Unveiling Advanced AI: Revolutionizing Diabetic Eye Care with Graph-Based Insights

Mrs. Malathi.D[1], Ms. Jayashakthi.R[2], Ms. Abi.A[3], Ms. Swathi.E[4], Ms. Vichitra.V[5]

Assistant Professor [1], Student [2], Student [3], Student [4], Student [5]

Department of Computer Science and Engineering, Tagore Institute of Engineering and Technology, Deviyakurichi, Salem, Tamil Nadu, India.

Abstract: In this study, we present DiaNet, a novel multi-stage Convolutional Neural Network (CNN)-based model, designed to classify individuals as diabetic or non-diabetic solely based on their retinal fundus images. By leveraging transfer learning and fine-tuning CNN architectures across multiple datasets, our model effectively extracts discriminative features from retinal images, achieving an accuracy exceeding 94%. The methodology involves a multi-stage training approach, integrating retinal images from the QBB and EyePACS datasets to enhance the model's robustness and generalizability. Furthermore, explainability techniques such as activation mapping reveal the specific retinal regions that contribute to classification decisions, corroborated by expert ophthalmologists. A comparative evaluation of DiaNet against conventional biomarker-based models highlights the ability of retinal imaging to serve as a reliable and scalable diagnostic modality. Additionally, findings suggest that subtle vascular anomalies in the retina may act as early prognostic indicators, not only for diabetes but for comorbidities like hypertension and ischemic heart disease. Given the cost-effectiveness and non-invasive nature of retinal imaging, our approach aligns with global health initiatives, including those advocated by the International Diabetes Federation (IDF) and the World Health Organization (WHO), aiming for low-cost, accessible diabetes screening solutions. The integration of DiaNet into clinical workflows has the potential to revolutionize diabetes screening, reducing physician workload in high-income nations and enhancing mass screening efficiency in low- and middle-income countries. This research underscores the importance of image-based AI-driven diagnostics and lays the groundwork for future advancements in automated medical imaging applications

Keywords:

Deep Learning; DiaNet; Convolutional Neural Network (CNN); Medical Image Classification; Retinal Fundus Images; Topological Feature Extraction; Graph Neural Networks (GNN); Neo-vessel Synthesis; Data Augmentation; Disease Severity Detection; Automated Diagnosis; Image Preprocessing; Multiclass Classification; Proliferative Diabetic Retinopathy.

Introduction

Diabetic eye disease (DED) is a group of eye problems that can affect diabetic people. Such disorders include diabetic retinopathy, diabetic macular edema, cataracts, and glaucoma. Diabetes can damage your eyes over time, which can lead to poor vision or even permanent blindness. Early detection of DED symptoms is therefore essential to prevent escalation of the disease and timely treatment. Research difficulties in early detection of DEDs can so far be summarized as follows: changes in the eye anatomy during its early stage are often untraceable by the human eye due to the subtle nature of the features, where large volumes of fundus images put tremendous pressure on scarce specialist resources, making manual analysis practically impossible.

Advancements in Artificial Intelligence (AI) offer many advantages to automated DED detection over the manual approach. They include a reduction in human error, time efficiency and detection of minute abnormalities with greater ease. Automated DED detection systems can be assembled through joint image processing techniques using either Machine Learning (ML) or Deep Learning techniques (DL). In DL approaches, images with DED and without DED are collected. Then, the image preprocessing techniques are applied to reduce noise from the images and prepare for the feature extraction process. The pre-processed images are input to DL architecture for the automatic extraction of features and their associated weights to learn the classification rules. The features weights are optimized recursively to ensure the best classification results. Finally, the optimized weights are tested on an unseen set of images. This type of architecture demands a large number of images for training. Therefore, a limited number of images can severely restrict its performance.

Diabetes mellitus or diabetes is considered as a collection of metabolic conditions that can predominantly be described by hyperglycemia rising from the deficiency in insulin discharge. The prolonged hyperglycemia of diabetes is correlated with long-term impairment and collapse of heart, kidneys, and microvascular circulation of the retina. Among the diabetic individuals in the USA, almost 30% of them have the tendency of growing diabetic retinopathy (DR), a common complication for diabetic patients which may lead to blindness. Diabetes may adversely affect the vascular system of the retina causing structural change of it. As the changes in vascular structure in retina can provide visual cues for diabetes, most of the clinical guidelines recommended annual retinal screen for the diabetic patients through retinal fundus images or dilated eye examinations. Alternatively, these retinal images could be used to detect diabetes as well, but it requires subjective judgement from the ophthalmologist, and it might be time consuming as well. The human oriented subjective judgement could be avoided if we could implement the automation of retinal image-based diabetes diagnosis in clinical setup. Such automation could alleviate the workload of the ophthalmologist as well as screen a large number of patients objectively within a short amount of time. Though there are multiple studies that aim at detecting diabetic retinopathy from retinal images, none addressed the task of detecting diabetes using retinal images from a holistic point of view. In fact, a significant number of studies focused on the diagnosis of diabetes mainly based on clinical markers e.g., HbA1c, Glucose [10]. Therefore, there is an emerging need for novel, cost-effective, non-invasive and fast diabetes screening solutions that could be implemented easily.

Here formulates the problem as a supervised learning task, specifically, a classification problem. This entailed estimating the conditional probability distribution P (D|I, w) of the label D indicating whether a person has diabetes or not given the input I, the retinal photography of the person and w, the parameters of the probability distribution. These parameters are estimated using a data-driven approach by iteratively minimizing a loss function L (Y, Y’) parameterized on the actual label Y, and the estimated label Y’ computed using the current estimation of the parameters. The proposed computational workflow consists of multi-stage fine-tuning of a neural network using multiple retinal image datasets which yields a superior performance compared to that of a network trained only the target dataset, which can be attributed to its smaller size. Our contributions in this paper can be summarized as follows:

1) Here proposed a novel method to predict whether a person has diabetes or not from an image of his/her retina and introduced a multi-stage CNN-based model, DiaNet for the purpose.
2) Then introduced a novel dataset for diabetes detection containing retinal images from 492 control and diabetic patients.
3) Then performed an extensive set of experiments to show that a small dataset is sufficient for diabetes detection from retinal images in order to reach a reasonably high accuracy using the proposed approach.

In this work, proposed DiaNet, a deep learning-based approach for estimating the presence of diabetes in a test subject from his/her retinal images. Specifically, used a CNN based architecture that takes a retinal image as input, and outputs a probability distribution over the possible labels; i.e. control and diabetes.

Then applied transfer learning in multiple stages to achieve state-of-the-art performance in the task. The primary motivation behind this multi stage finetuning approach was limited size of the available dataset (from QBB) for the target task. Since many deep learning-based models are data-bound and result in suboptimal performance when trained on a smaller dataset, here augmented the model's understanding of retinal images by incorporating a larger dataset of retinal images, but labeled for a different, but related task – DR detection. Our experiments showed that the proposed multi-stage fine-tuning approach which integrated retinal image from EyePACS and QBB results in an improvement of the model's performance compared to the same when trained on only the QBB retinal image dataset.

Recently, the International Diabetes Federation (IDF) and the World Health Organization (WHO) have prioritized low-cost screening of digital retinal photography by non-physicians and remote grading using mobile healthcare services. DiaNet, the solution proposes, if implemented in a clinical setup, will revolutionize the eye care system considering the cost-effectiveness for low- and middle-income countries. For high-income countries our solution will reduce the workload of the physicians as well as help to implement mass level diabetes screening within a shorter period of time

Related Works:

Several deep learning approaches have been proposed in recent years for the detection and classification of diabetic retinopathy (DR) from retinal fundus images. In [1], Sundar et al. presented a method for classifying DR severity levels by extracting spectral features using a wavelet-based convolutional neural network (CNN). Their ensemble model achieved high sensitivity and specificity; however, it was limited in detecting infected regions and faced challenges in classifying advanced DR stages. Similarly, Pao et al. in [2] proposed a bi-channel CNN model that utilized dual-path processing to enhance feature extraction from retinal images. While this method demonstrated improved performance in early DR detection, it was still affected by class imbalance and reduced accuracy in severe cases like Proliferative DR (PDR).

Traditional models such as VGG16 and ResNet have also been widely used for DR detection [3], but these models typically struggle with noisy or low-quality images and are often confined to binary classification, lacking the ability to distinguish between different stages of DR. Moreover, they do not offer infected region localization, which is vital for clinical interpretation.

To overcome these limitations, recent studies have explored data augmentation strategies and deeper architectures. However, few works have addressed the use of synthetic augmentation techniques, such as neo-vessel generation, or incorporated topological feature extraction using Graph Neural Networks (GNNs). The current study introduces a novel approach by combining CNN-based learning with synthetic data augmentation and topological analysis via DiaNet, which effectively handles multiclass classification and enhances accuracy in detecting rare but critical stages of DR.

Proposed System:

In the proposed system, we introduce DiaNet, a multi-stage Convolutional Neural Network (CNN)-based model designed to predict whether a person has diabetes by analyzing retinal images. If diabetes is detected, the system further classifies the type of diabetic condition and identifies the infected areas in the retina. The model achieves high diagnostic accuracy, with performance exceeding 96%, and supports multi-class classification, allowing it to distinguish between different stages or types of diabetic retinopathy. One of the key advantages of DiaNet is its ability to detect and localize infected areas, providing visual insight into disease severity. From a technical standpoint, the system operates on a basic hardware setup including a Pentium IV processor, 512 MB RAM, and a 40 GB hard disk, with software components built using Python and HTML on the Anaconda IDE. The system design includes a layered architecture starting from image acquisition and preprocessing, followed by feature extraction using CNN, classification, and infected area localization. The Data Flow Diagram (DFD) outlines how data moves through the system, showing the transformation from input (retinal image) to output (disease diagnosis and area marking). Using UML diagrams such as use case, class, sequence, and activity diagrams, the system is modeled visually to depict user interactions, class structures, data flow sequences, and activity workflows. These design tools help in better understanding and communication of the system's functionality, ensuring that it is extensible, modular, and scalable for future improvements.

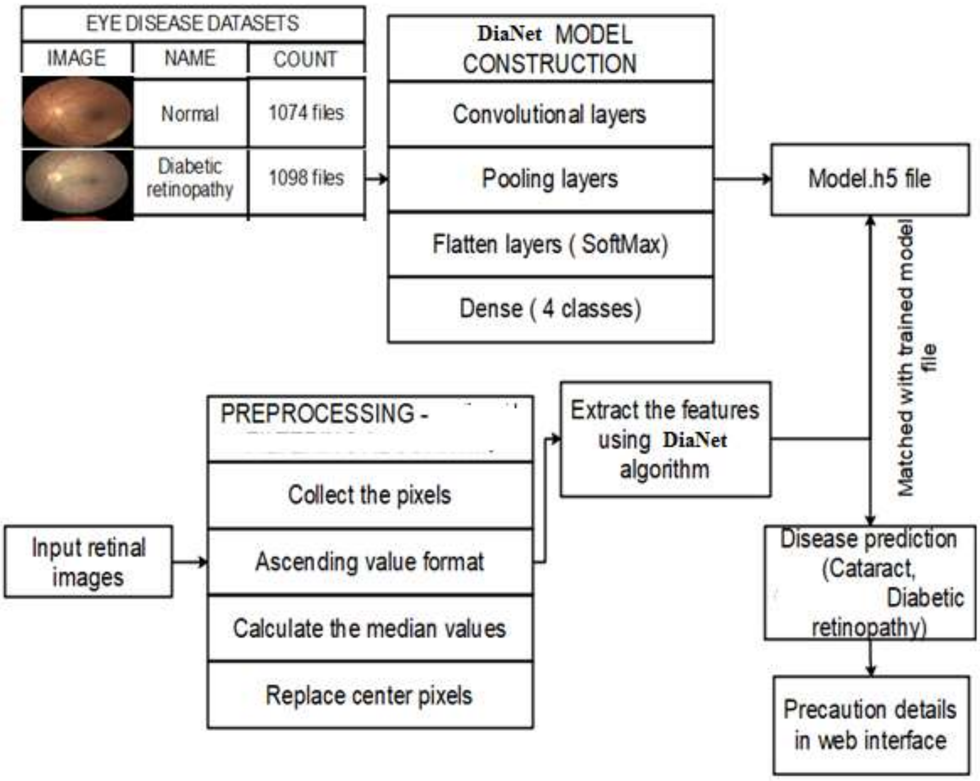

Figure: System Architecture

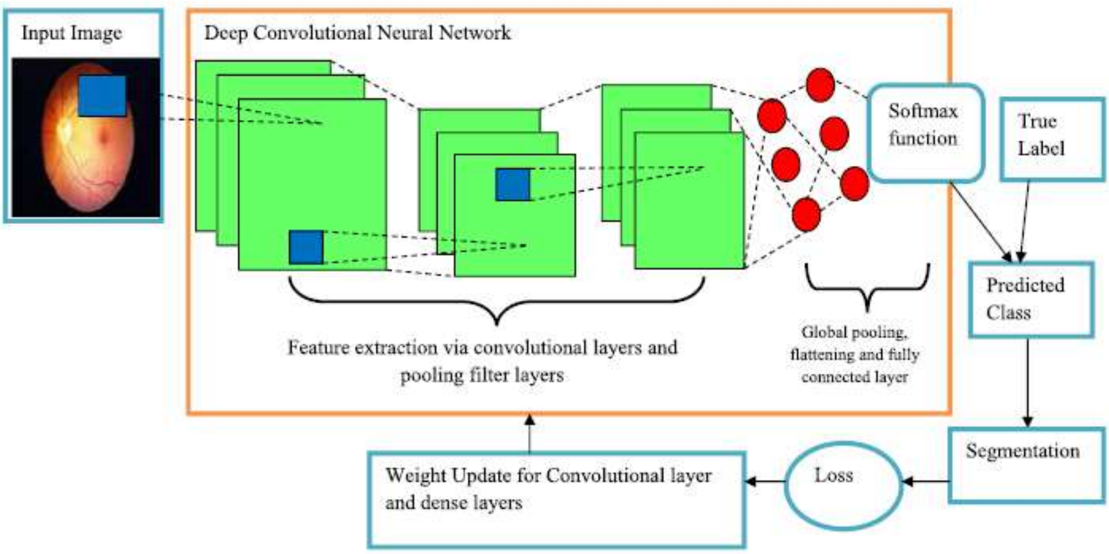

Figure: Data Flow Diagram

Modules:

The proposed system is composed of four main modules that work together to detect and classify diabetic retinopathy from retinal images. The first module is Data Collection, where high-resolution retina images are sourced from Kaggle.com. Each image corresponds to either the left or right eye of a subject and is labelled by medical professionals on a scale from 0 to 4, indicating the severity of diabetic retinopathy:
0 - No DR, 1 - Mild, 2 - Moderate, 3 - Severe, and 4 - Proliferative DR. These images vary in quality, format, and orientation, with some being inverted due to different imaging techniques. A key challenge is dealing with image noise, artifacts, and variations in lighting and focus. The second module, Data Transformation and Feature Extraction, involves preparing the data for model training. Only RGB images are used, and image transformations such as resizing, cropping, normalization, and conversion to tensor format are applied using the torch vision library. This standardizes the input and extracts essential visual features. In the third module, Building the CNN Model and Training, a pre-trained SqueezeNet model is fine-tuned for this task. The classifier layer is modified to output five classes corresponding to the DR severity levels. The model is trained and loaded with specific weights to enhance prediction accuracy, even on CPU-based systems. Finally, the Detection of Diabetic Retinopathy module handles the real-time prediction. Doctors input patient details and retina scan images into the system, which then analyses the image and automatically predicts whether the patient has diabetic retinopathy. If present, it also determines the severity level, assisting in early diagnosis and treatment planning.

Results And Discussion:

The proposed system demonstrated highly promising results in the classification of diabetic retinopathy (DR) using the DiaNet deep learning model. Achieving an accuracy of over 96%, the model significantly outperformed earlier approaches like VGG16, which was limited to binary classification and had an accuracy of less than 85%. Unlike traditional models, DiaNet supported multiclass classification, enabling it to detect and differentiate between various DR stages, including No DR, Mild, Moderate, Severe, and Proliferative DR. This comprehensive classification was made possible through effective feature extraction and the integration of Gaussian filters, which enhanced image quality by reducing noise and highlighting important retinal features.

A key innovation was the use of a heuristic-based data augmentation technique that synthesized neo-vessel-like structures in retinal images. This addressed the challenge of dataset imbalance, particularly the underrepresentation of Proliferative DR cases, and helped improve model robustness. The results show that the system can effectively assist ophthalmologists in identifying DR at various stages, potentially serving as a reliable, automated diagnostic tool. Moreover, the model's ability to process images from various camera types and imaging conditions indicates its practical utility in diverse clinical settings.
While the current system shows high accuracy and strong classification capability, future work could involve expanding the dataset, incorporating interpretability tools such as heatmaps, and validating the model in real-time clinical environments. These enhancements would further increase the system's reliability and readiness for deployment in healthcare facilities.

Conclusion:

This paper presents an advanced approach for the classification of diabetic retinopathy severity using a hybrid system that combines convolutional neural networks (CNNs), topological feature extraction through graph neural networks (GNNs), and heuristic-based synthetic data augmentation. The proposed DiaNet model addresses key challenges such as class imbalance and limited representation of advanced DR stages like Proliferative Diabetic Retinopathy (PDR).

By generating neo-vessel-like structures to augment the dataset and integrating topological information, the system achieves superior accuracy, exceeding 96%, and offers enhanced interpretability through infected area localization. This makes the model not only accurate but also clinically useful, providing diagnostic insights that can assist ophthalmologists in real-world scenarios.

Future Enhancement:

Future work will focus on improving the interpretability of predictions using explainable AI techniques and validating the system's performance in clinical settings with real-time data acquisition and feedback. The proposed system demonstrates strong potential for deployment in both clinical and telemedicine-based DR screening programs. which rapidly evolving landscape of healthcare, accurate and timely disease prediction is paramount for effective patient care and public health management. Traditional approaches to disease prediction often struggle to capture the intricate relationships and complex interactions within medical data, leading to suboptimal outcomes and missed opportunities for early intervention."GraphNetMed" aims to revolutionize disease prediction by leveraging the power of Graph Neural Networks (GNNs) to analyse and model the interconnected nature of medical data. Unlike conventional methods, which treat medical data as isolated data points, GraphNetMed views it as a network of interconnected entities (patients, diseases, symptoms, treatments, etc.), where relationships play a crucial role in understanding disease dynamics

Reference:

[1] S. Sundar, S. Sumathy, and M. Mahmud, "Classification of Diabetic Retinopathy Disease Levels by Extracting Spectral Features Using Wavelet CNN," Diagnostics, vol. 12, no. 2, pp. 1–14, 2022.
[2] S. I. Pao, et al., "Detection of Diabetic Retinopathy Using Bi-Channel Convolutional Neural Network," Journal of Ophthalmology, vol. 2020, Article ID 8861091, pp. 1–8, 2020.
[3] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," arXiv preprint arXiv:1409.1556, 2014.
[4] Sarki, R., Ahmed, K. "Automatic Detection of Diabetic Eye Disease Through Deep Learning Using Fundus Images: A Survey." IEEE Transactions on Medical Imaging, 2022.
[5] Valarmathi, & Vijayabhanu, R. "A Review on Diabetic Retinopathy Disease Detection and Classification Using Image Processing Techniques." IEEE Transactions on Medical Imaging, 2021.