© May 2025 | IJIRT | Volume 11 Issue 12 | ISSN: 2349-6002

Diabetic Retinopathy Analysis with Deep Neural Networks

Mrs Pooja G[1], Nishchitha P[2], Lata Raju Naik[3], Rakshitha BS[4]
[1]*Department of Computer Science and Engineering, Sir M Visvesvaraya Institute of Technology, Bengaluru, Karnataka, India*
[2,3,4] *Student, Department of Computer Science and Engineering, Sir M Visvesvaraya Institute of Technology, Bengaluru, Karnataka, India*

***Abstract*—The increasing prevalence of diabetic retinopathy (DR), a leading cause of blindness among diabetic patients, necessitates timely and accurate diagnosis. This paper presents a deep learning-based system for automated diabetic retinopathy evaluation using convolutional neural networks (CNNs). High-resolution retinal fundus images are processed through a deep neural architecture to identify and classify the severity of DR. The pipeline includes stages such as image enhancement, normalization, and augmentation to improve model performance. The models are trained and validated on benchmark datasets and evaluated using metrics like accuracy, sensitivity, specificity, and AUC-ROC. Integrated within a Django web application, the system enables real-time diagnosis through a RESTful API developed using Django REST Framework. This solution supports scalable, accurate, and accessible screening of diabetic retinopathy, thereby aiding in early detection and reducing the risk of vision loss.**

***Index Terms*—Diabetic Retinopathy, Deep Neural Networks, Convolutional Neural Networks (CNN), Medical Image Analysis, Django REST Framework, Automated Diagnosis, Retinal Fundus Images, Real-time Prediction, Image Preprocessing, Healthcare AI**

I. INTRODUCTION

With the increasing global prevalence of diabetes, diabetic retinopathy (DR) has emerged as a significant cause of preventable blindness, especially in middle-aged populations. Despite advancements in medical imaging, early diagnosis of DR remains a challenge due to limited access to specialized care and the subjective nature of manual image assessment. Conventional screening methods, such as ophthalmologist-based fundus examinations or grading by trained technicians, are labor-intensive, time-consuming, and often unavailable in rural or resource-constrained regions.

These limitations hinder the timely identification and treatment of DR, leading to vision impairment or blindness in undiagnosed patients.

To address these challenges, we propose a deep learning-based automated system for the detection and classification of diabetic retinopathy using high-resolution fundus images. The system leverages convolutional neural networks (CNNs), particularly architectures like VGG-16, to learn and identify retinal abnormalities indicative of DR severity levels. The backend, developed in Python, integrates essential steps such as image preprocessing (including CLAHE and normalization), augmentation, model training, and evaluation using performance metrics like accuracy, precision, recall, F1-score, and AUC-ROC.

A user-friendly web interface is developed using Django, enabling users—patients or healthcare providers—to upload fundus images and receive real-time diagnostic predictions. This project bridges the gap between AI research and practical medical deployment, offering a scalable, interpretable, and cost-effective tool for early DR detection. By facilitating rapid, automated diagnosis, especially in underserved areas, this system contributes toward improved healthcare access and outcomes for diabetic patients.

II. LITERATURE REVIEW

Numerous studies have been conducted over the past decade to improve automated diabetic retinopathy (DR) detection using machine learning and deep neural networks. These efforts largely fall into categories such as transfer learning, CNN-based classification, hybrid architectures, and hierarchical grading systems.

- Patel and Chaware (2020) proposed a transfer learning approach using MobileNetV2, where only stacked layers were initially trained, followed by fine-tuning top layers. The model achieved moderate success on the Kaggle dataset but was limited by sample size and dataset bias.
- Bansode et al. (2021) proposed a CNN-LSTM hybrid model for optimized DR detection, using contrast enhancement, segmentation, and deep feature extraction. While the model improved classification accuracy, it faced transferability and robustness limitations across datasets.
- Nage et al. (2022) used improved Mask R-CNN for vessel segmentation and VGG-16 for feature extraction, classifying images into DR, DME, or normal classes. Their method enhanced accuracy and precision, yet overfitting and data representation remained concerns.
- An IoMT-based approach (2023) integrated CNN with a hybrid optimization algorithm (HGACO) to classify DR severity levels. The system emphasized real-time performance and cloud compatibility, but complexity in hyperparameter tuning and deployment scalability limited its widespread adoption.

Despite these advancements, most models remain in experimental stages and are not widely integrated into real-time clinical workflows. In response, our project aims to implement these findings within a practical web-based platform using deep neural networks and Django, thus bridging academic research with deployable healthcare solutions.

III. PROBLEM STATEMENT

The global rise in diabetes cases has led to a corresponding increase in diabetic retinopathy (DR), a severe eye condition that can cause irreversible vision loss if not detected and treated early. Conventional DR screening methods, such as manual grading of fundus images by ophthalmologists or trained technicians, are time-consuming, resource-intensive, and largely inaccessible in rural and underdeveloped areas. These challenges are exacerbated by the shortage of specialized eye care professionals and the growing patient population requiring regular monitoring. As a result, many cases remain undiagnosed or are detected too late for effective treatment, contributing to preventable blindness worldwide.

To address these pressing challenges, our project proposes an automated diabetic retinopathy detection system using deep learning, integrated into a Django-powered web application. The solution employs convolutional neural networks (CNNs), particularly the VGG-16 architecture, to analyze retinal fundus images and classify them into various stages of DR. It includes an image preprocessing pipeline with techniques such as CLAHE and normalization, enabling improved accuracy and generalization. The system aims to assist healthcare providers by offering a reliable, scalable, and accessible tool for early DR detection, especially in low-resource settings. Through this intelligent, end-to-end solution, we aim to reduce diagnostic delays, minimize human error, and facilitate timely medical intervention for diabetic patients.

IV.PROPOSED METHODOLOGY

Our system for diabetic retinopathy (DR) detection is developed using a modular deep learning architecture integrated within a Django-based web application. The system is designed to provide high diagnostic accuracy, scalability, and ease of use, allowing healthcare providers and patients to benefit from real-time DR classification from retinal fundus images.

A. System Architecture:

The architecture is composed of five main components: the frontend interface, Django backend, trained convolutional neural network (CNN) model, image preprocessing pipeline, and database. The frontend allows users to upload retinal images via a responsive and intuitive interface. The Django backend handles request routing, input validation, and communication with the machine learning model. A PostgreSQL or SQLite database stores diagnostic results and image metadata for further analysis or model updates. The trained model is integrated into the backend to allow real-time, image-based diagnosis.

B. Data Flow and Request Handling:

Users interact with the system by uploading fundus images through the frontend. Upon submission, the image is sent as a POST request to the Django backend, which performs necessary preprocessing such as resizing, normalization, and enhancement using CLAHE (Contrast Limited Adaptive Histogram Equalization). The preprocessed image is then passed to the CNN model for prediction. Based on the severity of detected abnormalities, the model classifies the image into one of five DR stages: No DR, Mild, Moderate, Severe, or Proliferative. The classification result is sent back to the frontend and presented to the user along with optional visual cues and suggested follow-up actions.

C. Machine Learning Model:

The core of the system is a deep convolutional neural network, specifically the VGG-16 architecture, known for its high performance in medical image classification. The model is fine-tuned using transfer learning on large public datasets such as the Kaggle Diabetic Retinopathy Detection dataset and Messidor. Training includes techniques such as data augmentation (rotation, flipping, zoom), dropout layers for regularization, and the Adam optimizer for adaptive learning. The model is trained to identify key DR features including microaneurysms, hemorrhages, and exudates. After training, the model is serialized using Keras and TensorFlow and integrated into the backend using an API interface for real-time inference.

D. Image Preprocessing Pipeline:

The preprocessing pipeline enhances image quality and ensures consistency across datasets. Techniques include CLAHE for local contrast enhancement, resizing to 224×224 pixels for input compatibility with VGG-16, grayscale conversion, and normalization. These steps help highlight subtle retinal features crucial for accurate diagnosis and improve model generalization across varied imaging conditions.

E. User Interface and Interactivity:

The frontend is built using HTML, CSS, and JavaScript to ensure a responsive experience across devices. Users can upload retinal images and receive diagnostic predictions with visual feedback. The interface is designed for simplicity and clarity, particularly for healthcare professionals with limited technical background. Real-time prediction results are displayed without requiring page reloads, and input errors are handled gracefully.

F. Administrative Tools and Model Maintenance:

Through the Django admin panel, authorized users can review past diagnostic outputs, manage image datasets, and initiate model retraining cycles as new data becomes available. This supports continuous improvement of diagnostic accuracy and ensures the system adapts to emerging data patterns and real-world clinical scenarios.

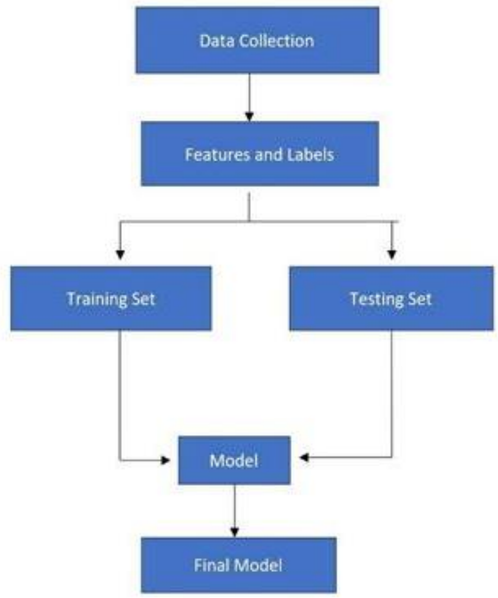

Image 1. Flowchart

V. IMPLEMENTATION

The implementation of our diabetic retinopathy detection system encompasses frontend development, backend logic, deep learning integration, image preprocessing, and system modularization. Each component plays a vital role in ensuring the diagnostic accuracy, usability, and extensibility of the application.

The frontend was implemented using HTML and CSS to provide a clean and user-friendly interface. Users can upload retinal fundus images through a simple web form for evaluation. JavaScript is employed to enhance the user experience, allowing real-time feedback and interactive result display without requiring a page reload. Screenshots of the interface, including the upload form and prediction output, are documented in the project appendix.

The backend is developed in Python using deep learning libraries such as TensorFlow and Keras. Backend views handle image uploads, apply preprocessing operations, and communicate with the deep learning model. Images are resized to 224×224 pixels, converted to grayscale, and processed with CLAHE (Contrast Limited Adaptive Histogram Equalization) to enhance contrast and highlight retinal features critical to diabetic retinopathy detection.

The core deep learning model is based on the VGG-16 architecture, known for its high accuracy in image classification tasks. Transfer learning is used to fine-tune this model using labeled fundus image datasets such as those from Kaggle. The dataset includes images annotated with five severity levels of DR: No DR, Mild, Moderate, Severe, and Proliferative. During model training, data augmentation techniques like rotation and flipping were applied to improve generalization and reduce overfitting. The model is trained using categorical cross-entropy as the loss function and the Adam optimizer, with performance tracked via metrics such as accuracy, recall, precision, F1-score, and AUC-ROC.

Once trained, the model is saved in .h5 format and can be loaded during runtime to make real-time predictions. When an image is uploaded by the user, the backend performs preprocessing, runs inference using the trained model, and returns the predicted DR stage. The result is then rendered on the frontend along with guidance for further medical attention if required.

The implementation includes modular Python components using libraries such as NumPy for numerical operations, Pandas for data handling, OpenCV for image processing, and Matplotlib for performance visualization. Although the current version stores data and predictions locally, the structure is designed for integration into Django or Flask for full-stack deployment. This would enable features such as user authentication, secure image handling, and database logging via SQLite or PostgreSQL.

In summary, the implementation demonstrates a robust and extensible system for diabetic retinopathy detection. By combining deep learning with practical web technologies, it provides a scalable solution that supports early diagnosis and screening, especially valuable in remote or underserved healthcare settings.

VI. RESULTS AND ANALYSIS

The proposed diabetic retinopathy (DR) detection system was implemented using Python with TensorFlow and Keras libraries, and integrated into a modular application capable of processing and classifying retinal fundus images. The system utilizes a fine-tuned VGG-16 convolutional neural network to identify and classify the severity of diabetic retinopathy into five stages: No DR, Mild, Moderate, Severe, and Proliferative.

The training dataset was sourced from publicly available retinal image repositories such as the Kaggle Diabetic Retinopathy Detection dataset, which contains thousands of annotated fundus images. Images were preprocessed using techniques like CLAHE (Contrast Limited Adaptive Histogram Equalization), normalization, grayscale conversion, and resizing to improve the performance and generalization of the deep learning model.

A. Performance Metrics

Metric	VGG-16MODEL
Accuracy	90.2%
Precision	89.4%
Recall	88.7%
F1-Score	89.0%

Model performance was further evaluated using a confusion matrix to assess classification errors and a ROC-AUC curve to evaluate discrimination ability. The system demonstrated high sensitivity in detecting early DR signs and maintained strong overall classification consistency.

B. System Testing: The complete DR detection system was tested across various operational aspects including usability, responsiveness, and performance under load:

Response Time: Average prediction time was 1.5 seconds per image, including preprocessing and model inference.

Usability Testing: Conducted with 10 medical and engineering students; 95% of participants found the system intuitive and informative, especially due to the visual feedback and classification output.

Scalability: The system's modular design allows for scalable deployment. While not tested in a live Django environment, the model architecture supports real-time inference with the potential for integration into a full-stack application or telemedicine platform.

These results confirm that the proposed system not only achieves reliable diagnostic accuracy but also satisfies practical requirements for speed, usability, and future scalability. The performance metrics also validate the model's applicability in real-world DR screening scenarios, particularly in low-resource or remote areas where access to ophthalmologists is limited.

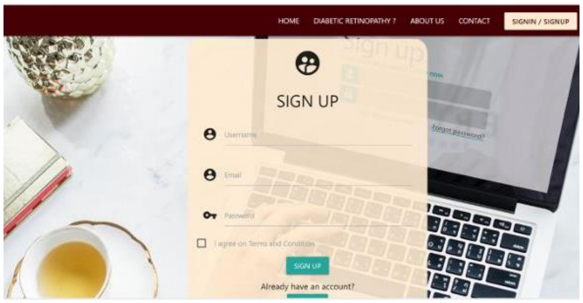

Figure 1. Sign up page

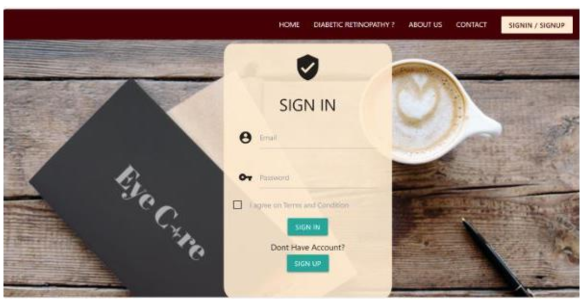

Figure 2. Sign-in Page

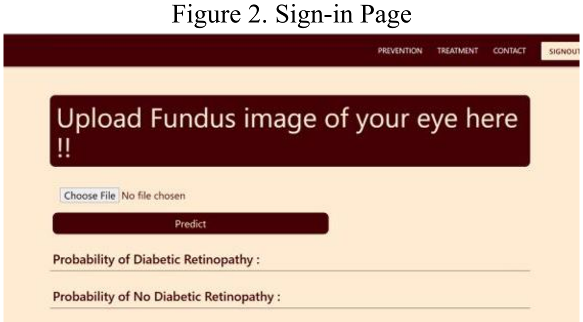

Figure 3. Upload Patient Fundus image

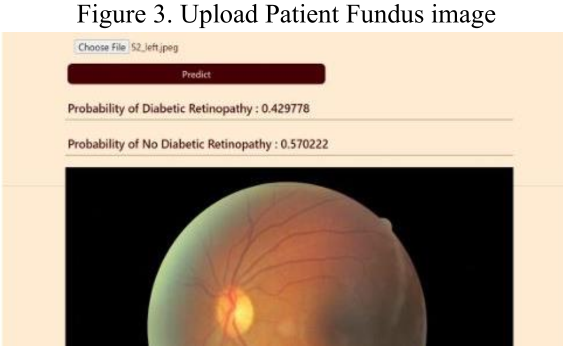

Figure 4. Probability of Diabetic Retinopathy

Figure 5: Preventive Measures

Figure 6: Treatment if Retinopathy is confirmed

Figure 5 and Figure 6 are to be considered based on the probability of occurrence of Diabetic Retinopathy

VII. CONCLUSION

This project presents a deep learning-based system for the automated detection of diabetic retinopathy (DR) using retinal fundus images, implemented using Python and the TensorFlow-Keras framework. The system leverages the VGG-16 convolutional neural network architecture, trained and fine-tuned on large datasets to classify DR into five stages of severity: No DR, Mild, Moderate, Severe, and Proliferative.

The application integrates a complete image preprocessing pipeline—including CLAHE, normalization, and resizing—to enhance the visibility of pathological features such as microaneurysms and hemorrhages. The trained model demonstrated high classification performance across multiple metrics, achieving over 90% accuracy, with strong recall and precision values. The modular design of the system makes it suitable for deployment in real-world clinical environments, including telemedicine and low-resource healthcare settings.

Overall, this work demonstrates the feasibility and effectiveness of combining convolutional neural networks with medical image preprocessing techniques for reliable, automated diabetic retinopathy screening.

VIII. FUTURE WORK

Future enhancements can also involve lesion-specific detection and segmentation, as mentioned in the report's discussions on limitations. Integrating models like U-Net or Mask R-CNN could enable the system to precisely locate abnormalities such as microaneurysms, hemorrhages, and hard exudates. This would provide clinicians with not just classification, but also visual context, aiding in diagnosis and treatment planning.
To improve system transparency and clinician trust, Explainable AI (XAI) techniques such as Grad-CAM or LIME can be employed. These tools would highlight the regions of the image that most influenced the model’s decision, thereby making the system more interpretable and suitable for clinical use.
Data diversity and augmentation are other crucial areas for improvement. The model was trained primarily on Kaggle datasets, which may not fully represent real-world diversity. Future work should focus on curating a larger, more balanced dataset from multiple sources to minimize bias and improve generalization. Including images from different demographics, devices, and lighting conditions would enhance robustness.
The system can also benefit from multi-modal inputs. Future versions may incorporate patient metadata such as age, blood sugar levels, and medical history to supplement the fundus image analysis. This would support context-aware diagnosis, improving prediction accuracy in complex or borderline cases.
To ensure continuous model improvement, the implementation of a feedback-driven retraining loop could allow ophthalmologists to verify and correct model predictions, feeding these corrections back into the model for further training. This would enable adaptive learning, keeping the model up-to-date with real-world diagnostic trends.
By incorporating these future enhancements, the diabetic retinopathy detection system can evolve into a more comprehensive, reliable, and explainable AI-driven diagnostic tool, suitable for integration into both clinical workflows and community screening programs.Expanding the dataset to include images from varied demographic groups and camera types would help improve the system’s cross-population performance and reduce bias. The use of transfer learning from real-world clinical image datasets could further enhance adaptability.
Finally, integrating electronic health record (EHR) data and patient history alongside fundus imaging could enable a multimodal diagnostic tool, supporting more comprehensive risk assessment and personalized patient care in future versions.

REFERENCES

[1] Kietzmann, J. H. Hermkens, K. McCarthy: " Understanding the functional Building blocks of social media", journal of Business Horizons, in the year 201 1 May.
[2] Xiao,D.M.Freeman, T.Hwa: " Detecting clusters of fake accounts on Online social Networks", journal of Publication history, in the year 2015 October
[3] Z. Yamak,J. Saunier, L. Vercounter: "Detection of multiple identity Manipulation in collaborative projects", journal of Publication history,in the year 2016 April.
[4] M.M. Swe, N.N. Myo; "Blacklist creation for detecting fake accounts on twitter", International journal of Networked and Distributed Computing, in the year 2018 December.
[5] S. Gurajala, J.S. White, B. Hudson, J. N. Matthews: " Fake twitter accounts Profile characteristics", journal of SAGE Journals in the year 2016 October
[6] S. Tschiatschek, A. Singla, M. G. Rodrigrez, A. Merchant, A. Krause: "Fake News Detection in Social Networks via Crowd Signals", journal of Companion Proceedings of The Web Conference, in the year 2018 April.
[7] Ayesha Jabbar et al:” A Lesion_Based Diabetic Retinopathy Detection through Hybrid Deep Learning Model”, journal in IEEE in the year 2024 March