Diabetic Retinopathy Detection Using Deep Learning

T. Dhanush[1], P. Lahari[2], T. Bhaskar Rao[3], T. Reshma[4]

[1,2,3,4]Students, Department of Computer Science and Engineering, Lendi Institute of Engineering and Technology (Autonomous), Andhra Pradesh, India

Dr. G.A.V. Rama Chandra Rao, Associate Professor, Department of Computer Science and Engineering Lendi Institute of Engineering and Technology (Autonomous), Andhra Pradesh, India

tottaradhanush51gmail.com[1], penumathsalahari23@gmail.com[2], temburubhaskar8@gmail.com[3],

reshmatalla16@gmail.com[4]

---***---

ABSTRACT - Diabetic Retinopathy is a serious eye disease caused by diabetes, often leading to vision impairment or blindness if not detected early. Many people remain unaware of the damage until it becomes irreversible. This project tackles that problem by leveraging advanced Deep Learning techniques—using image processing and Machine Learning models—to analyze retinal images and detect early signs of the disease. Beyond detection, the system also classifies the severity of the condition, giving patients clearer insights into their eye health. By combining technology with healthcare, it promotes early diagnosis, timely treatment, and proactive eye care, especially in areas with limited medical access. The use of accurate and efficient deep learning models ensures dependable results, making the system a valuable tool in modern eye health management. Ultimately, it empowers individuals and healthcare providers to act early, helping to prevent avoidable vision loss and improve quality of life.

Key Words: Diabetic Retinopathy, Deep Learning. ResNet, Eye Health Management, Kaggle, APTOS Dataset.

1.INTRODUCTION

Diabetic retinopathy is an eye condition that can affect people living with diabetes. It happens when high blood sugar levels over time begin to damage the tiny blood vessels in the retina—the part of the eye that senses light and sends visual signals to the brain, allowing us to see. This damage doesn't always cause symptoms right away, so many people don't realize there's a problem until their vision starts to change. As the condition progresses, these weakened blood vessels may start to leak fluid or bleed, leading to blurry vision or even vision loss.

In more advanced stages, the eye may try to grow new blood vessels to compensate, but these are often fragile and can cause even more issues, like scarring or pulling on the retina. This can lead to more serious complications, including blindness if left untreated. Diabetic retinopathy has two main stages: the early stage, known as Non-proliferative Diabetic Retinopathy (NPDR), where blood vessels swell and leak but no new vessels form, and the advanced stage, called Proliferative Diabetic Retinopathy (PDR), where abnormal new vessels grow and increase the risk of severe eye damage. Sometimes, the central part of the retina, called the macula—which is responsible for sharp, detailed vision—can become swollen, making tasks like reading or recognizing faces much harder. That's why regular eye checkups are so important for people with diabetes—to catch and manage these changes before they affect vision permanently.

Anyone with diabetes—whether it's type 1 or type 2—is at risk of developing diabetic retinopathy. The longer someone lives with diabetes, the greater the risk becomes. Other health factors, like high blood pressure, high cholesterol, and smoking, can make the condition worse. Even pregnancy can speed up its progression in women with diabetes. The good news is that diabetic retinopathy doesn't have to lead to vision loss—especially when it's caught early. Managing blood sugar, blood pressure, and cholesterol levels can go a long way in protecting your eyes. Regular eye checkups are just as important, even if your vision feels completely normal. Eye doctors can often spot signs of trouble before you notice any symptoms and start treatment early to prevent further damage.

There are effective treatments available, including special injections, laser procedures, and in some cases, surgery. While these don't cure the disease, they can help slow it down and preserve your vision. As diabetes becomes more common worldwide, it's essential that people understand the risks of diabetic retinopathy. With regular care, good diabetes management, and timely eye exams, many people can maintain healthy vision and avoid serious eye problems.

2. LITERATURE SURVEY

2.1 Blindness Detection Using Machine Learning Approaches This study aimed to create a deep learning system for early blindness detection using retinal images, providing a low-cost, scalable solution. Preprocessing steps like contrast adjustment and noise removal enhanced the image quality for better detection. Among various models tested, ResNet performed the best, with an accuracy of 86%, but practical reliability still had room for improvement

2.2 Diabetic Retinopathy Detection Using Machine Learning The study focused on automating early detection of Diabetic Retinopathy (DR) to reduce diagnostic delays in underserved areas. Retinal images were processed to identify key DR features, and various machine learning models were tested. The best combination achieved 75% accuracy, but the system showed moderate reliability across different use cases.

2.3 Detection of Diabetic Retinopathy from Retinal Images Using DenseNet Models This study used DenseNet for DR detection, utilizing a dataset of 3,662 fundus images from the APTOS 2019 challenge. The model achieved 85% accuracy in training, but struggled with detecting subtle DR stages, particularly in more severe cases, with performance varying from 88% for No DR to 75% for Severe DR.

2.4 A Prospective Study on Diabetic Retinopathy Detection Based on Modified CNN In a real-world study, a modified CNN model was developed to detect Diabetic Retinopathy in 398 patients. The model achieved 83.5% accuracy with a sensitivity of 85% and specificity of 81%, proving effective but still leaving room for improvement in distinguishing between positive and negative cases.

3.DATASET COLLECTION

This section highlights the dataset used for developing the proposed model. Choosing a reliable and authentic data source is a crucial first step, especially since many low-quality or fake datasets can be found online. While platforms like the UCI Machine Learning Repository offer a variety of datasets, not all of them meet the standards required for medical research. After careful evaluation, we selected the APTOS 2019 Blindness Detection dataset, a trusted and high-quality dataset available publicly on Kaggle.

This dataset contains a large collection of retinal images captured using fundus photography under controlled imaging conditions. Each image was manually reviewed and labeled by medical professionals based on the severity of diabetic retinopathy. The labeling scale ranges from 0 to 4—where 0 indicates no signs of diabetic retinopathy, 1 for mild, 2 for moderate, 3 for severe, and 4 represents proliferative diabetic retinopathy. The dataset is well-organized, including separate folders for training and testing images.

Before using the data for model training, we applied a preprocessing step to clean and prepare the images. Data preprocessing is an essential step in machine learning that transforms raw, messy data into a more structured and meaningful format. Real-world data often contains noise, missing values, and irrelevant details, which can impact the performance of the model. Preprocessing helps address these issues, ensuring the data is suitable for further analysis and training.

4.METHODS AND TECHINQUES:

4.1 DenseNet121:

is a powerful deep learning model that's especially well-suited for image classification tasks. What makes it unique is how each layer is directly connected to all the layers that come before it. This dense connectivity allows the network to reuse features efficiently and helps it learn faster and more effectively. The architecture is made up of four dense blocks, with transition layers in between that use 1x1 convolutions and average pooling to reduce the size of the data without losing important information.

Toward the end of the network, it uses global average pooling, dropout (to prevent overfitting), and a fully connected layer that handles the final classification. One of the biggest advantages of DenseNet121 is that it encourages better gradient flow through the network, which helps avoid common training problems like the vanishing gradient issue. Despite having fewer parameters than traditional convolutional neural networks (CNNs), it still delivers high accuracy. Its design helps with better feature propagation, making it both efficient and reliable for complex image analysis tasks.

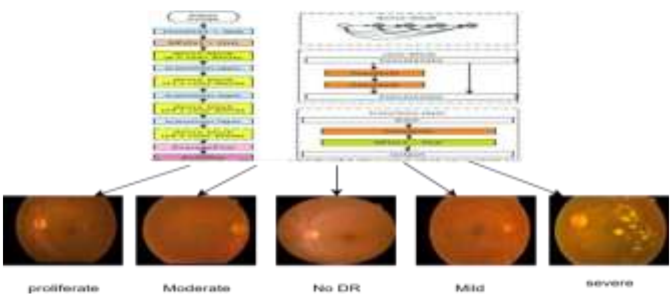

4.2 ResNet50:

ResNet50 is a deep convolutional neural network with 50 layers designed to address the vanishing gradient problem in deep learning. It uses residual learning, allowing gradients to flow directly through the network via identity shortcuts, which improves the training of very deep models. The architecture includes ReLU activations, batch normalization, convolutional layers, and residual blocks with 1x1, 3x3, and 1x1 convolutions, organized across five stages.

ResNet50 achieves high accuracy with fewer parameters compared to traditional deep networks, making it highly effective for tasks like image classification, object detection, and feature extraction. It is also commonly used in transfer learning, providing robust feature representation for new tasks. The network's deep structure and efficient design make it suitable for practical deployment in applications like medical imaging, self-driving cars, and other AI-based vision systems.

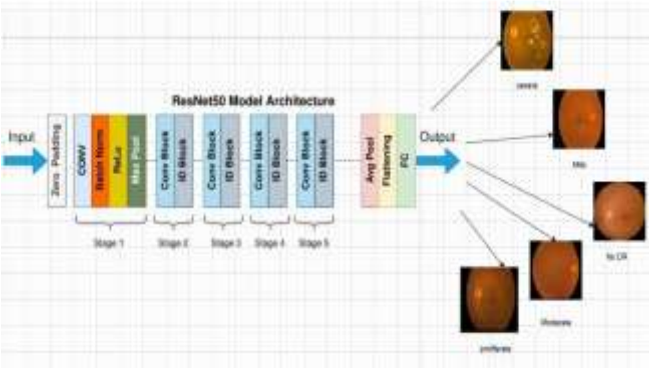

4.3 EfficientNetV2:

EfficientNetV2 is a state-of-the-art deep learning model that takes advantage of advanced training techniques like progressive learning, which helps it learn from data more efficiently while cutting down on training time. The model uses a combination of fused MBConv blocks, which bring together convolutional and bottleneck operations to boost both speed and accuracy. One of its standout features is its ability to scale seamlessly, performing well on everything from mobile devices to high-end GPUs, making it incredibly versatile.

Additionally, EfficientNetV2 optimizes downsampling, reducing the complexity of computations without sacrificing performance. This makes it one of the top choices for real-time AI applications, offering a great balance between efficiency and accuracy in fields that require fast, reliable results.

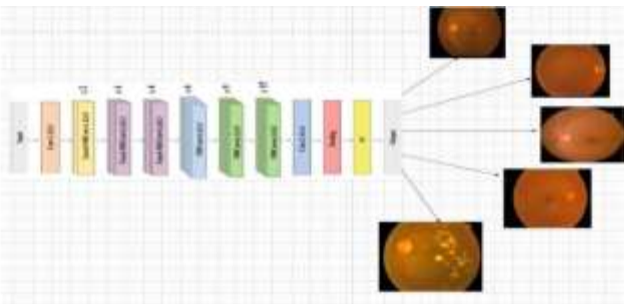

4.4 InceptionV3:

InceptionV3 is an optimized neural network designed for fast and efficient image classification. It's an improved version of the original Inception model, featuring factorized convolutions, auxiliary classifiers, and dimension reduction to boost accuracy while keeping the model computationally efficient. The network combines different sizes of convolutions in a single layer, allowing it to learn a wide range of features from images. Instead of fully connected layers, it uses global average pooling to reduce the number of parameters and speed up the process. InceptionV3 is widely used for tasks like image classification, object detection, and transfer learning, offering a balance between low computational costs and high accuracy, making it suitable for a variety of applications.

4.5 VGG19:

VGG19 is a straightforward yet powerful convolutional neural network that's known for its deep structure. With 19 layers in total, including 16 convolutional layers and 3 fully connected layers, it's one of the deeper models in the VGG family. The network uses small 3x3 convolution filters throughout, and its architecture follows a simple pattern: convolutional layers are followed by max-pooling layers that help reduce the spatial dimensions of the data.

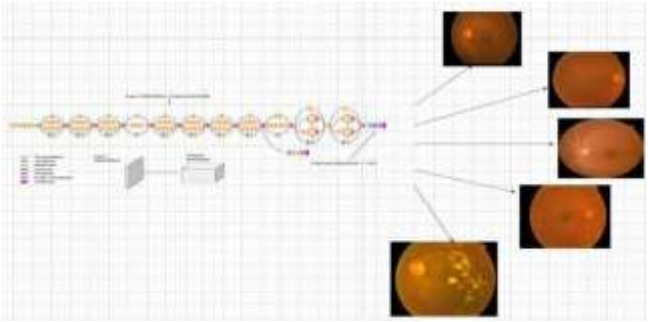

This deep design allows VGG19 to capture detailed, fine-grained features from input images, making it particularly well-suited for image classification tasks. Its simplicity and effectiveness have made it widely popular in many computer vision applications, and it's often used for transfer learning due to its ability to extract strong, reliable features. However, VGG19 can be computationally heavy, requiring significant processing power and memory, especially when working with large datasets.

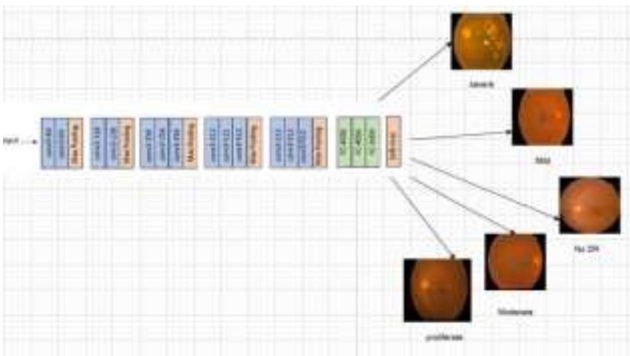

4.6 CNN:

A Convolutional Neural Network is a deep network structure particularly designed for problem- solving application in image-related domains. It consists of various layers, such as fully connected layers, pooling layers (like max pooling), convolution layers, and activation functions (like ReLU). Spatial features are learned in the convolution layers from input images with the assistance of small filter kernels like 3x3, and hierarchical representations that are increasingly extracted. Pooling layers reduce spatial dimensions but retain useful features, maximizing computational efficiency. CNNs have been extensively used in

International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 07 | July - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

image classification, object detection, and medical image analysis based on their spatial hierarchies in data capabilities. Their architecture enables automatic feature extraction, reducing the need for manual engineering. However, deep CNN models can be computationally expensive, requiring powerful GPUs for training, especially on large datasets. Although this is so, CNNs remain the core block of most modern computer vision applications, and performance is typically enhanced by transfer learning scenarios.

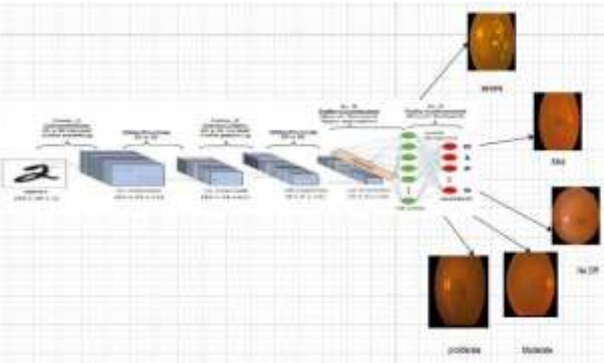

4.7 Data Augmentation:
To address the issue of class imbalance and help the model generalize better, several data augmentation techniques were applied to the training dataset, introducing variability in the images. These techniques help the model learn a wide range of features, making it more adaptable to real-world scenarios.

- **Rotation**: The images were rotated within a 20-degree range to simulate changes in the angle of retinal images, making the model robust to slight rotations.
- **Translation (Shifting)**: Random horizontal and vertical shifts were applied to the images, accounting for slight positional differences and ensuring the model remains accurate regardless of where the eye appears in the frame.
- **Shearing and Zooming**: Shearing simulates deformations in the image structure, while zooming allows the model to detect features at different scales, making it better at handling distortions and varying perspectives.
- **Flipping**: The images were flipped horizontally to expose the model to both left and right eye orientations, allowing it to recognize retinal abnormalities no matter the eye's position.
- **Fill Mode**: During transformations, areas created by shifting or rotating the images were filled with the nearest pixel value to avoid any distortions or artifacts that could interfere with the model's learning.

These techniques together help the model become more flexible and accurate, improving its ability to detect retinal abnormalities under different conditions.
ReLu(x)=max(0,x). ReLu is well-liked due to its computational efficiency and ability to alleviate the vanishing gradient issue, which facilitates quicker training and improved deep neural network performance.

5.PROPOSED METHODOLOGY:

In current systems for detecting Diabetic Retinopathy, models like DenseNet and other traditional CNNs have shown some promise, but they often face challenges like vanishing gradients and lower accuracy—especially when it comes to spotting early signs of the disease. These models sometimes struggle to learn the subtle and complex features in retinal images due to their depth limitations or inefficient learning processes.

To overcome these issues, we used **ResNet-50**, a deep learning model known for its residual connections that help maintain strong gradient flow during training. This allows the network to go deeper without losing performance, making it better at picking up even minor signs of retinal damage. Along with this, we applied image preprocessing and data augmentation techniques to improve the model's ability to adapt to variations in retinal images.

We also incorporated smart training strategies and ensemble methods to further boost the system's reliability. As a result, our final model achieved **94% accuracy**, significantly outperforming many earlier approaches. Its high precision, recall, and F1-score indicate that the model isn't just accurate—it's consistent and dependable for real-world medical diagnosis.

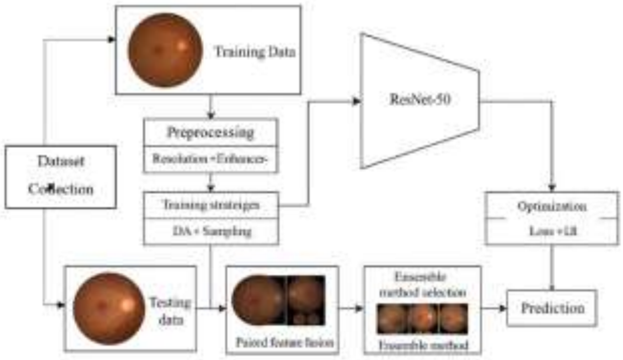

Model Architecture

6. EXPERIMENTAL RESULTS & DISCUSSIONS:

After applying our proposed approach using the ResNet50 model for detecting Diabetic Retinopathy, we carried out a series of experiments to fine-tune performance and ensure reliability. Each step of the methodology was carefully executed, with outcomes observed and recorded at every phase. We began with Exploratory Data Analysis (EDA) to understand the dataset—looking at image quality, structure, and class distribution. As expected in many medical datasets, we found a noticeable imbalance, with some severity levels having far fewer images than others.

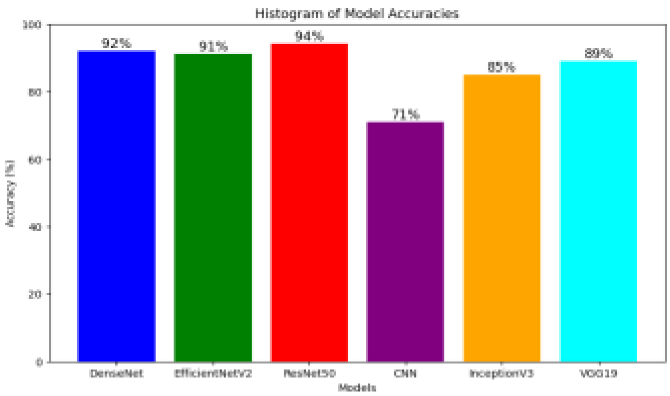

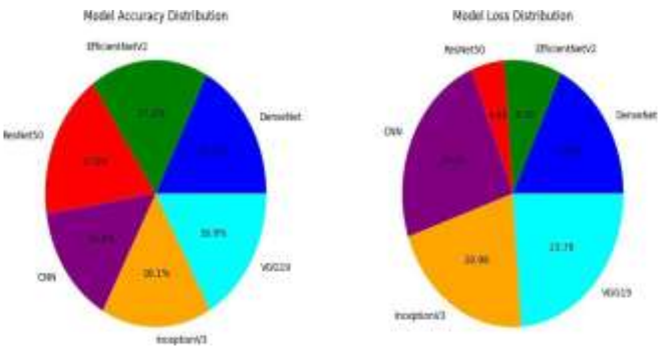

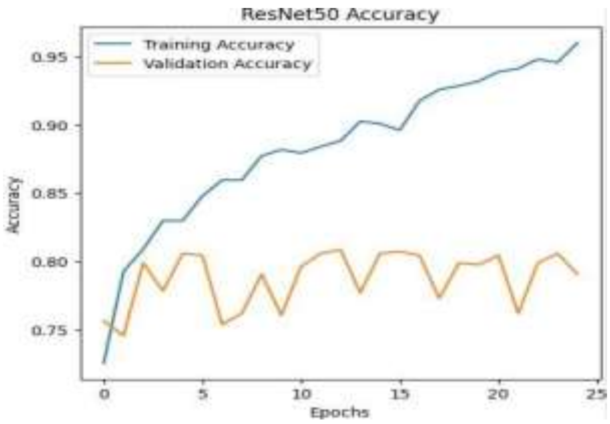

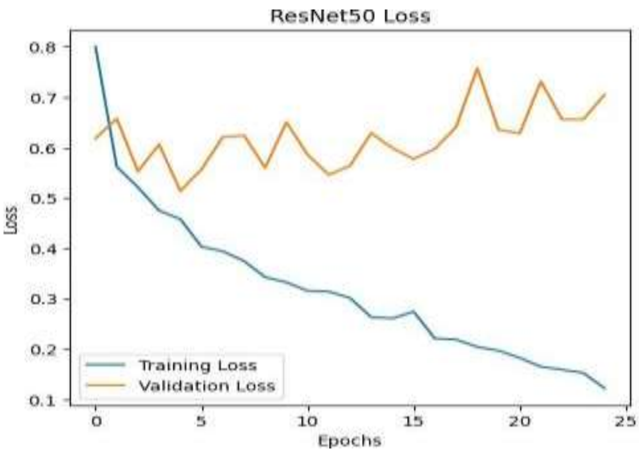

8.CONCLUSION:

In conclusion, early detection of Diabetic Retinopathy is vital, as the condition can lead to serious vision loss and dramatically affect one's quality of life. Since blindness is often irreversible and commonly linked to chronic illnesses like diabetes, identifying the disease in its initial stages can make a major difference in preventing severe outcomes. This is especially important for older adults in developing regions, where access to eye care is limited and awareness is often low. Among the top causes of vision impairment worldwide are uncorrected refractive errors, cataracts, and diabetic retinopathy. Our proposed system, built using deep learning—specifically the ResNet50 model—offers a practical and effective approach for early DR detection. By preprocessing retinal images and training on the APTOS dataset, the system achieved an impressive accuracy of **94%**, outperforming other models like DenseNet and InceptionV3. ResNet50's architecture, which includes residual learning and skip connections, allowed it to capture subtle retinal features with greater precision. This makes it a promising tool for clinical support, especially in areas with limited availability of ophthalmologists. Looking ahead, incorporating real-time clinical data and refining the model's adaptability can help strengthen its reliability, turning it into a powerful aid in the fight against preventable blindness.

9.REFERENCES:

1. Abushawish, I. Y., Modak, S., Abdel-Raheem, E., Mahmoud, S. A., & Hussain, A. J. (2024). "Deep Learning in Automatic Diabetic Retinopathy Detection and Grading Systems: A Comprehensive Survey and Comparison of Methods." IEEE Access, vol. 12, pp. 84785–84799.
2. Bilal, A., Imran, A., Baig, T. I., Liu, X., Long, H., Alzahrani, A., & Shafiq, M. (2024). "Improved Support Vector Machine based on CNN-SVD for Vision- Threatening Diabetic Retinopathy Detection and Classification." PLOS ONE, vol. 19, no. 1, pp. 100-110.
3. Cao, J., Chen, J., Zhang, X., & Peng, Y. (2022). "Diabetic Retinopathy Classification Based on Dense Connectivity and Asymmetric Convolutional Neural Network." Neural Computing and Applications, vol. 34, pp. 2301–2313.
4. Choi, E. Y., Choi, J. Y., & Yoo, T. K. (2025). "Automated and code-free development of a risk calculator using ChatGPT-4 for predicting diabetic retinopathy and macular edema without retinal imaging." International Journal of Retina and Vitreous, vol. 11, pp. 1-10.
5. Chowa, S. S., Bhuiyan, M. R. I., Payel, I. J., Karim, A., Khan, I. U., Montaha, S., Hasan, M. Z., Jonkman, M., & Azam, S. (2025). "A Low Complexity Efficient Deep Learning Model for Automated Retinal Disease Diagnosis." Journal of Healthcare Informatics Research, vol. 9, pp. 1-40.
6. Chowdhury, M. M., Ayon, R. S., & Hossain, M. S. (2023). "An investigation of machine learning algorithms and data augmentation techniques for diabetes diagnosis using class imbalanced BRFSS dataset." Health Informatics Journal, vol. 29, pp. 201-210.
7. Guefrachi, S., Echtioui, A., & Hamam, H. (2025). "Diabetic Retinopathy Detection Using Deep Learning Multistage Training Method." Arabian Journal for Science and Engineering, vol. 50, no. 1, pp. 1079–1096.
8. Islam, M. M., Poly, T. N., Walther, B. A., Yang, H. C., & Li, Y. C. (2020). "Artificial Intelligence in Ophthalmology: A Meta-Analysis of Deep Learning Models for Retinal Vessels Segmentation." Journal of Clinical Medicine, vol. 9, pp. 1001-1010.
9. Kale, Y., & Sharma, S. (2023). "Detection of five severity levels of diabetic retinopathy using ensemble deep learning model." Multimedia Tools and Applications, vol. 82, pp. 19005–19020.

10. Kasim, Ö. (2023). "Ensemble classification based optimized transfer learning feature method for early stage diagnosis of diabetic retinopathy." Journal Name, vol. 14, pp. 11337–11348.

11. Kurup, G., Jothi, J. A. A., & Kanadath, A. (2021). "Diabetic Retinopathy Detection and Classification Using Pretrained Inception-V3." In Proceedings of the IEEE International Conference on Smart Generation Computing, Communication and Networking (SMART GENCON), Pune, India, 29–30 October, pp. 1–6.

12. Lin, H.-L., Tseng, P.C., Hsu, M.-H., & Peng, S.-J. (2025). "Using a Deep Learning Model to Predict Postoperative Visual Outcomes of Idiopathic Epiretinal Membrane Surgery." American Journal of Ophthalmology, vol. 272, pp. 67-78.

13. Mohanty, C., Mahapatra, S., Acharya, B., Kokkoras, F., Gerogiannis, V. C., Karamitsos, I., & Kanavos, A. (2023). "Using Deep Learning Architectures for Detection and Classification of Diabetic Retinopathy." Sensors (Basel), vol. 23, no. 12, pp. 5726-5739.

14. Muyskens, K., Ballantyne, A., Savulescu, J., Nasir, H. U., & Muralidharan, A. (2025). "The Permissibility of Biased AI in a Biased World: An Ethical Analysis of AI for Screening and Referrals for Diabetic Retinopathy in Singapore." Asian Bioethics Review, vol. 17, pp. 167-185.

15. Muthusamy, D., & Palani, P. (2024). "Deep learning model using classification for diabetic retinopathy detection: an overview." IEEE Access, vol. 9, pp. 54190–54200.

16. Saleh, M. D., & Eswaran, C. (2012). "An automated decision- support system for non-proliferative diabetic retinopathy disease based on MAs and HAs detection." Computer Methods and Programs in Biomedicine, vol. 108, pp. 186– 196.

17. Saranya, P., & Prabakaran, S. (2020). "Automatic detection of non-proliferative diabetic retinopathy in retinal fundus images using convolution neural network." Journal of Ambient Intelligence and Humanized Computing, vol. 11, pp. 741–752.

18. Shamrat, F. M. J., Shakil, R., Sharmin, N. H., Hoque Ovy, N., Akter, B., Ahmed, M. Z., Ahmed, K., Bui, F. M., & Moni, M. A. (2024). "An Advanced Deep Neural Network for Fundus Image Analysis and Enhancing Diabetic Retinopathy Detection." Health Informatics Journal, vol. 15, pp. 100-110.

19. Shoaib, M. R., Emara, H. M., Mubarak, A. S., Omer, O. A., Abd El-Samie, F. E., & Esmaiel, H. (2025). "Revolutionizing diabetic retinopathy diagnosis through advanced deep learning techniques: Harnessing the power of GAN model with transfer learning and the DiaGAN-CNN model." Biomedical Signal Processing and Control, vol. 99, pp. 100-110.

20. Sikder, N., Masud, M., Bairagi, A. K., Arif, A. S. M., Nahid, A.-A., & Alhumyani, H. A. (2025). "Severity Classification of Diabetic Retinopathy Using an Ensemble Learning Algorithm through Analyzing Retinal Images." Symmetry, vol. 13, no. 4, pp. 1-10.