Anusandhanvallari
Vol 2024, No.1
January 2024
ISSN 2229-3388

A Deep Learning Framework with Multiple Stages for Better Diabetic Retinopathy Detection

Ahmad Talha Siddiqui8[1, 4], Harleen Kaur[2], Safdar Tanveer[3]
Research Scholar*[1], Department of Computer Science & Engineering
School of Engineering & Technology
Jamia Hamdard (Hamdard University)
Email-id: ahmadtalha2007@gmail.com
Professor[2], Department of Computer Science & Engineering
School of Engineering & Technology
Jamia Hamdard (Hamdard University)
Email-id: harleen@jamiahamdard.ac.in
Associate Professor[3], Department of Computer Science & Engineering
School of Engineering & Technology
Jamia Hamdard (Hamdard University)
Email-id:safdartanveer@jamiahamdard.ac.in
Assistant Professor[4], Department of CS & IT
Maulana Azad National Urdu University
Email-id: ahmadtalhasiddiqui@manuu.edu.in

ABSTRACT

One of the leading causes of blindness worldwide is diabetic retinopathy (DR), the most common diabetic eye disease. Retinal image classification by hand is currently a laborious procedure that calls for specific knowledge. In this paper, we looked into the possibility of automating this task using convolutional neural networks (CNNs). A number of potent deep learning models, such as InceptionV3, ResNet50, ResNet50V2 and DenseNet201, were chosen. We used data augmentation to prevent overfitting and increase the resilience of the models. Two stages comprised our training strategy: a first round of feature extraction using a custom classifier, and a fine-tuning step where we unfroze several layers to make the models better fit our data.

We tested our approach on the PubMed, Messidor, and kaggle diabetic retinopathy dataset and discovered that it worked quite well. With a remarkable accuracy of 94.51%, the improved InceptionV3 model was the most notable.

KEYWORDS

Diabetic Retinopathy, InceptionV3, Resnet50, Resnet50V2, Densenet201, Feature extraction, Fine-tuning, Convolutional Neural Network.

1. Introduction

Around the world, one of the main causes of vision loss and blindness [1] is eye illnesses that impact the fundus, the back of the eye. Among the most prevalent of these are diabetic retinopathy (DR), cataracts, and age-related macular degeneration (AMD). Many of these illnesses currently have no known cure, and when they advance to later stages, they can significantly impair vision.

Diabetic retinopathy is a serious issue. As diabetes affects a large number of people worldwide, DR has emerged as one of its most common side effects. It's a cunning illness that frequently doesn't cause any symptoms at first but can eventually cause blindness by slowly destroying the blood vessels in the retina. It is indeed one of the top four

causes of blindness [2, 31] in the world. In clinical settings, fundus scans are used by physicians to diagnose and track DR, a disorder that affects hundreds of millions of people [3]. Because of this, combating DR is of utmost importance and calls for a trifecta of biomedical research, legislative actions, and technological technology.

DR has becoming the most common cause of blindness in many Western nations as population's age. Timely treatment and early detection are essential for preventing severe vision loss. Without it, vision can quickly and permanently decline. Artificial intelligence (AI) has promising answers for this problem. AI can help physicians by using vast amounts of medical data to provide effective tools for early diagnosis and detection, particularly in primary care settings. Working together, AI and ophthalmologists could greatly enhance patient treatment for the increasing number of individuals suffering from fundus illnesses [4].

A branch of artificial intelligence called deep learning has demonstrated impressive success in numerous domains due to its capacity to extract intricate patterns from data. In particular, deep convolutional neural networks, or CNNs, have demonstrated remarkable performance in image categorization tasks [5,28]. Medicine greatly benefits from the ability to adapt a model trained on a big dataset for different tasks using techniques like transfer learning.

CNN application to fundus images, however, presents a unique set of difficulties. Since a single retinal image may exhibit symptoms of several different diseases in the real world, multi-label classification is a more realistic yet challenging issue. It is also challenging together big, high-quality datasets, particularly for rarer illnesses. A significant challenge is building a trustworthy model from start when dealing with sparse and occasionally noisy data.

In order to address these problems, we created a multi-phase training plan. Initial feature extraction [6] and fine-tuning [7] are two potent deep learning approaches that are combined in this method. To get over the restrictions of smaller, more specialized medical datasets, the main concept is to reuse knowledge from larger datasets. By methodically enhancing a number of models, we essentially create a powerful, precise diagnostic tool. To help with the diagnosis of DR, we created and assessed four different deep learning frameworks in this study.

2 Related Works

2.1 Traditional Approaches

Computer-aided diagnosis was the gold standard for automated eye disease detection until the deep learning revolution [29,30]. These systems successfully divided the diagnosis procedure into parts that resembled those of a clinician: they segmented lesions [8] to isolate important areas, preprocessed images to enhance quality, extracted significant features from those areas, and then used machine learning to classify the disease [9].

2.2 Deep Learning Methods

A lot of progress has been made in using deep learning to identify diabetic retinopathy (DR). Key milestones have been identified by research from top universities. The turning point came in 2016 when a Google research team released a study in JAMA showing that a deep neural network could accurately find DR in fundus images, just like a real eye doctor. Soon after, researchers from Stanford University published work in Ophthalmology that showed even more promise for AI in ophthalmology and got a lot of media notice [10].

Since these early discoveries, study has grown to include many more complex methods that aim to make things more accurate and useful. Dealing with the problem of not having enough information has been a main theme. A lot of researchers have been able to successfully use transfer learning [11], which involves changing powerful models that have already been trained, such as InceptionV3 [12], ResNet50 [13] , ResNet50V2 [14] and DenseNet201 [15] to do the specific job of DR classification. With this method, you can get good results even with smaller samples. Advanced data augmentation, color normalization, and methods like synthetic minority over-sampling to deal with class imbalance are some other ways to deal with a lack of data [16].

Another important trend is the creation of neural network designs that are getting more complicated and specialized. Researchers are making custom answers instead of using standard models [17]. Some people have come up with hybrid models that combine the best parts of different designs [18], while others have made completely new networks, such as AC-DenseNet201 [19], which uses attention and dense connectivity to improve feature extraction and work better on rotated images. Groups of models have also been shown to work well for getting high accuracy in public datasets such as Messidor, PubMed, and Kaggle

The effects are really great. Accuracy rates for multi-class classification (finding the different steps of DR) have been steadily rising. Many studies have reported accuracy rates above 90%, and some have even reported results above 98% or 99% on certain datasets. Not only are these models tested for accuracy, but also for important clinical metrics such as sensitivity (recall), specificity, and F1-score, which are necessary for a diagnostic tool to be dependable [20].
In short, early methods depended on features that were engineered by hand. Today's deep learning is powerful because it can automatically learn complex features from images [21]. The field has gone from proving that the idea is possible to improving models for greater accuracy, robustness, and clinical usefulness [22]. This shows a huge amount of promise for automating early diagnosis and making healthcare systems less busy.

3 Proposed Method

A novel approach to diabetic retinopathy (DR) detection and classification utilizing deep learning algorithms based on actual eye scans is investigated in this paper. We have a three-stage process. We started by gathering fundus pictures to use as a dataset. A total of four deep learning models were subsequently trained using this dataset: InceptionV3, DenseNet201, ResNet50, and ResNet50V2. Finally, we ran tests on all of the models to determine which one worked best in each scenario. In this part, we will go over each of these steps in depth. In Figure 1 we can see the whole procedure.

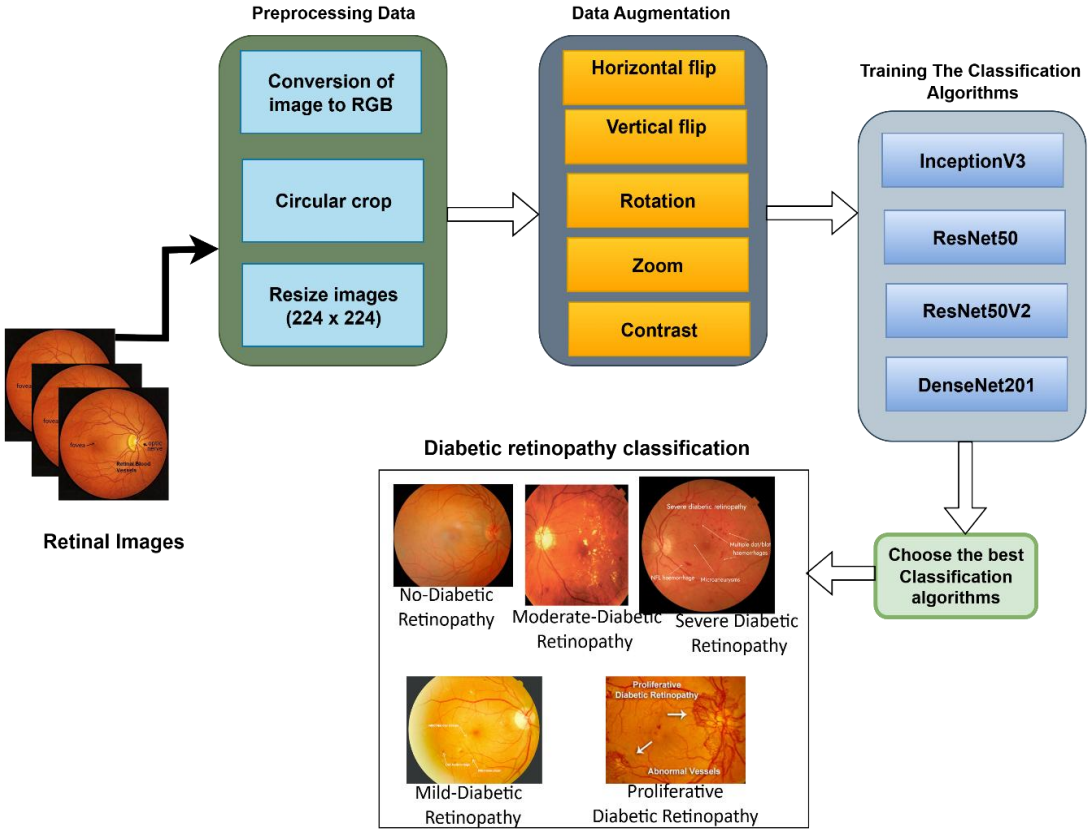

Figure 1: The following describes our suggested process.

3.1 Dataset

We used a set of real-life retina pictures to train our models. Each image had already been given a clinical severity score for diabetic retinopathy (DR) on the standard ICDR scale from 0 to 4. Our dataset [23] was built on top of a public set of 3,662 photos on different datasets. We did a careful cleaning

process, though, to make sure that our results were accurate. We got rid of entries that were already there and pictures that were too noisy or broken to be useful. This made a refined dataset of 2,631 high-quality images that we used for all training and tests that came after. As you look at Figure 2, you can see cases from each severity class.

Figure 2: The progression of diabetic retinopathy

3.2 Data Pre-Processing

We were required to go through a number of preparation processes in order to have our dataset ready for the deep learning models. This portion of the document provides a step-by-step guide to the procedure, beginning with the handling of the raw data and continuing through the application of augmentation techniques.

The original pictures of the retina are quite huge in size, with proportions of 2,896 by 1,944 pixels. Nevertheless, we were confronted with a frequent problem in real-world data: a considerable imbalance in the number of samples in each class. There are a significantly greater number of "No-DR" photos than "Severe" cases, as is demonstrated in Figure 3. Because the photographs included in the dataset were gathered from a variety of clinics over a period of time and were taken with a number of different cameras, it comes as no surprise that the collection also included its fair share of noise, which refers to images that were blurry, overexposed, or excessively dark.

The standardization of the photographs was the first thing that we did. Every single one of them was transformed into RGB color channels. After that, we centered a circular crop on the image and then applied it in order to concentrate on the retina itself. The black corners were effectively removed by this process, while the circular field of view, which was of utmost importance, was maintained.

We adjusted the size of all photos to 224x224 pixels in order to make training more efficient. We were able to utilize a bigger batch size of 64, which sped up the training process, because of this smaller, consistent size.

In order to increase the resilience of our model, we made use of data augmentation. We discovered that when the photos were zoomed in, certain light patches could be misconstrued by the model, as seen in Figure 4. The possibility of overfitting is significantly reduced by the model's ability to concentrate on the genuine indicators of illness. This ability is made possible by the model's use of augmentation techniques such as random zooming, flipping, and rotation, which enable it learn to disregard these changes that are not relevant.

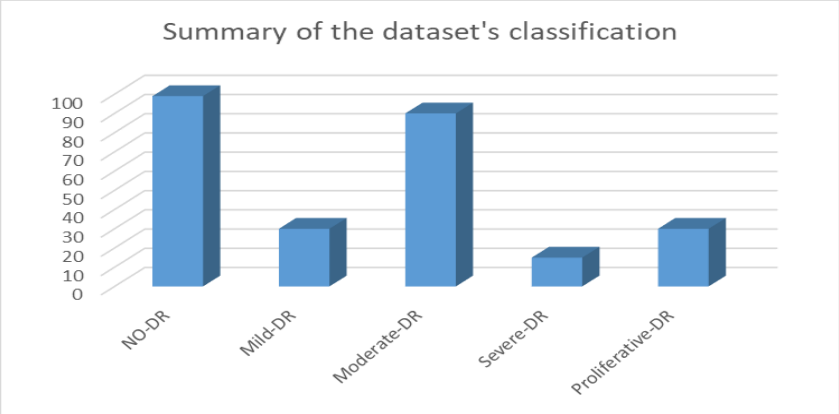

Figure 3: Summary of the dataset's classifications

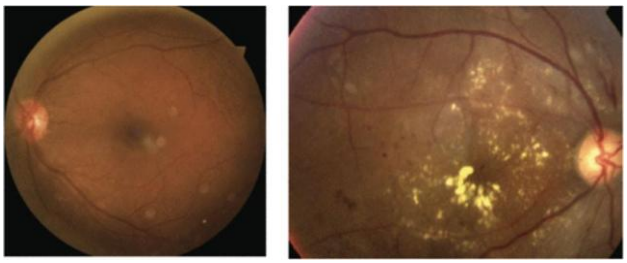

Figure 4: The image on the right presents a magnified perspective

3.3 Data Augmentation

We employed a number of augmentation approaches to increase the versatility of our dataset and aid in the model's improved generalization. We trained our algorithm on an updated dataset to distinguish retinas in different situations. Camera angles and difficult lighting were simulated by randomly zooming, flipping, rotating, and adjusting contrast Figure 5. Also, we found a way to remove the irritating black borders around circular retina scans.

Figure 5: Overview of the data augmentation methodologies utilized

Initially, we clipped out the black borders automatically. A technique known as "polar unrolling," which turns the circular image into a rectangular strip, was then employed. We discovered that this unrolling procedure decreased the requirement for additional rotation augmentations because it in itself produces a new perspective on the image.

4 Training of the Selected Algorithms

Transfer learning was smarter than constructing an AI model from scratch, a major challenge in medical AI. Our unique diabetic retinopathy diagnosis method starts with a model that recognizes broad images and carefully adjusts it to recognize disease symptoms.

4.1 Approach and Adopted Models

We selected four of the most recent image recognition models for our investigation: InceptionV3, ResNet50 ResNet50V2, and DenseNet201. According to a recent assessment by [24], these are some of the most widely utilized and practical designs in the sector today. Our objective was to determine the most accurate approach for diagnosing diabetic retinopathy, irrespective of the underlying model architecture, by testing this heterogeneous population.

Each model's training phase began with weights it had learned from the enormous ImageNet database. The benefit of this "pre-training" is enormous, despite the fact that our own retinal scan dataset is far smaller. This pre-training equips the models with a foundational comprehension of prevalent visual patterns. Although a retina scan differs significantly from typical images in ImageNet, this basic information facilitates the models' ability to recognize the unique indicators of diabetic retinopathy more rapidly and efficiently

4.1.1. InceptionV3

Over a million photos were used to teach InceptionV3, a powerful image recognition model, to recognize 1,000 different kinds of objects. You send it a 299x299 pixel picture to use it, and it tells you what it thinks the picture shows.

Its smart design uses "Inception" modules, which look at a picture at many sizes at once for patterns. This system also uses "skip connections," which are like shortcuts that let data flow easily through all 48 of its layers. That fixes a typical issue in deep networks, which lets the model learn faster and do better. Figure 6 displays how all of these parts fit together.

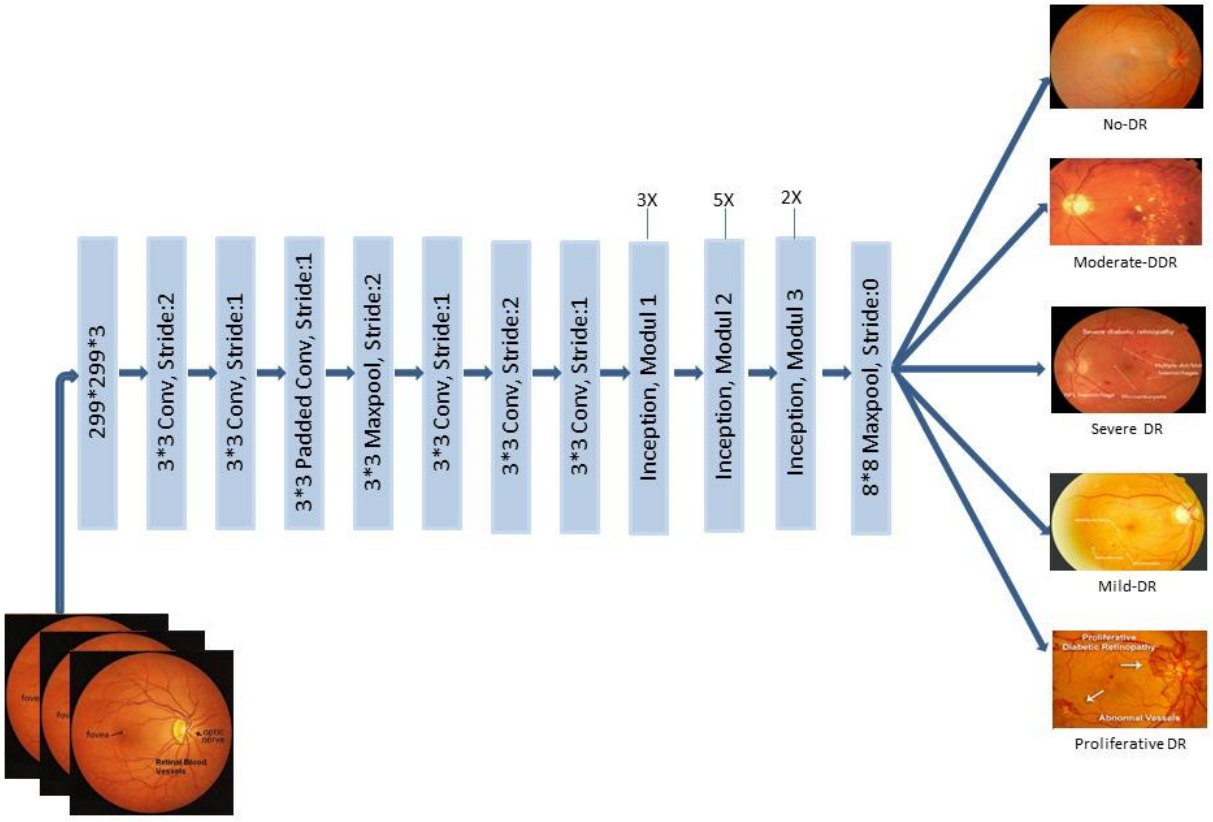

Figure 6: InceptionV3 architecture

4.1.2. ResNet50

The 224x224-pixel standard color picture that ResNet50 is meant to look at is what it actually does. The number "50" in its name might make you think it's a smaller model, but it's actually quite powerful and complicated. It has

fifty layers and an astonishing 138 million parameters, which makes it a big network even in today's world. Sorting photos into one thousand separate groups was the original purpose of its training.
The fact that ResNet50 has such an incredibly simple design is what makes it so effective. The architecture is based on a fairly logical pattern: convolutional layers are always followed by pooling layers, which gradually lower the size of the image.
There is a clear and regular evolution in the number of filters that it employs to "look" at the data, beginning with 64 and then systematically doubling until it reaches 128, then 256, and eventually 512 in the deepest sections of the network. As a result of this structured increase, it is able to learn features that are ever more complicated. Figure 7 provides a visual representation of this simplified design that you can observe for yourself.

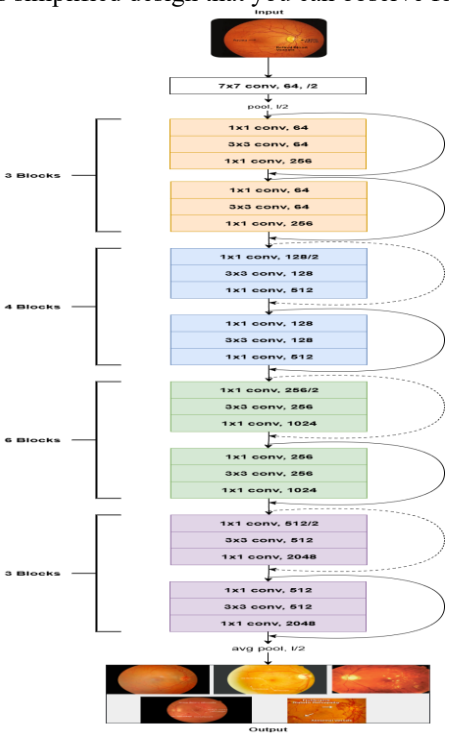

Figure 7: ResNet50 architecture

4.1.3. ResNet50V2

ResNet50V2 presents a more methodical approach to augmenting a neural network. Rather than indiscriminately augmenting a network's depth, breadth, or input resolution—which might be speculative—ResNet50V2 proportionately and systematically enhances all three dimensions.

Instead of stacking rooms, it's like adding a clever extension to a house. Making the network deeper, wider, and able to see finer details is coordinated. This collaboration between components produces a more effective and superior model.

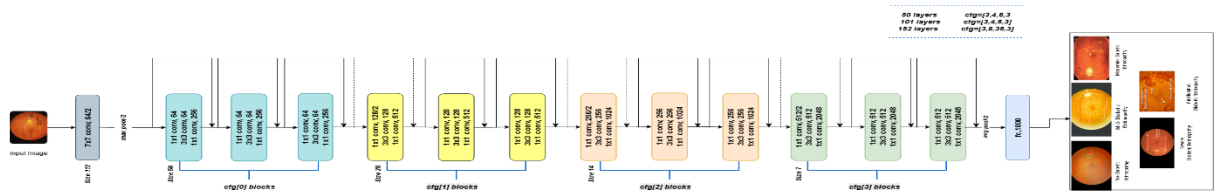

Figure 8: ResNet50V2 architecture

4.1.4. DenseNet201

DenseNet201, represented a significant advancement in CNN architecture due to its "dense" connectedness. In contrast to conventional models, where a layer communicates solely with the subsequent layer, DenseNet201 allows each layer to accept input from all preceding layers.

Envision a team in which each member possesses access to the entirety of work completed by all other members up to that moment. This framework facilitates an exceptionally efficient information flow. A primary benefit is that the network is not need to re-acquire identical characteristics, hence considerably diminishing the parameter count and streamlining the overall architecture. The particular DenseNet201 architecture employed in our research is illustrated in Figure 9.

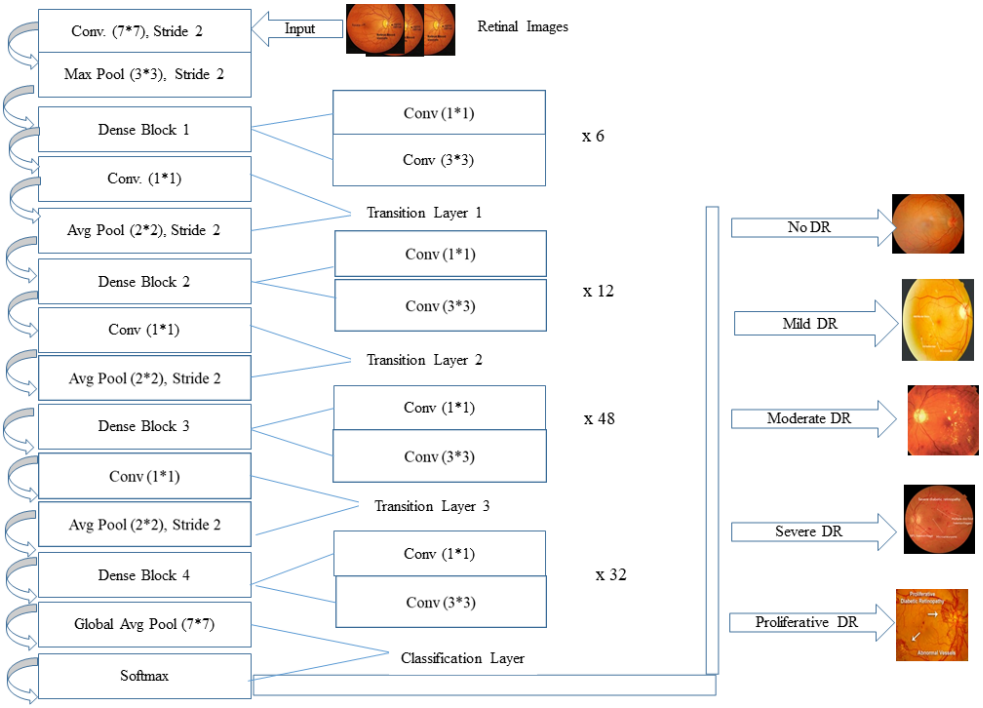

Figure 9: DenseNet201 architecture

4.2 Experimental Setup and Training

Our training strategy was divided into two primary stages. All four types were initially regarded as pre-made feature extractors. After that was finished, we proceeded to the fine-tuning phase, where we modified each model separately to improve accuracy.

4.3 First Stage: Feature Extraction

During this first stage, we aim to utilize the knowledge that is already there in our pre-trained models. Consider it

like looking at our new data through the eyes of an expert; we maintain their focus and sharpness by freezing their fundamental layers (the "convolutional base"). A specialized feature extractor is created from this basis.
We build a new brain—a unique classifier that we train from scratch—on top of this expert base. As a result, the model may learn our particular task without changing its broad understanding. Compared to training an entire model, this is far faster.
To down sample the features and make the data easier to handle, a 5x5 average pooling layer is used first. Next comes a dropout layer (set at 20%) that improves the model's generalization by reducing its reliance on any one feature. Then comes a batch normalization layer to expedite training and guarantee steady data flow. Lastly, the output layer's Softmax activation function provides our last predictions. This whole configuration is depicted in Figure 10.

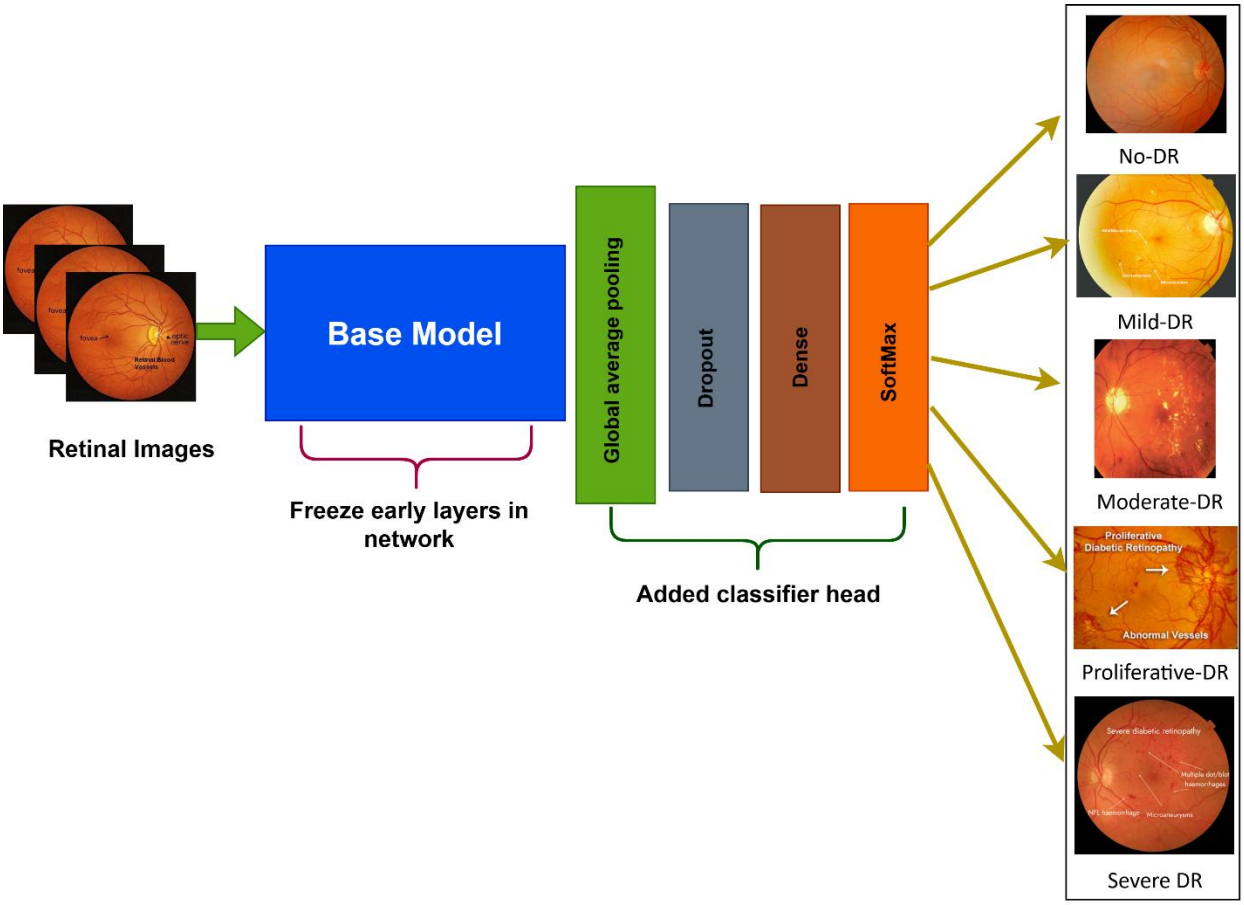

Figure 10: First training stage: Features extraction

4.4 Second Stage: Fine-Tuning

We simply trained a new classifier on top of the pre-trained model in our first experiment, leaving it locked. Consider it similar to following the recipes of a skilled chef without altering them.
We're fine-tuning now to go one step farther. To allow them to learn and adjust, we're releasing the final few layers of the original model. The model can improve its current knowledge to meet our particular requirements by training these layers in conjunction with our new classifier. The chef seems to be modifying their traditional recipes to make

them ideal for the ingredients we have. This whole configuration is depicted in Figure 11.

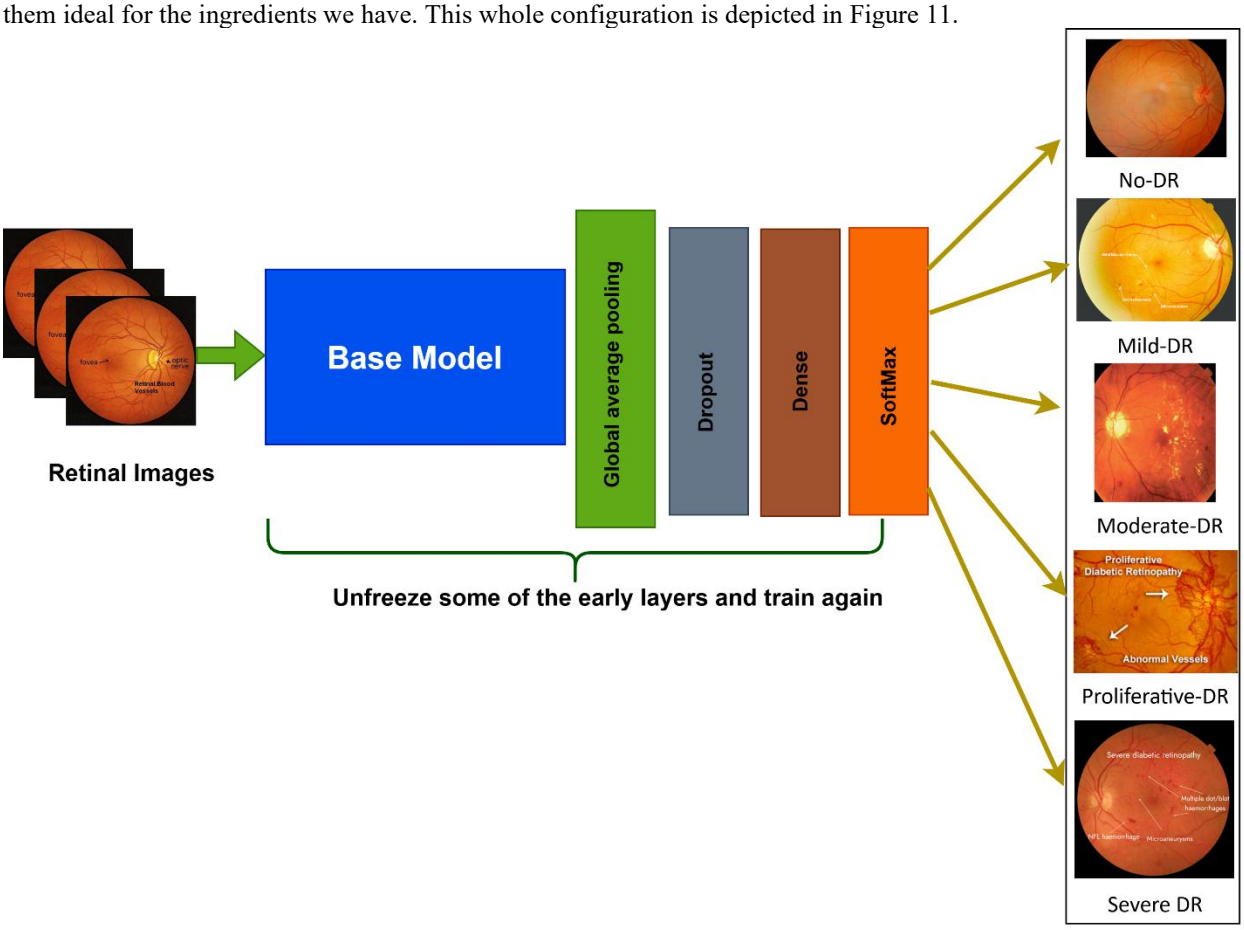

Figure 11: Second training stage: Fine-tuning

4.4.1 Experimental Setup

We utilized TensorFlow 2.1 [25] and its integrated Keras [26] API to construct and train our deep learning models [27]. Our trials were conducted on a Google Colab Pro account, which supplied us with 100 GB of RAM, P100 and T4 GPUs, and all code was written in Python [28]. We trained all of the models using 224x224 pixel pictures and a cross-entropy loss function using the Adam optimizer. Two stages of training were conducted: Feature extraction: 25 epochs of training were conducted with a learning rate of 1.00e-3.

For the fine-tuning step, we kept training for 25 more epochs while reducing the learning rate to 1.00e-5. The term "epoch" refers to the number of times a training image is shown to the model.

Based on the size of the model and the memory available in our GPU, we modified the batch size, which is the number of photos processed simultaneously, for each model. In Table 1 you can see a summary of all the model details, such as the number of layers and when fine-tuning started.

Model Name	Training Stage		Learning Stage		Number of total layers in base model	Fine-tuning start at layer number:	Batch Size
	Stage 1	Stage 2	Stage 1	Stage 2			
InceptionV3	25	25	1.00e-03	1.00e-05	780	100	64
ResNet50	25	25	1.00e-03	1.00e-03	19	3	64
ResNet50V2	25	25	1.00e-03	1.00e-05	22	4	64
DenseNet201	23	27	1.00e-03	1.00e-05	427	50	64

Table 1: Parameters in the multi stage training stage

4.4.2 Training Output

Training and validation accuracy and loss for all four models are shown in Figure 12 After a few training rounds, all models learnt quickly and performed well. Accuracy was 72%–77% during initial feature extraction. We improved performance by fine-tuning, achieving 90% to 95% accuracy. Lost values fell to a stable level.

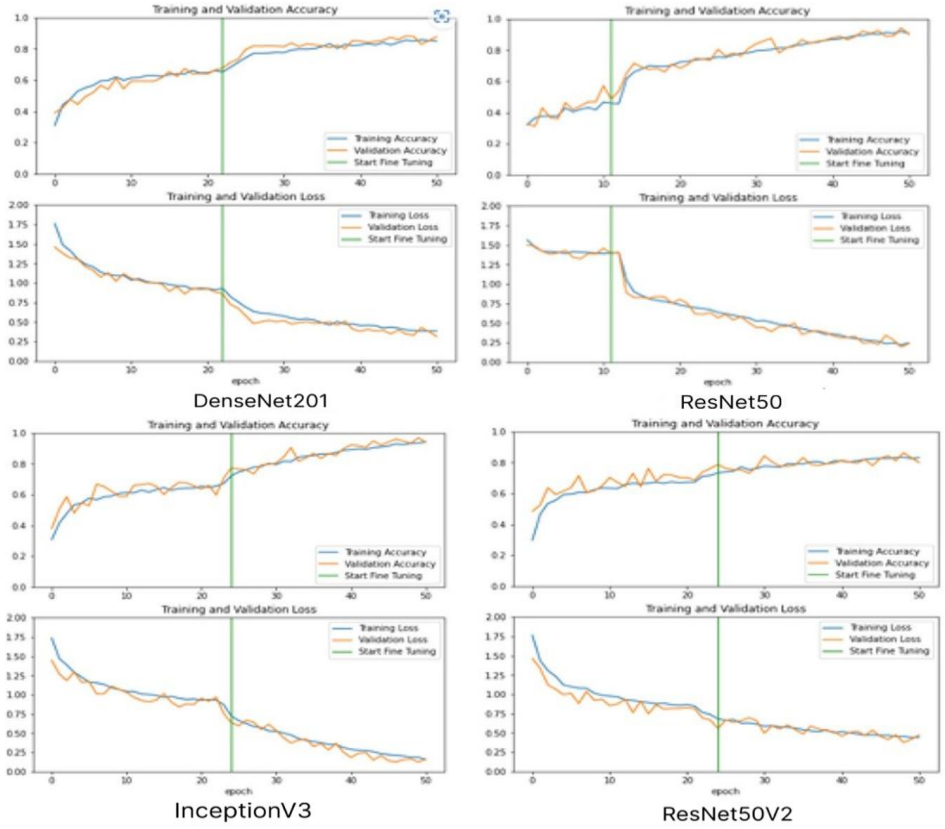

Figure 12: Training results

The models performed equally well on the validation set as they did on the training data, indicating no overfitting. Several aspects must be considered. These models must be tested with diverse retinal pictures and patient populations. It wasn't overfitting here, but more diverse data could cause it. Testing these models on external

datasets in clinical situations would give us confidence.
Table 2 shows all four models surpassed 90% accuracy after training. Real test is still to come. High scores are promising, but their value depends on the unknown test set. A model with good training data but inadequate test data is overfit and unsuitable for real-world use.

Model	Overall Accuracy	Overall Loss
InceptionV3	0.94	0.15
ResNet50	0.91	0.33
ResNet50V2	0.92	0.20
DenseNet201	0.93	0.31

Table 2: Overall accuracy and loss resulting from the training

5 Evaluations and Discussion

The outcomes of the training phase are discussed in this section. Next, in order to assess the suggested methods, we talk about the results.

5.1 Performance Evaluation and Metrics

We utilized four main metrics—accuracy, precision, recall, and the F1 score—to fully assess our suggested algorithms. Even though accuracy gives us a general idea of how right something is, we put more weight on precision and recall because they have direct clinical consequences.
Precision counts how trustworthy the model is when it marks a case as "No-DR." A model that isn't very accurate would mistakenly warn too many healthy people that they might have the disease. Important professional resources are being wasted, and patients are being made anxious for no reason.
In contrast, memory (or sensitivity) is very important for the stages of disease (Mild, Moderate, Severe, and Proliferate DR). Testing how well the model can find all real events. A low recall is a more serious failure because it means the model misses people who do have diabetic retinopathy, which could mean they don't get the life-saving treatment they need.

To find the best model for properly identifying each class, we can use the F1 score, which takes these two concerns into account. Using standard methods, we found these metrics:

$$Accuracy = \frac{True\ Positive + True\ Negative}{True\ Positive + True\ Negative + False\ Positive + False\ Negative}$$

$$Precision = \frac{True\ Positive}{True\ Positive + False\ Negative}$$

$$Recall = \frac{True\ Positive}{True\ Positive + False\ Negative}$$

$$F1\ Score = 2 * \frac{Precision * Recall}{Precision + Recall}$$

5.2 Experimental Results

The performance of the four models is presented in Figures 13 and Figure 14 and the results are promising. Most models demonstrated exceptional performance.
Analysis of the confusion matrices (Figures 13-14) indicates that nearly all models attained a perfect score of 100% across the four stages of diabetic retinopathy. For the critical task of differentiating between a "Disease" and "No-Disease" scenario, three of the four models attained a perfect accuracy of 100%.

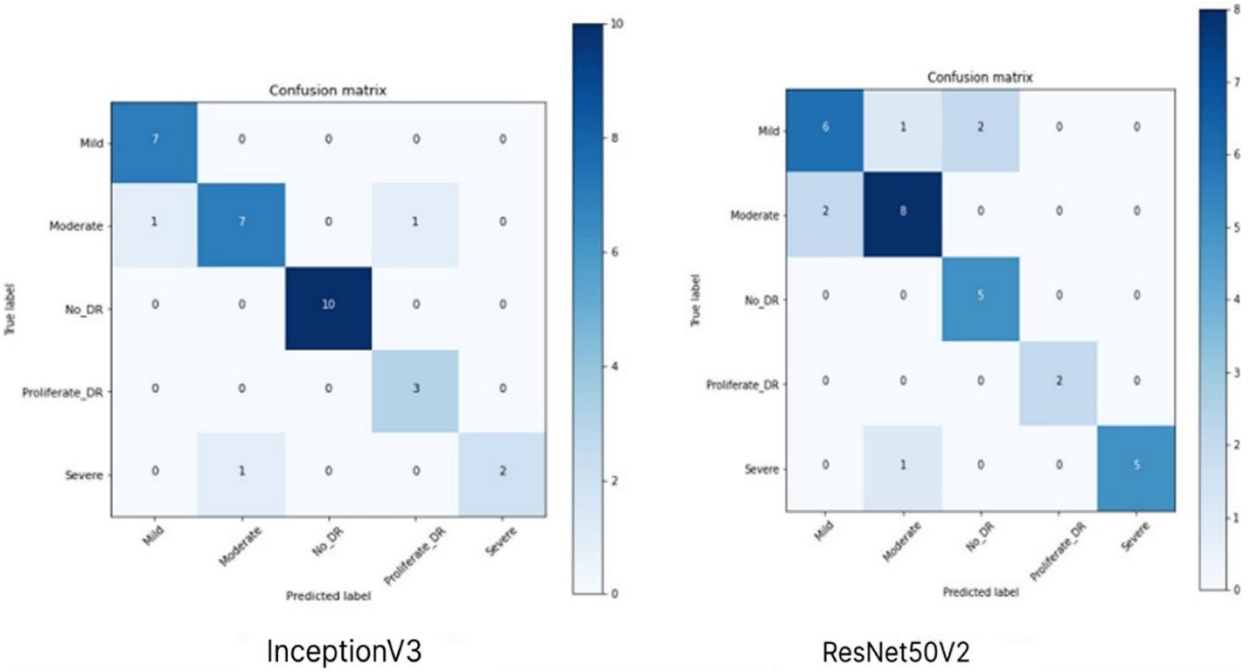

Figure 13: Confusion matrices for InceptionV3 and Resnet50V2

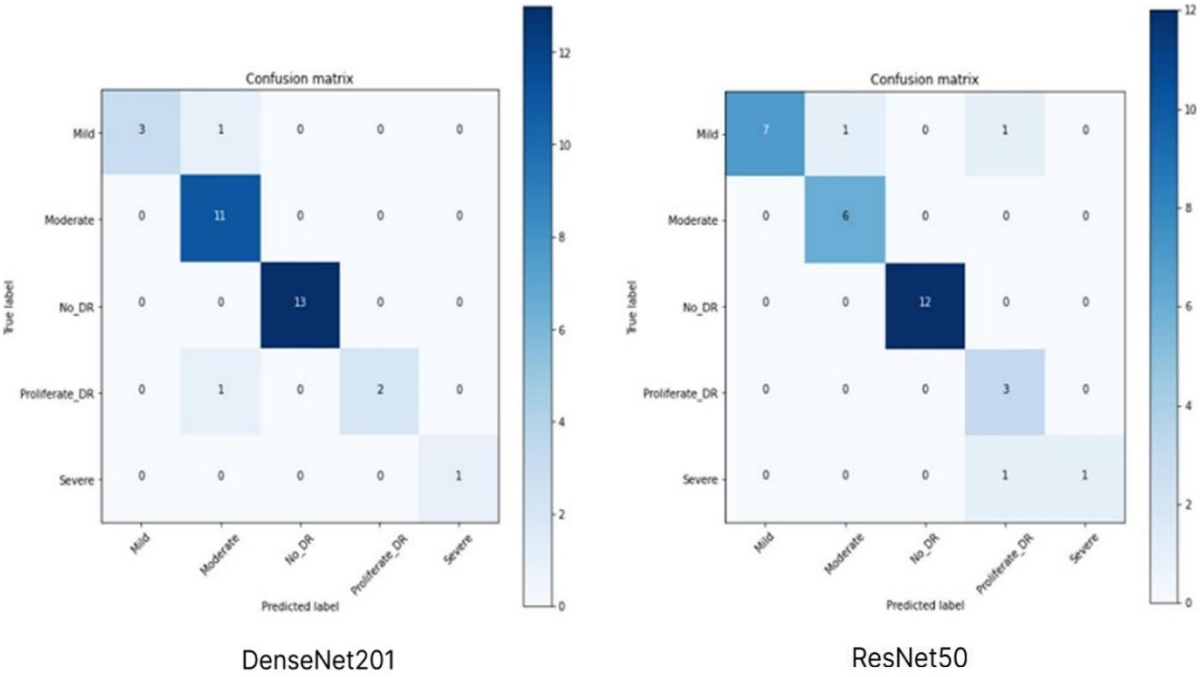

Figure 14: Confusion matrices for and Resnet50 and DenseNet201

We trained and tested all four algorithms on a dataset comprising 2,631 retinal images, employing an 80/10/10 split for training, validation, and testing, respectively. Figure 15 illustrates the very encouraging results. At 94.51%, InceptionV3 had the highest overall accuracy, followed closely by DenseNet201 at 93.20%.
More significantly, they all performed almost flawlessly on the three main clinical parameters of recall, accuracy, and F1-score.

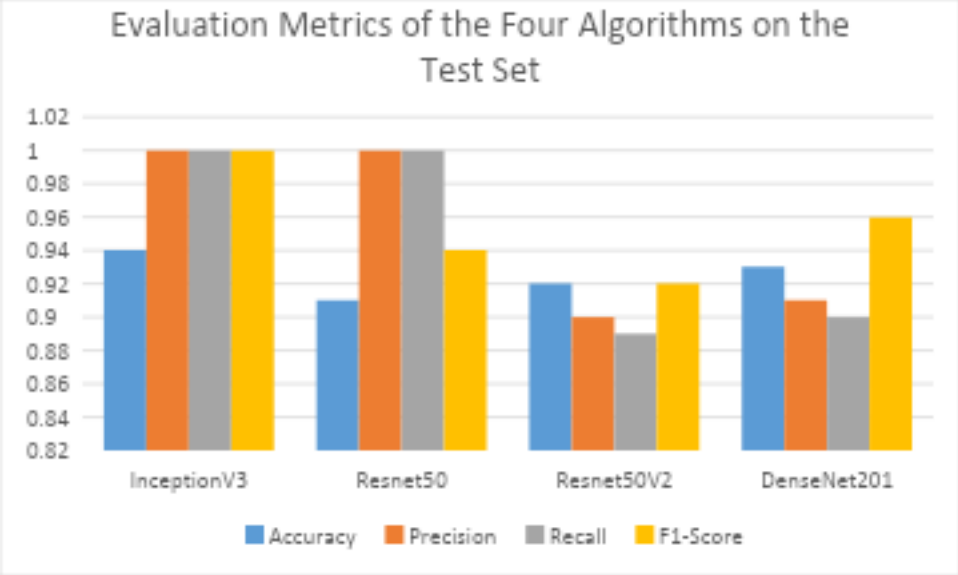

Figure 15: Evaluation Metrics of the four algorithms on the test set

For example, almost every model achieved a precision of 1.00. This is a significant practical advantage: we can be extremely certain that the model is accurate when it provides a "No-DR" diagnosis. This effectively allocates resources and, equally important, keeps healthy patients from experiencing needless stress.
Recall, or the model's capacity to identify each and every instance of the illness, is the most important statistic in this context. Remarkably, our best models, DenseNet201 and InceptionV3, obtained a flawless recall of 1.00. Significant medical implications result from this: there is almost no chance of inadvertently assuring a patient with DR that they are "safe." This guarantees that nearly all patients are discovered and able to get the necessary treatment to save their sight in a timely manner.
In conclusion, DenseNet201 and InceptionV3 both demonstrate remarkable efficacy in this challenge. They show great promise for practical clinical use since it is impossible to miss even a single case due to their exceptional accuracy, perfect recall, and precision.

6 Discussions

Our great results came from a well-thought-out process in which each step was very important. We started by carefully getting the retinal pictures ready. The choice to get rid of duplicate images from the original dataset was a big one. These cuts down on the overall number of photos, but they greatly improved the model's performance by getting rid of bias. To be able to extend to new images, the model had to learn the real features of diabetic retinopathy instead of remembering specific visual artifacts. Additionally, we used random rotations and made all of the pictures the same size (224x224 pixels) to avoid overfitting even more.

Within the classification job itself, we chose four well-known image recognition neural networks: InceptionV3, ResNet50, ResNet50V2, and DenseNet201. It was different for each one because we added a new top layer with average pooling, batch normalization, a dropout layer for regularization, and a SoftMax layer for classification. The results were very positive. All models were very accurate, with accuracy levels above 90%. The models that did the best on the test set were InceptionV3 and DenseNet201, which both got 100% to remember things. This is a very important accomplishment for a medical application because it means that the models properly identified all cases of diabetic retinopathy, lowering the frightening risk of false negatives.
Leading the group with a test accuracy of 94.51%, InceptionV3 was closely followed by DenseNet201, which scored 93.20%. Additionally, the other models also showed impressive efficiency. These outcomes are very good, especially when compared to earlier published methods that usually recorded accuracy levels between 79.50% and 87.04%.
Looking ahead, we know that accuracy alone isn't the only thing that matters for clinical use in the real world. Deployment depends on a lot of different factors, all of which are important. To fully understand each model's strengths and weaknesses, the next step will be to do a more detailed test on a variety of datasets. Excellent work is expected in the future. To improve their performance, we're going to train our models on bigger and more varied clinical samples and use advanced optimization methods. Additionally, we see a chance to use deep learning segmentation techniques to find the exact parts of the retina that are damaged by the disease. This could help us figure out how bad it is. Lastly, this diagnostic system could be changed to find other eye diseases like glaucoma or age-related macular degeneration. This would make it possible for a full AI-assisted ophthalmology platform to function.

7 Conclusions

We created a method in this study that uses pictures of the human retina to automatically identify the phases of diabetic retinopathy. Our strategy was based on sophisticated deep learning networks that we trained in a special two-step procedure after customizing them. This approach builds a final classifier that is more accurate and dependable by first tailoring the models to our particular goal and then fine-tuning them.
The outcomes are self-explanatory. Our approach continuously produced excellent results. Although all of our classifiers demonstrated remarkable accuracy, our top models—InceptionV3, ResNet50, and DenseNet201—achieved a flawless 100% recall, which is more significant. This indicates that every single instance of diabetic retinopathy was effectively detected by the system throughout our tests. Missing an actual case can have major health effects, therefore this is a critical accomplishment for a medical tool.
We think our method has a lot of promise to be a useful tool for medical professionals, allowing for quicker and more accurate diabetic retinopathy detection. It might improve medical outcomes by streamlining patient triage and cutting down on needless referrals to specialists if widely used in clinics.
Notwithstanding the great potential of these results, we acknowledge the study's limits and identify a number of fascinating avenues for further research to expand on this work

References

[1] Abadi, Martín, Paul Barham, Jianmin Chen, et al. 2016. “TensorFlow: A System for Large-Scale Machine Learning.” Version 2. Preprint, arXiv. https://doi.org/10.48550/ARXIV.1605.08695.
[2] Alwakid, Ghadah, Walaa Gouda, Mamoona Humayun, and Noor Zaman Jhanjhi. 2023. “Deep Learning-Enhanced Diabetic Retinopathy Image Classification.” DIGITAL HEALTH 9 (January): 20552076231194942. https://doi.org/10.1177/20552076231194942.
[3] Bodapati, Jyostna Devi, Nagur Shareef Shaik, and Veeranjaneyulu Naralasetti. 2021. “Deep Convolution Feature Aggregation: An Application to Diabetic Retinopathy Severity Level Prediction.” Signal, Image and Video Processing 15 (5): 923–30. https://doi.org/10.1007/s11760-020-01816-y.
[4] Bressler, Neil M. 2004. “Age-Related Macular Degeneration Is the Leading Cause of Blindness . . .” JAMA 291 (15): 1900. https://doi.org/10.1001/jama.291.15.1900.

[5] Cao, Juan, Jiaran Chen, Xinying Zhang, and Yang Peng. 2025. "Diabetic Retinopathy Classification Based on Dense Connectivity and Asymmetric Convolutional Neural Network." Neural Computing and Applications 37 (11): 7527–40. https://doi.org/10.1007/s00521-022-07952-5.

[6] Cui, Hainan, Shuhan Shen, Wei Gao, Hongmin Liu, and Zhiheng Wang. 2019. "Efficient and Robust Large-Scale Structure-from-Motion via Track Selection and Camera Prioritization." ISPRS Journal of Photogrammetry and Remote Sensing 156 (October): 202–14. https://doi.org/10.1016/j.isprsjprs.2019.08.005.

[7] Diabetic Retinopathy 224 x 224 (2019 Data). | Kaggle : https:// www.kaggle.com/datasets/sovitrath/diabetic-retinopathy-224x 224-2019-data 34.

[8] "G. Van Rossum; F.L. Drake Jr. Python Tutorial; Centrum Voor Wiskunde En Informatica: Amsterdam, The Netherlands, 1995." n.d.

[9] Gargeya, Rishab, and Theodore Leng. 2017. "Automated Identification of Diabetic Retinopathy Using Deep Learning." Ophthalmology 124 (7): 962–69. https://doi.org/10.1016/j.ophtha.2017.02.008.

[10] Gour, Neha, and Pritee Khanna. 2021. "Multi-Class Multi-Label Ophthalmological Disease Detection Using Transfer Learning Based Convolutional Neural Network." Biomedical Signal Processing and Control 66 (April): 102329. https://doi.org/10.1016/j.bspc.2020.102329.

[11] Gulshan, Varun, Lily Peng, Marc Coram, et al. 2016. "Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs." JAMA 316 (22): 2402. https://doi.org/10.1001/jama.2016.17216.

[12] Han, Baoru, Jinglong Du, Yuanyuan Jia, and Huazheng Zhu. 2021. "Zero-Watermarking Algorithm for Medical Image Based on VGG19 Deep Convolution Neural Network." Journal of Healthcare Engineering 2021 (July): 1–12. https://doi.org/10.1155/2021/5551520.

[13] Jonas, Jost B., Rupert R.A. Bourne, Richard A. White, et al. 2014. "Visual Impairment and Blindness Due to Macular Diseases Globally: A Systematic Review and Meta-Analysis." American Journal of Ophthalmology 158 (4): 808–15. https://doi.org/10.1016/j.ajo.2014.06.012.

[14] Kasim, Ömer. 2023. "Ensemble Classification Based Optimized Transfer Learning Feature Method for Early Stage Diagnosis of Diabetic Retinopathy." Journal of Ambient Intelligence and Humanized Computing 14 (8): 11337–48. https://doi.org/10.1007/s12652-023-04648-z.

[15] Kurup, Greeshma, J. Angel Arul Jothi, and Anusree Kanadath. 2021. "Diabetic Retinopathy Detection and Classification Using Pretrained Inception-V3." 2021 International Conference on Smart Generation Computing, Communication and Networking (SMART GENCON), October 29, 1–6. https://doi.org/10.1109/SMARTGENCON51891.2021.9645914.

[16] Leasher, Janet L., Rupert R.A. Bourne, Seth R. Flaxman, et al. 2016. "Global Estimates on the Number of People Blind or Visually Impaired by Diabetic Retinopathy: A Meta-Analysis From 1990 to 2010." Diabetes Care 39 (9): 1643–49. https://doi.org/10.2337/dc15-2171.

[17] Li, Xiaogang, Tiantian Pang, Biao Xiong, Weixiang Liu, Ping Liang, and Tianfu Wang. 2017. "Convolutional Neural Networks Based Transfer Learning for Diabetic Retinopathy Fundus Image Classification." 2017 10th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI), October, 1–11. https://doi.org/10.1109/CISP-BMEI.2017.8301998.

[18] Menaouer, Brahami, Zoulikha Dermane, Nour El Houda Kebir, and Nada Matta. 2022. "Diabetic Retinopathy Classification Using Hybrid Deep Learning Approach." SN Computer Science 3 (5): 357. https://doi.org/10.1007/s42979-022-01240-8.

[19] Molodij, G., E.N. Ribak, M. Glanc, and G. Chenegros. 2014. "Enhancing Retinal Images by Extracting Structural Information." Optics Communications 313 (February): 321–28. https://doi.org/10.1016/j.optcom.2013.10.011.

[20] Peli, E., and T. Peli. 1989. "Restoration of Retinal Images Obtained through Cataracts." IEEE Transactions on Medical Imaging 8 (4): 401–6. https://doi.org/10.1109/42.41493.

[21] Pires Dias, João Miguel, Carlos Manta Oliveira, and Luís A. Da Silva Cruz. 2014. "Retinal Image Quality

Assessment Using Generic Image Quality Indicators." Information Fusion 19 (September): 73–90. https://doi.org/10.1016/j.inffus.2012.08.001.

[22] Poostchi, Mahdieh, Kamolrat Silamut, Richard J. Maude, Stefan Jaeger, and George Thoma. 2018. "Image Analysis and Machine Learning for Detecting Malaria." Translational Research 194 (April): 36–55. https://doi.org/10.1016/j.trsl.2017.12.004.

[23] Pratt, Harry, Frans Coenen, Deborah M. Broadbent, Simon P. Harding, and Yalin Zheng. 2016. "Convolutional Neural Networks for Diabetic Retinopathy." Procedia Computer Science 90: 200–205. https://doi.org/10.1016/j.procs.2016.07.014.

[24] Qassim, Hussam, David Feinzimer, and Abhishek Verma. 2017. "Residual Squeeze VGG16." Version 1. Preprint, arXiv. https://doi.org/10.48550/ARXIV.1705.03004.

[25] Sandler, Mark, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. 2018. "MobileNetV2: Inverted Residuals and Linear Bottlenecks." 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, June, 4510–20. https://doi.org/10.1109/CVPR.2018.00474.

[26] Song, Wenai, Ying Cao, Zhiqiang Qiao, Qing Wang, and Ji-Jiang Yang. 2019. "An Improved Semi-Supervised Learning Method on Cataract Fundus Image Classification." 2019 IEEE 43rd Annual Computer Software and Applications Conference (COMPSAC), July, 362–67. https://doi.org/10.1109/COMPSAC.2019.10233.

[27] Soomro, Toufique Ahmed, Ahmed J. Afifi, Lihong Zheng, et al. 2019. "Deep Learning Models for Retinal Blood Vessels Segmentation: A Review." IEEE Access 7: 71696–717. https://doi.org/10.1109/ACCESS.2019.2920616.

[28] Kaur, G., Adhikari, N., Krishnapriya, S., Wawale, S. G., Malik, R. Q., Zamani, A. S., ... & Osei-Owusu, J. (2023). Recent advancements in deep learning frameworks for precision fish farming opportunities, challenges, and applications. Journal of Food Quality, 2023(1), 4399512.

[29] Mengash, Hanan Abdullah, Mohammad Alamgeer, Mashael Maashi, Mahmoud Othman, Manar Ahmed Hamza, Sara Saadeldeen Ibrahim, Abu Sarwar Zamani, and Ishfaq Yaseen. "Leveraging marine predators algorithm with deep learning for lung and colon cancer diagnosis." Cancers 15, no. 5 (2023): 1591.

[30] Mustafa Hilal, Anwer, Jaber S. Alzahrani, Ibrahim Abunadi, Nadhem Nemri, Fahd N. Al-Wesabi, Abdelwahed Motwakel, Ishfaq Yaseen, and Abu Sarwar Zamani. "Intelligent Deep Learning Model for Privacy Preserving IIoT on 6G Environment." Computers, Materials & Continua 72, no. 1 (2022).

[31] Obayya, Marwa, Nadhem Nemri, Mohamed K. Nour, Mesfer Al Duhayyim, Heba Mohsen, Mohammed Rizwanullah, Abu Sarwar Zamani, and Abdelwahed Motwakel. "Explainable artificial intelligence enabled TeleOphthalmology for diabetic retinopathy grading and classification." Applied Sciences 12, no. 17 (2022): 8749.