

Early Detection of Diabetic Retinopathy using Machine Learning

Shamika Kadav
MIT ADT UNIVERSITY
PUNE
+91 7030677003
shamikadav@gmail.com

Apurva Patil
MIT ADT UNIVERSITY
PUNE
+91 78210 93399
apurva122001@gmail.com

Aryan Tanwar
MIT ADT UNIVERSITY
PUNE
+91 83027 56425
aryan.taanwar2001@gmail.com

Sachin Takale
MIT ADT UNIVERSITY
PUNE
+91 96899 34042
sachin.takale@mituniversity.edu

ABSTRACT

Diabetes is a widespread condition that can cause significant side effects, such as Diabetic Retinopathy (DR), which is damage to the blood vessels in the retina. This disease is a major cause of vision loss in advanced nations, where there is a considerable risk of blindness. By 2030, it is predicted that 44% of the world's population would have diabetes, affecting an estimated 366 million people globally. Traditionally, qualified medical professionals have used manual screening of color fundus pictures to diagnose DR. However this approach takes a lot of time and is prone to human mistake, especially when handling a multitude of diabetes patient photos. A rising number of people are interested in using machine learning (ML) techniques for automated screening as a solution to these problems. Automated machine learning-based screening has the potential to optimize the detection process, lessen the workload for medical practitioners, and enhance diagnosis precision and consistency. These systems can evaluate retinal pictures more effectively and accurately than prior methods by using sophisticated algorithms to identify possible lesions suggestive of DR.

The aim of this study is to develop a robust system that can precisely analyze retinal images for early indicators of diabetic retinopathy by utilizing machine learning techniques. This would enable swift treatment and medical care. The technology attempts to improve the precision of identifying various phases of retinopathy by utilizing cutting-edge image processing algorithms, giving clinicians important information for patient care. In the end, our strategy seeks to reduce the variation in diagnoses based on human variables, guaranteeing accurate and consistent detection of diabetic retinopathy.

Keywords

Diabetes; Disease; Diabetic Retinopathy; Retina; Retinal Blood Vessels; Blindness; Machine Learning; Algorithms; Automated Screening; Manual Detections; Color Fundus Images; Early Intervention; Accuracy; Efficiency; Diagnosis.

1. INTRODUCTION

The severe side effect of diabetes that affects the eyes is called diabetic retinopathy. Damage to the blood vessels in the retina, the light-sensitive tissue at the back of the eye, is the root cause of it. If dealt with, this damage can cause blindness in addition to a variety of vision issues.

Blurred vision, floaters (spots or lines that appear to float in the field of vision), poor color vision, and the sense of black or empty areas in the visual field are some of the early symptoms of diabetic

retinopathy. Furthermore, some people may lose vision in particular fields of vision.

Diabetic retinopathy can have serious adverse consequences, such as blindness and visual impairment. In severe situations, retinal detachment—the separation of the retina from the posterior part of the eye—may occur. The development of mutant blood vessels in the retina, known as neovascularization, may result in more issues. Eventually, diabetic retinopathy may cause central vision loss, which could hamper an individual's clarity of vision and capacity to carry out routine tasks. To prevent or mitigate these effects, early detection and treatments are essential.

Each pixel in a digital image has a limited amount of grayscale or intensity. Spatial coordinates (x, y) on the x- and y-axes, respectively, define these values. Whether a picture has a set resolution determines the type of image it is.

Image Types -

1. Binary Images: Typically, these basic images have two values, 0 and 1, which are black and white. It is a 1-bit image since each pixel is represented by a single binary digit. These pictures are frequently utilized in processes when simply a general shape or contour is required, like optical character recognition (OCR). By applying a threshold operation on grayscale images, binary images are frequently produced. This process turns pixels from black ('0') to white ('1') depending on whether they are above or below a threshold level.
2. Grayscale Images: Monochrome images lack any color information; they are composed entirely of grayscale data. The range of gray levels that are available is determined by the bits per pixel. With 8 bits per pixel on average, 256 distinct shades of gray are conceivable in a grayscale image. Images in grayscale, with 12 or 16 bits per pixel.
3. Color Images: The actual information included in the digital picture data is represented by distinct hues assigned to each pixel. Color representation most commonly takes the form of Red, Green, and Blue (RGB) components. An equivalent color image would have 24 bits per pixel (eight bits for each of the three color bands) if the 8-bit monochrome standard served as the reference.
4. Multispectral Images: These images, which are typically the result of sensors such as radar, infrared, ultraviolet, X-ray, or acoustic ones, contain data that is beyond the range of human vision. Multispectral images are different from traditional images since human systems are unable to directly perceive this content. To show the data, the different spectral bands are usually transformed into RGB components.

Image Processing Steps -

1. Image acquisition: is the first stage in acquiring digital images; scaling and formatting may be necessary for preprocessing. It might involve everything from using sensors to take pictures to having access to previously taken digital pictures.

2. Image enhancement: refers to the technique of modifying an image to fit a particular purpose. The goal of enhancement techniques is to increase the image's visual quality for human perception; they are problem-oriented and subjective.
3. Image Restoration: Based on mathematical or probabilistic models of degradation, restoration approaches are objective and aim to improve the appearance of images, in contrast to augmentation. The goal of restoration is to restore a degraded image—caused by noise, blur, or other factors—to its original state.
4. Color Image Processing: As digital images are utilized more frequently, it is important to comprehend color models and processing. Various color spaces are used to represent color images, and processing methods are used to efficiently modify color information.
5. Wavelets: Wavelets enable effective compression and pyramidal representation by expressing images at different resolutions. Wavelets simplify operations like compression and feature extraction by breaking down decomposed images into separate frequency components.
6. Compression: Compression methods lower the amount of storage space or bandwidth needed to send an image. Image compression methods make it more effective to store or transfer photos by minimizing unnecessary information while maintaining image quality.
7. Morphological Processing: Operations that extract usable image components for shape representation are known as morphological operations. These operations, which are used to examine and modify the structure of pictures, include dilatation, erosion, opening, and closing.
8. Segmentation: This technique divides an image into its individual objects or sections. Tasks like item detection and recognition require accurate segmentation. The process entails segmenting an image into significant areas by utilizing pixel intensities, colors, or textures.
9. Feature extraction: is the process of identifying and characterizing quantitative characteristics of picture components after segmentation. These characteristics, which are employed for additional research or categorization, may include corners, edges, textures, or other patterns.
10. Image Pattern Classification: The technique of labeling or classifying things based on their feature descriptors is known as image pattern classification. Classification algorithms come in a variety of forms, from conventional minimum-distance and correlation classifiers to the more advanced deep neural networks, especially the deep convolutional neural networks (CNNs) that are appropriate for image processing applications.

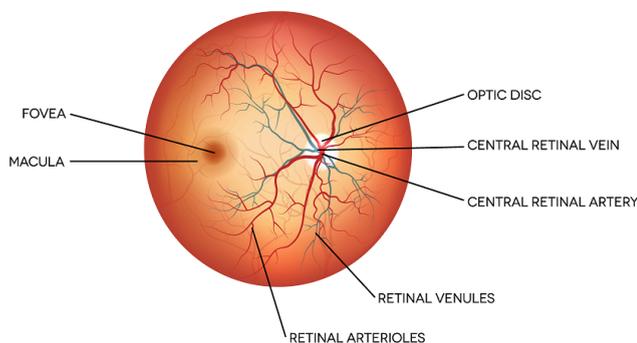


Figure 1. Normal Retina

When the retina is in its original condition, it is a thin layer of tissue lining the back of the eye that is involved in the conversion of light into neural signals that the brain interprets as vision. Blood vessels in a healthy retina typically possess a uniform appearance and show no symptoms of abnormalities.

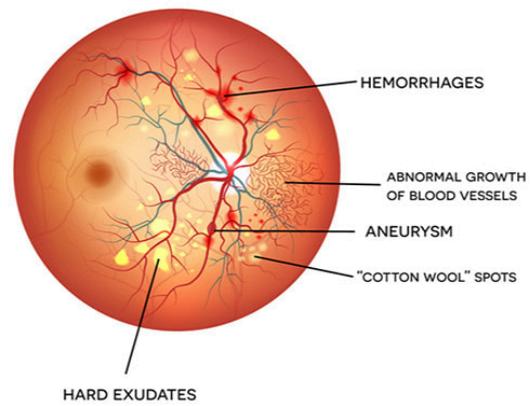


Figure 2. Retina affected by Diabetic Retinopathy

Diabetes damages the retinal blood vessels, resulting in a variety of pathological alterations in a retina damaged by diabetic retinopathy (DR). In the retina, hemorrhages, which are defined by bleeding from the weaker blood vessels, cause blotches or dark areas. Conversely, exudates—yellowish deposits that frequently take the form of tiny, spherical lesions—accumulate in the retina as a result of blood vessels leaking.

Cotton wool spots, which are fluffy, white patches brought on by infarctions in the nerve fiber layer as a result of insufficient blood flow, can also appear. Due to the development of new, irregular blood vessels in response to retinal ischemia, thickened blood vessels—a defining feature of diabetic retinopathy—appear as larger and twisted vessels.

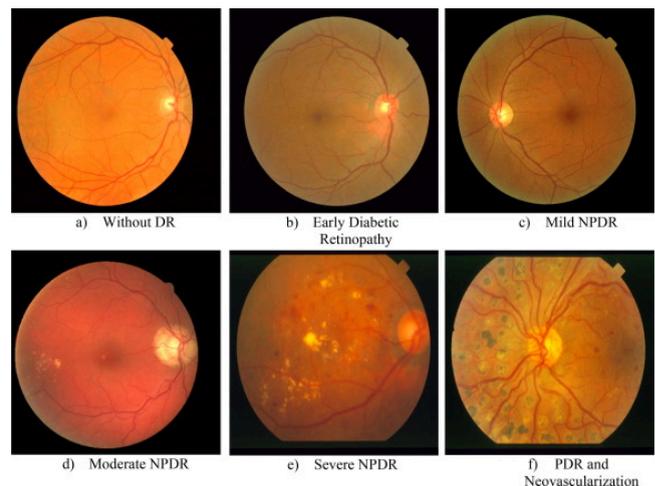


Figure 3. Stage of DR

Stage 1: Diabetic retinopathy with mild nonproliferative damage - Small enlargements in the retina's blood vessels are indicative of the initial stages of diabetic retinopathy. We refer to these swollen regions as micro aneurysms.

At this point, small amounts of fluid leakage into the retina might cause the macula to enlarge. This region is close to the retina's center.

Stage 2: Diabetic retinopathy with moderate nonproliferative changes -

An increase in the size of small blood vessels causes abnormalities in the blood supply to the retina, impeding its normal feeding. Blood and other fluids build up in the macula as a result.

Stage 3: Diabetic retinopathy with severe nonproliferative changes -

The amount of blood flowing to the retina is significantly reduced as a result of a greater portion of its blood vessels becoming clogged. At this stage, the retina starts to produce new blood vessels as a result of signals the body receives.

Stage 4: Diabetic retinopathy with proliferation -

At this point in the disease's progression, the retina is starting to develop new blood vessels. There is an increased chance of fluid leakage since these blood vessels are frequently brittle. This causes a variety of visual issues, including fuzziness, a narrower range of vision, and even blindness.

2. LITERATURE REVIEW

Diabetic retinopathy (DR) is an increasing concern in the field of medicine, especially in developed countries where it is the leading cause of blindness and visual impairment. Since diabetes is predicted to affect 4.4% of the world's population by 2030, the incidence of DR is expected to rise sharply and affect an estimated 366 million people globally. Timely and accurate diagnosis is critical, and is typically achieved by skilled doctors manually examining color fundus images. Nevertheless, this traditional method is tedious and prone to mistakes, especially when dealing with large image datasets. Automated machine learning (ML) approaches have surfaced as viable solutions to address these issues, including more efficient detection procedures, reduced workloads for professionals, and improved diagnostic accuracy.

The computational methodologies established for quality assessment include the use of histogram analysis, both structural and generic visual quality criteria, and local and global generic image quality parameters. The methods typically include segmenting the anatomical structures of the retina, extracting the properties of the anatomical structures, producing generic features that are local and global and characterize the quality factors, and classifying the quality using techniques for classification applied to the features that are extracted from retinal images.

The efficacy of convolutional neural network (CNN) modalities in DR lesion diagnosis has been demonstrated by recent research studies. These CNN architectures provide significant fundus image classification abilities, especially in distinguishing proliferative DR cases and those without apparent DR pathology, although with subtle sensitivity-specificity trade-offs. Furthermore, hybridized approaches that combine statistical models with ensemble classifiers show promise for enhancing diagnostic accuracy and computational effectiveness in drug discovery efforts.

Furthermore, the introduction of computer-aided diagnostic (CAD) systems presents a compelling pathway for disease screening over large population samples, potentially avoiding the time limits associated with manual evaluations. These frameworks enable comprehensive DR categorization and grading based on derived image features by utilizing advanced approaches like kernel support vector machines and K-means clustering. To enhance image fidelity and highlight important characteristics, these methods require careful exudates detection procedures, preliminary preprocessing stages, and post-analysis processing techniques.

Reforms in methodologies for image processing, such as texture feature extraction and wavelet-based edge amplification, significantly improve the diagnostic ability of DR evaluations. CNN models significantly outperform custom-crafted feature-based algorithms, especially when it comes to categorized DR severity based on the presence of exudate. Combining these CNN frameworks with machine learning classifiers, including decision trees and random forests, results in highly precise diagnoses that often exceed 98%.

To summarize, the incorporation of automated machine learning modalities, particularly CNN frameworks, holds great potential for improving the speed and accuracy of diagnosis for drug discovery. These frameworks, which utilize advanced image processing techniques and machine learning algorithms, have the potential to transform DR screening by enabling early intervention and improving patient outcomes while also relieving the workload of medical professionals. However, ongoing research efforts are necessary to tackle ongoing issues, such as reducing false positive rates and strengthening the reliability of automated diagnostic techniques in clinical settings.

3. METHODOLOGY

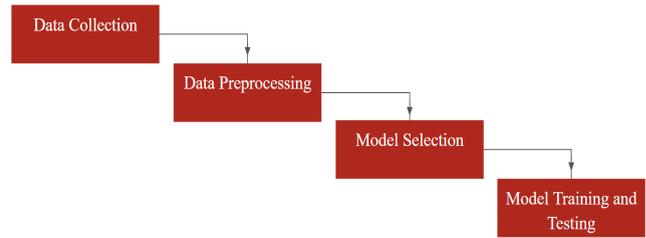


Figure 3. FLOW DIAGRAM - ML MODEL

1. Data collection:

- This stage involves compiling a dataset of retinal fundus images linked to diabetic retinopathy. If untreated, diabetic retinopathy—a consequence of diabetes—affects the eyes and can result in blindness.
- In the current study, a CNN architecture was created to classify fundus images of diabetic retinopathy (DR). The EYEPACS Dataset provided the fundus images. For this study, a total of 11208 fundus pictures were used.

2. Data preprocessing:

- The dataset is preprocessed in a number of ways to improve the efficiency of model training:
- Cleaning: To stop noisy or corrupted images from impairing the model's performance, they can be eliminated from the dataset.
- Normalization: To promote convergence during training and guarantee consistency between images, pixel values in the images are normalized to a standard range (e.g., [0, 1]).
- Data Augmentation: To improve the model's ability to generalize to unknown data, techniques like rotation, flipping, zooming, and shifting can be used to artificially increase the size of the dataset.
- Scaling: To guarantee that the images meet the input specifications of the selected machine learning model, they are resized to a standard size.

3. Model Selection :

- A machine learning model is chosen that is appropriate for the purpose of diagnosing diabetic retinopathy. Convolutional neural networks (CNNs) are widely utilized in image classification problems because of their capacity to automatically extract relevant data from images.
- The CNN model's architecture is determined by a number of criteria, including the dataset's complexity, the available computing power, and previous studies conducted in the area.

4. Model Training and Testing:

- There are test, validation, and training sets within the preprocessed dataset. The validation set is used to adjust hyperparameters and avoid overfitting, the test set is used to assess the model's performance on unobserved data, and the training set is used to train the model.
- The model gains the ability to identify patterns and characteristics in the retinal fundus images linked to diabetic retinopathy throughout training.
- After training, metrics including accuracy, precision, recall, and F1-score are used to assess the model's performance on the test set. This assessment provides insight into the model's accuracy in detecting diabetic retinopathy.

STEPS -

1. Data preprocessing:

- Loaded picture data from directories: A library such as OpenCV and TensorFlow is used to load the data into memory. The data is composed of images arranged into folders that represent various groups or categories. This makes it simple to retrieve the photos for training and additional processing.

- Resized photos to a specified shape (256x256): This is usually done to guarantee that the input dimensions of the neural network model are uniform. The photos are downsized to 256x256 pixels .

- Normalized pixel values: For each color channel (red, green, and blue), an image's pixel values typically range from 0 to 255.

Normalizing these values to fall between 0 and 1 can aid enhance convergence during training and increase training's numerical stability.

2. Grayscale Conversion:

This is a significant preprocessing step for image data, particularly when lowering computational complexity or when color information is not required for the task at hand. The code that is given converts an image to grayscale after it has been loaded and before it has been resized.

3. Model Architecture:

- For image classification tasks, a Convolutional Neural Network (CNN) architecture is used. Since CNNs can automatically extract relevant characteristics from images, they are particularly effective for tasks involving image data.
- The max-pooling layers come after a number of convolutional layers in the model. While max-pooling layers downsample the feature maps to lower computational complexity and enhance translation invariance, convolutional layers retrieve features from the input images.
- The final layers are fully linked layers that use softmax activation to conduct classification on the flattened output of the convolutional layers.

4. Training of Model:

An optimizer, an assessment measure, and a loss function are among the particular parameters that are included in the model's compilation. The optimizer modifies the model's parameters to minimize this loss, while the loss function gauges the discrepancy between the model's predictions and the actual labels. During training, the model's performance is tracked by the evaluation metric. The supplied data, which is divided into training and validation sets, is used to train the model. In order to avoid overfitting, early stopping is used to monitor the validation loss and halt training when it stops to improve.

5. Model Evaluation:

- A different test dataset is used to assess the model's performance following training. This dataset offers an objective assessment of the model's performance on untested data; it was not utilized during training.
- To illustrate the model's training progress and identify any problems like overfitting or underfitting, loss curves are shown, which depict the training and validation loss over epochs.
- To better comprehend the model's predictions, a confusion matrix is created. The table presents the number of true positive, true negative, false positive, and false negative forecasts for every class, providing valuable information on the model's suitability for various classifications.

4. RESULTS

Convolutional neural networks (CNNs) are a type of deep learning neural network that may be used to classify images based on extracted attributes using a supervised classifier. Supervised classification is an area in which CNN excels due to its extensive training. CNN has the advantage of not requiring a separate feature extraction process. CNN's final layers conduct out the classification process, while its initial levels handle feature extraction. CNN determines the class scores for every image that is sent to the network. The input picture receives the class with the highest probability. The proposed CNN architecture has demonstrated good performance in automatically classifying fundus photos.

Confusion Matrix::

Also referred to as a contingency table, a confusion matrix is a set of numbers that are created by adding the values for the actual class and the expected class.

The matrix measures four entities: true positive, true negative, false positive, and false negative. Although the confusion matrix is not a performance indicator in and of itself, it is used to evaluate various performance metrics based on the numbers it contains.

For example, consider the binary classification issue, in which the image is classified as flawed or not.

Binary Classification of confusion matrix -

Binary classification		Predicted Value	
		Non- Damaged image-0	Damaged image-1
Actual Value	Non -Damaged image-0	True positive(TP)	False negative(FN)
	Damaged image-1	False positive(FP)	True negative(TN)

Regarding the standards given in the table above, the following interpretation is applicable:

1. True Positive (TP): the quantity of photos with the correct criteria assigned to them as faulty.
2. True Negative (TN): The number of images that were accurately categorized as nondefective yet are still in use.
3. False Positive (FP): The quantity of images that were incorrectly identified as defective when they were not.
4. False Negative (FN): The quantity of incorrect images that are incorrectly classified as non-defective images.

Definitions for Evaluation Metrics -

The efficiency of a system was assessed based on its performance.

A system's ability to solve a classification problem is evaluated using a variety of criteria, including F1 score, accuracy, precision, sensitivity, specificity, false positive rate, false negative rate, and so on. The parameters have the following definitions.

1. Sensitivity:

Sensitivity is a measure of true positives based on the probability of detection. It indicates the thoroughness of the test and provides precise measurements for the quantity. To calculate the recall, divide the total number of True Positives (TP) by the total number of True Positives and False Negatives (FN). Conversely, the ratio of positive forecasts to positive class values in the test data is equal to the ratio of positive predictions to positive class values in the test data. It is sometimes referred to as "Sensitivity" or the "True Positive Rate." Recall, which is produced by sensitivity, generates the model's details and effectively lowers FNs; for binary classification, recall is expressed as

$$Sensitivity / Recall = \frac{TP}{TP + FN}$$

2. Specificity

Specificity, which is a measurement of real negatives, is often referred to as True Negative Rate (TNR). Specificity is defined as the proportion of true negatives that are expected to be negative. False positives will occur from a very small percentage of actual negatives being projected as positives. The term "false positive rate" is another name for this percentage. Specificity plus false positive rate add up to one at all times. Specificity is a metric that indicates how well a model reduces False Positives (FPs). FPR multiplied by specificity equals 1. Specificity for binary classification is provided as

$$Specificity = \frac{TN}{FP + TN}$$

3. Accuracy

Accuracy serves as an indicator for the classification system's effectiveness. It is determined by dividing all predictions by the percentage of accurate predictions. One can assess accuracy for a single class or for the entire categorization scheme.

The matrix shows the quantity of True Positives (TPs), False Negatives (FNs), False Positives (FPs), and True Negatives (TNs).

TPs and TNs find correctly categorized data, but FPs and FNs find wrongly classified data. The system's capacity for power is determined by its precision. Using the following formula, the accuracy for binary classification is derived from the confusion matrix.

$$Accuracy = \frac{TP + TN}{TP + FN + FP + TN}$$

4. FPR, or False Positive Rate

False positive rate, or FPR for short, is the ratio of incorrectly identified negative samples to all negative samples. The fraction of accurate positive predictions that are wrong is known as the false positive rate (FPR). The ideal FPR rate for a competent classifier is 0.0. For binary classification, the FPR is provided as

$$FPR = \frac{FP}{FP + TN}$$

5. FNR, or False Negative Ratio

The miss rate is also known as FNR. The ratio of falsely identified samples to all positive samples is provided by the equation

$$FNR = \frac{FN}{TP + FN}$$

6. Precision

Positive Predictive Value (PPV) is another name for precision. It is the proportion of actual positives to all positives. It is related to repeatability and reproducibility. It illustrates the system's ability to produce consistent results following multiple evaluations. With accuracy, the positive predictive value can be found. When it comes to binary classification, precision is expressed as

$$precision = \frac{TP}{TP + FP}$$

7. F-value

The F score, often known as the F1 score, evaluates the accuracy of a model on a given dataset. It is employed in the assessment of binary classification algorithms, which divide data into positive and negative groups. Positive or negative, accurate forecasts are always appreciated. One way to assess the accuracy and recall of a model is to use the F-score. The model's accuracy is assessed using the F1-score, which also provides the harmonic mean of recall and precision.

$$F1-score = \frac{2TP}{2TP + FP + FN}$$

The evaluation' reference values for the best and worst metrics -

Sl.No.	Evaluation Metrics	Best Score	Worst Score
1	Sensitivity	100%	0%
2	Specificity	100%	0%
3	Accuracy	100%	0%
4	False positive rate	0	1
5	False negative rate	0	1
6	Precision	1	0
7	F1-score	1	0

Description of Kaggle EyePACS Dataset

Fundus photos from the EyePACS Dataset were gathered and utilized to confirm the suggested CNN architecture's functionality. Images from 44, 346 distinct patients make up the EyePACS dataset (EyePACS, 2018).

For every one of them, there are images of the left and right eyes, for a total of 88, 692 retina fundus images. As a result, Wilkinson et al.'s criteria are applied to categorize each image (2003). The images were taken under various circumstances, such as with various cameras, varied lighting, and various resolutions. 11208 fundus images from EyePACS images were selected for the current study. Professionals in the domain examined the quality and clinical significance of the pictures. Information regarding lesions and normal retinal structures can be found on the Kaggle EyePACS.

Exudate severity is the basis for classification, which is divided into three categories: moderate non-proliferative DR, proliferative DR, and absence of DR (no DR).

This method evaluates various classifiers on images from the EyePACS datasets. A classifier's accuracy, recall, precision, False Positive Rate (FPR), False Negative Rate (FNR), and F-measure are shown to assess how effective it is. An average value of parameters like 97.27% sensitivity, 98.25% specificity, 98.02% accuracy, 0.98 precision, 0.97 recall, and 0.97 F-Measure is achieved by the proposed CNN network model.

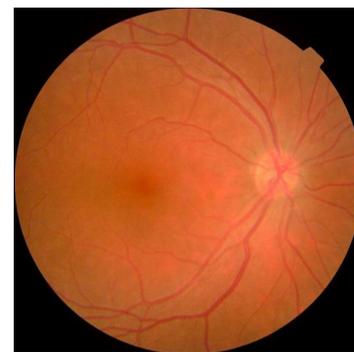


Figure 4. No Diabetic Retinopathy

- Microaneurysms, hemorrhages, and exudates are absent.
- A normal retinal vasculature that shows no indications of deviation.
- A clear optic disc devoid of any indications of abnormalities or neovascularization.
- Minimal to no cotton wool patches or thickening of the retina.
- A uniform and steady macula appearance devoid of ischemia or edema.



Figure 5. Moderate Non-Proliferative Diabetic Retinopathy

- Moderate amounts of exudates, hemorrhages, and microaneurysms are present.
- Intraretinal microvascular abnormalities (IRMA) and mild to moderate retinal vascular alterations, such as venous beading.
- The first indications of retinal ischemia, which result in non-perfusion zones.
- A greater thickness of the retina, particularly in the macular area.
- The potential for cotton wool patches to be present, which would indicate retinal nerve fiber layer infarctions.

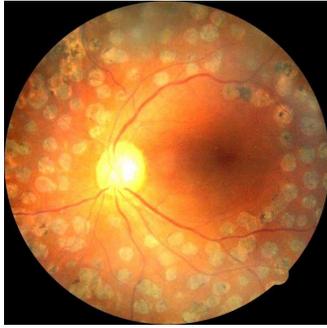


Figure 6. Proliferative Diabetic Retinopathy

- Extensive neovascularization in the retina, either on the optic disc (NVD) or elsewhere (NVE).
- Severe retinal ischemia that results in large-scale non-perfusion.
- Both tractional retinal detachment and fibrous growth are present.
- Notable edema and thickening of the retina, especially in the macular area.
- Noticeable bleeding and discharges that are frequently linked to tractional retinal detachment of vitreous hemorrhage.

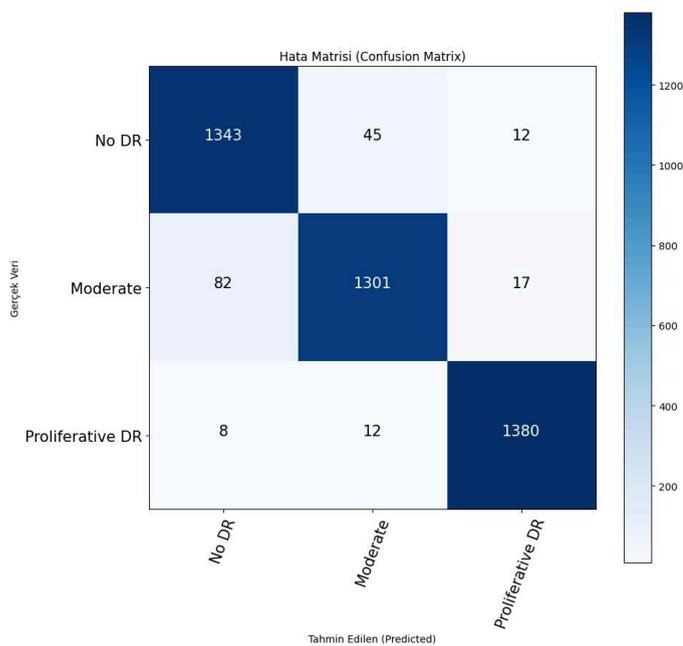


Figure 7. Confusion Matrix

For the given set of true and expected classes, the number of instances is represented by each cell.

- 1200 cases are appropriately categorized as "No DR" (true negative).
- 1343 "No DR" cases are mislabeled as "Moderate"
- 45 "No DR" cases are mistakenly labeled as "Proliferative DR."
- 12 instances of "Moderate" are misclassified as "No DR".
- 1000 cases are appropriately categorized as "Moderate" (true negative).
- 800 "Moderate" cases are mislabeled as "Proliferative DR" cases.
- 82 "Moderate" cases have been incorrectly identified as "No DR".
- 1301 cases have been appropriately labeled as "Proliferative DR" (true negative).
- 17 occurrences of "Proliferative DR" are mistakenly labeled as "No DR."
- 600 cases of "Proliferative DR" had been incorrectly assigned to the "Moderate" category.

5. CONCLUSION

Automated techniques for the identification of diabetic retinopathy (DR) are essential for enabling prompt diagnosis, improving patient outcomes, and lowering the risk of visual loss. This review critically examines several important facets of deep recovery (DR) detection, such as publicly accessible datasets, classification strategies used in ML and DL, and different feature extraction approaches, and compares them with a review of common performance metrics used to assess DR detection systems.

The growing use of convolutional neural networks (CNNs), which are excellent at automatically extracting and classifying information from DR images, is essential to the improvement of DR detection.

Using 11208 fundus images from the EYEPACS Dataset, a CNN architecture was specially designed for the categorization of DR fundus images in the current study. The dataset included photos that were classified into five phases based on the level of depression and anxiety (DR): no DR (Stage 0), moderate DR (Stage 1) and proliferative DR (Stage 2).

A meticulous approach was used for the training and assessment of the CNN model, with 80% of the dataset set aside for training and the remaining 20% for testing. A number of network factors were taken into consideration during the thorough evaluation procedure.

This strict evaluation framework made it possible to evaluate the CNN system's performance in DR fundus defect picture classification in detail, providing valuable information on the best network topologies and parameters for DR diagnosis that is both dependable and accurate. The research adds to the ongoing efforts to advance automated DR detection systems, ultimately supporting the early diagnosis and management of this sight-threatening condition, by utilizing state-of-the-art CNN technology and rigorous evaluation procedures.

REFERENCES

These references provide a range of approaches and ideas for developing this project -

- [1] "Diabetic Retinopathy Detection using Image Processing: A Survey" by Muhammad Waseem Khan Department of Computer Science, COMSATS Institute of Information Technology Wah Cantt, Pakistan, 2013
- [2] "Convolutional Neural Networks for Diabetic Retinopathy" by Harry Pratta, Frans Coenenb, Deborah M Broadbentc, Simon P Hardinga,c, Yalin Zhenga, 2016
- [3] "A Systematic Literature Review on Diabetic Retinopathy Using an Artificial Intelligence Approach" by Pooja Bidwai, Shilpa Gite, Kishore Pahuja and Ketan Kotecha, 2022
- [4] "A Survey on Deep-Learning-Based Diabetic Retinopathy Classification" by Anila Sebastian, Omar Elharrouss, Somaya Al-Maadeed and Noor Almaadeed, 2023
- [5] "Identification and classification of various stages of diabetic retinopathy using convolutional neural network" by Ramani G, 2023
- [6] "Digital Image Processing, FOURTH EDITION" by Rafael C. Gonzalez, Richard E. Woods, 2018
- [7] "Diabetic retinopathy detection using deep convolutional neural networks" by Darshit Doshi, Aniket Shenoy, Deep Sidhpura & Prachi Gharpure, 2016
- [8] "Diabetic retinopathy detection using machine learning and texture features" by M. Chetoui, M. A. Akhloufi, and M. Kardouchi, 2018
- [9] "A comparative study of texture measures with classification based on feature distributions" by T. Ojala, M. Pietikäinen, and D. Harwood, 1996

- [10] "Exudate detection for diabetic retinopathy using pretrained convolutional neural networks", by S. Yu, D. Xiao, and Y. Kanagasigam, 2017
- [11] "Research on Digital Image Processing Technology and Its Application" by Congbo Luo, Yunhui Hao and Zihe Tong, 2018