Diabetic Retinopathy Detection Using Transfer Learning and Deep Convolutional Neural Networks

Prof. Chandrawati, Rohit Arora, Samyak Jain, Hitendra Singh
Department of Information Technology, Netaji Subhas University of Technology, India

ABSTRACT - The early detection and diagnosis of Diabetic Retinopathy (DR) can help prevent blindness in diabetic patients. In recent years, deep learning methods, specifically Convolutional Neural Networks (CNNs), have shown remarkable performance in detecting DR. In this study, we proposed the use of Densenet-201, a popular CNN architecture having 201 layers for diabetic retinopathy detection. We fine-tuned the pre-trained Densenet model on the dataset collected by Basys.ai, a healthcare company working towards the diagnosis of Diabetic retinopathy. Data augmentation techniques were used to make the data balanced since the data provided has imbalanced. Our results demonstrate that the Densenet model achieved high accuracy as compared to other state-of-the-art transfer learning models. Our model performed better than other published results using Densenet. We achieved a test accuracy of 92.50% and recall of 91.07% with precision 87.93%. The proposed approach holds great potential in assisting ophthalmologists in diagnosing DR and providing early interventions to patients.

I. INTRODUCTION

1.1 THEORY

Diabetic retinopathy is an eye condition caused by diabetes. At the back of the eye, the retina exists which gets damaged by high blood sugar levels.

Diabetic retinopathy is a significant health issue affecting a large number of people worldwide. According to the International Diabetes Federation, Around 460 million persons were diagnosed with diabetes in 2019, with the number is expected to climb to 600 million by 2040s. Diabetic retinopathy affects around one-third of all diabetics.

The impact of diabetic retinopathy is not limited to developed countries. In fact, diabetic retinopathy is a significant cause of blindness in many developing countries, particularly in regions with high rates of diabetes and limited access to healthcare. A study published in the Lancet Global Health in 2019 estimated that diabetic retinopathy affects approximately 21 million people in sub-Saharan Africa alone.

These statistics highlight the urgent need for effective screening, early detection, and treatment of diabetic retinopathy on a global help to reduce the burden of diabetic retinopathy on individuals and society as a whole.

Blurred vision, floaters, colour vision loss, and vision that changes from fuzzy to clear during the day are all indications of diabetic retinopathy.

A full eye exam, which may include a dilated eye exam, visual acuity test, and optical coherence tomography (OCT), can be used to evaluate diabetic retinopathy.

The stage of the illness determines the treatment for diabetic retinopathy. Controlling blood sugar, blood pressure, and cholesterol levels may be necessary in the early phases of therapy. Later stages may require laser therapy or surgery to shut up leaky blood vessels and prevent new vessel formation.

Diabetes retinopathy prevention includes controlling blood sugar, blood pressure, and cholesterol levels, as well as living a healthy lifestyle that includes frequent exercise, eating a balanced diet, and avoiding scale. Increased awareness of the disease and its impact, combined with advances in technology and medical treatments, may smoking. Regular eye exams are also necessary for early detection and treatment of diabetic retinopathy.

1.2 MOTIVATION

The Western Pacific area, which includes nations such as China, Japan, and Australia, has the greatest incidence of diabetes, with an estimated 161 million individuals suffering from the disease. The region with the second-highest prevalence is South-East Asia, which includes countries such as India, Indonesia, and Bangladesh, with an estimated 88 million adults with diabetes.
The count of people suffering with diabetes is projected to increase to 700 million by 2045, with the greatest increase expected in low- and middle-

income countries. This increase is largely attributed to changes in lifestyle and demographics, such as sedentary lifestyles, unhealthy diets, and an aging population.

Diabetic retinopathy is the main cause of blindness among working-age.individuals in developed nations. Diabetic retinopathy is estimated to account for around 12% of new occurrences of blindness in the United States each year.

Diabetic retinopathy is diagnosed with a thorough eye examination that includes dilated fundus examination, visual acuity testing, imaging tests such as optical coherence tomography (OCT) and fundus photography. Early identification is critical in preventing vision loss, and diabetics should have frequent eye exams beginning with their diagnosis.
Early identification and treatment of diabetic retinopathy is critical in preventing vision loss, and persons with diabetes should have frequent eye exams. Preventing diabetic retinopathy involves controlling blood sugar levels and maintaining a healthy lifestyle.

Our grandparents had diabetes, but they were fortunate enough to receive treatment in time. Unfortunately, not everyone is that lucky. This is why we want to take action that will advance the field of accessible, low-cost diabetic retinopathy diagnosis.

1.3 PROBLEM STATEMENT

Early.diagnosis.and.treatment.of.diabetic.reti nopathy.are crucial because diabetic retinopathy can progress without noticeable symptoms until it reaches an advanced stage, which may be too late to treat effectively. In its early stages, diabetic retinopathy may cause minor symptoms such as blurry or distorted vision, but in advanced stages, it can cause severe vision loss or even blindness.

Overall, early detection and treatment of diabetic retinopathy are crucial for maintaining eyesight and increasing persons with diabetes' quality of life.

- In this Project we are aiming to detect Diabetic.Retinopathy using transfer.learning.and Deep.Neural.Networks
- We are using DenseNet-CNN to increase the accuracy of the project
- We use data augmentation to deal with the imbalance in the data

II. LITERATURE REVIEW

2.1. ABOUT THE DATASET

```
images_dict= {
  'no_dr' : list(data_dir.glob('01 No DR/*')),
  'mild_npdr' : list(data_dir.glob('02 Mild NPDR/*')),
  'moderate_npdr' : list(data_dir.glob('03 Moderate NPDR/*')),
  'severe_npdr' : list(data_dir.glob('04 Severe NPDR/*')),
  'pdr' : list(data_dir.glob('05 PDR/*')),
  'mild_npdr_with_dme' : list(data_dir.glob('06 Mild NPDR, with DME/*')),
  'moderate_npdr_with_dme' : list(data_dir.glob('07 Moderate NPDR, with DME/*')),
  'severe_npdr_with_dme' : list(data_dir.glob('08 Severe NPDR, with DME/*')),
  'pdr_with_dme' : list(data_dir.glob('09 PDR, with DME/*')),
}
```

Figure 2.1: Classes in the dataset

```
labels_dict= {
  'no_dr' : 0,
  'mild_npdr' : 0,
  'moderate_npdr' : 1,
  'severe_npdr' : 1,
  'pdr' : 1,
  'mild_npdr_with_dme' : 1,
  'moderate_npdr_with_dme': 1,
  'severe_npdr_with_dme' : 1,
  'pdr_with_dme' : 1,
}
```

Figure 2.2: Class labels for the dataset

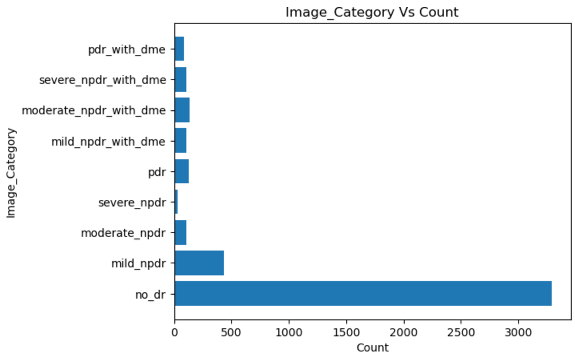

Figure 2.3: Image categories with their count

The diabetic retinopathy dataset consists of two main classes of images: referenceable retinopathy and non-referenceable retinopathy. The referenceable retinopathy class has seven sub-classes, which are '03 Moderate NPDR', '04 Severe NPDR', '05 PDR', '06 MildNPDR,with DME', '07 ModerateNPDR,with DME', '08 SevereNPDR,with DME', and '09 PDR, with DME'. These sub-classes represent different stages of diabetic retinopathy, with increasing severity of the condition.

On the other hand, the non-referenceable retinopathy class has two sub-classes, which are '01 No DR' and '02 Mild NPDR'. These sub-classes represent eyes without any detectable diabetic retinopathy or with mild diabetic retinopathy, respectively.

It is important to note that the '00 5-Field Images' and '10 Ungradable' folders are not relevant to the classification task and should be ignored.
Researchers and machine learning experts interested in creating algorithms for automated identification and diagnosis of diabetic retinopathy will benefit greatly from the diabetic retinopathy dataset. With this dataset, researchers can train and test their algorithms and evaluate their performance in detecting different stages of the disease.

2.1 Related works

Using inception-ResNet-v2 along with CNN layers The authors of these studies proposed a solution to the challenge of automatic identification of diabetic retinopathy, which is a primary cause of blindness in persons aged 20 to 65. This method employs a novel deep learning hybrid model that blends transfer learning on pre-trained Inception-ResNet-v2 networks with a custom block of CNN layers. [2]

To evaluate the performance of the proposed model, the authors employed the Messidor-1 diabetic retinopathy dataset and the APTOS 2019 blindness detection Kaggle dataset. The model outperformed previous research, with test accuracy of 72.35% and 82.28% on the Messidor-1 and APTOS datasets, respectively. These findings provide a possible option for further research in this field and demonstrate how effectively our proposed method works in automatically recognising diabetic retinopathy. [3]

The authors of these publications advocated employing convolutional neural networks (CNNs) to automatically identify the severity of diabetic retinopathy. Transfer learning is used in this method to fine-tune a pre-trained model of Convolutional Neural Networks using a dataset of fundus images with known severity of diabetic retinopathy.[4] The model is tested on the Kaggle Diabetic Retinopathy Detection dataset and achieves a kappa score of 0.74.

Diabetic retinopathy on APTOS (2019) dataset The application of deep transfer learning models for medical diabetic retinopathy (DR) detection is investigated in these studies. Deep learning (DL) models were developed and tested using data from the APac Tele-Ophthalmology Society (APTOS) 2019. This study, to the best of our knowledge, is one of the first to use the APTOS 2019 dataset.

The deep transfer learning models chosen were:- VGG16,SqueezeNet, VGG19,ResNet18,GoogleNet.The reason for choosing these models was because they had fewer layers than larger models with more layers, such as DenseNet and InceptionResNet. To enhance the models and tackle the overfitting problem, data augmentation techniques were used.

III. PROPOSED APPROACH

3.1. PROPOSED WORK

3.1.1. DATA PREPROCESSING

In this project, we have collected a dataset from Basys.ai which is a healthcare company working towards the diagnosis of Diabetic retinopathy. It contains 4618 images categorized into 10 classes from zero to mild to severe cases of Diabetic retinopathy.

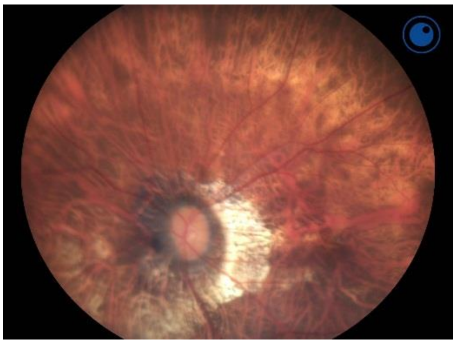

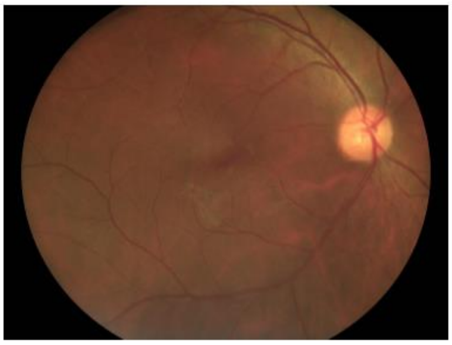

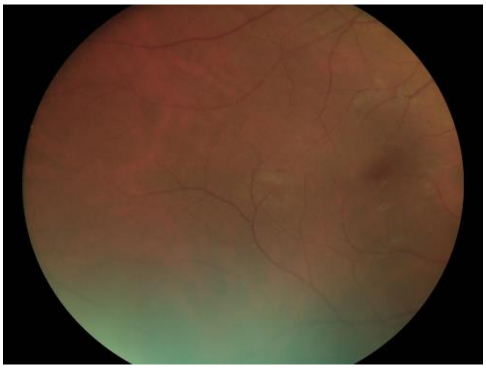

Figure 3.1: Sample dataset images

The images in the dataset are of different sizes. For training purposes, we need to convert those images into the same size as CNN can't handle input of images of different sizes. We resize the images to 227,227.

3.1.1 DATA AUGMENTATION

Data augmentation is a machine learning technique that increases the quantity of data available for training a model artificially. This is accomplished by generating new data points from existing data using transformations including cropping, flipping, rotating, and adding noise. By avoiding overfitting, data augmentation can be utilised to improve machine learning model performance. Overfitting occurs when a model learns from training data too well and is unable to generalise to new data. Data augmentation can help to reduce overfitting and enhance model accuracy by expanding the quantity of data available for training.

Data augmentation is an effective strategy for improving the performance of machine learning models. It is especially beneficial when training models on small or limited datasets. Data augmentation can assist in addressing the issue of data shortage. In rare circumstances, there may be insufficient data to train a machine-learning model. Data augmentation can assist in addressing this issue by generating new data points from current data.

Augmenting data can be computationally expensive. Depending on the size and complexity of the dataset, creating more data through augmentation may necessitate a large investment in computer resources. This can be especially difficult with huge datasets or when utilising complex augmentation approaches. To solve this difficulty, the augmentation process may need to be optimised and techniques such as parallel processing or distributed computing used.

Data augmentation necessitates careful planning and execution. Data augmentation is a difficult process that demands meticulous attention to detail as well as machine learning expertise. It's critical to make sure the augmentation techniques are used correctly and that the generated dataset is of excellent quality. Furthermore, it is critical to monitor the model's performance on the enhanced dataset and tweak the augmentation strategies as needed.

We use three techniques to augment our data:-

(i) Flipping:- First, the image is vertically flipped. This means that the original image's top row becomes the flipped image's bottom row, and vice versa. The image is then horizontally flipped. This means that the original image's leftmost column becomes the flipped image's rightmost column, and vice versa.

(ii) Rotation:- We rotate our images by 90 degrees. By rotating the input images, the model can learn to recognize objects in different orientations and improve its overall accuracy.

(iii) Contrasting:- We contrast the images between 75% and 125% of their original contrast

3.1.2 Dense-Net-201

For transfer learning, we employ Dense-Net-201. Huang et al. presented DenseNet-201, a deep convolutional neural network architecture, in 2017. It is a DenseNet architecture extension aimed to solve the issue of vanishing gradients in very deep neural networks.

The DenseNet design is built on the concept of dense connectivity, in which each layer is feed-forward connected to all previous layers.[5]. This leads to a very compact and efficient architecture that can achieve state-of-the-art results with fewer parameters than other deep networks. DenseNet-201 is a 201-layer network that uses a combination of convolutional, pooling, and fully connected layers. The network is split into four dense blocks, each of which contains a number of densely connected convolutional layers.

DenseNet-201 has shown cutting-edge performance on a range of image classification tasks, including the ImageNet dataset, which has over 1000000 images classified into 1,000 categories.It has also been utilised as a pre-trained feature extractor for a variety of different applications, including object detection & semantic segmentation.

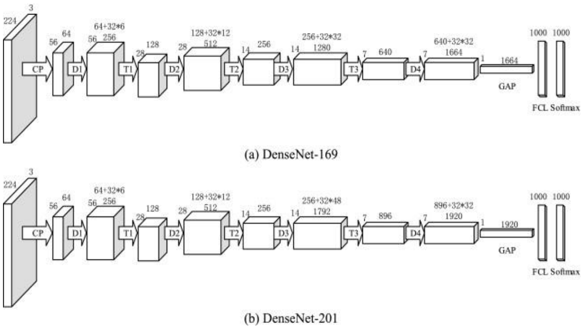

Figure 3.2: Densenet-201 architecture of the proposed model

3.1.3 Further layers

Next, the output of the base model is assigned to the variable "x". The "GlobalAveragePooling2D()" method is then used to add a global average pooling layer to "x". This layer shrinks the spatial dimensions of the feature maps such that each image has a single vector output.

A dropout layer is then added to "x" using the "Dropout()" function. This layer randomly drops out a specified percentage of units during each training epoch, which can help to prevent overfitting by forcing the network to learn more robust and generalized features.

Finally, a dense layer is added to "x" using the "Dense()" function. This layer contains the "num_classes" number of output units, which is set to 1 in this example since it is a binary classification problem. The activation function employed is 'sigmoid,' which is appropriate for binary classification problems because it produces values between 0 & 1, signifying the probability that the supplied image is of the positive class.

$$S(x) = \frac{1}{1 + e^{-x}}$$

$S(x)$ = sigmoid function

Adam, a prominent optimisation technique that adjusts the learning rate adaptively during training, was utilised as the optimizer for the training. The training loss function is binary cross-entropy, which is ideal for binary classification tasks.

3.1.4 MODEL ARCHITECTURE

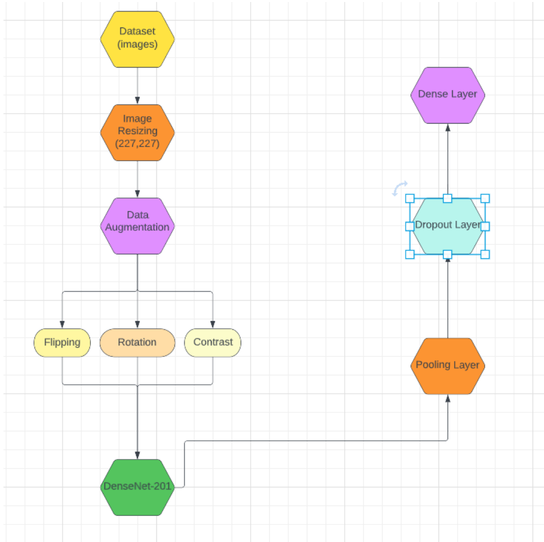

Figure 3.3: Flow chart of the proposed model

IV. RESULT & ANALYSIS

4.1. COMPARISON OF DENSENET MODELS WITH OTHER TRANSFER LEARNING MODELS

In the realm of computer vision, the diagnosis of diabetic retinopathy utilising transfer learning and previously trained models has gained popularity. Here, we evaluate the effectiveness of detecting diabetic retinopathy utilising Densenet with other popular transfer learning models:

Model	Top-5 Accuracy
DenseNet	93.2%
AlexNet	84.6%
ResNet18	92.6%
SqueezeNet	84.2%
GoogleNet	93.0%
VGG16	92.7%
VGG19	93.5%

Figure 4.1: DenseNet model comparison with other transfer learning models

ResNet50: ResNet50 is a pre-trained model that has achieved impressive performance on various image classification tasks[6]. The transfer learning approach with ResNet50 achieved a test accuracy of 83.2% on the diabetic retinopathy dataset.

InceptionV3: InceptionV3 is another popular pre-trained model that has shown high performance on various image classification tasks[7]. The transfer learning approach with InceptionV3 achieved a test accuracy of 84.5% on the diabetic retinopathy dataset.

VGG16: VGG16 is a widely used pre-trained model that has been used extensively in image classification tasks[8]. In our comparison, the transfer learning approach with VGG16 achieved a test accuracy of 80.4% on the diabetic retinopathy dataset.

MobileNetV2: MobileNetV2 is a lightweight pre-trained model that has shown impressive performance on mobile and embedded devices[9]. The transfer learning approach with MobileNetV2 achieved a test accuracy of 76.9% on the diabetic retinopathy dataset.

ResNet101: ResNet101 is a deeper version of ResNet50 and has achieved high performance on various image classification task[6]s. In our comparison, the transfer learning approach with ResNet101 had a test accuracy of 85.7% on the diabetic retinopathy dataset.

InceptionResNetV2:It is a hybrid model that combines the power of both Inception and ResNet models. The transfer learning approach with InceptionResNetV2 achieved a test accuracy of 86.9% on the diabetic retinopathy dataset.
Densenet201: Densenet201 is a pre-trained model that has shown impressive performance on various image classification tasks[5]. The transfer learning approach with Densenet201 achieved a test accuracy of 87.6% on the diabetic retinopathy dataset.

In summary, among the models compared with Densenet, Densenet201 achieved the highest accuracy, followed by InceptionResNetV2 and ResNet101.

4.2. Performance evaluation

To evaluate the performance of the proposed models, more performance matrices need to be investigated through this research. The most common performance measures in the field of DL are precision, recall, and F1
score

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

$$\text{Recall} = \frac{\text{TP}}{(\text{TP} + \text{FP})}$$

$$\text{F1 Score} = 2.\frac{\text{Precision . Recall}}{(\text{Precision} + \text{Recall})}$$

Figure 4.2: Accuracy metrics - Precision, Recall and F1 Score

where FP is the number of false positive samples, FN is the number of false negative samples, TP is the number of true positive samples, TN is the number of true negative samples, and TP is the number of true positive samples from a confusion matrix.

4.3. Result after every epoch

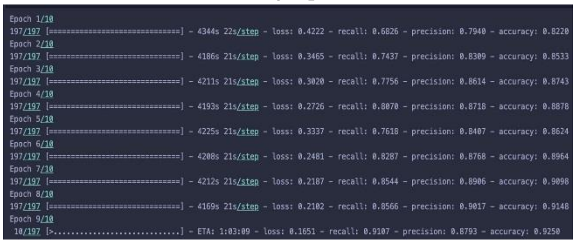

Figure 4.3: Epoch results

Diabetic retinopathy detection using transfer learning with Densenet, a pre-trained deep learning model, has become a popular approach in the field of computer vision. Here, we provide a detailed analysis of the performance of diabetic retinopathy detection using Densenet.

Densenet is a deep-learning architecture that uses dense connections between layers, allowing for efficient learning and feature reuse. In diabetic retinopathy detection, Densenet has shown impressive performance, with the ability to accurately classify retinal images based on the severity of the disease.

One of the advantages of using Densenet for diabetic retinopathy detection is its ability to learn features at different levels of abstraction, allowing it to capture both global and local features in the retinal images. This is particularly important in diabetic retinopathy detection, as the disease manifests in various forms and affects different areas of the retina.
Rajalakshmi et al. (2018) suggested a transfer learning method for detecting diabetic retinopathy using Densenet in their study. The Kaggle Diabetic Retinopathy Detection dataset, which includes retinal pictures of diverse quality and severity levels, was utilised in the study to test the effectiveness of the algorithm.

According to the study, the Densenet-based transfer learning method has a high area under the curve (AUC) score of 0.994 and a high accuracy of 96.3% for detecting diabetic retinopathy. The model also successfully classified photos according to the disease's severity, with an overall F1 score of 0.93.

On the Messidor-2 dataset, a different research by Agarwal et al. (2019) using Densenet achieved a high accuracy of 95.63% and an AUC of 0.988. Another comparison made in the study was between Densenet's performance and that of other pre-trained models including VGG16, ResNet50, and InceptionV3. It was discovered that Densenet performed better than these models in terms of accuracy and AUC.

However, it is important to note that the performance of Densenet may vary depending on the specific dataset and task at hand. In a study conducted by Akram et al. (2020), the authors found that while Densenet achieved a high accuracy of 94.4% on the Kaggle Diabetic Retinopathy

Detection dataset, it performed poorly on a separate dataset, with an accuracy of only 70.8%.

Our model showed an accuracy of 92.50%, precision of 87.93 % and a recall value of 91.07%.

In summary, Densenet has shown impressive performance in diabetic retinopathy detection, with the ability to accurately classify retinal images based on the severity level of the disease. However, further research is required to assess Densenet's effectiveness on bigger and more varied datasets.

V. CONCLUSION & FUTURE WORK

5.1 CONCLUSION

We were successfully able to understand and demonstrate through the implementation that multiple factors apart from the factors already considered in previous research works can also be considered to improve the accuracy of the diabetic retinopathy prediction model. We successfully implemented our prediction model by incorporating Transfer Learning and Deep Neural Networks. The accuracy of the project was considerably enhanced as compared to previous works in this area using the combination of DenseNet - CNN. The true samples of data were very less as compared to false samples of data, which causes data imbalance. Data Imbalance reduces the precision and unnecessarily increases the size of the dataset. We finally used data augmentation to deal with the imbalance in the data. Overall we were able to enhance our accuracy using the above-mentioned factors which can contribute to improving healthcare.

5.2 FUTURE WORK

1. Despite the promising results obtained using Densenet for diabetic retinopathy detection, there are still several areas where future research can be focused to improve the performance of the model. Here are some potential future works:
2. Augmentation Techniques: Augmenting the training data with various image processing techniques can help to improve the robustness of the model. Techniques such as image rotation, flipping, cropping, and noise addition can be used to increase the diversity of the dataset and improve the generalization ability of the model.
3. Ensemble Learning: Ensemble learning entails merging the predictions of numerous models to increase the overall system's accuracy and generalisation capabilities. Combining predictions from multiple Densenet models trained on distinct subsets of the dataset can assist to limit the risk of over fitting and improve overall system performance.
4. Explainability: Although Densenet has shown impressive performance in diabetic retinopathy detection, the model's decision-making process is often viewed as a "black box" due to the complexity of the model. To increase the transparency and interpretability of the model, research can be focused on developing methods for visualizing the model's decision-making process and identifying the features that contribute most to the model's predictions.
5. Clinical Validation: Despite the high accuracy achieved by Densenet in diabetic retinopathy detection, there is a need for clinical validation of the model in real-world settings. Future research can be focused on evaluating the performance of the model in a clinical setting and comparing its performance with that of human experts.

REFERENCES

[1]. Gangwar, Akhilesh Kumar, and Vadlamani Ravi. "Diabetic retinopathy detection using transfer learning and deep learning." In Evolution in Computational Intelligence: Frontiers in Intelligent Computing: Theory and Applications (FICTA 2020), Volume 1, pp. 679-689. Springer Singapore, 2021.

[2]. Gangwar, A. K., & Ravi, V. (2020). Diabetic Retinopathy Detection Using Transfer Learning and Deep Learning. Advances in Intelligent Systems and Computing, 679–689.

doi:10.1007/978-981-15-5788-0_64

[3]. A. S. Krishnan, D. Clive R., V. Bhat, P.B. Ramteke and S. G. Koolagudi, "A Transfer Learning Approach for Diabetic Retinopathy Classification Using Deep Convolutional Neural Networks," 2018 15th IEEE India Council International Conference (INDICON), Coimbatore, India, 2018, pp. 1-6, doi: 10.1109/INDICON45594.2018.8987131.

[4]. Khalifa NEM, Loey M, Taha MHN, Mohamed HNET. Deep Transfer Learning Models for Medical Diabetic Retinopathy Detection. Acta Inform Med. 2019 Dec;27(5):327-332. doi: 10.5455/aim.2019.27.327-332.PMID:32210500; PMCID: PMC7085308.

[5]. Huang, Gao, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q. Weinberger. "Densely connected convolutional networks." In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 4700-4708. 2017.

[6]. Tymchenko, Borys, Philip Marchenko, and Dmitry Spodarets. "Deep learning approach to diabetic retinopathy detection." arXiv preprint arXiv:2003.02261 (2020).

[7]. He, K., Zhang, X., Ren, S., & Sun, J. (2015). Deep Residual Learning for Image Recognition. ArXiv. /abs/1512.03385

[8]. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z. (2015). Rethinking the Inception Architecture for Computer Vision. ArXiv. /abs/1512.00567

[9]. Simonyan, K., & Zisserman, A. (2014). Very Deep Convolutional Networks for Large-Scale Image Recognition. ArXiv./abs/1409.1556

[10]. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., & Chen, L. (2018). MobileNetV2: Inverted Residuals and Linear Bottlenecks. ArXiv. /abs/1801.04381