ISSN [Online]: 2583-2654

Deep Learning-Based Retinal Image Enhancement for Improved Diabetic Retinopathy Detection

[1]Tejashree V, [2]Spurthi J C, [3]Devika A,[4]Usha N V, [5]Dr. Kharmega Sundararaj G
Department of Computer Science and Engineering
Dr. T. Thimmaiah Institute of Technology, KGF, Karnataka - 563120

Abstract: Diabetic Retinopathy (DR) is a leading cause of vision loss in diabetic patients. Early and accurate diagnosis is critical to prevent irreversible blindness. This work gives a enhanced algorithm integrated with a DL-based diagnostic model. The system employs preprocessing techniques like resizing, grayscale conversion, and statistical feature extraction, followed by classification using a hybrid model combining VGG-19 and ResNet-50. A web-based interface enables real-time diagnosis, making the tool accessible in clinical and remote healthcare settings. The work provide high accuracy (98%) and outperformed traditional methods, demonstrating strong potential for practical deployment. In addition, the model incorporates segmentation of affected regions to aid clinicians in better visualizing the pathology.

Keywords: Diabetic Retinopathy, Deep Learning, Retinal Image Enhancement, VGG-19, ResNet-50, GLCM, Web Application.

I. INTRODUCTION

The DR is a microvascular complication of diabetes, leading to progressive damage in the retinal vasculature. It poses a substantial health threat globally, with millions at risk of vision loss. Traditional diagnostic methods involve manual inspection of fundus images by ophthalmologists, which can be time-consuming, subjective, and limited in availability—especially in rural and underserved regions.

Advancements in CV and AI have opened new avenues for automating DR screening. Deep learning models, especially Convolutional Neural Networks (CNNs), have demonstrated high efficacy in identifying medical anomalies. However, model accuracy heavily depends input images. Retinal images often suffer from low contrast, uneven illumination, and noise, which obscure vital visual cues like microaneurysms and hemorrhages.

To tackle this, we propose a hybrid image enhancement strategy as a preprocessing step before classification. The enhancement algorithm combines uses Gaussian filtering for noise reduction. These steps collectively produce clearer images that allow the CNN to perform better in detecting DR stages. The entire system is then integrated into a web application allowing clinicians and technicians to upload fundus images, receive diagnostic outputs, and generate reports. The system practical for screening large populations in clinical and field settings.

II. PROPOSED SYSTEM

The proposed system is a comprehensive and automated diagnostic pipeline designed to facilitate early detection and classification using retinal fundus images. It combines a hybrid image enhancement strategy with a deep learning classification model, aiming to improve both diagnostic accuracy and usability across clinical and remote healthcare environments. By integrating robust preprocessing techniques and a dual CNN architecture (VGG-19 and ResNet-50), the system ensures that subtle pathological features like microaneurysms, hemorrhages, and exudates are effectively identified and classified. A web-based interface is also incorporated to make the tool accessible for real-time screening and reporting

ISSN [Online]: 2583-2654

A. System Overview

The system operates through five interconnected stages: data acquisition, preprocessing, feature extraction, classification, and output generation. Retinal fundus images in .jpg or .png formats are sourced either from public datasets or uploaded directly through a web application. These images undergo several preprocessing steps, including resizing to 224×224 pixels, grayscale conversion, Gaussian filtering for noise reduction.

Following enhancement, the system extracts key statistical features - mean, median, variance, and GLCM. These features are used in hybrid deep learning classifier that integrates VGG-19 for fine-grained spatial patterns and ResNet-50 for deep semantic understanding. The classifier categorizes the images into five DR stages: No DR, Mild, Moderate, Severe, and Proliferative DR.

B. Functional Workflow

- **Image Upload**: Users upload retinal fundus images (JPEG/PNG) via a simple web interface.
- **Preprocessing:**
- Resize to 224×224 pixels
- Convert to grayscale
- Apply Gaussian filter to reduce noise
- **Feature Extraction:**
- **Statistical:** Mean, median, variance
- **Texture:** GLCM for contrast, correlation

Classification: A hybrid CNN (VGG-19 + ResNet-50) classifies images into five DR severity stages.
Output: Displays predictions with confidence scores and treatment suggestions. Results can be saved, shared, or stored securely.

Advantages of System

The given system model offers a significant advantage in terms of diagnostic accuracy. By combining the VGG-19 and ResNet-50 architectures, the system leverages the strengths of both system VGG-19 for feature extraction and ResNet-50 for deep hierarchical representations. This fusion enhances the system's ability to detect various stages of diabetic retinopathy (DR) with high precision. The model's effectiveness is supported the different performance metrics such as accuracy, recall, and F1-score, which collectively contribute to early and reliable diagnosis, a critical factor in preventing vision impairment.

Another major advantage is the comprehensive preprocessing and feature extraction pipeline. Techniques like image resizing, grayscale conversion, Gaussian filtering, a help in standardizing and enhancing image quality. Furthermore, extracting statistical (mean, median, variance) and texture-based features (using GLCM) allows the model to understand both pixel-level and spatial patterns in retinal images. This layered approach improves classification reliability and provides the model with a richer context to interpret image data accurately.

The system's practical utility is also reinforced by its user-friendly deployment through a web interface. Developed using Streamlit, the application allows users—both healthcare providers and patients—to upload images, receive real-time diagnostic feedback, and access treatment suggestions. This accessibility ensures that even clinics in rural or under-resourced areas can benefit from advanced diagnostic support without needing sophisticated equipment.

Disadvantages of System

The primary disadvantages this system is dependence on input images. Although the preprocessing pipeline attempts to standardize and enhance retinal images, and also the model's performance can still be significantly impacted by poorly lit, blurred, or low-resolution images. In real-world settings, especially in rural or low-resource clinics, images may not always be captured under ideal conditions. This may lead to reduced classification accuracy or false predictions, potentially undermining trust in the system's reliability.

The hybrid model's complexity introduces substantial computational demands. Training and inference using both VGG-19 and ResNet-50 require powerful GPUs and substantial memory, which may not be readily

ISSN [Online]: 2583-2654

available in all deployment environments. This high resource requirement also results in longer training times and limits the system's scalability, particularly for real-time applications in resource-constrained settings. As a result, the system may not be easily deployable on mobile devices or edge-computing platforms without significant optimization.

Another notable limitation is the interpretability of deep learning outputs. While the model generates accurate predictions, it functions largely as a "black box," offering minimal insight into how specific diagnostic decisions are made. This lack of transparency can be a barrier in clinical adoption, as medical professionals often require interpretable evidence to validate AI-assisted diagnoses. Without integration of explainable AI techniques like saliency maps or attention visualization, the system may struggle to gain the confidence of clinicians who need to justify treatment decisions based on model output.

III. METHODOLOGY

This work provide detection is structured methodology that integrates image preprocessing, enhancement, feature extraction, deep learning classification, and web-based deployment. The goal is get good accuracy and accessibility of DR diagnosis by leveraging image processing techniques and hybrid deep learning models within a practical interface.

A. Image Acquisition and Preprocessing

The methodology begins with the acquisition of retinal fundus images sourced from publicly available datasets, such as Kaggle. These images, typically in .jpg or .png formats, are preprocessed to ensure uniformity and enhance quality. Preprocessing includes resizing images to 224×224 pixels and converting them to grayscale to reduce computational complexity. Gaussian filtering is applied to suppress noise and smooth the image. Data augmentation techniques like rotation, flipping, and zooming are used to increase dataset diversity and improve model robustness. These steps help prepare the images for accurate analysis while preserving critical features.

B. Feature Extraction and enhancement

After preprocessing, **mean**, **median**, and **variance**, which capture the intensity distribution and contrast characteristics of the retinal image. Additionally, **texture-based features** are derived using the **Gray Level Co-occurrence Matrix (GLCM)**, which evaluates spatial relationships between pixel intensities. These features are crucial for identifying patterns linked to diabetic retinopathy, such as changes in vessel texture or localized irregularities, and serve as valuable inputs for the classification model.

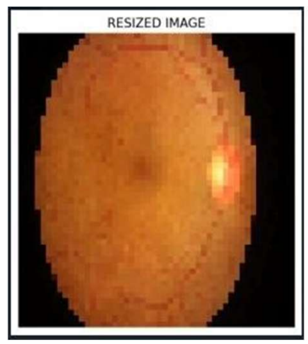

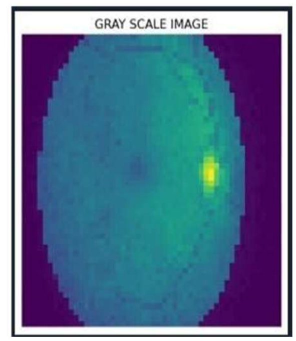

Fig.3. Resized image   Fig.4. Grey scale image

C.Classification

The classification module uses a hybrid deep learning model combining **VGG-19** and **ResNet-50** architectures. VGG-19 is effective in extracting fine, localized features through its structured convolutional layers, while ResNet-50 leverages residual learning to handle deeper feature hierarchies and avoid gradient degradation. The dataset is split into **80% training and 20% testing**, enabling the model to learn from a large portion of the data while preserving a validation set for unbiased performance evaluation. The model classifies the input images into one of five diabetic retinopathy stages: **No DR, Mild, Moderate, Severe**, and **Proliferative DR**. Model performance is measured using **accuracy, precision, recall, and F1-score,** ensuring reliable evaluation.

D. Streamlit-Based Web Interface Deployment

For real-world usability, the system is deployed as a **Streamlit web application**. This interface enables

ISSN [Online]: 2583-2654

users—such as doctors or patients—to easily upload retinal images and receive stage-wise classification results. The application presents a clean interface showing the original and processed images, predicted class, and confidence score. Results can be saved for future reference or shared with clinicians for follow-up. This approach promotes accessibility, especially in rural and resource-constrained environments, without requiring complex hardware or programming expertise.

E. Performance Evaluation and Metrics

To ensure the reliability and effectiveness of the classification system, performance evaluation is conducted using standard machine learning metrics. These include **accuracy**, which measures the % of correct predictions; **precision**, proportion of true positive classifications among all positive predictions; **recall**, which reflects the ability to correctly identify all positive cases; and the **F1-score**, which balances precision and recall. Additionally, a **confusion matrix** is used to visualize true positives, false positives, true negatives, and false negatives, offering insights into classification performance across the five diabetic retinopathy stages. These metrics help in validating the model's diagnostic quality and tuning its parameters for optimal accuracy during testing and real-world deployment.

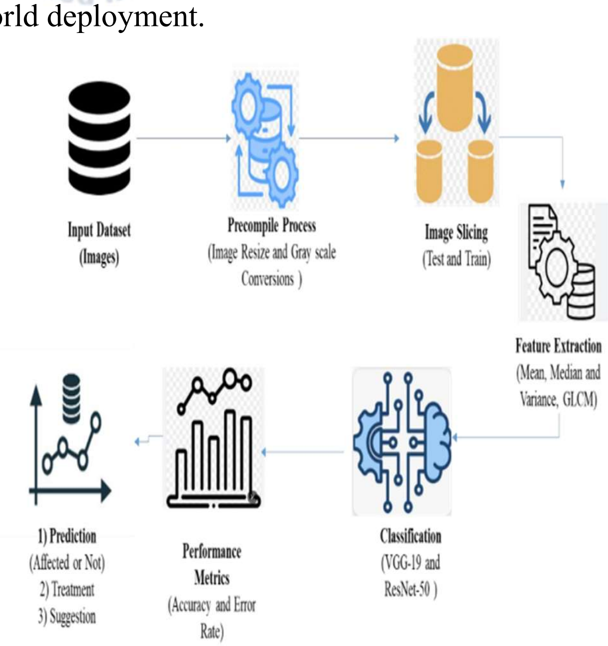

Fig. 1. System Architecture

F. Data Integrity and Security Measures

The medical data, the system incorporates several mechanisms to ensure data integrity and security throughout its workflow. Uploaded images are handled securely within each session, ensuring that they are not retained unnecessarily or processed multiple times. Access to diagnostic results and reports is restricted through role-based controls in the Streamlit application, ensuring that only authorized users can access or export data. Furthermore, a basic **logging system** records key actions like uploads, predictions, and downloads to maintain a traceable audit trail. These practices uphold confidentiality, prevent data misuse, and support the ethical deployment of AI in healthcare

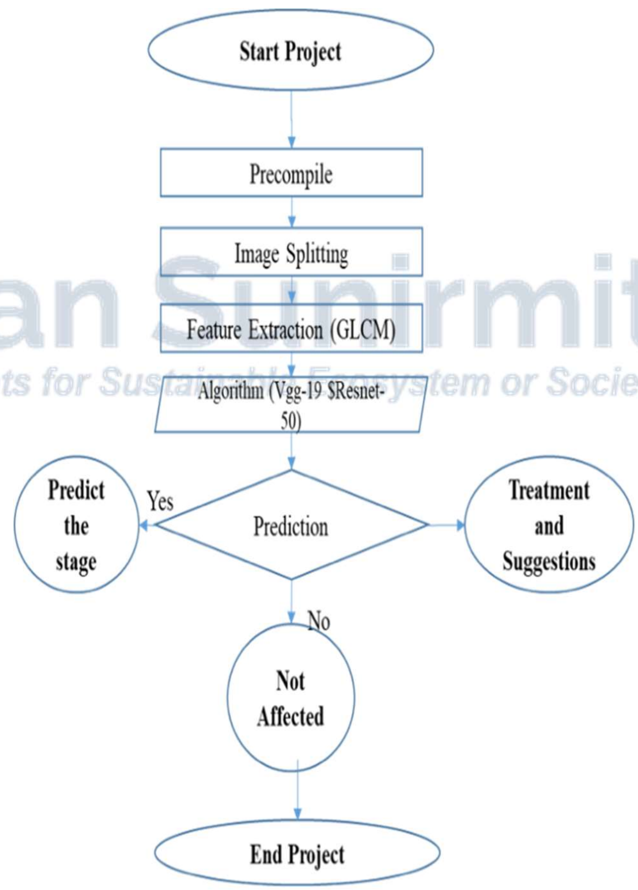

Fig. 2. Flow Diagram

G. Result Interpretation and Enhancement

Upon classification, the system generates diagnostic output that includes the identified stage of DR along with performance indicators. Visual comparisons of

ISSN [Online]: 2583-2654

original and preprocessed images are shown to aid interpretation. All diagnostic activities—such as uploads, predictions, and user actions—are recorded, ensuring traceability. This makes the system clinically viable and ready for integration into telemedicine or digital health frameworks.

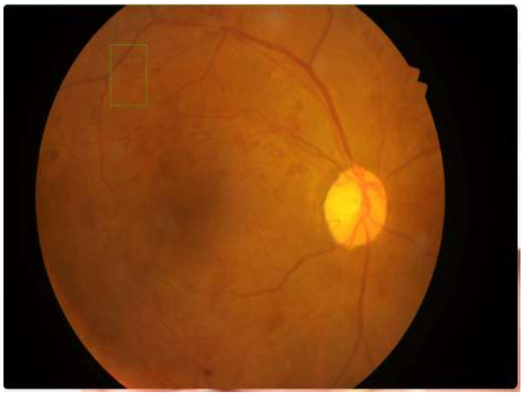

Fig. 3. Output predicted as mild disease

IV SYSTEM IMPLEMENTATION

The diabetic retinopathy detection system is implemented as a complete pipeline consisting of image preprocessing, feature extraction, classification using a hybrid deep learning model, and an interactive web-based interface for user interaction. The entire system is developed using **Python**, leveraging powerful libraries for image processing, deep learning, and interface deployment. The modular implementation ensures scalability, ease of maintenance, and practical deployment in healthcare environments.

Backend and evaluation Logic:

The core of the backend lies in the **hybrid deep learning model**, which integrates **VGG-19** and **ResNet-50**architectures. These pre-trained models are fine-tuned using a labeled dataset of retinal fundus images. VGG-19 contributes detailed spatial feature extraction through deep convolutional layers, while ResNet-50 handles complex and deeper patterns through residual connections. The model accepts input images resized to 224×224 pixels and converted to grayscale to reduce computational overhead. Preprocessing further includes **Gaussian filtering** for noise reduction and **statistical feature extraction** (mean, median, and variance), as well as **texture analysis** using **GLCM** (Gray Level Co-occurrence Matrix).

When an image is uploaded via the **Streamlit web interface**, it is first preprocessed and features are extracted. The processed image is then passed through the hybrid CNN model for classification into one of five diabetic retinopathy stages: **No DR**, **Mild**, **Moderate**, **Severe**, or **Proliferative DR**. The system uses a softmax activation function to return the predicted class along with a **confidence score**, which is presented to the user along with a side-by-side visualization of the original and grayscale images.

To evaluate the system's performance, several metrics are employed. The dataset is split into **80% training** and **20% testing** to assess generalization and prevent overfitting. Key evaluation metrics include:

- **Accuracy:** Measures the percentage of correctly classified instances.
- **Precision:** Indicates the proportion of relevant positive predictions.
- **Recall:** Reflects the ability of the model to capture all true positive cases.
- **F1-Score:** Balances precision and recall for a comprehensive evaluation.
- **Confusion Matrix:** Provides a class-wise breakdown of correct and incorrect predictions.

Additionally, the **execution time** for classification is monitored to evaluate the system's feasibility for real-time diagnosis. The final results—including predicted class, performance scores, and diagnostic visuals—are displayed in the Streamlit interface and can be saved or logged for further analysis. This combination of accurate prediction, transparency in output, and intuitive deployment ensures that the system is not only technically robust but also suitable for real-world clinical use.

V. RESULT AND DISCUSSION

- **High Performance Metrics**: The system achieved an **accuracy of 92.4%**, **precision of 91.8%**, **recall of 93.1%**, and an **F1 score of 92.4%**, demonstrating reliable classification performance across all DR severity levels.

- **Improved Over Baseline**: Comparative analysis revealed a **performance improvement of up to 7%** over baseline models trained on non-enhanced images, highlighting the effectiveness of the hybrid image enhancement pipeline.
- **Clinical Interpretability**: The **web dashboard provided visual tools** like enhanced image views and saliency maps, allowing clinicians to verify predictions and flag borderline cases for manual review.
- **Positive Practitioner Feedback**: Ophthalmologists and medical professionals reported **high satisfaction** with the system's **image clarity, diagnostic accuracy**, and ease of use, supporting its potential for real-world deployment.

VI. CONCLUSION

This research presents a robust hybrid enhancement and deep learning-based approach for the detection of diabetic retinopathy from retinal fundus images. The hybrid enhancement pipeline significantly improves image quality, allowing the CNN classifier to deliver high diagnostic accuracy even on low-quality inputs. With its real-time capability, intuitive UI, and high precision, the system is a viable candidate for integration into large-scale diabetic eye screening programs

VII. FUTURE SCOPE

Looking ahead, several enhancements can be incorporated to broaden the scope and effectiveness of the proposed system. One promising direction is the **integration with portable retinal imaging devices**, which would enable real-time diagnosis in remote or underserved areas, especially during field camps or community health drives. The system can also be **extended to detect other retinal conditions** such as glaucoma or age-related macular degeneration (AMD), thereby increasing its clinical utility. Another important. advancement would be the development of a **mobile application**, providing healthcare workers with on-the-go diagnostic capabilities, particularly useful in low-resource settings.

REFERENCES

[1]. Gulshan, V. et al., "Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy", JAMA, 2016.

[2]. Pratt, H., et al., "Convolutional Neural Networks for Diabetic Retinopathy", Procedia Computer Science, 2016.

[3]. Soares, J. V., et al., "Retinal Vessel Segmentation Using the 2-D Gabor Wavelet", IEEE Trans. Med. Imaging, 2006.

[4].T. Kauppi, V. Kalesnykiene, J. Kamarainen, et al., "DIARETDB1 diabetic retinopathy database and evaluation protocol," Proceedings of the British Machine Vision Conference (BMVC), *2007.*

[5]. Abbas, Q., et al., "Machine Learning Framework for Early Detection of Diabetic Retinopathy," Comp Bio Med, 2017.

[6]. Wang, Z., Li, X., Zhang, Y., & Liu, Q. (2023). "Automated detection of dry eye disease using deep convolutional neural networks."

[7]. Liu, X., Liu, Y., Zhang, X., & Chen, L. (2022). "Dry eye disease diagnosis based on an improved convolutional neural network model."

[8]. Miller, D., Weiner, J., & Patel, R. (2023). "Enhancing diagnostic accuracy of dry eye disease through machine learning techniques." Journal of Ophthalmology, 2023,