

Early Detection of Diabetic Retinopathy from Retinal Fundus Images Using Deep Learning Models

Samreen Sulthana¹, Nimmaraju Rajesh², Malipatel Anusha³

^{1,2,3}Assistant Professor, Department of CSE (AI & ML)

Guru Nanak Institute of Technology, Hyderabad, Telangana state, India

Abstract

Diabetic Retinopathy (DR) is a sight-threatening micro-vascular complication of diabetes mellitus and remains the leading cause of preventable blindness among adults of working age. Manual screening of retinal fundus images by ophthalmologists is time consuming, requires expert interpretation, and is not scalable to the rapidly growing diabetic population. In this paper, we propose an automated deep learning framework for the early detection and severity grading of DR using retinal fundus images. Four deep learning models—a baseline convolutional neural network (CNN), VGG-16, ResNet-50, EfficientNet-B3 and a Vision Transformer (ViT-B/16)—are trained and compared on the publicly available APTOS-2019 and EyePACS datasets. Images are pre-processed using Contrast Limited Adaptive Histogram Equalisation (CLAHE) and Gaussian blur removal, followed by extensive data augmentation. Transfer learning with the Adam optimiser and categorical cross-entropy loss is employed to fine-tune the pre-trained backbones over thirty epochs. Experimental results show that the Vision Transformer attains the highest classification accuracy of 95.6 %, F1-score of 95.0 % and AUC of 97.7 %, narrowly outperforming EfficientNet-B3 (95.1 %) and ResNet-50 (94.2 %). The proposed pipeline therefore demonstrates the strong potential of modern deep learning architectures in supporting real-time, low-cost DR screening systems and can be readily deployed in tele-ophthalmology platforms, particularly in resource-constrained regions.

Keywords: Deep Learning, Diabetic Retinopathy, Convolutional Neural Network, Transfer Learning, Vision Transformer, Medical Image Analysis, Computer-Aided Diagnosis.

1. Introduction

The prevalence of diabetes mellitus has reached pandemic proportions, with the International Diabetes Federation reporting more than 537 million affected adults in 2021 and a projected rise to 783 million by 2045 [1]. Diabetic Retinopathy (DR) is one of the most common micro-vascular complications of long-standing diabetes and is responsible for approximately 4.8 % of the 37 million cases of blindness worldwide [2]. Early detection of DR is critical because timely laser photo-coagulation, intravitreal injection or pan-retinal therapy can prevent up to 95 % of severe vision loss [3].

Traditionally, DR screening is performed by ophthalmologists who manually examine high-resolution retinal fundus photographs for the presence of micro-aneurysms, haemorrhages, hard exudates, soft exudates (cotton-wool spots), venous beading and neovascularisation. This manual process is laborious, subjective, depends on expert availability and is inherently non-scalable to the growing diabetic population, particularly in rural and low-income areas [4]. Consequently, the development of accurate, automated and inexpensive computer-aided diagnosis (CAD) systems has become an urgent research priority.

Over the last decade, Deep Learning (DL) – and in particular Convolutional Neural Networks (CNN) – has revolutionised the field of medical image analysis. Models such as AlexNet [5], VGGNet [6], ResNet [7], Inception [8], DenseNet [9] and EfficientNet [10] have achieved expert-level performance on a wide variety of classification tasks, including breast cancer detection, lung nodule classification and skin lesion analysis. More recently, the introduction of the Transformer architecture into computer vision through the Vision Transformer (ViT) [11] and its derivatives has further pushed the state of the art.

In this paper we present a comprehensive comparative study of five widely used deep learning architectures—a custom baseline CNN, VGG-16, ResNet-50, EfficientNet-B3 and ViT-B/16—for the five-class severity grading of DR

(No DR, Mild, Moderate, Severe and Proliferative). The novelty of the proposed work lies in (i) a unified end-to-end pipeline that combines CLAHE based pre-processing with aggressive data augmentation, (ii) the use of transfer learning to mitigate the limited size of medical datasets, and (iii) a rigorous comparative evaluation that includes classical CNNs as well as the latest attention-based ViT. The principal contributions of this paper are:

- An end-to-end deep learning framework for early DR detection from retinal fundus images.
- A systematic comparison of five state-of-the-art deep learning models on two benchmark datasets (APTOS-2019 and EyePACS).
- An ablation study quantifying the impact of CLAHE pre-processing, augmentation and transfer learning.
- Detailed performance analysis using accuracy, precision, recall, F1-score, AUC and confusion matrices.
- A discussion of deployment feasibility for tele-ophthalmology platforms in low-resource settings.

The remainder of the paper is organised as follows. Section 2 reviews the related literature. Section 3 describes the dataset, pre-processing pipeline and the architectures of the five models. Section 4 presents the experimental results, tabular comparisons and graphical analysis. Section 5 concludes the paper and outlines avenues for future research.

2. Literature Review

Automated DR detection has progressed through three broad eras: (i) classical image-processing pipelines, (ii) hand-crafted feature plus shallow classifier approaches and (iii) deep learning based end-to-end methods. This section reviews the most influential contributions, focusing on deep learning.

Gulshan et al. [12] presented one of the earliest large-scale studies, training an Inception-v3 network on 128,175 retinal images and achieving a sensitivity of 97.5 % and specificity of 93.4 % for referable DR detection. This work, published in JAMA, is widely regarded as a milestone that demonstrated the clinical feasibility of deep learning for ophthalmology.

Pratt et al. [13] proposed a customised CNN trained on the Kaggle EyePACS dataset and reported an accuracy of 75 % across five severity classes, highlighting the inherent difficulty of fine-grained DR grading. Wang et al. [14] introduced a Zoom-in-Net that fuses global and local features through three sub-networks and obtained an accuracy of 91.1 %.

ResNet-based architectures, originally proposed by He et al. [7], have proved particularly successful for DR. Li et al. [15] fine-tuned a ResNet-50 with focal loss to mitigate class imbalance and reported a quadratic weighted kappa of 0.84 on the APTOS-2019 dataset. Bodapati et al. [16] proposed a composite deep convolutional network with gated attention and achieved 82.5 % accuracy on the same dataset.

EfficientNet, introduced by Tan and Le [10], unifies depth, width and resolution scaling in a principled manner. Liu et al. [17] applied EfficientNet-B5 to DR grading and reported an AUC of 0.974, outperforming both ResNet and DenseNet baselines while requiring fewer parameters.

Attention mechanisms and Transformers have most recently entered the field. Dosovitskiy et al. [11] showed that pure Transformer models, when pre-trained on very large datasets, surpass CNNs on ImageNet classification. Sun et al. [18] successfully transferred this idea to fundus image analysis with a Vision Transformer pre-trained on JFT-300M and reported an accuracy of 95.4 % on EyePACS. Wu et al. [19] proposed Swin-Transformer-based pipelines that capture both local and global context for joint DR and diabetic macular edema grading.

Despite this rapid progress, three principal challenges remain. First, class imbalance in publicly available DR datasets degrades the performance on severe classes. Second, fundus images vary considerably in illumination, field of view and resolution, requiring robust pre-processing. Third, model interpretability is essential for clinical adoption; recent works such as those by Selvaraju et al. [20] (Grad-CAM) and Chattopadhyay et al. [21] (Grad-CAM++) provide visual explanations and are increasingly being integrated into DR pipelines. Surveys by Asiri et al. [22], Litjens et al. [23] and Tsiknakis et al. [24] provide a broader overview of the state of the art.

3. Methodology

This section describes the proposed end-to-end pipeline. Figure 1 provides a high-level overview, while Figure 2 details the architecture of the best performing CNN backbone.

Figure 1. Proposed Methodology Pipeline




Figure 1. Block diagram of the proposed methodology for diabetic retinopathy detection.

3.1 Dataset

Two publicly available benchmark datasets are used. (i) The APTOS-2019 Blindness Detection dataset [25] contains 3,662 high-resolution retinal fundus images labelled into five severity classes by expert clinicians of Aravind Eye Hospital, India. (ii) The Kaggle EyePACS dataset [26] contains 88,702 fundus images. The class distribution is highly imbalanced, with approximately 49 % images belonging to the No-DR class. Class weighting and oversampling are therefore applied during training to mitigate this skew.

3.2 Image Pre-processing

Each retinal image is centre-cropped to remove black borders and resized to 224×224 pixels to match the input dimension of the pre-trained models. Contrast Limited Adaptive Histogram Equalisation (CLAHE) is applied on the green channel, which contains the maximum vascular information, to enhance the visibility of small lesions such as micro-aneurysms. A Gaussian blur subtraction is then used to normalise illumination differences across images captured by different cameras. Finally, pixel intensities are normalised to the range $[0, 1]$ and standardised using ImageNet mean and standard deviation values.

3.3 Data Augmentation

To improve generalisation and reduce overfitting on the relatively small training set, the following augmentations are applied on the fly during training: random horizontal and vertical flipping, random rotation in $[-30^\circ, +30^\circ]$, random zoom ($\pm 15\%$), random brightness and contrast jitter ($\pm 20\%$), and random shifting up to $\pm 10\%$ of the image dimensions. Each augmentation is applied with probability 0.5.

3.4 Deep Learning Architectures

Five deep learning models are evaluated. (a) A baseline CNN with four convolutional blocks followed by two fully connected layers; (b) VGG-16 [6]; (c) ResNet-50 [7] which introduces residual connections to mitigate vanishing

gradients; (d) EfficientNet-B3 [10] which uses compound scaling and depth-wise separable convolutions; and (e) Vision Transformer ViT-B/16 [11] which divides the input image into 16×16 patches and uses self-attention to model long-range dependencies. For models (b)–(e), the backbone is initialised with ImageNet pre-trained weights, the final classification layer is replaced with a global average pooling layer followed by a 512-unit fully connected layer, dropout ($p = 0.4$) and a 5-way soft-max output.

Figure 2. Deep Learning Architecture (ResNet-50 backbone with custom classifier head)



Figure 2. Architecture of the ResNet-50 backbone with the custom classification head used in this study.

3.5 Training Configuration

All models are trained for 30 epochs with a mini-batch size of 32 using the Adam optimiser (initial learning rate $1e-4$, $\beta_1 = 0.9$, $\beta_2 = 0.999$). A ReduceLROnPlateau scheduler reduces the learning rate by a factor of 0.5 if the validation loss does not decrease for three consecutive epochs. Categorical cross-entropy is used as the loss function with class weights inversely proportional to class frequency. Training is performed on an NVIDIA RTX-3090 GPU with 24 GB of memory.

3.6 Evaluation Metrics

Five standard metrics are reported: Accuracy, macro-averaged Precision, Recall, F1-score, and the Area Under the Receiver Operating Characteristic Curve (AUC). Additionally, the Quadratic Weighted Kappa (QWK) is reported to penalise severe misclassifications more heavily, since DR labels are ordinal. Confusion matrices are used to visualise per-class behaviour.

4. Results and Discussion

4.1 Quantitative comparison

Table 1 summarises the test-set performance of the five models. The Vision Transformer obtains the highest accuracy (95.6 %), F1-score (95.0 %) and AUC (97.7 %), narrowly outperforming EfficientNet-B3 and ResNet-50. The custom baseline CNN, despite being lightweight, attains a respectable 82.4 % accuracy, demonstrating that even modest networks can learn discriminative DR features when combined with appropriate pre-processing.

Table 1. Test-set performance of the five deep learning models (%).

Model	Accuracy	Precision	Recall	F1-score	AUC	QWK
Baseline CNN	82.4	80.7	79.9	80.1	85.6	0.74
VGG-16	88.7	87.6	87.0	87.3	91.2	0.81
ResNet-50	94.2	93.8	93.3	93.5	96.8	0.88
EfficientNet-B3	95.1	94.7	94.5	94.6	97.4	0.90
ViT-B/16	95.6	95.2	94.9	95.0	97.7	0.91




Figure 5. Bar-chart comparison of Accuracy, F1-score and AUC across the five deep learning models.

4.2 Convergence behaviour

Figure 3 shows the training and validation curves of the best-performing CNN backbone (ResNet-50). Both training and validation accuracy rise steeply during the first ten epochs and plateau around epoch 25. The small gap between the training and validation losses (≈ 0.08) indicates that the model generalises well and does not significantly overfit, owing to the combination of CLAHE pre-processing, augmentation and dropout regularisation.




Figure 3. Accuracy (a) and loss (b) curves for training and validation across 30 epochs.

4.3 Confusion matrix and per-class analysis

Figure 4 presents the confusion matrix of the ResNet-50 model on the test set. The classifier is highly accurate for the No-DR (510/537 correct) and Moderate (290/324 correct) classes. The largest source of confusion occurs between the Mild and Moderate classes, which is consistent with the clinically subtle differences between these grades. Severe and Proliferative cases are correctly identified in 132/153 and 93/104 instances respectively, which is encouraging from a safety-critical perspective.




Figure 4. Confusion matrix of the ResNet-50 model on the held-out test set.

Table 2. Per-class precision, recall and F1-score of the ResNet-50 model.

Class	Support	Precision	Recall	F1-score
No DR	537	0.962	0.950	0.956
Mild	212	0.840	0.840	0.840
Moderate	324	0.893	0.895	0.894
Severe	153	0.832	0.863	0.847
Proliferative	104	0.894	0.894	0.894

4.4 Ablation study

Table 3 reports an ablation study that quantifies the individual contribution of CLAHE pre-processing, data augmentation and transfer learning. Removing transfer learning causes the largest drop (-9.8 %), confirming that ImageNet-pretrained features generalise well to fundus images. CLAHE accounts for a further +2.3 % improvement, while augmentation contributes +1.6 %.

Table 3. Ablation study on the ResNet-50 model (test accuracy %).

Configuration	Accuracy
w/o Transfer Learning (random init)	84.4
w/o CLAHE pre-processing	91.9
w/o Data Augmentation	92.6
Full pipeline (proposed)	94.2

4.5 Discussion

The results demonstrate that modern deep learning architectures can grade DR with an accuracy that approaches expert-level performance. The Vision Transformer marginally outperforms convolutional backbones, which we attribute to its global self-attention mechanism that captures long-range vascular patterns characteristic of advanced DR. However, ViTs require larger datasets and pre-training to reach this performance, which may not always be available. EfficientNet-B3 offers the best accuracy-to-parameter ratio (12 M parameters vs 86 M for ViT-B/16) and is therefore recommended for mobile deployment. The confusion-matrix analysis shows that the most clinically dangerous misclassifications—severe DR predicted as mild—occur in fewer than 2 % of cases, which is acceptable for a screening tool that triages patients to specialists. These findings align with the recent surveys [22-24] and confirm that DL based DR screening is now sufficiently mature for pilot deployment.

5. Conclusion and Future Scope

This paper presented an end-to-end deep learning framework for early detection and severity grading of Diabetic Retinopathy from retinal fundus images. Five models—a custom CNN, VGG-16, ResNet-50, EfficientNet-B3 and ViT-B/16—were trained and compared on the APTOS-2019 and EyePACS datasets. The Vision Transformer achieved the highest accuracy of 95.6 % and AUC of 97.7 %, while EfficientNet-B3 offered the best accuracy-to-parameter trade-off. CLAHE pre-processing, aggressive augmentation and transfer learning were shown to be essential components of the pipeline through a quantitative ablation study.

Future research directions include:

- Extending the framework to multi-disease classification (DR, glaucoma, age-related macular degeneration).
- Incorporating Grad-CAM and other explainable-AI methods to provide visual interpretability of model predictions for clinicians.
- Federated learning across hospitals to leverage larger and more diverse datasets while preserving patient privacy.
- Self-supervised pre-training on millions of unlabelled fundus images to further reduce dependence on labelled data.
- Deployment of a lightweight EfficientNet-B0 / MobileViT variant on smartphones for low-cost field screening in remote regions.
- Integration of multimodal data (Optical Coherence Tomography, patient metadata) to improve diagnostic robustness.
- Real-time tele-ophthalmology platforms connected to nationwide diabetic registries.

References

- [1] International Diabetes Federation, IDF Diabetes Atlas, 10th Edition, Brussels, 2021.
- [2] World Health Organization, "World Report on Vision," Geneva, 2019.
- [3] American Academy of Ophthalmology, "Diabetic Retinopathy Preferred Practice Pattern," Ophthalmology, vol. 127, no. 1, pp. 66–145, 2020.
- [4] T. Y. Wong, J. Sun, R. Kawasaki, et al., "Guidelines on Diabetic Eye Care: The International Council of Ophthalmology Recommendations for Screening, Follow-up, Referral, and Treatment Based on Resource Settings," Ophthalmology, vol. 125, no. 10, pp. 1608–1622, 2018.
- [5] A. Krizhevsky, I. Sutskever and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," NIPS, 2012.
- [6] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," ICLR, 2015.
- [7] K. He, X. Zhang, S. Ren and J. Sun, "Deep Residual Learning for Image Recognition," CVPR, pp. 770–778, 2016.

- [8] C. Szegedy, V. Vanhoucke, S. Ioffe, et al., "Rethinking the Inception Architecture for Computer Vision," CVPR, pp. 2818–2826, 2016.
- [9] G. Huang, Z. Liu, L. van der Maaten and K. Q. Weinberger, "Densely Connected Convolutional Networks," CVPR, pp. 4700–4708, 2017.
- [10] M. Tan and Q. V. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," ICML, 2019.
- [11] A. Dosovitskiy, L. Beyer, A. Kolesnikov, et al., "An Image is Worth 16×16 Words: Transformers for Image Recognition at Scale," ICLR, 2021.
- [12] V. Gulshan, L. Peng, M. Coram, 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.
- [13] 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.
- [14] Z. Wang, Y. Yin, J. Shi, et al., "Zoom-in-Net: Deep Mining Lesions for Diabetic Retinopathy Detection," MICCAI, pp. 267–275, 2017.
- [15] X. Li, X. Hu, L. Yu, et al., "Canet: Cross-disease Attention Network for Joint Diabetic Retinopathy and Diabetic Macular Edema Grading," IEEE TMI, vol. 39, no. 5, pp. 1483–1493, 2020.
- [16] J. D. Bodapati, V. Naralasetti, S. N. Shareef, et al., "Blended Multi-Modal Deep ConvNet Features for Diabetic Retinopathy Severity Prediction," Electronics, vol. 9, no. 6, p. 914, 2020.
- [17] Y. Liu, H. Yip, Y. Zheng, et al., "Efficient Diabetic Retinopathy Grading with Deep Learning," Scientific Reports, vol. 12, p. 18897, 2022.
- [18] R. Sun, Y. Li, T. Zhang, et al., "Lesion-aware Transformer for Diabetic Retinopathy Grading," IEEE Transactions on Medical Imaging, vol. 42, no. 7, pp. 2017–2029, 2023.
- [19] J. Wu, R. Hu, Z. Xiao, et al., "Vision Transformer-based Recognition of Diabetic Retinopathy Grade," Medical Physics, vol. 48, no. 12, pp. 7850–7863, 2021.
- [20] R. R. Selvaraju, M. Cogswell, A. Das, et al., "Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization," ICCV, pp. 618–626, 2017.
- [21] A. Chattopadhyay, A. Sarkar, P. Howlader and V. N. Balasubramanian, "Grad-CAM++: Improved Visual Explanations for Deep Convolutional Networks," WACV, pp. 839–847, 2018.
- [22] N. Asiri, M. Hussain, F. Al Adel and N. Alzaidi, "Deep Learning Based Computer-Aided Diagnosis Systems for Diabetic Retinopathy: A Survey," Artificial Intelligence in Medicine, vol. 99, p. 101701, 2019.
- [23] G. Litjens, T. Kooi, B. E. Bejnordi, et al., "A Survey on Deep Learning in Medical Image Analysis," Medical Image Analysis, vol. 42, pp. 60–88, 2017.
- [24] N. Tsiknakis, D. Theodoropoulos, G. Manikis, et al., "Deep Learning for Diabetic Retinopathy Detection and Classification Based on Fundus Images: A Review," Computers in Biology and Medicine, vol. 135, p. 104599, 2021.
- [25] APTOS 2019 Blindness Detection, Kaggle Competition Dataset, Asia Pacific Tele-Ophthalmology Society, 2019. Available: <https://www.kaggle.com/c/aptos2019-blindness-detection>.
- [26] EyePACS, Diabetic Retinopathy Detection Dataset, Kaggle Competition, California Healthcare Foundation, 2015. Available: <https://www.kaggle.com/c/diabetic-retinopathy-detection>.
- [27] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," ICLR, 2015.
- [28] I. Goodfellow, Y. Bengio and A. Courville, Deep Learning, MIT Press, 2016.
- [29] S. Ioffe and C. Szegedy, "Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift," ICML, pp. 448–456, 2015.
- [30] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever and R. Salakhutdinov, "Dropout: A Simple Way to Prevent Neural Networks from Overfitting," JMLR, vol. 15, no. 1, pp. 1929–1958, 2014.