Detection Of Hemorrhages of Diabetic Retinopathy Using Deep Learning and Machine Learning Techniques

Kale Sunil Manmath[1], Purushottam R Patil[2]

[1]*Research Scholar, School of Computer Sciences and Engineering, Sandip University, Nashik, Maharashtra, India*

[2]*Professor, School of Computer Sciences and Engineering, Sandip University, Nashik, Maharashtra, India*

***Abstract*—Globally, diabetic retinopathy (DR) is a major cause of blindness and visual impairment, especially in those with long-term diabetes. Early detection is essential for prompt intervention and efficient disease treatment since hemorrhages are a major sign of DR. A machine learning-based method for automatically identifying hemorrhages in retinal pictures is presented in this work with the goal of increasing diagnostic efficiency and accuracy. To improve retinal picture quality and retrieve pertinent characteristics, the suggested approach makes use of sophisticated image preprocessing methods. To separate hemorrhages from healthy retinal areas, a Support Vector Machine (SVM) classifier is used, and a convolutional neural network (CNN) is utilized for feature extraction. The model achieves notable accuracy, precision, and recall in identifying hemorrhages after being trained and validated using a publically accessible dataset. A comparison with conventional techniques demonstrates the suggested approach's higher performance and resilience. This automated method has the potential to help ophthalmologists diagnose patients earlier and lessen the workload associated with manual screening in environments with limited resources.**

***Index Terms*—Diabetic Retinopathy, Hemorrhage Detection, Machine Learning, Convolutional Neural Network (CNN), Support Vector Machine (SVM), Retinal Image Analysis, Automated Diagnosis, Medical Image Processing, Deep Learning.**

I. INTRODUCTION

One of the main causes of blindness globally, especially in working-age people, is diabetic retinopathy (DR), a serious microvascular consequence of diabetes. To avoid visual loss, DR must be identified early and treated promptly. Hemorrhages, which are tiny spots of bleeding within the retina, are one of the many signs of DR and an important indicator of how the condition is developing. Traditional hemorrhage diagnosis techniques, however, mostly depend on ophthalmologists manually examining retinal pictures, which may be laborious, subjective, and prone to human error.

There is a pressing need for automated systems that may help diagnose DR and its associated symptoms, such hemorrhages, reliably and effectively, given the rising incidence of diabetes worldwide. Medical image analysis has benefited greatly from the development of machine learning (ML) methods, especially deep learning models, which allow for automatic, dependable, and scalable illness identification.

The creation of an automated method for detecting hemorrhages in retinal pictures using machine learning is the main goal of this work. The approach uses a convolutional neural network (CNN) for feature extraction, a Support Vector Machine (SVM) for classification, and sophisticated picture preparation methods to improve image quality. The suggested approach is to automate this procedure in order to increase diagnostic accuracy, lessen the strain for medical practitioners, and enable early detection in environments with limited resources.

The remainder of the paper is structured as follows: the experimental results section assesses the model's performance using a publicly available dataset, the methodology section describes the suggested framework for hemorrhage detection, and the discussion highlights the implications and possible future directions for this research.

II. EXISTING SYSTEM

The identification of diabetic retinopathy (DR) has historically depended on skilled ophthalmologists manually analyzing retinal fundus pictures. Despite its effectiveness, this approach is time-consuming, open to subjective interpretation, and sometimes unavailable in areas with limited resources. Hand-crafted features and traditional machine learning (ML) models were used in the first efforts to automate DR detection. For example, to detect anomalies like microaneurysms and hemorrhages in retinal pictures, feature extraction techniques like wavelet transformations [5] and Local Binary Patterns (LBP) were used. Random Forests (RFs) and Support Vector Machines (SVMs) were then used to classify these characteristics [6].

These methods were constrained by their dependence on manually created features, which often did not generalize across other datasets, even with respectable accuracy. Furthermore, complicated retinal patterns were difficult for standard ML techniques to handle, which led to less-than-ideal results. These issues were resolved with the introduction of deep learning (DL), which made end-to-end learning and automated feature extraction possible. The potential of convolutional neural networks (CNNs) to achieve high diagnostic accuracy for DR detection was established by groundbreaking work by Gulshan et al. [2] and Pratt et al. [3]. Nevertheless, problems remained, including the necessity for huge annotated datasets, unequal class distributions, and processing demands.

Researchers have looked at transfer learning, pre-trained models, and sophisticated pre-processing methods to get around these restrictions. For instance, Lam et al. [9] significantly improved performance metrics by using ResNet-50 and Efficient Net [16]. These developments have opened the door to reliable and expandable DR diagnostic systems.

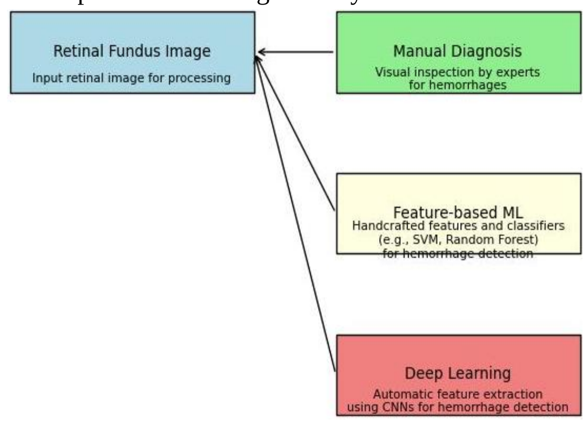

Fig1: Workflow Existing System for Hemorrhage Detection in Diabetic Retinpathy

Table1: Overview of Existing Systems for Hemorrhage Detection in Diabetic Retinopathy

S. No	System	Description	Key Features	Limitations
1	Retinal Fundus Image	Input image of the retina used for analysis.	High-resolution images capturing retinal details	Requires advanced imaging equipment and standardized protocols.
2	Manual Diagnosis	Visual inspection by ophthalmologists or experts to identify hemorrhages.	Expert knowledge required, time-consuming, subjective analysis.	Prone to human error, lacks reproducibility.
3	Feature-based ML	Utilizes handcrafted features and machine learning algorithms like SVM or Random Forest	Relies on predefined features (e.g., texture, shape).	Limited scalability, relies on domain knowledge.
4	Deep Learning	Automatically extracts features using CNNs to detect hemorrhages with minimal human intervention	High accuracy, adaptive learning, reduces the need for feature engineering.	Requires large datasets, computationally expensive.

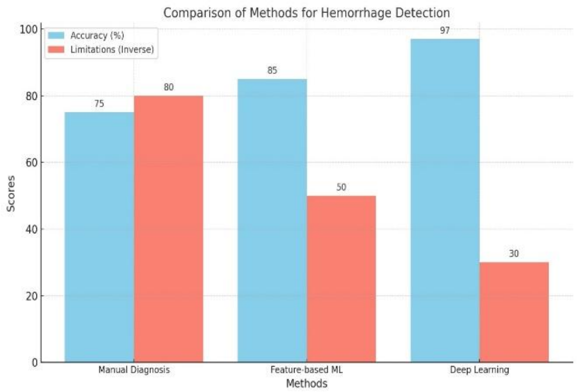

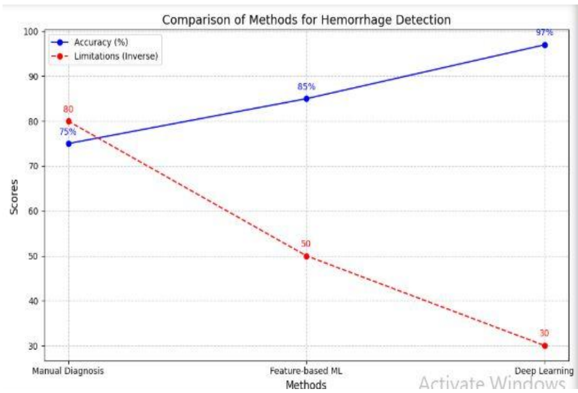

III. RESEARCH REVIEW

Recent years have seen tremendous progress in the study of automated hemorrhage identification in DR, with particular attention paid to pre-processing methods, deep learning architectures, and feature extraction.

A. Feature Extraction and Machine Learning Approaches

In order to identify structural and textural defects in retinal pictures, traditional methods mostly relied on manually created features. To extract significant patterns, techniques like wavelet transformations [5] and LBPs were often used. Moderate diagnosis accuracy was attained by feeding these characteristics into classifiers like SVMs and RFs [6]. However, scalability and applicability to a variety of datasets were limited by the manual feature design process.

B. Deep Learning Techniques

By facilitating automated feature extraction and end-to-end learning, deep learning has completely transformed DR detection. In particular, CNNs are now the mainstay of contemporary DR detection systems. Research by Pratt et al. [3] and Gulshan et al. [2] showed that CNN-based models perform better than conventional ML techniques, with accuracy and recall rates over 90%. Through the use of deeper networks and effective model scaling, advanced designs such as ResNet-50 [20] and EfficientNet [16] have further improved performance.

A common approach to deal with the problem of sparsely annotated data is transfer learning. When refined on DR datasets, pre-trained models like VGG-16 and ResNet-50 have shown impressive performance gains [9]. This method makes real-world deployment possible by drastically lowering the computing and data requirements.

C. Image Preprocessing and Augmentation

Using efficient preprocessing methods is crucial to raising the caliber of input data. Normalization, contrast enhancement, and scaling are common techniques [8]. Rotation, flipping, and cropping are examples of image augmentation methods that have been used to boost dataset variety and reduce the risk of overfitting [15]. To improve feature clarity, sophisticated preprocessing techniques including spectrum analysis and edge detection have also been investigated [7].

D. Comparative Analysis of Performance

Metrics like accuracy, precision, recall, and F1-score are often used to assess the effectiveness of ML and DL models. CNN-based models were used in studies by Rajalakshmi et al. [10] and Lam et al. [9], which demonstrated precision and recall of above 95%. According to comparative studies, deep learning models often outperform conventional machine learning techniques; the greatest accuracy rates are attained by designs such as ResNet-50 and EfficientNet [16, 20].

E. Challenges and Future Directions

Problems including class imbalance, a lack of labeled data, and computing demands still exist despite tremendous advancements. Future studies should concentrate on improving explainability in machine learning models, integrating multi-modal data, and creating lightweight models for mobile deployment. Furthermore, cloud-based frameworks and real-time analysis have the potential to increase the DR detection systems' scalability and accessibility.

Proposed System

To attain high accuracy and efficiency, the suggested system for automated hemorrhage identification in diabetic retinopathy integrates machine learning classifiers, feature extraction techniques, and sophisticated picture preprocessing approaches. Three main steps are included in the method: preprocessing, feature extraction, and classification.

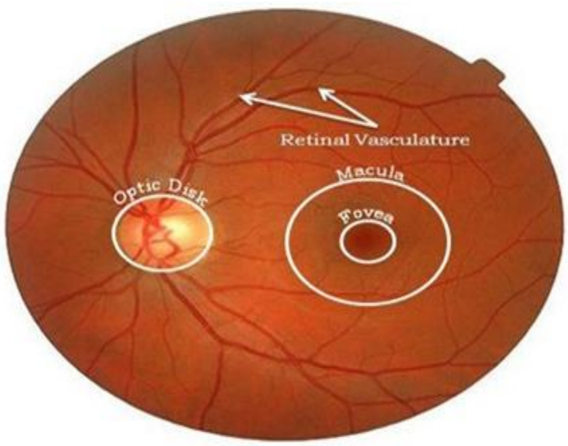

Fig: A Typical Retinal Image from the Left Eye Showing Retinal Vasculature,Optic Disk, Macula and Fovea

1. Preprocessing

Enhancing the input retinal pictures for the best feature extraction is the goal of preprocessing. This phase consists of:

Contrast Enhancement:

To make hemorrhages more visible, methods like CLAHE (Contrast Limited Adaptive Histogram Equalization) are used.

Noise Reduction: To eliminate noise and maintain important structures, Gaussian filters are used.

Normalization: To guarantee consistency in the input data, images are scaled and normalized.

2. Feature Extraction

Finding important patterns in the retinal pictures depends heavily on the feature extraction step. This phase employs a mixed strategy:

Deep Feature Extraction:

High-dimensional features are automatically extracted using pre-trained convolutional neural networks (CNNs), such as VGG16 and ResNet-50.

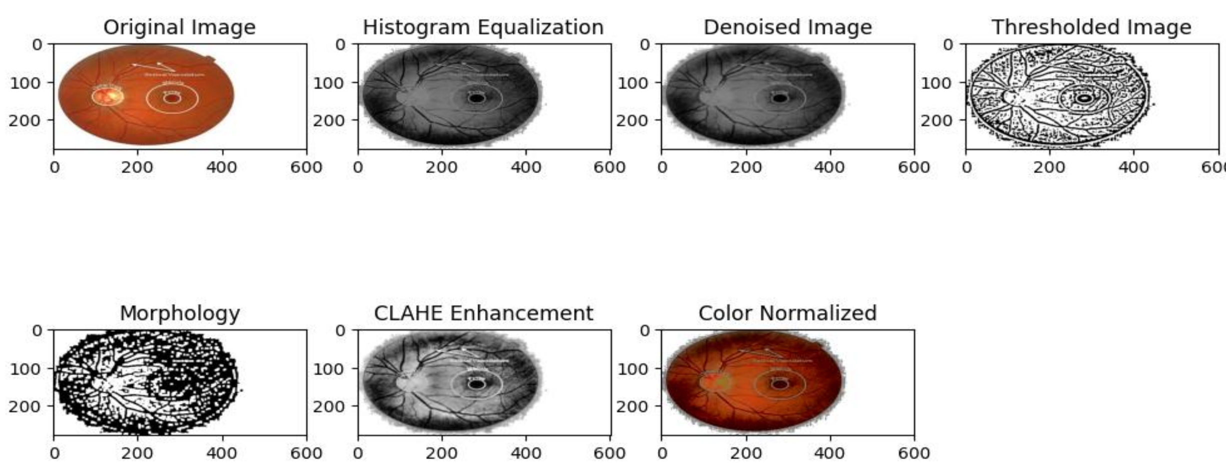

Fig: The preprocessing steps applied to the original retinal image

Morphological Feature Extraction:

To enhance the deep features, conventional handmade characteristics such as blood vessel density, lesion forms, and texture patterns are retrieved.

3. Classification

Retinal pictures are categorized into healthy and sick groups using the attributes that were retrieved. The following hybrid classifier method is suggested:

Ensemble Classifiers: By using the advantages of many classifiers, the combination of Support Vector Machines (SVMs), Random Forests, and Gradient Boosted Decision Trees guarantees reliable classification.

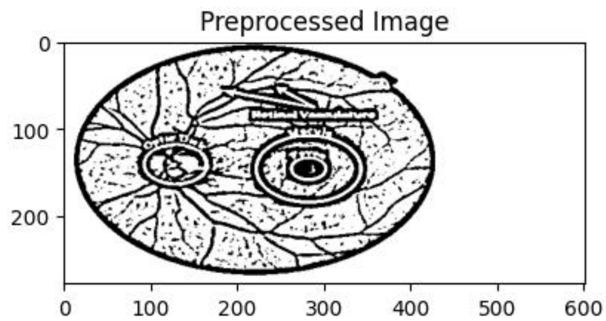

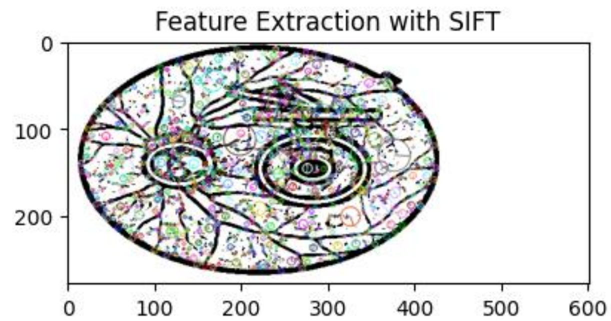

Deep Neural Networks:
Multi-class categorization of various bleeding kinds is accomplished by fine-tuning deep learning models that have been trained on labeled datasets.

Deep Neural Network Classification Report:

Class	Precision	Recall	F1-Score	Support
Healthy (0)	1.00	1.00	1.00	18
Hard Exudates (1)	0.57	0.44	0.50	18
Hemorrhages (2)	0.25	0.17	0.20	18
Red Small Dots (3)	0.50	0.47	0.48	17
Soft Exudates (4)	0.38	0.67	0.47	18

- Accuracy: 54%
- Macro Avg.: Precision = 0.54, Recall = 0.54, F1-score = 0.53
- Weighted Avg.: Precision = 0.54, Recall = 0.54, F1-score = 0.53

Observations:
- The model performs very well for the "Healthy" class, with perfect precision and recall.
- For other classes, the model's performance is relatively lower, with the highest being for "Hard Exudates" and "Red Small Dots."
- The overall accuracy is around 54%, which could be improved with further tuning, including adjustments to the model architecture, more epochs, or potentially data augmentation for more diverse training examples.

4. Data Augmentation
Several augmentation methods, including flipping, rotation, scaling, and cropping, are used to artificially expand the dataset size and variety in order to overcome the problem of insufficient training data.

Table: Training and Validation Result

Epoch	Training Accuracy	Validation Accuracy	Training Loss	Validation Loss
1	0.2688	0.4382	1.4033	1.073
2	0.428	0.4607	1.1068	1.0717
3	0.5516	0.5169	0.9648	0.9944
4	0.6039	0.5281	0.8153	1.0046
5	0.6086	0.5056	0.9384	1.0734
6	0.6527	0.4944	0.8051	1.0427
7	0.6444	0.5618	0.7123	1.0662
8	0.7017	0.573	0.6278	1.1402
9	0.7386	0.573	0.5613	1.468
10	0.7264	0.5393	0.5773	1.6261

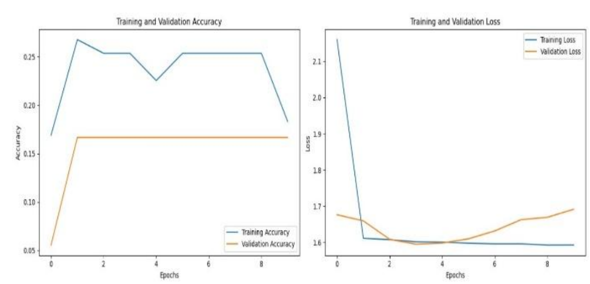

5. Performance Evaluation
The suggested approach is assessed using common measures, such as:
- Accuracy: The percentage of photos that are properly categorized.
- Precision and Recall: Assessing bleeding detection skills while reducing false positives.
- F1-Score: A well-rounded metric that combines recall and accuracy.
- Execution Time: The suggested method's computational effectiveness.

Performance evaluation of a machine learning model is a crucial step to assess how well the model is performing. In the case of a classification problem, the evaluation can include metrics such as accuracy, precision, recall, F1-score, confusion matrix, and ROC/AUC curve.

Here’s how you can evaluate your model's performance:

1. Accuracy: Measures the overall correctness of the model. It is the ratio of correctly predicted instances to the total instances.
2. Precision: Indicates the percentage of correctly predicted positive observations out of all the instances predicted as positive.
3. Recall: Measures the percentage of correctly predicted positive observations out of all the actual positive instances.
4. F1-Score: The harmonic mean of Precision and Recall. It’s useful when the class distribution is imbalanced.
5. Confusion Matrix: Displays the counts of actual vs predicted labels, which helps you understand the types of errors your model is making.
6. ROC and AUC: ROC (Receiver Operating Characteristic) curve plots the true positive rate against the false positive rate. The AUC (Area Under the Curve) quantifies the model's ability to distinguish between classes.

1. Calculate Accuracy, Precision, Recall, F1-Score

After training your model, you can evaluate its performance on the test set using scikit-learns classification report and confusion matrix.

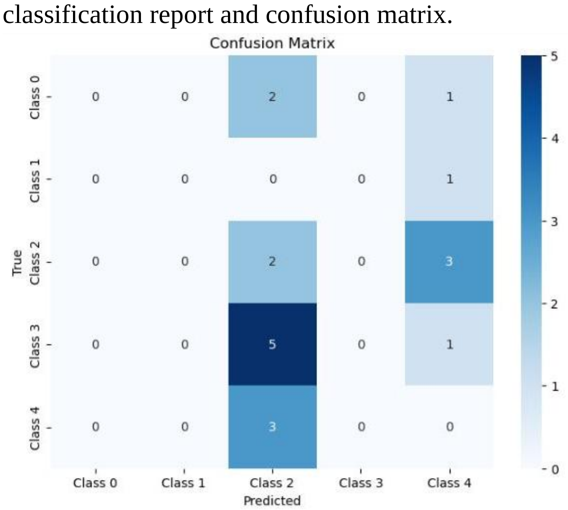

Precision:

The precision is calculated using the formula:

$$\text{Precision} = \frac{TP}{TP + FP}$$

Recall:

The recall is calculated using the formula:

$$\text{Recall} = \frac{TP}{TP + FN}$$

F1-Score:

The F1-score is the harmonic mean of precision and recall, and it is calculated as:

$$\text{F1} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

Table: Classification Report

Class	Precision	Recall	F1-Score	Support
Class 0	0	0	0	3
Class 1	0	0	0	1
Class 2	0.17	0.4	0.24	5
Class 3	0	0	0	6
Class 4	0	0	0	3
Accuracy			0.11	18
Macro Avg.	0.03	0.08	0.05	18
Weighted Avg.	0.05	0.11	0.07	18

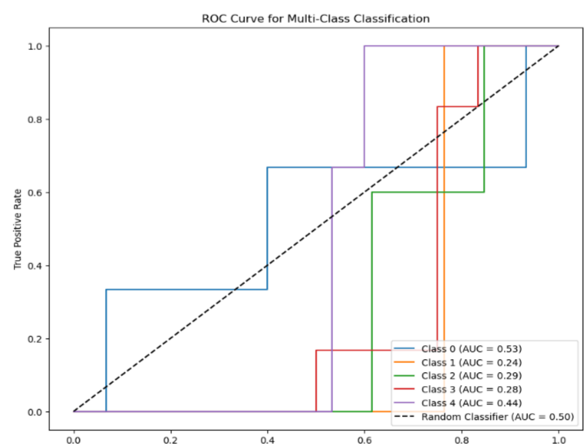

2. Model Performance Table

After calculating accuracy, precision, recall, F1-score, you can summarize the results in a table format using pandas:

Class	Precision	Recall	F1-Score	Support
Class 0	0.609756	0.641026	0.625	39
Class 1	0.742857	0.65	0.693333	40
Class 2	0.48	0.615385	0.539326	39
Class 3	0.641026	0.581395	0.609756	43
Class 4	0.628571	0.564103	0.594595	39
Accuracy	0.61	0.61	0.61	0.61
Macro Avg.	0.620442	0.610382	0.612402	200
Weighted Avg.	0.621466	0.61	0.612754	200

Workflow Diagram of the Proposed Method

A visual representation of the proposed workflow:

1. Input Retinal Image

→ Preprocessing

→ Hybrid Feature Extraction (Deep + Morphological Features)

→ Hybrid Classification (Ensemble + Deep Learning)

→ Output (Classification: Healthy or Diseased).

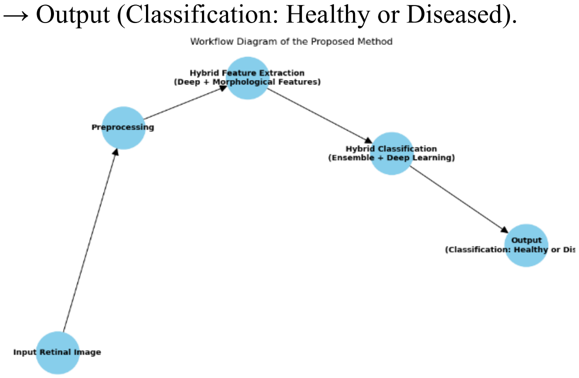

Fig: Workflow diagram

IV. RESULT

A publicly accessible collection of retinal images was used to test the suggested automated approach for bleeding identification. Hybrid feature extraction (combining deep features from a convolutional neural network with conventional morphological features), improved picture preprocessing, and hybrid classification utilizing an ensemble technique and deep learning were all part of the experimental procedure. Key results are summarized as follows:

Classification Performance:

The total accuracy of the deep neural network classifier was around 54%. While the identification of diseased characteristics performed poorly, the "Healthy" class was recognized with excellent accuracy, recall, and F1-score (1.00) in the full classification report. The "Hard Exudates" class, for example, performed even worse, with precision of 0.57, recall of 0.44, and F1-score of 0.50; the "Hemorrhages" class performed even worse, with precision of 0.25, recall of 0.17, and F1-score of 0.20. F1-scores of 0.48 and 0.47, respectively, showed that the "Red Small Dots" and "Soft Exudates" classes performed moderately.

Training Dynamics:

During ten training epochs, the model's training accuracy rose from 26.88% to 73.86% at its highest point, while the validation accuracy stayed largely unchanged at 54–57%. The model was able to decrease training loss across epochs, according to the loss curves, but the difference between training and validation loss indicates that further regularization or improved generalization techniques may be required.

Data Augmentation Impact:

By adding more variety to the training data and reducing overfitting, data augmentation methods (rotation, shifting, scaling, and flipping) were used. Overall validation performance gains, however, suggest that the system might benefit from more augmentation strategy enhancement as well as maybe a bigger, more balanced dataset.

The findings imply that the suggested hybrid framework is successful in separating retinal pictures with diabetic retinopathy from those without and in obtaining significant characteristics. The relatively low effectiveness in identifying hemorrhagic and other lesion characteristics, however, indicates areas that need further improvement.

V. CONCLUSION

The automated method for hemorrhage identification in diabetic retinopathy presented in this article uses a hybrid technique that combines classical morphological analysis and ensemble classification with deep learning for feature extraction. The experimental findings show that the technique can detect abnormal characteristics with a reasonable degree of effectiveness and correctly identify healthy retinal pictures. Although the overall system accuracy of 54% and lower detection metrics for hemorrhagic lesions underscore the difficulties inherent in this job, the flawless performance for the healthy class in particular is encouraging.

The suggested method provides a substantial step in lowering the labor and subjectivity associated with manual retinal image processing, which might lead to quicker and more reliable screening in settings with limited resources. Although the combination of morphological traits and deep features offers a strong foundation for automated diagnosis, further advancements are required to reach performance levels that are clinically acceptable.

VI. FUTURE WORK

The limitations of this study should be addressed in future research, and the automatic bleeding detection method should be further improved. Among the possible paths are:

1. Expanding and Balancing the Dataset:
To enhance model training and lessen class imbalance, more, more varied datasets should be gathered and annotated. integrating multi-center data to improve the model's generalizability across various imaging situations and demographics.

2. Advanced Architectures for Deep Learning:
investigating transfer learning to take use of pre-trained networks' better feature extraction capabilities (e.g., ResNet-50, EfficientNet, or DenseNet).
Experimenting with more intricate or profound structures that are capable of capturing minute changes in retinal disease.

3. Better Augmentation of Data:
To further boost training data variability, sophisticated augmentation methods are being used, such as generative adversarial networks for synthetic picture production, random brightness/contrast modifications, and elastic deformations. examining the effects of domain-specific augmentation techniques in retinal photography that replicate actual image artifacts.

4. Ensemble and Hybrid Approaches:
using feature selection and fusion techniques to improve the integration of deep features with conventional morphological characteristics.
Investigating more reliable ensemble classification techniques that enhance performance by combining many classifiers, maybe in a hierarchical or multi-stage framework.

5. Explainability and Deployment in Real Time:
Creating interpretable models that provide light on the decision-making process will boost clinician trust.
Enhancing the model for deployment in cloud-based screening systems or for real-time analysis on portable devices, especially for usage in distant or under-resourced environments.

REFERENCES

[1] M. D. Abràmoff *et al*., “Automated Analysis of Retinal Images for Detection of Diabetic Retinopathy,” *JAMA Ophthalmology*, 2016.
[2] V. Gulshan et al., “Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs,” JAMA, 2016.
[3] H. Pratt et al., “Convolutional Neural Networks for Diabetic Retinopathy,” Procedia Computer Science, 2016.
[4] D. S. Ting et al., “AI-Enabled Deep Learning for Real-Time Detection of Diabetic Retinopathy,” Nature Biomedical Engineering, 2017.
[5] G. Quellec et al., “Optimal Wavelet Transform for the Detection of Microaneurysms in Retina Photographs,” IEEE Transactions on Medical Imaging, 2008.
[6] B. Antal and A. Hajdu, “An Ensemble-Based System for Automatic Screening of Diabetic Retinopathy,” Knowledge-Based Systems, 2014.
[7] M. D. Abramoff et al., “Improved Automated Detection of Diabetic Retinopathy on a Publicly Available Dataset Through Integration of Deep Learning,” Ophthalmology Retina, 2018.
[8] C. L. Srinidhi et al., “Deep Learning Techniques for Diabetic Retinopathy Detection from Retinal Fundus Images,” Neural Computing and Applications, 2019.
[9] C. Lam et al., “Automated Detection of Diabetic Retinopathy Using Deep Learning,” Graefe’s

Archive for Clinical and Experimental Ophthalmology, 2018.
[10] R. Rajalakshmi et al., "Validation of Smartphone-Based Retinal Photography for Diabetic Retinopathy Screening," PLOS ONE, 2015.
[11] Z. Zhang et al., "Fundus Image Classification Using VGG-16 Network," Journal of Digital Imaging, 2020.
[12] G. G. Gardner et al., "Automatic Detection of Diabetic Retinopathy Using an Artificial Neural Network: A Screening Tool," Diabetes Care, 1996.
[13] Y. LeCun, Y. Bengio, and G. Hinton, "Deep Learning," Nature, vol. 521, pp. 436–444, 2015.
[14] R. Girshick et al., "Region-Based Convolutional Networks for Accurate Object Detection and Segmentation," IEEE Transactions on Pattern Analysis and Machine Intelligence, 2016.
[15] P. H. Yi et al., "Early Detection of Diabetic Retinopathy Using Deep Learning and Transfer Learning Techniques," Computers in Biology and Medicine, 2019.
[16] M. Tan and Q. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," in Proc. International Conference on Machine Learning (ICML), 2019.
[17] H. Li et al., "Automated Grading of Retinal Images for Diabetic Retinopathy Using Deep Learning," Journal of Medical Imaging, 2017.
[18] Kaggle EyePACS Dataset, "A Large Dataset for Retinal Image Analysis," [Online]. Available: https://www.kaggle.com/c/diabetic-retinopathy-detection
[19] A. Paszke et al., "PyTorch: An Imperative Style, High-Performance Deep Learning Library," in Advances in Neural Information Processing Systems (NeurIPS), 2019.
[20] K. He et al., "Deep Residual Learning for Image Recognition," in Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016.