

A Novel Deep Learning Based Architecture for Measuring Diabetes

Shaima Sharaf,

shaimasharaf75@gmail.com

Department of Computer Science, TKM Institute of Technology, Kerala, India

Abstract

Diabetes is a chronic condition that happens when the pancreas fails to produce enough insulin or when the body's insulin is ineffectively used. Uncontrolled diabetes causes hyperglycaemia, or high blood sugar, which causes catastrophic damage to many of the body's systems, including the neurons and blood vessels, over time. The burden of disease on the global healthcare system is enormous. As a result, early diabetes diagnosis is critical in saving many lives. Current methods for determining whether a person has diabetes or is at risk of acquiring diabetes, on the other hand, rely heavily on clinical biomarkers. This research presents a unique deep learning architecture for predicting whether or not a person has diabetes and the severity levels of diabetes from the person's retinal image. This study incorporates datasets such as EyePACS and IDRID, which comprise Diabetic Retinopathy (DR) images and uses Dense-121 as the base due to its improved performance.

Keywords:

Convolutional neural network, deep learning, diabetes, machine learning, diabetic retinopathy

1. Introduction

A collection of metabolic illnesses known as diabetes mellitus, or simply diabetes, is characterised by hyperglycemia that develops when there is insufficient insulin in the body. The heart, kidneys, and microvascular circulation of the retina all suffer long-term harm and collapse as a result of diabetes. diabetes retinopathy (DR), a frequent problem for diabetes patients that can result in blindness, is at risk for developing in a significant portion of diabetic patients. The vascular system of the retina may be negatively affected by diabetes, leading to structural alterations. Because changes in vascular structure in the retina can provide visual indicators for diabetes, the majority of clinical guidelines recommend a yearly retinal screen for diabetic patients utilising retinal fundus imaging or dilated eye examinations. Although this would need to be done subjectively by an ophthalmologist and would probably take some time, these retinal images could also be used to detect diabetes. Automating retinal image-based diabetes diagnosis in clinical settings would allow us to do away with subjective human evaluation. The burden on ophthalmologists may be lessened, and a large number of patients could undergo objective screening quickly. Few studies have examined the task of detecting diabetes from a

comprehensive standpoint, despite the fact that many studies have focused on using retinal images to diagnose diabetic retinopathy. In fact, a large number of research have focused on the diagnosis of diabetes using clinical markers such as HbA1c and glucose. As a result, there is a growing demand for novel, cost-effective, non-invasive, and quick diabetes screening technologies that are simple to apply.

Deep learning-based methods have been used in numerous research to grade DR (e.g., "mild," "moderate," "severe," and "proliferate") and detect it. While ischemic optic neuropathy, cataract, glaucoma, macular degeneration, and other retinal and non-retinal disorders can all result in vision loss, diabetic retinopathy is the most prevalent cause of vision loss. As a result, integrating retinal images in diabetes diagnosis will give diabetic patients' treatment regimens more information.

In Qatar, there is only one study on the comprehensive diagnosis of diabetes based on retinal imaging. However, it is less precise than previous DR detection investigations and only captures a binary categorization issue, whether the subject has diabetes or not. I describe a novel deep learning-based architecture for assessing diabetes in order to close this gap. This design addresses the challenge of determining a person's diabetes status and the severity of the illness from just a single retinal image. There are other research that concentrate on detecting DR grade, but this one is distinct in that it aims to use retinal pictures to predict diabetes and its stages whether or not DR is present.

The problem is formulated as a supervised learning task, specifically a classification problem, in this paper. This paper's contributions can be summarised as follows:

- The proposed novel deep learning model predicts whether a person has diabetes or not and its severity levels from an image of his/her retina and introduced a multi-stage CNN based model for the purpose.
- The layers have been increased in the dense block in order to achieve better performance.

2. LITERATURE SURVEY

Tymchenko et al. [1] developed a multistage technique to transfer learning and an automatic deep learning-based system for stage diagnosis of diabetic retinopathy using a single fundus image of the retina. They employed numerous datasets for their research, including EYEPACS, IDRiD, MES-SIDOR, and others, and the evaluation was done in kaggle APTOS 2019 with Cohen's kappa score as the evaluation metric. They used an ensemble of three CNN architectures in their investigation, including EfficientNet-B4, EfficientNet-B5, and SEResNeXt 50, and used transfer learning for the final answer. For this model, a multistage training strategy was adopted. Even with an unstable measure, the experimental findings show that this produces high and stable outcomes. The key benefit of their model was that it used an ensemble of networks that were pretrained on a broad dataset and fine-tuned on a target dataset to boost generalisation and reduce variance. Based on the results, the model could achieve a sensitivity and specificity of 0.99 and was ranked 54 of 2943 competing methods.

Bora et al. [2] used three-field or one-field colour fundus pictures as input to develop two versions of a deep learning system to predict the probability of DR within two years. They primarily employed two datasets in their research: the first was EyePACS (which comprised 362283 people), and the second was from 13 medical centres across Thailand (included 6791 patients). They performed a survival analysis with Kaplan-Meier curves, log-rank tests, and Cox proportional hazards regression models to evaluate the system in terms of outcomes that happened after two years. The following were the main benefits of their model: To begin, they might stratify the largest group of patients to solve the problem of optimising four screening intervals. Second, when controlling for available risk factors, the strong prognostic value is maintained. Finally, the system was validated using two validation datasets from two countries on two distinct continents. However, their study contains flaws, such as the lack of some risk indicators in data sets, the lack of device information on a per-image basis, and other issues that were crucial to increasing the system's performance.

Qiao et al. [3] created a system that analyses the existence of MAs in fundus images utilising CNN with deep learning as the core component and GPU acceleration for

medical picture identification and segmentation. The fundus images were classified using a semantic segmentation technique in this investigation. Convolutional layers were used to analyse the input images, and each unit of a convolutional layer in an input image binds to a small area termed the receptive field, which extends across the whole image range. A convolutional layer was used to insert the units in feature maps. Through the use of more convolutional layers, the network was able to learn deeper features. ReLU activated the convolution neurons, and the final layer connects to a Softmax feature that alters the performance in different categories of input picture probabilities. To improve the images in the dataset, preprocessing methods such as curve transform, morphological closure, and others were used. This work created a sparse Principal Component Analysis-based unregulated classification strategy for detecting microaneurysms.

Qummar et al. [4] used a publicly accessible retina image dataset to train an ensemble of five deep CNN models, including Resnet50, Inceptionv3, Xception, Dense121, and Dense169, to encode the rich information and enhance categorization for various phases of DR. They used the Kaggle dataset for their research and applied preprocessing techniques including scaling, up sampling, down sampling, flipping, and so on. They used stacking to improve the prediction of the suggested model because it outperforms individual models. Stacking shows where each base model excels and discredits where it falls short. As a result, stacking is most effective when the base models are significantly diverse, as evidenced by their model's results. Despite the model's success in detecting all stages of DR, it had flaws such as model complexity, transparency loss, and the fact that it didn't produce findings for a single model.

To speed up training and model convergence, Khan et al. [5] concentrates on classifying DRs into distinct stages with the fewest learnable parameters possible. VGG-16, Spatial Pyramid Pooling Layer (SPP), and Network in Network (NiN) are stacked in this study to create the VGG-NiN model, a highly non-linear scale invariant deep model. Due to the virtues of SPP layers, it was able to process a DR picture at any size, and the stacking of NiN added more non linearity to the model, resulting in better categorization. This study used only the labelled images from the Kaggle EyePacs dataset. The Stochastic Gradient Decent (SGD) optimizer was used to carry out the training. The proposed

model produced superior computation results and reduced the number of learnable parameters, which is normally a disadvantage of ensemble models.

Bilal et al. [6] proposed an innovative and hybrid approach for prior DR detection and classification, in which the final result is determined by a majority voting system. Different machine learning models, such as SVM, KNN, and Binary Tree, were utilised as classifiers in their research. This work was evaluated using the IDRID dataset, which contains 516 photos. Adaptive histogram equalisation and contrast stretching were used for pre-processing, followed by a median filtering approach. The researchers employed two separate techniques (bright and red) to detect the lesions, with the retrieved features being concatenated into a feature vector. The classification was done using an increasing severity threshold, and the voting system's total performance was recorded. The study's sensitivity was lower than some of the other research, despite the fact that it produced superior results than similar investigations.

Islam et al.'s [7] innovative deep learning architecture proposed a method for determining whether or not a person has diabetes from a picture of their retina. Using a relatively small dataset, they successfully identified the regions on the retina images that contribute to its decision-making process and created a multi-stage convolutional neural network (CNN)-based model DiaNet that can accomplish an accuracy level of over 84 percent on this task, as confirmed by medical experts in the field. They employed the EyePACS dataset and the QBB retina dataset for their research, which included retinal scans from 492 control and diabetic patients in Qatar.

The same pre-processing was applied to both datasets. They have carried out a large number of trials to show that the suggested method can detect diabetes from retinal pictures with a reasonable high accuracy with a little dataset. This model was able to identify between diabetes and nondiabetic patients from retinal images. In their proposed model, they used DenseNet-121 as the base CNN. They compared DiaNet's performance to that of existing clinical data-based machine learning algorithms and came to the conclusion that the retinal images contain sufficient information to distinguish the Qatari diabetic cohort from the control group. However, compared to earlier research that concentrated on DR detection, this study was less precise.

3. MATERIALS AND METHODS

This section discusses the process involved in the collection, curation, and pre-processing of the dataset as it pertains to this work and the development of the proposed solution for the problem at hand. To attain state-of-the-art in detecting diabetes from retinal pictures, the proposed method employs two datasets.

1. THE EyePACS DR DATASET

The EyePACS dataset, which contains over 80,000 retinal images acquired from various cameras, is the largest dataset of diabetic retinopathy fundus images. This dataset includes a large number of high-resolution retinal fundus images collected under various imaging settings [8],[9].




Figure 1. Few examples from the EyePACS dataset

It comprises of labels showing one of the five possible gradings of diabetic retinopathy based on the severity of diabetic retinopathy, such "none," "mild," "moderate," "severe," and "proliferative." Figure 1 depicts some of the dataset's samples in each category.

2. INDIAN DIABETIC RETINOPATHY IMAGE DATASET

The IDRiD dataset was developed using real clinical examinations at the Nanded (M.S.) eye clinic in India and is open to the public. Figure 2 depicts some of the dataset's samples in each category. A Kowa VX-10 fundus camera was used to capture these images, which were focused on the macula. Before imaging, both subjects' pupils were

dilated with tropicamide at a concentration of 0.5 percent. The final dataset is a collection of 516 images in jpg format with 50 fields of view and a resolution of 4288*2848 pixels. It is the only dataset that displays typical diabetic retinopathy lesions as well as normal retinal structures down to the pixel level. These results are useful for developing and testing DR image analysis methods for prior detection.




Figure 2. Few examples from the IDRiD dataset

The dataset is made up of five DR (0-4) and three DME (0-3) groups with well-defined characteristics that adhere to international clinical relevance standards. A grading score ranges from 0 to 4, representing the severity of the disease. With a grade 0 patient, there is no DR.

3. IMAGE PREPROCESSING

Multiple pre-processing approaches are used on the data set in this study. In the pre-processing stage, the

following steps were carried out. First, each image's circular section is extracted and resized so that the retina's radius is 300 pixels.

Finally, the photos are oriented such that the backgrounds are similar. Techniques such as random brightness and contrast perturbation, as well as horizontal flipping, were used to enhance the data.




Figure 3. Images before(left) and after(right) pre-processing in IDRiD dataset

Fig.3 shows the image after preprocessing. Each image in the dataset was changed from a varied sized image to a uniform

size image with a black background using this sequence of pre-processing processes, resulting in all pre-processed

images being aligned and having a similar background. Both the EyePACS and IDRID datasets went through the same pre-processing and augmentation methods.

4. SMOTE

SMOTE (synthetic minority oversampling technique) is one of the most commonly used oversampling methods to solve the imbalance problem. Its goal is to achieve a more balanced distribution of classes by replicating minority class examples at random.

SMOTE is a statistical method for evenly raising the number of observations in the data set. It operates by creating new instances from existing minority cases, which must be provided as input. The majority of cases do not change as a result of this implementation. Because the approach gathers examples of all the features for each target class and its nearest neighbours, the newly created instances are not simply duplicates of the current minority class. This method expands the number of features available to each class and broadens the scope of the examples.

SMOTE then takes the dataset as an input and only increases the percentage of the minority class in the data. In this study, SMOTE has been applied to resolve the imbalanced problem in the dataset. Fig 4 shows the plots of EyePacs dataset before and after the application of SMOTE.




Figure 4. Images before(above) and after(below) the application of SMOTE in EyePacs dataset

5. DEVELOPMENT OF PROPOSED MODEL

5.1 DENSENET

DenseNet (Dense Convolutional Network) is a network architecture that focuses on deepening deep learning networks while also making them more effective to train by employing shorter connections between layers [11]. DenseNet is a convolutional neural network in which each layer is connected to all other layers deeper in the network; for example, the first layer is connected to the second, third, fourth, and so on, while the second layer is connected to the third, fourth, fifth, and so on. This is done to ensure that the maximum amount of data can move between the network's layers. To maintain the feed-forward character of the system, each layer receives input from all previous layers and passes on its own feature maps to all subsequent layers. It does not integrate characteristics by summation, like Resnets do, but rather by concatenating them. As a result, the 'ith' layer contains 'i' inputs and is made up of feature maps from all of the convolutional blocks before it. All of the subsequent 'I-I' layers receive their own feature maps. Instead of only 'I' connections, like in classic deep learning designs, the network now has $(I(I+1))/2$ connections. As a result, it requires fewer parameters than standard convolutional neural networks because no insignificant feature maps should be learned.

Aside from the basic convolutional and pooling layers, DenseNet has two important blocks. They are the Transition Layers and the Dense Blocks. DenseNet was created specifically to improve the vanishing gradient-induced decline in accuracy in high-level neural networks. In simpler terms, the information vanishes before reaching its destination due to the longer path between the input layer and the output layer. DenseNet begins with a basic layer of convolution and pooling. Then there's a dense block followed by a transition layer, another dense block followed by a transition layer, another dense block followed by a transition layer, and lastly a dense block followed by a classification layer.

5.2 PROPOSED APPROACH

To develop the proposed model, begin with a CNN model M1 that has been pre-trained on the ImageNet dataset [10],[12]. Then, before the final layer, add a few more layers to the network to improve its ability to recognize more complex features from the input. We first fine-tune this augmented network on the EyePacs dataset in order to equip it with retinal image interpretation skills. We have a model M2 at the end of this fine-tuning stage. The model M2 is then used to fine-tune it using the IDRID dataset. As a result, we have M3, a model that can assess the severity of diabetics using retinal pictures. DenseNet121 is chosen

as the basis CNN in the suggested model. Fig.5 shows the high-level view of the proposed model.

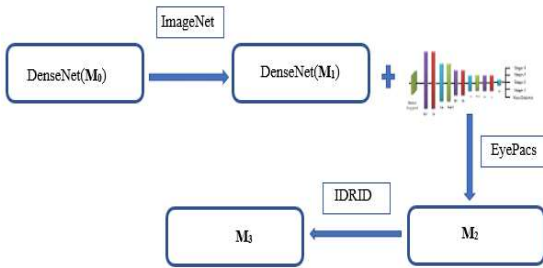


Figure 5. A high-level view of the proposed model

5.3 CNN ARCHITECTURE

In this study, due to its superior performance, diabetic detection is implemented using DenseNet-121, which consists of four dense blocks. Despite the network being

121 layers deep, overfitting was not an issue due to the facts such that the dense connections acting as regularizers and the network having less parameters than a non-dense CNN. DenseNet121 is employed as the backbone in the suggested model, which has 121 layers and eliminates overfitting. By altering the number of neurons in the output layer, the suggested model makes a multiclass prediction. This will reveal the diabetes' severity. i.e., non-diabetic, stage 1, stage 2, stage 3, and stage 4 diabetes. Variants of Dense can also be created by altering parameters such as the number of layers, padding value, and stride. By varying the number of dense layers and activation function, the suggested method creates a new variation that gives optimal performance while using fewer computational resources.

The proposed architecture is depicted in Fig.6. It is made up of Densenet-121 and a few additional layers, including a pooling layer and a sequence of batch normalisation (BN), dropout (Dr), linear (Lin), and ReLU activation layers.




Figure 6. The architecture of the proposed model

The pooling layers in the CNN are used to summarize the features from the feature map. Global Average Pooling is a function that computes the average output of each feature map in the preceding layer. This relatively simple operation significantly reduces the data and prepares the model for the final classification layer. It is designed to replace fully connected layers in classical CNNs. It takes the average of each feature map and feeds the resulting vector directly into the next layer, rather than building completely connected layers on top of the feature maps. One advantage of global average pooling over fully connected layers is that it enforces correspondences between feature maps and categories, making it more native to the convolution structure. As a result, the feature maps can be simply interpreted as categories confidence maps. Another benefit of global average pooling is that there are no parameters to tune, therefore overfitting is prevented at this layer.

Batch normalisation is a technique for making deep neural networks faster and more stable by adding more layers to the network. The standardising and normalising

procedures are performed by the new layer on the input of a previous layer.

Dropout is a strategy in which neurons are chosen at random to be disregarded during training, or "dropped-out." This means that on the forward pass, their contribution to downstream neuron activation is momentarily removed, and any weight modifications are not applied to the neuron on the backward pass.

In deep learning models, the Rectified Linear Unit is the most widely employed activation function. If the function receives any negative input, it returns 0; however, if the function receives any positive value x , it returns that value. As a result, it can be written as

$$f(x) = \max(0, x) \tag{1}$$

The ReLU function is straightforward and does not require any sophisticated calculations. As a result, the model can be trained or run in less time, and another

essential attribute of the ReLU activation function is sparsity.

5.4 HARDWARE AND SOFTWARE SETUP

The proposed work was conducted in the following configuration: Core i5 based CPU with 8GB RAM and Python 3.9 used as the base language. The Experiment was carried out on Jupiter notebook. Pandas, NumPy, Keras, TensorFlow, Scikit-learn, Matplotlib packages will also be used in the experiment for data pre-processing, manipulation, etc.

6. EXPERIMENTAL RESULTS

The results of the experiments are presented in this part, along with an analysis of the findings.

6.1 EVALUATION METRICS

Accuracy, precision, recall, sensitivity, specificity and also the F1 score were employed for quantitative performance reporting.

$$\text{Accuracy} = \frac{TP+TN}{TP+TN+FP+FN} \quad (2)$$

$$\text{Precision} = \frac{TP}{TP+FP} \quad (3)$$

$$\text{Recall} = \frac{TP}{TP+FN} \quad (4)$$

$$\text{F1-Score} = \frac{2 * \text{Precision} * \text{Recall}}{\text{Precision} + \text{Recall}} \quad (5)$$

$$\text{Sensitivity} = \frac{TP}{TP+FN} \quad (6)$$

$$\text{Specificity} = \frac{TN}{TN+FN} \quad (7)$$

where TP is the number of true cases, TN is the number of true negative cases, FP is the number of false-positive case and FN is the number of false-negative cases.

6.2 MODEL TRAINING RESULTS

Densenet121 is used as the base model M0 in this study. M1, the CNN model pre-trained on the ImageNet dataset. Then, before the final layer, add a few more layers to the network to improve its ability to recognise more complex features from the input. We first fine-tune this augmented network on the EyePacs dataset to give it retinal image interpretation abilities. At the end of this stage of fine-tuning, we have a model M2. The IDRID dataset is then employed to fine-tune the model M2. As a result, we now have M3, a model that uses retinal images to measure the severity of diabetics. The weights of a model are saved once it has been built. The weights that have been saved are

subsequently loaded into the following model. Each model was applied with 10 epochs with a batch size of 20. The accuracy of both models increased with each epoch, while the loss decreased.

6.2 MODEL TRAINED ON EYEPACS DATASET(M2)

When given a test image of the retina, the model classifies the image into any one of the diabetes severity levels. The model trained on EyePacs dataset yielded good performance. With each epoch, the accuracy of the model trained on the EyePacs dataset improved. After two epochs, an accuracy of 0.7932 and a loss of 0.2440 were recorded which was 0.7245 ,0.4336 respectively in the previous epoch. Finally, the model achieved an average accuracy of 84.5. It is visible that with each epoch the model accuracy kept increasing. Meanwhile, the model loss kept decreasing with each epoch. The blue line in the figure shows the training part and the orange line indicates the validation part. The model was run multiple time for estimating the average accuracy and it was able to achieve an accuracy of 84.5 percentage.

6.3 MODEL TRAINED ON EYEPACS AND IDRID DATASET(M3)

This model was first trained on the EyePacs dataset, after which the model weights were saved and loaded into the next model. That is, the model was trained on the EyePACS and IDRID datasets, and this model (M3) outperformed the previous model.

The model was also given 10 epochs and a batch size of 10. After each epoch, the accuracy of the model trained on the IDRID dataset improved. The first epoch yielded an accuracy of 0.6074 and a loss of 0.4921. M3's accuracy has improved and loss has decreased, similar to the model trained on EyePacs. Finally, the model achieved an average accuracy of 92 percentage. After each epoch the model accuracy kept increasing. Meanwhile, the model loss kept decreasing with each epoch. The blue line in the figure shows the training part and the orange line indicates the validation part.

When compared the model M2 which was only trained on the EyePacs dataset, the model M3, that is trained on both EyePacs and IDRID dataset has shown better results than the prior one. The evaluation metrics such as precision, recall, F1-score, etc can be calculated from the confusion matrix obtained based on the values true positive, true negative, false positive and false negative. Based on the results, it is obvious that the model trained on both EyePacs and IDRID dataset performs better than the model that is trained only on EyePacs dataset.

From retinal pictures, the suggested deep learning-based approach is used to estimate the existence of diabetes in a test subject. Here, employed a CNN-based design that uses

a retinal image as an input. Because of the small size of the IDRID dataset, a fine-tuning strategy is used in order to obtain good performance. As a result, incorporating a larger dataset of retinal images improved the model's understanding of retinal images. When images from the EyePACS dataset were used, the proposed method produced better outcomes.

7 PRACTICAL APPLICATION OF THE PROPOSED MODEL

The World Health Organization (WHO) and the International Diabetes Federation (IDF) have recently highlighted low-cost digital retinal photography screening by non-physicians and remote grading using mobile healthcare services. If applied in a clinical setting, the suggested technology will change the eye care system in low- and middle-income nations due to its cost-effectiveness. For high-income countries, this technology will reduce physician workload while also assisting in the implementation of mass diabetes screening in a shorter period of time.

8 FUTURE ENHANCEMENT

The current study concentrated on training the Dense model with two datasets, adding weights in each phase. Though the system's performance measures are satisfactory, it can be improved further by including LSTM functions in the training phase. When the models are run in a task parallel mode system, it can help to maintain memory cells for a long time. Aside from detecting and grading diabetes, retina images can be used to identify other diseases such as ischemia, ischemic heart disease, hypertension, glaucoma, age-related macular degeneration, retinal detachment, cancer, and so on. The datasets on such diagnoses can be safely queued up in the current sequential model for better disease identification. If more pre-processing functions are added, it can also be fed into a live monitoring system.

9. CONCLUSION

Diabetes is a group of metabolic disorders caused by an increase in blood sugar levels, and Diabetic Retinopathy is a disease that can develop in people who have diabetes (DR). This proposed model uses simply the patient's retinal picture to determine whether or not a person has diabetes and the severity of the disease. The model trained on the EyePACS and IDRID datasets has an accuracy of 92 percent and also provides a specificity of 98.6 percent. Since the results are satisfactory, the model can be used in a clinical setting to determine whether or not a patient has diabetes. This can reduce the manual effort required to diagnose diabetes patients using retinal fundus images.

REFERENCES

- [1] B.Tymchenko, P. Marchenko, and D. Spodarets, "Deep learning approach to diabetic retinopathy detection," in Proc. 9th Int. Conf. Pattern Recognit. Appl. Methods, 2020, pp. 501-509.
- [2] A. Bora, S. Balasubramanian, B. Babenko, S. Virmani, S. Venugopalan, A. Mitani, G. de Oliveira Marinho, J. Cuadros, P. Ruamviboonsuk, G. S. Corrado, L. Peng, D. R. Webster, A. V. Varadarajan, N. Hammel, Y. Liu, and P. Bavishi, "Predicting the risk of developing diabetic retinopathy using deep learning," *Lancet Digit. Health*, vol. 3, no. 1, pp. e10-e19, Jan. 2021.
- [3] L. Qiao, Y. Zhu and H. Zhou, "Diabetic Retinopathy Detection Using Prognosis of Microaneurysm and Early Diagnosis System for Non-Proliferative Diabetic Retinopathy Based on Deep Learning Algorithms," in *IEEE Access*, vol. 8, pp. 104292-104302, 2020, doi: 10.1109/ACCESS.2020.2993937.
- [4] S. Qummar et al., "A Deep Learning Ensemble Approach for Diabetic Retinopathy Detection," in *IEEE Access*, vol. 7, pp. 150530-150539, 2019, doi: 10.1109/ACCESS.2019.2947484.
- [5] Z. Khan et al., "Diabetic Retinopathy Detection Using VGG-NIN a Deep Learning Architecture," in *IEEE Access*, vol. 9, pp. 61408-61416, 2021, doi: 10.1109/ACCESS.2021.3074422.
- [6] A. Bilal, G. Sun, Y. Li, S. Mazhar and A. Q. Khan, "Diabetic Retinopathy Detection and Classification Using Mixed Models for a Disease Grading Database," in *IEEE Access*, vol. 9, pp. 23544-23553, 2021, doi: 10.1109/ACCESS.2021.3056186
- [7] M. T. Islam, H. R. H. Al-Absi, E. A. Ruagh and T. Alam, "DiaNet: A Deep Learning Based Architecture to Diagnose Diabetes Using Retinal Images Only," in *IEEE Access*, vol. 9, pp. 15686-15695, 2021, doi: 10.1109/ACCESS.2021.3052477.
- [8] J. Cuadros and G. Bresnick, "EyePACS: An adaptable telemedicine system for diabetic retinopathy screening," *J. Diabetes Sci. Technol.*, vol. 3, no. 3, pp. 509-516, May 2009.
- [9] B. Graham, "Kaggle diabetic retinopathy detection competition report," Univ. Warwick, Coventry, U.K., Tech. Rep., 2015.
- [10] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., Jun. 2009, pp. 248-255.
- [11] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), Jul. 2017, pp. 4700-4708
- [12] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in Proc. Adv. Neural Inf. Pro- cess. Syst. (NIPS), Stateline, NV, USA, vol. 25, Dec. 2012, pp. 1097-1105