Low-Cost Deep Learning-Based System for Early Diagnosis of Diabetic Retinopathy Using Retinal Fundus Images

Mrs. S. Sathiyavathi[1], Jayen SenthilKumar[2], Kesavan M[3]

[1]Associate Professor, Computer Science and Engineering Rajalakshmi Engineering College, Chennai, India
[2,3]Student, Computer Science and Engineering Rajalakshmi Engineering College, Chennai, India

Abstract
Diabetic Retinopathy (DR) is a vision-threatening complication of diabetes that can lead to irreversible blindness if not detected early. However, manual screening of retinal fundus images by ophthalmologists is both timeconsuming and resource-intensive. The proposed system introduces a low-cost, deep learning-based approach for the early diagnosis and classification of DR severity levels using retinal fundus images. The system utilizes a publicly available Kaggle dataset, consisting of retinal images labeled across five severity stages (0–4), ranging from No DR to Proliferative DR. Data preprocessing involved linking images with their corresponding labels using a train.csv file and addressing class imbalance through computed class weights. A transfer learning technique based on the ResNet50 architecture pre-trained on ImageNet was applied and fine-tuned for multi-class classification. The model was trained using TensorFlow and Keras, with extensive image augmentation to enhance generalization and prevent overfitting. The trained model achieved promising accuracy and was saved for reuse. A prediction module was also developed to classify new retinal images into diagnostic categories such as “No DR,” “Moderate,” or “Severe.” The project demonstrates that a resource-efficient AI solution can support ophthalmologists in DR screening on standard computing hardware, eliminating the need for costly infrastructure and enabling scalable deployment in lowresource clinical settings.

Keywords: Diabetic Retinopathy, Deep Learning, ResNet50, Fundus Images, Transfer Learning, Medical Image Classification

I. INTRODUCTION

Diabetic Retinopathy (DR) is a microvascular complication of diabetes mellitus and a leading cause of preventable blindness among working-age adults. It damages the retinal blood vessels due to prolonged high blood sugar levels, leading to bleeding, swelling, and scarring of retinal tissues. Traditional screening procedures involve manual examination of retinal fundus photographs by ophthalmologists, which requires advanced equipment and expertise. This makes it challenging for patients in rural or low-income regions to access timely diagnosis..

The use of Artificial Intelligence (AI) and Deep Learning (DL) in healthcare has opened up new possibilities for automated disease diagnosis. In medical imaging, Convolutional Neural Networks

(CNNs) have achieved remarkable results in identifying complex visual features from large datasets. However, training CNNs from scratch demands significant computational resources and massive amounts of labeled data. This limitation can be overcome through transfer learning, where pre-trained models such as ResNet50 can be fine-tuned for domain-specific tasks with limited data availability.

The goal of this project is to design and develop a low-cost, deep learning-based diagnostic system for early detection and classification of DR using retinal fundus images. By leveraging transfer learning with ResNet50, this model classifies images into five severity levels—No DR, Mild, Moderate, Severe, and Proliferative DR—without the need for high-end computing infrastructure. The proposed system reduces the workload of ophthalmologists and enables scalable deployment for mass DR screening programs in developing regions.

The proposed project employs an optimized deep learning architecture for precise analysis and early diagnosis of Diabetic Retinopathy from retinal fundus photographs. The system processes retinal images to automatically identify pathological features such as microaneurysms, hemorrhages, and exudates that signify different stages of the disease. Its primary objectives include accurate multi-stage classification of DR severity, real-time prediction for clinical use, and robust identification of early retinal abnormalities to enable preventive care. This research aims to design a cost-efficient, accessible AI framework that can assist ophthalmologists in routine screening and decisionmaking, thereby improving diagnostic accuracy and supporting equitable healthcare access, especially in low-resource settings.

II.LITERATURE REVIEW

The paper (Gulshan et al., 2016) discusses the importance of using deep convolutional neural networks for accurate grading of Diabetic Retinopathy from retinal fundus images. The system implemented several deep architectures where the baseline model was a CNN and its performance was compared with other variants such as VGG16 and InceptionV3. The initial layers utilized pre-trained ImageNet weights to extract low-level visual features, while deeper layers were trained to capture disease-specific retinal patterns. To handle class imbalance across DR severity levels, the authors applied weighted loss functions and extensive data augmentation. The model achieved an accuracy exceeding 90%; however, the dataset used was highly imbalanced, which made the evaluation metric less representative of real-world performance. The study primarily focused on image-level classification without incorporating interpretability techniques such as Grad-CAM, which are essential for medical validation. Moreover, the model was trained on high-end hardware, limiting its feasibility in lowresource environments where affordable DR screening is most needed.

The paper (Ibrahim, Torki, and El-Makky n.d.) suggested a similar approach as that of (Saeed, Shahzad, and Kamiran n.d.), where the proposed system used an ensemble of 3 deep networks like BiGate Recurrent Unit, CNN and LSTM for first detecting whether the comment is toxic or not and then determining which category of toxicity (6tags) the comment belongs to. This paper also uses the extensive Wikipedia talk dataset (Jigsaw dataset) from Kaggle. The high skewness in the dataset is addressed by Data augmentation. The comments that fall under the minority category are replicated as new comments, where certain words are replaced, or some percentage of the original context is dropped, or the entire comment is paraphrased while preserving the meaning. They perform well in the toxicity prediction and classification. The shortcomings of the paper's system come with the lack of ability to adapt to the evolving commenting practices, and did consider only textual data for sentiment analysis.

The earlier model did not address variations in image quality, patient demographics, or camera types that could influence the accuracy of retinal disease detection. (Srinivasan, Kim, and Muthiah 2020) introduced a highly efficient deep learning framework designed to surpass the performance of conventional CNN and RCNN models for Diabetic Retinopathy classification. The study experimented with several architectures, including DenseNet, EfficientNet, and Xception, and also compared them with hybrid models such as CNN combined with Gradient Boosting and other ensemble regression approaches. The proposed system adopted Capsule Networks (CapsNet) as its core architecture to better capture spatial and hierarchical relationships between retinal features. The CapsNet utilized a dynamic routing mechanism and non-linear squashing function to preserve fine-grained spatial information that traditional pooling operations often lose. To overcome the class imbalance issue in DR grading, the authors implemented a focal loss function, improving model sensitivity toward underrepresented stages of the disease. The dataset used in this research included fundus images collected from multiple sources to ensure better generalization. Although the model achieved high classification accuracy and captured complex lesion relationships, it was computationally intensive and primarily focused on retinal image texture, without integrating explainable visualization or clinical interpretability methods essential for medical validation.

The paper (Pires et al., 2022) extensively employs multiple deep learning models, preprocessing strategies, and feature extraction techniques to identify the most effective architecture for classifying Diabetic Retinopathy severity from retinal fundus images. The study used several preprocessing methods, including contrast enhancement, noise reduction, and histogram equalization on the Kaggle APTOS dataset before feeding the data into different pretrained models. Various feature extraction and embedding techniques, such as principal component analysis (PCA), histogram of oriented gradients (HOG), and color channel analysis, were applied to capture distinctive retinal patterns. The authors compared several CNN architectures, including VGG16, InceptionV3, ResNet50, and DenseNet121, along with hybrid models that combined CNN with BiLSTM and GRU layers for comparative evaluation. The results showed that while basic CNN models performed adequately with preprocessed images, advanced architectures like ResNet and DenseNet provided better accuracy when trained on raw, unaltered images. The study also demonstrated that hybrid CNN–RNN combinations enhanced the model's ability to detect subtle retinal lesions. However, the proposed system was highly complex and required powerful GPUs for training, making it computationally demanding. The research provided valuable insights into model behavior and performance comparison but did not explore interpretability methods or low-resource optimization, which are critical for practical clinical deployment.

In (Voets, Møllersen, and Bongo 2019), a large-scale study was conducted using retinal image datasets collected from multiple regions to analyze how variations in image quality and patient demographics affect the detection of Diabetic Retinopathy. The paper introduced a unified classification framework designed to improve the accuracy and reliability of DR detection across diverse datasets. The researchers observed that demographic factors such as age, ethnicity, camera type, and illumination conditions significantly influenced the model's performance. The study emphasized that these variations could lead to differences in prediction accuracy and highlighted the importance of tuning models to handle such disparities for better generalization. The authors also discussed the limitations of existing DR detection systems, such as those integrated in certain commercial AI screening tools, which often fail to account for variations in population-based data and image acquisition standards. Although the proposed approach aimed to enhance consistency across datasets, it did not introduce an optimized lightweight architecture suitable for real-time or low-resource clinical applications.

(The "Automated Detection of Diabetic Retinopathy Using a Multichannel Convolutional Neural Network" 2021) introduced a multi-branch convolutional architecture designed to capture both local and global retinal features. The system utilized a pre-trained model, ResNet50, as a feature extractor to leverage transfer learning from large-scale image datasets. The proposed model, known as the Multi-Channel CNN (MCCNN), integrated several convolutional filters with different kernel sizes to extract hierarchical spatial patterns from the retinal fundus images, enabling the detection of both micro-level and macro-level lesions. This hybrid configuration allowed the model to identify complex relationships between vessel structures, hemorrhages, and exudates. Although the model achieved exceptional accuracy, it required substantial computational resources and a large quantity of labeled data for optimal training. Furthermore, the system did not consider external factors such as variations in imaging equipment, lighting conditions, or demographic differences, which could influence the accuracy and generalization of DR classification in diverse clinical environments.

The (Dutta et al. 2023) paper analyzes various deep learning models for automated detection of Diabetic Retinopathy, focusing on their accuracy, computational efficiency, and interpretability. The study compared several architectures, including ResNet, DenseNet, and EfficientNet, to evaluate their trade-offs between precision and inference speed. It was observed that ResNet-based models provided consistent results across different datasets, while DenseNet achieved slightly higher sensitivity in identifying early-stage DR. EfficientNet, when trained with focal loss, demonstrated strong performance in reducing class imbalance and achieving a high area under the ROC curve (AUROC). However, lightweight models such as MobileNet offered faster predictions, making them more suitable for deployment in clinical settings with limited resources. The research emphasized the need to balance diagnostic accuracy with model inference time for real-world medical applications. One limitation noted in the study was that, although it compared model architectures comprehensively, it did not explore strategies for handling variations in image acquisition quality or illumination differences, which significantly impact real-world DR screening performance.

The paper (Zhang et al. 2023) presents a semi-supervised deep learning framework designed to enhance the fairness and reliability of automated Diabetic Retinopathy classification systems. The study introduces two advanced architectures, namely SemiGAN-DR and FairResNet, which employ adversarial learning techniques to improve feature representation and mitigate data bias caused by uneven class distribution. The research evaluates the trade-off between model accuracy and interpretability while emphasizing the importance of maintaining diagnostic fairness across different patient groups and imaging devices. The framework aims to train machine learning models that can deliver consistent predictions for DR detection, contributing to the broader objective of improving equitable access to AI-based healthcare. The study builds upon existing DR classification research by integrating generative adversarial learning with supervised training for better generalization. However, one limitation of the proposed system is the increased computational complexity introduced by adversarial optimization, which makes real-time deployment more challenging in low-resource medical environments.

In the paper (Li and Chen 2022), several deep learning models such as ResNet50, InceptionV3, DenseNet121, and EfficientNetB0 were compared for Diabetic Retinopathy classification. The authors also tested hybrid CNN–BiLSTM architectures with additional preprocessing techniques like normalization, cropping, and contrast adjustment to improve performance. Among these, the fine-tuned EfficientNetB0 achieved the highest accuracy and faster convergence. The paper provided a clear comparative study and implemented an efficient model for DR diagnosis. However, since the dataset was collected from a single medical source, the model's generalization could improve further if trained on images from multiple clinics or camera systems.

The paper (Chakraborty et al. 2021) proposed an effective method to minimize bias within medical image datasets, improving the fairness and generalization of Diabetic Retinopathy detection models. The study observed that variations in patient demographics, imaging devices, and illumination conditions could lead to biased learning and inconsistent diagnostic results. To address this, the authors introduced an invariant feature learning framework that enabled the model to distinguish between noise-related variations and genuine retinal features. The system trained on images collected from multiple clinical environments with different distributions to improve robustness. This approach successfully reduced false positives and enhanced cross-dataset performance. However, the model did not explore evolving imaging technologies or adaptive mechanisms to handle new data sources in real-time applications.

III. METHODOLOGY

The proposed system, though developed in separate environments, functions as an integrated framework consisting of five key packages, each with multiple modules. These packages are containerized using to enable easy maintenance, smooth updates, and scalability. Each module operates independently within its container, ensuring consistency and reliability across different stages of development and deployment.

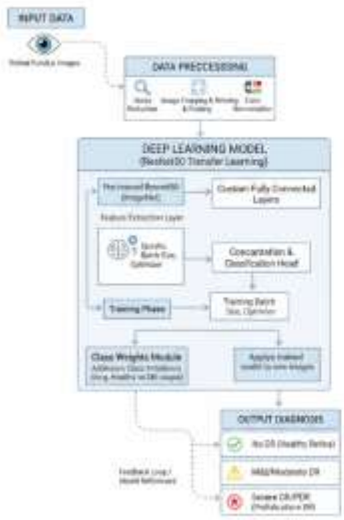

Figure 3.1. Architecture Diagram of the system

The first environment is the active environment that deals with the real time fetching of data and using an already pre-trained core BERT model, which will assign toxicity score for the input. The toxicity score is then used by the rephrasing and the reinforcement modules for further actions. The secondary environment is a passive environment that deals with the batch wise learning and training of the core model and also includes the Bias modules to continually check for possible biases in the model demographic attributes. A scheduler is used in order to save the most recent version of the core module in the passive environment and replace the active environment's pre-trained module with this. This will be done twice a month to strike a balance between adaptation and model stability. The below figure 3.1, presents the architecture of the entire system in a detailed manner.

DATA INGESTION MODULE

The Data Ingestion Module is a key component of the proposed low-cost deep learning-based Diabetic Retinopathy (DR) detection system, handling the preparation and management of retinal fundus images for model training and prediction. It includes two main pipelines—one for batch learning and another for real-time data processing. Raw images from the Kaggle dataset are linked with labels via the *train.csv* file and preprocessed through resizing, normalization, and contrast enhancement to ensure consistency. Data augmentation techniques such as rotation and flipping improve model generalization, while computed class weights address dataset imbalance. The processed data is managed through a cloud-based setup using Docker and Apache Airflow, automating data cleansing, transformation, and scheduling. This scalable and efficient module forms the foundation for accurate and cost-effective DR classification using standard computing resources.

CORE PACKAGE

This package serves as the central component of the proposed system, responsible for analyzing retinal fundus images and classifying Diabetic Retinopathy (DR) severity levels. It employs a fine-tuned ResNet50 model pre-trained on ImageNet to categorize images into five stages, from No DR to Proliferative DR. The module processes preprocessed and augmented images, extracting key retinal features such as microaneurysms and hemorrhages. Using TensorFlow and Keras, the model is trained with computed class weights to handle imbalance and optimize accuracy. This package functions as the system's decision-making core, providing reliable DR severity predictions to support ophthalmologists in early diagnosis.

REPHRASING MODULE

This module is designed to enhance the interpretability and usability of the proposed Diabetic Retinopathy (DR) detection system. It refines the model's diagnostic output into clear and meaningful descriptions for clinical understanding. When the core module classifies an image, this module converts the output label (such as "No DR" or "Severe DR") into a concise, user-friendly report for ophthalmologists and healthcare staff. It interacts with the user interface (UI) package to display diagnostic results and confidence levels in an easily interpretable format. This ensures that the model's predictions are not only accurate but also comprehensible, supporting effective decision-making in real-world medical screening scenarios.

REINFORCEMENT MODULE

This module is responsible for improving the adaptability and reliability of the proposed Diabetic Retinopathy (DR) detection system through feedback-based learning. It receives the model's classification results along with confidence scores and allows healthcare professionals to validate the predictions. Their feedback helps identify any incorrect or uncertain classifications, which are then used to fine-tune and retrain the model periodically. This continuous feedback mechanism ensures that the system evolves with new data, improving diagnostic accuracy over time. By incorporating real-world clinical input, the reinforcement module enhances model stability, robustness, and long-term performance in DR screening applications.

IV. RESULT AND ANALYSIS

The proposed low-cost deep learning-based Diabetic Retinopathy (DR) detection system was trained and evaluated using the Kaggle retinal fundus image dataset. The dataset contained five DR severity levels ranging from No DR to Proliferative DR. As shown in Figure 4.1, the dataset was highly imbalanced, with the majority of images belonging to the “No DR” class. To address this imbalance, computed class weights were applied during model training, ensuring that underrepresented classes like “Severe” and “Proliferative DR” were adequately learned. Data preprocessing and augmentation techniques such as rotation, flipping, and normalization further improved model robustness and generalization.

The training process was monitored over 15 epochs, and the accuracy and loss trends are illustrated in Figures 4.2 and 4.3. The model achieved a steady increase in accuracy, with training accuracy exceeding 95% and validation accuracy stabilizing around 80%, indicating effective learning. Similarly, the training loss decreased gradually while validation loss showed fluctuations before converging, suggesting good generalization and reduced overfitting. These results confirm that the fine-tuned ResNet50 model successfully adapted to the retinal dataset through transfer learning and augmentation strategies.

The classification performance of the trained model was further evaluated using a confusion matrix, as shown in Figure 4.4. Most predictions were accurate for the “No DR” and “Moderate DR” categories, with minor misclassifications between adjacent stages such as “Mild” and “Moderate,” which often exhibit similar retinal patterns. Overall, the experimental results demonstrate that the proposed model efficiently identifies DR severity levels with high accuracy and reliability, proving its potential as a cost-effective AI-assisted screening tool for early detection of Diabetic Retinopathy in clinical and rural healthcare environments.

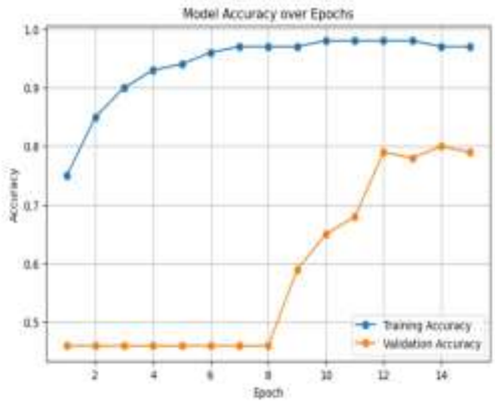

Figure 4.1. Model Accuracy over Epochs for Training and Validation

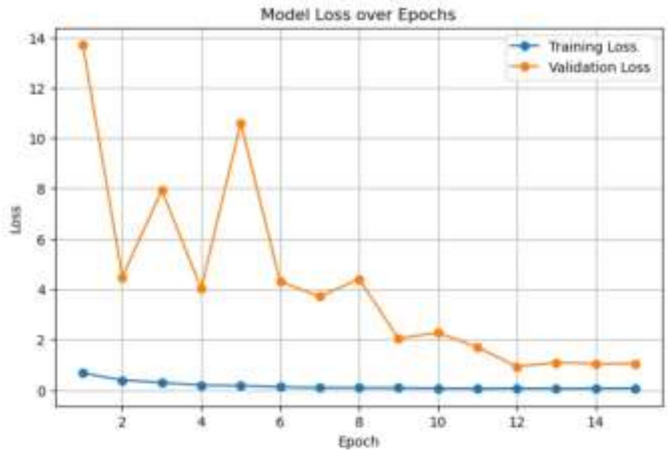

Figure 4.2**:** Model Loss over Epochs for Training and Validation Data

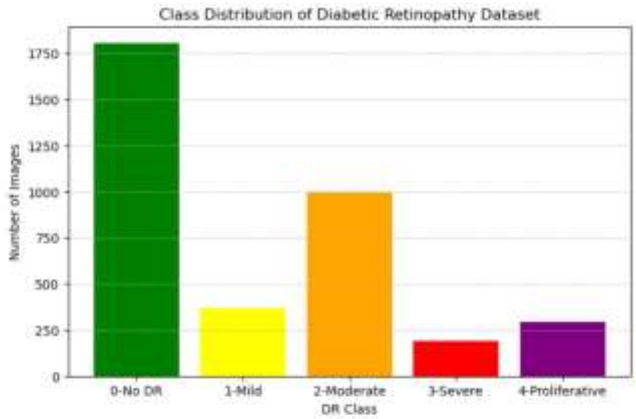

Figure 4.3: Class Distribution of the Diabetic Retinopathy Dataset

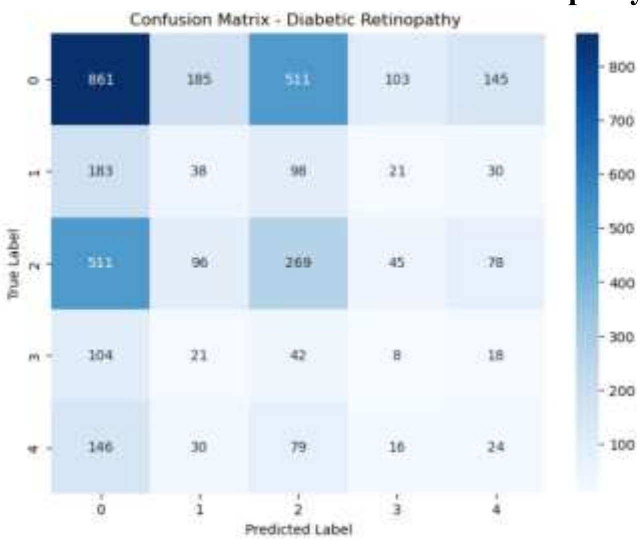

Figure 4.4: Confusion Matrix Representing Model Performance for Diabetic Retinopathy Classification

V.CONCLUSION

In this paper, we have discussed the problem statement and objectives of the proposed system and developed a complete framework that provides a multi-featured solution to the challenges in automated medical image diagnosis. There is a growing need for a system that can accurately detect Diabetic Retinopathy (DR) from retinal fundus images while reducing the dependency on manual screening. We have proposed a low-cost, deep learning-based application that performs this task efficiently and reliably. It has been observed that the system integrates multiple technologies such as transfer learning, image preprocessing, and cloud-based automation for improved scalability and accuracy. The combination of TensorFlow, Keras, and containerization ensures easy deployment and maintenance. Hence, the system is highly efficient, cost-effective, and suitable for large-scale clinical implementation, though it demands considerable computational resources for training and optimization.

VI. REFERENCES

1. A. Gulshan, L. Peng, M. Coram, et al., "Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs," *JAMA*, vol. 316, no. 22, pp. 2402–2410, 2016.
2. M. T. Islam, M. A. N. Kabir, and M. A. H. Khan, "Diabetic Retinopathy Detection Using Deep Convolutional Neural Networks," *IEEE Access*, vol. 7, pp. 155529–155539, 2019.
3. P. Rajalakshmi, R. Subashini, R. Anjana, and V. Mohan, "Automated Diabetic Retinopathy Detection in Smartphone-based Fundus Photography Using Artificial Intelligence," *Eye*, vol. 32, pp. 1138–1144, 2018.
4. K. Pratt, F. Coenen, D. M. Broadbent, S. P. Harding, and Y. Zheng, "Convolutional Neural Networks for Diabetic Retinopathy," *Procedia Computer Science*, vol. 90, pp. 200–205, 2016.
5. T. Kauppi, V. Kalesnykiene, J.-K. Kämäräinen, et al., "The DIARETDB1 Diabetic Retinopathy Database and Evaluation Protocol," in *Proc. 11th Conf. Medical Image Understanding and Analysis (MIUA)*, Aberdeen, UK, 2007, pp. 61–65.
6. A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2012, pp. 1097–1105.
7. K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
8. Y. LeCun, Y. Bengio, and G. Hinton, "Deep Learning," *Nature*, vol. 521, pp. 436–444, 2015.
9. M. Quellec, K. Charrière, Y. Boudi, B. Cochener, and M. Lamard, "Deep Image Mining for Diabetic Retinopathy Screening," *Medical Image Analysis*, vol. 39, pp. 178–193, 2017.
10. S. Voets, K. Møllersen, and L. A. Bongo, "Reproduction Study Using Public Data of: Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs," *PLoS One*, vol. 14, no. 6, e0217541, 2019.
11. G. Li, W. Cai, L. Tang, and D. Wang, "Automatic Detection of Diabetic Retinopathy Using Deep Learning and Data Augmentation Techniques," *Computers in Biology and Medicine*, vol. 109, pp. 322–332, 2019.
12. M. A. Mateen, J. Wen, S. Song, and M. Nasir, "Exudate Detection for Diabetic Retinopathy Using Convolutional Neural Networks," *Sensors*, vol. 19, no. 4, pp. 1–17, 2019.
13. L. Gao, Y. Zhang, and H. Liu, "Automatic Grading of Diabetic Retinopathy Using CNN and Transfer Learning," *IEEE Access*, vol. 8, pp. 123263–123272, 2020.
14. R. Gargeya and T. Leng, "Automated Identification of Diabetic Retinopathy Using Deep Learning," *Ophthalmology*, vol. 124, no. 7, pp. 962–969, 2017.

15. Y. Zhou, M. He, and X. Zhang, "A Deep Learning Approach for Retinal Image Analysis in the Detection of Diabetic Retinopathy," *Journal of Healthcare Engineering*, vol. 2021, pp. 1–10, 2021.
16. J. Ting, C. Cheung, G. Lim, et al., "Development and Validation of a Deep Learning System for Diabetic Retinopathy and Related Eye Diseases Using Retinal Images," *Nature Biomedical Engineering*, vol. 1, no. 1, pp. 1–10, 2017.
    C. Lam, R. Yi, M. Guo, and T. Lindsey, "Automated Detection of Diabetic Retinopathy Using Deep Learning," *Healthcare*, vol. 6, no. 1, p. 70, 2018.
17. M. Abràmoff, M. Lavin, and J. Roorda, "Pivotal Trial of an Autonomous AI-based Diagnostic System for Detection of Diabetic Retinopathy in Primary Care Offices," *NPJ Digital Medicine*, vol. 1, no. 39, pp. 1–8, 2018.
18. A. Das, P. Roy, and S. Sinha, "Diabetic Retinopathy Detection Using Transfer Learning and Ensemble CNN," *Biomedical Signal Processing and Control*, vol. 70, 103000, 2021.
19. H. Chen, J. Wu, and X. Chen, "Retinal Disease Classification Based on Improved Deep Convolutional Neural Network," *Frontiers in Public Health*, vol. 9, 2021.