Point-of-Care Diabetic Retinopathy Screening Using Edge-Optimized Deep Learning

Yashwanth Patel GJ*, Irfan Nabilal Ron*, Manoj C R*, Yashaswin M*, *Shaziya Banu***
**-Students, VII Sem, Dept. of AIML, Jawaharlal Nehru New College of Engineering, Shivamogga*
***-Assistant Professor, Dept. of AIML, Jawaharlal Nehru New College of Engineering, Shivamogga*

***Abstract*—Diabetic Retinopathy (DR) remains one of the leading causes of avoidable vision loss among individuals with diabetes, making early diagnosis and precise severity grading essential for effective clinical intervention. Conventional DR screening relies heavily on manual assessment of retinal fundus images by experienced ophthalmologists, resulting in high operational costs, limited scalability, and delayed diagnosis, particularly in resource-constrained healthcare settings. To address these challenges, this paper presents an automated Diabetic Retinopathy screening system based on a hybrid deep learning framework that combines DenseNet201 and ResNet50 convolutional neural networks. Retinal fundus images obtained from the Kaggle Diabetic Retinopathy dataset are utilized and categorized into five clinical stages: No DR, Mild, Moderate, Severe, and Proliferative DR. Transfer learning is employed by initializing both networks with pretrained ImageNet weights, while selectively freezing early layers to preserve generalized feature representations. High-level feature maps extracted from both architectures are fused through feature concatenation and processed by fully connected layers with a softmax classifier to enable multi-class severity prediction. The trained model is deployed using a FastAPI-based backend integrated with a web-based interface, facilitating real-time retinal image analysis. Experimental outcomes demonstrate that the proposed system offers an efficient, accurate, and practical solution for automated DR screening, supporting timely diagnosis and enhancing accessibility in clinical and point-of-care environments.**

***Keywords*—Diabetic Retinopathy, Edge AI, Deep Learning, DenseNet, ResNet, Medical Image Analysis, CNN**

I. INTRODUCTION

Diabetic Retinopathy (DR) is a serious eye disease caused due to prolonged diabetes, which affects the blood vessels of the retina. It is one of the leading causes of vision impairment and blindness among working-age adults worldwide. In its early stages, diabetic retinopathy often shows no noticeable symptoms, but as the condition progresses, it can lead to irreversible vision loss if not diagnosed and treated in time.

Traditional diagnosis of diabetic retinopathy requires retinal fundus images to be examined manually by trained ophthalmologists. This process is time-consuming, requires specialized expertise, and may be subject to human error due to fatigue or subjective interpretation. With the increasing number of diabetic patients, there is a growing demand for automated and efficient screening systems that can assist medical professionals in early detection.

The importance of this work lies in its ability to support early detection and diagnosis of diabetic retinopathy, which is crucial in preventing vision loss. Early identification of retinal abnormalities allows timely medical intervention, thereby reducing the risk of severe complications such as retinal detachment or complete blindness.

Automating the detection process significantly reduces the workload on ophthalmologists and healthcare professionals. Instead of manually screening large volumes of retinal images, the system can quickly analyze images and provide preliminary diagnostic results. This is especially beneficial in rural or under-resourced areas where access to expert eye care specialists is limited. The system enhances consistency and reliability in diagnosis by minimizing human subjectivity. Deep learning models trained on large datasets can identify subtle patterns and features in retinal images that may be difficult to detect through manual examination. This leads to more accurate and standardized screening outcomes. Furthermore, the project demonstrates the effective application of modern deep learning techniques in the medical domain. By integrating artificial intelligence with healthcare, the system highlights how technology can improve disease screening, reduce costs, and support large-scale population health management.

II. RELATED WORK

Several studies have explored the use of traditional machine learning techniques for diabetic retinopathy (DR) detection using retinal fundus images. Early research proposed hybrid ensemble-based frameworks that combined classifiers such as Support Vector Machine (SVM), K-Nearest Neighbors (KNN), Random Forest, Logistic Regression, and Multi-Layer Perceptron (MLP) to improve classification accuracy. These approaches relied on preprocessing steps including noise removal, image enhancement, and segmentation, followed by handcrafted feature extraction targeting retinal abnormalities such as microaneurysms, hemorrhages, and exudates. Experimental results demonstrated moderate performance, with ensemble models achieving approximately 82% accuracy and Random Forest classifiers reaching up to 90%. However, the dependence on manually engineered features and the use of very small datasets significantly limited robustness, scalability, and generalization to real-world clinical settings.

With advancements in artificial intelligence, deep learning-based approaches, particularly Convolutional Neural Networks (CNNs), have gained prominence for automated DR detection. Unlike traditional methods, CNN-based systems perform end-to-end learning by automatically extracting discriminative features directly from retinal images. Several studies adopted standardized five-stage DR classification aligned with clinical guidelines and demonstrated improved diagnostic performance compared to classical machine learning models. Preprocessing techniques such as resizing and normalization were commonly employed, and training behavior was analyzed using validation curves to address overfitting and underfitting. Although these studies highlighted the superior accuracy of CNNs, challenges such as the requirement for large annotated datasets, high computational costs, and lack of real-time deployment remained unresolved.

Further research extended deep learning-based DR detection into IoT-enabled and telemedicine healthcare environments. Comparative analyses between traditional classifiers (KNN, SVM, Random Forest) and CNN models showed that deep learning approaches consistently outperformed classical methods, achieving classification accuracies in the range of 90% to 95%, along with improved sensitivity and specificity. These studies emphasized the suitability of CNNs for remote screening and teleophthalmology applications. However, the reliance on high computational resources posed challenges for deployment on low-power IoT and edge devices. Additionally, most works focused solely on fundus images and did not incorporate multimodal imaging techniques such as Optical Coherence Tomography (OCT), limiting diagnostic completeness.

Comprehensive review studies have systematically analyzed hundreds of research works on DR detection using machine learning and deep learning techniques. These reviews highlighted trends in preprocessing methods, feature extraction strategies, classification algorithms, and evaluation metrics. Reported accuracies across studies ranged from 90% to 99.8%, with hybrid models combining CNNs and traditional classifiers achieving particularly high sensitivity. Despite these promising results, major challenges such as dataset imbalance, imaging device bias, limited demographic diversity, and insufficient large-scale clinical validation were identified. The reviews strongly recommended hybrid, multimodal, and self-supervised learning approaches to enhance robustness and real-world applicability.

Some studies focused exclusively on traditional machine learning classifiers for DR prediction using fundus images, comparing models such as SVM, Bayesian classifiers, and Probabilistic Neural Networks (PNN). These approaches demonstrated high performance under controlled conditions, with SVM achieving up to 95% testing accuracy and high sensitivity and specificity. Nevertheless, their reliance on handcrafted features and sensitivity to imaging conditions restricted adaptability and scalability. Other works explored non-image-based machine learning models using structured clinical and demographic data to predict DR risk. While such models offered interpretability and moderate accuracy, they lacked direct retinal analysis capability and were unsuitable as standalone diagnostic tools.

Overall, existing literature indicates that while traditional machine learning approaches laid the foundation for automated DR detection, they are limited by feature engineering and dataset constraints. Deep learning-based CNN models have significantly improved detection accuracy and reliability; however, challenges related to computational efficiency, data availability, and deployment persist. These limitations motivate the development of hybrid deep learning architectures optimized for edge deployment, enabling accurate, real-time, and privacy-preserving diabetic retinopathy screening in practical healthcare environments.

III. SYSTEM ARCHITECTURE

The architectural components of the proposed Diabetic Retinopathy Detection System are designed in a modular manner to ensure scalability, reliability, and ease of maintenance. The frontend component serves as the user interaction layer and is developed using HTML, CSS, and JavaScript. It provides a simple and intuitive interface through which users can upload retinal fundus images and view diagnostic results. This component handles client-side validation, initiates communication with the backend through HTTP requests, and dynamically displays the predicted diabetic retinopathy stage and confidence score, ensuring a smooth and user-friendly experience.

The backend component acts as the central processing and coordination layer of the system and is implemented using FastAPI. It is responsible for handling incoming requests from the frontend, validating uploaded image files, and managing temporary storage of images. The backend also controls the flow of data between the frontend and the deep learning model. By exposing RESTful API endpoints, it ensures secure and efficient communication, handles error conditions gracefully, and formats the prediction results into JSON responses that can be easily interpreted by the frontend interface.

The deep learning model component forms the core analytical unit of the architecture. This component is developed using TensorFlow and Keras and employs a hybrid Convolutional Neural Network architecture combining DenseNet201 and ResNet50 through transfer learning. It performs automated feature extraction and classification of retinal images into different diabetic retinopathy stages. Preprocessing operations such as resizing, normalization, and RGB conversion are applied before inference to maintain input consistency. Together, these architectural components work in a coordinated manner to deliver accurate, real-time diabetic retinopathy detection suitable for clinical screening environments.

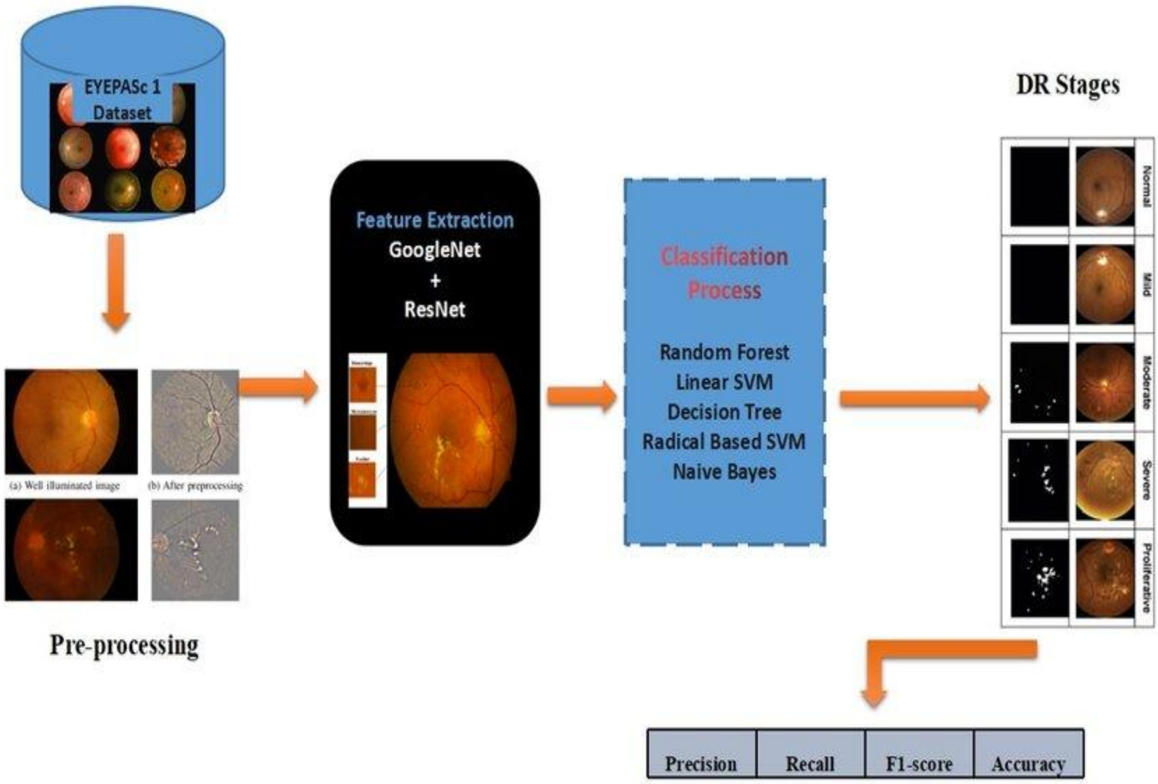

Figure 1: Proposed System Architecture

IV. IMPLEMENTATION

The implementation follows a modular and sequential workflow beginning with image acquisition and ending with result visualization. All processing and inference operations are executed locally on the edge device, ensuring low latency, data privacy, and real-time performance. The modular design allows easy maintenance, scalability, and future enhancements such as model updates or additional diagnostic features.

Algorithm 1: System Initialization

```
BEGIN
   Load trained deep learning model into Edge device
   Initialize image acquisition module
   Initialize preprocessing module
   Initialize inference engine
   Initialize result display interface
END
```

Algorithm 2: Retinal Image Acquisition

```
BEGIN
   Capture retinal fundus image from camera
   IF image capture fails THEN
      Display error message
      Request image re-capture

   END IF
   Store captured image temporarily in memory
END
```

Algorithm 3: Image Preprocessing

```
BEGIN
   Input: Raw retinal image
   Resize image to required input dimensions
   Normalize pixel intensity values
   Remove noise and artifacts
   Enhance contrast and illumination
   Output: Preprocessed retinal image
END
```

Algorithm 4: Diabetic Retinopathy Detection and Classification

```
BEGIN
   Input: Preprocessed retinal image
   Feed image to trained CNN model
   Perform forward propagation
   Extract prediction probabilities
   Determine DR class with maximum probability
   Output: Predicted DR stage and confidence score
END
```

Algorithm 5: Edge-Based Inference Control

```
BEGIN
   Load preprocessed image into inference engine
   Start inference timer
   Execute model inference on Edge device
   Stop inference timer
   Calculate inference latency
   Store inference time and prediction result
END
```

Algorithm 6: Result Visualization and Reporting

```
BEGIN
   Display predicted DR stage on user interface
   Display confidence score
   Display inference time
   IF DR stage is Moderate or Severe or Proliferative
THEN
      Display warning message for medical
consultation
   END IF
END
```

Algorithm 7: Secure Data Handling

```
BEGIN
   Do not transmit image data to external servers
   Clear temporary image storage after inference
   Store only non-identifiable performance logs if
required
END
```

Algorithm 8: Complete System Workflow

```
BEGIN
   Initialize system
   WHILE system is active DO
      Acquire retinal image
      Preprocess image
      Perform DR detection and classification
      Execute inference on Edge device
      Display results to user
      Securely handle and clear data
   END WHILE
END
```

V. RESULTS AND DISCUSSIONS

The home screen of the application serves as the entry point for users. It provides a clean and user-friendly interface that introduces the purpose of the system. The home screen displays the project title and basic instructions guiding users on how to upload retinal images for analysis. The simple design ensures ease of navigation and accessibility for healthcare professionals. This screen plays an important role in ensuring that users understand the system's functionality before proceeding with image upload and prediction.

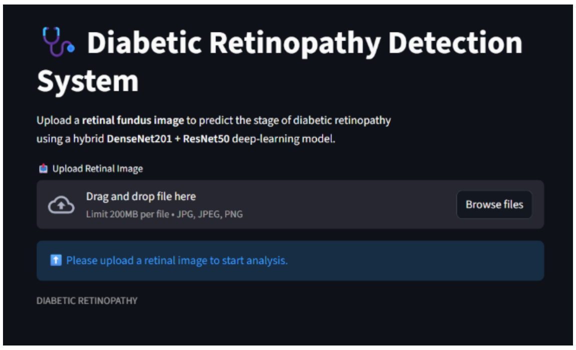

Figure 2: Home Screen

The upload page allows users to select and upload retinal fundus images in supported formats such as JPG, JPEG, and PNG. The system validates the selected image to ensure compatibility before processing. Once the image is uploaded, the system confirms successful upload and prepares the image for analysis. This stage ensures that only valid retinal images are passed to the deep learning model, reducing the chances of incorrect predictions due to invalid input.

Figure 3: Upload Page

After successful image upload and preprocessing, the image is passed to the trained deep learning model for prediction. The model analyzes the retinal features and classifies the image into one of the five diabetic retinopathy stages: No DR, Mild, Moderate, Severe, or Proliferative DR. The predicted class label is displayed along with the corresponding confidence score. This result provides a clear and immediate diagnostic output, helping users understand the severity level of diabetic retinopathy present in the uploaded image.

The visualization of results in the proposed Diabetic Retinopathy Detection System plays a crucial role in presenting the model's prediction in an interpretable and user-friendly manner. After the retinal fundus image is uploaded and analyzed by the hybrid DenseNet201 + ResNet50 deep learning model, the system visually displays both the predicted class and the confidence levels associated with each diabetic retinopathy stage. As shown in the result screen, the uploaded retinal image is first displayed to confirm successful input. Below the image, the system presents the final prediction result, highlighting the detected diabetic retinopathy stage along with the confidence percentage.

In the given example, the system predicts Severe Diabetic Retinopathy with a confidence of 57.05%, clearly indicating the severity level to the user. To enhance interpretability, the system provides a class probability visualization using horizontal progress bars. These bars represent the probability distribution across all five diabetic retinopathy classes: No DR, Mild, Moderate, Severe, and Proliferative DR. Each bar length corresponds to the confidence score of the respective class, allowing users to visually compare the likelihood of different disease stages. This probabilistic visualization helps medical professionals understand not only the final diagnosis but also the uncertainty and closeness between different classes.

Additionally, the final diagnosis is emphasized using color-coded highlighting, where severe conditions are displayed prominently to draw immediate attention. This visual emphasis assists healthcare professionals in quickly identifying critical cases that may require urgent medical intervention.

Overall, the result visualization improves transparency, usability, and trust in the system by combining numerical confidence values with graphical representation. It ensures that predictions are easy to interpret, making the system suitable for real-world clinical screening and decision-support applications.

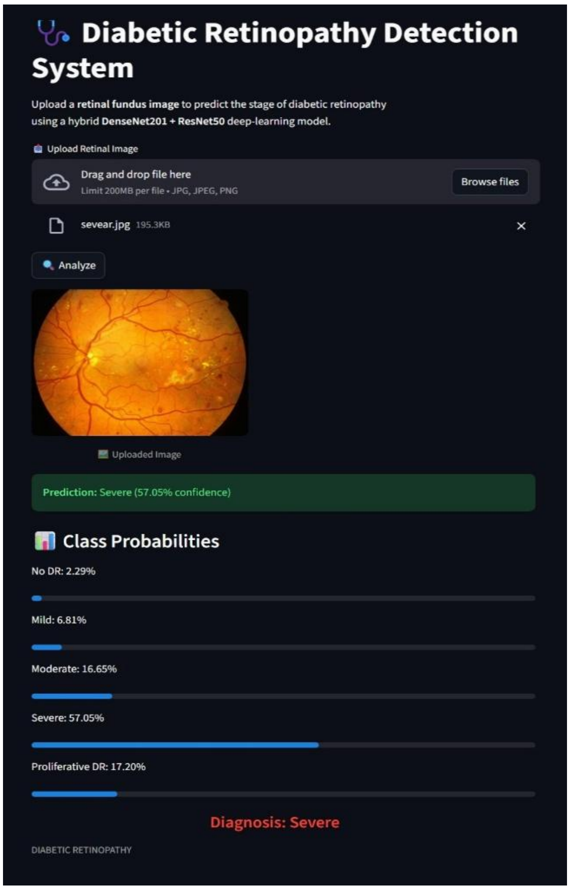

Figure 4: Result Page

VI. CONCLUSION AND FUTURE SCOPE

This work successfully demonstrates the development of an automated Diabetic Retinopathy detection system using deep learning techniques, aimed at providing an accurate, efficient, and user-friendly solution for retinal image analysis. By employing a hybrid DenseNet201 and ResNet50 architecture with transfer learning, the system effectively extracts both fine-grained and high-level retinal features, resulting in improved classification accuracy and robustness compared to single-model approaches. Image preprocessing techniques such as resizing, normalization, and RGB conversion ensure consistent input quality and stable model performance. The integration of the trained model into a web-based application with a FastAPI backend enables seamless image upload, real-time prediction, and clear visualization of confidence scores, enhancing result interpretability for clinical support. Experimental results confirm reliable classification of retinal images into five stages: No DR, Mild, Moderate, Severe, and Proliferative DR. Although not a replacement for ophthalmologists, the system significantly reduces screening time and workload, making it suitable for large-scale screening and remote healthcare environments.

The proposed Diabetic Retinopathy Detection System offers strong potential for future improvements and real-world clinical deployment. While the current system successfully classifies retinal fundus images into multiple stages of diabetic retinopathy using a hybrid deep learning approach, further enhancements can significantly improve its accuracy, reliability, and clinical acceptance. One important future direction is the integration of explainable artificial intelligence (XAI) techniques such as Grad-CAM or saliency maps, which can visually highlight retinal regions influencing predictions and help ophthalmologists better understand and trust the system's decisions.

BIBLIOGRAPHY

[1] A. Verma, R. Singh, and P. Sharma, "Diabetic Retinopathy Detection Using Machine Learning Techniques," *International Journal of Engineering Research & Technology (IJERT)*, vol. 9, no. 6, pp. 1120–1125, 2020. Available: http://www.ijert.org

[2] S. Kumar, A. Gupta, and N. Jain, "Diabetic Retinopathy Detection Using Deep Learning," *Journal of Emerging Technologies and Innovative Research (JETIR)*, vol. 7, no. 8, pp. 124–130, 2020. Available: https://www.jetir.org

[3] R. Mehra, V. Patel, and S. Desai, "Diabetic Retinopathy Detection Using Machine Learning in IoT-Based Healthcare Systems," *EAI Endorsed Transactions on Internet of Things*, vol. 10, pp. 1–9, 2024.

[4] G. Quellec, K. Charrière, Y. Boudi, B. Cochener, and M. Lamard, "A Comprehensive Review on Machine Learning and Deep Learning Techniques for Diabetic Retinopathy Detection," *Big Data and Cognitive Computing*, vol. 6, no. 154, 2022. DOI: 10.3390/bdcc6040154

[5] M. Rahman, T. Islam, and M. Hossain, "Predicting Diabetic Retinopathy Using Machine Learning," *Journal of Scientific Research (JSR)*, vol. 12, no. 3, pp. 563–575, 2020. Available: https://www.jsr.org

[6] K. Prasad, V. Babu, and R. Reddy, "Automated Diagnosis of Diabetic Retinopathy Using Machine Learning Techniques," *ICTACT Journal on Soft Computing*, vol. 3, no. 4, pp. 498–503,

July 2013. DOI: 10.21917/ijsc.2013.0083
[7] V. 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, 2016.
[8] H. Pratt, F. Coenen, D. Broadbent, S. Harding, and Y. Zheng, "Convolutional Neural Networks for Diabetic Retinopathy," *Procedia Computer Science*, vol. 90, pp. 200–205, 2016.
[9] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," *Proceedings of CVPR*, pp. 770–778, 2016.
[10] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely Connected Convolutional Networks," *Proceedings of CVPR*, pp. 4700–4708, 2017.
[11] R. R. Selvaraju et al., "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization," *Proceedings of ICCV*, pp. 618–626, 2017.
[12] G. Litjens et al., "A Survey on Deep Learning in Medical Image Analysis," *Medical Image Analysis*, vol. 42, pp. 60–88, 2017.
[13] J. Ting et al., "Artificial Intelligence and Deep Learning in Ophthalmology," *British Journal of Ophthalmology*, vol. 103, no. 2, pp. 167–175, 2019.
[14] S. Sengupta et al., "Artificial Intelligence in Diabetic Retinopathy Screening," *Indian Journal of Ophthalmology*, vol. 68, pp. 152–158, 2020.
[15] L. Li et al., "Automated Detection of Diabetic Retinopathy Using Deep Learning," *IEEE Access*, vol. 7, pp. 12494–12503, 2019.
[16] Y. Lecun, Y. Bengio, and G. Hinton, "Deep Learning," *Nature*, vol. 521, pp. 436–444, 2015.
[17] A. Krizhevsky, I. Sutskever, and G. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," *NIPS*, pp. 1097–1105, 2012.
[18] S. Quellec et al., "Deep Image Mining for Diabetic Retinopathy Screening," *Medical Image Analysis*, vol. 39, pp. 178–193, 2017.
[19] T. Shanthi and R. S. Sabeenian, "Modified AlexNet Architecture for DR Detection," *Computers in Biology and Medicine*, vol. 120, 2020.
[20] P. Costa et al., "End-to-End Deep Learning for DR Grading," *IEEE Journal of Biomedical and Health Informatics*, vol. 22, no. 3, pp. 1–11, 2018.