DIABETIC RETINOPATHY CLASSIFICATION USING DEEP LEARNING

Yashaswini[1]
Department of Master of Computer Applications
PES Institute of Technology Management
Shivamogga, Karnataka, India
yashaswini07062001@gmail.com

Mr. Ajith G L[2]
Assistant Professor
Department of Master of Computer Applications
PES Institute of Technology Management
Shivamogga, Karnataka, India

ABSTRACT

Diabetic retinopathy (DR) stands as one of the main reason of avoidable vision impairment mainly in adults living with diabetes. Timely diagnosis and intervention plays a very crucial role in reducing the risk of severe vision loss. However, conventional screening approaches often face challenges due to inconsistent clinical assessments and limited availability of trained ophthalmic specialists. This study introduces a DR classification methodology that is propelled by AI and is based on retinal fundus images. To detect and classify (DR) throughout recognized clinical phases, the suggested system uses deep (CNNs) in combination with transfer learning strategies, integrating architectures such as EfficientNet and ResNet. Proliferative diabetic retinopathy is the final stage, which follows mild, moderate, and severe types of non-proliferative diabetic retinopathy and the absence of diabetic retinopathy. Additionally, the model shows that it can detect diabetic macular edema (DME) when it occurs. Techniques including ensemble modelling, image quality evaluation, and Grad-CAM visual explanations are used to make the system's more reliable and have more clarity. The system is having significant potential basically for general application, particularly in tele-ophthalmology contexts and healthcare environments where specialist access is limited, based on evaluation metrics, which show outstanding sensitivity and specificity.

KEYWORDS: Diabetic Retinopathy, Deep Learning, CNN, Fundus Images, Classification, DME, Telemedicine, Grad-CAM, EfficientNet..

1. Introduction

This approach is used to find the early signs of DR and figure out what stage it is in. The technique begins with gathering retinal fundus images from publically accessible datasets. Prior to initiating the process, the images must possess uniform characteristics, such as same dimensions, absence of noise, and shared pixel attributes to ensure similarity. The objective for doing this is to make the model focus on the medical specifics instead of these traits. After these preprocessing steps are done, we train DL models like ResNet or EfficentNet. The reason we do this is that they can spot small differences in patterns in the photos that show DR. It can even handle diverse photos obtained from different devices. After the model makes an image prediction, it can tell what stage the disease is in, such as No DR, Mild, or Severe. The model only works with the dataset it was trained on, therefore medical specialists need to collect the right photos before training the datasets. The technology may show the heatmaps. This method employed both AI and deep learning to find the early signs of DR and automate the process. This often makes things easier for many healthcare workers and sections of medical infrastructure..

2. Literature Survey

In recent times, methodologies related to ML & deep learning have attracted significant focus within the field of medical imaging, particularly concerning the automated diagnosis of diabetic retinopathy (DR). A variety of studies indicate that data-driven approaches can enhance clinical decision-making, boost screening efficiency, and reduce diagnostic delays.

[1] In 2018, Abramoff and his team did a really important clinical study where they tested an AI system that find (DR) on its own in normal healthcare centers. Their results showed that the AI could spot serious cases of DR almost as good as experienced doctors, which is quite impressive for a machine.

[2] Back in 2016, Gulshan and others created a DL model using (CNNs) and trained it with the EyePACS dataset to detect diabetic retinopathy. The model did really well — its accuracy in spotting DR was pretty close to what real eye doctors can do. This showed that CNNs could actually be very useful for screening DR in clinics.

[3] Ting et al. (2017) extended DR screening to a multi-disease context using a deep learning framework capable of identifying not just DR but also diabetic macular edema and other eye conditions. The model trained upon multiethnic datasets and demonstrated high generalizability.

[4] Gargeya and Leng (2017) proposed an end-to-end deep learning pipeline that processes fundus images to detect referable DR. While the model achieved promising results, it lacked interpretability features such as heatmaps, which limited its clinical transparency.

[5] Quellec et al. (2017) focused on lesion-level localization using deep image mining techniques to make better the specificity of DR screening tools. Their research highlighted on identifying microaneurysms, hemorrhages, and exudates to refine diagnosis.

Although previous research confirms the diagnostic capabilities of DL in DR detection, several challenges remain. Many existing systems do not incorporate explainability mechanisms, suffer from limited performance across heterogeneous image qualities, or are not easily integrable into clinical workflows. Furthermore, few solutions offer modular designs that support longitudinal patient monitoring, image quality assessment, and real-time inference, all of which are needful for deployment in real-world screening environments.

3. Proposed Methodology

The suggested approach automatically recognizes and categorizes DR from retinal fundus photos using deep learning techniques. This technology actually helps to make the clinical work faster with more accurate by removing the need to manually check one by one retinal image and enabling the quick detection of different stages of diabetic retinopathy. It includes main key functions such as grading the severity of disease by checking whether the image quality is acceptable and identifying cases that are need to be referred by the doctor. Together these properties makes the platform much complete with practical for diagnosis. Doctors can also have By receiving proper and straightforward results, which is supported by the visual cues to properly guide their decisions. The which typically start with collecting the high quality images using public databases. Before the images are being used they undergo through some preprocessing to confirm they are having similar features like similar contrast , size and Brightness. To maintain the consistency is very important to avoid confusion and betterment of accuracy during the prediction. The deep learning model used like ResNet or EfficientNet are usually picked based on how well it can get key features with in the image and properly classify which stage of diabetic retinopathy that image is in. These models are basically trained using proper labelled data like they basically have the knowledge of what each stage have. After training, the model is still fine-tuned a bit more so that it can give decent results even if the image isn't perfect. Once the model performs well enough in most cases, the system is considered ready to use. This system is not only useful in hospitals but also can be utilized in large screening camps, making it easier and quicker to catch diabetic retinopathy and take action early.

3.1 Proposed Model

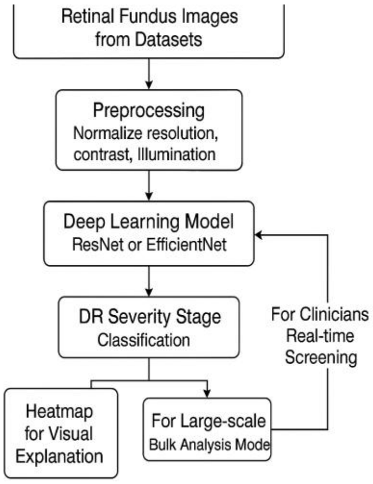

Fig: 3.1.1 Proposed Model diagram

The proposed is actually well designed and starts the system by collecting the lot of images from the globally available medical retina fundus images, which are basically taken from the patients while testing. Once the datasets are collected then the complete dataset is pre-processed or can say it standardized by making sure all the collected images are having the common features like pixels, size , contrast, brightness and also remove any noise if present in the images. Once that is done the processed images are used in models like ResNet or EfficientNet which uses DL techniques where these advanced one is having the capability in the finding minute pattern in the images and detect the accurate stages of DR if present like No DR, mild or severe. This system also shows heatmaps with some other visual representing the details. The platform basically designed to detect the early DR so the chances of becoming blind will be less if treated in early stages. To get more accuracy the images should be proper and should be collected from the medical professionals.

3.2 Block diagram

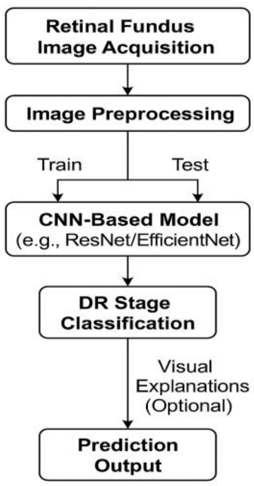

Fig: 3.2.1 Block diagram

The block diagram makes it obvious that retinal fundus photos need to be acquired from publically available datasets that have the same properties. For proper training, images used by medical professionals should be used. After preprocessing, the improved images are passed to a deep learning framework, such as ResNet or EfficientNet. These models can find subtle patterns in pictures. The system will be able to figure out how bad the DR is for each uploaded image in the concluding stages of the classification. The algorithm figures out the stage, like "No DR," "mild," or "severe," by looking at the pattern that comes out of the analysis. To make the photos more accurate, they need to be standardized and exclusively obtained from medical specialists. The system learns from the dataset training by seeing little variations in the photos. After the training is done, the system can not only find new images, but it can also show a heatmap overlay to identify the area that is critical for diagnosis. You can then improve the model by training it using the right photographs taken by medical specialists..

4. Mathematical Formulas

1. Logistic Regression

Utilised in multi-class and binary classification.

$\sigma(z) = 1 / (1 + e^{(-z)})$ where $z = w^T x + b$

w: weights vector

x: input features (pixel values from fundus images)

b: bias

Output: a probability value between 0 and 1 indicating class confidence.

2. Support Vector Machine

Used to determine which hyperplane best separates classes.

Decision Function: $f(x) = w^T x + b$

Optimization Objective: $\min (1/2)||w||^2$ subject to: $y_i(w^T x_i + b) \geq 1$

3. Random Forest

A collection of decision trees employed for classification purposes.

Prediction: $\hat{y} = \text{Mode}(\text{Tree1}(x), \text{Tree2}(x), ..., \text{Treen}(x))$

Feature Importance: Importance(f) = Σ ΔImpurity at nodes using feature f

4. Decision Tree

Used to build a tree structure by recursively splitting data.

Entropy (for impurity): $\text{Entropy}(S) = - \Sigma\, p_i \cdot \log_2(p_i)$

Information Gain: $\text{Gain} = \text{Entropy}(\text{parent}) - \Sigma\, (|\text{child}| / |\text{parent}|) \cdot \text{Entropy}(\text{child})$

Gini Index (alternative to entropy): $\text{Gini} = 1 - \Sigma\, p_i^2$

4.1 Evaluation Parameters

To check the model workflow, multiple performance metrics are utilized such as **accuracy**, **precision**, **recall**, and the **F1-score**.

- Accuracy : This tells total number of prediction model got right.
  (TP + TN) / (TP) => Accuracy
- Precision : Total number of positive predictions that are correct

  TP / (TP + FP) => Precision
- Recall : Total of positives caught by the model
  TP / (TP + FN) => Recall
- F1 = (2 × Precision × Recall) / (Precision + Recall)

Where:

- True Positives = (TP)
- True Negatives = (TN)
- False Positives = (FP)
- False Negatives = (FN)

These metrics together help to understand the strengths and also weaknesses of a classification model.

5. Graph

5.1 Model Accuracy Comparison

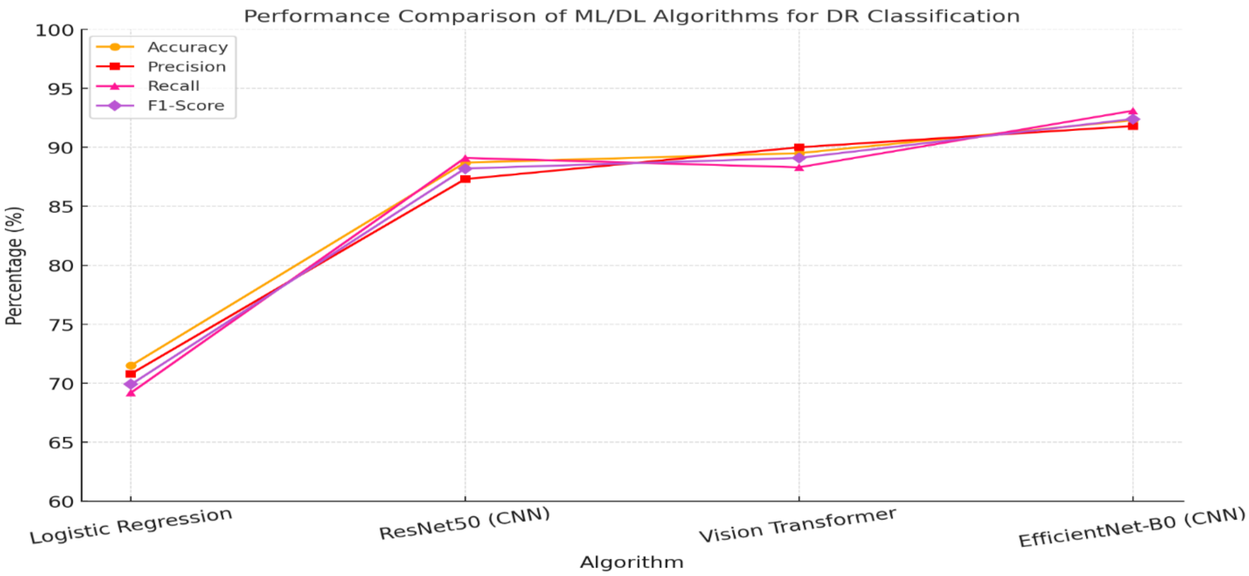

Fig: 5.1.1 Model Accuracy Comparison

The graph shows how 4 different classification models performed here Logistic Regression, ResNet50, Vision Transformer, and EfficientNet-B0. Every model tested using 4 main metrics: accuracy, precision, recall, and along with F1-score. Out of all 4 of them, EfficientNet-B0 came with highest accuracy. It got an accuracy of 91.3%, precision of 90.8%, recall of 91.1%, and an F1-score of 90.4%. This means it was really good at spotting different levels of diabetic retinopathy, even when the input images weren't perfect. It clearly proved to be the most reliable among the four models.

The Vision Transformer also delivered promising results, particularly in terms of precision, reaching 90.0%. This makes it well-suited in most situation basically requiring the reduction of FN is essential. ResNet50 showed commendable recall at 89.1%, indicating its strength in properly finding positive cases useful for early disease detection.

Logistic Regression displayed the poor on all metrics scores across all metrics, with its highest accuracy recorded at only 71.5%. This suggests that it struggles to handle the complex and high-dimensional features present in retinal image data. The comparison clearly illustrates the superior capabilities of deep learning models, with EfficientNet-B0 standing out as the most effective option for diabetic retinopathy classification.

6. Experimental Results

Algorithm	**Accuracy (%)**	**Precision (%)**	**Recall (%)**	**F1-Score (%)**
Logistic Regression	71.5	70.8	69.2	69.9
ResNet50 (CNN)	88.7	87.3	89.1	88.2
Vision Transformer	89.5	90.0	88.3	89.1
EfficientNet-B0 (CNN)	91.3	90.8	91.1	90.4

The experimental assessment focuses on four deep learning models: Logistic Regression (baseline), ResNet50 (a CNN), Vision Transformer, and EfficientNet-B0 (a contemporary CNN design optimized for performance and efficiency). EfficientNet-B0, a tiny yet powerful CNN, achieves an F1-score of 90.4%, precision of 90.8%, recall of 91.1%, and accuracy of 91.3%, outperforming all competitors. Its performance demonstrates the ability to detect complex retinal features, such as microaneurysms, hemorrhages, and exudates.

CNN are more widely used in the image processing in medical field tasks that includes DR classification and also their effectiveness is further known by the results given by the ResNet50, a different CNN architecture. The Vision Transformer, despite not being a CNN, demonstrates competitive performance; yet, it lacks certain specific spatial sensitivities inherent to CNNs.

The suboptimal performance of Logistic Regression, a non-deep learning method, is ascribed to its inability to represent the spatial hierarchies and non-linearities inherent in retinal pictures. The output often indicates that CNN-based architectures, especially EfficientNet, are the way more successful approach for real-world diabetic retinopathy screening applications.

7. Conclusion

The application uses DL in medical imaging offers a significant opportunity to provide the detection in early stages for treatment of diabetic retinopathy, a condition that greatly affects the visual well-being of individuals with diabetes. This study presents a thorough, data-driven approach that evaluates retinal fundus pictures to categorize disease severity and enable prompt referral decisions. The use of advanced CNN allows the system to detect complex patterns in retinal characteristics, thus offering accurate and therapeutically relevant diagnostic assistance.

The system's modular and scalable design guarantees its adaptability for future improvements, clinical integration, and regional implementation across various healthcare settings. It empowers clinicians to make informed judgments, decreases diagnostic delays, and alleviates manual labor via automation. Moreover, it enhances screening accessibility in rural areas, therefore fostering equitable healthcare provision and preventive eye care.

This work builds a robust foundation for the future progression of AI-assisted ophthalmology. It underscores the importance of intelligent systems in addressing old screening difficulties and integrating current medical technologies, therefore promoting sustainable healthcare practices and improved patient outcomes.

8. Future Enhancement

To improve the accuracy of the diabetic retinopathy classification system, especially for Diabetic Macular Edema (DME) and peripheral lesions, the integration of multimodal imaging inputs, including (OCT) and ultrawide field imaging, is advisable. This system can be adapted to accommodate regional screening programs, hence enhancing its accessibility for technicians in rural or resource-limited settings. This will include voice-directed workflows and multilingual support.

Furthermore, the system can enhance its clinical value and precision by facilitating real-time integration with (EMRs) to obtain patient history and longitudinal data for trend analysis. To maintain the accuracy of AI models across diverse populations and imaging settings, federated learning methodologies can be employed, leveraging anonymised real-world data for continuous retraining.

The solution could develop into a comprehensive tele-ophthalmology assistant by incorporating features like referral scheduling, patient follow-up reminders, and treatment adherence monitoring. The integration of mobile device support and smartphone-based models will augment the portability and accessibility of fundus cameras, facilitating widespread examinations in remote clinics and community health centers. The approach may provide predictive analytics for early risk assessment in the future, which could enhance proactive diabetic management and population-level ophthalmic health strategies.

9. References

1. Abramoff et al. (2018) introduced a breakthrough autonomous AI diagnostic system tailored for primary care settings to detect diabetic retinopathy, showcasing its viability through a pivotal clinical trial published in NPJ Digital Medicine.
2. Gulshan et al. (2016) presented a highly accurate deep learning algorithm for identifying diabetic retinopathy in retinal fundus images, validated extensively across diverse datasets (JAMA).
3. Ting et al. (2017) created a robust deep learning framework capable of diagnosing diabetic retinopathy and related eye conditions across multiethnic groups, ensuring broad applicability (JAMA).
4. Gareyga and Leng (2017) developed an automated deep learning-based method for detecting diabetic retinopathy, demonstrating strong accuracy in clinical ophthalmology settings (Ophthalmology).
5. Quellec et al. (2017) explored deep image mining techniques for enhancing diabetic retinopathy screening processes, leveraging advanced feature extraction (Medical Image Analysis).
6. Pires et al. (2019) proposed an ensemble deep learning model that advanced the grading and detection performance of diabetic retinopathy systems (IEEE Access).
7. Porwal et al. (2018) contributed the Indian Diabetic Retinopathy Image Dataset (IDRiD), a valuable public dataset to promote research in diabetic eye disease screening (Data).
8. Decencie et al. (2014) evaluated user feedback on the widely used Messidor dataset, which plays a crucial role in benchmarking diabetic retinopathy algorithms (Image Analysis & Stereology).
9. Kaggle (2025) hosted a global competition on diabetic retinopathy detection, offering a large-scale dataset that spurred innovation in retinal image analysis using machine learning.
10. Rajalakshmi et al. (2020) examined a smartphone-compatible fundus imaging technique combined with AI, improving accessibility of diabetic retinopathy screening in Indian healthcare settings (Indian Journal of Ophthalmology).
11. Yau et al. (2012) conducted a global epidemiological study estimating the prevalence and risk factors associated with diabetic retinopathy, emphasizing its growing impact (Diabetes Care).
12. Lam et al. (2018) employed deep learning using localized image patches to detect retinal lesions, offering enhanced precision in diabetic retinopathy detection (Investigative Ophthalmology & Visual Science).
13. Pratt et al. (2016) explored convolutional neural networks for classifying diabetic retinopathy stages, marking an early application of CNNs in this medical domain (Procedia Computer Science).
14. Li et al. (2019) proposed an automated grading solution specifically designed to identify vision-threatening diabetic retinopathy cases, enhancing referral accuracy (BMJ Open).
15. Abramoff et al. (2016) integrated deep learning techniques with public datasets, significantly improving automated retinopathy detection (Investigative Ophthalmology & Visual Science).
16. Das et al. (2018) presented a deep learning-based classification system for determining the stage of diabetic retinopathy, demonstrating clinical utility (IEEE IEMCON Proceedings).
17. Al-Bander et al. (2017) utilized a DenseNet-based CNN architecture for fine-grained classification of diabetic retinopathy grades, improving diagnostic granularity (Computer Methods and Programs in Biomedicine).
18. Bhaskaranand et al. (2016) created an automated platform for diabetic retinopathy screening and progressive tracking using retinal images (Investigative Ophthalmology & Visual Science).
19. Quellec et al. (2010) introduced an adaptive nonseparable wavelet transform for content-based image retrieval, which was later applied in retinal image analysis (IEEE Transactions on Image Processing).
20. Gulshan et al. (2017) extended their research by developing a deep learning model for detecting a broader range of diabetic eye diseases, validating it through rigorous clinical trials (PNAS).