Lesion-Aware Binary Diabetic Retinopathy Detection in an Imbalanced Fundus Image Dataset

Hafsah Mahmood, Nadeem Anjum
Capital University of Science and Technology, Islamabad.
*Correspondence: Hafsah.mahmood@cust.edu.pk, Nadeem.anjum@cust.edu.pk

Citation | Mahmood. H, Anjum. N, "Lesion-Aware Binary Diabetic Retinopathy Detection in Imbalanced Fundus Image Dataset", IJIST, Special Issue pp 674-693, May 2026
Received | April 08, 2026 **Revised |** May 14, 2026 **Accepted |** May 17, 2026 **Published |** May 19, 2026.

Diabetic retinopathy (DR) is one of the leading causes of vision impairment in the world. Early diagnosis is needed to prevent permanent vision loss. Despite being shown to perform well on benchmark DR datasets, class imbalance and loss of smaller lesion features in preprocessing often lead to performance degradation of convolutional neural networks (CNNs) on real-world, imbalanced, multi-disease datasets. Precise lesion identification is particularly important is particularly important in multi-disease datasets, in which DR-related lesions can be small and easily overlooked, since it has a direct impact on sensitivity and clinical reliability. This work uses the imbalanced multi-disease Retinal Fundus Multi-disease Image Dataset (RFMiD) to evaluate a baseline CNN pipeline, originally designed on the balanced Messidor dataset. We propose two targeted augmentation techniques: rotation and photometric transformations (contrast adjustment, sharpening, color shifting), along with an improved preprocessing pipeline that incorporates Difference of Gaussian (DoG) and Dilated DoG (DDoG) to address decreased sensitivity and improve lesion visibility. The proposed lesion-aware preprocessing approach was tested using six different pretrained CNN architectures, such as AlexNet, ResNet-18, VGGNet-S, VGGNet-16, VGGNet-19, and GoogleNet. The application of rotation augmentation along with the proposed preprocessing pipeline resulted in considerable improvements in the performance metrics for each network architecture, where sensitivity increased to 95.16% (VGGNet-19), showing an increase of 13.7%. At the same time, the increase in accuracy (91.72%) and AUC (0.9605) values also increased. Moreover, the results obtained through photometric augmentation and proposed preprocessing showed significant improvement in the robustness of the models, with maximum sensitivity recorded as 91.94% (VGGNet-16), increased by 8.9%, and AUC scores increasing to 0.9646 (VGGNet-19). The PSNR value of 31.26dB and SSIM of 0.88 further confirm the efficacy of the proposed preprocessing pipeline. In summary, lesion-aware image preprocessing alongside rotations and photometric data augmentation methods has significantly increased the accuracy of various pretrained CNN models in diagnosing diabetic retinopathy.

Keywords: Diabetic Retinopathy; Retinal Fundus Images; Convolutional Neural Networks; Data Augmentation; Adaptive Preprocessing.

Introduction:

Diabetic Retinopathy (DR) is a serious chronic condition of type 2 diabetic patients, which affects the blood vessels of the retina due to prolonged elevated blood sugar levels. It can result in irreparable vision loss if left untreated [1]. To prevent disease progression and maintain the vision of millions of people, early and precise detection is essential. DR weakens the blood vessels of the retina and creating tiny bulges in capillaries called microaneurysms. In some cases, the blood vessels start leaking blood, causing hemorrhages to appear on the retina. In severe cases, the proteins get deposited around the blood vessels, resulting in the formation of exudates, which are yellow spots appearing on the surface of the retina [2]. These are collectively called the lesions of DR and are the key indicators of the disease. Figures 1, 2, 3, and 4 visually show the DR lesions on the fundus images.

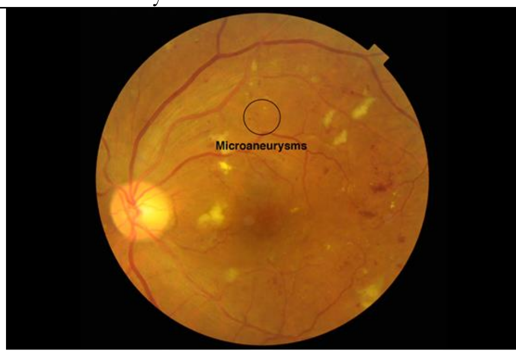

Figure 1. Microaneurysms in DR [3]

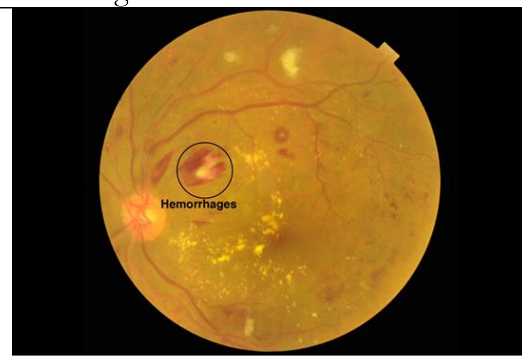

Figure 2. Hemorrhages in DR [3]

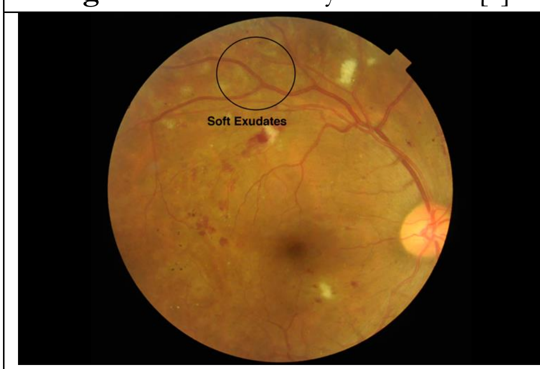

Figure 3. Soft Exudates in DR [3]

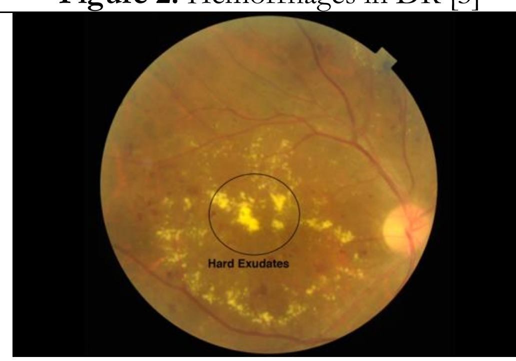

Figure 4. Hard Exudates in DR [3]

DR is usually detected by capturing the fundus images of the retina, which contain the details, including the macula, optic disc, and blood vessels. Fundus images are captured through various camera devices, which introduce numerous camera artifacts, lighting conditions, and noise that are associated with these images [4]. At the initial stages of the disease, the DR lesions are small and sometimes overlooked due to the inherent issues of fundus photography, which make detection challenging. To address this issue, there are various preprocessing techniques that enhance the contrast of the images so that these small lesions can become visible. These techniques include histogram equalization, its variants, Gaussian filtering, adaptive gamma correction, and many others to enhance the overall image quality [5]. Although global contrast is enhanced with normal preprocessing pipelines, they can also inadvertently reduce these small lesions, potentially leading to poorer sensitivity and wrong classification of DR cases. After preprocessing, the images are usually fed to a model for automated detection.

Deep learning models, including Convolutional Neural Networks, have emerged because of their capacity to learn hierarchical visual features and achieve good classification performance on disease-specific datasets for DR detection, including Messidor, APTOS2019, DDR, and EyePACS, which contain healthy and DR images. But their performance drops when detecting DR in the presence of other similar retinal pathologies present in multi-disease datasets. The examples of such datasets are MURED and RFMiD.

The reason for the performance degradation is the overlapping features of DR with other similar ocular conditions, including Branch Retinal Vein Occlusion (BRVO) and Central Retinal Vein Occlusion (CRVO), which have lesions that are quite similar to those of DR. Therefore, the model's sensitivity, which reflects recall for the disease class, drops, leading to an increased number of false negatives and missed cases of DR. Although in such scenarios, the model may present high accuracy, this increase is due to the increase in specificity, which measures the proportion of true negatives (meaning healthy or non-DR cases). Several studies investigated the reason behind the sensitivity drop in such datasets. Apart from the problem of overlapping features, which is discussed above, the other issue is the under-representation of the DR class in these multi-disease datasets. To address this issue, existing literature has presented several augmentation techniques to handle class imbalance and to make the generalization of the model better. There are improvements that are shown by these approaches, but it is noted that the sensitivity is not more that 87% when detecting DR using CNNs alone. Hence, there is a need to increase the sensitivity and reduce false negatives for accurate detection of DR in the presence of other retinal diseases.

Therefore, this work suggests combining targeted data augmentation with adaptive lesion-aware preprocessing that is specific to RFMiD to overcome these constraints. Only DR-class images are subjected to two augmentation techniques: photometric transformations (contrast adjustment, sharpening, and color shifting) to replicate real-world variations in illumination and color distribution, and rotation-based augmentation to increase intra-class diversity without changing retinal anatomy. The preprocessing pipeline includes the addition of difference of Gaussian (DoG) and Dilated DoG (DDoG) filters to accentuate and keep small lesion information without compromising the general context of the image. We use six pretrained CNN architectures (AlexNet, ResNet-18, VGGNet-s, VGGNet-16, VGGNet19, GoogleNet) on RFMiD to measure the impact of the proposed augmentation and preprocessing methods. The results of the experiments indicate a significant improvement in sensitivity, which underscores the importance of dataset-specific augmentation and lesion-preserving preprocessing in ensuring sensitive DR identification in clinically challenging, multi-disease datasets.

Objectives of the Study:

The main objectives of this study are listed below.

To create and build a lesion-aware preprocessing framework capable of enhancing the features of DR retinal lesions (e.g., microaneurysms, haemorrhages, and exudates) to improve diagnostic accuracy.

To examine how data augmentation methods, such as rotation-based and photometric transformations, can be used to meet the goal of mitigating the effects of class imbalance and enhancing model generalization in multi-disease scenarios.

To determine the effectiveness of the proposed methodology with the help of various pretrained CNN architectures, emphasizing the sensitivity and overall classification when applied to complex and multi-disease settings.

Novelty of the Study:

The study makes the following contributions.

The paper introduces a novel lesion-aware preprocessing pipeline with the specific purpose of improving the appearance of important retinal lesions in fundus images, as it is a major drawback of the traditional preprocessing tools, which tend to obscure minor pathological details. The proposed method is tested in a more difficult multi-disease scenario, unlike the existing strategies, which mainly deal with single-disease datasets, with overlapping lesion features and a class imbalance that greatly influences performance.

Moreover, the combination of focused data augmentation techniques and lesion-aware preprocessing offers a single pipeline that enhances sensitivity and balances performance across various pretrained CNN models. This lesion-specific enhancement and augmentation-based generalization is a unique addition to enhancing the accuracy of automated diabetic retinopathy detection in clinically realistic conditions.

The rest of this document is structured as follows: Pre-processing and augmentation techniques are highlighted in Section 2's review of related research on deep learning-based DR detection. The dataset, baseline pipeline reproduction, suggested augmentation methods, and the improved preprocessing procedure are all covered in Section 3. The experimental findings and their implications for DR detection on multi-disease datasets are presented in Section 4. The work is finally concluded in Section 5, which also suggests future research areas.

Literature Review:

CNN-based DR Detection:

The latest developments of CNN-based diabetic retinopathy detection have shown promising results, but there are still a number of limitations. As an example, [1]. suggested an adaptive deep CNN model that enhances detection performance by optimizing its architecture, but the methodology assumes only enhancement of performance, without considering lesion-level preservation of features, which is vital to early detection. In the same way, [2] proposed a multi-label model to address joint prevalent retinal diseases; such models, however, tend to fail on the issue of class imbalance and overlapping lesion features in multi-disease data.

Attention-based models, like HIRD-Net by [4], are more interpretable and able to localize lesions, although they are dependent on superior preprocessing (e.g., CLAHE-DOG) that can suppress small lesions or cause artifacts. In comparison, [6] used background removal and augmentation methods to enhance the performance, but did not use these methods as an integrated pipeline, but as separate preprocessing and augmentation steps.

Survey-based research [7] also indicates that even with high-reported accuracies, CNN-based systems show poor generalization between datasets and unreliable sensitivity, especially in clinically heterogeneous conditions. Also, previous architectures like VGG-NIN [8] and updated ResNet-based models [9] have high capabilities in feature extraction but are mostly tested within the framework of single-disease datasets, which restricts their usability in practice in a multi-disease setting.

These findings suggest a significant research gap: current approaches are either based on architectural enhancements, or preprocessing, or augmentation options and strategies, but the importance of lesion-aware preprocessing together with augmentation has not been given due attention, especially in the context of multi-disease settings where the overlap of lesions and class imbalance are critical issues.

Preprocessing Techniques in DR Detection:

Preprocessing is important in improving the quality of fundus images, and the detection of diabetic retinopathy using CNNs, but current methods have several limitations. An example is that [10] suggested a hybrid retinal image enhancement strategy that enhances contrast and visibility of retinal structures, but global enhancement strategies can inadvertently boost noise or reduce subtle lesion patterns. On the same note, [5] conducted a review of different image enhancement approaches and emphasized the fact that most image enhancement methods are aimed at improving the quality of the overall image but not preserving the lesions, which is not effective in clinical diagnosis.

Recent studies by [11] provided a contrast enhancement method to classify DR, which has shown enhanced performance in detection; still, the algorithm does not directly consider lesion variation across disease stages. Similarly, [12] conducted quantitative research on the methods of enhancement, revealing that preprocessing can enhance the CNN performance, but it is not always able to adapt to various datasets and imaging situations.

According to these studies, the traditional methods of preprocessing are mostly global and non-selective, which may result in the loss or distortion of the important pathological features. This, in turn, necessitates lesion-conscious preprocessing methods that would be able to selectively boost regions of interest to the diagnostician and retain finer lesion features.

Data Augmentation Techniques:

Data augmentation has gained massive application to deal with data scarcity and class imbalance in medical image analysis, but its performance in diabetic retinopathy detection is limited by various factors. Conventional augmentation methods, like rotation, flipping, and scaling, as explained by [13], can enhance the diversity of the dataset, yet not necessarily provide lesion-specific representation. On the same note, generative algorithms that use GANs, as developed, can be used to generate synthetic data, but they can produce unrealistic artifacts, and they need a lot of training materials.

Recent studies, including [14], highlight the fact that although augmentation enhances generalization, its effect on clinically significant aspects, especially small and barely noticeable lesions, is not uniform in all cases. Moreover, most current methods use augmentation without preprocessing, and they do not consider how image transformations interplay with feature enhancement methods.

This inability to integrate restricts the overall performance of augmentation strategies, particularly in multi-disease datasets where the imbalance in classes and overlapping patterns of lesions are more evident. Thus, an integrated strategy involving the combination of augmentation and lesion-aware preprocessing is necessary to attain a higher sensitivity and strong performance of the model.

Research Gap:

Although there are major strides in CNN-based diabetic retinopathy detection, the current research has several very critical limitations. Most methods aim at improving architecture or global preprocessing methods that increase the quality of the final image, but do not allow for preserving fine lesion-specific details to achieve early and accurate diagnosis. Also, most of the techniques are tested using single-disease datasets, which are not representative of clinical practice, where patients may have several retinal pathologies. The overlapping of lesions and class imbalance in such multi-disease environments have tremendous negative effects on model sensitivity and reliability.

Also, data augmentation methods are frequently used in isolation, without considering their interaction with preprocessing mechanisms, leading to an inefficient augmentation of diagnostically relevant features. Therefore, the current frameworks do not have an integrated approach that would effectively tackle lesion preservation, data imbalance, and the complexity of multi-diseases.

To overcome these shortcomings, the current paper presents a lesion-sensitive preprocessing model, which selectively amplifies the important retinal lesions without altering their structural features. The suggested method is also supplemented with specific data augmentation methods, such as rotation-based and photometric transformations, to enhance the diversity of the data and reduce the class imbalance.

In contrast to the current approaches, the proposed framework is tested in a multi-disease environment with the use of several pretrained CNN architectures, with particular attention paid to the increase in sensitivity and the general diagnostic performance. This combined method offers a more powerful and clinically applicable solution to automated detection of diabetic retinopathy.

Methodology:

This section describes the dataset, experimental setup, and methodological framework used to detect diabetic retinopathy on the RFMiD dataset in an automated manner. Figure 5 represents the complete methodology of the study.

The suggested framework applies to the initially class-imbalanced RFMiD fundus image dataset. To compensate for it, two augmentation techniques are used: rotation $(-25° \ to \ +25°)$ and photometric changes (color, contrast, and sharpening) to enhance the data variety and strength.

The augmented dataset is further subjected to a lesion-aware pipeline, which improves features of the retina with color and intensity analysis, adaptive gamma correction, histogram-based transformations, and quantile partitioning. Further refinement algorithms, such as normalization, histogram equalization, retinal masking, and Difference of Gaussian filtering, are used to emphasize lesion structures, and the feature fusion and HSV-based reconstruction are done.

Lastly, the processed data are used to train six CNN models (AlexNet, ResNet-152, VGGNet-S, VGGNet-16, VGGNet-19, and GoogleNet), and the performance of the models is assessed through accuracy, sensitivity, specificity, and AUC.

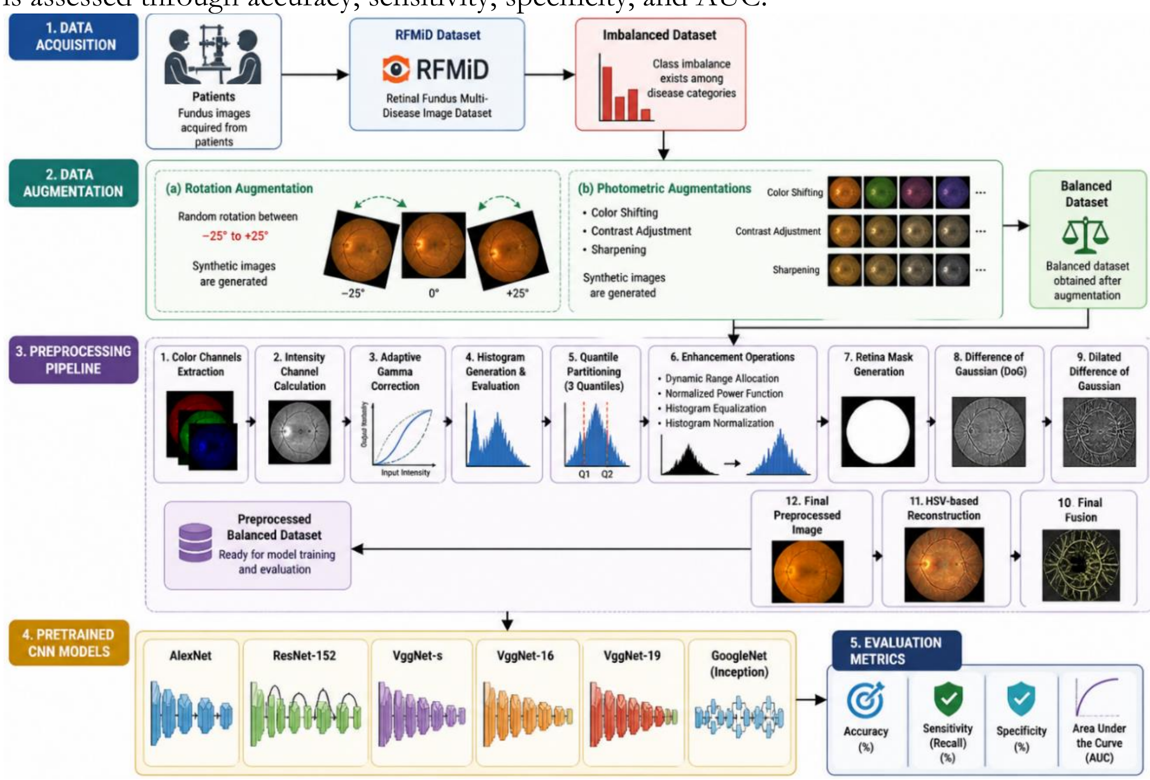

Figure 1. Proposed Methodology

Dataset Description:

In this work, diabetic retinopathy (DR) identification under realistic and clinically challenging situations was assessed using the Retinal Fundus Multi-disease Image Dataset (RFMiD). The 3,200 color retinal fundus images in RFMiD are separated into 1,920 training, 640 validation, and 640 testing images. The collection is representative of actual ophthalmic screening scenarios because each image is labelled for the presence of various retinal disorders, including diabetic retinopathy.

The training dataset has only contains only 376 DR images among images representing 46 other retinal pathologies, which is 19.5% of the entire training set. The extreme class imbalance of RFMiD, in which DR-positive images make up a very small portion of the total dataset, is one of its distinguishing features. In actual clinical practice, DR cases are less common than non-DR or other retinal disorders, which is reflected in this imbalance.

Moreover, RFMiD is highly heterogeneous with changes in field-of-view, illumination, contrast, image quality, and demographics of the patient. These heterogeneities present significant challenges to automated DR detection techniques due to the difficulty of identifying small lesions associated with DR that may be hard to detect or may be masked by the presence of other retinal conditions. RFMiD was specifically selected due to the limitations of existing CNN-based DR detection pipelines that can perform well on single-disease and balanced datasets, but fail in multi-disease and unbalanced settings. DR lesions are often small compared to other diseases with significant pathological patterns in multi-disease datasets, leading to low sensitivity and increased false negative diagnoses. RFMiD can be an ideal research testbed to study lesion-preserving preprocessing and class-specific augmentation methods with the aim of boosting sensitivity without reducing the overall performance. The study focuses more on clinical realism than benchmark optimization since it only considers RFMiD, and the evaluation of how targeted augmentation and adaptive preprocessing can enhance DR detection under real-world screening is more insightful.

Baseline Paper Reproduction:

The baseline study's technique [15] was first replicated on the RFMiD dataset to create a trustworthy performance reference. Six pretrained convolutional neural network architectures were used in the baseline method for binary DR classification:

AlexNet
ResNet-18
VggNet-S
VggNet-16
VggNet-19
GoogleNet

All models were fine-tuned on RFMiD, using the same training and evaluation protocols as in the baseline study, with ImageNet pretrained weights.

Most architectures exhibited a gradual decline in sensitivity, although the overall accuracy and AUC of the reproduced baseline were reasonable. The primary reason behind this reduction in sensitivity is the imbalanced nature of RFMiD, in which the underrepresentation of DR-positive samples skews the learning process to non-DR predictions. Due to this reason, models tend to give preference to the majority group to maximize accuracy, a phenomenon that leads to the loss of DR cases, an outcome that is not desirable clinically.

In addition, analysis of the baseline preprocessing pipeline revealed that it was effective in improving the global image contrast, but it had undesirable smoothing effects. This smoothing prevented fine-grained lesion focus, like microaneurysms and small hemorrhages. The loss of subtle characteristics is particularly detrimental in multi-disease datasets such as RFMiD, in which the presence of DR-related lesions might be faint and easily overlooked.

Lesion-preserving preprocessing and focused data augmentation methods were necessary to make the classification more sensitive without deteriorating overall classification performance.

Data Augmentation:

The diversity of DR samples and the lack of classes were addressed using targeted data augmentation to alleviate the imbalanced nature of classes. Two of the augmentation procedures were selected as the most stable in terms of performance increase, based on a comprehensive review of recent retinal image augmentation methods, as indicated in Figure 6.

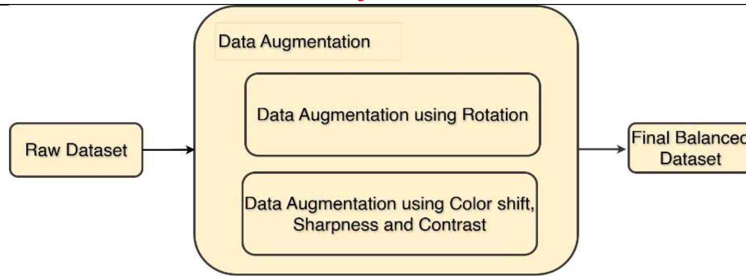

Figure 2. Data Augmentation

Rotation-Based Augmentation:

Rotation-based augmentation was employed to simulate realistic changes in fundus orientation and preserve anatomical accuracy:

Angles of rotation were randomly selected between $-25°$ and $+25°$.

This range was selected to avoid anatomically unreasonable distortions that were anatomically unreasonable.

An additional ten images were generated using clockwise rotations at various angles for each original DR image.

Besides contributing to the intra-class variability, this technique helps the network to learn rotationally invariant properties, which are essential for effective DR detection.

Photometric Augmentation: Contrast Adjustment, Sharpening, and Color-Shifting:

To further enhance the visibility of lesions and simulate natural acquisition variability, the following photometric adjustments were made:

Color Shifting: To reproduce the changes in illumination and color balance that are common in clinical imaging, integer offsets were applied to the RGB channels

Sharpening: To improve edge information, each image was first smoothed using a Gaussian filter with variance $\alpha = 1$. The resulting blurred image was then subtracted from the original.

Contrast adjustment was performed by linearly scaling pixel intensities between predefined thresholds (i_1, i_2), with out-of-range values clipped to 0 or 255. For each image, three contrast-enhanced variants were generated.

It should be noted that only DR-class images were used for all augmentation techniques. This class-specific augmentation strategy maintains the original distribution of non-DR images while directly addressing dataset imbalance. Representative examples of the employed augmentation techniques are shown in Figure 7 and Figure 8.

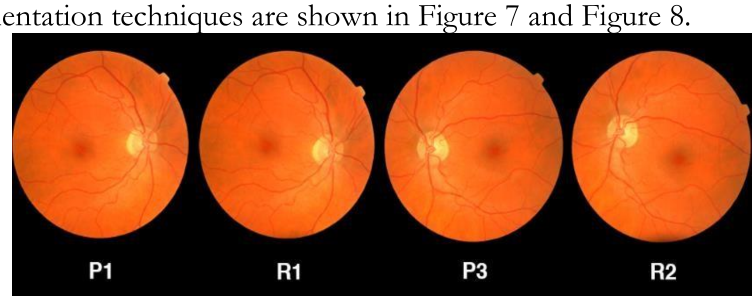

Figure 3. Rotation Augmentation: P1 and P2 are original images and R1 and R2 are rotated images

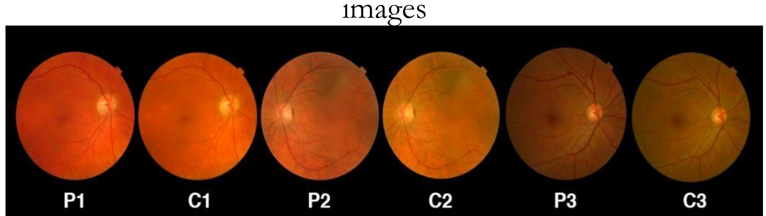

Figure 4. Photometric Augmentation: P1, P2, P3 are Original Images and C1, C2 and C3 are photometric

Proposed Preprocessing Pipeline:

A lesion-aware preprocessing pipeline was implemented after augmentation to improve lesion visibility while maintaining global retinal structure. The proposed preprocessing is shown in Figure 9.

Input Image and Color Channel Decomposition:

The proposed preprocessing pipeline receives a fundus image I_{RGB} as input and makes a sequence of transformations to improve retinal lesions without eliminating clinically valuable visual information. The RGB image is first broken down into its three constituent channels:

$$I_R, I_G, I_B = split(I_{RGB}) \ (1)$$

This is an important step since diabetic retinopathy (DR) lesions like microaneurysms and hemorrhages have characteristic color variations. Channel separation provides chromatic information to be used later in analysis.

Luminance Computation:

A conventional luminance conversion is used to model perceptual brightness:

$$Y = 0.299R + 0.587G + 0.114B \ (2)$$

This equation represents a weighted combination of RGB channels, according to the human visual system's sensitivity to various wavelengths. The green channel contributes most strongly to brightness, then red and blue. The resultant luminance image Y gives a perceptually significant intensity image of retinal structures.

Gamma Correction using Adaptive Histogram:

Histogram-based gamma correction is done to improve contrast adaptively. To begin with, we compute the histogram of the luminance image $H(i)$, and add it to a uniform histogram. $H_u(i)$ as.

$$H_m(i) = \alpha H(i) + (1 - \alpha)H_u(i) \ (3)$$

Where $\alpha = 0.7$. The weighted formulation is an equal weight prior plus actual intensity distribution, avoiding extreme histogram bias. Then, the cumulative distribution function (CDF) is calculated:

$$CDF(i) = \sum H_m(i) \ (4)$$

The CDF is the cumulative distribution of intensity levels, and it is used to inform adaptive enhancement. The value of adaptive gamma is, then, defined as:

$$\gamma(x, y) = \frac{1 - CDF(I(x, y))}{0.7} \ (5)$$

This formulation guarantees spatially varying gamma correction. Pixels of lower intensity will be enhanced more, enhancing the visibility of darker parts of the retina and avoiding overexposure in brighter ones.

Intensity Gain and Dynamic Range Partitioning by Quantile:

Following gamma correction, the intensity gain is calculated to measure the improvement effects and reduce nonlinear distortions. A second histogram is created to recalibrate the intensity distribution. The image is subsequently subdivided into three intensity areas, which are quantile-based:

Low intensity region

Mid-intensity region

High intensity region

This division allows adaptive processing of various brightness levels, whereby dark lesions and bright anatomical structures are maintained.

Normalized Power Transformation and Global Enhancement:

A normalized power transformation is used to smooth out the variations in intensity:

$$I_{norm} = \left(I_{quant}\right)^{\beta} \ (6)$$

where β is the intensity scaling. This is a shift that minimizes the drastic differences and enhances international uniformity. This is followed by histogram equalization and normalization to boost the global contrast and make sure that there is even distribution of intensity throughout the image.

The output of these steps is named a weighted image.

Retinal Mask Generation:

To ensure that subsequent preprocessing stages only concentrate on clinically significant regions, a retinal mask was created to eliminate background artifacts and isolate the retinal region.

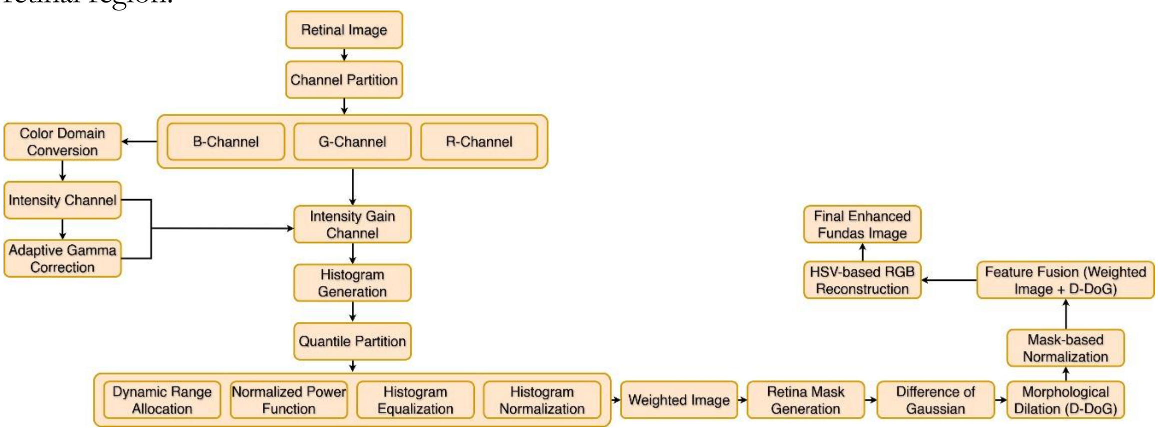

Figure 5. Proposed Preprocessing Pipeline

Difference of Gaussian (DoG):

The Difference of Gaussian (DoG) filter is used to sharpen high-frequency lesion images like microaneurysms and hemorrhages by removing the low-frequency background data. DoG operation can be mathematically expressed as:

$$I_{DoG}(x,y) = G_{\sigma 1}(x,y) * I_w(x,y) - G_{\sigma 2}(x,y) * I_w(x,y) \quad (7)$$

Where

$I_{DoG}(x,y)$ is a weighted input image

$G_\sigma(x,y)$ is the Gaussian kernel of standard deviation σ.

$*$ denotes a convolutional operation

A 9×9 kernel is chosen, which provides adequate local spatial coverage and keeps small retinal structures like microaneurysms and small hemorrhages. Though larger kernels may provide a more accurate approximation of full Gaussian support, a smaller kernel is preferred to prevent excessive smoothing and loss of clinically relevant details. The selected kernel offers an effective compromise between computational cost and preservation of features.

The standard deviation $\sigma_1 = 2.0$ is chosen to smooth the data at the fine level, which can remove noise effectively while preserving lesion-level structures. The definition of the second scale $\sigma_2 = \alpha \cdot \sigma_1$, where $\alpha = 2$ value, is sufficient to provide a good separation between small structural details and larger background light. This scale ratio is typical in Difference of Gaussian frameworks since it offers a good approximation to Laplacian of Gaussian behaviour, but still offers good feature enhancement.

Dilated Difference of Gaussian (DDoG):

Morphological dilation is then done to the Difference of Gaussian (DoG) result to further improve lesion visibility, as well as structural continuity in the response map. This is a more refined version of the formulation, known as the Dilated Difference of Gaussian (DDoG), which enhances lesion areas by boosting weak activations and spatial coherence. The DDoG operation can be defined as:

$$I_{DDoG} = I_{DoG} \oplus S \quad (8)$$

Where I_{DDoG} represents the Difference of Gaussian response map, S the structuring element, and $\oplus$ the morphological dilation operator. A small 2×2 structuring element is used, which is as follows:

$$S = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \quad (9)$$

By applying dilation once, fine structural boundaries are maintained, yet connectivity is increased. Morphological dilation works by increasing the size of the identified lesion areas based on the prescribed structuring element. This operation adds the neighboring pixels to the lesion regions and enhances the weak or fragmented responses of the DoG output. Consequently, the DDoG transformation has a few benefits: first, it improves the boundaries of faint lesions, maintains structural consistency, and minimizes the risk of over-segmentation; second, it makes the subsequent deep learning models more sensitive to small-scale pathological images. In general, the proposed DDoG representation enhances the intensity of lesions and preserves anatomical integrity, which is more apt for downstream classification and analysis using convolutional neural networks.

Fusion:

In order to incorporate improved lesion information while preserving contextual retinal features, the DDoG mask was replicated across RGB channels and additively fused with the weighted image.

HSV-based RGB Reconstruction:

HSV-based RGB reconstruction was used in the last stage to enhance contrast balance and color fidelity. By taking this step, the natural appearance of retinal structures is not distorted by lesion amplification. Figure 10 shows a visual comparison of the original images, baseline preprocessing, and the suggested preprocessing workflow.

Figure 11 shows another zoomed shot of the comparison among raw, baseline preprocessing, and proposed preprocessing pipeline in terms of lesion detection.

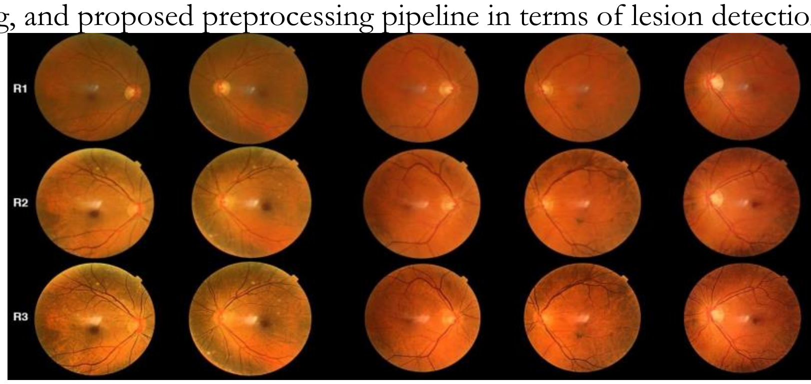

Figure 6. Preprocessing Comparison, Row 1 corresponding to Original Images, Row 2 corresponding to Baseline Preprocessed Images, and Row 3 corresponding to Proposed Preprocessed Images.

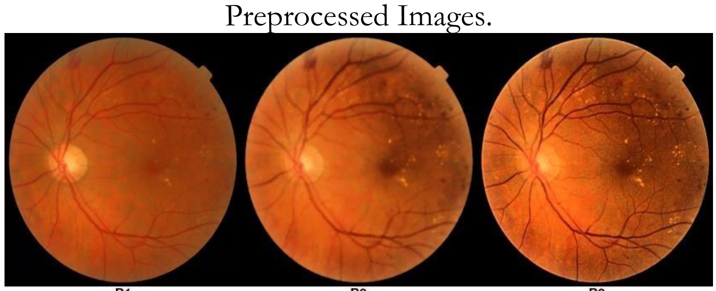

Figure 7. P1: Raw Image, P2: Baseline Preprocessed Image, P3: Proposed Preprocessed Image

The algorithm of our proposed lesion-aware preprocessing is shown in Figure 12.

Algorithm 1: Proposed Lesion-Aware Retinal Image Preprocessing Pipeline

Input : RGB fundus image $I_{RGB} \in \mathbb{R}^{H \times W \times 3}$
Output : Preprocessed image $I_{final} \in \mathbb{R}^{H \times W \times 3}$

1 **Channel Partitioning**
$I_R, I_G, I_B = \text{split}(I_{RGB})$ // *Split RGB image into R, G, B channels*

2 **Luminance Computation**
$Y = 0.299R + 0.587G + 0.114B$ // *Compute luminance using standard formula*

3 **Histogram Computation**
Compute histogram of $Y \rightarrow H(i)$ // *Histogram of luminance image*

4 **Modified Histogram**
$H_m(i) = \alpha H(i) + (1-\alpha)H_u(i),\ \alpha = 0.7$ // *$H_u(i)$ is uniform histogram*

5 **CDF Calculation**
$CDF(i) = \sum H_m(i)$ // *Cumulative distribution function*

6 **Adaptive Gamma Estimation**
$\gamma(x, y) = \frac{1 - CDF(I(x,y))}{0.7}$ // *Compute adaptive gamma map*

7 **Gamma Correction**
$I_\gamma = Y^{\gamma(x,y)}$ // *Apply adaptive gamma correction*

8 **Intensity Gain & Re-histogram**
Compute intensity gain and histogram of I_γ // *Compensate non-linear effects*

9 **Quantile-Based Partitioning**
Partition intensities into: Low, Mid, High // *Three quantile levels*

10 **Normalized Power Transformation**
$I_{norm} = (I_{quant})^\beta$ // *β : normalization parameter*

11 **Histogram Equalization & Normalization**
$I_{enhanced} = \text{normalize}(\text{equalize}(I_{norm}))$ // *Enhance global contrast*

12 **Difference of Gaussian (DoG) Filtering**
$I_{DoG} = G_{\sigma_1}(I_{enhanced}) - G_{\sigma_2}(I_{enhanced}),\ \alpha = 2$ // *$\alpha = \sigma_2/\sigma_1 = 2$ (scale factor)*

13 **Morphological Dilation**
$I_{dilated} = \text{dilate}(I_{DoG}, SE_{2\times 2})$ // *Dilation with 2×2 structuring element*

14 **Feature Fusion (Channel Replication)**
$I_{fusion} = [\ I_{dilated}, I_{dilated}, I_{dilated}\ ]$ // *Replicate across 3 channels*

15 **HSV-Based Reconstruction**
$I_{final} = \text{HSV}^{-1}(H, S, V)$ from I_{fusion} // *Reconstruct RGB image using HSV*

Return: I_{final} // *Final preprocessed image*

Figure 8. Algorithm of Proposed Preprocessing Pipeline

Implementation Details and Evaluation Metrics:

To be reproducible, the entire experimental setup can be summarised in Table 1, given below:

Table 1. Experimental setup Details

Component	Configuration
Framework	PyTorch (Python 3.9)
Hardware	NVIDIA GPU (CUDA-enabled, 16–24 GB VRAM), Intel Xeon/AMD CPU
Optimizer	AdamW
Loss Function	Weighted Loss (to handle class imbalance)
Batch Size	16
Epochs	15
Training Strategy	Early Stopping (based on validation loss)
Learning Rate	Empirically tuned (validated on validation set)

To evaluate the performance of the models, the following evaluation metrics were employed:
Accuracy
Sensitivity
Specificity
Area Under the Curve (AUC)

Since the clinical significance of early Diabetic Retinopathy (DR) identification is critical, the sensitivity aspect is given more priority so that the pathological cases are not overlooked.

The choice of hyperparameters was informed by empirical validation and conventional wisdom on hyperparameters in transfer learning-based medical image classification. A batch size of 16 was selected because of the GPU memory limitation, but with a stable gradient estimation. The epochs were restricted to 15 since no significant improvement was observed in the additional training, and early stopping was used to avoid overfitting.

AdamW optimizer was chosen because it has a better regularization property compared to Adam, especially when overfitting in small-to-moderate medical datasets. To balance the issues of class imbalance in the RFMiD dataset, a weighted loss function was added, so that minor classes (samples with lesions) play a sufficient role in the optimization.

The learning rate was empirically tuned using validation, where candidate values were evaluated and the best one was selected based on convergence stability and validation performance trends, as is common in the transfer learning literature in medical imaging.

Results and Discussion:

The suggested lesion-preserving preprocessing and class-specific augmentation techniques for diabetic retinopathy (DR) identification on the severely imbalanced RFMiD dataset are thoroughly evaluated in this section. We compared the results of our proposed preprocessing with the baseline results. To create reference metrics, we first replicate the baseline performance of six pretrained CNN models, emphasizing sensitivity constraints brought on by class imbalance and loss of small lesion features. We then evaluate how rotation and contrast, sharpening, and color-shift augmentations, along with the suggested preprocessing pipeline, affect model performance. The results are discussed in terms of accuracy, sensitivity, specificity, and area under the ROC curve (AUC), although special attention is paid to the sensitivity, which is of great importance when it comes to clinical screening use. A comparison between models and augmentation techniques is provided to demonstrate the effectiveness and clinical applicability of the proposed methodology.

Experimental Setup:

The proposed methods, class-specific augmentation, and adaptive preprocessing were evaluated on the Retinal Fundus Multi-disease Image Dataset (RFMiD), which is an extremely imbalanced dataset consisting of many instances of co-occurring retinal diseases along with rare cases of DR. Three experimental designs were adopted to do binary DR classification using six pretrained CNNs: AlexNet, ResNet-18, VGGNet-s, VGGNet-16, VGGNet-19, and GoogleNet.

Baseline preprocessing: Original baseline pipeline without any modification.

Rotation-based augmentation with proposed preprocessing: Geometric augmentation on DR class images only to augment intra-class diversity.

Contrast, sharpening, and color-shift augmentation with proposed preprocessing: Photometric transformations applied to DR-class images to simulate real-world variations in lighting, contrast, and color distribution.

Accuracy, sensitivity, specificity, and area under the ROC curve (AUC) were used to assess the model's performance. Sensitivity was given special attention because it is crucial in clinical DR screening to reduce missed diagnoses.

Evaluation of Lesion Aware Preprocessing:

To measure the success of the suggested lesion-aware preprocessing framework, image quality metrics were employed to measure structural preservation and reconstruction quality. Specifically, Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM) were used since these measures are popular in evaluating the trade-off between noise reduction and structural fidelity in medical image enhancement applications. Table 2 shows the PSNR and SSIM values of the proposed preprocessing compared with baseline values.

Table 2. PSNR and SSIM comparison of the proposed preprocessing with the baseline preprocessing

Preprocessing Method	PSNR (dB)	SSIM
Adaptive Deep Convolutional Preprocessing [1]	28.72	0.72
Proposed Preprocessing Framework	**31.26**	**0.88**

A larger PSNR corresponds to less distortion, and a larger SSIM corresponds to a better structural similarity to the reference image.

The suggested preprocessing framework shows a remarkable enhancement over the baseline adaptive deep convolutional preprocessing approach. In particular, the PSNR growth of 28.72 dB to 31.26 dB means that the reconstruction error decreases, and the signal fidelity increases. The same can be said about the enhancement of SSIM between 0.72 and 0.88, as this is an indication of superior structural information to retain lesion integrity in retinal images.

These findings validate that the proposed lesion-aware preprocessing is effective to improve the quality of the image and, at the same time, reduce noise and maintain clinically useful structural information, thus establishing its appropriateness in the downstream Diabetic Retinopathy classification tasks and hence, objective 1 of this study is achieved.

Baseline Performance on RFMiD:

The reproduced baseline results are summarized in Table 3. Sensitivity varied and was typically subpar, ranging from 81.45% to 91.13% across CNN architectures, whereas overall accuracy and AUC were moderate to high.

Table 3. Classification Results of RFMiD on Baseline Methodology

Model	Accuracy	Sensitivity	Specificity	AUC
AlexNet	84.38%	83.06%	84.69%	0.9157
ResNet-18	91.56%	85.48%	93.02%	0.9655
VggNet-s	90.94%	82.26%	93.02%	0.9497
VggNet-16	86.56%	91.13%	85.47%	0.9491
VggNet-19	90.62%	81.45%	92.83%	0.9531
GoogleNet	91.25%	90.32%	91.47%	0.9716

The highest sensitivity scores (90.32% and 91.13%, respectively) were notably attained by GoogleNet and VGGNet-16, but frequently at the expense of specificity, suggesting a trade-off between false positives and false negatives. Despite increased global contrast, visual examination of baseline-pre-processed images showed that fine-lesion-level features, such as microaneurysms and tiny hemorrhages, were smoothed out. The observed decrease in sensitivity can be explained by this over-smoothing, which has a substantial impact on DR detection in multi-disease datasets where lesions may be minor and easily missed. Figure 13 shows the confusion matrix of the best-performing model in baseline classification results.

These results emphasize the necessity for lesion-preserving improvement techniques and the limits of conventional preprocessing processes when applied to real-world, imbalanced, Multi-disease datasets.

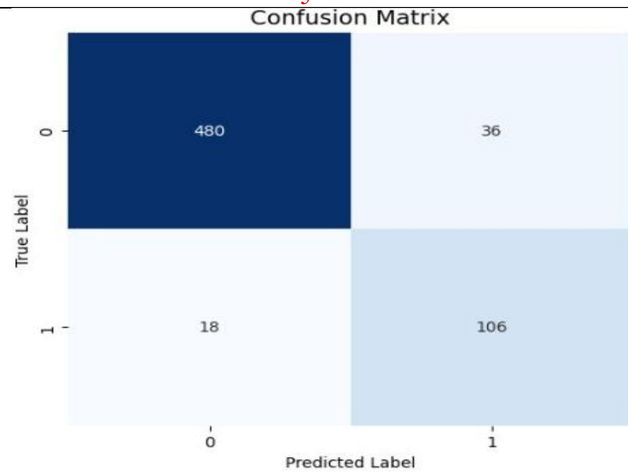

Figure 13. ResNet-18 Confusion Matrix on Baseline Methodology

Effects of Rotation-Based Augmentation and Proposed Preprocessing Pipeline:

All six CNN architectures showed notable increases in sensitivity when rotation-based augmentation was introduced and applied preferentially to DR-class images in conjunction with the suggested preprocessing workflow (Table 4). GoogleNet achieved a balanced improvement with sensitivity going from 90.32% to 91.94% and a maintained AUC of 0.9583, while VGGNet19 demonstrated the biggest gain, with sensitivity rising from 81.45% to 95.16%.

Table 4. Classification results of RFMiD after rotation-based augmentation and proposed preprocessing

Model	Accuracy	Sensitivity	Specificity	AUC
AlexNet	89.53%	86.29%	90.31%	0.9379
ResNet-18	91.25%	84.68%	92.83%	0.9510
VggNet-s	90.00%	83.06%	91.67%	0.9430
VggNet-16	91.72%	88.71%	92.44%	0.9605
VggNet-19	88.59%	95.16%	87.02%	0.9596
GoogleNet	89.06%	91.94%	88.37%	0.9583

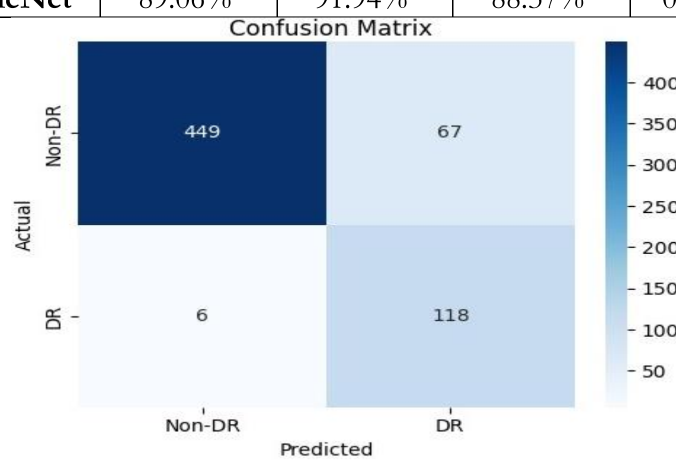

Figure 14. VggNet-19 Confusion Matrix on Rotation Augmentation and Proposed Preprocessing

Geometric augmentation increases intra-class variation, which improves the models' ability to generalize to different DR lesion orientations while maintaining anatomical integrity. This trade-off is clinically acceptable because decreasing false negatives is significantly more important than slightly raising false positives in DR screening, even when sensitivity advances occasionally coincided with a slight drop in specificity (e.g., AlexNet and VGGNet-s). Figure

14 shows the confusion matrix of the best-performing model's classification results with rotation augmentation and proposed preprocessing.

Effects of Contrast, Sharpening, and Color-Shift Augmentation with Proposed Preprocessing:

To mimic real-world differences in lighting, contrast, and color distribution, photometric augmentation was used. The related findings are shown in Table 5. This technique increased overall feature discrimination, as evidenced by continuously high AUC values (up to 0.9646 for VGGNet-19) and improved robustness against color and contrast changes, even if sensitivity increases were not as significant as those obtained by rotation-based augmentation.

Table 5. Classification Results of RFMiD after photometric augmentation and proposed preprocessing

Model	Accuracy	Sensitivity	Specificity	AUC
AlexNet	85.62%	83.87%	86.05%	0.9155
ResNet-18	86.56%	89.52%	85.85%	0.9395
VggNet-s	89.06%	87.90%	89.34%	0.9475
VggNet-16	86.25%	91.94%	84.88%	0.9548
VggNet-19	91.41%	90.32%	91.67%	0.9646
GoogleNet	90.62%	90.32%	90.70%	0.9548

These findings suggest that contrast and color-based augmentation improve generalization in multi-disease contexts by strengthening the model's capacity to catch minute photometric differences in DR lesions. However, class imbalance cannot be entirely compensated for by photometric augmentation alone, underscoring the complementary significance of rotation-based augmentation. Figure 15 shows the confusion matrix of the best-performing model classification results with proposed preprocessing and photometric augmentation.

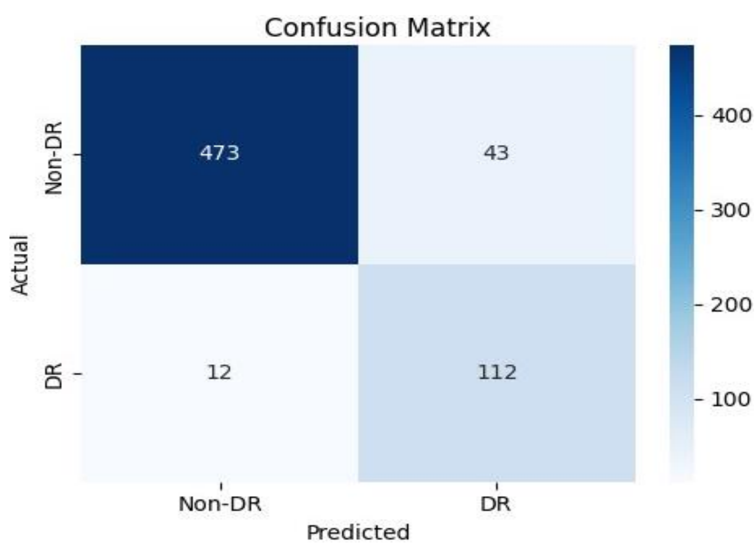

Figure 15. VggNet-19 Confusion Matrix on Photometric Augmentation and Proposed Preprocessing

Comparative analysis across Augmentation Strategies:

The comparative analysis shows that:

Rotation-based augmentation is especially useful for screening-focused applications since it continuously maximizes sensitivity.

Photometric augmentation provides robustness against illumination and color fluctuations while stabilizing AUC and enhancing discrimination across models.

The most robust architectures were VGGNet-19 and GoogleNet, which benefited from both augmentation techniques and lesion-preserving adaptive preprocessing by consistently improving metrics.

A significant drawback of baseline pipelines in multi-disease, imbalanced datasets is addressed by the combination of lesion-preserving adaptive preprocessing and tailored augmentation, which lowers false negatives. The clinical difficulty of identifying small DR lesions, which are frequently disregarded in traditional preprocessing techniques, is immediately addressed by this dual strategy. Hence objective 2 and 3 are achieved as well.

Discussions and Clinical Implications:

The experimental findings verify that baseline preparation procedures designed for balanced datasets are inadequate for real-world, multi-disease scenarios such as RFMiD. Without considerably sacrificing sensitivity or AUC, the suggested approach, which combines class-specific augmentation and lesion-aware preprocessing, significantly improves lesion visibility and sensitivity. Figures 16, 17, and 18 show the comparison of accuracy, sensitivity, and specificity.

The proposed framework exhibits better sensitivity and AUC than the recent state-of-the-art methods that are reported in the literature, which mainly utilize standard augmentation and global feature learning. Although current approaches claim to be highly accurate in cases of balanced or single-disease datasets, they exhibit moderate performance in multi-disease settings like RFMiD. Compared to this, our approach offers consistent performance between CNN architectures, which is a strength and allows clinical use.

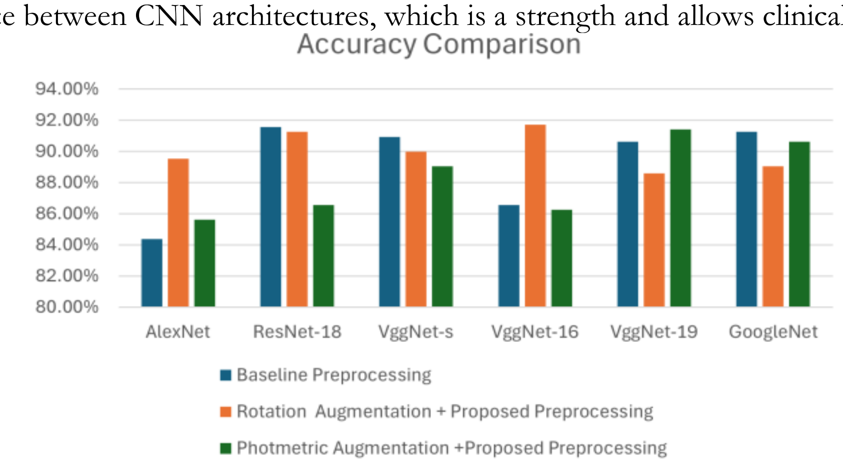

Figure 16. Accuracy Comparison

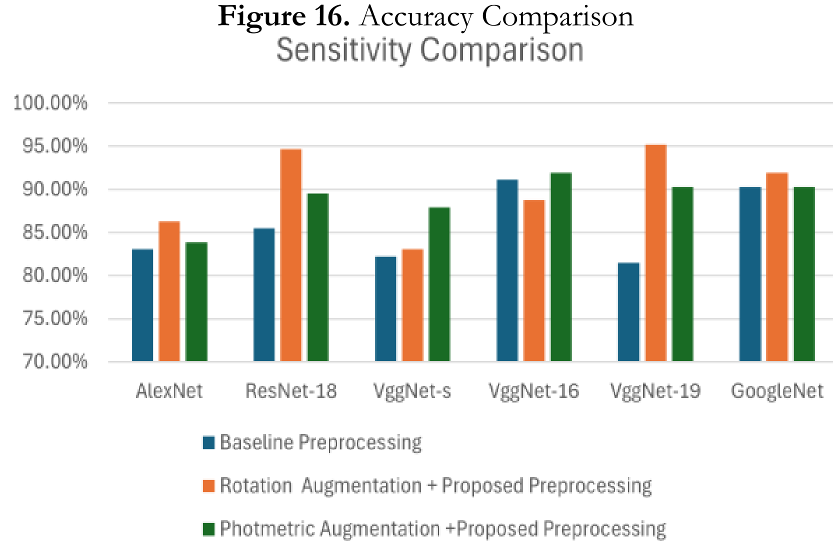

Figure 17. Sensitivity Comparison

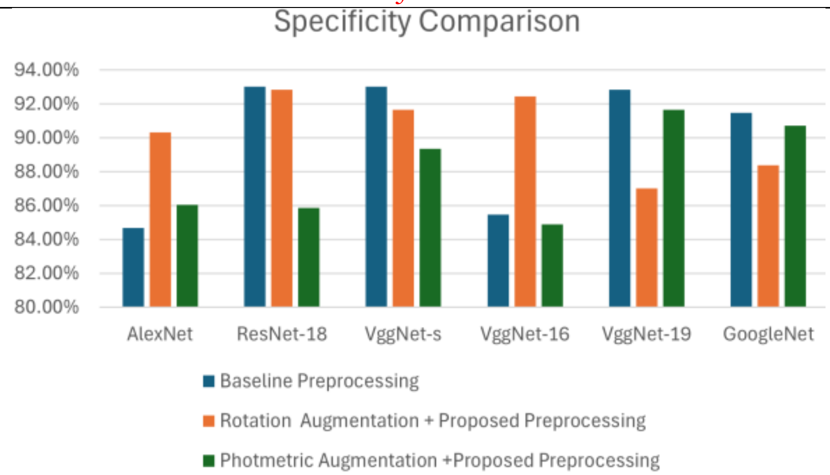

Figure 18. Specificity Comparison

This is pertinent from a clinical standpoint:

Decreased false negatives: The suggested pipeline reduces missed diagnoses, which are vital in averting irreversible vision loss, by enhancing the detection of small DR lesions.

Balanced performance: In DR screening, where sensitivity is given priority, slight increases in false positives are acceptable.

Generalizability: The technique demonstrates that preprocessing and augmentation can be customized to aspects unique to a dataset, providing a practical system for DR screening in practice.

In general, by emphasizing the essential applicability of lesion-preserving preprocessing to multi-disease samples and ensuring that small signals related to DR are retained to be accurately classified, the present study makes substantial contributions to the area of automated diabetic retinopathy detection. Moreover, class-specific augmentation enhances the detection of underrepresented cases of DR by effectively dealing with the challenge of extreme dataset imbalance. The proposed design demonstrates great and clinically relevant improvements in sensitivity, accuracy, and AUC by empirical testing over various pretrained CNN models, and indicates its capacity to be reliably applied in clinical practice in the context of DR screening.

Implications of the Study:

The intended research has clinical, technical, and research implications.

Clinical implications: The increased lesion visibility and sensitivity minimize false negatives, which is important in the early detection of Diabetic Retinopathy and in preventing permanent loss of vision. The suggested framework thus helps the trustworthy screening with AI assistance in practical ophthalmic conditions.

Technical implications: Feature-preserving preprocessing (when combined with lesion-aware pre-processing) is shown to substantially enhance the performance of deep learning, especially when working with imbalanced multi-disease datasets.

Research implications: The work provides a generalizable preprocessing and augmentation pipeline that can easily be scaled to other medical imaging tasks, emphasizing the role of lesion-sensitive enhancement and dataset-specific augmentation approaches in enhancing CNN robustness.

Recommendations:

Though the suggested framework shows excellent results in terms of improving lesion visibility and accuracy in classifications, there are still several directions that can be followed in the future. First, the test of statistical significance (e.g., t-test or ANOVA) can be included to further confirm the performance gains in various configurations. Second, the present research can be expanded to transformer-based models and CNN-Transformer hybrids to enhance global context learning in retinal images.

Moreover, future research could also examine domain adaptation methods to improve generalization in various retinal datasets obtained with various imaging devices. Lastly, the clinical applicability and strength of the proposed framework would be further enhanced by real-time implementation in clinical screening systems and testing on a larger multi-disease dataset.

Conclusion:

Using the highly imbalanced, multi-disease RFMiD dataset, this study examined the limits of a baseline CNN pipeline for diabetic retinopathy (DR) identification. The baseline models showed reasonable accuracy (84.38%–91.56%) and AUC (0.9157–0.9716); however, sensitivity was much lower for several designs, ranging from 81.45% to 91.13%, indicating missing DR instances because of class imbalance and over-smoothing of small lesion markers. To overcome these difficulties, we suggested two specific augmentation techniques: rotation and photometric treatments (color shifting, contrast adjustment, and sharpening) in conjunction with an improved lesion-preserving preprocessing pipeline that incorporates Difference of Gaussian (DoG) and Dilated DoG (DDoG).

Rotation-based augmentation with the suggested pre-processing enhanced sensitivity by up to 13.71% (VGGNet 19: from 81.45% to 95.16%) and balanced AUC gains (up to 0.9596), according to experiments using six pretrained CNN models (AlexNet, ResNet-18, VGGNet-s, VGGNet-16, VGGNet-19, GoogleNet). While retaining excellent accuracy and specificity, photometric augmentation improved feature discrimination by 8–10% for deeper networks like ResNet 18 and VGGNet-16. These findings unequivocally show that the combined approach of lesion-preserving preprocessing and class-specific augmentation successfully tackles the difficulties of imbalanced, multi-disease datasets, resulting in increased DR classification that is clinically significant.

The suggested methodology can be expanded for future research to include multi-class DR severity grading and the identification of additional retinal illnesses found in multi-disease datasets, which would improve its clinical usefulness. Furthermore, including transformer topologies or attention-based deep learning models should improve the identification of small lesions while lowering false positives. Model generalization may also be enhanced by investigating automated, adaptive augmentation techniques that dynamically adapt to dataset properties. Lastly, the pipeline's resilience and preparedness for actual clinical deployment would be confirmed by testing it on bigger, heterogeneous, multi-center datasets.

References:

[1] R. Abbasi *et al.*, "Diabetic retinopathy detection using adaptive deep convolutional neural networks on fundus images," *Sci. Reports 2025 151*, vol. 15, no. 1, pp. 24647-, Jul. 2025, doi: 10.1038/s41598-025-09394-0.

[2] R. Kumar, V. Kohli, R. K. Singh, and R. K. Ratnesh, "Multi-label Deep Learning Framework for Early Detection of Diabetic Retinopathy Diseases," *IEEE Int. Conf. Next Gener. Inf. Syst. Eng. NGISE 2025*, 2025, doi: 10.1109/NGISE64126.2025.11085237.

[3] A. R. Majeed, W. A. Awan, N. Ul Hassan, M. A. Asghar, and M. J. Khan, "Retinal Fundus Image Refinement with Contrast Limited Adaptive Histogram Equalization, Noise Filtration and Intensity Adjustment," *Proc. - 2020 23rd IEEE Int. Multi-Topic Conf. INMIC 2020*, Nov. 2020, doi: 10.1109/INMIC50486.2020.9318104.

[4] Muhammad Hassaan Ashraf, Muhammad Nabeel Mehmood, "HIRD-Net: An Explainable CNN-Based Framework with Attention Mechanism for Diabetic Retinopathy Diagnosis Using CLAHE-D-DoG Enhanced Fundus Images," *Life*, vol. 15, no. 9, p. 1411, 2025, doi: https://doi.org/10.3390/life15091411.

[5] J. R. Balashunmugam, M. M. R. Sindha, A. Makkie, and U. M. Pandiyan, "Image enhancement techniques for fundus images - A review," *AIP Conf. Proc.*, vol. 2857, no. 1, Aug. 2023, doi: 10.1063/5.0164304.
[6] Chaichana Suedumrong, Suriya Phongmoo, "Diabetic Retinopathy Detection Using Convolutional Neural Networks with Background Removal, and Data Augmentation," *Appl. Sci.*, vol. 14, no. 19, p. 8823, 2024, doi: https://doi.org/10.3390/app14198823.
[7] Israa Y. Abushawish, Sudipta Modak, "Deep Learning in Automatic Diabetic Retinopathy Detection and Grading Systems: A Comprehensive Survey and Comparison of Methods," *IEEE Access*, 2024, doi: 10.1109/ACCESS.2024.3415617.
[8] Zubair Khan, Fiaz Gul Khan, "Diabetic Retinopathy Detection Using VGG-NIN a Deep Learning Architecture," *IEEE Access*, 2021, doi: 10.1109/ACCESS.2021.3074422.
[9] Chun Ling Lin, Kun Chi Wu, "Development of revised ResNet-50 for diabetic retinopathy detection," *BMC Bioinformatics*, vol. 24, no. 1, p. 157, 2023, [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/37076790/
[10] S. Das, A. Lasker, M. Ghosh, S. M. Obaidullah, and K. Roy, "A Deep Learning-based Approach for Detecting Diabetic Retinopathy in Retina Images," *Internet Things-Based Mach. Learn. Healthc. Technol. Appl.*, pp. 85–95, Jan. 2024, doi: 10.1201/9781003391456-5/DEEP-LEARNING-BASED-APPROACH-DETECTING-DIABETIC-RETINOPATHY-RETINA-IMAGES-SAHANA-DAS-ASIFUZZAMAN-LASKER-MRIDUL-GHOSH-SK-MD-OBAIDULLAH-KAUSHIK-ROY.
[11] H. Naz and N. J. Ahuja, "A novel contrast enhancement technique for diabetic retinal image pre-processing and classification," *Int. Ophthalmol.*, vol. 45, no. 1, Dec. 2024, doi: 10.1007/S10792-024-03377-2.
[12] R. Alaguselvi and K. Murugan, "Quantitative analysis of Fundus Image Enhancement in the Detection of Diabetic Retinopathy Using Deep Convolutional Neural Network," *IETE J. Res.*, vol. 69, no. 9, pp. 6315–6325, Sep. 2023, doi: 10.1080/03772063.2021.1997356.
[13] C. Shorten and T. M. Khoshgoftaar, "A survey on Image Data Augmentation for Deep Learning," *J. Big Data*, vol. 6, no. 1, pp. 1–48, Dec. 2019, doi: 10.1186/S40537-019-0197-0/FIGURES/33.
[14] E. Goceri, "Medical image data augmentation: techniques, comparisons and interpretations," *Artif. Intell. Rev.*, vol. 56, no. 11, pp. 12561–12605, Nov. 2023, doi: 10.1007/s10462-023-10453-z.
[15] F. M.Javed Mehedi Shamrat, Rashiduzzaman Shakil, "An advanced deep neural network for fundus image analysis and enhancing diabetic retinopathy detection," *Healthc. Anal.*, vol. 5, p. 100303, 2024, doi: https://doi.org/10.1016/j.health.2024.100303.