

ABSTRACT

Diabetic Retinopathy (DR) is a progressive microvascular complication of diabetes and a leading cause of vision impairment and blindness worldwide. Early detection and accurate classification of DR are essential for timely medical intervention and prevention of severe visual deterioration. Traditional diagnostic approaches rely on manual grading by ophthalmologists, which is time-consuming, labor-intensive, and susceptible to inter-observer variability. Recent advancements in Deep Learning (DL) have demonstrated significant potential in automating DR classification; however, existing models face challenges related to overfitting, computational complexity, and generalization across diverse datasets. This research aims to develop a robust and computationally efficient DR classification system capable of accurately detecting all four stages of DR like normal, mild, moderate and severe while addressing these challenges.

The proposed methodology follows a three step approach to enhance model performance and adaptability. In the first phase, Transfer Learning (TL) is employed to fine-tune state-of-the-art pretrained models, particularly EfficientNetV2L, using the Asia Pacific Tele-Ophthalmology Society (APTOS) dataset. The objective is to utilize deep feature extraction capabilities while ensuring improved generalization. Data augmentation and scaling techniques are incorporated to mitigate overfitting and enhance model robustness. In the second phase, a custom Convolutional Neural Network (CNN) architecture is designed to optimize feature extraction and computational efficiency. The architecture incorporates structured convolutional layers, batch normalization, and dropout mechanisms to enhance learning stability while reducing the risk of overfitting.

In the third phase, a hybrid model is developed by integrating the EfficientNetV2L with the custom CNN architecture, thereby utilising the strengths of both architectures. Extensive hyperparameter tuning is performed, including learning rate scheduling, dropout regularization, batch size optimization, and adaptive optimization using Adam optimizer, to ensure efficient training and convergence. The hybrid model is designed to provide a balance between high classification accuracy and computational feasibility, facilitating its deployment in clinical applications.

The proposed system is evaluated using key performance metrics, including accuracy, precision, recall, F1-score, and computational complexity (trainable parameters).

The hybrid model (EfficientNetV2L + Custom CNN Model) achieves the highest classification accuracy of 94%, with a precision of 94%, recall of 93%, and F1-score of 94%, utilizing 441,087 trainable parameters. Comparative analysis with benchmark models demonstrates the superiority of the proposed approach, with EfficientNetV2L achieving 92% accuracy, the Custom CNN Model reaching 88% accuracy, while conventional architectures such as VGG16 (85%), ResNet-50 (87%), InceptionV3 (89%), and DenseNet121 (90%) exhibit relatively lower performance. These findings confirm the effectiveness of the hybrid approach in achieving state-of-the-art classification accuracy while maintaining computational performance.

The developed hybrid model is evaluated against an existing architecture that integrates DenseNet121, Xception, and EfficientNetB3, comprising 2,361,860 trainable parameters. The existing model achieves an accuracy of 75%, precision of 73%, recall of 75%, and an F1 score of 71%. In contrast, the proposed model, with only 441,087 parameters, demonstrates superior performance achieving a 19% increase in accuracy, 21% in precision, 18% in recall, and 23% in F1 score highlighting its effectiveness despite a significantly reduced parameter count.

This research contributes to the advancement of automated DR classification by integrating fine-tuned pretrained models with a lightweight custom CNN, resulting in a scalable, efficient, and clinically viable diagnostic framework. The system is validated using clinical data to ensure its applicability in practical healthcare settings, particularly in resource-constrained environments where access to high-end computational resources is limited. The proposed model has the potential to significantly enhance early DR detection, improve diagnostic consistency, and ultimately contribute to better patient outcomes in ophthalmic disease management.