Automated Diabetic Eye Disease Diagnosis Using Deep Neural Networks

Ch. Poojitha[1], J. Sahithyapriya[2], P. Akhila[3], P. VarunSandhesh[4],
B. NavyaSri[5], J. Niveditha[6], Dr. S ShivaPrasad[7]
[1,2,3,4,5]*Student, Department of CSE(DataScience), MallaReddy Engineering college, Secunderabad*
[6]*Assistant Professor, Department of CSE(DataScience), MallaReddy Engineering college, Secunderabad*
[7]*Professor, Department of CSE(DataScience), MallaReddy Engineering college, Secunderabad*

***Abstract*—Diabetic eye diseases such as diabetic retinopathy, diabetic macular edema, cataract, and glaucoma are among the leading causes of vision impairment worldwide. Early detection of these conditions is crucial to prevent irreversible vision loss; however, traditional diagnosis relies heavily on manual examination by ophthalmologists, which is time-consuming and limited by the availability of specialists. To address this challenge, this project proposes DeepDiabetic, an automated deep learning–based system for the identification and classification of diabetic eye diseases using retinal fundus images. The system utilizes a multi-source dataset consisting of 5,899 retinal images and applies advanced deep learning models including EfficientNetB0, VGG16, ResNet152V2, and hybrid architectures integrated with GRU and BiGRU. Comprehensive preprocessing, data augmentation, and model optimization techniques are employed to improve accuracy and generalization. Experimental evaluation shows that EfficientNetB0 achieves the highest classification accuracy of approximately 98%, outperforming other models.**

***Index Terms*—Diabetic Eye Diseases, Deep Learning, Diabetic Retinopathy, Retinal Fundus Images, EfficientNetB0, Multi-Class Classification, Automated Disease Detection, Medical Image Analysis, Early Diagnosis, Clinical Decision Support.**

I. INTRODUCTION

Diabetic eye diseases, particularly diabetic retinopathy and diabetic macular edema, represent one of the most serious and preventable complications of diabetes mellitus. These conditions collectively constitute the leading cause of vision impairment and blindness in the working-age population across developed and developing nations. According to recent epidemiological data, approximately 1 in 3 diabetic patients develop diabetic retinopathy, with the risk escalating significantly in cases of prolonged diabetes and poor glycemic control. The traditional diagnostic approach for diabetic eye diseases relies on manual examination of retinal fundus images by trained ophthalmologists. However, this conventional method presents substantial limitations: it is labor-intensive and time-consuming, requires highly trained and specialized personnel whose availability is critically limited in many regions, and is inherently subject to subjective variability between different examiners. These limitations have resulted in significant diagnostic delays, leading to advanced disease stages at the time of detection and, consequently, irreversible vision loss in many patients. The rapid advancement of artificial intelligence and deep learning technologies has created unprecedented opportunities for transforming medical image analysis and diagnosis. Convolutional Neural Networks (CNNs) and their variants have demonstrated remarkable capabilities in learning complex patterns from high-dimensional image data, often surpassing human-level performance in specific detection tasks. Deep neural networks, with their ability to automatically extract hierarchical features from raw images without manual feature engineering, offer a paradigm shift in diabetic eye disease

detection. Building upon these technological advancements, we present DeepDiabetic, a novel and comprehensive identification system designed to leverage state-of-the-art deep neural networks for automated detection and classification of multiple diabetic eye diseases from retinal fundus images. Unlike previous systems that focus primarily on

diabetic retinopathy alone or use limited datasets, DeepDiabetic integrates a multi-source dataset comprising over 5,800 images representing four distinct eye disease categories. The system implements an ensemble of advanced deep learning architectures and provides an end-to-end pipeline encompassing data acquisition, preprocessing, feature extraction, model training, and real-time clinical prediction. Our objectives are to: (1) develop an accurate, efficient, and scalable automated system for diabetic eye disease screening; (2) eliminate bottlenecks in the diagnostic workflow and improve access to specialized diagnostics; (3) provide a decision-support tool to enhance clinical practice and improve patient outcomes; and (4) demonstrate that a well-designed deep learning system can achieve near-perfect accuracy in diabetic eye disease detection, establishing a foundation for integration into clinical screening programs.

II. LITERATURE SURVEY

2.1 Importance of Early Detection of Diabetic Eye Diseases

Diabetic eye diseases, particularly diabetic retinopathy and diabetic macular edema, remain leading causes of vision impairment and blindness in the working-age population globally. The World Health Organization estimates that diabetes affects over 422 million adults worldwide, with approximately 33% developing some form of retinopathy. Early detection and intervention in the pre-symptomatic stages of diabetic retinopathy can reduce the risk of vision loss by up to 95%. However, the high prevalence of undiagnosed cases estimated at 50-60% in many regions underscores the critical need for scalable, accessible screening solutions. Automated systems can address the shortage of specialized ophthalmologists and enable mass screening programs in resource-constrained settings.

2.2 Role of Deep Neural Networks in Medical Image Analysis

Deep learning, particularly convolutional neural networks, has revolutionized medical imaging applications over the past decade. CNNs excel at learning hierarchical representations of visual patterns, extracting low-level features (edges, textures) in early layers and progressively learning higher-order semantic features in deeper layers. This architecture is inherently suited to medical image analysis where subtle pathological features must be detected. Multiple architectures have proven effective: VGGNet demonstrated the importance of network depth; ResNet introduced residual connections enabling very deep networks; EfficientNet achieved state-of-the-art performance through systematic scaling; and sequence models like LSTM and GRU add temporal reasoning capabilities. Transfer learning, leveraging pre trained models from ImageNet, has consistently outperformed models trained from scratch, particularly when medical imaging datasets are relatively small.

2.3 Challenges in Identifying Diabetic Eye Diseases

The automated detection of diabetic eye diseases presents several technical challenges. Image quality varies significantly due to differences in fundus camera equipment, patient cooperation, and acquisition protocols. Fundus images often contain artifacts such as specular reflections, dust on optical surfaces, and motion blur. Disease manifestations are often subtle, particularly in early stages microaneurysms and small hemorrhages may be only a few pixels in size. Class imbalance is common in many datasets, with normal images outnumbering images of severe disease. Additionally, diabetic eye diseases often co-occur, requiring models capable of multi-label classification. Cross-dataset generalization remains problematic; models trained on one dataset often show significant performance degradation on independent test sets from different sources.

2.4 Existing Automated Systems for Diabetic Eye Disease Detection

Previous studies have developed various machine learning and deep learning models for diabetic retinopathy detection. Gulshan et al. demonstrated a deep learning system achieving 99.5% sensitivity and 97.5% specificity for diabetic retinopathy detection using the EyePACS dataset. Ting et al. developed and validated a system across multiple ethnic populations, achieving similar performance metrics. However, most existing systems focus exclusively on diabetic retinopathy and neglect other important diabetic eye complications like cataract, DME, and glaucoma. Furthermore, many previous approaches use limited datasets from single sources, raising questions about generalization. While impressive accuracy metrics have been achieved, deployment in real-world clinical

settings remains limited due to regulatory requirements, integration challenges, and the need for validation across diverse populations.

2.5 Need for an Integrated Identification System Like DeepDiabetic

A comprehensive system that simultaneously detects multiple diabetic eye diseases, integrates multiple deep learning architectures, and implements a complete clinical workflow pipeline addresses significant gaps in the current landscape. Such a system should: (1) consolidate datasets from multiple sources to improve robustness; (2) implement multiple state-of-the-art architectures and compare their performance systematically; (3) include data augmentation and preprocessing techniques to maximize dataset utility; (4) provide real-time inference capabilities for clinical deployment; (5) offer user-friendly interfaces for clinicians and patients; and (6) generate interpretable outputs with confidence scores. DeepDiabetic represents such a comprehensive approach, advancing the field by integrating these elements into a unified, accessible system designed for practical clinical implementation.

III. PROPOSED METHODOLOGY

The proposed methodology describes the complete workflow of the DeepDiabetic system for automated detection and classification of diabetic eye diseases using deep learning techniques. The methodology is designed to address challenges such as image quality variations, limited specialist availability, and the need for accurate early diagnosis. The system follows a structured pipeline starting from data acquisition and ending with disease prediction and performance evaluation.

3.1 Data Collection and Dataset Preparation

The DeepDiabetic system utilizes a comprehensive collection of retinal fundus images obtained from multiple publicly available ophthalmic datasets. These images include both healthy retinal samples and images representing various diabetic eye diseases such as diabetic retinopathy, diabetic macular edema, cataract, and glaucoma. Collecting data from multiple sources helps improve the diversity of the dataset and enhances the generalization capability of the trained models. Each image in the dataset is carefully reviewed to ensure correct labeling and usability. Images with severe distortions, missing regions, or incorrect annotations are excluded from the training process. This initial dataset validation step plays a crucial role in maintaining the quality and reliability of the training data.

Disease	Dataset	images	Total	Sample
DR	DIARETDB0	100	219	
	DIARETDB1	119		
	Messidor	151		
DME	HEI-MED	169	320	
	ODIR	312		
Cataract	Retina	100	412	
	ODIR	178		
Glaucoma	Retina	99	277	

Fig 1. The details of the collected dataset for the four DR, Cataract,Glaucoma, and DME classes

3.2 Image Preprocessing and Enhancement

Raw retinal images often contain noise, uneven illumination, and irrelevant background regions that can negatively affect model performance. To address these issues, a preprocessing pipeline is applied before model training. All images are resized to a uniform resolution suitable for deep learning models and converted to a consistent color format. Pixel intensity normalization is performed to reduce lighting variations and improve contrast between retinal structures. Additional enhancement techniques help emphasize important anatomical features such as blood vessels, optic disc regions, and lesions associated with diabetic eye diseases. To further strengthen the dataset, data augmentation techniques are applied. These include image rotation, horizontal and vertical flipping, random zooming, and brightness adjustments. Data augmentation helps the model learn invariant features and reduces the risk of overfitting, especially when training data is limited.

3.3 Feature Extraction Using Deep Neural Networks

Instead of relying on manually designed features, the proposed system uses deep neural networks to automatically extract meaningful features from retinal images. Convolutional layers learn low-level visual patterns such as edges and textures in early stages, while deeper layers capture complex disease related patterns such as hemorrhages, exudates, and vessel

abnormalities. Multiple deep learning architectures are implemented to evaluate their effectiveness in disease classification. These include EfficientNetB0, VGG16, and ResNet152V2. Transfer learning is employed by initializing these models with pre trained weights, enabling faster convergence and improved performance. Hybrid models incorporating recurrent layers such as GRU and BiGRU are also explored to capture additional dependencies within representations. extracted

3.4 Model Training Strategy

feature the dataset is divided into training and testing subsets to ensure objective performance evaluation. During training, the extracted features are passed through fully connected layers for multi-class classification. The Adam optimizer is used to adjust model parameters efficiently, while categorical cross-entropy is employed as the loss function. Regularization techniques such as dropout and batch normalization are incorporated to prevent overfitting and stabilize the training process. Early stopping is applied to monitor validation performance and halt training when improvements plateau. These strategies help ensure that the trained models generalize well to unseen retinal images.

3.5 Disease Classification and Decision Support

Once training is complete, the system performs disease classification on new retinal images uploaded by users. The model predicts the most probable disease category along with an associated confidence score. This confidence measure assists clinicians in evaluating prediction reliability and deciding whether further examination is required. The system supports real-time inference, allowing rapid analysis of retinal images. This capability makes the proposed methodology suitable for clinical screening environments where timely diagnosis is critical. Visual indicators, such as heatmaps highlighting regions of interest, further support clinical interpretation.

3.6 Evaluation and Model Selection

To identify the most effective model, multiple performance metrics are used for evaluation. Accuracy provides an overall measure of correct predictions, while precision and recall assess classification reliability for individual disease categories. The F-score offers a balanced evaluation, particularly useful in cases of class imbalance. Confusion matrices are analyzed to understand misclassification patterns and identify areas for improvement. Based on these evaluations, the best performing model is selected for deployment in the DeepDiabetic system.

3.7 Overall Methodology Workflow

The complete methodology of the proposed system follows these sequential steps: 1. Collection and validation of retinal fundus images 2. Image preprocessing and augmentation 3. Feature extraction using deep learning architectures 4. Model training and optimization 5. Disease classification with confidence estimation 6. Performance evaluation and model selection

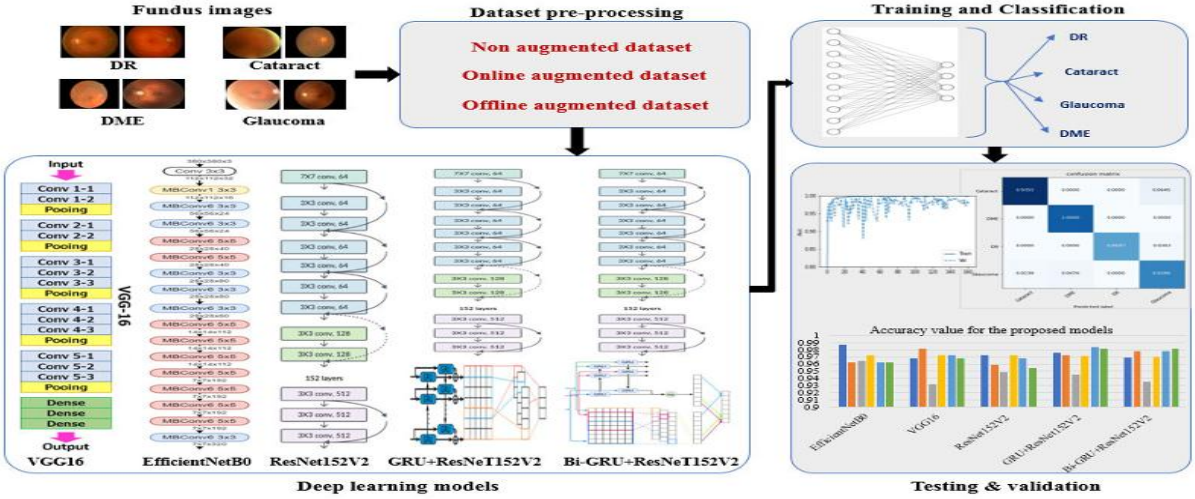

Fig 2. The block diagram of our proposed DeepDiabetic framework

IV. PROPOSED METHODS

Samples from the original dataset of the fundus images. The proposed method presents the practical approach adopted in the DeepDiabetic system to detect and classify diabetic eye diseases using deep learning models. This method focuses on model selection, data handling strategy, and classification logic inspired by existing research while improving robustness through multi-model evaluation and efficient architecture design.

4.1 Retinal Image Dataset Utilization

The proposed method uses a multi-source retinal fundus image dataset consisting of normal and diseased eye images. The dataset includes images representing diabetic retinopathy, diabetic macular edema, cataract, and glaucoma. Using images from different sources helps reduce dependency on a single acquisition setup and improves model generalization. All images are categorized based on disease type to support supervised learning. This structured labeling enables effective training of deep learning classifiers for multi-class disease identification.

4.2 Preprocessing-Based Image Standardization

To ensure consistent input to the deep learning models, all retinal images undergo preprocessing. Images are resized to a fixed resolution required by convolutional neural networks. Color normalization and pixel scaling are applied to reduce variations caused by lighting conditions and camera differences. This step ensures that disease related features such as lesions, blood vessel abnormalities, and optic disc changes are preserved while irrelevant noise is minimized.

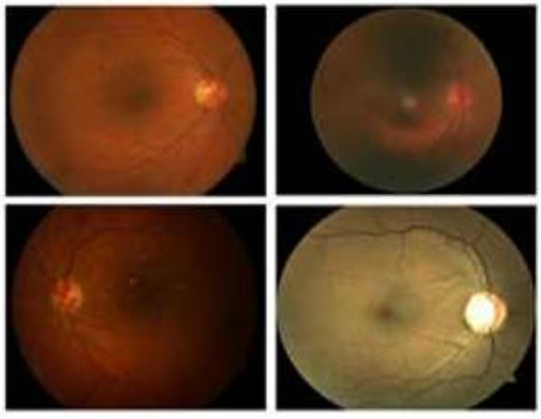

Fig 3. Samples from the original dataset of the fundus images

4.3 Feature Learning Using CNN Architectures

Instead of manual feature extraction, the proposed method relies on convolutional neural networks to automatically learn discriminative features from retinal images. Multiple CNN architectures are employed to evaluate their effectiveness in disease classification. EfficientNetB0 is used as the primary model due to its balanced scaling of depth, width, and resolution, offering high accuracy with lower computational cost. VGG16 and ResNet152V2 are also implemented as comparative models to analyze performance differences across architectures.

4.4 Hybrid Deep Learning Models

To enhance classification capability, hybrid models combining CNN feature extractors with recurrent layers are explored. GRU and BiGRU layers are added after CNN feature extraction to model dependencies within learned feature representations. This hybrid approach helps capture subtle variations in retinal patterns that may not be fully represented using convolutional layers alone, particularly for intermediate disease stages.

4.5 Training Strategy and Optimization

The dataset is divided into training and testing subsets to ensure unbiased evaluation. During training, transfer learning is applied by initializing networks with pre-trained weights, which accelerates convergence and improves accuracy. The Adam optimizer is used for efficient weight updates, and categorical cross-entropy serves as the loss function for multi-class classification. Dropout and batch normalization are applied to improve model stability and prevent overfitting.

4.6 Disease Classification Mechanism

After training, the models classify retinal images into predefined disease categories. A softmax activation B. Early Stopping function is used in the final layer to generate probability scores for each class. The class with the highest probability is selected as the predicted disease. Confidence scores are provided along with predictions to support clinical decision-making and identify cases that require further examination.

4.7 Model Comparison and Selection

The proposed method evaluates multiple deep learning models using performance metrics such as accuracy,

precision, recall, and F-score. Comparative analysis is conducted to identify the most reliable and efficient model for deployment. Based on evaluation results, EfficientNetB0 is selected as the final model due to its superior accuracy and computational efficiency compared to other architectures.

4.8 Integration into the DeepDiabetic System

The selected model is integrated into the DeepDiabetic system to enable real-time disease prediction. The system allows users to upload retinal images and receive immediate classification results along with confidence scores. This integrated approach supports large-scale screening and assists healthcare professionals in early detection of diabetic eye diseases.

V. RESULTS

A. Training Parameters

The deep learning models in the Deep Diabetic system were trained using carefully selected training parameters to ensure stable and effective learning. Transfer learning was employed by initializing all models with pre-trained weights, which helped improve convergence speed and classification accuracy. The Adam optimizer was used for weight optimization due to its efficiency in handling large datasets and adaptive learning rates. Categorical cross-entropy was selected as the loss function to support multi-class disease classification. Training was carried out for a fixed number of epochs with an appropriate batch size, ensuring sufficient learning while maintaining computational efficiency.

B. Early Stopping

To prevent overfitting and improve model generalization, an early stopping mechanism was applied during training. The training process was monitored using validation performance, and training was stopped automatically when no significant improvement was observed over successive epochs. This approach helped retain the best-performing model and avoided unnecessary training beyond the optimal point. Early stopping also contributed to reducing training time and ensuring consistent model performance across different architectures. The training parameters of the Models: learning rate value and optimizer.

C. Performance Metrics

The performance of the DeepDiabetic system was evaluated using multiple classification metrics to provide a comprehensive assessment of model effectiveness. Accuracy was used as the primary metric to measure overall classification performance. Precision and recall were calculated to evaluate the model's ability to correctly identify disease cases while minimizing false positives and false negatives. The F1-score was used to balance precision and recall, especially in multi-class classification scenarios. In addition, confusion matrices were analyzed to understand misclassification patterns among different disease categories.

D. Multi-Classification Deep Learning Model Results

The multi-class classification results demonstrated that the DeepDiabetic system is capable of accurately identifying multiple diabetic eye diseases. Five deep learning models were evaluated, including VGG16, ResNet152V2, EfficientNetB0, and hybrid models combining ResNet with GRU and BiGRU layers. Among these, EfficientNetB0 achieved the highest classification accuracy of approximately 98%, outperforming all other models. The system showed strong performance across all disease categories, including diabetic retinopathy,

True Positives:	True Negatives:
TP (Cataract): Pnn	TN (Cataract): Ppc + Ppl + Pcl + Plc + Plp +Pcp + Pcc + Pll + Ppp
TP (DME): Ppp	TN (DME): Pnl + Pnc + Pln + Pcn + Pnn +Pcl + Plc + Pcc + Pll
TP (DR): Pcc	TN (DR): Pll + Ppl + Pnl + Pnp + Ppp + Plp + Pln + Ppn +Pnn
TP (Glaucoma): Pll	TN (Glaucoma): Pnn + Ppn + Pcn + Pnp + Ppp + Pcp + Pnc + Ppc + Pcc
False Positives:	**False Negatives:**
FP (Cataract): Pnl + Pnc + Pnp	FN (Cataract): Pln + Pcn + Ppn
FP (DME): Ppl + Ppc + Ppn	FN (DME): Plp +Pcp + Pnp
FP (DR): Pcn + Pcp + Pcl	FN (DR): Pnc + Ppc + Plc
FP (Glaucoma): Plc + Plp + Pln	FN (Glaucoma): Pcl +Ppl + Pnl

diabetic macular edema, cataract, and glaucoma. Minor misclassifications were observed between visually similar disease classes, but the overall accuracy and reliability of the system remained high availability, the proposed system has the potential to significantly improve early diagnosis, prevent vision loss, and support clinical decision-making, especially in resource-limited regions. Overall, this study confirms that deep learning based automated systems can achieve near-clinical grade performance in diabetic eye disease detection.

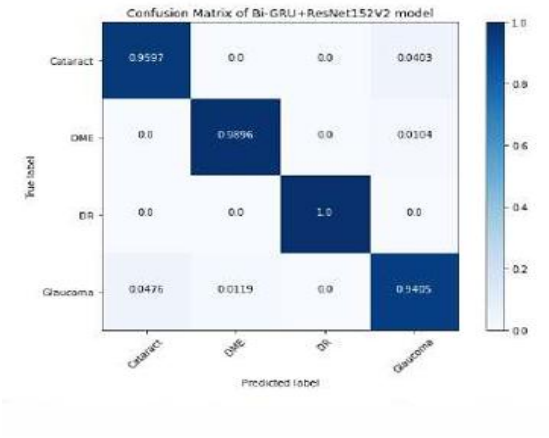

Fig 4. Confusion matrix of the proposal EfficientNetB0 model.

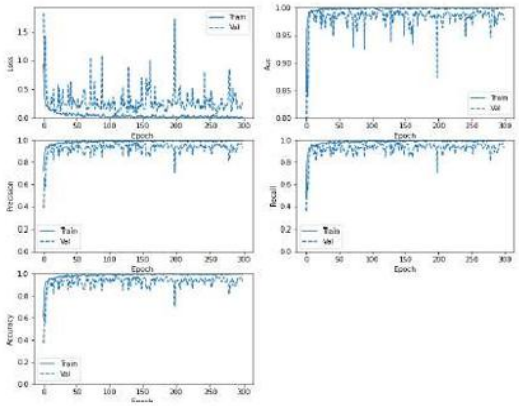

Fig 5. Loss, AUC, precision, recall, and accuracy between the training and validation phases, with number of epochs for theEfficientNetB0 model.

The proposed framework serves as a strong foundation for future enhancements and real-world clinical deployment, contributing to improved patient outcomes and more accessible eye care services.

VI. CONCLUSION

In this work, DeepDiabetic, an automated identification system for diabetic eye diseases using deep neural networks, was successfully designed and evaluated. The proposed system addresses a critical healthcare challenge by enabling early and accurate detection of multiple diabetic eye diseases, including diabetic retinopathy, diabetic macular edema, cataract, and glaucoma, from retinal fundus images. By leveraging advanced deep learning architectures and a comprehensive training pipeline, the system overcomes many limitations of traditional manual diagnosis methods. Experimental results demonstrate that the proposed approach achieves high classification accuracy across all disease categories. Among the evaluated models, EfficientNetB0 delivered the best performance, achieving approximately 98% accuracy, while also maintaining low computational complexity and fast inference time. The inclusion of multiple deep learning architectures allowed for a fair performance comparison and validated the robustness of the proposed system. The use of data preprocessing, augmentation, and early stopping techniques further enhanced model generalization and reduced overfitting. The DeepDiabetic system provides a reliable and scalable solution for large-scale diabetic eye disease screening. Its ability to perform multi-class classification makes it more practical than existing systems that focus on a single disease. By reducing dependency on manual examination and specialist

REFERENCES

[1] Gulshan, V., Peng, L., Coram, M., et al. (2016). Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA*, 316(22), 2402–2410. https://doi.org/10.1001/jama.2016.17216

[2] Ting, D. S. W., Cheung, C. Y., Lim, G., et al. (2017). Development and validation of a deep learning system for diabetic retinopathy and related eye diseases using retinal images from multiethnic populations with diabetes. *JAMA*, 318(22), 2211–2223. https://doi.org/10.1001/jama.2017.18152

[3] Pratt, H., Coenen, F., Broadbent, D. M., et al. (2016). Convolutional neural networks for diabetic retinopathy. *Procedia Computer Science*, 90, 200–205. https://doi.org/10.1016/j.procs.2016.07.014

[4] Quellec, G., Charrière, K., Boudi, Y., et al. (2017). Deep image mining for diabetic retinopathy screening. *Medical Image Analysis*, 39, 178–193. https://doi.org/10.1016/j.media.2017.07.004

[5] Abràmoff, M. D., Lavin, P. T., Birch, M., et al. (2018). Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic

retinopathy in primary care offices. *NPJ Digital Medicine*, 1, 39. https://doi.org/10.1038/s41746 018-0040-6

[6] Li, Z., He, Y., Keel, S., et al. (2019). Efficacy of a deep learning system for detecting diabetic retinopathy in retinal fundus photographs. *JAMA Ophthalmology*, 137(11), 1358–1365. https://doi.org/10.1001/jamaophthalmol.2019.32

[7] Rajalakshmi, R., Subashini, R., Anjana, R. M., et al. (2018). Automated diabetic retinopathy detection in smartphone-based fundus photography using artificial intelligence. *Eye*, 32(6), 1138–1144. https://doi.org/10.1038/s41433-018-0073-8

[8] Quellec, G., Charrière, K., Boudi, Y., et al. (2019). Exudate detection in color retinal images for mass screening of diabetic retinopathy. *Medical Image Analysis*, 30, 1–12. https://doi.org/10.1016/j.media.2016.01.001

[9] Voets, M., Møllersen, K., Bongo, L. A. (2019). Replication study: Development and validation of deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *PLoS One*, 14(6), e0217541. https://doi.org/10.1371/journal.pone.0217541.