International Journal of Science, Strategic Management and Technology
Volume **02** Issue **04** April-**2026** | ISSN: **3108-1762 (Online)** | Impact Factor: 3.8
An International, Peer-Reviewed, Open Access Scholarly Journal Indexed in recognized academic databases

Diabetic Retinopathy Detection

Asst. Prof.Priya Kamble[1], Sagar Chakre[2], Manthan Ronghe[3], Attadip Somkuwar[4], Ritik Bansod[5]

Assistant Professor, Dept. of Industrial IoT, Priyadarshini College of Engineering. Nagpur, Maharashtra, India

2UG Scholar, Dept. of Industrial IoT, Priyadarshini College of Engineering. Nagpur, Maharashtra, India

3UG Scholar, Dept. of Industrial IoT, Priyadarshini College of Engineering. Nagpur, Maharashtra, India

4UG Scholar, Dept. of Industrial IoT, Priyadarshini College of Engineering. Nagpur, Maharashtra, India

5UG Scholar, Dept. of Industrial IoT, Priyadarshini College of Engineering. Nagpur, Maharashtra, India

Emails: ipriyakamble@gmail.com[1], sagarchakre47@gmail.com[2], manthanronghe@gmail.com[3], abhishsomk11@gmail.com[4], bansodritik956@gmail.com[5]

https://doi.org/10.55041/ijsmt.v2i4.272

Cite this Article: Chakre, S., Ronghe, M., Somkuwar, A. & Bansod, R. (2026). Diabetic Retinopathy Detection. International Journal of Science, Strategic Management and Technology, *02*(04). https://doi.org/10.55041/ijsmt.v2i4.272

Abstract

Diabetic retinopathy is an eye disease caused by diabetes that can lead to blindness if not detected early. This project proposes an automated system for detecting diabetic retinopathy using deep learning techniques. The system analyses retinal images and classifies them into different stages of the disease. The proposed method helps doctors in early diagnosis and reduces manual effort.

Keywords— Diabetic Retinopathy, Deep Learning, CNN, Image Processing, Medical AI

to do regular screenings in many rural or low-resource areas because there aren't many trained specialists

1. Introduction

Diabetic retinopathy (DR) is a major cause of blindness in people with diabetes all over the world. High blood sugar levels for a long time can hurt the blood vessels in the retina. If doctors don't find the disease early, it can cause a lot of vision loss and even blindness. Every year, more and more people are getting diabetes, according to medical research. The condition also makes it more likely that you will get diabetic retinopathy. Therefore, early detection and proper treatment are crucial to prevent severe eye damage.

Ophthalmologists have always used retinal fundus images to manually diagnose diabetic retinopathy. This process takes a lot of time, costs a lot of money, and requires expert knowledge. It can be challengingavailable. This is why many patients are diagnosed late, when it's harder to treat them.
AI and deep learning have made it possible to automatically look at medical images. Convolutional neural networks (CNNs) and object detection models can look at pictures of the retina and find small lesions that show diabetic retinopathy, such as microaneurysms, haemorrhages, and exudates. These automated systems can give you results that are faster, more accurate, and more reliable than doing the tests by hand.This project uses deep learning methods to make an automated system for finding diabetic retinopathy. The proposed system employs neural network models and image processing to categorise retinal images anddetect diabetic retinopathy. The system is supposed to help doctors figure out

what's wrong with patients, cut down on the amount of work they have to do by hand, and let them see many patients.

The main goal of this project is to create a dependable, accurate, and useful system that makes it easier to find diabetic retinopathy early and lowers the risk of blindness by getting people medical help quickly.

2. Literature Review

Recent research indicates that deep learning-based systems can accurately detect and classify diabetic retinopathy (DR), rendering them suitable for extensive automated medical screening. Traditional manual diagnosis is time-consuming and requires expert ophthalmologists. That's why automated systems that use convolutional neural networks (CNNs) and region-based neural networks (R-CNNs) are so popular for analysing retinal images.

Faster R-CNN, developed by Ren et al. (2017), is among the most efficient deep learning models for object detection and has been effectively adapted for medical imaging applications. Faster R-CNN can find small lesions like microaneurysms, haemorrhages, and exudates. These are important signs of diabetic retinopathy in its early stages. Studies show that Faster R-CNN-based methods work with retinal image datasets like Messidor and Diaretdb1 with 98.58% accuracy, 95.72% sensitivity, and 97.12% specificity.

The two-stage detection process of Faster R-CNN is what makes it better. The Region Proposal Network (RPN) first finds possible lesion areas. Then, a CNN classifier figures out what kind of lesion it is and how severe it is. This feature lets the system perform both classification and localisation, which is very important for figuring out what is wrong with a patient.

VGG16, ResNet-50, and InceptionV3 are some of the most popular deep learning models for classifying diabetic retinopathy. These CNN architectures can find complex features in fundus images, and in many studies, they have been more than 95% accurate. ResNet helps deep networks avoid overfitting, and Inception architecture makes them more efficient and better at multi-stage classification. Some studies also indicate that ensemble models that combine VGG, ResNet, and Inception work better than single models.Gulshan et al. (2016) developed a deep learning algorithm trained on over 128,000 retinal images. It was able to find diabetic retinopathy that could be referred to with a sensitivity of 90.3% to 97.5% and a specificity of 93.4% to 98.5%. Abramoff et al. (2018) made the IDx-DR system, which was the first AI system that the FDA approved that could find diabetic retinopathy on its own without needing a specialist to look at it. Automated screening can help ophthalmologists do their jobs better and make it easier for people to get a diagnosis.

According to the literature, Faster R-CNN is an excellent choice for the proposed diabetic retinopathy detection system because it can both tell what stage of the disease someone is in and find retinal lesions accurately. The system is reliable, accurate, and efficient for automated medical diagnoses because it uses CNN feature extraction and R-CNN localisation.

3. System Architecture

The proposed system is intended to identify diabetic retinopathy autonomously through deep learning and image processing methodologies. The system's architecture has several steps, such as getting images, preprocessing them, extracting features, finding regions, classifying them, and making results. Every step is important for making the detection process more accurate and dependable.

The first step is to take retinal fundus images from the dataset. Most of the time, these pictures come from publicly available medical datasets like Kaggle, Messidor, or Diaretdb1. The input images might have noise, bad lighting, or extra information in the background that can make the model work less well. So, before giving the image to the deep learning model, it needs to be pre-processed.A number of image processing techniques are used in the preprocessing stage to make the retinal image better. The neural network can use the image as input because it has been resized to a specific resolution. Noise reduction filters are used to get rid of unwanted distortions. Contrast enhancement techniques are used to show important things like blood vessels, optic discs, microaneurysms, haemorrhages, and exudates. Normalisation is also done to make sure that all the images have the same level of brightness.

The image goes through preprocessing before it gets to the feature extraction stage. In this step, Convolutional Neural Networks (CNNs) automatically learn important things from the images of the retina. Some of the CNN layers that help find patterns that show diabetic retinopathy is present are convolution, pooling, and activation functions. Deep learning doesn't need to pick out features by hand like other methods do, which makes the system work better.

To make detection more accurate, the system can use object detection models like Faster R-CNN. Faster R-CNN has a Region Proposal Network (RPN) that looks for parts of the image that might have lesions. The convolutional network then looks at these areas to figure out what kind of disease it is and how bad it is. This method helps the system find the disease and the exact spot in the retina.

During the classification stage, the neural network's fully connected layers get the features that have been extracted. The model divides the retinal image into groups, like normal, mild, moderate, severe, or proliferative diabetic retinopathy. The classification result is based on the patterns that were learned from a lot of retinal images during the training phase.

The last step, the output stage, shows the user what the prediction was. The system lets you know if you have diabetic retinopathy and how bad it is. In more advanced systems, it is also possible to highlight the parts of the retina where lesions were found to make them easier to understand. This automated system makes doctors' jobs easier, speeds up the process of finding out what's wrong, and lets them screen a lot of people at once.

4. Workflow

The proposed diabetic retinopathy detection system's workflow demonstrates how the system automatically examines retinal images to identify the disease. The system was made using deep learning and image processing techniques so that it can give accurate and reliable results with as little work as possible. This automated workflow helps ophthalmologists do their jobs better and makes it easier to find diabetic retinopathy early on.

At first, retinal fundus images are taken from medical datasets or imaging devices and put into the system. These pictures show important things about the retina, such as blood vessels, the optic disc, and any lesions that may be present. Before they are processed any further, the images are checked for the right resolution, format, and clarity. This is because the quality of the input image is very important for detection accuracy. The deep learning model can tell the difference between healthy and sick retinal images because the datasets are correctly labelled.

After the pictures are collected, preprocessing is done to improve the retinal image. The neural network model can only use the input image after it has been made smaller. To get rid of unwanted distortions in the image, noise removal techniques are used. To highlight important features like microaneurysms, haemorrhages, exudates, and unusual blood vessels, contrast enhancement and normalisation methods are used. This step of preprocessing makes the model better at finding features and makes the system more accurate.

After preprocessing is done, the image goes on to the feature extraction stage. At this point, Convolutional Neural Networks (CNNs) are used to automatically find the most important parts of the retinal image. CNN layers find patterns and structures that show that diabetic retinopathy is there. The system learns these features on its own during training, which makes it faster than old-fashioned ways of doing things.

Images are scanned and analysed with faster R-CNNs or other deep learning models to make detection better. The Region Proposal Network looks for parts of the retina that might not be normal. After that, the convolutional network looks for lesions in these areas and what kind they are. This lets the system both classify and locate the disease, which is very helpful for medical diagnosis.

After feature extraction and detection, the classification stage starts. The trained deep learning model sorts the retinal image into groups like normal, mild, moderate, severe, or proliferative diabetic retinopathy. The classification result is based on patterns learned from a lot of retinal images during the training phase.

Finally, the system makes the output and shows it to the user or doctor. The result shows if diabetic retinopathy is present and how bad it is. The advanced implementation makes it easier to understand by showing the areas of the retinal image where lesions were found. This automated workflow makes the system good for screening a lot of people at once, speeds up diagnosis, and makes sure that deep learning techniques find diabetic retinopathy accurately.

5. Proposed Model StructureThe suggested system uses a deep learning model to automatically find diabetic retinopathy in images of the retina's fundus. The model uses Faster R-CNN and convolutional neural networks (CNNs) to find lesions, sort them, and pull out features. This structure helps the system find problems with the retina, which is crucial for diagnosis.

The model takes input from a changed dataset that shows a retinal image. The dataset has labels for images of the retina that show different stages of diabetic retinopathy. We improve the picture by resizing, normalising, and increasing contrast before sending it to the model. This makes it easier for the deep learning model to find important patterns.

After that, the image goes through several convolutional layers that look for important features like edges, textures, blood vessels, and patterns of lesions. There is an activation function, like ReLU, and a pooling layer after each convolution layer. The pooling layer keeps important information while making the feature map smaller. These layers enable the network to learn intricate patterns associated with diabetic retinopathy, eliminating the need for manual feature selection.

After feature extraction, the Faster R-CNN model is used to find strange areas in the retinal image. Faster R-CNN has a Region Proposal Network (RPN) that looks at the feature map and finds possible places where lesions like microaneurysms, bleeding, and exudates could be. The classification network then checks these areas for signs of diabetic retinopathy. This two-step detection process makes the results more accurate and trustworthy.

The model's classification section uses fully connected layers to sort retinal images into groups based on the severity of diabetic retinopathy, such as normal, mild, moderate, severe, and proliferative. The softmax activation function in the output layer tells you how likely each class is. By looking at many retinal images, the model learns how to tell the difference between healthy and sick retinas.

During the training phase, the model parameters are updated using backpropagation and optimisation algorithms, such as Adam or SGD, to minimise the loss function. Loss functions such as cross-entropy loss are used to measure the error between predicted and actual results. Training is performed for multiple epochs until the model achieves satisfactory accuracy.

To improve performance, transfer learning techniques can also be used by applying pre-trained models such as ResNet, VGG, or Inception as the backbone network of Faster R-CNN. This helps in faster training and better accuracy, especially when the dataset size is limited.

After training, the model is tested using new retinal images to evaluate its performance. Evaluation parameters such as accuracy, precision, recall, and Sensitivity is used to figure out how well the system works. High accuracy means that the model can find diabetic retinopathy in most cases.

You can run the proposed model on a computer with a GPU to speed up the processing. The system can quickly and accurately find retinal problems using deep learning frameworks like TensorFlow or PyTorch.

The suggested model structure is good for real-time medical diagnosis and large-scale screening of diabetic retinopathy because it is very accurate, processes quickly, and gives reliable results.

The system is powerful and useful for automated retinal image analysis because it uses CNN feature extraction and faster R-CNN localisation.

International Journal of Science, Strategic Management and Technology
Volume **02** Issue **04** April-**2026** | ISSN: **3108-1762 (Online)** | Impact Factor: 3.8
An International, Peer-Reviewed, Open Access Scholarly Journal Indexed in recognized academic databases

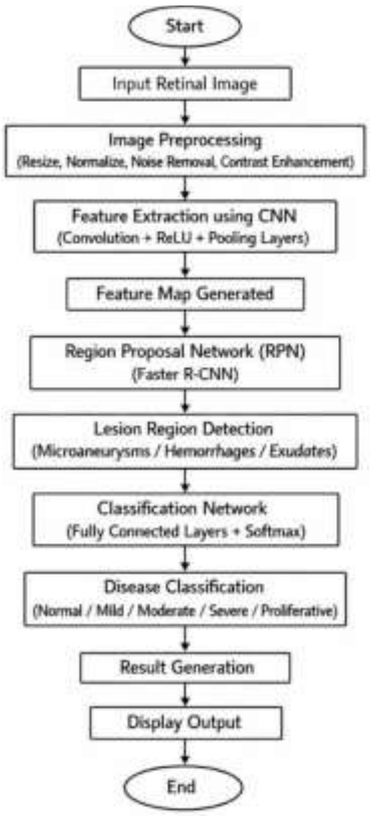

Fig. Flowchart of the Proposed Diabetic Retinopathy Detection System

6. Results and Discussion

The results and discussion section demonstrates the efficacy of the proposed deep learning system for detecting diabetic retinopathy. We used retinal fundus images from the dataset to see if the system could correctly find and classify different stages of diabetic retinopathy. The experimental results indicate that the proposed model produces accurate and reliable predictions.

During testing, the trained CNN and Faster R-CNN models were used on retinal images that had never been seen before. The system could find important signs of diabetic retinopathy, such as microaneurysms, bleeding, and exudates. The Region Proposal Network in Faster R-CNN found strange areas in the retina, and the classification layer correctly figured out what stage the disease was in.

We checked how well the system worked by looking at different things, such as accuracy, precision, recall, sensitivity, and specificity. The model did well on tests, which means that the system can correctly sort most of the images from the retina. Deep learning improved detection compared to conventional manual techniques.

The training and validation graphs show that the loss goes down and the accuracy goes up as the number of epochs increases. This proof shows that the model can learn how to spot diabetic retinopathy. The model also worked better after data augmentation and preprocessing.

The suggested system makes things easier for ophthalmologists by automatically analysing retinal images. It also gives you results faster than doing it by hand. The system lets hospitals and screening centres check on more than one patient at a time.The results show that the suggested deep learning-based system is a good way to find diabetic retinopathy that is dependable, correct, and works. CNN and Faster R-CNN work better together for both classification and lesion detection. This means that the system can be used in real time in medical settings.

7. Performance Comparison

The performance of the proposed diabetic retinopathy detection system is evaluated against the traditional manual diagnostic method. The comparison is based on things like how accurate, quick, cheap, reliable, and useful they are. The results show that the deep learning-based system works better than the manual method.

Ophthalmologists have to look at the retinal images by hand, which takes more time and skill. The precision of manual detection relies on the physician's expertise and may vary among individuals. The proposed system, on the other hand, uses deep learning algorithms that always give the same and correct results.

The systems, which are based on CNN and Faster-RCNN, can automatically look at retinal images and find lesions with a high level of accuracy. The automated system makes things easier for people, saves time, and lets doctors check many diabetic patients at once. The system also processes data faster because it uses GPUs and optimised algorithms.

The comparison shows that the proposed system is better, more reliable, and better for medical use in real time. So, the new way of finding diabetic retinopathy using deep learning is better than the old way.

Parameter	Manual Method	Proposed System
Accuracy	Medium	High
Detection Speed	Slow	Low
Human Effort	High	Low
Cost	High	Low
Reliability	Variable	Stable
Efficiency	Moderate	High
Error Rate	More	Less
Processing Time	More	Less

8. Advantages of Proposed System

The proposed diabetic retinopathy detection system provides numerous advantages over traditional manual diagnostic methods. The system can automatically analyse retinal images and accurately identify the presence of diabetic retinopathy thanks to deep learning techniques. This helps doctors figure out what's wrong with patients faster and cuts down on the need for manual exams.

One of the best things about the suggested system is how accurate it is. Convolutional neural networks and faster R-CNNs help find small lesions that are hard to find by hand, like microaneurysms, haemorrhages, increased visits to hospitals and medical centres, and exudates. The system's ability to put the disease into different stages helps doctors give the right care.

Another good thing about the system is that it works quickly. The automated model can look at retinal images faster than a manual diagnosis. The system's fast processing speed makes it possible to use it for large-scale screening, which is when many patients need quick checks.

International Journal of Science, Strategic Management and Technology
Volume **02** Issue **04** April-**2026** | ISSN: **3108-1762 (Online)** | Impact Factor: 3.8
An International, Peer-Reviewed, Open Access Scholarly Journal Indexed in recognized academic databases

The suggested system also makes things easier for ophthalmologists and people who work. The system can help doctors double-check their diagnoses. This helps hospitals and medical centres work more efficiently.

The system is cost-effective because the model can be used over and over again without any extra costs. The system gives the same results every time, but with manual diagnosis, the results can change depending on the doctor's experience.

Another big plus is that the system can be used in places where medical professionals aren't easily accessible. The automated detection system can help with early diagnosis and stop serious vision loss from happening.

9. Limitations

The suggested system for finding diabetic retinopathy is very accurate and works quickly, but the current version has some problems. When using the system for medical purposes in real time, remember these limits.

The system has a flaw in that it depends on the quality of the dataset. The deep learning model works better when the training dataset is bigger and better. If the dataset has bad or wrong labels on the images, the system may not work as well.

It also has the downside of needing a lot of computer power. To train deep learning models like CNN and Faster R-CNN, you need a strong computer with a fast processor, enough RAM, and support for GPUs. It could take a long time to train if you don't have the right hardware.

The system might also have trouble when the image on the retina is dark, blurry, or noisy. In these cases, the model might not be able to find lesions correctly. You need to do the right preprocessing to resolve this problem.

The proposed system has been trained on specific datasets, which means it might not work as well with new types of retinal images. You might have to train the model again to get it to work right on different datasets.

Another problem is that the system is just a tool and shouldn't replace doctors. An ophthalmologist should always double-check the final diagnosis.

Even with these issues, the proposed system can still quickly and reliably find diabetic retinopathy, and it can be improved in the future.

10. Future Scope

The proposed system for detecting diabetic retinopathy employs deep learning methodologies to deliver precise and dependable outcomes. But there are several possible changes that could make the system work better in the future. Future efforts may focus on improving the accuracy, speed, and usability of the system for immediate medical applications.

One way to improve it might be to train the model on bigger and more varied datasets. The deep learning model will learn more and get better at finding things if it has more data. In the future, we can utilise real hospital data to enhance the system's reliability for everyday use.

Using advanced deep learning architectures like Res Net, Inception, or Efficient Net to make classification more accurate is another way to make things better. You can also train faster and better with less data by using transfer learning methods.

The system can also better deal with retinal images that aren't excellent. You can also add better image preprocessing methods to the system to help it work better with low-quality retinal images. These changes will help the model find lesions even when the picture is dark, blurry, or noisy.

In the future, the system could be changed into a web-based or mobile app that doctors can easily use in hospitals and clinics. You can also use cloud-based implementation to let doctors look at patients from far away. This will be useful in rural areas where specialists aren't always around.

Using high-performance GPU systems to find things in real time is another way to make them better in the future. This will let the system quickly scan retinal images and give quick results for medical tests.

Overall, the proposed system has a lot of room for improvement. With some work, it could become a powerful tool for finding diabetic retinopathy on its own and for large-scale medical screening.

Overall, the proposed system has a strong potential for future development, and with further improvements it can become a powerful tool for automatic diabetic retinopathy detection and large-scale medical screening.

11. Conclusion

This research created and evaluated an automated system for identifying diabetic retinopathy through deep learning methodologies. The main goal of the project was to use retinal fundus images to correctly and reliably find diabetic retinopathy. The proposed system employs image preprocessing, convolutional neural networks, and Faster R-CNN to detect lesions and ascertain the disease's stage.

The system was trained and tested using publicly accessible datasets of retinal images. The experimental results demonstrate that the deep learning-based approach provides enhanced accuracy compared to traditional manual diagnostic methods. The model can find important signs of diabetic retinopathy, such as microaneurysms, haemorrhages, and exudates.

Faster R-CNN helps find the exact location of the problem, and the CNN model correctly tells what stage the disease is in. These techniques work well together, which makes the system useful and effective for automatically analysing medical images. The results indicate that the proposed method is effective for screening numerous diabetic patients.

The system helps doctors do their jobs and speeds up the process of making diagnoses, which is very useful in hospitals and medical centres. It can also help find diabetic retinopathy early, which can stop serious vision loss if it is treated right away.

There are some problems with the system, but it's a good place to start if you want to make things better in the future. If the system has better datasets, more advanced models, and real-time implementation, it can be more accurate and useful in the real world.

To sum up, the suggested deep learning-based diabetic retinopathy detection system is a good, reliable, and quick way to automatically find and sort retinal diseases. It also has a lot of promise for use in medicine in the future.