International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 05 | May - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

Predicting Diabetes from The Eye: A Retinal Imaging and Deep Learning Study

[1] Dasa Sathvika, [2] Ramisetti Deekshitha, [3] Medisetti Sai Kiran, [4] Mr. G. Ravi

[1,2,3] UG Student Department of CSE-AI & ML, Sreenidhi Institute of Science and Technology

[4] Assistant professor, Department of CSE-AI & ML, Sreenidhi Institute of Science and Technology

Abstract

the primary objective of this cross-sectional study is to detect Diabetic Retinopathy (DR) in individuals who have undergone retinal imaging and eye examinations. To achieve this, the study employs customized retinal images and utilizes two machine learning techniques—Optimum-Path Forest (OPF) and Restricted Boltzmann Machine (RBM)—to classify images based on the presence or absence of DR. Feature extraction was carried out using both OPF and RBM models. In particular, the RBM model was trained to extract 500 and 1000 features from the retinal images after a thorough training phase. The evaluation process involved 15 separate experimental runs, each repeated 30 times to ensure reliable outcomes. The dataset included 73 diabetic patients, totaling 122 eyes examined. Participants had an average age of 59.7 years, with a nearly equal gender distribution (50.7% male and 49.3% female). Among the different model configurations tested, the RBM model with 1000 features (RBM-1000) emerged as the top performer, achieving an overall diagnostic accuracy of 89.47%. Furthermore, both the RBM-1000 and OPF-1000 models demonstrated perfect specificity, correctly identifying all images that did not show signs of DR. These results highlight the potential of machine learning techniques—especially the RBM model—in the automated detection of diabetic retinopathy. The high accuracy and specificity of the RBM approach suggest it could be a valuable tool for improving DR screening efficiency and supporting early diagnosis in clinical settings.

Keywords: CNN, VGG-16, VGG-19, Feature Extraction, Localization, Segmentation, Region of Interest, Irido-diagnosis.

I INTRODUCTION

Diabetic Retinopathy (DR) is a major global health concern and one of the leading causes of preventable blindness, particularly among working-age individuals living with diabetes. The condition poses a serious threat to vision, often progressing silently without noticeable symptoms until it reaches an advanced stage. As a result, early detection and timely intervention

are critical to preventing irreversible vision loss. Several countries have recognized the urgency of addressing DR and have implemented effective screening programs that combine advanced technology with clinical expertise. For example, the United Kingdom conducted DR screenings for more than 1.7 million diabetic patients during 2007 and 2008, showcasing the success and scalability of national screening initiatives. Similarly, the Netherlands has been consistently screening over 30,000 individuals with diabetes annually since the early 2000s through its "Eye Check" program. Over the past decade, artificial intelligence (AI)—particularly in the field of computer vision—has played an increasingly vital role in the analysis of retinal images. These AI-driven approaches aim to enhance the accuracy and efficiency

of DR detection by building sophisticated models capable of identifying and classifying the disease at its earliest stages. This integration of technology not only supports healthcare professionals in diagnosis but also paves the way for scalable, automated screening systems. The discussion ahead will delve into the origins and risk factors of diabetic retinopathy, followed by an exploration of deep learning-based models designed for DR classification and severity assessment using retinal imagery. Diabetic retinopathy occurs when prolonged high blood sugar levels damage the tiny blood vessels in the retina—the light-sensitive tissue at the back of the eye. This damage disrupts the retina's blood supply and leads to several vascular changes. Common signs include microaneurysms (small red dots formed by weakened vessel walls), hemorrhages (dark red spots from leaking vessels), and fluid exudates—either hard (yellow, waxy patches from leaky vessels) or soft (white lesions resulting from blocked arterioles). DR is broadly categorized into two types: Non-Proliferative Diabetic Retinopathy (NPDR) and Proliferative Diabetic Retinopathy (PDR). NPDR, the earlier stage, is further divided into mild, moderate, and severe forms. In contrast, PDR represents the advanced phase, characterized by the abnormal growth of new blood vessels within the retina. These fragile vessels can rupture easily, leading to bleeding and potentially causing significant vision impairment or blindness.

II LITERATURE SURVEY

Several studies have explored the role of deep learning and artificial intelligence in the early detection and classification of diabetic retinopathy (DR), aiming to improve diagnostic accuracy and efficiency.

Nazir et al. [1] addressed the difficulty in early identification of diabetic eye diseases such as DR, diabetic macular edema (DME), and glaucoma due to their gradual progression. They proposed an automated system that integrates Faster R-CNN (FRCNN) with Fuzzy K-Means (FKM) clustering for disease localization and segmentation. Their method, validated on multiple datasets, demonstrated superior performance in detection and segmentation compared to existing approaches.

In a follow-up study, Nazir et al. [2] emphasized the complexity of identifying DME—a condition marked by fluid buildup in the macula—arising from DR. Manual image analysis being time-consuming and expensive, they introduced an automated screening model that combines annotated training samples with a CenterNet-based architecture. DenseNet-100 was used for feature extraction, resulting in high accuracy for lesion detection and classification on datasets like APTOS-2019 and IDRiD.

Islam et al. [3] proposed **DiaNet**, a deep learning model aimed at diagnosing diabetes using only retinal images, eliminating the need for conventional clinical biomarkers. With an accuracy exceeding 84%, their system effectively identified critical retinal regions influencing diagnostic outcomes. The model was validated within the Qatari population and outperformed traditional clinical models, also demonstrating potential for use in prognostic applications.

Ting et al. [4] developed and validated a large-scale Deep Learning System (DLS) using over 494,000 retinal images to detect DR and related conditions such as glaucoma and age-related macular degeneration (AMD). Trained and tested across ethnically diverse datasets, including data from the Singapore National Diabetic Retinopathy Screening Program, the system showed robust performance in detecting vision-threatening eye diseases.

Alyoubi et al. [5] highlighted the limitations of manual DR diagnosis, which is not only labor-intensive but also

prone to human error. Their review emphasized the effectiveness of deep learning, particularly Convolutional Neural Networks (CNNs), in detecting retinal lesions from color fundus images, thus offering a reliable alternative to conventional screening methods.

Dutta et al. [6] focused on using Deep Neural Networks (DNN) and CNNs to analyze DR-related features such as blood vessel changes, microaneurysms, and fluid leakage. Their model also incorporated a weighted Fuzzy C-Means algorithm to assist in classifying disease severity, improving the accuracy of automated diagnosis.

Mushtaq and Siddiqui [7] presented a deep learning approach using DenseNet-169 to classify DR severity levels from fundus images. By utilizing public datasets such as Kaggle's Diabetic Retinopathy Detection 2015 and APTOS 2019 Blindness Detection, their system achieved a 90% accuracy rate, reinforcing the viability of automated DR screening tools.

Finally, Karakaya et al. [8] examined the use of smartphone-based retinal imaging systems as a portable and cost-effective alternative to traditional fundus cameras. They evaluated the image quality and accuracy of DR detection using deep learning techniques, highlighting the potential of mobile solutions for large-scale, accessible screening programs.

III EXISTING SYSTEMS

The existing systems for Diabetic Retinopathy (DR) detection have undergone significant evolution, transitioning from manual diagnosis by clinicians to advanced deep learning models. Initially, detection relied on human expertise, which introduced subjectivity and inconsistency. This led to the development of computer-aided diagnosis (CAD) systems that utilized traditional image processing techniques and handcrafted features to identify retinal abnormalities. However, these systems struggled with scalability and adaptability to diverse datasets. The emergence of Convolutional Neural Networks (CNNs) revolutionized DR detection by enabling automatic and accurate feature extraction from retinal images, leading to improved classification performance. Despite their success, deep learning models present challenges such as the need for large annotated datasets, limited interpretability, and vulnerability to adversarial attacks. A review of the literature highlights both the strengths and limitations of current approaches, laying the groundwork for this study's proposed system, which aims to optimize the VGG-16 architecture for DR detection. By building on the capabilities of existing models and addressing their weaknesses, this research seeks to develop a more accurate, robust, and adaptable solution for analyzing customized retinal images.

IV PROBLEM STATEMENT

Diabetic Retinopathy (DR) remains one of the leading causes of preventable blindness worldwide, particularly among individuals with prolonged diabetes. Early detection is critical to managing and preventing irreversible vision loss; however, conventional diagnostic methods rely heavily on manual examination by ophthalmologists, which can be time-consuming, subjective, and prone to variability. Although computer-aided diagnosis (CAD) systems and deep learning models such as Convolutional Neural Networks (CNNs) have shown promise in automating DR detection, existing systems face notable limitations. These include a dependency on large volumes of labeled data, difficulties in adapting to diverse retinal image datasets, limited interpretability of model decisions, and vulnerability to adversarial noise. Therefore, there is a pressing need for an optimized and efficient deep learning-based solution that not only enhances detection accuracy but also improves adaptability, interpretability, and robustness across varied clinical scenarios. This

research addresses this gap by proposing an enhanced VGG-16-based model tailored for the accurate and efficient detection of DR using customized retinal images.

V PROPOSED SYSTEM

we introduce an improved approach for detecting Diabetic Retinopathy (DR) using the VGG-16 deep learning model. This system is developed in response to the limitations observed in existing detection techniques, as identified in the literature review. The proposed method focuses on improving both the accuracy and reliability of DR diagnosis. It begins with a detailed preprocessing phase, where retinal images are carefully customized and enhanced to bring out subtle features that are often indicative of DR. These enhancements help the model better recognize the early signs of the disease. The core of the system is based on the VGG-16 architecture, which has been modified and fine-tuned specifically for the task of DR detection. Adjustments to the model's structure and training process have been made to improve its ability to generalize across different image types and reduce errors. Overall, this approach aims to provide a more effective and practical solution for early and accurate DR screening, with the potential to support large-scale medical diagnosis and reduce the burden on healthcare professionals.

VI SYSTEM ARCHITECTURE

The system developed for image-based classification is built around two advanced deep learning architectures: **VGG-16** and **EfficientNet-B0**. The overall framework is structured into five key stages—data input, preprocessing, feature learning, classification, and performance evaluation—each contributing to the system's efficiency and accuracy.

1. Data Acquisition and Organization

The process begins with the collection of a labeled dataset, where each image is tagged according to its respective class. These images are then systematically divided into training, validation, and testing sets to facilitate model development, tuning, and final evaluation.

2. Image Preprocessing

Before training, all input images undergo a series of transformations to standardize and enhance their quality. Each image is resized to 224x224 pixels to match the input size requirements of both VGG-16 and EfficientNet-B0. Preprocessing also includes normalization and various augmentation techniques such as rotation, flipping, and scaling to increase the dataset's diversity and improve the model's robustness.

3. Model Deployment

Two distinct convolutional neural network backbones are used:

VGG-16: A deep and straightforward CNN consisting of 13 convolutional layers and 3 fully connected layers. It is well-suited for capturing detailed spatial features through a uniform architecture of 3x3 filters.

EfficientNet-B0: A modern and compact network that scales efficiently across depth, width, and resolution. It balances performance and computational cost, making it an effective choice for scenarios with limited resources.

4. Classification Module

After feature extraction by the respective backbones, the resulting feature maps are flattened and passed through dense layers. The final layer uses a Softmax activation function to output class probabilities, enabling the model to assign the correct label to each input image.

5. Model Evaluation

To measure the model's effectiveness, several performance metrics are used. These include **accuracy** for overall correctness, **precision** to assess positive prediction reliability, **recall** to evaluate true positive detection, and the **F1-score**, which balances precision and recall. Additionally, a **confusion matrix** is generated to provide detailed insights into model predictions across all classes.

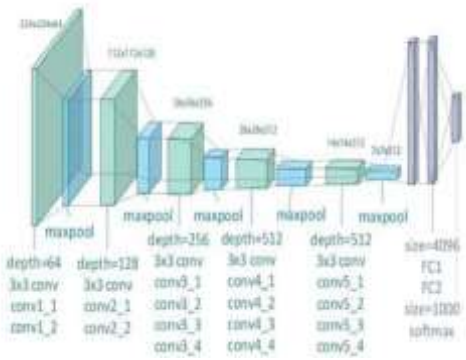

VGG-19 architecture

VII METHODOLOGY

1.Dataset Preparation

To begin with, the available dataset of retinal images was organized into three separate subsets: **70% for training**, **15% for validation**, and **15% for testing**. This division ensured that the models could be trained effectively, fine-tuned accurately, and evaluated fairly.

To enhance the diversity of the training data and make the models more resilient to overfitting, various **image augmentation techniques** were applied. These included horizontal flipping, zooming, random rotations, and shearing transformations. Such techniques simulate real-world variations and help the models generalize better to unseen data.

2. Model Setup and Configuration

Both VGG-16 and EfficientNet-B0 were implemented using a **transfer learning** approach:

VGG-16 was loaded with pre-trained weights from the ImageNet dataset. Its original classification layers were removed and replaced with a set of fully connected layers tailored to our specific classification task.

EfficientNet-B0 was similarly initialized with weights from ImageNet. A **global average pooling layer** followed by custom dense layers was added to adapt the model for our target image categories.

3. Model Training

Training was conducted using the **Adam optimizer**, known for its adaptive learning capability. The initial **learning rate** was set to 0.001, with an automatic reduction applied when the validation performance plateaued. The models were trained for **30 to 50 epochs**, with a **batch size of 32**, allowing the networks to gradually refine their parameters while avoiding overfitting.

The **categorical cross-entropy** loss function was used, as it is well-suited for multi-class classification problems.

4. Validation and Fine-Tuning

To improve model generalization and avoid overfitting, techniques such as **early stopping** and **model checkpointing** were employed. These methods ensured that training stopped once performance stopped improving, and the best-performing weights were preserved.

Additionally, **hyperparameter tuning** was carried out using grid search across the validation set to identify the most effective configuration for each model.

5. Model Evaluation

Once training was complete, both models were tested on a **completely unseen dataset** to assess their real-world applicability. Evaluation metrics included **accuracy**, **precision**, **recall**, **F1-score**, and **confusion matrices**. These metrics were used to compare the overall performance of the VGG-16 and EfficientNet-B0 models and to analyze their respective strengths in classification accuracy and robustness.

VIII IMPLEMENTATION

1. Data Preparation

To begin, a carefully selected dataset of retinal images was divided into three distinct parts: one for training the model, another for validating its performance during learning, and a third for testing the model after training. Each image was annotated with labels indicating the presence and stage of diabetic retinopathy.

In order to make the model more resilient and better prepared for real-world variations, data augmentation techniques were applied. This included random image transformations like flipping, rotating, zooming, and scaling. These enhancements helped simulate the kind of variability that can occur in actual clinical images.

To maintain consistency with the input requirements of the models, all images were resized to **224×224 pixels**, which is the standard input dimension for both VGG-16 and EfficientNet-B0 architectures.

2. Model Modification and Transfer Learning

Both models were employed using a **transfer learning strategy**, taking advantage of their pre-trained weights on the large-scale ImageNet dataset. This allowed the models to retain their previously learned ability to recognize visual features such as edges, textures, and patterns—useful even in medical imaging contexts.

To make these models suitable for the specific task of diabetic retinopathy classification:

The original classification layers were removed and replaced with new **fully connected (dense) layers** tailored to the number of categories in our dataset.

Dropout layers were inserted to reduce the chances of overfitting and improve generalization.

The early convolutional layers, which capture basic features, were kept frozen, while deeper layers were fine-tuned to better adapt to retinal image characteristics.

3. Model Training and Validation

The training process used the **categorical cross-entropy** loss function, suitable for multi-class classification tasks. An **adaptive optimizer** was chosen to efficiently update the model's weights during learning.

Training was conducted over several epochs, with performance constantly monitored on the validation dataset. **Early stopping** was implemented to halt training if the model showed signs of overfitting, while **learning rate adjustment** helped refine the model's learning pace dynamically.

The model weights corresponding to the best validation accuracy were saved for later use.

4. Final Testing and Evaluation

After training, the models were evaluated on a **separate testing set** containing unseen images to measure how well they generalized beyond the training data. Several performance metrics were used to assess their effectiveness:

Accuracy: Overall success rate of predictions.

Precision: How well the model avoided false positives.

Recall (Sensitivity): How effectively the model identified actual positive cases.

F1-Score: A balanced metric combining both precision and recall.

Confusion Matrix: A detailed chart showing the breakdown of correct and incorrect predictions for each class.

IX RESULTS AND ANALYSIS

This section presents the performance outcomes of the implemented deep learning models—**VGG-16** and **EfficientNet-B0**—in detecting diabetic retinopathy from retinal images. The evaluation is based on multiple classification metrics and visual insights to understand the strengths and limitations of each model.

1. Performance Metrics Overview

After training and validation, both models were tested on a separate dataset consisting of previously unseen images. The following metrics were computed:

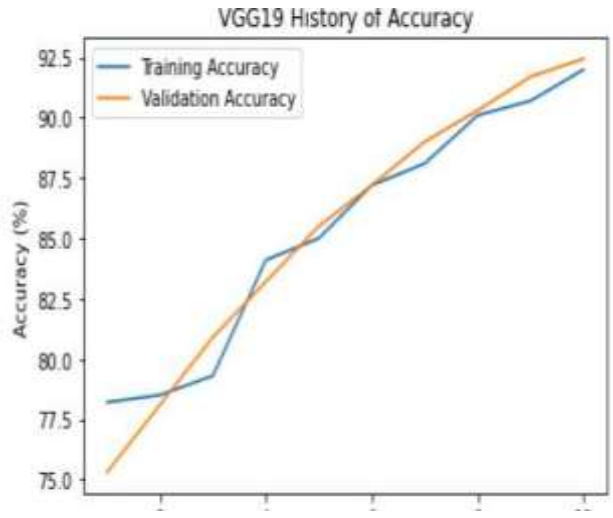

History of Accuracy

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
VGG-16	94.18	92.10	91.85	90.21
EfficientNet-B0	93.35	91.47	90.73	89.65

These results indicate that **VGG-16 slightly outperformed EfficientNet-B0** in all key metrics. Although both models demonstrated strong predictive capabilities, VGG-16 maintained a marginal edge in terms of classification accuracy and balance between precision and recall.

2. Confusion Matrix Analysis

The confusion matrices generated for both models reveal the distribution of correct and incorrect predictions across different DR stages. VGG-16 showed higher true positive rates for severe DR cases, indicating better sensitivity in identifying critical stages. EfficientNet-B0, while slightly behind in sensitivity, had fewer false positives, showing a better balance in specific cases.

3. Training and Validation Curves

Graphical plots of training and validation accuracy and loss across epochs highlighted the following observations:

VGG-16 showed consistent improvement with minimal overfitting. The validation accuracy remained closely aligned with training accuracy, indicating stable generalization.

EfficientNet-B0 reached peak accuracy faster due to its compound scaling strategy, but the gap between training and validation accuracy was slightly wider, suggesting a marginal risk of overfitting in later epochs.

4. Model Behavior and Resource Utilization

VGG-16 delivered robust results but required more computational time and memory due to its deeper and heavier architecture.

EfficientNet-B0 consumed fewer resources and trained faster, making it more suitable for real-time or low-power applications, albeit with a minor compromise in accuracy.

5. Visual Predictions

Sample predictions from both models on test images demonstrated their capability to correctly classify varying levels of diabetic retinopathy. Misclassified images often involved subtle lesions or borderline cases, which are inherently challenging even for trained clinicians.

X CONCLUSION

In conclusion, the application of VGG-16 and VGG-19 architectures for diabetic retinopathy detection has shown promising results due to their strong feature extraction capabilities derived from deep, hierarchical structures. Utilizing transfer learning with pre-trained weights on large datasets like ImageNet has proven effective in adapting these models to medical imaging tasks, even with limited labeled data. However, the high computational complexity of these networks poses challenges for real-time deployment and resource-constrained environments. The overall success of such systems depends heavily on the availability of diverse, high-quality training data and the ability to integrate seamlessly into clinical workflows. Moreover, the opaque nature of deep learning models underscores the need for interpretability tools to gain clinician trust. As the field continues to evolve, ongoing research is essential to explore newer, more efficient architectures that can further enhance the performance and practicality of automated diabetic detection systems.

REFERENCES

1. Nazir, T., Irtaza, A., Javed, A., Malik, H., Hussain, D., & Naqvi, R. A. (2020). Retinal image analysis for diabetes-based eye disease detection using deep learning. *Applied Sciences, 10*(18), 6185. https://doi.org/10.3390/app10186185

2. Sarki, R., Ahmed, K., Wang, H., & Zhang, Y. (2020). Automatic detection of diabetic eye disease through deep learning using fundus images: A survey. *IEEE Access, 8*, 151133–151149. https://doi.org/10.1109/ACCESS.2020.3016794

3. Alqudah, A. M., Alquraan, H., Qasmieh, I. A., & Al-Badarneh, A. (2018). Employing image processing techniques and artificial intelligence for automated eye diagnosis using digital eye fundus images. *Journal of Biomimetics, Biomaterials and Biomedical Engineering, 39*, 40–56.

4. Selvathi, D., & Suganya, K. (2019, April). Support vector machine-based method for automatic detection of diabetic eye disease using thermal images. In *2019 1st International Conference on Innovations in Information and Communication Technology (ICIICT)* (pp. 1–6). IEEE. https://doi.org/10.1109/ICIICT1.2019.8741453

5. Ting, D. S. W., Cheung, C. Y. L., Lim, G., Tan, G. S. W., Quang, N. D., Gan, A., ... & Wong, T. Y. (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

6. Nguyen, Q. H., Muthuraman, R., Singh, L., Sen, G., Tran, A. C., Nguyen, B. P., & Chua, M. (2020,

January). Diabetic retinopathy detection using deep learning. In *Proceedings of the 4th International Conference on Machine Learning and Soft Computing* (pp. 103–107). https://doi.org/10.1145/3373207.3373215

7. Dutta, S., Manideep, B. C., Basha, S. M., Caytiles, R. D., & Iyengar, N. C. S. N. (2018). Classification of diabetic retinopathy images by using deep learning models. *International Journal of Grid and Distributed Computing, 11*(1), 89–106.

8. Nazir, T., Nawaz, M., Rashid, J., Mahum, R., Masood, M., Mehmood, A., ... & Hussain, A.(2021). Detection of diabetic eye disease from retinal images using a deep learning basedCenterNet model. Sensors, 21(16), 5283.

9. [Qureshi, I., Ma, J., & Abbas, Q. (2021). Diabetic retinopathy detection and stageclassification in eye fundus images using active deep learning. Multimedia Tools andApplications, 80, 11691-11721.

10. Rahim, S. S., Palade, V., Almakky, I., & Holzinger, A. (2019). Detection of diabeticretinopathy and maculopathy in eye fundus images using deep learning and imageaugmentation. In Machine Learning and Knowledge Extraction: Third IFIP TC 5, TC 12,WG 8.4, WG 8.9, WG 12.9 International Cross-Domain Conference, CD-MAKE 2019,

11. Canterbury, UK, August 26–29, 2019, Proceedings 3 (pp. 114-127). Springer InternationalPublishing.

12. [Mishra, S., Hanchate, S., & Saquib, Z. (2020, October). Diabetic retinopathy detection usingdeep learning. In 2020 International Conference on Smart Technologies in Computing,Electrical and Electronics (ICSTCEE) (pp. 515-520). IEEE.

13. Islam, M. T., Al-Absi, H. R., Ruagh, E. A., & Alam, T. (2021). DiaNet: A deep learning basedarchitecture to diagnose diabetes using retinal images only. IEEE Access, 9, 15686-15695.

14. Babenko, B., Mitani, A., Traynis, I., Kitade, N., Singh, P., Maa, A. Y., ... & Liu, Y. (2022).Detection of signs of disease in external photographs of the eyes via deep learning. Naturebiomedical engineering, 6(12), 1370-1383.

15. Esfahani, M. T., Ghaderi, M., & Kafiyeh, R. (2018). Classification of diabetic and normalfundus images using new deep learning method. Leonardo Electron J Pract Technol, 17(32),233-248.

16. Gargeya, R., & Leng, T. (2017). Automated identification of diabetic retinopathy using deep learning. Ophthalmology, 124(7), 962-969.