



IJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** IV **Month of publication:** April 2026

DOI: <https://doi.org/10.22214/ijraset.2026.80027>

www.ijraset.com

Call:  08813907089

E-mail ID: ijraset@gmail.com

Deep Learning for Diabetic Retinopathy Detection from Retinal Images

Maku Akhila¹, Maanas Kaniti², Induri Aishwarya³, Dr. S. Sreekanth⁴

^{1, 2, 3}Dept. of CSE (Data Science), Institute of Aeronautical Engineering Hyderabad, India

⁴Associate Professor, Dept. of CSE (Data Science) Institute of Aeronautical Engineering Hyderabad, India

Abstract: Diabetic retinopathy (DR) is a progressive microvascular complication of diabetes and is today considered a major cause of avoidable blindness on a global scale. This should be detected at a young age, but the conventional approach to screening uses manual interpretation of the retina fundus images, which is laborious, subjective, and hard to automatically do on a large population. The proposed project presents an automated deep learning model based on the use of sophisticated convolutional neural networks (CNNs) to identify stage of DR precisely using the retinal images. It starts with the creation of a heterogeneous dataset and proceeds to a powerful preprocessing step comprised of normalization, contrast enhancement, image re-scaling, and illumination removal to enhance the appearance of lesions. More data augmentation methods including rotation, flipping, zooming, noise infusion also improve the generalization of the models and deal with the issue of class imbalance. A discriminative retinal feature extractor with a small CNN architecture, built on the latest models, like EfficientNet or ResNet, is trained to recognize images according to various levels of severity of DR. The diagnostic reliability is evaluated with rigor by accuracy, precision, recall, F1-score, sensitivity, specificity, and ROCAUC on model performance to be sure that the model performance is reliable. The resulting trained model is embedded into a user-friendly web-based application which enables clinicians to provide retinal images and get quick and evidence-based predictions with heatmap visualizations to read between the lines. This system provides better screening of DR before symptoms, encourages diagnosing ophthalmology at scale, and shows a promising influence of deep learning on medical image processing.

Index Terms: Diabetic Retinopathy, Deep Learning, EfficientNetB3, Retinal Fundus Images, Medical Image Classification, APTOS 2019, Grad-CAM, Computer-Aided Diagnosis.

I. INTRODUCTION

DIABETIC Retinopathy (DR) is One of the most frequent and vision threatening microvascular complications of diabetes mellitus is known as retinopathy (DR), and is considered to be one of the leading causes of preventable blindness in working-age adults around the globe. With the current and expected increase in the prevalence of diabetes both globally and within the individual countries, there has been a significant increase in the burden of DR screening and diagnosis, posing significant challenges to the healthcare systems across many countries especially with limited resources. One of the biggest clinical issues is the fact that DR does not always have any symptoms during the early phases, and patients did not notice any significant change in their vision until the disease reached more serious conditions. Therefore, early detection and early intervention is critical to minimize the potential of irretrievable retinal damages and permanent blindness. Traditional DR screening requires the use of a manual analysis of retinal fundus images, as an evaluation process that may be lengthy, costly, and prone to inter-rater disagreement. The limitations are even more prone in large-scale screening programs, rural medical settings, and those in developing nations in which availability of specialists is scarce. Hence, the automated, precise and scalable computer-aided diagnosis systems have emerged as a significant research focus in the ophthalmic image diagnosis [1], [9], [12], [17], [25], [35].

Deep learning based on the rapid progress has significantly helped revolutionize the concept of medical image analysis and demonstrated impressive promise in screening of retinal disease, localization of lesions, screening of diseases, and prediction of medical prognosis. The initial development of computer vision was propelled by deep convolutional neural network (CNN) models like AlexNet that showed the strength of hierarchical feature learning using huge amounts of visual evidence [5]. Better architectures that followed like ResNet and DenseNet also featured over features propagation, residual learning, and dense connectivity, that provided deeper and stronger networks which have played a significant role in bettering performance on challenging image classification challenges [6], [7]. Those developments formed the basis of the implementation of CNN-based methods in ophthalmology, where minor pathological structures, such as microaneurysms, hemorrhages, hard exudates, and cotton wool spots, have to be detected in the retinal fundus. Deep learning models, including the study by Gulshan and colleagues, have

shown that retinal fundus images could be identified as having referable diabetic retinopathy with near-perfect sensitivity and specificity, thereby setting one of the first and the most impactful milestones in the research area regarding to deep learning models [1]. Similarly, Ting et al. had already confirmed a deep learning platform to DR and other retinal diseases in multiethnic sample subjects, which further confirms that clinically applicable AI-based retinal screening systems are a possibility [9].

The subsequent developments of these foundational incidences led to further investigations in the field of DR fighting beyond binary screening to advanced multi-class grading, lesion-sensitive classification, disease progression modeling and real-world clinical practice. Widescale models like DeepDR have shown the validity of pipelines which combine the image quality appraisal, lesion detection, and tissue level severity in the entire diabetic retinopathy [12]. To a great extent, such systems are necessary due to the fact that real-world data on retinal screening can be heavily heterogeneous in terms of image quality, type of cameras, field of view, and demographics of a patient population. Besides, new extensive validity studies have revealed that deep learning-based DR algorithms may display good results in laboratory conditions, as well as in real-world screening procedures. As an illustration, clinical deployment-based tests in India showed that automated DR screening systems could be highly sensitive and specific to severe DR and had an infinitely low risk of clinically important missed cases, thus portraying the adaptability of AI to large population screening [35]. All these studies lead to the overall consideration that deep learning has grown out of the proof-of-concept models and became clinically relevant decision-support systems [1], [9], [12], [35].

A large literature base in the recent past also supports the fact that nowadays the paradigm of deep learning has become the leading paradigm of DR detection and grading on the basis of retinal fundus images. Numerous reviews and survey papers all comment that transfer learning and deep CNN-based tools achieve superior performance to much of more traditional handcrafted-feature pipelines, particularly when trained on large public datasets like EyePACS, APTOS 2019, Messidor, IDRiD, DIARETDB, and others [13], [19], [20]. The systematic reviews have also indicated that the current research in DR is at an earlier stage of exploring ways to enhance its robustness, generalization, computational efficiency, and clinical interpretability as opposed to standing on high quality on a single dataset [13], [19], [20], [25]. The recent surveys are also focusing more on the advanced preprocessing, data augmentation, optimized feature selection, federated learning, lightweight mobile-friendly architecture, and multimodal fusion strategies, which integrate retinal imaging with one of the other sources of data, e.g., OCT or structured clinical data [17], [19], [20], [25], [27]. This development covers a more general transformation into deployable, explainable AI systems of ophthalmic care in the form of more experimental models.

Within the framework of the methodological development, there have been a number of current studies that have put forth various deep learning strategies that can enhance the performance of DR classification. The combination of both handcrafted and deep representations has been investigated in some works and demonstrates improvements on discrimination with combined feature pipelines across the stages of DR severity [16]. There are other studies which have suggested optimized feature selection schemes to eliminate replicas and improve on classifications [21]. Better activation functionality is also examined in order to boost the levels of discriminative ability of CNN-based DR models, and it is observed that their performance is strong on the Kaggle retinal image datasets [14]. More complicated spatial or contextual patterns on retinal images have been proposed by using hybrid architecture e.g.: CNN-RNN models and encoder-decoder or attention-based models [18], [29]. The systems with sequential preprocessing pipelines enabled by transfer learning have demonstrated that image enhancement and finetuning of models can be significantly enhanced to become far more robust in the context of multi-class DR grading [28]. Further advances have also studied the hybrid segmentation-classification systems and refined architecture mix to implement the stage-wise DR identification [26], [30]. Taken altogether, these studies indicate the fact that research in DR detection is no longer in the standard CNN classification but in more feature learning paradigm, architecture optimization and the design with clinical focus.

The second trend that is also significant in the literature is the diversification of the imaging setting and application situations. Although traditional color fundus photography has been the center of most studies, more studies have lately advanced to smartphone-based retinal screening which is most notable in low-cost screening and being portable in the underserved areas [10]. The given direction is especially important in the countries where the availability of special fundus cameras and specialists in retina is low. Likewise, the attention has been laid on the ultra-widefield fundus visualization and diabetic macular edema detection as the researchers seek to present a wider field in the retina and deal with its related retinal complications [24]. There are also some recent systems which have explored the joint analysis of diabetic retinopathy and glaucoma and realize that co-occurring ocular conditions can affect screening processes and decision support systems [31]. In addition to the use of images in grading, newer research is looking at predicting disease outcomes and individual screening periods. DeepDR Plus is a time-to-progression model that uses baseline retinal images to directly model the risk of future DR progression which provides the potential of an individualized follow-up schedule and better screening resource allocation [22].

In addition, multimodal concepts involving integration of fundus photographs alongside OCT or any other clinical modalities are also developing as the future trend in the context of early-stage prediction and increased diagnostic sensitivity [27], [37]. Although the development of the DR detection based on deep learning is impressive, there are a variety of challenges. First, retinal datasets are in many cases, unbalanced in the disease severity classes with mild and severe stages being underserved compared to normal or moderate ones. Second, generalization to different datasets can be reduced because of differences in image quality, illumination, and acquisition equipment, and field of view. Third, black-box behavior is another strong forcing factor that does not facilitate clinical adoption, ophthalmologists are looking for interpretable evidence to place their trust on automated predictions. Fourth, even high-performing models can still be computationally expensive which means that they can not be used in real-time or low-resource settings. The challenges drive the desire of architectures providing high levels of representational power at the same time being computationally efficient and transfer learnable. In this regard, EfficientNet has also become a very appealing model family. EfficientNet presents a compound scaling approach, which to the network, combines depth, width, and input resolution for a better accuracy-efficiency tradeoff than traditional CNN designs [2]. The property is particularly useful in medical image analysis, where representational sensitivity and computational practicality is required due to the high-resolution images and the requirement of subtle features of lesions. Variants of EfficientNet have thus become progressively popular in the fields of ophthalmic imaging, or any other medical use of models where the efficiency of the model is no less significant than the predictive performance [2]. EfficientNetB3 is one of the EfficientNet family that has a good trade off between accuracy, number of parameters as well as the computation cost and the effect of this is applicable in retinal fundus image classification. This is specifically relevant in DR detection since the retinal lesions can be small, diffuse and can be hidden in visual detail and the network must be able to maintain discriminative detail at different resolutions. EfficientNetB3 can successfully acquire disease-relevant retinal features and can be trained in a way that is feasible to deploy at large scale since used together with appropriate preprocessing and augmentation strategies. Secondly, explainability has also become a fundamental element of the modern medical AI systems. Such methods as Gradient-weighted Class Activation Mapping (Grad-CAM) allow visualising the areas in an image that have the strongest effect on network predictions and, in this way, allow clinicians to verify that the network is targeting pathologically relevant structures in the retina [4]. These explainable AI systems are increasingly being considered as a requirement in establishing trust between clinicians and to facilitate model validation as well as to coordinate automated systems with human diagnostic processes. Inspired by these advancements, this work will provide a proposal of an automated diabetic retinopathy system based on EfficientNetB3 to detect retinal fundus images in the APTOS 2019 Blindness Detection dataset [3]. APTOS 2019 data is not only popular on recent DR research due to labeled retinal fundus images at various stages of severity, but also as form of a practical benchmark score on multi-class classification [3], [19], [20]. The suggested structure builds on the great representative power and parameter efficiency of EfficientNetB3, as well as the corresponding preprocessing and augmentation considerations, to enhance the classification and generalization. In addition, there is the inclusion of Grad-CAM-based visualization in order to improve the interpretability of the model and to give the qualitative information regarding the retinal regions contributing to the model predictions [4]. This study will apply the thorough foundational clinical research, large-scale deployment-oriented systems, recent systematic review research, and modern deep learning-based models in the detection of diabetic retinopathy at its onset to provide a robust, efficient and clinically meaningful solution to the current problem. It is hoped that the proposed approach does not just enhance the performance of automated DR classification but contributes to a wider objective, namely scalable, interpretable, and practically deployable AI-assisted ophthalmic diagnosis.

II. LITERATURE SURVEY

Deep learning in diabetic retinopathy (DR) detection has expanded swiftly in the last ten years and has developed over time to include binary screening systems, multi-class grading, lesion-sensitive classification, and clinically applicable decision-support systems. The success of deep convolutional neural networks (CNNs) in large-scale image recognition tasks had a strong impact on early breakthroughs in the field. They established the power of hierarchical feature learning through foundational architectures like AlexNet, ResNet, and DenseNet that had a substantial impact on the creation of the retinal image analysis systems in ophthalmic diagnosis [5]– [7]. Concerning DR, the study by Gulshan et al. was the first published work to present a clinically relevant deep learning model that identifies referable diabetic retinopathy with high sensitivity and specificity on retinal fundus photographs and formed a benchmark in future studies [1]. On the same note, Ting et al. showed that deep learning models could be generalized to multiethnic groups and similar retinal pathology, which once again confirms the clinical potential of AI-based retinal screening [9].

After these initial investigations, scientists started to pay more attention to enhancing robustness, scalability, and disease-stage discrimination. Large systems like DeepDR went beyond mere classification to include image quality, lesion and severity grading throughout the entire DR spectrum [12]. These systems emphasized the need to have end-to-end automated pipelines that can process heterogeneous real-world fundus images. More recent studies of real-world deployment have also demonstrated a high level of clinical relevance. As an example, the extensive screening assessment in India has shown that the deep learning-based DR algorithm can attain high sensitivity and specificity under the real-world screening setting and can be used in the screening of the population in terms of ophthalmic care [35]. All these studies signify that deep learning is no longer a phase of proving of concept but a system of clinically significant diagnostic assistance [1], [9], [12], [35].

A significant portion of recent survey and review articles offers an extensive review of the current situation in DR detection with deep learning. In the review articles, it is always reported that the deep CNNs and transfer learning models are dominating the world as they can learn the discriminative retinal features using small medical data [13], [19], [20]. EyePACS, APTOS 2019, Messidor, IDRiD and DIARETDB are the common datasets that can be utilized as the benchmarks to binary and multi-class DR grading [3], [13], [19], [20]. Other significant directions of research mentioned in these reviews include transfer learning, hybrid feature extraction, optimized preprocessing, federated learning, lightweight architecture, multimodal fusion as well as explainable AI [17], [19], [20], [25], [27]. Systematic studies also observe that although most of these methods report high accuracy on benchmark datasets, performance across different acquisition devices, patient groups, and image quality conditions are a major problem [13], [19], [25].

Recent original work has covered a vast variety of model architectures and optimization strategies to enhance the performance of DR classification. Other researchers have also used handcrafted retinal features along with deep CNN features to improve stage-based discrimination, and they have claimed high accuracy in their classification among different levels of DR severity [16]. Other publications presented optimization techniques in feature selection to enhance classification efficiency and minimize the redundant or noisy representations [21]. Better activation functions have also been suggested to enhance the discriminative power of the deep CNN models, and good results are achieved on Kaggle retinal image datasets [14]. Complex spatial dependencies and patterns of retinal lesions have been learned using hybrid architectures, including CNN-RNN models and encoder-decoder or attention-based architectures [18], [29]. Pipelines based on transfer learning and sequential processing and refining of features have shown that image enhancement and fine-tuning approaches can be used to enhance performance on multi-class grading of challenging retinal datasets [28]. Moreover, recent works have suggested hybrid deep learning systems, enhanced segmentation-classification hybrids, and multi-stage detection pipelines to more effectively address the complexity of retinal pathology [26], [30]. These articles show that recent DR research is increasingly focused on architectural art, enhanced representation of features, and resilience to real-life variation. The other notable trend in the literature is the expansion of DR detection to go beyond standard fundus image classification. Deep learning-assisted diagnosis in less-resourced and underserved areas has also been explored through smartphone-based retinal imaging as a low-cost and portable alternative to screening in resource-constrained areas [10]. The use of ultrawidefield fundus imaging and diabetic macular edema related activities have also been considered, where a wider retinal coverage can enhance the identification of peripheral lesions and other related retinal abnormalities [24]. Recent publications have also examined the concomitant analysis of diabetic retinopathy with other ocular diseases like glaucoma and have identified the significance of integrated ophthalmic decision-support systems [31]. Furthermore, the concept of predicting future DR progression and personalized screening intervals based on baseline retinal images has been presented in the progression-aware system like DeepDR Plus, which is a major step forward compared to the state-of-the-art in terms of static severity classification [22]. Multimodal deep learning models that integrate fundus images with OCT or structured clinical data are also beginning to be implemented as a promising solution to early-stage detection and better clinical decision-making [27], [36], [37].

EfficientNet is among the newly popularized deep learning models due to its compound scaling mechanism, which equalizes the network depth, width, and resolution to providing the additional accuracy with the modified computation rate in comparison to the majority of conventional CNN paradigms [2]. This is what makes EfficientNet particularly attractive to the medical image analysis, where high-resolution retinal images, and fine-pattern lesions require not only the representational capacity but also the computing efficiency. EfficientNet variants, like EfficientNetB3, are more beneficial to feature DRs, whereas they are able to discover fine-grained pathological detail including microaneurysms, hemorrhages and exudates, at a workable scale of the model. Good performance and deployment capabilities also give EfficientNetB3 a good trade-off, especially with multi-class classification tasks on datasets such as APTOS 2019 [2], [3], when compared to heavier CNN backbones. Moreover, the interpretability of models has also been gaining popularity in medical AI.

The Grad-CAM has emerged as one of the most widely used post-hoc explainability algorithms to visualise discriminative areas using retinal images hence contributing to support validation that the model is interested in clinically significant lesions and increasing confidence in automated predictions [4]. Despite the significant gains made, there are still a number of gaps in research. Most of the literature has high accuracy on curated benchmark datasets, but fail to adequately cover class imbalance, inter-dataset generalization, computational efficiency, and interpretability in one unified framework [13], [19], [25]. There are also high-performing approaches based on hybrid compute architectures or computationally costly pipelines, potentially constraining their use in practice in real- world screening systems [18], [29], [30]. More so, not every model is accompanied with visual explanations that can be used to validate predictions in clinical settings [4]. Hence, an effective and computationally efficient DR detection system, capable of carrying out an effective multi-class classification and simultaneously deliver interpretable results is still required. This gap inspired the current study, which uses EfficientNetB3 on APTOS 2019 Blindness Detection data with proper preprocessing and augmentation and uses Grad-CAM to obtain explainable diabetic retinopathy classification [2]- [4]. The combination is designed to offer a practical tradeoff between accuracy, efficiency and interpretability to real world computer-aided retinal screening.

III. PROPOSED METHODOLOGY

This section describes the overall methodology followed for automated diabetic retinopathy (DR) detection using the EfficientNetB3 deep learning model. The proposed framework is designed to classify retinal fundus images into multiple diabetic retinopathy severity levels by combining image preprocessing, data augmentation, transfer learning, and explainable AI techniques. The complete workflow of the proposed system is illustrated in Fig. 1.




Fig. 1. Overall workflow of the proposed diabetic retinopathy detection system.

A. Dataset Description

The given model is trained and tested with the help of the APTOS 2019 Blindness Detection dataset, which is a publicly accessible retinal fundus image dataset that is commonly utilized in the research on diabetic retinopathy classification. The data comprises color fundus photos recorded in different imaging conditions, which is why it can be used to assess the resilience of deep learning models when it comes to real-life screening. All the images in the dataset are categorized into five types of diabetic retinopathy severity no DR, mild, moderate, severe and proliferative DR. As the dataset has images of various resolution, illumination and quality variations, it is a demanding benchmark on multi-class classification problems. To develop the model, the dataset is divided into training and validation sets. Supervised learning is performed with the help of the class labels, and the distribution of the images among the severity classes is taken into account during preprocessing and augmentation so that the impact of the class imbalance could be minimized. The reason why the APTOS 2019 dataset is chosen in this work is the fact that it is well-known in the recent literature and can be utilized to benchmark the transfer learning-based DR detection models [3].

B. Image Preprocessing

The images of the retinal fundus in the APTOS 2019 dataset are very variable in the size, brightness, contrast, and noise level since the imaging devices and contexts vary. A preprocessing pipeline is used to refine the input data and increase the quality of the features of the retina, prior to the process of training the model. To begin with, any input images will be downsized to a constant resolution that can be used by the EfficientNetB3 model. This creates consistency in the dimensions of the input and makes the training less complex to compute. The next stage is the normalization of pixel values to make the image intensity values within a similar range, and this will enhance the convergence during optimization. Moreover, contrast enhancing and noise removing algorithms can be used to enhance visibility of the relevant retinal structures like blood vessels, exudates, hemorrhages and microaneurysms. These steps are used in the preprocessing to ensure that the model concentrates on retinal patterns that are relevant in the disease but minimizes effects of irrelevant changes in the background. Circular cropping or region-of-interest centering may also be employed, as needed, to highlight the retinal disc area and eliminate the superfluous black margins that are so prevalent in the fundus images. Fig can be used to provide a sample of original retinal images as well as preprocessed retinal images




Fig. 2. Sample retinal fundus images before and after preprocessing.

C. Data Augmentation

The deep learning models usually need a vast and varied training to generalize effectively. However, medical image data sets like APTOS 2019 can be imbalanced in terms of classes, and minimal fluctuation in some phases of disease. To address data augmentation is used to the training in these challenges, data to the extent that it is artificial and tends to de-diversify data, overfitting.

Other common augmentation methods are random rotation, horizontal flipping, zooming, shifting, making the picture brighter, and minor affine transformations. These transformations simulate realistic variations of image acquisition and maintain the critical structures of the retina and pathology. Data augmentation assists the model in growing stronger to change in direction, light, and size, and thus enhancing its capacity to generalize to invisible retinal images. By increasing the powerful heterogeneity of the training data, enrichment as well, helps reduce the effect of class imbalance and aids greater learning of minority DR severity classes.

D. EfficientNetB3 Architecture

The central component of the suggested system is the EfficientNetB3. The primary is referred to as convolutional neural network, diabetic retinopathy severity feature extractor and classifier, prediction. EfficientNet is selected due to the fact that it employs a compound scaling strategy which scales the network depth, width, jointly, and inject resolution with a balanced effect, causing it to outperform most parameters with very minimal parameters, conventional CNN architectures [2].

In this work, the EfficientNetB3 is used by means of transfer learning, in which the model is primed using pre-trained ImageNet weights. However, with pretrained weights, the network utilizes to capitalize on already acquired low and mid-level visual characteristics, especially useful in case of working with medical datasets of limited size. The last classification layer the pretrained network is substituted with task-specific fully associated layer with the five diabetic retinopathy severity classes.

The softmax activation function is employed in the. output layer to give multi-class class probabilities. classification. To improve regularization and reduce overfitting, additional layers such as dropout may be incorporated before the final classification layer. The architecture enables the model to learn discriminative retinal features while maintaining computational efficiency, making it suitable for practical deployment in computer-aided screening systems. A simplified architecture diagram of the proposed EfficientNetB3-based classifier can be shown in Fig. 3.

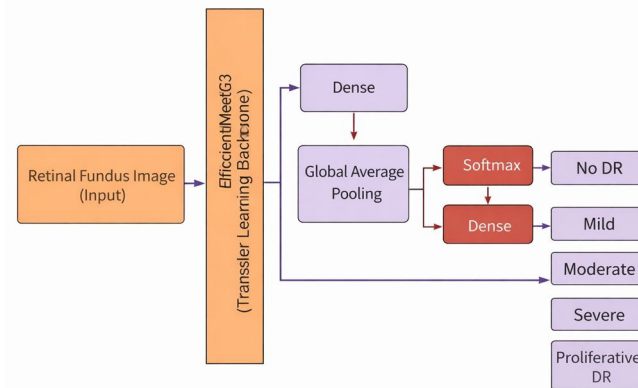


Fig. 3. Simplified architecture of the proposed EfficientNetB3-based diabetic retinopathy classifier.

E. Training Procedure

In order to better regularize and minimize overfitting, more layers can be added like the dropout before the last classification layer. The model is able to do so with the help of the architecture. discriminate retinal features with computational efficiency, and thus applicable to practical application. in computer aided screening systems. A simplified architecture Training Procedure diagram of the proposed EfficientNetB3- based classifier may be.

The model is a multi-class classification system that is trained as a supervised one. cation system with the labeled retinal fundus images of. the APTOS 2019 dataset. The training is done on the processed. and augmented images are processed with the EfficientNetB3. network, and the estimated class probability are compared. train the ground truth labels with an appropriate categorical loss. function, e.g. categorical cross-entropy or sparse categorical cross-entropy, based on the format of label encoding. Adaptive optimization algorithm like Adam is applied. to update the model parameters in the process of backpropagation. The The learning process is implemented in several epochs with a. batch size is defined and performance of validation is measured. to evaluate generalization. To avoid stability and overfitting. lize training, strategies such as early stopping, learning rate scheduling, and model checkpointing can be used. These techniques can be used to make the final model keep the best. doing weights using validation accuracy or validation loss.

Evaluation of the trained model is based on the performance. standard classification measures like accuracy, precision, recall, F1-score, and analysis of confusion matrix. Since diabetic retinopathy is a multi-class medical classification problem, class-wise analysis is significant to estimate the model. is consistent at all levels of severity.

F. Grad-CAM Explainability

Model interpretability is needed to analyze medical images. inculcate trust in automated forecasts. To improve the openness of the proposed system, Grade-based Class. The post-hoc explainability technique is called Activation Mapping (Grad-CAM). plainability technique [4]. The heatmaps produced by grad-cam are used to show the areas of the retinal image that are the most significant to the model in its decision-making using a predicted class. In the proposed structure, the Grad-CAM is implemented reasoned on the model inference on the selected retinal fundus. images.

The resulting activation maps are superimposed on original fundus images in a visual manner to determine whether the model is concentrated on clinically relevant retinal distortions such as microaneurysms, exudates, hemorrhages, or other foci of lesions. This is in order to guarantee. that as the model is learning, it is learning good pathological patterns as against the employment of the haphazard image artifacts. Grad-CAM can increase system interpretability and facilitate its application is a possibly explainable computer aided diagnostic tool in clinical practice. Sample Grad can be represented with the help of Fig. 4. CAM visualizations.

IV. RESULTS AND DISCUSSION

Proposed EffectNetB3-based diabetic retinopathy (DR) class signification framework of the APTOS 2019 Blindness Detection dataset. The model is tested with the help of standard multi-class. classification measures, such as accuracy




Fig. 4. Sample Grad-CAM visualization showing the retinal regions influencing model prediction.

precision, recall, F1. score, confusion matrix analysis and visual interpretability. using Grad-CAM. In addition, the training and validation per accuracy and loss are used to analyse performance of the model. curves to measure convergence behavior and generalization. potential of the suggested framework.

G. Training and Validation Performance

The training behaviour of the proposed EfficientNetB3 model. was examined in terms of epoch-wise accuracy and loss curves. The compound training and validation accuracy are depicted in Fig. 5. and loss trends, which give the learning progression and convergence properties of the model in training. It can be seen that the accuracy and loss plots indicate convergence behavior of the model during training.




Fig. 5. Combined training and validation accuracy and loss curves of the proposed EfficientNetB3 model.

That the model is only learnt with progressively more training epochs. and approaches the level of steady performance. A close congruence between the training and validation curves suggests that the model is reasonably well generalized and not. suffer from severe overfitting. The following observations indicate that the preprocessing, data augmentation, and transfer applied. learning strategies add a positive impact on the learning stability. of the EfficientNetB3-based classifier

H. Overall Classification Performance

The suggested EfficientNetB3 model obtained a total. On the test set, classification accuracy of 91 on a cumulative basis of 1434 retinal fundus images. In addition to accuracy, the model performed a macro-average precision of Micro-average recall of 0.91, macro-average recall of 0.92 and macro-average F1. score of 0.92. On the same note, the weighted-average precision, recall, and F1-score were found to be 0.92, 0.91, and 0.92, respectively. Such findings show that the model works coherence between classes and has high over all predictive capability Table I.

TABLE I
OVERALL PERFORMANCE METRICS OF THE PROPOSED EFFICIENTNETB3 MODEL

Metric	Value
Accuracy	91%
Macro Precision	0.92
Macro Recall	0.91
Macro F1-Score	0.92
Weighted Precision	0.92
Weighted Recall	0.91
Weighted F1-Score	0.92
Test Samples	1434

I. Class-wise Performance Analysis

An overview of the total performance metrics is shown in Table I. To gain a better insight into the behavior of the proposed model, with varying levels of diabetic retinopathy severity, class-wise. The values of precision, recall, F1-score and support were evaluated, using the classification report. The APTOS 2019 dataset Classes were viewed in the following way: Class 0- No DR, Class;—human—; Classes were viewed in the following way: Class 0- No DR, Class;—human—; 1- Mild, Class 2- Moderate, Class 3- Severe, and Class 4- Proliferative DR.

The class-wise performance metrics are shown in Table II.

TABLE II
CLASS-WISE PERFORMANCE METRICS

Class	Precision	Recall	F1-Score	Support
No DR (0)	0.97	0.97	0.97	287
Mild (1)		0.90	0.90	286
Moderate (2)	0.79	0.90	0.84	287
Severe (3)	0.97	0.95	0.96	287
Proliferative DR (4)	0.95	0.86	0.90	287

Table II presents the performance metrics in classes. The findings indicate that this model is very good, in No DR and Severe DR, and the Class Moderate DR, is relatively difficult, owing to the visual resemblance, between adjacent severity stages. Nonetheless, the high recall for Moderate DR indicates that the model is effective in determining the majority of clinically relevant cases.

J. Confusion Matrix Analysis

To further assess the prediction performance of the classes of A confusion matrix was analyzed in the proposed model. Fig. 6 displays the confounding matrix of the suggested EfficientNetB3 based classifier.

The confusion matrix confirms that the majority of samples are correctly classified along the diagonal, indicating strong overall performance. However, some misclassifications occur between visually similar stages such as Mild and Moderate DR, as well as Moderate and Proliferative DR. This is consistent with the class-wise performance metrics and reflects the inherent difficulty of multi-class retinal disease grading.

K. Grad-CAM Visualization and Interpretability

In order to enhance the transparency of the proposed framework, Two images of retinal fundus were selected and subjected to Grad-CAM. Grad-CAM produces heatmaps which




Fig. 6. Confusion matrix of the proposed EfficientNetB3-based diabetic retinopathy classification model.

point to the re-retinal. gions replying most to the predicted class, permit visual interpretation of the decision-making of the model. process.

The Grad-CAM images show that the model fo- cuses on the areas of the retina with lesions, abnormal areas included. related to exudates, hemorrhages, microaneurysms, and other pathological patterns. This enhances the interpretability. of the system and raises its practical applicability in clinical. when it is significant to have explainable predictions in environments. building trust in AI-assisted diagnosis.

L. Summary of Findings

Altogether, the experimental outcomes can prove that the proposed EfficientNetB3-based diabetic retinopathy detection framework is efficient, powerful, and applicable to classify multi-classes of retinal images. The model is significantly per- forming well, the training and validation curves are stable and give interpretable results using Grad-CAM. Although there is still some muddiness between the intermediate severity levels, the findings show that the suggested framework and framework have a high potential to be used as a workable computer-aided diagnostics tool in the screening and measurement of diabetic retinopathy severity at its early stages.

V. CONCLUSION

In this paper, an EfficientNetB3 deep learning model was proposed. to diagnose and categorize diabetic retinopathy in retinal fundus. sample images of the APTOS 2019 Blindness Detection dataset. The suggested system integrated image pre- processing, data transfer learning and Grad-CAM-based inter augmentation. pretability in order to introduce an effective and trustworthy multi- class. severe diabetic retinopathy severity classification model assess. ment. The experimental findings showed that the proposed model has been shown to do well in terms of classification accuracy, having an overall accuracy at 91At 0.92, macro and weighted F1-scores at 0.92. The model demon had a high performance in the classification of No DR and Severe. DR categories, as well as retaining a high level of success in the remaining phases of diabetic retinopathy. Despite the fact that the precision of the Moderate DR class was rather low due to the following reason the levels of severity were visually similar to each other, the model was able to recall many cases of the category, which means that it was able to capture the vast majority of clinically meaningful cases.

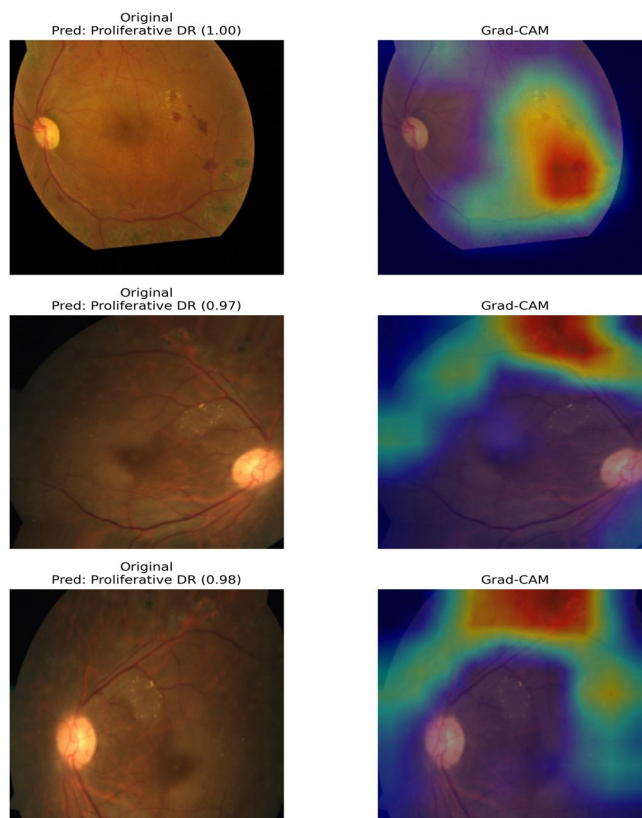


Fig. 7. Sample Grad-CAM visualizations highlighting retinal regions influencing model predictions.

EfficientNetB3 was found to be effective due to its rich features of extracting retinal features and at the same time be computationally efficient. Also, Grad-CAM integration enhanced the explanation of the framework as the retinal areas that affected the model predictions were emphasized and the system was more transparent and clinically significant. All in all, the proposed solution proves that EfficientNetB3 is an appropriate and viable architecture when it comes to multi-class diabetic retinopathy detection, and can be a viable computer-aided diagnostic solution to early screening of retinal diseases.

VI. FUTURE SCOPE

Even though the proposed EfficientNetB3-based system of diabetic retinopathy detection showed good results, there are a number of directions, in which this work can be enhanced and expanded. The use of larger and more heterogeneous retinal data sets is one of the aspects of future research that can be used to enhance the generalization of the model across a variety of imaging devices, populations, and clinical conditions. Testing the model on external data sets like EyePACS, Messidor or IDRiD would give a more solid insight on its practical strength.

Additional potential enhancements could be to use more sophisticated class balancing methods like focal loss, weighted loss functions, oversampling or synthetic image generation to further enhance performance on visually challenging intermediate classes like Moderate DR. Ensemble learning methods involving EfficientNetB3 with other deep learning models could also be used to enhance classification stability and overall predictive performance.

It is also possible to investigate multimodal diabetic retinopathy detection in the future with retinal fundus images and other clinical data or OCT images. These multimodal systems can enhance timely detection and aid in the more in-depth evaluation of illness. Moreover, the system may be more appropriate in rural or resource-limited healthcare facilities due to lightweight model optimization and deployment to mobile or edge devices.

Lastly, the explainability aspect can be further improved by including more sophisticated visual interpretation and uncertainty estimation methods than Grad-CAM. This would also enhance the trust that clinicians have and contribute to the implementation of AI-assisted retinal screening systems in the real-world medical setting. These advancements allow building the proposed framework into a more robust, scalable, and clinically deployable automated diabetic retinopathy diagnosis solution.

REFERENCES

- [1] 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.
- [2] M. Tan and Q. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. International Conference on Machine Learning (ICML)*, 2019.
- [3] APTOS 2019 Blindness Detection Dataset, Kaggle, 2019. [Online]. Available: <https://www.kaggle.com/competitions/aptos2019-blindness-detection>
- [4] R. R. Selvaraju et al., "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 618–626.
- [5] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2012.
- [6] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [7] G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 4700–4708.
- [8] T. Voets, K. Møllersen, and L. A. Bongo, "Reproduction study using public data of: Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs," *PLOS ONE*, vol. 14, no. 6, 2019.
- [9] D. S. W. Ting et al., "Development and validation of a deep learning system for diabetic retinopathy and related eye diseases using retinal images from multiethnic populations with diabetes," *JAMA*, vol. 318, no. 22, pp. 2211–2223, 2017.
- [10] "Deep learning frameworks for diabetic retinopathy detection with smartphone-based retinal imaging systems," *Pattern Recognition Letters*, vol. 137, pp. 409–417, 2020.
- [11] "Diabetic retinopathy detection through deep learning," *MethodsX*, 2020.
- [12] "A deep learning system for detecting diabetic retinopathy across the disease spectrum," *Nature Communications*, vol. 12, no. 1, 2021.
- [13] "Deep learning for diabetic retinopathy detection and classification based on fundus images: A review," *Computers in Biology and Medicine*, 2021.
- [14] "Deep learning for the detection and classification of diabetic retinopathy with an improved activation function," *Healthcare*, vol. 11, no. 1, 2022.
- [15] "Diabetic retinopathy detection from fundus images of the human retina using deep learning," 2022.
- [16] "Predicting of diabetic retinopathy development stages of fundus images using deep learning based on combined features," *PLOS ONE*, 2023.
- [17] "Deep learning and medical image processing techniques for diabetic retinopathy: A survey of applications, challenges, and future trends," 2023.
- [18] "A new hybrid deep learning model for diabetic retinopathy detection," 2023.
- [19] "A systematic review on diabetic retinopathy detection and classification based on deep learning techniques using fundus images," *PeerJ Computer Science*, 2024.
- [20] "Comprehensive survey of the state-of-the-art deep learning models for diabetic retinopathy detection and grading using retinal fundus photography," 2024.



- [21] "Diabetic retinopathy detection using deep learning with optimized feature selection," 2024.
- [22] "A deep learning system for predicting time to progression of diabetic retinopathy," *Nature Medicine*, 2024.
- [23] "Enhancing diabetic retinopathy diagnosis: A lightweight CNN for efficient classification," arXiv preprint, 2024.
- [24] "Deep learning-based detection of referable diabetic retinopathy and diabetic macular edema in ultra-widefield fundus images," arXiv preprint/ MICCAI Challenge paper, 2024.
- [25] "Systematic review and meta-analysis of regulator-approved deep learning systems for fundus diabetic retinopathy detections," *npj Digital Medicine*, 2025.
- [26] "Diabetic retinopathy screening using machine learning," 2025.
- [27] "Deep learning-based optical coherence tomography and retinal imaging for diabetic retinopathy: A systematic review and meta-analysis," 2025.
- [28] "New accurate deep learning model for diabetic retinopathy detection utilizing sequential pre-processing and transfer learning," *Biomedical Signal Processing and Control*, 2025.
- [29] "A hybrid deep learning framework for early detection of diabetic retinopathy using retinal fundus images," *Scientific Reports*, 2025.
- [30] "Detection and diagnosis of diabetic retinopathy in retinal fundus images using agentic AI approaches," *Scientific Reports*, 2025.
- [31] "Deep learning-based joint analysis of diabetic retinopathy and glaucoma in retinal fundus images," *Scientific Reports*, 2025.
- [32] "A novel deep learning model for diabetic retinopathy detection," 2025.
- [33] "Enhancing diabetic retinopathy classification accuracy using deep learning," arXiv preprint, 2025.
- [34] "Development and validation of a deep learning model for early diabetic retinopathy detection," 2025.
- [35] "Performance of a deep learning diabetic retinopathy algorithm in India," *JAMA Network Open*, 2025.
- [36] "Deep learning for enhanced prediction of diabetic retinopathy: A comparative study on the diabetes complications data set," *Frontiers in Medicine*, 2025.
- [37] "Early prediction of diabetic retinopathy using a multimodal deep learning model," 2026.



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)