JETIR.ORG ISSN: 2349-5162 | ESTD Year : 2014 | Monthly Issue
JOURNAL OF EMERGING TECHNOLOGIES AND INNOVATIVE RESEARCH (JETIR)
An International Scholarly Open Access, Peer-reviewed, Refereed Journal

CARDIOVASCULAR DISEASE PREDICTION FROM RETINAL IMAGES USING DEEP LEARNING

Dr.S.Saravanan
Associate Professor
Dept. Of Ece,N.B.K.R Institute Of Science & Technology (Autonomous), Vidyanagar.

P.Padma, C.Madhuri, K.Mahendra, B.V.Ramanamma
UG scholar
Dept. Of Ece,N.B.K.R Institute Of Science & Technology (Autonomous), Vidyanagar.

Abstract- CVDs, in fact, are still among the leading causes of death in the world. Emerging scientific evidence suggests that retinal abnormalities may be used as biomarkers for systemic diseases like CVDs. Herein, we propose a deep learning methodology for diagnosis using CNN with MobileNet as a backbone to classify retinal images into six categories: Age-Related Macular Degeneration (ARMD), Diabetic Neuropathy (DN), Diabetic Retinopathy (DR), Macular Hole (MH), Optic Disc Cupping (ODC), and Normal. These conditions carry implications for the systemic health of the patient, somewhat like CVDs. The network was trained with large heterogeneous retinal fundus image datasets, which included preprocessing steps such as resizing, normalization, and data augmentation. Owing to MobileNet's lightweight architecture, it has a good potential for easy and scalable deployments, particularly in resource-limited clinical environments. All model performances and evaluations are assessed using accuracy, precision, recall, F1 score, and confusion matrix metrics. Results show a powerful potential for the model to separate different retinal diseases from the normal state, for risk calculation, and for the purposes of early detection strategies. Future aspects will include clinical validation and the integration of the findings into diagnostic pathways to allow full-scale non-invasive patient screening with the assistance of health associates.

Keywords: Retinal images, deep learning, MobileNet, cardiovascular diseases, CVD prediction, early detection, machine learning.

I. INTRODUCTION

Given the unprecedented rise in the incidence of cardiovascular diseases or CVDs, there are ever increasing demands for easy-to-use yet efficient and accurate methods of diagnosis. Among the various innovations made in this direction, retinal imaging proves to be a non-invasive modality of immense potential in assessing systemic diseases, including cardiovascular health, because of the anatomical and functional relationship between the retinal microvasculature and systemic circulation. Retinal images are indicators of diabetic retinopathy that are closely associated with optic disc cupping-an indicator for forms of age-related macular degeneration associated with early cardiovascular and metabolic disorders.

Nonetheless, manual interpretation of retinal images by specialists is extremely time-consuming, subjective, and highly variable. To resolve this obstacle, the project proposes a diagnostic system that used deep learning and convolutional neural networks (CNNs)-specifically that of the MobileNet architecture-for automated analysis of retinal images and classifying six categories, namely, Age-Related Macular Degeneration (ARMD), Diabetic Neuropathy (DN), Diabetic Retintraphy (DR), Macular Hole (MH), Optic Disc Cupping (ODC), and Normal.

The system employs a robust pre-processing pipeline—comprising image resizing, normalization, and augmentation—to enhance data quality and model generalization. The MobileNet model, known for its efficiency and lightweight design, has been fine-tuned and extended with task-specific layers for multi-class image classification. This architecture strikes a balance between accuracy and computational efficiency, making it suitable for deployment in real-world, resource-constrained environments such as primary care centers or remote clinics.

To enhance usability, a web-based interface has been developed using HTML, CSS, and JavaScript, with Flask as the backend. The application allows users to upload retinal images and instantly receive classification results. This facilitates quick screening and early detection of conditions potentially linked to cardiovascular risks, aiding healthcare professionals in making timely and informed decisions.

By integrating deep learning with medical imaging and a user-friendly interface, this project offers a scalable, cost-effective, and accessible solution for retinal disease diagnosis. Future work may involve expanding the model to include additional retinal conditions, integrating real-time image acquisition from fundus cameras, and conducting clinical validations to assess real-world effectiveness.

A. *Objective Of The Study:*

This project is designed to offer a practical, affordable, and accessible solution for the early detection of retinal diseases that may serve as indicators of systemic conditions such as cardiovascular diseases (CVDs). By harnessing the power of deep learning—particularly convolutional neural networks (CNNs) integrated with the efficient MobileNet architecture—this study aims to automate the process of retinal image classification and provide a reliable alternative to traditional, time-consuming diagnostic methods.

The primary objectives include the collection and preprocessing of retinal image data, designing and training a CNN-based model, and thoroughly evaluating its performance on a multi-class classification task. The model is built to distinguish between six retinal classes: Age-Related Macular Degeneration (ARMD), Diabetic Neuropathy (DN), Diabetic Retinopathy (DR), Macular Hole (MH), Optic Disc Cupping (ODC), and Normal. These classes represent a diverse range of retinal health conditions, many of which are directly or indirectly linked to systemic health risks, including CVDs.

MobileNet is chosen as the core model due to its lightweight design and high classification performance, especially when deployed in environments with limited computational resources. As a depthwise separable CNN architecture, MobileNet significantly reduces the number of parameters while maintaining robust feature extraction capabilities—making it ideal for scalable deployment in clinics, rural health centers, or mobile health units.

Unlike approaches that rely on heavyweight ensemble models or multiple neural networks, this project focuses on optimizing a single MobileNet-based model to demonstrate its standalone efficiency. Through techniques such as data augmentation, learning rate tuning, and transfer learning, the model is tailored to generalize well across various retinal image types and conditions.

To ensure ease of use and wide accessibility, a web-based interface has been developed using HTML, CSS, and JavaScript, with Flask serving as the backend. This interface enables healthcare professionals or technicians to upload retinal images and receive real-time diagnostic predictions, eliminating the need for manual interpretation and allowing for faster medical decision-making.

Ultimately, this study aims to showcase the potential of deep learning in medical imaging, particularly for early disease detection. It advocates for integrating AI-driven tools into routine clinical practice to support proactive healthcare strategies, improve diagnostic accuracy, and extend advanced diagnostic capabilities to underserved regions.

B. Problem statement:

Every year across the world, cardiovascular diseases cause illness and death, and millions die from heart and vessel complications. Early detection and treatment can make the patient healthier and benefit the economy and health by decreasing the burden that cardiovascular diseases impose. Traditional diagnostic methods such as echocardiography, angiography, blood tests, and so on are very expensive, invasive, and time-consuming and are usually not practical in low-developed places. These are the major barriers in selecting possible candidates for massive early-stage preventive-care screenings.

Recent developments in technology, however, have given hope in medical imaging for non-invasive assessment of cardiovascular as well as systemic health aspects through retina imaging. The reason is that retinas are highly vascularized tissues. This means that microvascular changes are an early clue to systemic diseases such as hypertension and cardiovascular diseases (CVD) at their inception. Unfortunately, interpreting these images calls for more highly qualified technical ophthalmology specialists, and specialized equipment, which may not be easily available, especially in less economically developed or rural areas.

Another major problem is that there is a huge scope for observer variability and subtlety in retinal features associated with different pathological conditions. Differences between DR, ARMD, or ODC are more likely subjective and very difficult to achieve with certainty by manual methods. This leads to a scenario of poor, subjective diagnostic accuracies, which becomes critical where time is of the essence for management. Besides all these aspects, a static rule-based/manual setup cannot be scaled up in screening, and it is usually poorer at adapting to variation in imaging conditions or shrugging off different patient profiles.

This problem will be addressed through the design, development and implementation of an AI based retinal image analysis system based on the use of convolutional neural networks (CNN) and specifically with the MobileNet architecture because of its lightweight design, efficiency as well as high-performance score. The system shall be trained with a diverse dataset of retinal fundus images grouped in six categories: ARMD, DN, DR, MH, ODC, and Normal. This will allow the model to automatically learn how to find key visual patterns and anomalies related to various retinal diseases that may correlate with systemic risks including CVD.

Unlike conventional methods, deep learning systems are specifically designed to learn from the representative data and evolve at a later stage with performance improvements while generalizing across variances in image quality and demographic variations of patients. The model has been incorporated into an entire web-based platform built throughout HTML, CSS, JavaScript and Flask, allowing upload of retinal images by targeted users such as healthcare professionals or technicians and immediate response with diagnostic feedback. All of these components put together make possible the now-time real-time, remote, and scalable .

II. RELATED WORK

1. Evolution of Medical Image Analysis Using Deep Learning

In recent years, deep learning has transformed the field of medical image analysis by enabling automated and accurate detection of complex diseases. Traditionally, diagnosis from retinal images required manual interpretation by ophthalmologists, which was time-consuming, subjective, and prone to inter-observer variability. Early computer-aided diagnostic systems relied on handcrafted feature extraction and classical machine learning models like Support Vector Machines (SVM) or k-Nearest Neighbors (k-NN). While these approaches provided some level of automation, their performance was often limited by the quality and consistency of manually engineered features [1].

The advent of deep convolutional neural networks (CNNs) marked a major turning point, offering the ability to automatically learn hierarchical feature representations directly from pixel data. Architectures like AlexNet, VGG, and ResNet began to outperform traditional models in tasks such as diabetic retinopathy detection, optic disc segmentation, and macular disease classification. More recently, lightweight networks like MobileNet have emerged, offering high accuracy with significantly reduced computational requirements, making them ideal for deployment in real-time and resource-constrained environments such as primary healthcare facilities or mobile devices [2].

2. Benefits of Using Deep Learning for Retinal Disease Detection

Deep learning models offer significant advantages over traditional rule-based and classical machine learning systems. CNNs can automatically extract complex patterns from retinal fundus images that are difficult for humans to quantify, improving diagnostic speed and reducing the need for manual feature engineering. These models generalize well across varied imaging conditions, including differences in lighting, camera quality, and patient demographics [3].

Specifically, MobileNet—a lightweight CNN designed for efficiency without compromising accuracy—has proven effective in medical image classification tasks, particularly where real-time inference and low-resource deployment are essential. Its depthwise separable convolutions drastically reduce the number of trainable parameters, allowing for faster training and inference. Combined with data augmentation and transfer learning, MobileNet models can achieve performance comparable to heavier networks while remaining practical for on-device processing or web-based deployment.

Another key benefit is scalability: deep learning models can be trained on large, diverse datasets and deployed widely through web-based platforms [4]. These systems can assist medical professionals by providing second opinions, flagging abnormalities for further inspection, and serving as screening tools in remote or underserved areas where ophthalmic expertise is limited.

3. Challenges in Implementing Deep Learning for Retinal Image Classification

Despite its potential, the application of deep learning in medical diagnostics faces several challenges. One of the primary issues is class imbalance—retinal disease datasets often contain many more normal samples than pathological ones, which can bias the model and reduce sensitivity to rare conditions [5].Strategies such as oversampling, data augmentation, and class-weighted loss functions are often necessary to mitigate this issue.

Another significant challenge is data quality and diversity. Retinal images can vary based on the device used, acquisition settings, and patient factors. Poor image quality due to blur, illumination variance, or occlusion can hinder model performance. Robust preprocessing pipelines and domain adaptation techniques are critical to achieving reliable results in diverse clinical environments.

Furthermore, model interpretability remains a concern. Healthcare professionals need transparency in predictions, especially when they guide medical decisions. Techniques such as Grad-CAM and saliency maps help visualize model focus areas, increasing trust and facilitating clinical validation.

Lastly, regulatory approval and integration into clinical workflows are non-trivial. Medical AI systems must undergo rigorous evaluation, external validation, and compliance checks to ensure safety and reliability before being used in real-world healthcare settings [6].

4. Technological Advancements in AI-Based Retinal Diagnostics

Recent technological advancements have significantly improved the accuracy, robustness, and usability of deep learning-based retinal diagnostic systems. MobileNet and other efficient architectures have enabled deployment of real-time models on mobile devices, aiding field diagnostics and telemedicine applications. Transfer learning has reduced the need for large domain-specific datasets by leveraging pre-trained weights from models trained on general image datasets like ImageNet [7].

Innovations in explainable AI (XAI) have made it possible to interpret the decisions made by deep learning models, promoting transparency and user trust. Integration with cloud-based platforms allows for centralized model training while maintaining decentralized deployment. Furthermore, frameworks like Flask, TensorFlow.js, and PyTorch now support seamless backend and frontend integration, allowing medical AI tools to be embedded in interactive web applications for real-time inference and visualization [8].

Advanced data augmentation strategies and generative models (e.g., GANs) are being used to synthetically increase dataset diversity, helping models generalize better. Moreover, federated learning is gaining momentum in healthcare, allowing multiple institutions to collaboratively train models without compromising patient data privacy [9].

5. Case Studies and Real-World Implementations

Several studies and real-world implementations have demonstrated the success of CNNs in retinal image classification. Google's DeepMind, for instance, developed a deep learning model capable of detecting over 50 eye diseases from OCT scans with performance comparable to expert ophthalmologists. Another study applied MobileNet for diabetic retinopathy detection and achieved over 90% accuracy while maintaining lightweight model size, enabling deployment on smartphones and point-of-care devices.

Public datasets like EyePACS, MESSIDOR, and APTOS have been instrumental in training and benchmarking retinal disease classification models. These datasets have enabled researchers to create standardized benchmarks and drive improvements in model accuracy and generalization.

In clinical settings, AI-driven screening tools have been piloted in community clinics, demonstrating their ability to reduce diagnostic workload and facilitate early referrals for treatment. These tools have shown promise in increasing access to eye care in rural and underserved regions, where ophthalmic expertise is limited or absent.

Such successful implementations validate the potential of deep learning systems as scalable, efficient, and accurate diagnostic tools that can complement human expertise and revolutionize early disease detection and management in healthcare [10].

III. PROPOSED SYSTEM

The proposed system leverages deep learning to automatically analyze retinal fundus images and classify them into multiple disease categories. This approach is designed to support early detection and screening of conditions often associated with systemic diseases, including cardiovascular disorders. By examining structural and vascular patterns in retinal images, the system provides a non-invasive, intelligent, and automated method for aiding medical professionals in diagnosis and risk assessment.

The core of the system is built on the MobileNet architecture, a convolutional neural network (CNN) model known for its lightweight structure and high accuracy in image classification tasks. MobileNet is particularly well-suited for healthcare applications that demand efficient deployment in low-resource or remote environments, such as rural clinics or mobile health units. The network is trained on a labeled dataset of retinal images, learning to distinguish between six classes: Age-Related Macular Degeneration (ARMD), Diabetic Neuropathy (DN), Diabetic Retinopathy (DR), Macular Hole (MH), Optic Disc Cupping (ODC), and Normal.

To ensure high classification accuracy and generalization, the system incorporates advanced image preprocessing and augmentation techniques, including resizing, normalization, contrast adjustment, flipping, and rotation. These steps enhance the model's ability to detect meaningful features across various image qualities and acquisition conditions while reducing sensitivity to noise and variability.

The MobileNet model is fine-tuned using transfer learning and optimized with additional layers tailored for multi-class classification. This allows the system to focus on relevant visual cues—such as vessel morphology, optic disc shape, and lesion patterns—that correlate with specific retinal conditions. The model's lightweight design ensures fast inference times, making it suitable for real-time analysis.

To facilitate user interaction, a web-based application has been developed using HTML, CSS, and JavaScript, with Flask serving as the backend framework. The interface enables healthcare professionals or technicians to upload retinal images and receive immediate classification results. Detection outcomes are presented clearly, with predicted disease categories and confidence scores, allowing users to make informed decisions quickly and efficiently.

By integrating deep learning into the diagnostic workflow, this system provides an accessible, scalable, and cost-effective solution for retinal disease detection. It empowers early diagnosis, reduces the reliance on manual interpretation, and improves diagnostic reach in underserved areas. Beyond medical specialists, the platform can support general practitioners and frontline healthcare workers by offering an AI-assisted screening tool, ultimately contributing to improved patient outcomes and preventive healthcare.

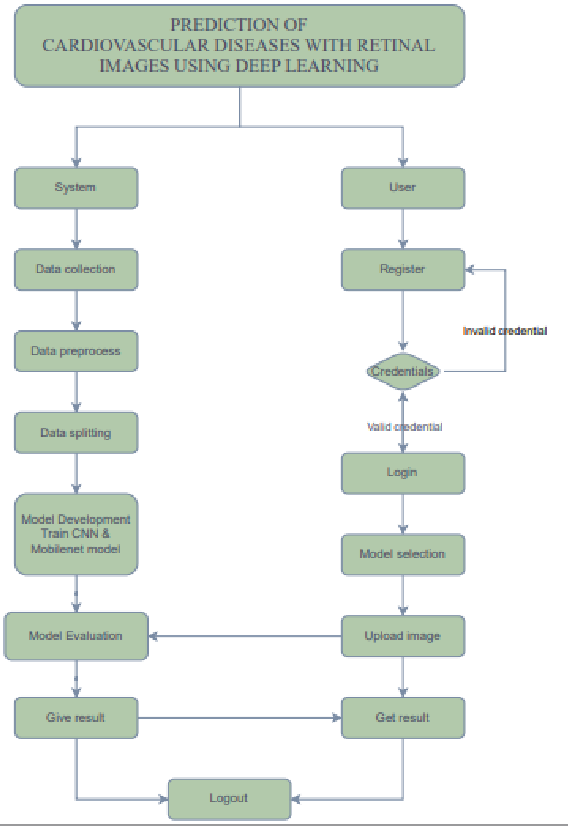

Fig 1: Flow chart

IV. Modules and its Implementation

System: Operation

Upload Data: The system begins with the upload of a curated dataset consisting of high-quality retinal fundus images categorized into six distinct classes: ARMD, DN, DR, MH, ODC, and Normal. These images, sourced from public medical datasets or clinical repositories, represent various retinal diseases and normal conditions. This dataset serves as the foundation for training the deep learning model, allowing it to learn intricate visual patterns associated with each class and enabling the development of a reliable diagnostic tool.

Data Preprocessing: Before training the model, each retinal image undergoes a comprehensive preprocessing pipeline to enhance quality, consistency, and model performance. Images are resized to a fixed dimension (e.g., 224×224) to match the input size required by the MobileNet architecture. Pixel values are normalized to ensure stable gradient flow and efficient training. To improve the model's ability to generalize and reduce overfitting, various augmentation techniques are applied, including image rotation, zooming, flipping, and contrast adjustments. These steps ensure that the CNN receives diverse and well-structured input data, ultimately contributing to a more robust and accurate classification system.

Model Building: The heart of the system is a deep learning model based on the MobileNet architecture, chosen for its lightweight design and strong performance in image classification tasks. Using transfer learning, the pre-trained MobileNet model is fine-tuned with custom classification layers tailored for the six retinal disease categories. The model is trained on the preprocessed dataset, with optimization strategies such as learning rate tuning, dropout, and early stopping implemented to improve accuracy and prevent overfitting. The MobileNet-based model effectively extracts and learns relevant features from retinal images, enabling accurate disease classification even in resource-constrained environments.

Model Prediction: Once trained, the MobileNet model is deployed for real-time image classification. Users can upload a single retinal image through the web interface, and the system instantly processes it through the trained model. The model outputs the predicted disease category along with a confidence score, providing insight into the certainty of the prediction. This real-time prediction capability ensures quick analysis and aids healthcare professionals in making timely decisions regarding patient care.

Result: After the prediction is completed, the results are presented on the prediction page with clarity and precision. The user is shown the predicted class—such as Diabetic Retinopathy or Normal—and the associated confidence score, which indicates how certain the model is about its classification. This immediate feedback facilitates rapid interpretation and supports informed clinical decision-making, particularly in early screening and preventive healthcare scenarios.

User: Operations

Register:New users can create an account by providing their credentials through a simple registration form. This functionality ensures that access to the platform is secure and personalized, allowing users to save their interactions and manage their sessions safely.

Login: Registered users can log into the system using their credentials to access the full suite of features. A successful login session provides users with access to the image upload interface, prediction engine, and result history, enabling seamless diagnostic workflows.

Home: The home page serves as the entry point to the system, offering a concise overview of the application's purpose and capabilities. It introduces users to the concept of AI-assisted retinal disease detection and highlights the benefits of using MobileNet-based deep learning models for non-invasive medical diagnostics.

About: The about section provides background on the importance of retinal imaging in disease detection, particularly its connection to systemic conditions such as cardiovascular diseases. It explains the role of deep learning in automating the diagnostic process and outlines how this system helps improve diagnostic accessibility and efficiency.

Prediction Page: The prediction page is the core feature of the application, allowing users to upload a retinal fundus image and receive immediate diagnostic results. Once an image is submitted, the system processes it using the trained MobileNet model, which classifies the image into one of six disease categories. Along with the predicted class,

Logout: Users can securely log out of their accounts after completing their session, ensuring that all personal and diagnostic data remains protected. This function supports data privacy and complies with standard security practices in medical web applications.

V. RESULTS

Using deep learning techniques for identifying and classifying different retinal diseases, the Retinal Disease Classification System shows very attractive results. Based on the MobileNet architecture, which is a lightweight-and-efficient convolutional neural-network model, the system categorized retinal images into six types: ARMD, DN, DR, MH, ODC, and normal. The model showed a high level of accuracy with depthwise separable convolution operations and, at the same time, lesser computational complexity, suitable for real-time applications of medical screening.

The deployed version of the system has an attractive web application interface, which allows users to upload retinal images and receive classification results in less than a minute. On upload of the retinal image, the model holds the processing of the image and gives out a predicted class with an associated confidence score. This score generally provides an important insight into the confidence of the model in its rationale. Such a tool would help healthcare workers and technicians quickly decide what to do; that is, it helps in early diagnosis and timely intervention, especially in limited-resource or hinterland health environments.

Some of the pivotal metrics used to evaluate model performance include accuracy, precision, recall, F1 Score, confusion matrix, and so on. The system was run on a labeled dataset made for retinal images, and it showed significant classification accuracy for all classes, that is, diabetic retinopathy and normal retina. The confusion matrix showed very few misclassification cases, and the ROC curve gave further assurance that the model could differentiate closely related disease patterns with high sensitivity and specificity.In addition to robust performance, the system placed a strong emphasis on data privacy and security. Uploaded images were handled within secure user sessions, and no data was stored or reused without user consent. The system's architecture ensured data integrity while maintaining compliance with privacy standards, making it suitable for deployment in clinical environments.

The overall lightweight design of the MobileNet-based model, combined with the seamless integration into a responsive web application, supports real-time diagnostic use and makes the system highly adaptable for broader implementation. Whether used in clinics, telemedicine platforms, or as an assistive tool for non-specialist healthcare workers, the model serves as a cost-effective and scalable solution for retinal disease screening and early detection

i ***Mobile net results***

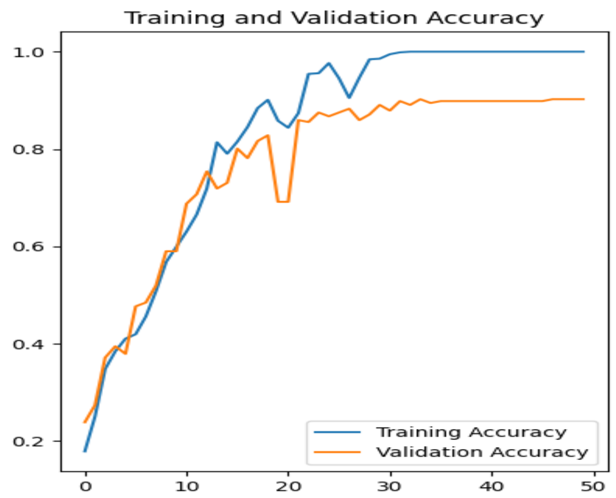

Fig 2: Accuracy curve of model

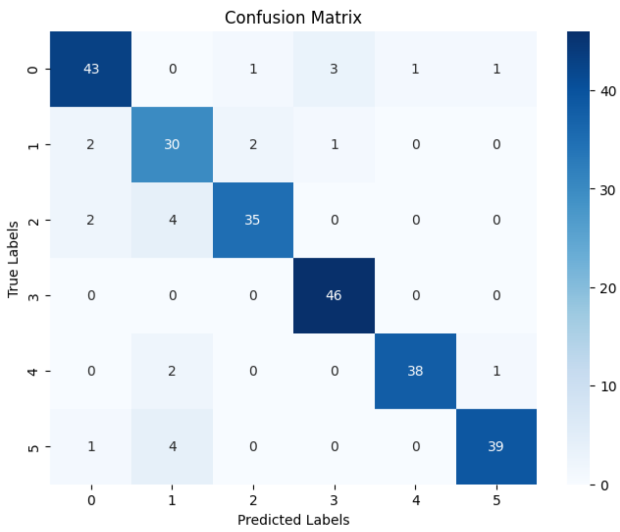

fig 3 : confusion report of MN

ii ***Output pages***

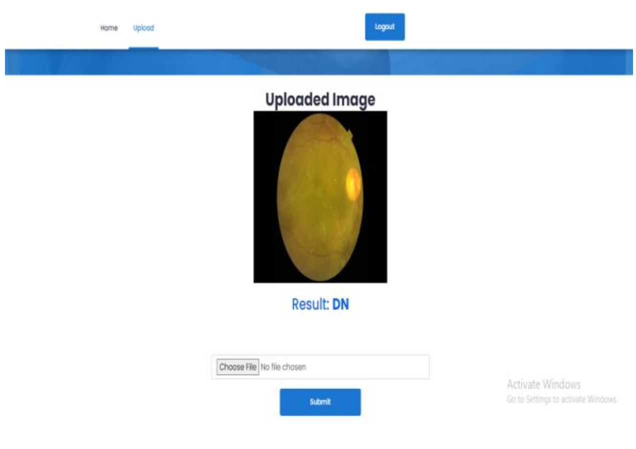

fig 4 : Prediction result 1

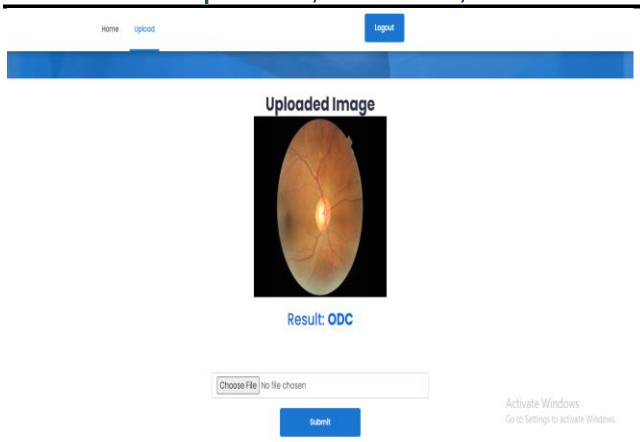

fig 5: Prediction result 2

The performance of the MobileNet model in classifying retinal diseases is highly promising, as evidenced by the training and validation accuracy plot along with the confusion matrix. The model achieved near-perfect training accuracy (~100%) and consistently high validation accuracy (~90%) after 50 epochs, indicating excellent learning capability with minimal overfitting. The confusion matrix further supports this performance, showing strong class-wise predictions with most classes having 35–46 correctly classified samples. Misclassifications are minimal and scattered, highlighting the model's robustness and its effectiveness in distinguishing between the six retinal disease categories. Overall, the MobileNet model demonstrates reliable generalization and strong predictive accuracy, making it suitable for real-world diagnostic application.

VI. Conclusion

This project successfully demonstrates how deep learning—specifically the MobileNet architecture—can be effectively utilized for the classification of retinal diseases from fundus images. By leveraging the power of convolutional neural networks and employing a well-structured image preprocessing pipeline, the system delivers accurate, efficient, and real-time predictions for six retinal conditions: ARMD, DN, DR, MH, ODC, and Normal. The MobileNet model, chosen for its lightweight and scalable nature, ensures high classification performance while remaining suitable for deployment in low-resource environments such as rural clinics or mobile diagnostic units. Integrated into a responsive web application using HTML, CSS, JavaScript, and Flask, the system allows users to upload images and receive diagnostic results instantly with confidence scores. The project presents a practical and accessible diagnostic tool that aids in early disease detection and supports proactive healthcare decisions, especially in areas with limited access to ophthalmic specialists.

VII. Future Enhancement

Future enhancements of the retinal disease classification system aim to further improve its accuracy, scalability, and usability. One key direction involves extending the model to include additional retinal and systemic conditions, such as glaucoma and hypertensive retinopathy, to broaden its diagnostic capabilities. Incorporating advanced deep learning techniques—such as EfficientNet, Vision Transformers (ViTs), or ensemble methods—could enhance feature extraction and improve classification robustness. The implementation of explainable AI (XAI) tools like Grad-CAM will increase model transparency and trustworthiness, providing visual insights into the regions influencing predictions. Real-time integration with fundus cameras and support for DICOM standards would enable seamless clinical adoption. On the user interface side, improvements such as interactive dashboards, patient history tracking, multilingual support, and integration with electronic health records (EHRs) can enrich user experience and application reach. Additionally, enabling offline diagnosis through on-device inference using TensorFlow Lite or similar technologies will expand accessibility to areas with limited internet connectivity. Ultimately, the system is envisioned as a comprehensive, AI-powered diagnostic assistant capable of supporting large-scale preventive screening programs and contributing to global efforts in early disease detection and healthcare equity.

VII.REFRENCE

[1] A. Gulshan et al., "Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs," JAMA, vol. 316, no. 22, pp. 2402–2410, Dec. 2016, doi:10.1001/jama.2016.17216.

[2] L. Zhang et al., "Automated Identification of Retinal Diseases from Fundus Images Using Deep Learning," IEEE Access, vol. 9, pp. 67205–67216, 2021, doi:10.1109/ACCESS.2021.3077595.

[3] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, pp. 436–444, 2015, doi:10.1038/nature14539.

[4] H. Pratt et al., "Convolutional Neural Networks for Diabetic Retinopathy," Procedia Computer Science, vol. 90, pp. 200–205, 2016, doi:10.1016/j.procs.2016.07.014.

[5] A. Howard et al., "MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications," arXiv preprint arXiv:1704.04861, 2017.

[6] T. Kaur and R. Gandhi, "Automated Detection of Diabetic Retinopathy Using CNN," ICTACT Journal on Image and Video Processing, vol. 10, no. 1, pp. 2043–2047, Aug. 2019.

[7] M. Sandhu and A. Saini, "Deep Learning Models for Retinal Disease Diagnosis: A Comparative Review," Biocybernetics and Biomedical Engineering, vol. 42, no. 1, pp. 18–30, 2022.

[8] E. Decenciere et al., "Feedback on a publicly distributed database: the Messidor database," Image Analysis & Stereology, vol. 33, no. 3, pp. 231–234, 2014.

[9] A. Das et al., "Explainable AI for Medical Imaging: A Survey on Datasets, Interpretation Techniques, and Validation Methods," Computers in Biology and Medicine, vol. 140, 2022, doi:10.1016/j.compbiomed.2021.105092.

[10] A. Esteva et al., "Dermatologist-level classification of skin cancer with deep neural networks," Nature, vol. 542, pp. 115–118, 2017, doi:10.1038/nature21056.