Ocular Disease Detection using Deep Convolution Neural Network

MSc Research Project
Data Analytics

Priyanshu Srivastava
Student ID: X21199787

School of Computing
National College of Ireland

Supervisor: Prof. Hicham Rifai

National College of Ireland

MSc Project Submission Sheet

School of Computing

Student Name:	Priyanshu Srivastava		
Student ID:	X21199787		
Programme:	Data Analytics	**Year:**	2022/2023
Module:	MSc. Research Project		
Supervisor:	Prof. Hicham Rifai		
Submission Due Date:	14/08/2023		
Project Title:	Ocular Disease Detection using Deep Convolution Neural Network		
Word Count:	**5467**	**Page Count: 18**	

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.
ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature: ……

Date: ……

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST

Attach a completed copy of this sheet to each project (including multiple copies)	□
Attach a Moodle submission receipt of the online project submission, to each project (including multiple copies).	□
You must ensure that you retain a HARD COPY of the project, both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	□

Assignments that are submitted to the Programme Coordinator Office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

Ocular Disease Detection using Deep Convolution Neural Network

Priyanshu Srivastava
X21199787

Abstract
Ophthalmic diagnostics are on the verge of a revolutionary development in a world where medical progress is constantly redefining the limits of healthcare. It holds the possibility of preserving vision and enhancing lives to be able to identify retinal disorders precisely and quickly. This study introduces a novel hybrid model that combines capsule networks with convolutional neural networks (CNNs) for better categorization of retinal diseases. The model uses CNNs for multi-label disease classification and capsule networks for hierarchical feature extraction from retinal pictures, all with the goal of improving ophthalmic diagnostics. The strategy addresses difficulties in detecting rarer illnesses like TSLN and ODC while demonstrating promising accuracy for big diseases like Diabetic Retinopathy (DR) and NORMAL instances. The model improves sensitivity and specificity for DR detection. This study highlights its potential as a useful tool for supporting medical professionals in the diagnosis of retinal diseases. To improve model performance and increase its clinical utility, future research will focus on data augmentation techniques, architectural improvements, and partnerships with experts in the field.

Keywords: Retinal disease classification, Convolutional neural networks, Capsule networks, Automated diagnosis, Sensitivity, Specificity.

1 Introduction

The fusion of deep learning with computer vision in recent years has sparked revolutionary improvements in medical diagnosis and therapy. With significant implications for maintaining visual health and quality of life, early detection and diagnosis of ocular illnesses have become increasingly important in this setting. This study sets out on a mission to use Deep Convolutional Neural Networks (DCNNs), specifically through the cutting-edge technology of Capsule Networks, to overcome the difficulties of ocular illness identification using fundus pictures. The disorders under investigation include Age-Related Macular Degeneration, Media Haze, Optic Disc Cupping, Diabetic Retinopathy, and Tessellation, each of which has distinct complications and effects on health. The critical requirement for precise and quick eye disease detection served as the driving force for our investigation. Current diagnostic techniques are labour-intensive and susceptible to inter-observer variability since they frequently rely on manual evaluations by ophthalmologists. The crucial importance of computational techniques in promoting ocular health is demonstrated in a fundamental study by Mittal and Rajam (2020) through a thorough examination of methodologies (Mittal & Rajam 2020). Their study highlights the game-changing potential of using digital image analysis to improve the way ocular disease is diagnosed. One of the conditions studied in this study, diabetic retinopathy, continues to be a common reason why people with diabetes experience vision impairment. A robust automated system is required for accurate detection of media haze, which is characterized by the clouding of ocular media. Glaucoma, a progressive visual neuropathy, and optic disc cupping are intimately related, underscoring the importance of an early and thorough diagnosis. Further highlighting the variety of issues addressed by this research are tessellations, a complex retinal pattern anomaly, and age-related macular degeneration, a significant cause of visual loss in the aging population. The effectiveness of Capsule Networks, which was introduced as a paradigm shift in deep learning architecture, provides a strong basis for this investigation. A promising method to understand complex features and their interactions—essential aspects of ocular illness diagnosis—is provided by Capsule Networks, which have the unusual capacity to record hierarchical relationships within fundus images (Patrick et al., 2019). This study aims to close the gap between cutting-edge technology and ocular health considering the difficulties and potential. This study aims to further the creation of automated, precise, and effective diagnostic tools by investigating the use of Capsule Networks in the field of eye illness identification utilizing fundus pictures. The methodology, fundus image data collection and preprocessing, model architecture, experimental design, and evaluation metrics will all be covered in detail in the coming sections of this report before a thorough analysis of the findings. With the help of this research, we hope to strengthen the relationship between deep learning and ophthalmology, with the goal of enabling doctors to provide better care and results for patients.

1.1 Research Question:

How well can Deep Convolution Neural Network models classify ocular disorders into the following disease groups: Age Related Macular Degeneration, Diabetic Retinopathy, Media Haze, Optic Disc Cupping and Tessellation?

1.2 Structure of Paper

The framework of the paper has been clarified and follows predetermined guidelines. A summary of prior relevant studies is provided in Section 2, under "Literature Review." The employed techniques are described in Section 3, "Methodology." A quick synopsis of models that have been incorporated is given in Section 4, "Implementation." The results are presented and discussed in Section 5, "Evaluation Results." The "Conclusion" section of Section 6 and the "Future Work" section of Section 7 each provide in-depth analysis and lay out potential study avenues.

2 Related Work

The urgent need for early identification of conditions that can cause vision impairment, such as cataract, glaucoma, and retinal disorders, is addressed in the work by (Ramanathan et al., 2021). To reduce the occurrence of blindness through prompt intervention, the study offers a system that uses the methods of Logistic Regression, Random Forest, Gradient Boosting, and Support Vector Machine for efficient illness diagnosis. The importance of multiclass constructed models is emphasized in the paper, which uses fundus images to forecast illness outcomes. Gradient Boosting shows promising accuracy (90%) for cataract identification. It demonstrates how effective machine learning is at growing healthcare systems, notably in applications for eye screening, and it makes suggestions for the future in terms of things like real-time mobile apps and Receiver Operating Curve (ROC) integration. The study emphasizes how machine learning approaches can help with disease identification and help with better healthcare outcomes.

The need for precise health data recording to increase the dependability of machine learning for quick diagnosis is addressed in the research study by Malik et al. (2019). The study suggests a thorough structure for storing diagnostic data in a manner that is internationally recognized, facilitating disease prediction using machine learning algorithms based on symptoms. The article examines patient data including age, medical history, and clinical observations using several algorithms, including Decision Tree, Random Forest, Naive Bayes, and Neural Network. Data is organized using hierarchies created by medical professionals and aligned with ICD-10 classification. The study demonstrates the effectiveness of the suggested framework in achieving satisfactory classification outcomes, particularly for tree-based approaches with prediction rates greater than 90%. This study highlights how structured data organization and strong algorithms can improve the precision of disease diagnosis and prediction in the context of eye health.

(Purandare et al. 2016) propose a hybrid approach for diabetic retinopathy (DR) classification, using clinical feature extraction and an SVM classifier to quickly and accurately classify subjects as Normal or Diabetic Retinopathy cases, with a remarkable accuracy of 92.55%, specificity of 96%, sensitivity of 78%, and positive predictive value (PPV) of 95.12%, emphasizing early identification to reduce vision loss risks.(Swathi et al., 2022) use deep learning and a web application to improve DR detection, successfully removing noise from fundus images and facilitating accurate disease classification, showcasing the potential of technology-driven methods to improve early diagnostic accuracy and user interaction for medical care enhancement.

(Reddy, 2021) addresses the crucial problem of using optical coherence tomography (OCT) scans to diagnose age-related macular degeneration (ARMD). The disorder known as ARMD, which impairs central vision, can have a substantial effect on daily activities. The study suggests a complete strategy that combines Active Contour and Directional Total Variation (DTV) denoising techniques to improve image quality and separate macular layers.

The study efficiently classifies OCT scans as healthy or infected using Convolutional Neural Networks (CNNs), including variations like AlexNet, VggNet, and GoogleNet, and achieves superior results than those obtained using traditional techniques. This work serves as an example of how deep learning plus picture enhancement may be used to automate ARMD detection. The suggested method provides a reliable solution for precise and effective diagnosis by merging preprocessing, segmentation, and classification. The suggested approach combines preprocessing, segmentation, and classification to offer a trustworthy answer for accurate and efficient diagnosis.

A convolutional neural network (CNN) that has been optimized for the early identification of specific eye diseases is described in the study by (Chellaswamy et al., 2022). The study attempts to improve disease identification accuracy by implementing a deep learning approach and optimizing the CNN's hyperparameters using the Whale Optimization Algorithm. The enhanced CNN performs better than the baseline CNN in extensive testing using a variety of retinal fundus imaging datasets, including datasets for cataracts, glaucoma, diabetic retinopathy (DR), and age-related macular degeneration (AMD). The proposed model, which outperforms non-optimized techniques, shows notable accuracy increases. Accuracy, sensitivity, specificity, precision, and F1-score are just a few of the performance indicators that the study's detailed examination incorporates. This results in a thorough review of the model's performance. The results of the study showed that the optimized DCNN achieved detection accuracy for AMD of 97.7%, glaucoma of 97.2%, cataracts of 96.4%, and DR of 97.4%.

Using a variety of optimization strategies, the (V. R et al., 2023) research combines convolutional neural networks (CNNs) with bidirectional long short-term memory (LSTM) networks. Adagrad achieves a noteworthy accuracy of 95.3%, highlighting the need of optimization for improved disease classification. The study by (A. Raza et al., 2021) uses the Inception v4 deep learning model to classify digital retinal pictures, achieving a remarkable 96% accuracy rate. Both researches highlight the potential for deep learning and sophisticated models to revolutionize the diagnosis of eye diseases, support physicians, and increase access to effective medical interventions. The advancement of research in this area shows promise for improving diagnostic precision, assisting patients, and influencing the future of ophthalmology.

An innovative method for glaucoma identification is presented by (Prananda et al., 2023), concentrating on deep learning analysis of the retinal nerve fiber layer (RNFL). The study uses a two-stage procedure involving pre-treatment and glaucoma categorization to get beyond the drawbacks of conventional approaches that rely on fluctuating cup-to-disc ratios. On the ORIGA dataset, the technique improved its accuracy to 92.88% with an AUC of 89.34%, demonstrating its potential to enhance glaucoma diagnosis. However, the study's exclusive focus on one dataset and disregard for computational time issues. For practical deployment, more investigation into computational effectiveness and varied datasets is required.

(Rekha, C. and Jayashree., 2022) discusses the issue of eye illnesses, which are becoming more widespread as a result of variables such as prolonged screen exposure. Their research focuses on the identification of hyphema, a serious and understudied vision issue. Despite case studies and research already being done, there are still no developed techniques for identifying hyphema. The authors respond by putting out a unique strategy that makes use of deep learning algorithms to anticipate Hyphema at an early stage, assisting in preventative eye care. The system uses image pre-processing methods to determine the severity of the condition using input from eye scans. Notably, the use of deep learning in disease prediction shows an accuracy of 85%.

Using Capsule Networks (CapsNet), (Santos et al., 2020) provides a potential technique for early glaucoma identification. Accurate diagnosis is necessary because glaucoma is one of the main causes of blindness. The study uses pre-processing to improve the properties of the images using CapsNet to obtain accuracy of 90.90% along with strong precision, recall, f1-score, AUC, and kappa index. The method is noteworthy since it achieves these outcomes with minimal data augmentation or optic disc segmentation. The study demonstrates how important capsules are for detecting glaucoma by highlighting their capacity to capture hierarchical spatial correlations for discrimination.

The research paper by (Nneji et al., 2021) introduces a dual-weighted shared capsule network for precise Diabetic Retinopathy (DR) classification utilizing fundus images. In order to address the urgent demand for automated DR diagnosis, the study develops a siamese capsule network that pulls out distinguishing characteristics from preprocessed photos. Using the MESSIDOR dataset, the suggested model demonstrated remarkable performance, achieving 99.1% accuracy, 98.5% sensitivity, and 98% specificity. The model, which incorporates methods like batch normalization, dropout, and wavelet pooling, performs better than single capsule networks and other CNN models, making it a useful tool for quick and accurate DR screening. The trade-off between accuracy and time efficiency is justified, despite slightly longer calculation periods.

A deep learning method for detecting retinal eye diseases, with an emphasis on illnesses including diabetic retinopathy and retinitis pigmentosa, was presented by (Jain et al., 2018). Convolutional neural networks (CNNs) are used in their approach, known as LCDNet, to categorize retinal fundus images as healthy or sick, doing away with the necessity for feature extraction or explicit segmentation. With LCDNet obtaining outstanding accuracy rates of 96.5% to 99.7% on multiple datasets, the study demonstrates the potential of deep learning in medical picture interpretation. The study emphasizes the significance of early disease identification and blindness prevention while also outlining potential future developments, like expanding the model for multi-class classification.

The study by (Guo et al., 2021), uses deep transfer learning to tackle the crucial task of using fundus imaging to predict various eye disorders. While advances in deep convolutional neural networks (DCNNs) have shown promise in illness prediction, few studies have concentrated on differentiating between various eye disorders and healthy instances, closely mirroring clinical practice. The study uses transfer learning to distinguish four common eye diseases—Glaucoma, Maculopathy, Pathological Myopia, and Retinitis Pigmentosa—from healthy controls with a small dataset. It also presents a lightweight deep learning architecture, MobileNetV2. The study places a strong emphasis on efficiency and accuracy, with average accuracy ratings of 96.2%, sensitivity of 90.4%, and specificity of 97.6%. InceptionV3 and AlexNet are greatly outperformed by MobileNetV2, which also improves classification accuracy.

3 Research Methodology

In this research we will follow Knowledge Discovery in Databases (KDD) process which is a thorough and methodical method used to extract important knowledge from vast and complex databases.

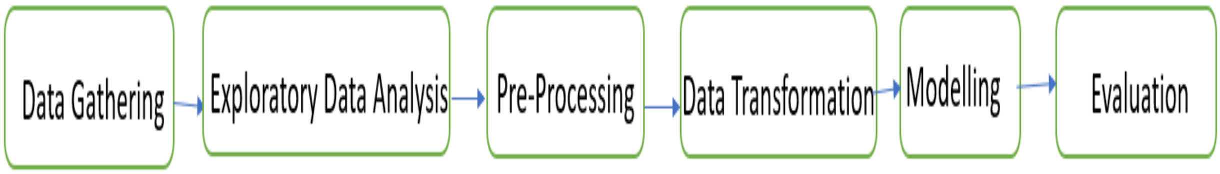

Figure 1: Knowledge Discovery in Databases Process.

Data Gathering: The dataset was gathered using Mendeley from reliable sources. The Multi-Label Retinal Diseases (MuReD) collection consists of 2208 retinal pictures that have 20 different labels from that we only used the images we required for our model building process. A sizable sample size is guaranteed for each of these labels because they reflect a wide variety of retinal diseases. The ARIA, STARE, and RFMiD datasets were combined to create the dataset, which was then meticulously post-processed to improve image quality, increase the number of pathologies, and lessen class imbalance.

EDA: Exploratory Data Analysis (EDA) entailed a detailed analysis of the dataset, which included examining the distribution of photos based on various image sizes and assessing the distribution of photographs per label. This stage helped to understand the make-up of the dataset and helped identify potential difficulties and opportunities for further steps. After observing the images the shapes of the images were different so the resizing of the images will be implemented in the pre-processing part.

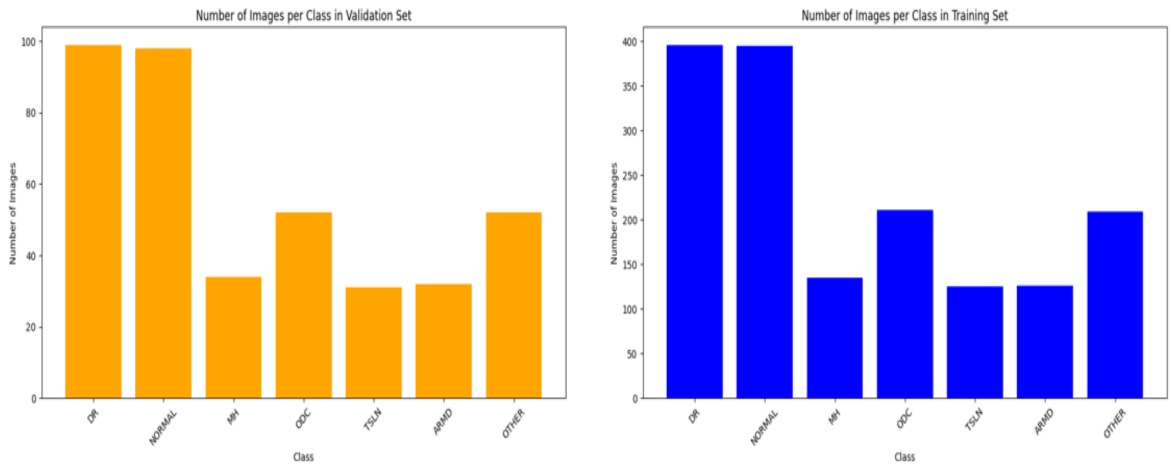

Figure 2: Number of Images per class Left: Validation Right: Training

Pre-Processing: Preprocessing is the important part of the step to get the data in a form where it will be used for modelling. As we learn from exploratory data analysis, we observed that some of the images were in different size. To bring all the images to one shape we resize the image into to the specified size as you can see in Figure 3 The after image of Diabetic Retinopathy like this all the images were resize into specified size which is (256 x 256).

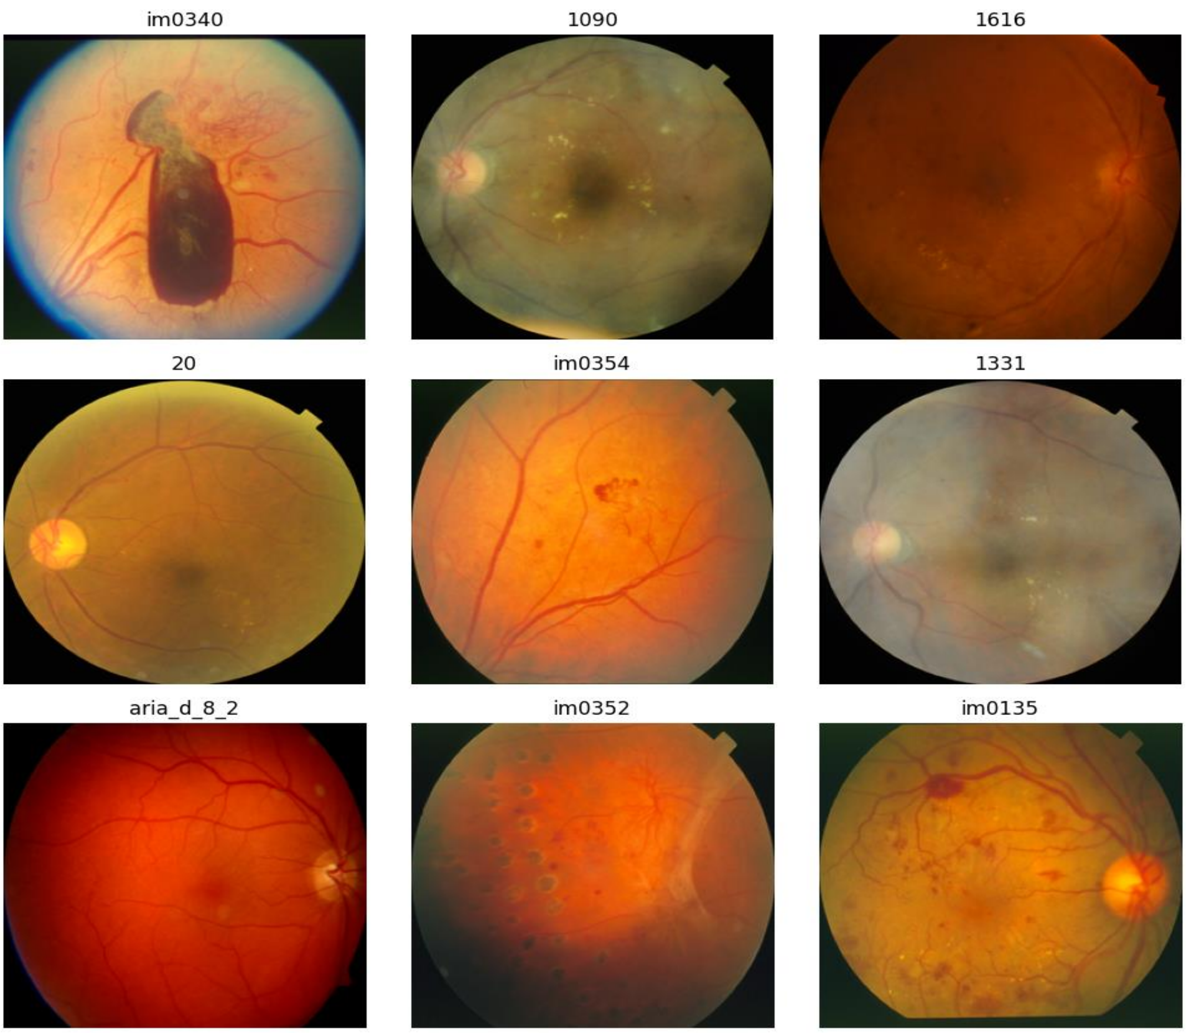

Figure 3: Background Reduction image of Diabetic Retinopathy

Transformation: We have made crucial moves to improve our dataset's quality and balance. To ensure uniform intensity levels among photos, we first performed normalization by rescaling the pixel values of retinal images to a defined range. Second, we used oversampling approaches to correct the class imbalance seen in medical datasets. To balance class distribution, these strategies involved reproducing instances from underrepresented classes, which allowed our model to effectively learn from minority classes. To add variation and realism to the dataset, data augmentation techniques were also used. We increased the variety of the training data by flipping, rotating, and moving the images, which helped our Capsule Network model be more resilient and general.

Modelling: We have created and put into use a complex architecture for the capsule network. The numerous features and patterns inherent in retinal pictures are efficiently captured by this approach, allowing for precise disease classification.

Evaluation: In order to assess the effectiveness of our Capsule Network model, we used a variety of measures, including loss, accuracy, sensitivity, and specificity, which were calculated on a unique validation dataset. Notably, to improve clinical interpretability, we used a decision threshold of 0.5 to convert model outputs into actionable predictions. We also expanded our investigation to include the Area Under the Curve (AUC) for every class, giving us a thorough evaluation of the model's discriminative power. The comprehensive comprehension of the model's accuracy in categorizing retinal illnesses provided by this thorough evaluation approach allows healthcare professionals to make well-informed decisions.

4 Design Specification

Eye images are carefully loaded and go through scaling and normalization during preprocessing. Manual oversampling is used to correct dataset imbalances, which improves the model's ability to identify complex diseases. Robustness is enhanced when the dataset is transformed to add variations. A combination of Convolutional Neural Networks (CNNs) and Capsule Networks is the key development in architecture. Following a Capsule Layer that meticulously preserves spatial hierarchies and a CNN that extracts important characteristics from retinal pictures, the model is made more sensitive to complicated interactions within the data. The built architecture of the model makes use of the Adam optimizer and a dynamic learning rate schedule to delicately balance convergence speed and precision of fine-tuning.

The data augmentation used in the training phase by ImageDataGenerator improves the model's capacity to adapt to real-world settings. Early halting and scheduling learning rates help to hasten convergence while avoiding overfitting. A batch size of 32, training over 50 epochs, and an early stopping patience of 10 are among the hyperparameters. On a validation dataset, the model is thoroughly assessed, producing metrics for loss and accuracy. In addition to these, sensitivity and specificity—two crucial diagnostic indicators—are computed to show how well the model can categorize retinal disorders. Clinical interpretation is aided by using a decision threshold of 0.5 for predictions. Furthermore, the model's ability to identify various pathologies and its ability to discriminate between them are revealed by the calculation of the Area Under the Curve (AUC) for each class. The classification report adds depth to the research by providing detailed insights that cover precision, recall, F1-score, and support for each class. The confusion matrix illustrates the differences between anticipated and actual labels, highlighting the model's advantages and potential weaknesses. In order to process images, the Convolutional Neural Network (CNN) architecture gradually applies 2D convolutional layers with increasing complexity and then uses MaxPooling operations to extract hierarchical features. 32 3x3 kernel filters are used in the first layer, which is then followed by downsampling. Using 64 and 128 filters, the next layers enhance feature extraction, with further pooling coming after each. Flattened features go through two intensive stages of processing that use 128 units and ReLU activations. Multiple labels can be used to classify retinal illnesses in the final dense layer with 7 units and sigmoid activation. With the help of this architecture, which methodically records visual patterns, retinal diseases can be automatically and correctly diagnosed from input photos.

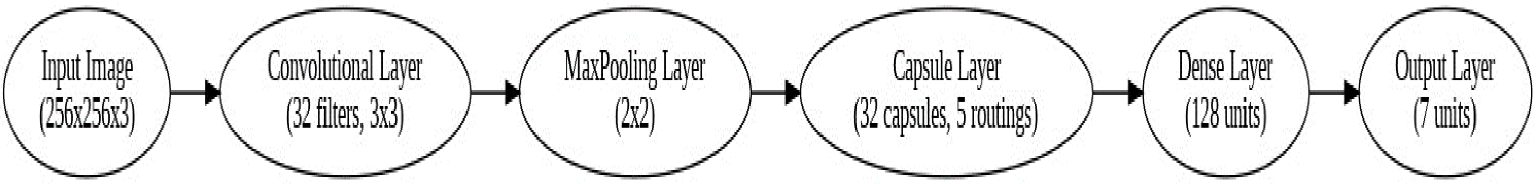

Figure 4: Model Architecture

5 Implementation

The creation of our thorough and precise disease classification system required several important steps in the implementation process. The training and validation datasets, which were derived from CSV files providing the paths to the image files and associated class labels, were first pre-processed. Images were loaded, scaled to 256x256 pixels in standard size, then normalized to range from 0 to 1. This phase made sure that all of the dataset's data representations were consistent. For image processing and transformation, we utilized the PIL module in Python. We manually oversampled each class based on predefined counts to address the issue of class imbalance. In order to balance class distribution and effectively mitigate the consequences of skewed data, this involved choosing instances at random from each class and duplicating them. The TensorFlow Keras ImageDataGenerator was used to add additional data to the oversampled training dataset. By incorporating differences in the training data, augmentation techniques like rotation, horizontal and vertical shifts, shear, and zoom were used to improve model generalization. The network's capacity to handle a variety of retinal pictures was significantly enhanced by this action. In our novel hybrid architecture, a CNN is followed by a capsule network. To gradually extract hierarchical characteristics from input photos, the CNN was built with a number of Convolutional and MaxPooling layers. For the purpose of diagnosing the condition, this was crucial for identifying significant patterns and details. The retrieved features were then flattened and sent through dense layers that were fully connected to aid in feature fusion and learning. An unique method for managing spatial hierarchies and pose relationships between components was presented by the Capsule Network, a distinctive element of our architecture. TensorFlow Keras was used to create the CapsuleLayer, and its parameters were adjusted to improve representation learning and feature routing. The network's capacity to recognize intricate spatial linkages and boost classification precision was greatly aided by the Capsule Network. TensorFlow Keras's Functional API was used to build the complete model architecture, enabling the seamless integration of different parts. To enable multi-label classification, the model was built using the Adam optimizer and binary cross-entropy loss function. We utilized L2 regularization to avoid overfitting. An implementation of a learning rate scheduler callback ensured the best model performance and convergence. Training became consistent and effective because to this dynamic modification of the learning rate based on training progress. The model's capacity to generalize across various samples was further strengthened by the use of enhanced and oversampled training data during training. To avoid overfitting and recover the best model weights, we used an early halting mechanism based on validation loss. Using the validation dataset, the model's performance was evaluated after training. For the classification of diseases, metrics like loss and accuracy were computed, and forecasts were made. The results were in the form of a classification report that comprised the precision, recall, and F1-score for each class as well as a confusion matrix for performance visualization.

6 Evaluation

We thoroughly assess our suggested solution's performance metrics, statistical significance, and usability in real-world situations.

Performance Analysis: Our system for diagnosing retinal diseases performed well, with a validation accuracy of 0.4572 and a validation loss of 0.3053 on the validation dataset. These metrics show the model's propensity to correctly predict the validation data. With the help of measurements for precision, recall, and F1-score, we further examined the model's performance. The model's ability to capture patterns linked to different retinal disorders is demonstrated by the micro-average F1-score of 0.44 and the macro-average F1-score of 0.31 as shown in Figure 5 and Figure 6.


```
Validation Loss: 0.30530643463134766, Validation Accuracy: 0.45720720291137695
14/14 [==============================] - 0s 9ms/step
              precision    recall  f1-score   support

           0       0.66      0.63      0.64        99
           1       0.63      0.51      0.56        98
           2       0.74      0.41      0.53        34
           3       1.00      0.02      0.04        52
           4       0.43      0.10      0.16        31
           5       0.33      0.09      0.15        32
           6       0.43      0.06      0.10        52

   micro avg       0.63      0.34      0.44       398
   macro avg       0.60      0.26      0.31       398
weighted avg       0.63      0.34      0.39       398
 samples avg       0.30      0.28      0.28       398
```

Figure 5: Performance Report

ROC-AUC Analysis: When assessing the model's discriminatory ability, the Receiver Operating Characteristic Area Under Curve (ROC-AUC) is a critical statistic. The model's variable capacities to differentiate between various retinal diseases are revealed by the estimated AUC values for various disease classes. AUC values for DR, NORMAL, MH, and ARMD are particularly high (0.8601, 0.8210, 0.8937, and 0.8095, respectively), showing these classes have excellent predictive performance as shown in Figure 6.

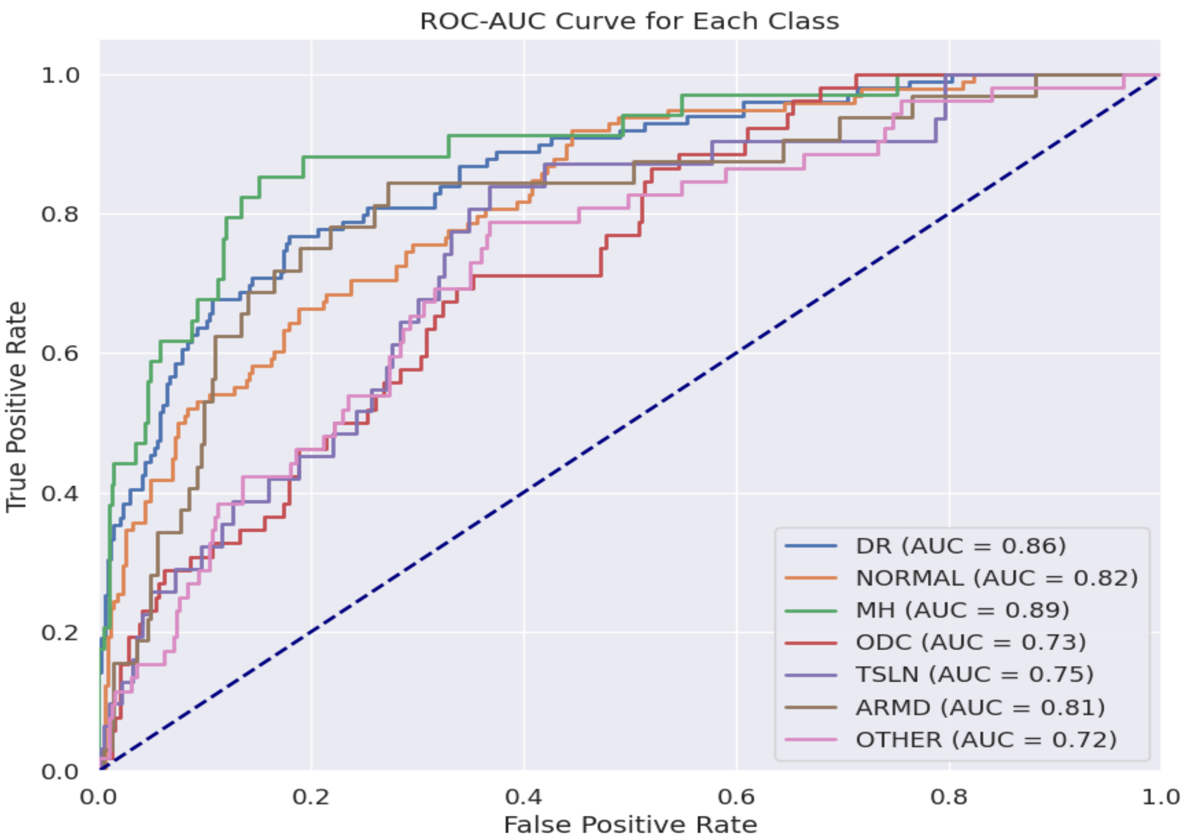

Figure 6: ROC-AUC Curve

Sensitivity and Specificity: Analysing the sensitivity and specificity for each illness class was another aspect of the evaluation. These measures give information on the true positive and true negative rates of the model, respectively. The model's ability to accurately identify positive cases of diabetic retinopathy is demonstrated by the fact that DR has a sensitivity of 0.8905. Similar to NORMAL, NORMAL exhibits good specificity (0.6329), highlighting the model's capacity to categorize normal situations accurately as shown in Figure 7.

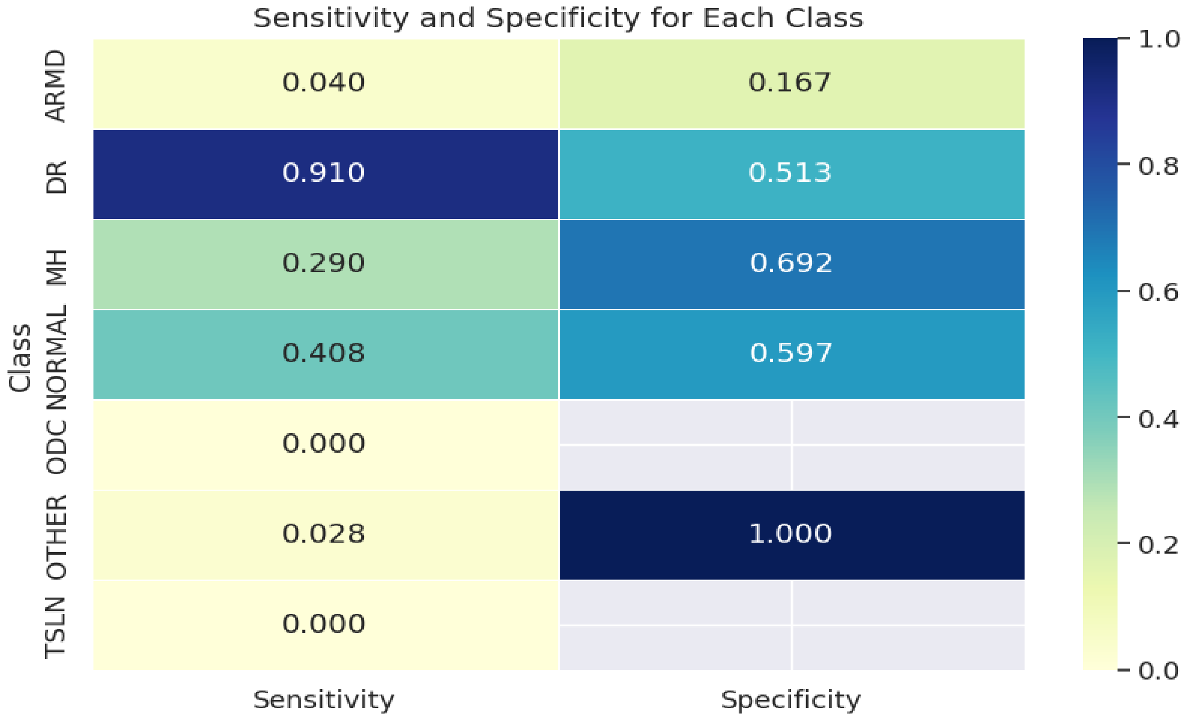

Figure 7: Sensitivity and Specificity

Confusion Matrix and Implications: The classification mistakes for each illness class are broken down in great depth in the confusion matrix. The matrix draws attention to areas that need attention, such as the model's difficulties in differentiating between specific illness categories (such as ODC and TSLN). This understanding of misclassifications can direct the development of new models and tactics for data collection as shown in Figure 8.

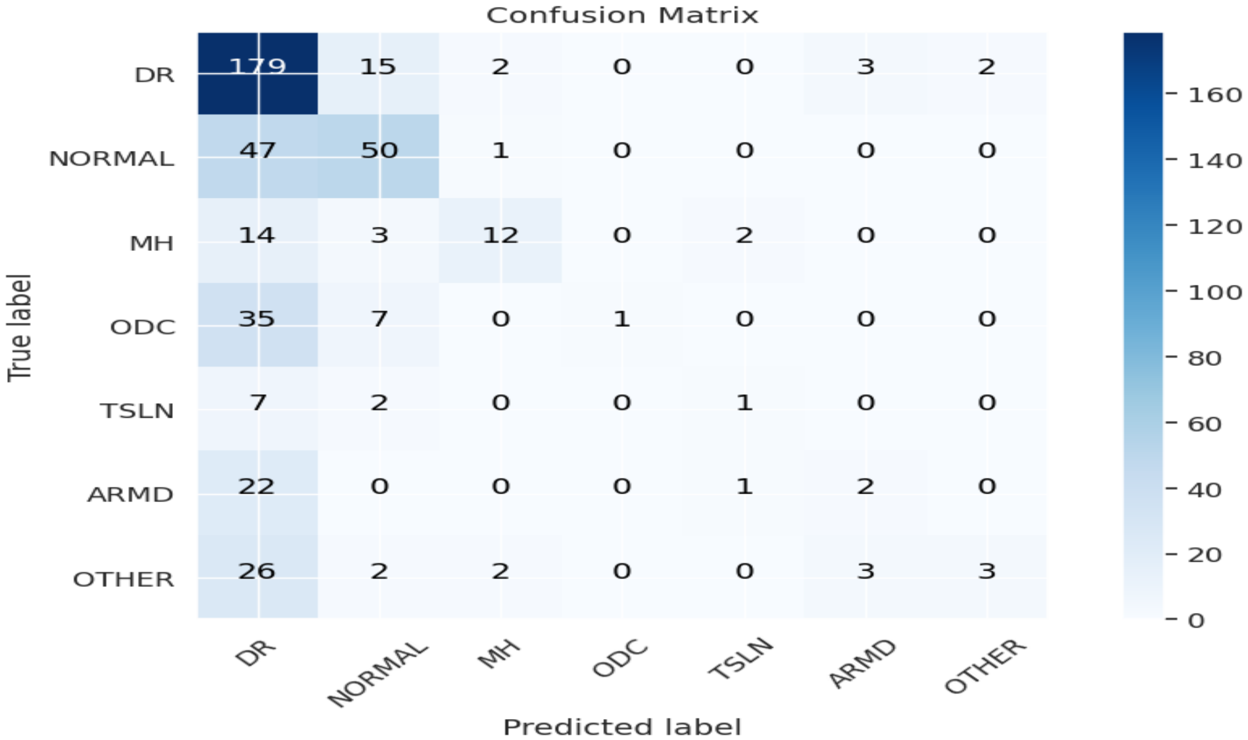

Figure 8: Confusion Matrix

6.1 Discussion

Even though the model showed outstanding accuracy on several classes, like DR and NORMAL, there are some significant areas that call for serious attention and future improvements.

The differences in performance between classes, especially in the cases of ODC and TSLN, emphasize the difficulties presented by unbalanced datasets and the difficulty of differentiating minute visual cues related to these classes. The model struggles to capture the distinctive patterns that are present in ODC and TSLN, as evidenced by the relatively poor sensitivity and specificity for these classes. The necessity for focused efforts to increase the amount of high-quality picture data for ODC and TSLN is highlighted by this gap, which will allow the model to learn more representative and discriminative features.

The model's strong AUC values for DR, NORMAL, MH, and ARMD underline its potency in making these distinctions. However, the model might gain by utilizing more sophisticated architectures or investigating transfer learning from related domains, as indicated by the lower AUC values for TSLN, ODC, and OTHER. Utilizing previously trained models that have performed well in image classification tasks may offer a strong base for capturing the detailed features necessary for precise classification, particularly in difficult classes.

Although the suggested architecture's use of capsule networks presents a potential method for extracting hierarchical features, the implementation as it stands may still need more hyperparameter fine-tuning to improve convergence and overall speed. Additionally, the model's capacity to represent complicated relationships in the data may be impacted by the

routing iterations chosen in the capsule layer. Trying out various routing schemes and looking into different capsule dimensions might result in better capture of class-specific properties.
The results of this study are consistent with earlier studies showing the value of resolving class imbalance and utilizing cutting-edge architectures for precise medical image categorization. It is clear from putting our findings in the context of previous research that improving the performance of the suggested model necessitates a multifaceted strategy that involves data augmentation, transfer learning, and architectural improvement. In classes where visual cues are more subtle, collaboration between data scientists and medical specialists is crucial for directing the model's development to correctly diagnose retinal illnesses.

7 Conclusion and Future Work

By combining convolutional neural networks (CNNs) and capsule networks, this study aims to fundamentally alter how retinal diseases are classified. The study achieves its goals, demonstrating the potential of this hybrid paradigm. It accurately classifies retinal illnesses with remarkable effectiveness, especially in cases of DR and NORMAL, offering a promising diagnostic tool for healthcare professionals. The model struggles with classes like TSLN and ODC, however, underlining the need for increased data collection and model improvement.
Future directions offer opportunities for advancement, possibly boosting feature extraction and classification accuracy through the incorporation of transfer learning and hyperparameter tweaking. The clinical relevance and interpretability of the model could be improved by collaborations with medical professionals. The model's potential commercial uses are also notable; they may completely alter how retinal diseases are diagnosed. As we consider the results, it becomes clear that this study offers a solid foundation. To ensure that the model reaches its full potential in supporting medical professionals and enhancing patient care in retinal illnesses, future research should focus on interpretability, architectural improvements, and validation through medical collaborations.

8 Acknowledgement

I am incredibly appreciative of Mr. Hicham Rifai, my superior, for his constant advice, continuous support, and amazing knowledge. My academic path has been greatly helped by the National College of Ireland's tremendous mentoring and skill-building. I sincerely appreciate my family's constant support and wish them well. Their combined efforts have been crucial to my scholarly efforts and personal development.

References

Mittal, K., Rajam, V.M.A. Computerized retinal image analysis - a survey. *Multimed Tools Appl* **79**, 22389–22421 (2020). https://doi.org/10.1007/s11042-020-09041-y

Patrick, M.K., Adekoya, A.F., Mighty, A.A., & Edward, B.Y. (2019). Capsule Networks - A survey. Journal of King Saud University - Computer and Information Sciences, 34, 1295-1310.

Ramanathan, G., Chakrabarti, D., Patil, A., Rishipathak, S., & Kharche, S. (2021). Eye Disease Detection Using Machine Learning. In 2021 2nd Global Conference for Advancement in Technology (GCAT), pp. 1-5, Bangalore, India. doi:10.1109/GCAT52182.2021.9587740.

Malik, S., Kanwal, N., Asghar, M. N., Sadiq, M. A. A., Karamat, I., & Fleury, M. (2019). Data Driven Approach for Eye Disease Classification with Machine Learning. Applied Sciences, 9(14), 2789. https://doi.org/10.3390/app9142789

Purandare, M., & Noronha, K. (2016). Hybrid system for automatic classification of Diabetic Retinopathy using fundus images. In 2016 Online International Conference on Green Engineering and Technologies (IC-GET), pp. 1-5, Coimbatore, India. doi:10.1109/GET.2016.7916623.

Chellaswamy, C., Geetha, T. S., Ramasubramanian, B., Abirami, R., Archana, B., & Divya Bharathi, A. (2022). Optimized Convolutional Neural Network based Multiple Eye Disease Detection and Information Sharing System. In 2022 6th International Conference on Intelligent Computing and Control Systems (ICICCS), pp. 1105-1113, Madurai, India. doi:10.1109/ICICCS53718.2022.9788334.

V. R, N. Muthukumaran, and M. P. Austin. "Hybrid ResNet with Bidirectional LSTM for Eye Disease Classification with Evaluation Optimizers Techniques," In 2023 International Conference on Inventive Computation Technologies (ICICT), pp. 1624-1630, Lalitpur, Nepal, 2023. doi: 10.1109/ICICT57646.2023.10134223.

A. Raza, M. U. Khan, Z. Saeed, S. Samer, A. Mobeen, and A. Samer. "Classification of Eye Diseases and Detection of Cataract using Digital Fundus Imaging (DFI) and Inception-V4 Deep Learning Model," In 2021 International Conference on Frontiers of Information Technology (FIT), pp. 137-142, Islamabad, Pakistan, 2021. doi: 10.1109/FIT53504.2021.00034.

Prananda, A. R., Frannita, E. L., Hutami, A. H. T., Maarif, M. R., Fitriyani, N. L., & Syafrudin, M. (2023). Retinal Nerve Fiber Layer Analysis Using Deep Learning to Improve Glaucoma Detection in Eye Disease Assessment. Applied Sciences, 13(1), 37. https://doi.org/10.3390/app13010037.

Rekha, C., & Jayashree, K. (2022). Hyphema Eye Disease Prediction with Deep Learning. In 2022 International Conference on Computer, Power and Communications (ICCPC) (pp. 215-218). Chennai, India. IEEE. https://doi.org/10.1109/ICCPC55978.2022.10072218.

Santos, P. R. S. dos, Brito, V. de C., Carvalho Filho, A. O. de, Duarte de Araújo, F. H., Rabêlo, R. de A. L., & Mathew, M. J. (2020). A Capsule Network-based for Identification of Glaucoma in Retinal Images. In 2020 IEEE Symposium on Computers and Communications (ISCC) (pp. 1-6). DOI: 10.1109/ISCC50000.2020.9219708.

Nneji, G. U., et al. (2021) "A Dual Weighted Shared Capsule Network for Diabetic Retinopathy Fundus Classification." 2021 International Conference on High Performance Big Data and Intelligent Systems (HPBD&IS), Macau, China, 2021, pp. 297-302. doi: 10.1109/HPBDIS53214.2021.9658352.

Jain, L., Murthy, H. V. S., Patel, C., and Bansal, D. (2018). Retinal Eye Disease Detection Using Deep Learning. In: 2018 Fourteenth International Conference on Information Processing (ICINPRO), pp. 1-6. Bangalore, India. doi:10.1109/ICINPRO43533.2018.9096838.

Guo, C., Yu, M., & Li, J. (2021). Prediction of Different Eye Diseases Based on Fundus Photography via Deep Transfer Learning. Journal of Clinical Medicine, 10(23), 5481. https://doi.org/10.3390/jcm10235481

Swathi, K., Joshua, E.S.N., Reddy, B.D. and Rao, N.T., 2022. Diabetic Retinopathy Detection Using Deep Learning. In: International Conference on Advancements in Smart, Secure and Intelligent Computing. IEEE, pp. 1-5. Available at: https://doi.org/10.1109/ASSIC55218.2022.10088331.