Diabetic Retinopathy Detection Using Meta Learning and Deep Learning Techniques

Muhammad Ammar Khan[1] , Ali Okatan[2]

[1]*Department of AI & Data Science, Istanbul Aydin University, Istanbul, mamarkhan@stu.aydin.edutr*

[2]*Prof. Dr. at Department of AI & Data Science, Istanbul Aydin University, Istanbul, aliokatan@aydin.edu.tr ORCID: 0000-0002-8893-9711*

ABSTRACT

&In the world of ocular health, diabetic retinopathy is a common condition that, if not recognized and treated promptly, can cause vision loss. In this study we present a meta learning stacking approach for the diagnosis and referral of ocular defects. Our approach demonstrates exceptional efficacy in detecting uncommon conditions by utilizing a combination of pre-trained convolutional neural networks (CNNs) and stacking meta-learning techniques. This novel approach improves the accuracy of results while significantly reducing the time required compared to conventional deep learning methods. The method demonstrates the promise of stacking meta learning in addressing data scarcity and improving early diagnosis of sight-threatening diseases by achieving an outstanding accuracy of 93%. Additionally, the solution beats problems brought on by a lack of readily available data which needs to be preprocessed. When compared to other deep learning models frequently used in ocular abnormality detection. These findings underscore the potential impact of our approach as an advanced computer-aided diagnosis tool for ocular anomalies, paving the way for significant advancements in the field. These valuable insights provide a solid foundation for future research, driving innovation and progress in computer-aided diagnosis tools for ocular health.&

Keywords: *Meta learning; Deep Learning; Retinopathy Detection; Diabetic Reinopathy*

1. INTRODUCTION

Diabetic retinopathy (DR) is a prevalent and serious complication of Diabetes Mellitus, characterized by progressive vascular disruptions in the retina caused by chronic hyperglycemia [1]. It has become a significant global cause of visual impairment, impacting a staggering 93 million individuals worldwide and often leading to blindness among working-age adults. Timely detection and appropriate management of DR are crucial in preventing vision loss, making regular eye screening for diabetic patients essential. The manual diagnosis of DR through fundus examination by skilled ophthalmologists poses challenges in terms of expertise, time, and accessibility, particularly in densely populated or remote areas with limited access to specialized healthcare professionals.

To address these challenges, researchers have increasingly relied on artificial intelligence (AI) and deep learning (DL) methodologies to create computer-aided diagnosis systems for diabetic retinopathy (DR). DL models, empowered by deep neural networks, have shown promising capabilities in analyzing fundus images for automated DR detection, classification, and severity grading [2]. These models can accurately identify retinal lesions and assess the progression of DR, providing valuable insights for timely intervention and treatment. Moreover, DL-based models offer the potential to bridge the gap in access to specialized eye care, particularly in regions with a shortage of ophthalmologists. In this journal paper, we present an integrated review of meta-learning based approach for DR detection and classification. Our aim is to provide a comprehensive overview of the advancements, challenges, and future prospects in this rapidly evolving research area. Specifically, we will examine the technical implementations, dataset characteristics, preprocessing techniques, and model performance of published studies in the field. By analyzing the complete data analysis pipeline, including dataset selection, preprocessing, DL model development, and real-world implementation, we aim to offer valuable insights into the current state and potential of meta-learning in diabetic retinopathy diagnosis. By objectively analyzing the benefits and drawbacks of various DL models used to study diabetic retinopathy, we want to provide insight on their relative merits and clinical applicability. Additionally, we will discuss the ethical considerations and interpretability of DL models in addressing the challenges posed by DR diagnosis, emphasizing the importance of fairness, transparency, and robustness in these AI-powered systems. By providing a comprehensive analysis of DL-based approaches for diabetic retinopathy, this journal paper aims to contribute to the advancement of meta-learning research in the field and support the development of accurate, efficient, and accessible diagnostic tools. Ultimately, our goal is to enhance the early detection and management of diabetic retinopathy, thereby reducing the burden of vision loss associated with this prevalent complication of diabetes.

2. RELATED WORK

Indeed, artificial intelligence technologies have greatly advanced the field of identifying diabetic eye disorders. Sarki et al., with its fine-tuned VGG16 model [3] and Nazir et al., The DenseNet-100 model along with a custom CenterNet model made significant progress by successfully using convolutional neural networks (CNNs) to analyze retinal pictures [4]. Their improved models and deep learning techniques significantly improved the detection and classification of glaucoma and diabetic eye disorders. Another important development in this area was a study that used meta-learning CNNs to identify diabetic retinopathy [5]. They suggested data resampling techniques for the selection of the training and validation set, then trained their CNNs using the chosen data. Additionally, they experimented with two ensemble learning strategies: meta-learner and unweighted average, and used straightforward data augmentation techniques to diversify the training data patterns. The meta-learner strategy demonstrated greater accuracy, and the ensemble methods regularly outperformed single CNNs. This investigation highlights how cutting-edge methods like meta-learning CNNs may improve the precision of diabetic retinopathy diagnosis.

3. DIABETIC RETINOPATHY

During the early stages of diabetic retinopathy, specific manifestations on the retina can be observed, indicating disease progression. Microaneurysms are one of the initial signs, this occurs as a consequence of pericytes' degeneration and loss, causing the dilation of capillary walls. Intraretinal hemorrhages occur when the walls of capillaries or micro aneurysms rupture. Non-proliferative diabetic retinopathy (NPDR) includes additional abnormalities like soft and hard exudates, intraretinal microvascular abnormalities (IRMA), venous beading, and venous loops or reduplication. IRMAs, characterized as large-caliber tortuous vessels, often emerge in areas of ischemia as a response to attempted vascular remodeling. Neovascularization, characterized by the growth of new retinal vessels due to ischemia, is also a notable manifestation, distinguishing proliferative diabetic retinopathy (PDR). Fig. 1 presents an illustrative fundus image displaying these indicative lesions [6].

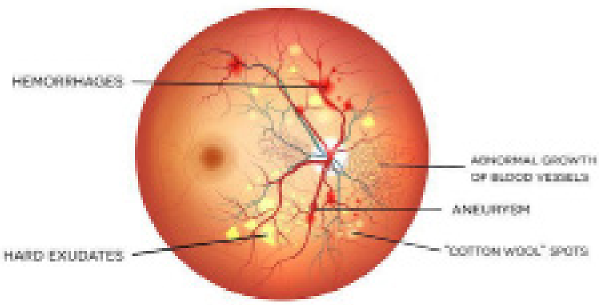

Figure 1. Diabetic retinopathy impacts on retinal blood vessels

At any stage of diabetic retinopathy, the occurrence of diabetic macular edema (DME) can significantly contribute to visual impairment. Notable abnormalities associated with DME include exudates within one disc diameter of the fovea's center, exudates within the macula, retinal thickening within one disc diameter of the fovea's center, and the presence of microaneurysms or hemorrhages within the same region. Different grading protocols have been devised to evaluate the clinical severity of diabetic retinopathy.Although the Early Treatment Diabetic Retinopathy Study (ETDRS) grading system is regarded as the gold standard, it has proven difficult to implement in everyday clinical practice. Alternative severity scales have been suggested as a result in various nations to improve patient screening and streamline communication among healthcare professionals. One such scale is the International Clinical Diabetic Retinopathy Disease Severity Scale as shown in Table 1, which classifies diabetic retinopathy into five severity levels based on specific findings observed during dilated ophthalmoscopy. These findings include the presence of microaneurysms, intraretinal hemorrhages, venous beading, intraretinal microvascular abnormalities (IRMA), and neovascularization [7].

Table 1. International Clinical Diabetic Retinopathy Disease Severity Scale

Stage	Dilated Ophthalmoscopy Observable Findings	Severity
I	No abnormalities	No DR
II	Micro-aneurysms only	Mild non-proliferative DR
III	Any of the following: - micro-aneurysms - retinal dot and blot haemorrhages - hard exudates or cotton wool spots No signs of severe non-proliferative diabetic retinopathy	Moderate non-proliferative DR
IV	Any of the following: - more than 20 intra-retinal hemorrhages in each of 4 quadrants - definite venous beading in 2 or more quadrants - prominent intra-retinal microvascular abnormality (IRMA) in 1 or more quadrants No signs of proliferative retinopathy	Severe non-proliferative DR
V	One or both of the following: - Neovascularization - Vitreous/pre-retinal hemorrhage	Proliferative DR

4. RESEARCH AND METHODOLOGY

4.1. Dataset

The Aravind Eye Hospital in rural India provided the data for the Kaggle APTOS 2019 Competition [8]. The goal was to provide reliable diagnostic tools for automated Diabetic Retinopathy Detection, giving the hospital the ability to quickly locate affected patients. The dataset consists of 5,590 DR images, which is the third-largest of its sort. It is important to note, however, that one drawback is the existence of a considerable class imbalance,particularly in the Severe NPDR class, which has only 193 photos. Aravind Eye Hospital technicians traveled to the rural areas of the country to capture the images and that's the main concern that APTOS dataset demonstrates variances brought on by various camera settings applied across many centers, much like the Kaggle EyePACS dataset does. As they were gathered in a real-world clinical scenario, the labels and data may also include noise like artifacts, focus issues, and differences in exposure (over or under) levels. The dataset is divided in 5 categories i.e 0-No DR 1-Mild 2-Moderate 3-Severe 4-Proliferative DR.

Figure 2. Classification of APTOS-2019 dataset

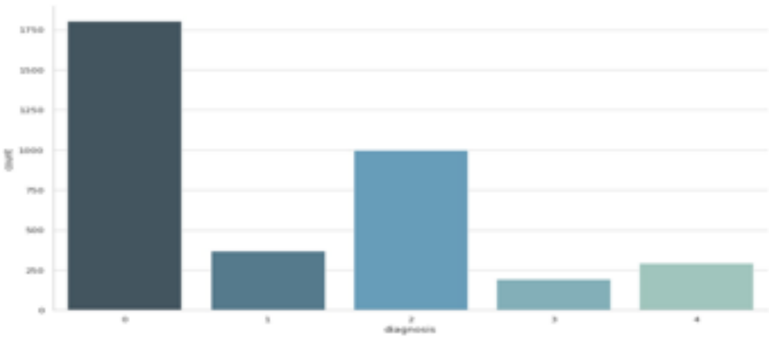

4.2. Deep learning

Deep learning is a branch of machine learning that utilizes neural networks with multiple layers to extract intricate patterns and representations from data. Inspired by the structure of the human brain, deep learning techniques rely on artificial neural networks to process and analyze complex information, enabling the model to learn and make predictions based on the extracted features. It involves automatically learning the mathematical representation of the underlying relationships within data, uncovering latent and intrinsic patterns in an automatic manner. Deep learning models easily pick up complex features by using a lot of labeled data andeasily pick up complex features by using a lot of labeled data and modifying network weights through back propagation,leading to significant advancements in artificial intelligence.

4.3. Convolutional Neural Network

With inspiration from human vision, convolutional neural networks (CNNs) are neural networks that process 2D arrays as input. They use convolution, a mathematical procedure, to take important details out of images. CNNs use filters or kernels to compute convolutions across the input data, resulting in feature maps, in contrast to traditional Deep Neural Networks (DNNs), where all neurons at a given layer contribute to the output of every neuron at the next layer. A window of pixels determined by the filter's size affects each unit in the feature map. The area in the input space that a specific CNN feature takes into account is determined by this receptive field. The convolutional part of the network is responsible for feature extraction, while the subsequent layers, often a DNN, perform classification based on the learned features.

4.4. Transfer Learning

Transfer learning describes the method of using the knowledge gained by using pre-trained models on one task to enhance performance on a related but distinct task. The goal of transfer learning is to move the learnt parameters or representations from the source task to the target task. The pre-trained model collects common patterns and features that are helpful for a variety of tasks and is often trained on a big dataset.The pre-trained model can be applied to new data more successfully and possiblyperform better than when learned from the start by adjusting or tuning it for the desired task. Transfer learning becomes very beneficial when the intended task has few labeled data points or when the target task and source task have similar properties. It aims to leverage knowledge from pre-trained models to improve performance on a related task.

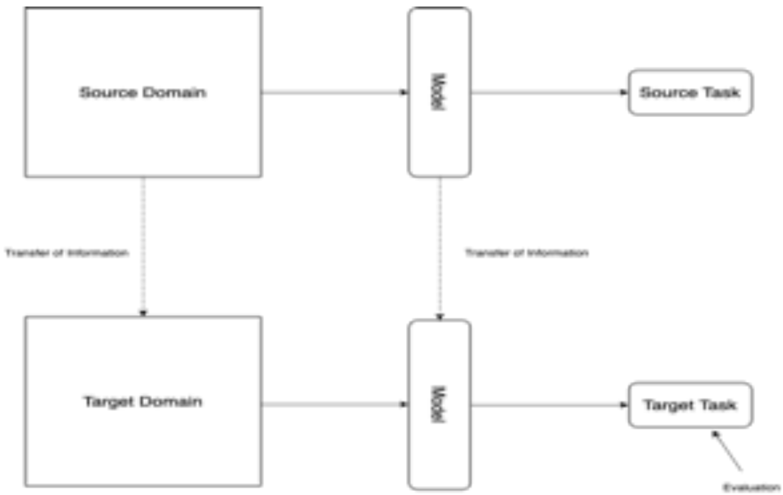

Figure 3. Basic structure of transfer learning

4.5. Meta Learning

Meta-learning, also known as learning to learn, focuses on developing algorithms or models that can learn how to learn or adapt quickly to new tasks. Meta-learning involves training a model on a distribution of tasks, where the model learns generalizable knowledge or prior knowledge that enables rapid learning on new, unseen tasks. The objective is to acquire knowledge or strategies that can be applied to new tasks, allowing the model to learn with fewer examples or adapt quickly to new scenarios. Meta-learning can involve techniques such as few-shot learning, where a model is trained to learn from only a few examples, or meta-optimization, where the model learns to optimize its own learning process.

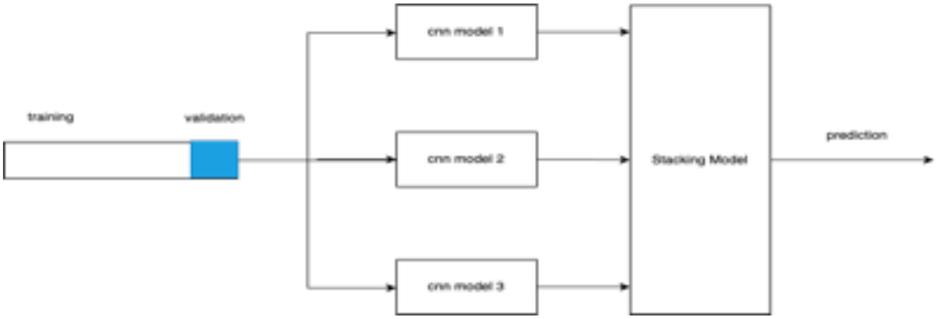

Figure 4. Basic structure of meta learning

5. IMPLEMENTATION DETAILS

5.1. Base Models

For training the stacking model in ensemble learning we first need to train and prepare the base models. After the model is trained on the dataset we use the outputs of these individual base models as the input of the stacking model. We have used only three base models for this experiment.

5.1.1. EfficientNet

EfficientNet models aim to achieve high accuracy while still using computational resources efficiently. Efficient Net-B5 has been used as the first base model. The scalable architecture of EfficientNet models makes it simple to change the model size in accordance with the resources at accessibility or your specific requirements [9].

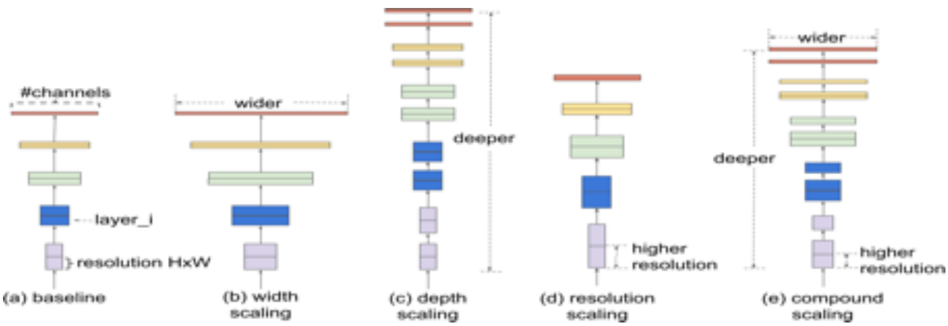

Figure 5. Basic architecture of EfficientNet

5.1.2. ResNet

ResNet models make use of residual connections to help solve the issue of training with vanishing gradients. Resnet-18 is used for the following implementation. This makes it possible to train deeper networks efficiently and raises the performance of the model as a whole. In comparison to deeper variants, ResNet18 is a relatively light network, which speeds up training and uses less computing power [9].

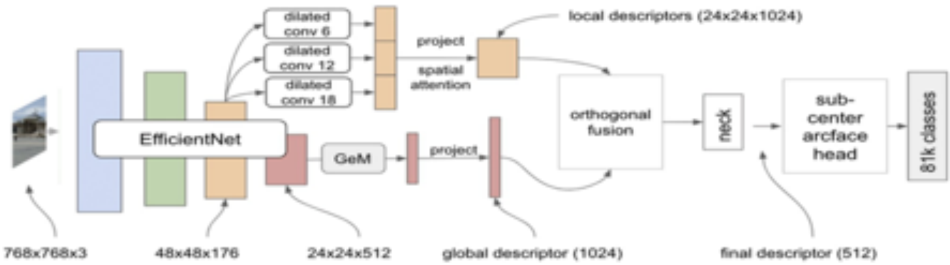

Figure 6. Basic architecture of ResNet

5.1.3. DenseNet

Each layer in a DenseNet model receives feature maps from all layers above it due to the dense connection used in these models. We have used Densenet-121 for this implementation. Better feature representations are the result of the network's dense connection, which stimulates feature reuse, enhances gradients, and encourages information flow. By promoting feature reuse and facilitating better gradient flow, DenseNet's dense connection helps prevent overfitting, which may lead to enhanced generalization [10].

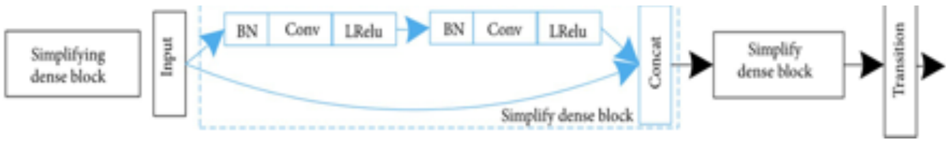

Figure 7. Basic architecture of DenseNet

5.2. Model Ensemble and Stacking

In this study, the challenge of predicting diabetic retinopathy was dealt through utilizing an ensemble of multiple deep learning models. ResNet18, DenseNet121, and EfficientNet-B5 are some of the models employed in the ensemble. By changing the final layers to correspond to the number of classes in our dataset, the mentioned three models which were originally trained on the ImageNet dataset were customized for our specific goal. Stacking is an approach which involves combination of various machine learning models in hopes of enhancing predictive performance is the core idea behind our ensemble method. Stacking works by teaching a meta-model to base its predictions on those of a number of base models. In our scenario, a new model known as the "stacking model" was created by concatenating the outputs from these three base models.

The stacking model used in this study was a straightforward linear model. A 15-element vector is created by concatenating the outputs of the base models, each of which generates a prediction vector of length 5 (corresponding to the 5 classes in the dataset). The model is able to learn the best combination of the predictions from the basic model by passing this vector through a linear layer with five outputs. By employing this stacking strategy, the model is able to fully capitalize on the advantages of each base model, possibly leading to a more reliable and precise predictive model. The stacking model is trained using the Adam optimizer and a common loss function (in our instance, Cross Entropy Loss).

5.3. Training and Validation Split

To guarantee that the models could be properly trained and evaluated in this research, the dataset was divided into training and validation subsets. By dividing the dataset in this form, the models can be trained on one part of the data (the training set), and their performance can then be assessed on a second subset of the data (the validation set), which they were not exposed to during training. In a real-world application, this helps in estimating how well the trained models will respond to new data. This was done using pandas' sample function, which takes a sample of data at random, according to the provided code. 80% of the data was chosen for the training set since the frac parameter was set to 0.8.

The random_state parameter was set to a specific value to ensure the reproducibility of the results. It ensures that each time the code is executed, the random selection of data will be the same. After choosing the training set, the validation set was made by removing its indices from the entire dataset, thereby erasing the 20% of data that wasn't part of the training set. This type of split was used for this study on the detection of diabetic retinopathy since it is frequently used in machine learning and deep learning projects to provide a balanced approach to training and testing models.

5.4. Image Preprocessing

In the world of deep learning and computer vision, image preprocessing is an important and crucial step. It puts a huge impact on the model's performance. Using PyTorch's transformations module, a set of common preprocessing procedures were applied to the images for this project. Before being input into the models, the photographs were initially downsized to 224x224 pixels in order to ensure that they were all the same size. The photos were then transformed into tensors, the necessary input format for PyTorch models. Additionally, the photos were standardized using a particular mean and standard deviation. These parameters were used to pre-train the models on the ImageNet dataset and are typical for images in the RGB format. This normalization step is essential since it speeds up the model's convergence during training and can improve performance.

5.5. Training Process

Standard deep learning techniques were used to handle the training of the models, including the base models and the stacking model. The Adam optimization technique was used for optimizing the parameters of the models. This adaptive learning rate method uses little memory and is computationally effective. Due to its effectiveness, it has been widely used in deep learning. A common approach for multi-class problems with classification, the Cross-Entropy Loss, was used to train the models. The variation between the expected probability and the actual labels is measured by the Cross-Entropy Loss. As a result, the model is recommended to output low probability for the incorrect classes and high probabilities for the correct classes.

The main models were initially trained separately. The stacking model was then trained using the Adam optimization algorithm and Cross-Entropy Loss, and its data inputs were concatenated and used in this process. The stacking model learns to correct for errors made by the base models by training it to base its predictions on their outputs, which will enhance the accuracy of all predictions.

6. RESULTS AND DISCUSSION

6.1. Base Model Performance

In the case of the individual base models, they were evaluated using the accuracy metric, which is the proportion of correct predictions out of the total predictions. For the EfficientNet-B5 model, the accuracy on the validation set was approximately 82.24%. The second model, ResNet18, achieved an accuracy of about 78.96% on the validation set. The third model, DenseNet121, performed comparably to the first model, with an accuracy of about 81.28% on the validation set. In general, these results suggest that all three base models were able to learn useful representations from the training data and generalize well to unseen data in the validation set. However, there were variations in their performances, and this further justifies the use of model ensemble and stacking techniques to leverage their complementary strengths and improve the overall system's

6.2. Stacking Model Training

The base models (efficientnet-b5, resnet-18, densenet-121) predictions were used to train the stacking model. The stacking model's input was a perfomance. combination of the predictions from these models. The typical method of using the original features as input (in this case, the image data) is different from this. The stacking model was trained using a standard supervised learning procedure. Binary Cross-Entropy with Logits Loss, a loss function appropriate for binary classification issues, was employed. Adam was used for training the model as its the widely used optimizer algorithm because of its effectiveness. Every 20 epochs, the loss was recorded during the 500 epochs of training in order to track its progress. The model parameters were preserved for later usage after the training.

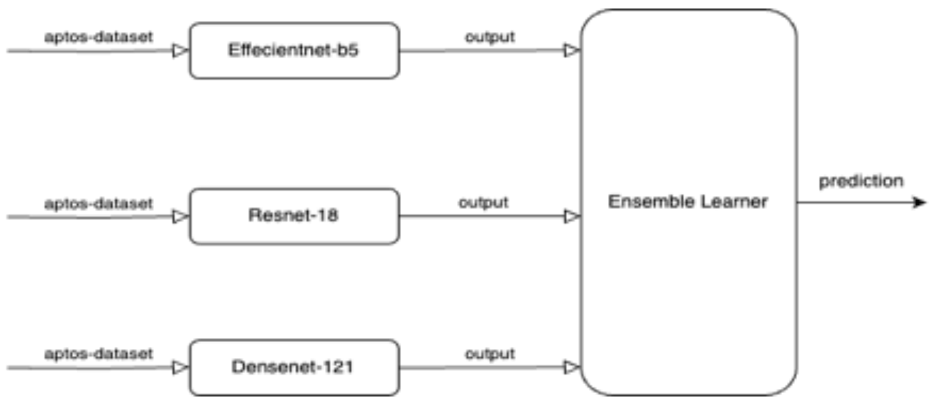

Figure 8. Base models as input for training Stacking Model

6.3. Stacking Model Performance

Multiple metrics were used to evaluate the stacking model's performance on the validation set. On the validation images, the stacking model's accuracy was 92.76%. This indicates that for approximately 93% of the photos in the validation set, the stacking model provided accurate predictions. The precision and recall score, known as the F1 score, was approximately 0.91. This shows that a decent balance between accurately detecting positive cases (recall) and reducing false positives (precision) was maintained by the stacking model.

Table 2. Showing the matrics scores of stacking model

Matrics	f1	accuracy	precision	recall	roc_auc
Score	0.92	0.90	0.90	0.91	0.92

Additionally, the precision score was approximately 0.90, indicating that the model correctly predicted positive cases 90% of the time. Recall was approximately 0.92, indicating that 92% of the positive cases in the validation set were properly identified by the model. Finally, the model's ROC-AUC score was close to 0.93, showing that it is very capable of differentiating between the classes. A high score on the ROC-AUC measure denotes a strong model in binary classification problems. These metrics indicate the effectiveness of the chosen ensemble and stacking technique overall, indicating that the stacking model performed well on the validation set

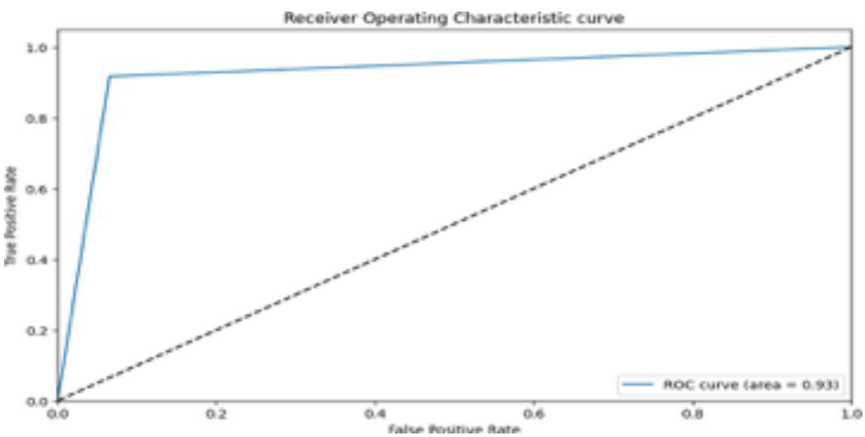

Figure 9. Showing the ROC curve area of the stacking model

7.EXPERIMENTS

7.1.Few-Shot Learning with APTOS-2019

We used few-shot learning as an experimental strategy with the APTOS-2019 dataset. This method seeks to develop models that can efficiently learn knowledge from a limited amount of data. Due to the APTOS-2019 dataset's unbalanced class distribution, the results, however, were not satisfactory. An unbalanced dataset can have a major impact on the effectiveness of a machine learning model, especially when utilizing a few-shot learning strategy that mainly relies on a tiny sample of data.

7.2.Classification Using Meta Learning Approach

Moving forward, we used the meta learning approach on the APTOS-2019 dataset for classification into multiple classes rather than binary classification. The performance of the stacking model on validation images was as follows:

Matrics	f1	accuracy	precision	recall
Score	0.833	0.833	0.835	0.833

The ROC-AUC scores were varying significantly across different classes:

Classes	Class 0	Class 1	Class 2	Class 3	Class 4
Score	0.978	0.726	0.874	0.709	0.716

The results suggest an acceptable level of performance, but there is room for improvement.

7.3.Binary Classification

After evaluating the multi-class model's performance, we modified the strategy for binary classification. This modification was developed in light of the difficulties that multi-class classification presents when dealing with imbalanced data. As a result, the findings significantly improved and the accuracy of the stacking model exceeded 90%. The experiments have demonstrated that although meta-learning and few-shot learning algorithms can be used to diagnose diabetic retinopathy, careful consideration of the class distribution in the dataset is necessary to ensure the best model performance.

8. CONCLUSIONS

8.1. Summary of Findings

This study employed deep learning techniques to classify the severity of Diabetic Retinopathy, a serious eye condition common in diabetic patients. Specifically, the study used an ensemble of pretrained convolutional neural networks (EfficientNet-B5, ResNet18, and DenseNet121) and a stacking approach to improve model performance. The base models achieved individual validation accuracies around 78% to 82%. However, the stacking model achieved a validation accuracy of approximately 93%, showcasing the effectiveness of the ensemble and stacking method. The stacking model also achieved impressive scores on other performance metrics, including precision, recall, F1-score, and ROC-AUC score, indicating its robustness in classifying the DR condition. These findings suggest that ensemble learning and stacking methods can be powerful tools in the context of medical image classification tasks, particularly when the goal is to leverage the strengths of multiple models to improve overall performance.

8.2. Potential Improvements and Future Work

Despite the promising results, there are several potential areas of improvement and directions for future work. Fine-tuning the base models on the specific task could potentially improve their individual performances and, by extension, the performance of the ensemble model. Additionally, exploring different architectures for the stacking model could lead to improved performance. Hyperparameter tuning, such as adjusting the learning rate or the number of epochs, could also yield better results. Another promising direction for future work could involve expanding the ensemble to include more diverse base models. The current study utilized models based on CNNs; including models based on different architectures, such as transformer models, could potentially increase the robustness of the ensemble.

8.3.Implications of the Study

This study has several implications, especially for the application of deep learning in the medical field. First off, it shows how deep learning models could be used to help with the identification of diseases like diabetic retinopathy. Accurate early detection can significantly enhance a patient's prognosis, resulting in better medical outcomes. Additionally, the application of ensemble and stacking methods demonstrates how we can capitalize on the complementing qualities of diverse models to enhance overall performance, a strategy that can be used in a variety of other healthcare situations. The study additionally emphasizes the value of ongoing model validation and evaluation to make sure the model works effectively on both training and real-world data. This is essential in the medical field because misdiagnosis may be quite expensive and the stakes are so high.

REFERENCES

[1]Atwany, M.Z.; Sahyoun, A.H.; Yaqub, M. Deep learning techniques for diabetic retinopathy classification: A survey. IEEE Access 2022, 10, 28642–28655.w

[2]Amin, J.; Sharif, M.; Yasmin, M. A review on recent developments for detection of diabetic retinopathy. Scientifica 2016, 2016, 6838976.

[3]Sarki R, Ahmed K, Wang H, Zhang Y. Automated detection of mild and multi-class diabetic eye diseases using deep learning. Health Inf Sci Syst. 2020 Oct 8;8(1):32. doi: 10.1007/s13755-020-00125-5. PMID: 33088488; PMCID: PMC7544802.

[4]T. Nazir, M. Nawaz, J. Rashid, R. Mahum, M. Masood, A. Mehmood, F. Ali, J. Kim, H.-Y. Kwon and A. Hussain, Detection of diabetic eye disease from retinal images using a deep learning based CenterNet model, Sensors, vol.21, no.16, DOI: 10.3390/s21165283, 2021.

[5]Enkvetchakul, Prem & Surinta, Olarik & Noppitak, Sangdaow. (2022). Effective Data Resampling and Meta-learning Convolutional Neural Networks for Diabetic Retinopathy recognition. ICIC Express Letters, Part B: Applications. 13. 939-948. 10.24507/icicelb.13.09.939.

[6]https://www.eyeops.com/contents/our-services/eye-diseases/diabetic-retinopathy.

[7]https://www.researchgate.net/publication/326854068_Diabetic_Retinopathy_Stage_Classification_Using_Convo lutional_Neural_Networks.

[8]https://www.kaggle.com/c/aptos2019-blindness-detection.

[9]Tan, M. & Le, Q.. (2019). EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. Proceedings of the 36th International Conference on Machine Learning, in Proceedings of Machine Learning Research 97:6105-6114 Available from https://proceedings.mlr.press/v97/tan19a.html.

[10]Gao Huang, Zhuang Liu, Laurens van der Maaten, Kilian Q. Weinberger; Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017, pp. 4700-4708.

[11]Gao Huang, Zhuang Liu, Laurens van der Maaten, Kilian Q. Weinberger; Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017, pp. 4700-4708.

[12] Ramzan, Farheen & Khan, Muhammad Usman & Rehmat, Asim & Iqbal, Sajid & Saba, Tanzila & Rehman, Amjad & Mehmood, Zahid. (2019). A Deep Learning Approach for Automated Diagnosis and Multi-Class Classification of Alzheimer's Disease Stages Using Resting-State fMRI and Residual Neural Networks. Journal of Medical Systems. 44. 10.1007/s10916-019-1475-2.