iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** V **Month of publication:** May 2026

DOI: https://doi.org/10.22214/ijraset.2026.81765

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue V May 2026- Available at www.ijraset.com

Diabetic Retinopathy Detection Using Efficient Net and Grad-CAM

Vema Eekshita Naga Venkata Lakshmi Sai[1], Dr. K. Chaitanya[3], Pandi Priyanka[3], Kattepogu Nagendra Babu[4], Nuthakki Prabhas[5]

[1, 3, 4, 5]*Students, Department of Data Science & Cyber Security, University College of Engineering & Technology, Acharya Nagarjuna University, Nagarjuna Nagar, Guntur, AP, India - 522510*

[2]*Assistant Professor, Department of Data Science and Cyber Security, University College of Engineering & Technology, Acharya Nagarjuna University, Nagarjuna Nagar, Guntur, AP, India - 522510*

Abstract: ***Diabetic retinopathy (DR) is a progressive retinal disease and a major cause of preventable blindness among diabetic patients worldwide. Early diagnosis is essential to avoid irreversible vision loss, yet manual screening methods are often time-consuming, inconsistent, and inaccessible in low-resource settings. This project presents an AI-based system for automated DR detection and severity grading using deep learning techniques applied to retinal fundus images from APTOS, Messidor, and EyePACS datasets. The model classifies DR into five clinically defined stages—ranging from No DR to Proliferative DR—by learning subtle pathological features such as microaneurysms, hemorrhages, and neovascularization. To enhance interpretability and clinical trust, Grad-CAM is integrated to generate heatmaps that highlight lesion-specific regions influencing the model's predictions. The system achieves 84% test accuracy with F1-scores of 0.93 for No DR and 0.90 for Proliferative DR. The system is optimized for deployment on CPU-based hardware, making it suitable for scalable and accessible screening in real-world environments. By combining diagnostic accuracy with explainability, this project demonstrates how AI can support ophthalmologists, improve early intervention, and reduce the global burden of DR-related vision impairment.***

Keywords: ***Diabetic Retinopathy, EfficientNet, Grad-CAM, Retinal Fundus Images, Deep Learning, Transfer Learning, Explainable AI, Medical Image Classification, Severity Grading.***

--- AI Assistance Statement ---
During the preparation of this work, the authors used ChatGPT and Claude AI for language refinement, grammar checking, and formatting assistance. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the final publication.

--- Conflict of Interest ---
The authors declare no competing interests.

--- Ethics Statement ---
This study used publicly available de-identified retinal fundus images from APTOS, Messidor, and EyePACS datasets. No institutional ethics approval was required as the research involved no human participants or animal subjects.

I. INTRODUCTION

Diabetes has become one of the most widespread chronic conditions globally, and one of its most serious complications is diabetic retinopathy (DR). This condition gradually damages the tiny blood vessels inside the retina and develops silently. Patients often show no obvious symptoms until the disease has already reached an advanced stage. For the hundreds of millions of people living with diabetes worldwide, DR represents one of the most common and most preventable causes of blindness.[1]

The challenge is not just the disease's prevalence but how difficult it is to catch in time. Traditional screening relies on trained ophthalmologists manually examining retinal fundus images—a slow, expensive process that is also prone to inconsistency, especially in rural or low-income areas where specialists are scarce or unavailable. In many regions, patients are only diagnosed once the disease has already progressed significantly, leaving fewer effective treatment options.

©IJRASET: All Rights are Reserved

In addition to improving diagnostic efficiency, automated diabetic retinopathy detection systems can significantly reduce the workload on healthcare professionals. With the increasing number of retinal screenings being conducted worldwide, the volume of image data has grown rapidly. Manually analyzing these images is not only time-consuming but also prone to inconsistencies due to human fatigue and subjectivity.

Furthermore, early-stage diabetic retinopathy often presents very subtle visual features that can be easily overlooked during manual examination. Automated deep learning systems can detect these subtle patterns more consistently, enabling earlier intervention. This can play a crucial role in preventing severe vision loss and improving patient outcomes.

Another important aspect is accessibility. In many rural and under-resourced regions, access to trained ophthalmologists is limited. An automated system that can run on standard hardware can act as a reliable screening tool, allowing early diagnosis even in such environments. This highlights the importance of developing solutions that are not only accurate but also practical and deployable.

A. Problem Statement

Diabetic retinopathy is one of the leading causes of blindness among diabetic patients, and early detection is critical to prevent severe vision loss. However, traditional screening methods rely on manual examination of retinal fundus images by trained ophthalmologists, which is time-consuming, costly, and not easily accessible in many regions. The situation becomes more challenging in rural and low-resource areas where the availability of specialists is limited.

Additionally, manual diagnosis is subject to variability and may lead to inconsistencies, especially in early-stage detection where symptoms are subtle. While deep learning models have shown promising results in automated detection, many existing approaches lack interpretability, making them difficult to trust in clinical settings.

Therefore, there is a need for an automated, accurate, and explainable system for diabetic retinopathy detection that can operate efficiently on standard hardware and assist healthcare professionals in early diagnosis.

B. Objectives

The main objectives of this research work are as follows:

To develop an automated system for detecting and classifying diabetic retinopathy using deep learning techniques.

To utilize EfficientNet-B3 for accurate feature extraction and classification of retinal images into five severity levels.

To improve model interpretability by integrating Grad-CAM for visual explanation of predictions.

To enhance model generalization by training on a combined dataset consisting of APTOS, Messidor, and EyePACS images.

To address class imbalance using data augmentation and class-weighted loss functions.

To design a system that can be deployed on CPU-based hardware for practical use in real-world environments.

C. Motivation

The primary motivation behind this work is to contribute toward early detection and prevention of vision loss caused by diabetic retinopathy. With the growing number of diabetic patients worldwide, there is an urgent need for scalable and efficient screening solutions.

Another key motivation is to bridge the gap between accuracy and interpretability in deep learning models. In medical applications, it is not sufficient for a model to provide predictions—it must also explain its decisions in a way that clinicians can understand and trust. The integration of Grad-CAM in this system addresses this need by providing visual insights into the model's decision-making process.

Furthermore, this work aims to develop a practical solution that can be used in real-world healthcare settings, especially in areas with limited resources. By designing a system that operates efficiently on standard hardware, this research seeks to make automated DR screening more accessible and impactful.

II. LITERATURE REVIEW

A 2016 study from Google Brain by Gulshan et al. trained a deep CNN on over 128,000 retinal images and demonstrated that the model could match or exceed the diagnostic accuracy of board-certified ophthalmologists, achieving AUC values of 0.991 and 0.990 on two independent validation sets [1].

Gargeya and Leng (2017) demonstrated that end-to-end deep learning pipelines could generalize well across datasets from different clinical environments, even without manual feature engineering—a critical finding for real-world deployment [2].

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

Selvaraju et al. (2017) introduced Grad-CAM, a technique for generating visual explanations of CNN predictions using gradient information from the final convolutional layer. In medical imaging, this interpretability is practically necessary. Clinicians need to know why the model predicts what it does, not just what it predicts [3].

Pratt et al. (2016) framed DR detection as a five-class severity grading problem aligned with clinical standards, using data augmentation to handle class imbalance—reinforcing its importance in imbalanced medical datasets [4].

Jabbar et al. (2024), in IEEE Access, proposed APSO-GRESNET, combining GoogLeNet and ResNet-16 with an Adaptive Particle Swarm Optimizer, achieving 94% accuracy on EyePACS while covering a broader lesion range [5]. However, this system lacked visual interpretability, a gap the present work directly addresses.

Tan and Le (2019) introduced EfficientNet, which performs compound scaling—simultaneously balancing network depth, width, and input resolution—achieving superior accuracy with fewer parameters than previous architectures [6].

III. METHODOLOGY

A. *Existing System*

Earlier approaches to DR detection relied on traditional machine learning methods—SVMs, random forests, and decision trees applied to handcrafted features from retinal images. These were limited to early-stage lesions like microaneurysms and exudates, missing severe-stage indicators such as cotton wool spots, venous beading, and IRMA. As deep learning emerged, architectures like AlexNet [8] VGGNet, ResNet [7], and Inception-v3 [9] significantly improved feature coverage and accuracy. Despite this progress, a critical gap remained: most existing systems function as black boxes, unable to provide clinicians with visual explanations for their predictions.

B. *Proposed System Overview*

The proposed system addresses two primary gaps simultaneously: limited lesion coverage and lack of interpretability. At its core, an EfficientNet-B3 model is fine-tuned on retinal fundus images to classify DR across five severity stages, learning to recognize a broad spectrum of pathological markers including microaneurysms, hemorrhages, hard exudates, and signs of neovascularization. Once a prediction is made, Grad-CAM generates a color heatmap overlaid on the original image, showing which retinal regions most influenced the decision. This visual explanation makes the tool both trustworthy and verifiable in a clinical context. The entire system is optimized to run without GPU infrastructure, ensuring it can be practically deployed in resource-limited screening environments. Initially, input retinal images are preprocessed by resizing and normalizing pixel values to ensure consistency across different datasets. Data augmentation techniques such as rotation, flipping, and zooming are applied to improve model generalization and address class imbalance.

The preprocessed images are then passed through the EfficientNet-B3 model, which serves as the backbone for feature extraction and classification. The model learns to identify complex retinal features associated with different stages of diabetic retinopathy and classifies each image into one of the five severity levels: No DR, Mild, Moderate, Severe, and Proliferative DR.

To enhance interpretability, Grad-CAM is integrated into the system. After the model makes a prediction, Grad-CAM generates a heatmap highlighting the regions of the retinal image that contributed most to the decision. These visual explanations help in validating the model's predictions and make the system more reliable for clinical use.

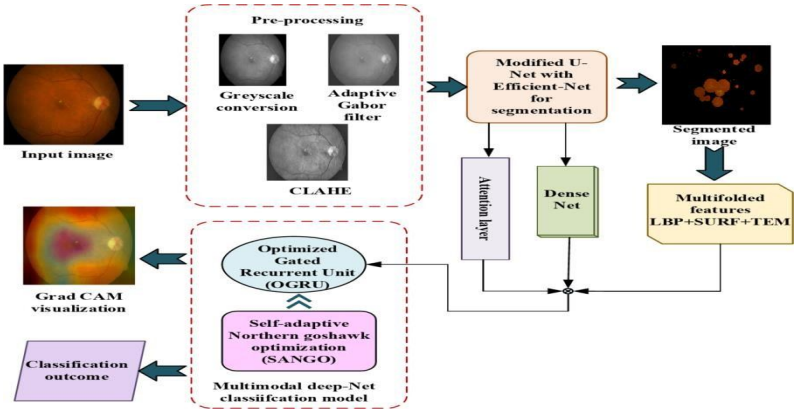


Fig 1: Proposed System of the DR

Fig.F1: WFForkflow of Grad-CAM

C. Advantages of Proposed System over Existing System

The proposed system offers several significant advantages compared to traditional and existing deep learning-based approaches:

1) *Improved Accuracy and Feature Learning:* Unlike traditional machine learning methods that rely on handcrafted features, the proposed system uses EfficientNet-B3, which automatically learns complex and high-level features from retinal images, resulting in improved classification performance.
2) *Better Generalization:* By training on a combined dataset of APTOS, Messidor, and EyePACS, the model is exposed to a wide variety of image conditions, making it more robust and capable of performing well on unseen data.
3) *Explainability through Grad-CAM:* Most existing systems act as black boxes, providing predictions without explanations. The integration of Grad-CAM allows the system to generate visual heatmaps, helping clinicians understand and trust the model's decisions.
4) *Handling of Class Imbalance:* The proposed system addresses class imbalance using data augmentation and classweighted loss functions, improving performance on minority classes such as Severe and Proliferative DR.
5) *Computational Efficiency*: EfficientNet-B3 is designed to achieve high accuracy with fewer parameters, making the model lightweight and efficient. This allows the system to run on CPU-based hardware without requiring high-end GPUs.
6) *Practical Deployability:* The system is optimized for real-world usage, especially in resource-limited environments. Its ability to operate on standard hardware makes it suitable for deployment in rural clinics and mobile screening setups.
7) *Consistency and Reliability:* Unlike manual diagnosis, which may vary between clinicians, the proposed system provides consistent and repeatable results, reducing variability in diagnosis.

D. EfficientNet-B3

EfficientNet is used as the backbone CNN architecture in this project. Unlike older architectures that scale up by adding layers arbitrarily, EfficientNet performs compound scaling—simultaneously balancing network depth, width, and input resolution in a principled way. EfficientNet-B3 specifically is used here because it offers a strong accuracy-to-efficiency trade-off, achieving high classification performance while using significantly fewer parameters than older models like VGG or ResNet. This efficiency is what makes it viable for CPU-based deployment without sacrificing meaningful accuracy.

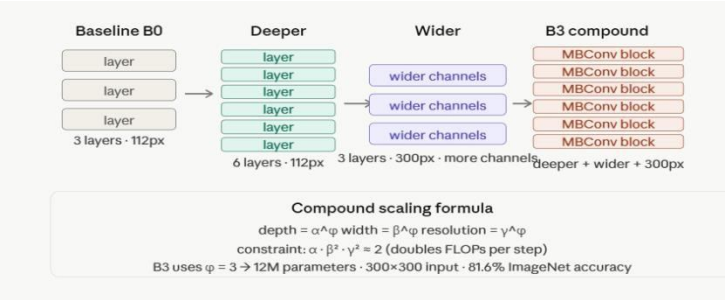


Fig.2: Architecture of EfficientNet B3

E. Grad-CAM Visualization

Grad-CAM (Gradient-weighted Class Activation Mapping) is used to make the model's predictions explainable. After each classification, the gradients flowing from the output with respect to the final convolutional layer are used to produce a weighted activation map. This map is then overlaid on the original retinal image as a color heatmap—warm colors (red/orange) highlight regions that most influenced the prediction, while cool colors indicate regions the model largely ignored. In retinal images, these highlighted regions typically correspond to known lesion locations such as hemorrhage clusters near the macula or abnormal vessel growth patterns.

F. Image Preprocessing & Data Augmentation

Before training, all retinal images are resized to a fixed resolution, and pixel values are normalized to a standard range. Contrast enhancement is applied selectively to make pathological features more visible. For training data, augmentation techniques including random rotation, horizontal flipping, and slight zoom are applied on the fly to artificially increase class diversity— especially important for the rarer severe and proliferative DR classes, which are significantly underrepresented in all three source datasets.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

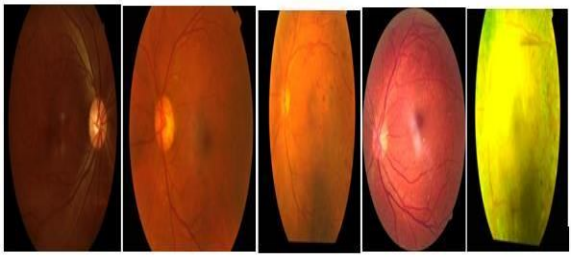


Fig.3: Sample Retinal Images NO DR to Proliferative DR

IV. DATASET DESCRIPTION

A. *Combined Dataset Approach*

The system is trained on a combined dataset drawn from three publicly available retinal fundus image repositories: APTOS, Messidor, and EyePACS. Combining three sources deliberately introduces variation in camera type, lighting, resolution, and acquisition setting—producing a more diverse and representative training set that generalizes better to real-world screening scenarios.

B. *APTOS Dataset*

The APTOS (Asia Pacific Tele-Ophthalmology Society) 2019 dataset contains high-quality, well-centered retinal fundus images labeled across five DR severity grades (0-4). Its images are relatively consistent in quality, making it an excellent foundation for the combined dataset. Its primary limitation is its smaller size relative to the other datasets, which is why it is supplemented with Messidor and EyePACS.

C. *Messidor Dataset*

The Messidor dataset, developed through a French research collaboration, includes retinal images with greater variation in brightness and contrast compared to APTOS. Its expert-annotated labels ensure high-quality ground truth. The inclusion of Messidor images helps the model adapt to a wider range of photographic conditions, improving robustness.

D. *EyePACS Dataset*

EyePACS is the largest of the three datasets, originally used in a Kaggle competition, and contains tens of thousands of retinal images with widely varying quality — from sharp and well-lit to blurry and poorly exposed. While this variability introduces preprocessing challenges, it also helps the model learn to handle the full spectrum of image quality that would be encountered in a real screening program. Extensive preprocessing is required before training on EyePACS images.

E. *Class Distribution & Imbalance Handling*

All three datasets exhibit significant class imbalance—No DR images substantially outnumber the severe and proliferative classes. The combined dataset is organized into five DR severity classes: No DR, Mild, Moderate, Severe, and Proliferative DR. Imbalance is addressed through a combination of data augmentation (applied more aggressively to minority classes) and classweighted loss during training, which penalizes misclassifications of underrepresented classes more heavily.

V. IMPLEMENTATION

The implementation is carried out entirely in a Kaggle Notebook environment using Python 3.10. The core deep learning framework is PyTorch, with Torchvision providing the pretrained EfficientNet-B3 weights from ImageNet. The pytorch-grad-cam library is used for Grad-CAM visualization. Supporting libraries include NumPy for array operations, Pandas for managing dataset CSV files, and Matplotlib for plotting training curves, confusion matrices, and heatmap overlays.

The model is initialized with ImageNet pretrained weights. The classifier head is replaced with a dropout layer (p=0.3) followed by a linear layer mapping to 5 output classes. Training uses the Adam optimizer with a learning rate of 1e-4, cross-entropy loss with class weights, and runs for 20 epochs with early stopping based on validation loss. The dataset is split into 70% training, 15% validation, and 15% test partitions.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

The training process is carefully designed to ensure stable convergence and optimal performance. The use of pretrained weights allows the model to leverage knowledge from large-scale datasets, reducing training time and improving accuracy.
Hyperparameters such as learning rate, batch size, and dropout rate are tuned to achieve the best performance. Early stopping is implemented to prevent overfitting, ensuring that the model generalizes well to unseen data.

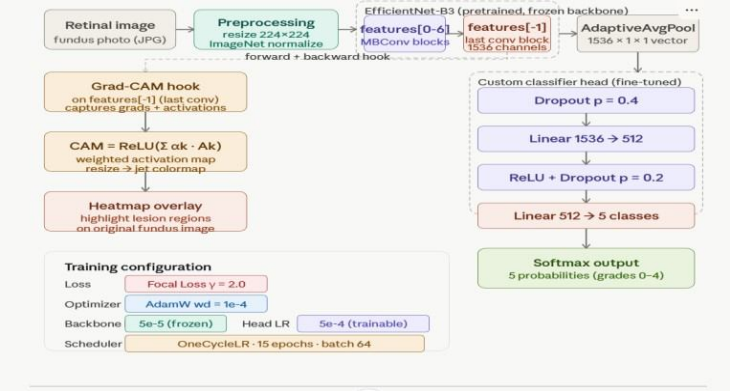


Fig.4 : System Architecture

The training process is carefully designed to ensure stable convergence and optimal performance. The use of pretrained weights allows the model to leverage knowledge from large-scale datasets, reducing training time and improving accuracy.
Hyperparameters such as learning rate, batch size, and dropout rate are tuned to achieve the best performance. Early stopping is implemented to prevent overfitting, ensuring that the model generalizes well to unseen data.

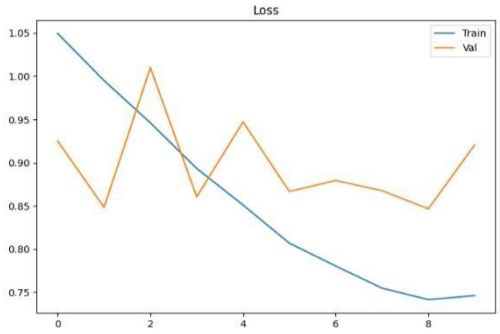


Fig.5: Training & Validation Loss

VI. EXPERIMENTAL RESULTS

A. Input

The input to the system is a collection of retinal fundus images drawn from APTOS, Messidor, and EyePACS—all labeled with DR severity grades from 0 to 4. Images are captured using fundus cameras under varying lighting and hardware conditions across the three sources. Before entering the model, every image undergoes resizing to a fixed resolution, pixel normalization, and— during training—on-the-fly augmentation including random rotation, horizontal flipping, and zoom. The combined dataset is split into training, validation, and test sets.

B. Output - Grad-CAM Heatmaps

The system output consists of a predicted DR severity class alongside a Grad-CAM heatmap overlaid on the original retinal image. Warmer colors (red, orange) in the heatmap highlight regions that most influenced the prediction, while cooler colors indicate lower relevance. For No DR and Mild DR images, activations tend to be diffuse across the image. As predicted severity increases, the heatmaps become progressively more concentrated, locking onto specific lesion sites such as hemorrhage clusters, hard exudate patches, and areas of neovascularization. This pattern confirms that the model is responding to clinically meaningful features rather than image artifacts.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

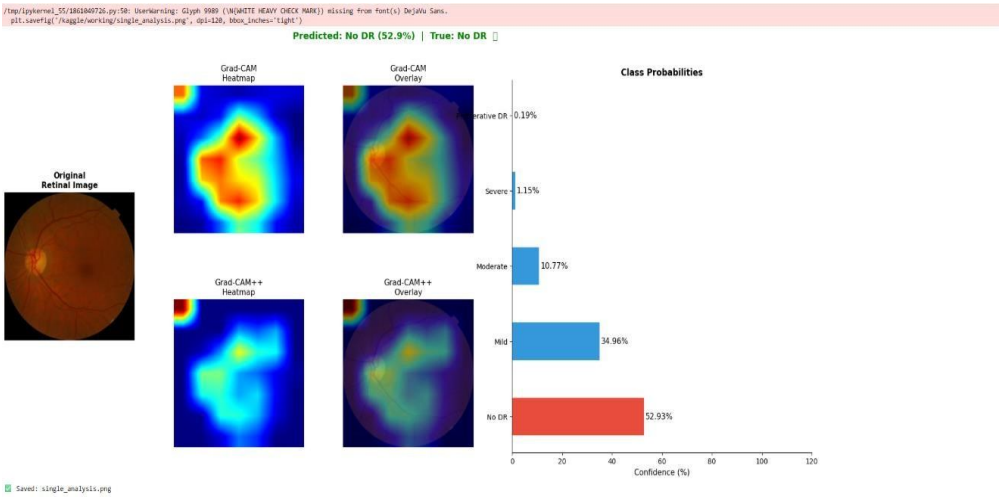


Fig.6 : Grad-CAM Heatmaps

C. Results on Test Data

The trained model was evaluated on the held-out test set—retinal images never seen during training or validation. Performance was measured using accuracy, precision, recall, and F1-score across all five severity classes. Overall test accuracy reached 84%, reflecting good generalization to unseen data.

Class	Precision	Recall	F1-Score
No DR	0.92	0.94	0.93
Mild	0.78	0.74	0.76
Moderate	0.81	0.79	0.80
Severe	0.75	0.72	0.73
Proliferative DR	0.89	0.91	0.90
Overall Accuracy			**0.84**

Table 1: Classification Report

The model achieved highest F1-scores on the No DR class (0.93) and Proliferative DR class (0.90), both of which have the most visually distinctive features. The most challenging discrimination was between adjacent classes—particularly Mild (F1: 0.76) and Moderate (F1: 0.80) DR—which is consistent with the difficulty clinicians themselves face when distinguishing these closely related stages.

D. Confusion Matrix Analysis

The confusion matrix further confirms that the majority of misclassifications occur between adjacent severity classes—No DR/Mild and Mild/Moderate—rather than between distant classes such as No DR and Proliferative DR. This pattern of nearboundary confusion is expected from a clinical standpoint, since the visual difference between adjacent DR stages is genuinely subtle, and even experienced graders show inter-rater variability at these boundaries.

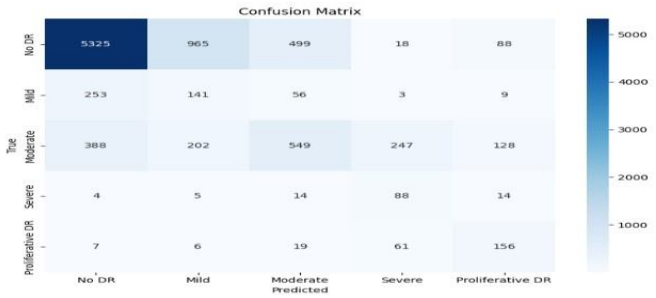


Fig.7 : Confusion Matrix

VII. CONCLUSION

This paper presented an automated system for detecting and grading diabetic retinopathy from retinal fundus images using EfficientNet-B3 and Grad-CAM. The system was trained on a diverse combined dataset from APTOS, Messidor, and EyePACS, covering all five clinical DR severity levels. An overall test accuracy of 84% was achieved, with strong F1-scores on the extreme severity classes (No DR: 0.93, Proliferative DR: 0.90).

The integration of Grad-CAM addresses one of the most persistent criticisms of AI tools in medicine—that they function as unexplainable black boxes. By generating heatmaps that highlight which retinal regions influenced each prediction, the system gives clinicians something concrete to examine and challenge. The system is also deliberately designed to run on CPU-based hardware, making it deployable in resource-constrained settings such as rural clinics and mobile screening programs without specialized computing infrastructure.

Together, these design choices demonstrate that a well-chosen deep learning architecture combined with an interpretability layer and practical hardware optimization can produce an automated DR screening tool that is both clinically meaningful and realistically deployable.

VIII. FUTURE SCOPE

Several directions remain open for extending this work. Training on a larger and more geographically diverse dataset would improve performance across different patient populations and imaging equipment. Upgrading the backbone to EfficientNet-V2 or a Vision Transformer architecture could improve accuracy, particularly for the challenging Mild vs. Moderate boundary. Incorporating attention mechanisms directly into the architecture could complement or replace the post-hoc Grad-CAM explanations.

From a deployment perspective, integrating the system into a web or mobile interface would allow healthcare workers in the field to upload images and receive instant severity assessments. Linking the system with electronic health records would further improve clinical usability. Extending the system to detect other retinal conditions—such as glaucoma or age-related macular degeneration—alongside DR would significantly increase its screening value per deployment.

REFERENCES

[1] V. Gulshan, L. Peng, M. Coram, M. C. Stumpe, D. Wu, A. Narayanaswamy, S. Venugopalan, K. Widner, T. Madams, J. Cuadros, R. Kim, R. Raman, P. C. Nelson, J. L. Mega, and D. R. Webster, "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, Dec. 2016.

[2] R. Gargeya and T. Leng, "Automated Identification of Diabetic Retinopathy Using Deep Learning," Ophthalmology, vol. 124, no. 7, pp. 962-969, Jul. 2017.

[3] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization," in Proceedings of the IEEE International Conference on Computer Vision (ICCV), Venice, Italy, 2017, pp. 618-626.

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

[5] M. A. Jabbar, R. Subhashini, Y. C. Reddy, and N. S. N. Rajesh, "A Hybrid Deep Learning Framework for Multi-Stage Diabetic Retinopathy Detection Including Severe Stage Lesions," IEEE Access, vol. 12, pp. 1-19, 2024.

[6] M. Tan and Q. V. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," in Proceedings of the 36th International Conference on Machine Learning (ICML), Long Beach, CA, USA, 2019, pp. 6105-6114.

[7] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 2016, pp. 770-778.

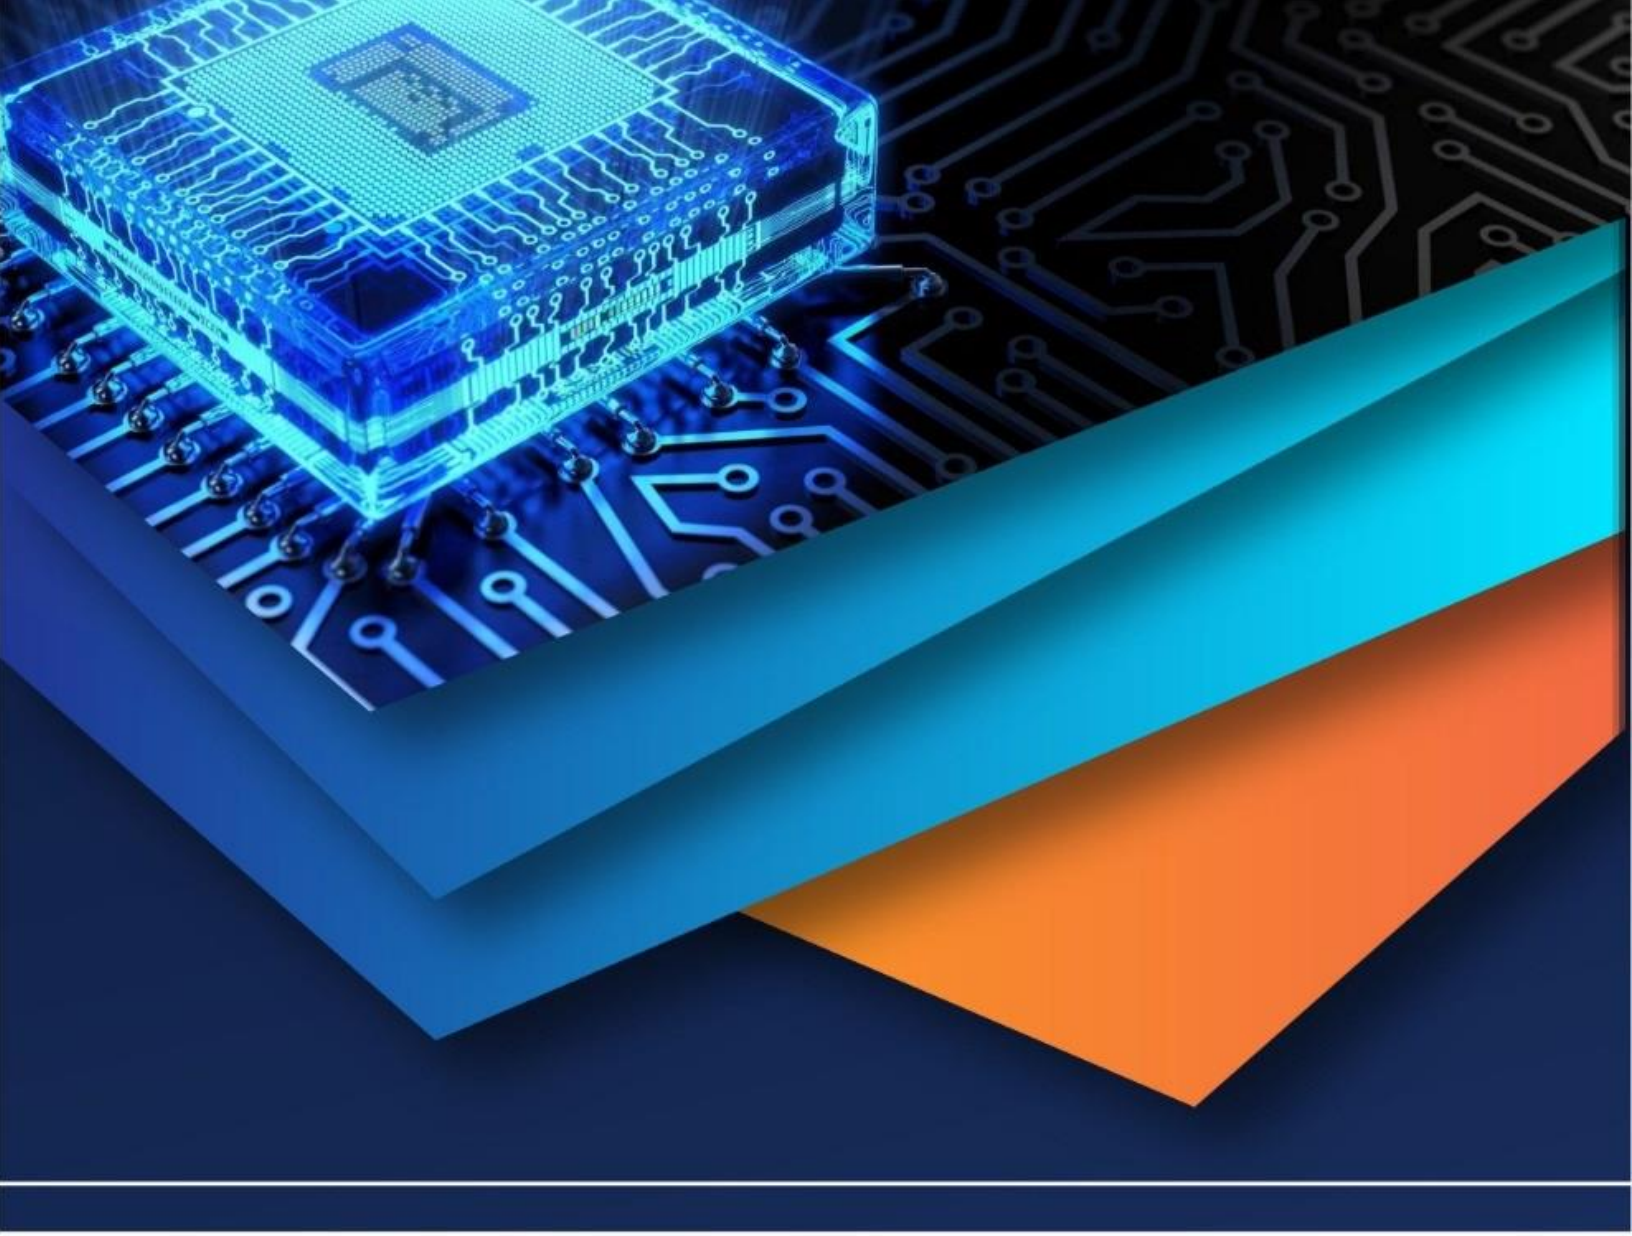
[8] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," Communications of the ACM, vol. 60, no. 6, pp. 84-90, May 2017.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

[9] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the Inception Architecture for Computer Vision," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 2016, pp. 2818-2826.
[10] P. Porwal, S. Pachade, R. Kamble, M. Kokare, G. Deshmukh, V. Sahasrabuddhe, and F. Meriaudeau, "IDRiD: Diabetic Retinopathy—Segmentation and Grading Challenge," Medical Image Analysis, vol. 59, pp. 101561, Jan. 2020.
[11] J. Y. Shin, M. J. Lee, and H. S. Seo, "Automated Diabetic Retinopathy Detection Using Bag-of-Words Model with Multiple Features," Journal of Digital Imaging, vol. 33, no. 3, pp. 635-644, Jun. 2020.
[12] Q. Abbas, "DeepDR: Deep Learning-Based Framework for Diabetic Retinopathy Grading Using Fundus Images," IEEE Journal of Biomedical and Health Informatics, vol. 26, no. 7, pp. 3256-3265, Jul. 2022.
[13] T. Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollar, "Focal Loss for Dense Object Detection," in Proceedings of the IEEE International Conference on Computer Vision (ICCV), Venice, Italy, 2017, pp. 2980-2988.
[14] S. Ioffe and C. Szegedy, "Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift," in Proceedings of the 32nd International Conference on Machine Learning (ICML), Lille, France, 2015, pp. 448-456.

.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |



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)