Diabetic Retinopathy Detection Using Deep Learning

K Skanda Shetty[1], Prof. Clitus Neil D'Souza[2], P S Chirag[3], Satvik Chandrakant Moger[4], Vilas Rathnakar Naik[5]

[1,3,4,5]*Department of Electronics and Communication Engineering, Srinivas Institute of Technology, Mangaluru – 574 143, Karnataka, India*

[2]*Assistant Professor, Department of Electronics and Communication Engineering, Srinivas Institute of Technology, Mangaluru – 574 143, Karnataka, India*

doi.org/10.64643/IJIRTV12I12-206794-459

***Abstract*—Diabetic Retinopathy (DR) is one of the leading causes of blindness among diabetic patients worldwide. Early detection is essential to prevent severe vision impairment, but traditional diagnostic methods are time-consuming and require expert ophthalmologists. This paper presents an automated system for detecting diabetic retinopathy using deep learning techniques. The proposed system utilizes Convolutional Neural Networks (CNN) and transfer learning models such as EfficientNetB2 to classify retinal fundus images into different severity levels. Image preprocessing techniques such as resizing, normalization, and contrast enhancement are applied to improve model performance. The trained model is deployed using a web-based interface for real-time prediction. The system demonstrates high accuracy and efficiency, making it a valuable tool for large-scale screening and assisting medical professionals in early diagnosis**

***Index Terms*—Diabetic Retinopathy, Deep Learning, CNN, EfficientNetB2, Image Processing, Medical Imaging, Artificial Intelligence**

I. INTRODUCTION

Diabetic Retinopathy (DR) is a progressive eye disease caused by damage to the blood vessels of the retina due to prolonged diabetes. It is one of the leading causes of blindness among working-age adults worldwide. As the number of diabetic patients continues to rise, the need for efficient and scalable screening methods has become increasingly important. In the early stages, diabetic retinopathy often does not show noticeable symptoms, making timely diagnosis challenging. If left untreated, it can progress to severe stages, leading to irreversible vision loss. Therefore, early detection and continuous monitoring are essential to prevent complications and improve patient outcomes.

Traditionally, the diagnosis of diabetic retinopathy is performed manually by ophthalmologists through the examination of retinal fundus images. However, this process is time-consuming, subjective, and requires highly trained professionals. In many rural and underdeveloped areas, access to such expertise is limited, which further delays diagnosis and treatment. With the rapid advancement of artificial intelligence and deep learning, automated medical image analysis has emerged as a promising solution. Convolutional Neural Networks (CNNs), in particular, have demonstrated remarkable performance in image classification tasks. These models are capable of learning complex features directly from raw images, eliminating the need for manual feature extraction.

In this paper, we propose an automated system for detecting and classifying diabetic retinopathy using deep learning techniques. The system utilizes both a custom CNN model and a transfer learning approach based on EfficientNetB2 to improve classification accuracy. The goal of this work is to develop a reliable, fast, and scalable solution that can assist healthcare professionals in early diagnosis and large-scale screening.

II. LITERATURE SURVEY

In recent years, significant research has been conducted in the field of automated diabetic retinopathy detection using machine learning and deep learning techniques. Various approaches have been proposed to improve the accuracy and efficiency of detecting retinal abnormalities from fundus images.

Gulshan et al. (2016) developed a deep learning algorithm for detecting diabetic retinopathy that achieved performance comparable to that of expert ophthalmologists. Their work demonstrated the potential of deep neural networks in medical image analysis and highlighted the feasibility of automated screening systems.

Pratt et al. (2016) proposed a Convolutional Neural Network (CNN) model to classify retinal images into different severity levels of diabetic retinopathy. Their model showed promising results, proving that CNNs can effectively extract features from fundus images and perform multi-class classification

Lam et al. (2018) focused on detecting specific retinal lesions such as microaneurysms and hemorrhages using deep learning techniques. Their approach improved the interpretability of the model by identifying specific disease-related features.

This paper aims to address these challenges by combining a robust deep learning model with a user-friendly web application for real-time diabetic retinopathy detection.

III. SYSTEM REQUIREMENT

A. Hardware Requirements

The hardware requirements for the proposed diabetic retinopathy detection system are minimal, as the system is primarily software-based. However, certain configurations are recommended to ensure efficient training and execution of the deep learning models.

A computer or laptop with a minimum of Intel i5 processor or equivalent

- Minimum 8 GB RAM (16 GB recommended for faster processing)
- GPU support (NVIDIA GPU with CUDA) for accelerated deep learning training
- Storage space of at least 20 GB to accommodate datasets and model files
- Stable internet connection for downloading datasets and libraries

Although GPU is not mandatory, it significantly reduces training time and improves overall system performance.

B. Software Requirements

The implementation of the system requires various software tools and libraries for data processing, model development, and deployment. The complete software stack used in this project is summarized below

Table 1 summarises the software components selected for this system.

Table 1. Software Requirements

Component	Role
Python	Core programming language used for implementation
TensorFlow / Keras	Deep learning framework for model development
OpenCV	Imagepreprocessing and enhancement
NumPy	Numerical computations
Pandas	Data handling and manipulation
Matplotlib	Visualization of graphs and results
Scikit-learn	Performance evaluation metrics
Flask / Gradio	Development of web-based user interface

C. Functional and Non-Functional Requirements

Functional Requirements

The system must perform the following operations:

- Accept retinal fundus images as input from the user
- Preprocess the input images for uniformity and clarity
- Analyze images using the trained deep learning model
- Classify the images into different diabetic retinopathy severity levels
- Display the prediction results along with confidence values
- Provide an easy-to-use interface for image upload and result visualization

Non-Functional Requirements

The system should also satisfy the following constraints:

- High accuracy in classification of retinal images
- Fast response time for real-time prediction
- Scalability to handle large datasets
- Reliability and consistency in output results
- User-friendly and intuitive interface
- Efficient memory and resource utilization

E. Dataset Requirements

The proposed system requires a retinal fundus image dataset containing labeled diabetic retinopathy images.

- Images should be categorized into five severity levels:
  - No DR
  - Mild DR
  - Moderate DR
  - Severe DR
  - Proliferative DR

- The dataset should contain sufficient image samples for training and validation.
- Images should be of good quality to improve feature extraction and classification performance.

Publicly available datasets such as EyePACS and APTOS can be used for training and evaluation.

IV. METHODOLOGY AND DESIGN

A. System Architecture

The system is designed using a layered architecture consisting of three main stages:

- Input Layer: Retinal fundus images are provided as input to the system through a dataset or user upload interface.
- Processing Layer: This stage includes image preprocessing, feature extraction using deep learning models, and classification.
- Output Layer: The predicted class label representing the severity of diabetic retinopathy is displayed to the user through a web interface.

B. Data Preprocessing

The retinal images obtained from the dataset vary in size, resolution, and quality. Therefore, preprocessing is essential to ensure uniformity and improve model performance.

The following preprocessing steps are applied:

- Resizing images to a fixed dimension (224 × 224 pixels)
- Removing noise and unwanted background regions
- Applying contrast enhancement techniques to highlight retinal features
- Normalizing pixel values to a standard range

These steps help the model focus on relevant features and improve classification accuracy.

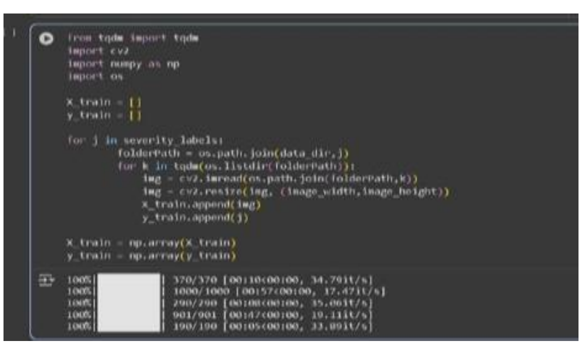

Fig.1. loading and preprocessing an image dataset for training deep learning model

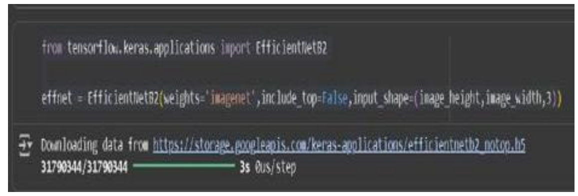

Fig.2. loading EfficientNetB2 as the base of the model

C. Model Design and Implementation

The implementation of the proposed system is carried out using Python and deep learning libraries such as TensorFlow and Keras. Two different approaches are used for model development.

The first approach involves designing a custom Convolutional Neural Network (CNN) consisting of multiple convolutional layers, activation functions, pooling layers, and fully connected layers. This model learns hierarchical features directly from the input images.

The second approach uses EfficientNetB2, a pre-trained deep learning model, as a base for transfer learning. The top layers of the model are modified by adding Global Average Pooling, Dense layers, Dropout layers, and a Softmax output layer for classification into five classes.

The models are compiled using the Adam optimizer and categorical crossentropy loss function. Training is performed over multiple epochs with a validation split to monitor performance.

Fig.3. visualizing training process

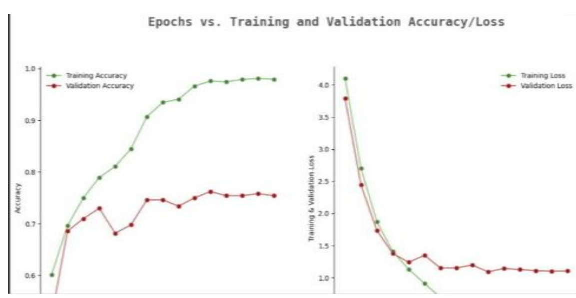

Fig.4. Accuracy and loss curves

D. Results and Discussion

The proposed system successfully classifies retinal images into five categories representing different severity levels of diabetic retinopathy.

The training and validation accuracy graphs show a steady increase in accuracy over epochs, indicating effective learning. The validation accuracy closely follows the training accuracy, suggesting good generalization of the model.

The loss graphs demonstrate a gradual decrease in both training and validation loss, confirming proper convergence of the model. Minor fluctuations in validation loss indicate slight variations in data but do not significantly affect overall performance.

Among the models tested, EfficientNetB2 achieved higher accuracy compared to the custom CNN model due to its ability to leverage pre-trained features. The system provides reliable predictions and performs efficiently in real-time scenarios.

Fig.5.Gradio interface

Fig. 6.1. DR Detection Output

Fig.6.2. DR Detection Output

E. Deployment (Web Application)

To make the system accessible, a web-based application is developed using Flask/Gradio. The application allows users to upload retinal images and obtain predictions instantly.

The interface is designed to be simple and user-friendly, ensuring ease of use for both medical professionals and general users. The integration of the trained model with the web application enables real-time diabetic retinopathy detection.

Fig .7. Confusion Matrix

```
from sklearn.metrics import classification_report

print(classification_report(y_test_new,pred))

              precision    recall  f1-score   support

           0       0.67      0.58      0.62        38
           1       0.91      0.97      0.94        91
           2       0.86      0.49      0.62        37
           3       0.72      0.85      0.78        89
           4       0.55      0.52      0.54        21

    accuracy                           0.78       276
   macro avg       0.74      0.68      0.70       276
weighted avg       0.78      0.78      0.77       276
```

Fig.8.Sklearn Metrics

V. CONCLUSION AND FUTURE SCOPE

A. Conclusion

In this paper, a deep learning-based system for the detection of diabetic retinopathy has been proposed. By utilizing Convolutional Neural Networks and transfer learning with EfficientNetB2, the system effectively analyzes retinal fundus images and classifies them into different severity levels. The approach reduces the need for manual screening and enables faster, more accurate diagnosis. Overall, the proposed system demonstrates the potential of artificial intelligence in improving early detection and supporting healthcare professionals in preventing vision loss.

B. Future Work

- Expansion of Dataset:

Increase the size and diversity of the dataset by including images from different sources, populations, and imaging conditions. This will help improve the robustness and generalization capability of the model.

- Handling Data Imbalance:

Apply advanced techniques such as oversampling, undersampling, and synthetic data generation (e.g., SMOTE) to address class imbalance and improve classification performance across all severity levels.

- Explainable AI Integration:

Incorporate Explainable Artificial Intelligence (XAI) methods such as Grad-CAM or saliency maps to visualize the regions of the retina influencing the model's prediction, thereby increasing transparency and trust in medical applications.

- Mobile and Edge Deployment:

Develop lightweight versions of the model that can be deployed on mobile devices or edge computing platforms, enabling real-time diabetic retinopathy screening in remote and rural areas.

- Integration with Healthcare Systems: Integrate the system with hospital management systems and electronic health records (EHR) to enable seamless data sharing, patient monitoring, and clinical decision support.

- Real-Time Screening System:

Extend the system to support real-time image capture and analysis using fundus cameras or smartphone-based retinal imaging devices for instant diagnosis.

- Advanced Deep Learning Models:

Explore the use of more advanced architectures such as Vision Transformers (ViT), hybrid CNN-Transformer models, and ensemble techniques to further improve accuracy and performance.

ACKNOWLEDGMENT

The authors gratefully acknowledge the management of Shama Rao Foundation for providing an environment conducive to completing this project. We express deep gratitude to our Principal, Dr. Shrinivasa Mayya D, for his kind co-operation and encouragement. Special thanks are extended to our project guide Prof. Clitus Neil Dsouza and Head of Department, Dr. Soorya Krishna K, for his invaluable support and guidance throughout the various stages of this work. We also thank Assistant Prof.
Deeksha, project coordinator, for being a source of inspiration and a constant driving force. Finally, we thank all the teaching and non-teaching staff of the Electronics and Communication Engineering Department for their assistance.

REFERENCES

[1] V. Gulshan, L. Peng, M. Coram, M. C. Stumpe, D. Wu, A. Narayanaswamy, *et al.*, "Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs," *Journal of the American Medical Association (JAMA)*, vol. 316, no. 22, pp. 2402–2410, 2016.

[2] H. Pratt, F. Coenen, D. M. Broadbent, S. P. Harding, and Y. Zheng, "Convolutional Neural Networks for Diabetic Retinopathy," *Procedia Computer Science*, vol. 90, pp. 200–205, 2016.

[3] Kaggle, *Diabetic Retinopathy Detection Challenge Dataset and Solutions*, 2019.

[4] Lam, C. Yu, L. Huang, and D. Rubin, "Retinal Lesion Detection with Deep Learning Using Image Patches," *Investigative Ophthalmology & Visual Science*, vol. 59, no. 1, pp. 590–596, 2018.

[5] Various Contributors, "Diabetic Retinopathy Detection Web Apps Using Flask & Keras," *GitHub*, 2020–2024.

[6] Revathy *et al.*, "Applied Machine Learning Techniques to Classify Diabetic Retinopathy into Different Severity Levels," 2020.

[7] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2015, pp. 234–241.

[8] Dosovitskiy *et al.*, "An Image Is Worth 16×16 Words: Transformers for Image Recognition at Scale," in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021.

[9] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov,

and L. C. Chen, "MobileNetV2: Inverted Residuals and Linear Bottlenecks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 4510–4520.

[10] Szegedy *et al*., "Going Deeper with Convolutions," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 1–9.