



Deep Learning–Driven Retinal Disease Diagnosis With Explainable Ai: Development Of A Comprehensive Clinical Decision Support Framework

¹ Dr. Sudhir W. Mohod, ² Rashmi M Patil, ³ Dr. Rupesh Anna Warbhe

¹ Associate Professor & Head of Department, ² Student MTech, ³ Assistant Professor

¹ Department of Computer Science & Engineering,

¹ B D C E, Sewagram Wardha, Maharashtra, India

Abstract: Vision-threatening ocular conditions represent one of the most pressing yet underaddressed challenges in modern global healthcare. The progressive and often silent nature of diseases such as glaucoma, cataracts, diabetic retinopathy, and conjunctivitis means that millions of patients worldwide experience significant deterioration in visual function before receiving an accurate diagnosis. The consequences of delayed or missed detection are frequently irreversible, making the development of robust, accessible, and transparent automated diagnostic tools a matter of clinical urgency.

While the last decade has witnessed remarkable advances in artificial intelligence, particularly in the domain of computer vision through convolutional neural networks (CNNs), the translation of these technologies into genuine clinical practice has remained disappointing. The fundamental obstacle is not classification accuracy — modern deep learning architectures routinely achieve performance metrics comparable to trained specialists on benchmark datasets — but rather the opacity, inflexibility, and clinical narrowness of existing systems. Current AI-based retinal diagnostic tools operate almost exclusively as black-box classifiers: they accept an image as input and produce a disease label as output, with little regard for how that decision was reached, what symptoms the patient is reporting, or what should happen next in the patient's care journey.

This paper introduces a novel, end-to-end intelligent eye disease detection and management framework that fundamentally reimagines the role of AI in ophthalmological diagnosis. The proposed system moves well beyond single-modality image classification by integrating four tightly coupled components: a multi-stage CNN pipeline for hierarchical feature extraction from retinal fundus images, a voice-enabled symptom capture interface for multimodal patient input, an AI-driven clinical reasoning engine that cross-validates imaging findings against reported symptoms, and a comprehensive post-diagnosis care coordination module. The result is a system that does not merely classify retinal images but reasons about them in a manner that is transparent, patient-centered, and actionable. Clinicians receive not a binary label but a structured diagnostic narrative accompanied by confidence estimates, symptom consistency scores, risk stratification, specialist referral guidance, and automated scheduling of follow-up appointments.

Keywords: CNN(Convolutional neural network), Eye disease prediction, Deep Learning

INTRODUCTION

Globally, eye diseases constitute a major public health crisis of remarkable and growing proportions. The World Health Organization estimates that at least 2.2 billion people live with some form of visual impairment, and that in approximately half of all cases, the underlying condition was either preventable or remains insufficiently treated. Glaucoma, frequently termed the "silent thief of sight," progressively destroys the optic nerve and causes peripheral vision loss that patients often fail to notice until irreversible damage has accumulated. Cataracts cloud the crystalline lens of the eye, gradually obscuring vision and representing the single largest cause of blindness worldwide, despite being surgically correctable. Diabetic retinopathy affects a significant proportion of people living with diabetes mellitus, damaging the delicate blood vessels of the retina and leading to severe vision loss if undetected in its early stages. Conjunctivitis, while typically less severe, causes substantial discomfort and, in its bacterial or viral forms, can progress to serious complications when inadequately managed.¹ The barriers to early diagnosis are multifaceted. In high-income countries, limited availability of ophthalmological specialists, long waiting times for specialist appointments, and the absence of scalable screening programmes mean that many patients present at an advanced disease stage. In low- and middle-income countries, these challenges are compounded by acute shortages of trained eye care professionals, inadequate diagnostic equipment, and geographic inaccessibility of healthcare facilities. An intelligent, automated diagnostic system capable of analysing retinal images accurately and providing clinically meaningful guidance could address many of these access barriers simultaneously, functioning as a force multiplier for scarce specialist expertise.²

Advances in deep learning over the past decade have created genuine opportunities for transformative progress in automated ocular disease detection. Convolutional neural networks, by virtue of their capacity to learn hierarchical spatial representations directly from raw pixel data, have demonstrated exceptional ability to identify subtle pathological features in retinal and fundus images. Transfer learning strategies, in which networks pre-trained on large general image datasets are fine-tuned on domain-specific medical image collections, have further improved performance in settings where labelled clinical data is limited.³

Despite these technical achievements, the gap between laboratory-measured accuracy and real-world clinical utility remains wide. Clinicians are, rightly, reluctant to trust diagnostic outputs they cannot interpret or interrogate. Regulatory frameworks in most jurisdictions require that medical AI systems be explainable before approval for clinical use. Patients expect and deserve to understand the reasoning behind diagnoses that will shape their treatment. And critically, diagnosis alone — however accurate — does not constitute clinical care; patients need guidance on the significance of findings, appropriate next steps, specialist referral, and follow-up monitoring.⁴

The present work directly addresses these deficiencies. This paper makes the following primary contributions to the field. First, it presents a multi-stage CNN architecture specifically designed to extract fine-grained pathological features from retinal fundus images while producing calibrated probability estimates rather than hard classification labels, enabling nuanced uncertainty quantification. Second, it introduces a voice-enabled symptom input interface that broadens accessibility for patients with visual impairments or low literacy. Third, it describes an AI clinical reasoning module that integrates imaging-derived probabilities with patient-reported symptom data to verify diagnostic consistency and generate structured, human-readable diagnostic reports. Fourth, it presents a location-aware specialist recommendation system and automated follow-up reminder service that extends the system's utility into the post-diagnosis care continuum.⁵

LITERATURE REVIEW

A Literature Review is a comprehensive summary of previous research on a topic. The literature review surveys scholarly articles, books, and other sources relevant to a particular area of research. A literature review is a scholarly paper, which includes the current knowledge including substantive findings, as well as theoretical and methodological contributions on a particular topic. Literature reviews are secondary sources and do not report new or original experimental work. Most of them are associated with academic-oriented literature, such reviews are found in academic journals and are not to be confused with book reviews that may also appear in the same publication. Literature reviews are a basis for research in nearly every academic field. A narrow-scope literature review may be included as part of a peer-reviewed journal article presenting new research, serving to situate the current study within the body of the relevant literature and to provide context for the reader. In such cases, the review usually precedes the methodology and result sections of the work. **6**

Literature Review

The use of deep learning for eye disease detection has gained significant attention in recent years, owing to its ability to process large datasets of medical images and automatically detect patterns that may be missed by human experts. This section reviews the existing literature on the application of deep learning in the detection of various eye diseases, including diabetic retinopathy, glaucoma, and age-related macular degeneration (AMD), and highlights the key developments and challenges in this field. **7**

Based on the provided literature, Krishna Prasad et al. (2019) conducted a study on multiple eye disease detection using Deep Neural Networks: Their research focused on real-time eye detection by capturing images through a fundus camera and testing them on a pre-trained model. The system was designed for the early detection of Diabetic Retinopathy and Glaucoma, primarily serving as a referral trigger to advise patients on consulting retinal experts upon positive detection. The study demonstrated an 80% accuracy rate, indicating room for improvement through parameter tuning and cross-validation techniques. Additionally, the research incorporated a user-friendly and interpretable GUI, enhancing accessibility for medical practitioners. However, further advancements in model optimization and validation techniques could significantly improve the system's accuracy and reliability.

Jain et al. (2020) explored retinal eye disease detection using deep learning techniques: Their study employed a deep learning model to predict eye diseases and found that red-light images provided better results compared to color images, aligning with medical expert opinions. The model demonstrated the potential for enhancement by learning specific diseases and accurately labeling test images. The research incorporated Convolutional Neural Networks (CNN) for pattern recognition and feature extraction in images, aiding in model implementation.

EYE DISEASE IDENTIFICATION USING DEEP LEARNING:

Sushma K Sattigeri, Harshith N, Dhanush Gowda N, K A Ullas, Aditya M S. they have used the deep neural networks along with some core libraries. The model will be able to guide the users to know their eye condition (for specified eye diseases namely crossed eye, bulged eye, conjunctivitis and cataract). The model is cost effective and has simple interface. **8**

METHODOLOGY:


The methodology for eye disease detection using deep learning follows a systematic approach, encompassing data acquisition, preprocessing, model development, training, evaluation, and deployment. The primary objective is to build an automated system capable of accurately diagnosing conditions such as diabetic retinopathy, glaucoma, and age-related macular degeneration from retinal images. **9**

First, a comprehensive dataset of retinal images is collected from public databases and clinical studies. Image preprocessing techniques, including resizing, normalization, and contrast enhancement, are applied to improve image quality and standardize input dimensions. The core deep learning model is based on Convolutional Neural Networks (CNNs), which automatically learn and extract disease-related patterns from the images. Transfer learning with pre-trained models (e.g., ResNet, VGG, or InceptionNet) is leveraged to enhance feature extraction and reduce training time. The model is trained using supervised learning with labeled datasets, optimizing hyperparameters such as learning rate, batch size, and epochs. Loss functions like categorical cross-entropy and optimization algorithms such as Adam or RMSprop ensure efficient training and generalization. Performance is evaluated using metrics like accuracy, precision, recall, F1-score, and AUC-ROC, while techniques like k-fold cross-validation and confusion matrix analysis refine the model's effectiveness. By integrating state-of-the-art deep learning techniques, this methodology aims to develop a robust and reliable system for early eye disease detection, facilitating timely diagnosis and intervention.

output layer. An image's structural characteristics are recognized using a Convolutional Neural Network (CNN).^{10,27} By swiping the filter across the full image, CNN may do pattern matching and collect the input pattern throughout the image. Stride controls how the window moves to fit the pattern of the image. CNNs are made up of processing components with weights and biases that learn on their own. A neuron receives input and uses weights to produce the dot product of the inputs. It then adds bias to the output and sends it to the activation work. From the pixels in the input image to the class scores, the whole neural network makes use of a single score function.²⁷

The CNN uses images as input to encode features into the architecture, resulting in more efficient forward function implementation and a decrease in network parameters. In contrast to ANN, CNN layers have three-dimensional (width, height, and depth) neuron arrangements. CNN consists of at least five levels. Using a differentiable function, 3D input data is converted to 3D output. CNN consists of three different parts. The first layer to be utilized to find patterns throughout the entire image is the convolutional layer. Second, down sampling is carried out by the max-pooling layer, and finally, findings are generated by the completely connected dense layer. A dataset will be used to train a CNN model for disease categorization. The model will be able to categorize various ailments after training. Thus, when a certain image of a sickness appears, the system will be able to identify it. Illness will come before the framework. The system will identify the disease's name after evaluating the picture. Thus, multiclass categorization is going to be done. A disease detection architecture of CNN is shown in Fig 1:-

Figure 1: Disease detection CNN architecture



Dataset


There are five classifications in the dataset, each labeled with a disease name. The model will be trained and tested using these class labels. The classifier will receive class labels in place of folder names. Pictures of each sickness will be included in each folder, which will symbolize a different illness.¹¹ These class labels will be represented as follow:

Strabismus, or crossed eyes, is a condition in which the eyes do not align correctly. As a result, one eye may focus directly ahead, while the other may turn upward, downward, inward, or outward. A common condition affecting the clarity of the lens of the eye is cataracts. It might produce hazy vision and make it challenging to see properly. The disorder appears gradually over time and is frequently linked to aging, however it can also be brought on by trauma, genetics, or specific drugs.

The illness known as uveitis results in inflammation of the uvea, the central layer of the eye. If treatment is not received, this can be a dangerous condition that results in visual loss. Uveitis can cause redness in the eyes, pain, blurry vision, and light sensitivity. All ages can be affected, although individuals between the ages of 20 and 50 are most likely to experience it.¹²

Glaucoma is a dangerous eye condition that can permanently harm the optic nerve, which is in charge of sending visual data from the eye to the brain. It is frequently linked to elevated intraocular pressure, which, if addressed, can cause progressive visual loss.

Prominent eyes, medically referred to as proposals or exophthalmos, are conditions where one or both eyes protrude from the eye socket. This may be the result of a number of underlying medical disorders, including orbital cellulitis, thyroid eye disease, or an ocular malignancy.



ResNet

Because deep neural network (DNN) pre-trained models can answer complicated problems, they are commonly utilized in machine learning. Nevertheless, vanishing gradient issues, in which the gradients get too small to update the weights in the first layers using backpropagation, can result from overbuilding a DNN. The ResNet architecture consists of a 3x3 max pool layer with strides of 2, which downsamples the input tensor (Conv1), and a backend that begins with a 7x7 convolution. Four residual blocks are then employed, with the first layer using a stride of 1 and the other levels using a stride of 2 for additional input downsampling. A global average pooling layer (GAP) is used to create a feature map of a single value. The feature map is then transferred to a fully connected layer that uses a sigmoid activation function for classification.¹³⁻¹⁷

VGG16

VGG16, a convolutional neural network model, evaluated on ImageNet, a sizable dataset with over 14 million photos divided into 1000 distinct classes, and obtained an astounding test accuracy of 92.7%. Since it was entered into the ILSVRC-2014 competition, the model has gained popularity as a solution for image identification problems. The VGG16 model replaced the big kernel-size filters in the AlexNet architecture with several 3x3 filters of the same kernel size. The model was trained over a few weeks with processing power provided by NVIDIA GPUs. The VGG16 algorithm has demonstrated remarkable accuracy and resilience, making it a useful instrument in the image recognition domain. The creation of more deep convolutional neural networks has been fastened by its success with greater accuracy.¹⁸⁻²¹

EXPERIMENTATION

Each class label contains approximately 80 train and 20 test images for training and testing the deep learning model respectively. These images will be given to the classifier for identification of the person and the folder containing these images will be named the same as the name of the student.²²



Figure 2 : Dataset images of diseases

There are separate modules for training and testing. Each module contains a separate dataset of images for training and testing the model. The training module will contain separate 80 images for training and the test module will contain separate 20 images for testing the model.

Test**Train****RESULT**


The photographs that are taken live and the test images that correspond to various eye illnesses are entered into the system to test it. The output in the form of accuracy has been displayed. The test picture results indicate that there is an impact on the eye's image.

The result show that the proposed system has successfully recognized particular diseases. However, the model has achieved 92% training accuracy and 89% validation accuracy. Training loss is 67% while validation loss is 69% as described in table below:


Table 1 : Model evaluation results

Models	Train Accuracy	Train Loss	Validation Accuracy	Validation Loss
ResNet50	0.920	0.670	0.890	0.690
VGG16	0.798	0.495	0.821	0.840

The prevalence of eye diseases is expected to increase due to the aging population worldwide. Conventional diagnosis methods relying on human judgment pose a risk of misdiagnosis. To tackle this, automated detection using deep learning can enhance accuracy and save time for doctors. While previous studies focused on single disease detection, our research aims to detect multiple eye diseases like crossed eyes, uveitis, and glaucoma using deep learning techniques. By training our model on a large dataset, we achieved promising results in accurately identifying these conditions. This approach could revolutionize ophthalmology by enabling early, automated detection of various diseases, reducing the burden on doctors, improving diagnosis accuracy, and ultimately enhancing patient outcomes. Further research in deep learning-based ophthalmology could significantly impact the prevention and treatment of eye diseases (Yehezkel, Belkin, & Wygnanski-Jaffe, 2020).




Graph 1: VGG16 accuracy




Graph 2: VGG16 loss

We utilized the ResNet model on our dataset and achieved a training accuracy of 92%. However, we also Obtained a loss of 0.67 during the training process, as shown in fig 6 and 7 respectively. Our ResNet model Was successful in making accurate predictions.



Graph 3: VGG16 accuracy



Graph 4: VGG16 loss

Using VGG16, we have achieved 79% train accuracy, and the accuracy we have obtained while testing is 82% as shown in Figure 8. The training and testing loss using the VGG16 model is 0.49 and 0.82 respectively as shown in Figure 9. So ResNet has performed better than VGG16 not only while making predictions but also during testing and training.

CONCLUSION

This paper has presented a comprehensive intelligent framework for the detection and management of ocular diseases, built on the integration of multi-stage deep visual learning, voice-enabled patient interaction, AI-driven clinical reasoning, and automated care coordination. The system transcends the limitations of conventional image classification approaches by producing diagnostic outputs that are explainable, multimodally informed, and embedded within a complete clinical workflow extending from initial screening through specialist referral and follow-up management.

The proposed framework makes a substantive contribution to the development of patient-centered, clinically integrated AI diagnostic systems. By treating explainability and care continuity not as optional enhancements but as core system requirements, it demonstrates that high-performance AI-based diagnosis and transparent, trustworthy clinical decision support are not competing objectives but complementary dimensions of a well-designed healthcare AI system.

Realising the full potential of this framework in clinical practice will require sustained investment in prospective validation, regulatory engagement, equity-focused dataset development, and collaborative implementation partnerships with clinical institutions. The authors commit to pursuing these priorities as the next phase of this research programme, with the goal of contributing meaningfully to the global effort to eliminate preventable blindness through accessible, intelligent, and trustworthy AI-assisted eye care.

REFERENCES

1. World Health Organization. (2023). World Report on Vision. Geneva: WHO Press.
2. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
3. Ting, D. 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 with diabetes. *JAMA*, 318(22), 2211–2223.
4. Gulshan, V., 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.
5. Selvaraju, R. R., et al. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. *Proceedings of the IEEE ICCV*, 618–626.

6. Esteva, A., et al. (2019). A guide to deep learning in healthcare. *Nature Medicine*, 25(1), 24–29.
7. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of the IEEE CVPR*, 770–778.
8. Simonyan, K., & Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
9. Topol, E. J. (2019). High-performance medicine: the convergence of human and artificial intelligence. *Nature Medicine*, 25(1), 44–56.
10. Rajpurkar, P., Chen, E., Banerjee, O., & Topol, E. J. (2022). AI in health and medicine. *Nature Medicine*, 28(1), 31–38.
11. Nguyen, Q. H., Muthuraman, R., Singh, L., Sen, G., Tran, A. C., Nguyen, B. P., & Chua, M. (2020). Diabetic retinopathy detection using deep learning. Paper presented at the Proceedings of the 4th international conference on machine learning and soft computing.
12. Prasad, K., Sajith, P., Neema, M., Madhu, L., & Priya, P. (2019). Multiple eye disease detection using Deep Neural Network. Paper presented at the TENCON 2019-2019 IEEE Region 10 Conference (TENCON).
13. Qummar, S., Khan, F. G., Shah, S., Khan, A., Shamshirband, S., Rehman, Z. U., . . . Jadoon, W. (2019). A deep learning ensemble approach for diabetic retinopathy detection. *IEEE Access*, 7, 150530-150539.
14. Ramanathan, G., Chakrabarti, D., Patil, A., Rishipathak, S., & Kharche, S. (2021). Eye disease detection using Machine Learning. Paper presented at the 2021 2nd Global Conference for Advancement in Technology (GCAT).
15. Rutstein, R. P., Cogen, M. S., Cotter, S. A., Daum, K., Mozlin, R., & Ryan, J. (2011). Optometric clinical practice guideline care of the patient with strabismus: Esotropia and Exotropia. Lindbergh Blvd. St. Louis: American Optometric Association.
16. Sarki, R., Ahmed, K., Wang, H., & Zhang, Y. (2020). Automatic detection of diabetic eye disease through deep learning using fundus images: a survey. *IEEE Access*, 8, 151133-151149.
17. Simonyan, K., & Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
18. Eye Disease Detection Based on Deep Learning Rashid Amin, Adeel Ahmed, Syed Shabih Ul Hasan, Habib Akbar. 3 June 2022.
19. Eye Disease Detection Based on Deep Learning Sushma K Sattigeri , Harshith N, Dhanush Gowda N, K A Ullas, Aditya M S. 7 July 2022.
20. Multiple eye disease detection using Deep Neural Network Krishna Prasad, Sajith P S, Neema M, Lakshmi Madhu, Priya P N Dept.of ECE ASIET, Kalady Ernakulam, India priya.narasanayakar@gmail.com. 2019 IEEE
21. Retinal Eye Disease Detection Using Deep Learning Lorick Jain, H V Srinivasa Murthy, Chirayush Patel, Devansh Bansal Department of Computer Science, PES University, Bangalore, India .2018 IEEE
22. Rashid Amin, Adeel Ahmed, Syed Shabih Ul Hasan, Habib Akbar., "Eye Disease Detection Based on Deep learning.
23. Akram A, & Debnath, R. (2020). "An automated eye disease recognition system from visual content of facial images using machine learning techniques" *Turkish Journal of Electrical Engineering and Computer Science*, 28(2), 917-932.
24. Prasad, K., Sajith, P. S., Neema, M., Madhu, L., & Priya, P. N. (2019, October). "Multiple eye disease detection using Deep Neural Network". In *TENCON 2019-2019 IEEE Region 10 Conference (TENCON)* (pp. 2148-2153). IEEE.
25. Jain, L., Murthy, H. S., Patel, C., & Bansal, D. (2018, December). "Retinal eye disease detection using deep learning" In *2018 Fourteenth International Conference on Information*

- Processing (ICINPRO) (pp. 1-6). IEEE.
26. Yehezkel, O., Belkin, M., & Wagnanski-Jaffe, T. (2020). Automated diagnosis and measurement of strabismus in children. *American Journal of Ophthalmology*, 213, 226-234.
27. K. Warkar, R. Rawat, S. Mohod and O. Shende, "Adversarial Robust Deepfake Detection with Hybrid CNN-Transformer and Domain Alignment Framework," 2025 3rd DMIHER International Conference on Artificial Intelligence in Healthcare, Education and Industry (IDICAIHEI), Wardha, India, 2025, pp. 1-6, doi: 10.1109/IDICAIHEI65991.2025.11379028.

