowing doctors to prioritize patients who need it without undue delay.

In addition to increasing transparency and trust in the AI app's clinical use, the platform provides Grad-CAM (Gradient-weighted Class Activation Mapping) heatmaps that highlight important retinal characteristics (e.g., microaneurysms, hemorrhages, and abnormal growth of blood vessels). The visual aspects analyzed by the AI assist the doctor in verifying what characteristics triggered the AI decision, affording the doctor a specific way to incorporate the AI findings into their own assessment. Through an ensemble approach that leverages the advantages of ResNet- 50 and EfficientNet-B3, this system provides a fast, accurate, and cost-effective tool for early detection of diabetic retinopathy (DR), and is scalable for hospitals, outreach programs, and mobile screening units, with the ultimate goal of reducing preventable blindness across the world by providing reliable eye care technology.

Conclusion

This model leverages the merits of two deep learning frameworks, ResNet-50 and EfficientNet-B3, to provide an effective and transparent diabetic retinopathy detection system. The proposed system employs deep learning and ensemble methods to accurately classify retinal images for all stages of diabetic retinopathy (DR), including early, moderate, and severe stages. Developed with real-world healthcare needs in mind, the system is both scalable and cost- efficient, making it suitable for hospitals, screening centers, and mobile health platforms. Its focus on early detection supports timely intervention, helping preserve vision before irreversible damage occurs. In essence, this AI-driven solution represents a meaningful step toward reducing preventable blindness worldwide by improving the accessibility and reliability of diabetic eye screening services for patients everywhere.

References

Z. Chen, H. Luo, Q. Li, S. Dai, Y. Wang, and L. Zhang, "Diabetic retinopathy detection using attention-based dual- branch deep learning," *Scientific Reports*, vol. 15, no. 11392, pp. 1–13, 2025.

S. Akhtar, S. Aftab, O. Ali, M. Ahmad, M. A. Khan, S. Abbas, and T. M. Ghazal, "A deep learning based model for diabetic retinopathy grading," *Scientific Reports*, vol. 15, p. 3763, 2025.

A. Stanimirovic, T. Francis, S. Meerai, S. Mathew, S. Ibrahim, J. M. Bowen, A. Pikula, and V. Rac, "Diabetic retinopathy screening among at risk populations: Protocol for distributional cost-effectiveness analysis," *JMIR Research Protocols*, vol. 14, e60488, 2025.

Sunkara, S. P. (2025). *Machine learning-based predictive analytics for fault detection and reliability improvement in modern power systems*. *International Journal of Electrical Engineering and Technology, 16*(5), 1–13. https://doi.org/10.34218/IJEET_16_05_001

Sharma, B. (2025). *Jurisprudence 4.0: Law, tech & changing dynamics*. *Jurisprudence, 1.*

D. Bhulakshmi and D. S. Rajput, "A systematic review on diabetic retinopathy detection and classification based on deep learning techniques using fundus images," *PeerJ Comput. Sci.*, vol. 10, p. e1947, 2024.

M. Vijayan and V. S., "A regression-based approach to diabetic retinopathy diagnosis using EfficientNet," *Diagnostics*, vol. 13, no. 774, pp. 1–16, 2023.

A. Bajwa, N. Nosheen, K. I. Talpur, and S. Akram, "A prospective study on diabetic retinopathy detection based on modify convolutional neural network using fundus images at Sindh Institute of Ophthalmology & Visual Sciences," *Diagnostics*, vol. 13, no. 393, pp. 1–11, 2023.

C.-L. Lin and K.-C. Wu, "Development of revised ResNet-50 for diabetic retinopathy detection," *BMC Bioinformatics*, vol. 24, no. 157, pp. 1–18, 2023.

Hazarika, I. (2022). *Digital transformation of the silk industry of Assam*. *Archives of Business Research, 10*(4), 110–119. https://doi.org/10.14738/abr.104.12261

A. ElMoufidi and H. Ammoun, "EfficientNetB3 architecture for diabetic retinopathy assessment using fundus images," Preprint, 2022.

S. Gothane, S. K. Raju, and N. Bhaskar, "Diabetic retinopathy detection using deep learning," in *Lecture Notes in Networks and Systems*, Springer, pp. 373–383, 2022.

A. B. Kadan and P. S. Subbian, "Diabetic retinopathy detection from fundus images using machine learning techniques: A review," *Wireless Personal Communications*, vol. 121, pp. 2225–2249, 2021.

M. Mohammedhasan and H. Uğuz, "A new early stage diabetic retinopathy diagnosis model using deep convolutional neural networks and principal component analysis," *Traitement du Signal*, vol. 37, no. 5, pp. 711–722, 2020.