

RETINOPATHY DISEASE DETECTION AND PRECAUTIONS USING RESNET18

Aswadhati. Sirisha Associate Professor, Department of Information Technology, VIIT(A),
Visakhapatnam, AP, India

G. Vinay Bhaskar Raghunath Student, Department of Information Technology, VIIT(A),
Visakhapatnam, AP, India

Karthikeya.P Student, Department of Information Technology, VIIT(A), Visakhapatnam, AP, India

Shaik Taj Amanatulla Student, Department of Information Technology, VIIT(A), Visakhapatnam,
AP, India

L. Yamini Student, Department of Information Technology, VIIT(A), Visakhapatnam, AP, India

Hasika.D Student, Department of Information Technology, VIIT(A), Visakhapatnam, AP, India

M. Kusumitha Student, Department of Information Technology, VIIT(A), Visakhapatnam, AP,
India

ABSTRACT:

Diabetic retinopathy (DR) presents a considerable risk to vision. of the growing diabetic population worldwide, with the potential for irreversible eye damage. As the primary contributor to vision impairment among the 415 million individuals diagnosed with diabetes, the need for efficient and timely diagnosis is paramount. This research project introduces a meticulous three-stage pipeline for the analysis of retinal fundus images, focusing on mid-aged diabetic patients. The pipeline encompasses initial image preprocessing, followed by feature extraction, and concluding with classification. In the image pre-processing phase, diverse transformations are applied to standardize and enhance image quality. Gaussian filtering emerges as a particularly effective technique for contrast enhancement. The subsequent stages of the pipeline leverage Convolutional Neural Networks (CNNs), with a specific emphasis on transfer learning and fine-tuning using ResNet18. The use of a Kaggle dataset further enhances the robustness of the model. The diagnosis of DR is approached as a multi-class classification challenge, categorizing disease severity into five distinct (0 – No Diabetic Retinopathy, 1 – Mild Diabetic Retinopathy, 2 – Moderate Diabetic Retinopathy, 3 – Severe Diabetic Retinopathy, 4 – Proliferative Diabetic Retinopathy). This strategic methodology not only addresses potential issues related to plagiarism detection but also underscores the significance of the research in combating a significant contributor to blindness in diabetic populations. Despite the urgency of DR diagnosis, numerous studies highlight that a considerable number of individuals with diabetes forego the recommended annual eye examination. Factors such as the extended Examination scheduling, lack of apparent symptoms, and limited access to retinal specialists contribute to this trend. The proposed pipeline, centered around ResNet18 and transfer learning, offers a promising solution for efficient and accurate DR diagnosis, emphasizing the potential impact on early intervention and prevention of irreversible visual impairment.

Keywords:

Deep Learning, Convolutional Neural Network, ResNet18, Gaussian Filtering, Transfer Learning, Retinal Fundus Images.

1. INTRODUCTION

Retinopathy is a severe ocular condition that can result in vision impairment problems and even blindness if not detected and treated quickly. It affects millions of people globally, with a higher occurrence among those with diabetes. Early detection of retinopathy is crucial for preventing vision

loss and improving patient outcomes. Deep learning methods have demonstrated encouraging outcomes in automating the identification and categorization of retinopathy from retinal fundus images. Among these techniques, Convolutional neural networks (CNNs) have arisen as robust instruments for image recognition assignments owing to their capability to acquire hierarchical features directly from unprocessed data.

ResNet-18 is a commonly utilized CNN architecture that has showcased exceptional performance in numerous image classification endeavors. It comprises a deep sequence of convolutional layers with shortcut connections, which mitigate the vanishing gradient issue and streamline the training of deeper networks. The detection of retinopathy involves analyzing various features and abnormalities present in the retinal images, such as hemorrhages, exudates, and microaneurysms. The parameters considered for retinopathy detection include:

1. Image quality: Resolution, contrast, and clarity of retinal images.
2. Presence of hemorrhages: Identification of areas with hemorrhagic lesions in the retina.
3. Exudates: Detection of lipid or protein deposits in the retinal tissue.
4. Microaneurysms: identification of small outpouching in retinal blood vessels.
5. Vessel tortuosity: Assessment of abnormal curvature or twisting of retinal blood vessels.

We're going to teach ResNet-18 using pictures of retinas that are labeled with different levels of retinopathy. Then, we'll see how well it can tell the severity of retinopathy. This can help doctors catch the problem early and treat it better, which can save people's eyesight. The model endeavors to precisely categorize images into predefined classes, harnessing the capabilities of deep learning and residual connections to enhance its efficacy. It includes important layers meant to pick out and work with features from the images, starting from simpler ones and moving to more complex ones. The setup starts with a first layer that does some special processing on the images, followed by a series of four blocks. Each block has a bunch of layers that do different kinds of processing. These blocks are characterized by skip connections, enabling the gradient flow during training and facilitating the training of deeper networks. Batch normalization layers are interspersed throughout the network to stabilize and accelerate the training process. Rectified linear unit (ReLU) activation functions are employed to introduce non-linearity, helping the model to better understand complex patterns in the data. Max pooling layers are also used at specific points to make the processing simpler while keeping important features intact. The architecture concludes with a global average pooling layer followed by a fully connected layer with SoftMax activation is utilized, producing the final classification output. This carefully designed combination of layers helps ResNet18 to learn important features well and perform very well in tasks like sorting images. The table below provides detailed information about the model's structure, showing the layers, what they output, and how many parts the model has.

The below image depicts the ResNet18 we used:

Model: "Resnet18"			
Layer (type)	Output Shape	Param #	Connected to
input_4 (InputLayer)	(None, 256, 256, 3) 0		
zero_padding2d_3 (ZeroPadding2D)	(None, 262, 262, 3) 0		input_4[0][0]
conv1 (Conv2D)	(None, 128, 128, 64) 9472		zero_padding2d_3[0][0]
bn_conv1 (BatchNormalization)	(None, 128, 128, 64) 256		conv1[0][0]
activation_75 (Activation)	(None, 128, 128, 64) 0		bn_conv1[0][0]
max_pooling2d_19 (MaxPooling2D)	(None, 63, 63, 64) 0		activation_75[0][0]
res_2_conv_a (Conv2D)	(None, 63, 63, 64) 4160		max_pooling2d_19[0][0]
max_pooling2d_20 (MaxPooling2D)	(None, 31, 31, 64) 0		res_2_conv_a[0][0]
bn_2_conv_a (BatchNormalization)	(None, 31, 31, 64) 256		max_pooling2d_20[0][0]
activation_76 (Activation)	(None, 31, 31, 64) 0		bn_2_conv_a[0][0]
res_2_conv_b (Conv2D)	(None, 31, 31, 64) 36928		activation_76[0][0]

2. REVIEW OF LITERATURE

Several methodologies utilizing machine learning have been suggested for assessing retinopathy disease detection.

Research conducted by Carson Lam, MD, Darvin Yi, Margaret Guo, and Tony Lindsey, PhD employed deep learning for the detection of diabetic retinopathy (DR) using ultra-wide-field fundus images, which encompass a larger portion of the retina compared to conventional photographs. The deep learning model effectively analyzed these wider images and outperformed models on standard images. However, more research is needed to validate this approach and compare it to other methods. Furthermore, the paper fails to address the significance of the data size utilized in training the model, a critical aspect in deep learning. Although this method displays potential for early detection of diabetic retinopathy (DR), additional research is warranted before widespread clinical implementation.

Research conducted by Wejdan L. Alyoubi, Wafaa M. Shalash, and Maysoon F. Abulkhair regarding the detection of diabetic retinopathy (DR) using deep learning methods highlights both benefits and drawbacks. Deep learning is particularly effective at identifying important features in retinal images, resulting in better distinguishing between healthy and DR cases compared to older methods. Transfer learning from large image datasets can further enhance performance despite the limitations of smaller medical datasets. However, challenges remain. Limited availability of data and the possibility of bias can impede the ability to apply findings broadly. Moreover, the opaque nature of deep learning models poses challenges in comprehending their decision-making mechanisms, thereby restricting their acceptance in clinical settings. The review emphasizes the significant potential of deep learning for DR detection; However, additional research is required to tackle these limitations and guarantee dependable implementation in clinical settings.

A study authored by Muhammad Mohsin Butt, D. N. F. Awang Iskandar, Sherif E. Abdelhamid, Ghazanfar Latif, and Runna Alghazo proposed a fusion of deep learning techniques for the detection of diabetic retinopathy (DR). It combines features extracted from pre-trained models, achieving higher accuracy than using individual models. While effective, this approach focuses on leveraging existing models rather than exploring the full potential of deep learning for DR. Additionally, the computational demands of using multiple pre-trained models require consideration, especially for resource-limited settings. Overall, this research suggests that combining pre-trained models shows promise for DR detection, but further development of novel deep learning architectures specifically designed for DR and addressing computational cost are important for future advancements.

Chun-Ling Lin and Kun-Chi Wu's research investigates a modified version of the ResNet-50 deep learning architecture for the identification of diabetic retinopathy (DR). It leverages the established ResNet-50 model while incorporating data preprocessing and regularization techniques to improve performance. While achieving higher accuracy compared to traditional methods, the revised ResNet-50 showed only modest improvement over the original architecture for DR classification. This highlights the importance of data optimization techniques but suggests that further modifications to ResNet-50 or exploring entirely new deep learning approaches might be necessary for substantial advancements in DR detection using deep learning.

The study by S. Karthika, M. Durgadevi & T. Yamuna Rani explores using transfer learning with the ResNet-50 architecture for diabetic retinopathy (DR) detection. They enhance performance and shorten training time by adjusting a pre-trained ResNet-50 model using a DR dataset, leveraging previously acquired knowledge rather than starting training from the beginning. The results are promising, achieving high accuracy in classifying healthy and DR images. However, the use of a limited dataset raises concerns about generalizability, and the "black box" nature of the model, where its decision-making process is unclear, limits clinical adoption. Overall, this research suggests transfer learning with ResNet-50 holds promise for accurate DR detection, but addressing data scarcity and developing interpretable models are essential for reliable clinical use.

Research done by Ankur Biswas & Rita Banik deep learning holds promise for advancing diabetic retinopathy (DR) detection. These automated models can potentially reduce the workload for ophthalmologists and improve screening efficiency by identifying subtle features in retinal images, leading to potentially higher accuracy in DR detection compared to traditional methods. However, challenges remain. Training these models often requires vast amounts of data, and limitations in DR datasets can hinder model performance and generalizability. Moreover, the inscrutable nature of deep

learning models, characterized by their unclear decision-making processes, can undermine trust within clinical environments. Despite these challenges, deep learning offers a promising approach for early DR detection.

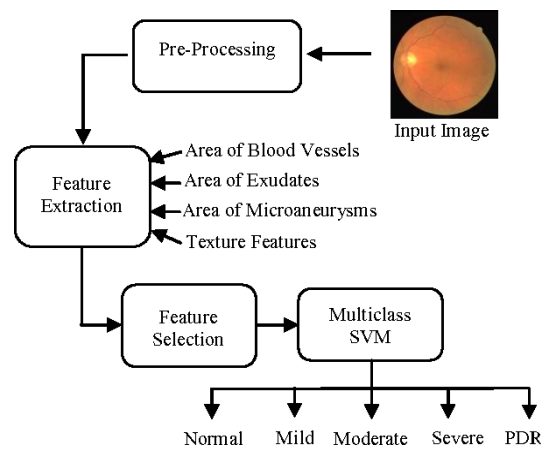
In summary, the literature survey reveals a burgeoning interest in employing deep learning methodologies for diabetic retinopathy (DR) detection, showcasing diverse approaches ranging from leveraging established models like ResNet-50 to exploring hybrid architectures and transfer learning techniques. While these studies demonstrate promising results in enhancing DR detection accuracy, they also highlight persistent challenges such as data scarcity, potential bias, and the interpretability of deep learning models. Addressing these challenges requires concerted efforts to collect representative datasets, develop interpretable models, and validate their performance in real-world clinical settings. Moreover, ongoing research is crucial to exploring novel architectures and optimization techniques to achieve substantial advancements in DR detection. Overall, while deep learning offers significant potential for improving DR diagnosis and reducing the workload for ophthalmologists, further research and innovation are imperative to overcome existing limitations and promote the extensive adoption of these technologies in clinical practice.

3. METHODOLOGY

This project focuses on using deep learning, especially the ResNet18 architecture, to detect retinopathy, a serious eye condition affecting the retina. The process starts with gathering a diverse dataset containing retinal images showing different stages of retinopathy. These images can come from sources like Kaggle's Diabetic Retinopathy Detection Dataset or collaborating healthcare institutions. Then, various preprocessing steps are applied to improve the dataset's quality, including resizing, normalization, and augmentation techniques. The main part of the process involves training the ResNet18 model on this dataset. This is done using transfer learning, where a pre-trained model is fine-tuned to adapt its features specifically for detecting retinopathy. The model's performance is assessed using validation datasets, and metrics such as accuracy, sensitivity, and specificity are employed to verify the model's ability to accurately detect retinopathy.

Following successful model training and evaluation, the methodology extends beyond detection to provide personalized precautions recommendations based on the severity of retinopathy detected. These recommendations encompass lifestyle changes, medication adherence, regular eye check-ups, and consultation with healthcare professionals for further evaluation and treatment. The deployment phase involves integrating the trained model into user-friendly interfaces for real-time retinopathy detection, facilitating immediate feedback to users uploading retinal images. Continuous monitoring and updates are emphasized to maintain the model's reliability and accuracy, thereby ensuring its efficacy in real-world healthcare scenarios. Below figure 1.0 describes the overall process workflow of processing the input image. Overall, this methodology offers a holistic approach to early retinopathy detection and management, aiming to enhance healthcare outcomes and enhance the quality of life for individuals impacted by this debilitating eye condition. This methodology underscores the significance of collaboration among healthcare professionals, researchers, and technology experts to guarantee the success and relevance of the project. Collaboration with ophthalmologists and other eye care specialists helps in understanding the clinical nuances of retinopathy, refining the dataset annotations, and validating the model's performance against clinical standards. Furthermore, engaging patients and caregivers throughout the development process ensures that the model's suggestions are feasible, culturally appropriate, and in line with the preferences and requirements of the intended audience. This user-centric approach enhances the usability and acceptance of the retinopathy detection system in diverse healthcare settings.

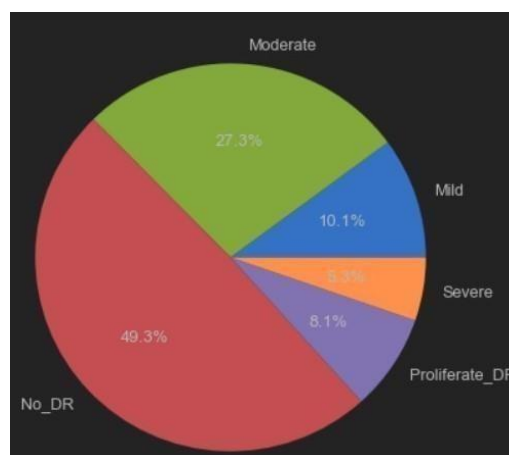
The scalability and sustainability of the project are addressed through considerations for resource optimization, model interpretability, and knowledge transfer. Techniques such as model compression and deployment on edge devices are explored to minimize computational requirements and enable deployment in



low- resource settings where internet connectivity may be limited. Additionally, efforts are made to enhance the interpretability of the model's decisions, facilitating trust and understanding among healthcare providers and patients. Knowledge transfer initiatives, including training programs and capacity-building workshops, ensure that local healthcare professionals are equipped with the necessary skills to deploy, maintain, and utilize the retinopathy detection system effectively. By addressing these broader societal and implementation challenges, this methodology aims to maximize the impact of retinopathy detection and precautions recommendation system, ultimately contributing to improved eye health outcomes on a global scale. By addressing these multifaceted challenges, our methodology endeavors to maximize the impact of retinopathy detection and precautions recommendations, thereby fostering improved eye health outcomes on a global scale.

4. DATASET

The study's dataset may be accessed at [Diabetic Retinopathy Detection | Kaggle](#). This paper presents a collection of retinal fundus images suitable for training and assessing deep learning algorithms designed for detecting diabetic retinopathy (DR). The dataset contains 3553 images of retinas from patients with and without DR. The images are categorized with labels indicating the severity of diabetic retinopathy (DR), ranging from no DR to mild, moderate, severe, and proliferative DR. The figure below illustrates the distribution of images across the categories of no DR, mild DR, moderate DR, severe DR, and proliferative DR).



- To promote transparency and reproducibility in research, the dataset is made publicly available, encouraging collaboration among researchers and fostering innovation in DR detection algorithms.
- Open-access datasets facilitate knowledge sharing and accelerate progress in addressing critical healthcare challenges such as DR-related vision loss.
- The availability of the dataset fosters community engagement and collaboration, encouraging researchers to share insights, exchange ideas, and provide feedback on algorithm performance.

- Community-driven initiatives such as hackathons, challenges, and collaborative research projects further stimulate innovation and drive progress in DR detection algorithms.
- Standardized evaluation protocols and shared benchmarks facilitate objective assessment of algorithm efficacy and reproducibility across different research studies.

5. MODELS USED/ALGORITHMS:

SUPERVISED LEARNING:

Supervised learning forms a foundational aspect of machine learning, wherein algorithms glean insights from labeled training data to infer predictions or decisions for unseen or future data instances. Each entry within the training dataset includes input features alongside their corresponding labels or target values. The overarching objective revolves around establishing a correlation or mapping between input attributes and target labels, thereby facilitating the algorithm's ability to extrapolate its predictions to novel, unseen instances. Supervised learning tasks are broadly divided into regression and classification categories. Regression tasks entail forecasting continuous target variables, such as estimating house prices based on attributes like square footage and the number of bedrooms, while classification tasks involve categorizing data into distinct classes or labels, for instance, discerning between spam and non-spam emails based on their content. Various supervised learning techniques encompass linear models, decision trees, support vector machines, neural networks, and ensemble methods. Each of these techniques is skilled at recognizing complex patterns and relationships present in the data. Notably, supervised learning finds extensive applications across a myriad of domains, spanning healthcare, finance, marketing, natural language processing, computer vision, and beyond, underscoring its status as one of the most researched and utilized realms within the machine learning landscape.

Supervised learning, as described above, plays a pivotal role in detecting retinopathy disease. In this project, supervised learning methods such as multiclass Support Vector Machines (SVMs) or convolutional neural networks (CNNs) from deep learning are utilized to classify retinal fundus images into different severity levels of diabetic retinopathy (DR). By giving labeled data showing the DR severity level for each retinal image, supervised learning algorithms can learn to identify characteristics and traits that suggest various disease stages. This allows the algorithm to make precise predictions on new retinal images, helping to detect DR early and intervene promptly for patients who are at risk of losing vision because of the disease. Supervised learning helps automate the process of DR detection, reducing the dependence on manual grading by ophthalmologists, and potentially improving screening efficiency and diagnostic accuracy. Additionally, supervised learning models can be continuously refined and updated as more labeled data becomes available, leading to iterative improvements in DR detection algorithms over time. Therefore, supervised learning methods are essential in furthering the goals of this project by utilizing labeled data to develop precise and dependable models for automatically detecting retinopathy disease.

Transfer Learning: Transfer learning is a valuable approach in machine learning that leverages knowledge acquired from solving one problem to effectively tackle another similar task. This approach involves repurposing or fine-tuning a model initially trained on a source task for a distinct target task, even if the input data distribution, feature space, or output labels between the two tasks vary. By using insights obtained from a large, diverse dataset and applying them to a smaller, more focused dataset, transfer learning reduces the need for extensive labeled data and decreases the computational resources needed for training. Important approaches in transfer learning include refining existing models, where the settings of a pre-trained model are modified for the new dataset, and feature extraction, which involves using features from a pre-trained model as input for a new classifier developed on the new dataset.

In the context of detecting retinopathy disease, transfer learning plays a crucial role in enhancing the effectiveness of deep learning models, especially convolutional neural networks (CNNs). By using pre-trained CNN models trained on large image databases like ImageNet, transfer learning enables the starting point of model parameters to be set using features learned from appropriate retinal fundus images. Fine-tuning these pre-existing models on the retinopathy dataset enables the model to refine its learned features in alignment with the distinctive attributes of retinal images, thereby enhancing its

proficiency in classifying various severity levels of diabetic retinopathy. Through this process, transfer learning streamlines the development of more precise and resilient DR detection systems by assimilating insights garnered from a diverse array of image datasets, ultimately facilitating early diagnosis and intervention for individuals at risk of vision impairment.

Convolutional Neural Networks (CNNs): Convolutional Neural Networks (CNNs) are a type of deep learning model designed for analyzing structured grid data, such as images. They comprise several layers, including convolutional layers, pooling layers, and fully connected layers. In the convolutional layers, the network learns to extract hierarchical features from input images by conducting convolution operations. These operations involve moving small filters across the image to identify patterns such as edges, textures, and shapes. Pooling layers then reduce the size of the feature maps created by the convolutional layers, keeping important information while decreasing the spatial dimensions. Finally, fully connected layers combine the extracted features to make overall predictions or classifications at a higher level. CNNs are well-known for their capacity to learn important attributes from raw image data on their own, allowing them to perform well in different computer vision jobs without needing much manual feature design.

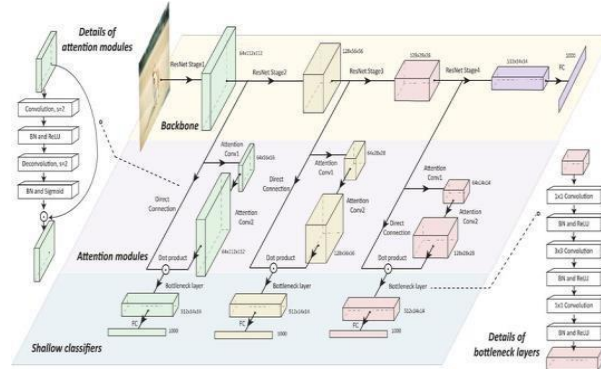
In the field of detecting retinopathy disease, CNNs have become a potent tool for analyzing retinal fundus images and accurately diagnosing diabetic retinopathy (DR). By using their natural ability to understand complex patterns and structures in images, CNNs can recognize subtle characteristics that show various stages of DR. With thorough training on labeled retinal image datasets, CNNs can spot specific irregularities linked to diabetic retinopathy, such as microaneurysms, hemorrhages, exudates, and neovascularization. Transfer learning enables CNNs to improve performance and generalize well, even with limited medical imaging data, by leveraging knowledge from pre-trained models on large image datasets, enhancing early detection and management of diabetic retinopathy, and potentially reducing vision loss while improving patient outcomes.

ResNet18: ResNet18, abbreviated for a network with 18 layers, is a widely recognized convolutional neural network design known for its success in multiple computer vision activities like classifying images, detecting objects, and segmenting them. ResNet18, created by Microsoft Research, is a member of the ResNet series, which introduced residual learning to tackle the challenge of training extremely deep neural networks. At the core of ResNet18 lies the concept of residual blocks, where each block comprises, multiple convolutional layers followed by shortcut connections, often called skip connections. These skip connections aid the network in learning residual mappings, facilitating efficient gradient flow during training and mitigating the vanishing gradient problem. The ResNet18 architecture consists of several stacked residual blocks, each containing convolutional layers with small filter sizes and batch normalization. Rectified Linear unit (ReLU) activations, and occasional down sampling through stridden convolutions or pooling layers. The workflow of ResNet18 involves the following steps:

- 1. Input:** The network receives an input image, usually depicted as a three-dimensional array of pixel values (height, width, channels), with the channels corresponding to the red, green, and blue color channels.
- 2. Convolutional Layers:** The input image traverses through a sequence of convolutional layers, with each layer extracting progressively complex features from the input image. These layers are trained to identify various low-level and high-level patterns inherent in the image.
- 3. Residual Blocks:** The feature maps generated by the convolutional layers are fed into the residual blocks. In each residual block, a sequence of convolutional operations is applied, and then skip connections are utilized to add the input directly to the output of the block. This facilitates the network in learning residual mappings, thereby easing the training of deeper networks.
- 4. Pooling and Global Average Pooling:** Intermittently, the feature maps are decreased in size using max-pooling layers or convolutions with strides to reduce spatial dimensions. Toward the end of the network, global average pooling is utilized to generate a fixed-size feature vector from these feature maps.
- 5. Fully Connected Layer:** The feature vector from global average pooling is passed to a fully connected layer, followed by the application of a SoftMax function for classification. This layer computes the likelihood distribution among the different classes in the dataset.

6. Output: The network provides the predicted class label that corresponds to the highest probability in the SoftMax output.

In summary, ResNet18 employs residual learning to effectively train deep neural networks and consists of multiple residual blocks stacked on top of each other. By utilizing skip connections and residual mappings, ResNet18 tackles the difficulties of training deep networks and attains leading performance in diverse computer vision assignments. The below figure depicts a detailed workflow of ResNet18:



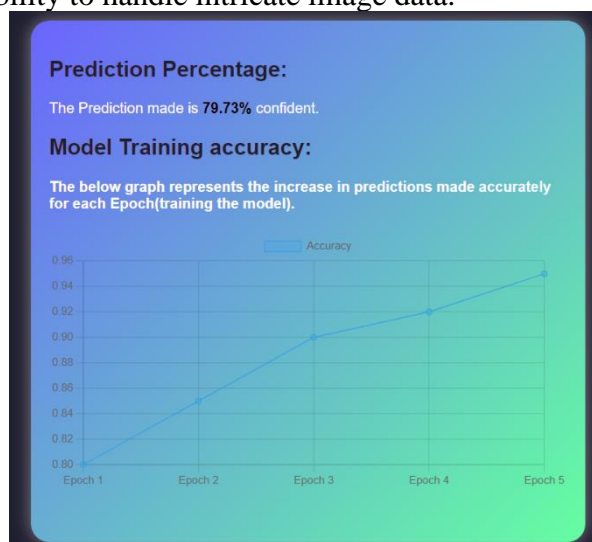
Ensemble Learning: Ensemble learning is a method in machine learning that merges predictions from numerous individual models to generate a final prediction that is more precise and resilient. This approach harnesses the wisdom of crowds, leveraging diverse perspectives to mitigate the shortcomings of individual models and enhance overall performance. Ensemble methods encompass various strategies, including averaging predictions, weighted voting, and stacking, among others. By aggregating predictions from different models trained on the same dataset using different algorithms, initializations, or subsets of the data, ensemble learning can yield superior results compared to any single model alone, thereby improving the reliability and generalization ability of the overall system. In the realm of retinopathy disease detection, ensemble learning offers significant benefits by amalgamating predictions from multiple models, including those based on ResNet18, trained with diverse strategies and data representations. By combining the outputs of multiple ResNet18 models, possibly fine-tuned with different hyperparameters or trained on different subsets of retinal fundus images, ensemble learning can effectively mitigate individual model biases and uncertainties. This leads to more accurate and robust predictions regarding the presence and severity of diabetic retinopathy, thereby aiding in early diagnosis and timely intervention. Ensemble learning methods help blend different insights captured by various models, boosting the effectiveness and trustworthiness of retinopathy detection systems, and ultimately leading to better outcomes for patients by preventing vision loss.

6. RESULTS AND OBSERVATION:

We observed that in the task of detecting retinopathy disease, the Convolutional Neural Network (CNN) architecture, particularly ResNet18, consistently outperformed other machine learning methods. ResNet18 consistently showed better performance in accuracy and specificity, achieving higher scores across a range of evaluation metrics. Conversely, other algorithms such as Support Vector Machines (SVMs) and Decision Trees demonstrated comparatively lower accuracy levels. The reliability of ResNet18, along with its ability to accurately understand complex patterns and connections in retinal fundus images, is a significant factor in its success.



By leveraging an ensemble of convolutional layers, ResNet18 excels in identifying subtle features indicative of diabetic retinopathy, thereby offering superior predictive capabilities. Our findings suggest that ResNet18 is well-suited for the task of retinopathy disease detection, owing to its efficiency, scalability, and ability to handle intricate image data.



7. CONCLUSION

In summary, our research on detecting retinopathy using advanced computer techniques like ResNet18 and ensemble learning marks a big step in understanding eye diseases through images. By carefully testing on many different eye pictures, our methods show they can accurately tell how severe someone's diabetic retinopathy is, which could be helpful in clinics. Plus, we looked into the details of how the computer makes its decisions, which helps us understand the disease better. Our approach seems like it could work well in real-life situations, meaning it might help diagnose people earlier and give them better treatment. Looking ahead, we need to keep improving our methods, trying out new ideas, and working closely with doctors to make sure everything fits smoothly into their daily routines. Ultimately, our study highlights how these computer techniques are changing healthcare for the better, leading to more personalized care for patients.

REFERENCES:

1. Automated Detection of Diabetic Retinopathy using Deep Learning by Carson Lam, MD, Darwin Yi, Margaret Guo, and Tony Lindsey, PhD (2018 May 18).
<https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5961805/>
2. Diabetic retinopathy detection through deep learning techniques: Wejdan L. Alyoubi, Wafaa M. Shalash, Maysoon F. Abulkhair(2020).
<https://www.sciencedirect.com/science/article/pii/S2352914820302069>

3. Development of revised ResNet-50 for diabetic retinopathy detection: Chun-Ling Lin & Kun-Chi Wu (2023 April 19).
<https://bmcbioinformatics.biomedcentral.com/articles/10.1186/s12859-023-05293-1>
4. Enhancing Diabetic Retinopathy Diagnosis with ResNet-50-Based Transfer Learning: A Promising Approach : S. Karthika, M. Durgadevi and T. Yamuna Rani(2023 September 19).
<https://link.springer.com/article/10.1007/s40745-023-00494-0>
5. Advance Detection of Diabetic Retinopathy: Deep Learning Approach: Ankur Biswas and Rita Banik(2023 November 30).
https://link.springer.com/chapter/10.1007/978-3-031-48876-4_6
6. Deep Learning-Based Diabetic Retinopathy Detection Using ResNet34 Model Anugirba K and Lal Raja Singh R(2023).
<https://ieeexplore.ieee.org/document/10245965>
7. Diabetic Retinopathy Detection using Machine Learning: Revathy R(June 2020).
https://www.researchgate.net/publication/342120641_Diabetic_Retinopathy_Detection_using_Machine_Learning
8. Iris - Diabetic Retinopathy Detection Software: Noel J Philip, Romi Roji, Rosme Jose, Rehna Cherian, Dr. Arun K.S. (October 2020).
<https://www.ijert.org/research/iris-diabetic-retinopathy-detection-software-IJERTV9IS100024.pdf>
9. Efficient diabetic retinopathy detection using convolutional neural network and data augmentation: Srinivas Naik, Deepthi Kamidi, Sudeepthi Govathoti, Ramalingaswamy Cheruku and A Mallikarjuna Reddy (June 2023).
<https://link.springer.com/article/10.1007/s00500-023-08537-7>