Architecture
Image Studies

Vol.7, Issue 1, pp.2238-2248, 2026
DOI: https://doi.org/10.62754/ais.v7i1.1203

https://journals.ap2.pt/index.php/ais/index

GBLR-Net: Graph-Based Lesion Relationship Network for Diabetic Retinopathy Classification through Segmented Retinal Features

Debasis Deb[1], R Murugan[2]

Abstract

Diabetic retinopathy [DR] is the leading cause of vision loss in humans, and disease severity is determined not only by the presence of retinal lesions but also by their spatial distribution and inter-lesion relationships. There is a reliance on deep learning (DL) approaches that use convolutional neural networks (CNNs), which primarily capture local pixel-level patterns, limiting their ability to model clinically meaningful lesion interactions. To mitigate this, the present study incorporates the two-stage graph network that integrates lesion-level segmentation with graph-based representation learning for diabetic retinopathy classification. In the first stage, a U-Net architecture with ResNet50 backbone is employed to perform precise multi-class segmentation of pathological retinal lesions from fundus images. In the second stage, individual lesion instances are extracted from the segmentation maps and represented as nodes in a graph, where node attributes encode morphological and spatial characteristics, and edges capture spatial proximity relationships among lesions. A graph neural network (GNN) is then applied for learning high-level features for disease grading. Experimental evaluations conducted on publicly available retinal datasets demonstrate that the proposed lesion-aware graph-based model achieves improved classification performance compared to conventional convolutional baselines, particularly in distinguishing intermediate disease stages. The results indicate that explicitly modeling lesion relationships enhances both diagnostic accuracy and interpretability, offering a clinically relevant and extensible solution for automated diabetic retinopathy assessment. Our model achieves state-of-the-art performance on the Messidor-2 dataset, outperforming most existing methods across accuracy, precision, recall, and F1 score. 99.08%, 99.40%, 99.20%, and 99.10%, respectively that demonstrating its superior effectiveness compared to prior works.

Keywords: *Diabetic Retinopathy, Convolutional Neural Network, Graph Neural Network, Segmentation, Feature Extraction.*

Introduction

Diabetic retinopathy [DR] is a complicated microvascular complication of the human eye and a major cause of vision loss worldwide. The disease manifests through characteristic retinal lesions, including microaneurysms, exudates, haemorrhages, etc., whose presence, severity, and spatial distribution determine the clinical stage of DR. Early and accurate detection and classification of these pathological changes help to identify and prevent vision loss. Recent advances in DL have significantly improved automated DR analysis from fundus images. The CNNs provided a strong impact on image-level DR grade; however, most existing models rely on implicit feature learning and treat fundus images as unstructured visual inputs. As a result, they often fail to capture lesion-specific information and struggle to discriminate between intermediate disease stages, where subtle variations in lesion distribution and inter-lesion relationships are clinically important [1].

Lesion-level modeling has emerged as a promising direction to improve both performance and interpretability. Semantic segmentation networks, particularly U-Net-based architectures, have demonstrated effectiveness in localizing and delineating retinal lesions at the pixel level. Nevertheless, segmentation outputs are typically used only for visualization or auxiliary supervision, without explicitly

[1] Biomedical Imaging Laboratory (BIOMIL) Department of Electronics and Communication Engineering National Institute of Technology Silchar, Assam 788010, India
[2] Biomedical Imaging Laboratory (BIOMIL) Department of Electronics and Communication Engineering National Institute of Technology Silchar, Assam 788010, India

exploiting the spatial and morphological relationships among individual lesions. This restricts ophthalmologists' structured diagnostic and lesion identification [2].

Graph-based learning provides a principled framework for modeling relational information in non-Euclidean domains. By representing lesion instances as nodes and their spatial or morphological relationships as edges, graph neural networks (GNNs) enable explicit reasoning over lesion interactions and spatial topology. Such representations are well-suited for retinal pathology analysis, where disease severity is influenced by lesion co-occurrence patterns and spatial dispersion rather than isolated visual features [3].

In this work, we propose a lesion-aware framework for diabetic retinopathy analysis that integrates U-Net–based (ResNet-50 backbone) multi-class lesion segmentation with graph neural network–based classification. The proposed approach first performs precise segmentation of retinal lesions from fundus images and subsequently constructs lesion graphs that encode morphological attributes and spatial relationships. A graph neural network is then employed to learn high-level relational representations for DR grading. Experiments are conducted on the publicly available Messidor-2 dataset for lesion detection and classification, demonstrating the effectiveness and generalizability of the proposed method.

Literature Review

In recent years, automated analysis of diabetic retinopathy (DR) has advanced considerably, largely driven by the success of CNNs in retinal image analysis. CNN-based architectures have demonstrated strong performance in learning discriminative visual features from fundus images. However, despite these strengths, conventional CNN models are limited in their capacity to explicitly model the relationships among localized retinal lesions, which are clinically important for reliable DR assessment. To address this shortcoming, graph-based learning paradigms have been introduced, where lesions or image regions are represented as nodes, and their spatial or contextual dependencies are encoded as edges. In this direction, Zedadra et al. [2] explored multimodal graph-driven approaches for retinal image classification, while Gupta et. al. proposed a CNN-based multiclass framework for the detection of DR.

Bhaskar Marapelli et al. [3] presented a GNN-based framework specifically designed for DR recognition. Collectively, these studies highlight the potential of GNNs in capturing lesion-level dependencies; however, they predominantly focus on overall disease grading rather than explicit lesion-type differentiation. Lei et al. [4] proposed a hybrid framework that integrates GNNs with Capsule Networks (CapsNet) and multi-head prediction mechanisms, enabling richer hierarchical feature representations and improved DR grading accuracy. In parallel, semi-supervised and weakly supervised strategies have been explored to reduce the reliance on extensive lesion annotations. Bhaumik et al. [5] further demonstrated that graph neural networks (GNNs) can effectively distinguish different stages of DR by leveraging relational information.

Building upon this concept, a graph convolutional network (GCN) is proposed by Wang et al. [6], the framework adaptively integrates multiscale features for DR severity grading. Their proposed multiscale dynamic GCN combines hierarchical feature extraction with lesion-level connectivity, enabling joint modeling of local and global retinal patterns.

Hossain et al. [7] proposed GDRNet, an AI-driven diagnostic framework that leverages graph theory to effectively select features for diabetic retinopathy (DR) grading. The method employs a deep graph correlation network (DGCN) to extract discriminative features from color fundus images by modeling intra-class relationships. An iterative random forest algorithm is then used for feature selection and ranking, identifying the most informative features derived from the DGCN. This iterative process enhances classification robustness by refining feature representations and integrating multi-scale contextual information. Finally, an extreme gradient boosting (XGBoost) classifier trained on the optimized feature set is used to predict DR severity.

Abushawish et al. [8] utilized Grad-CAM visualizations to interpret the decision-making process of deep learning models for diabetic retinopathy detection, thereby improving model transparency and bridging advanced artificial intelligence techniques with practical clinical applicability.

More recent efforts have aimed to enhance representation learning by combining graph models with other deep learning paradigms. Yu Xie et al. [9] investigated semi-supervised GNNs for graph classification, which is particularly relevant for DR scenarios where labeled lesion data are scarce. Additionally, deep graph correlation networks have been developed to perform automated DR grading

without expert annotations by exploiting correlations among CNN-extracted features through pseudo-contrastive learning objectives.

To further overcome the limitations of CNNs in modeling lesion interdependencies, several studies have explicitly formulated DR analysis as a graph learning problem. In these approaches, lesions or retinal regions serve as graph nodes, while spatial, morphological, or contextual relationships define the edges. Feng et al. [10] and Sumod Sundar et al. [11] demonstrated that incorporating topological and morphological relationships via GNNs yields more informative structural representations than purely CNN-based pipelines. Furthermore, Yuan et al. [12] emphasized the importance of explainability in graph-based DR models. Vignesh et al. [13] present an automated framework for the identification and classification of diabetic retinopathy manifestations using a pre-trained convolutional graph neural network. Chandran et al. [14] proposed an automated GNN optimized using meta-heuristic search strategies to jointly grade DR and lesions such as diabetic macular edema (DME). Despite these notable advances, most existing graph-based DR methods primarily target global severity grading or combined lesion analysis, rather than fine-grained discrimination between individual lesion types such as microaneurysms and exudates, etc. This limitation motivates the proposed GBLR-Net framework, which integrates morphological, intensity, and texture attributes into lesion-centric graph representations. By enabling detailed characterization of individual lesion categories, GBLR-Net facilitates a more nuanced and clinically meaningful assessment of diabetic retinopathy beyond conventional severity-level classification. Jingbo Hu et al. [15] introduced graph adversarial transfer learning techniques to improve cross-dataset generalization.

Methodology

The proposed framework adopts a structured, step-wise pipeline for retinal fundus image classification. Initially, retinal fundus images are acquired from publicly available datasets and systematically organized for subsequent analysis. The collected images are then preprocessed and normalized to achieve consistent resolution and intensity distribution across samples. Following this, lesion regions are segmented using an EfficientNet–U-Net architecture, and connected-component analysis is then used to localize individual lesions and generate corresponding bounding boxes. From these lesion-specific regions, deep feature representations are extracted using a ResNet-50 backbone and subsequently organized into lesion-level graph structures based on spatial proximity. Finally, the constructed graphs are processed through stacked Graph Isomorphism Network (GIN) layers with global mean pooling to classify images into respective lesion categories. Figure 1 illustrates the complete workflow of the proposed methodology.

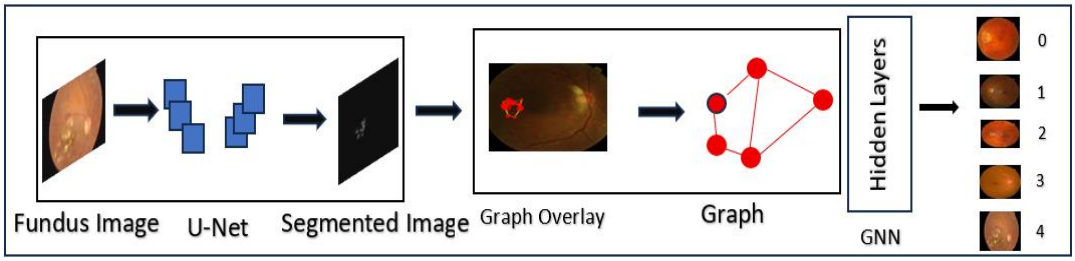

Fig 1. Generic Workflow of the Proposed Method

Dataset

This study uses the publicly available retinal fundus image dataset, Messidor-2, which is a widely used benchmark for diabetic retinopathy (DR) analysis. The Messidor-2 dataset [16] contains color fundus images captured under heterogeneous imaging conditions, with variations in resolution, illumination, and acquisition protocols that reflect real-world clinical settings. Messidor-2 comprises high-resolution images with spatial dimensions of 1440×960, 2240×1488, and 2304×1536 pixels, ensuring preservation of fine pathological details. For disease-level classification, five clinically relevant classes, Normal represented as ‘0’, Mild represented as ‘1’, Moderate represented as ‘2’, Severe represented as ‘3’, and PDR represented as ‘4’, these five classes were selected from the publicly available dataset to evaluate lesion-specific classification performance, with the class-wise distribution summarized in Table 1.

Table 1. DR Classes Distribution Based on the Number of Lesions.

Class	Number of Lesion Images	Grade
Normal	1017	0
Mild	347	1
Moderate	270	2
NPDR	75	3
PDR	35	4

Dataset Splitting

The dataset was split into three subsets using stratified sampling, with 70% for training, 10% for validation, and 20% for testing. This dataset's percentage distribution or splitting ensured that each subset contained a balanced representation of DR and non-DR retinopathy images. To learn the model parameters, a training set was used, while a validation sample was used to support hyperparameter tuning and prevent overfitting. The test set was reserved for objective evaluation of the model's generalization performance. Maintaining class balance across all splits mitigates the impact of data imbalance, promotes stable learning, and enables reliable and consistent performance assessment on previously unseen samples.

Preprocessing

The preprocessing pipeline begins with normalizing all fundus images to a fixed size and resolution, with pixel intensity values scaled to the range of zero to one to ensure consistency across samples. Subsequently, 2048-dimensional feature vectors are extracted from each candidate lesion or image patch using a pre-trained U-Net network, which serve as node embeddings in the graph representation. Lesion masks corresponding to lesion classes are obtained via binary segmentation and used solely for visualization, including segmented regions, bounding boxes, and graph overlays. Finally, connected component analysis is applied to treat each lesion as an independent component, enabling node-level representation within the constructed graph.

Segmentation

Image segmentation is performed using an EfficientNet-based U-Net architecture to accurately delineate lesions present in fundus images. The proposed framework employs a U-shaped encoder-decoder structure with residual connections to effectively capture subtle lesion features. This segmentation approach has been validated and then applied to fundus images for classification.

Lesion Detection and Bounding Box

Lesion detection is performed by identifying connected components in the U-Net-predicted segmentation masks, and bounding boxes for lesions are generated by enclosing each lesion region within its minimum axis-aligned rectangular boundary.

Graph Construction

Each retinal image is transformed into a graph representation, with lesion regions identified by the U-Net segmentation model treated as nodes. For each segmented lesion region, a 2048-dimensional deep feature embedding is extracted using a ResNet50 backbone and assigned as the node feature. Edges between nodes are established based on spatial proximity using k-nearest neighbour connectivity, enabling information exchange among nearby or clinically related lesions. This graph-based representation explicitly models inter-lesion relationships, which are critical for accurate analysis of diabetic retinopathy. Algorithm 1 describes the graph construction process and how these graphs are fed into the model.

Algorithm 1: Graph Construction and Feeding into the Model

Input: Retinal fundus images I, segmentation masks

Stage 1: Preprocessing and Lesion Segmentation

- Initialize image set $I \leftarrow \{I1, I2, \ldots, IN\}$.

for each image $Ii \in I$ do

Resize Ii to a fixed resolution and normalize pixel intensities.

Segment lesion regions using a U-Net model.

-Obtain lesion masks Li = {X}.

end for

Stage 2: Feature Extraction

- Initialize feature set F ← [ ].

for each lesion region rj ∈ Li do

Extract morphological features: area, perimeter, eccentricity.

Extract intensity-based and texture-based features.

Form lesion feature vector fj. Append fj to F.

end for

- Normalize all feature vectors in F.

Stage 3: Graph Construction

- Define node set V = {fj| fj ∈ F}. Each node represents a lesion

- Compute spatial coordinates pj for each lesion region.

- Construct edge set E using k-Nearest Neighbors (kNN) on coordinates pj

- Form the lesion graph G = (V, E).

Stage 4: GNN-Based Classification

- Apply first graph convolutional layer: H1 = ReLU(GCN1(V, E)).

- Apply second graph convolutional layer: H2 = ReLU(GCN2(H1, E)).

- Perform global mean pooling: z = GlobalMeanPool(H2).

- Apply dropout regularization: z′ = Dropout(z).

- Compute class probabilities using Softmax: ˆy = Softmax(Wz′ + b).

Output: Predicted class label ˆy ∈ {Normal=0, Mild=1, Moderate=2, Severe=3,PDR=4}

Layered Architecture

This proposed graph-based classification model comprises seven sequential layers that progressively extract structural representations from high-dimensional node features and perform final disease grading. Initially, the network's input consists of node-level feature vectors of dimension 2048. These features are processed by a Graph Isomorphism Network (GIN) convolutional layer, which reduces the input representation from 2048 to 128 dimensions.

This layer performs neighborhood aggregation and updates node embeddings using a learnable multilayer perceptron, enabling the model to capture discriminative structural patterns within the graph. A Rectified Linear Unit (ReLU) activation function follows the first GIN layer to introduce nonlinearity and enhance the network's expressive capacity.

Subsequently, a second GIN convolution layer further refines the node embeddings while maintaining the 128-dimensional feature space. This additional graph convolution allows the network to capture higher-order neighborhood interactions and more abstract structural dependencies. Another Rectified Linear Unit (ReLU) activation layer is applied to stabilize learning and prevent vanishing gradients.

After node-level feature extraction, a Global Mean Pooling layer is employed to aggregate all node embeddings into a single graph-level representation. This operation converts the node-wise feature matrix of shape (N, 128) into a fixed-dimensional vector of size (1, 128), enabling graph-level classification regardless of the number of nodes.

To reduce overfitting and improve generalization performance, a Dropout layer is applied to the pooled feature vector during training.

Finally, a Fully Connected layer maps the 128-dimensional graph representation to a 5-dimensional output vector corresponding to the five disease severity classes. This layer performs the final classification by generating class logits used for prediction.

Overall, the architecture progressively transforms high-dimensional node features into compact graph-level embeddings through hierarchical neighborhood aggregation and pooling, followed by a linear decision layer for multi-class classification.

Layered Architecture Description: The proposed graph-based classification model operates on node-level retinal representations and comprises graph convolutional layers, followed by graph-level aggregation and classification.

A) Input Layer (Graph Node Features): The model receives graph-structured input where each node is represented by a 25-dimensional feature vector derived from 5X5 image patches. These features encode localized structural information extracted from retinal regions. The input tensor has shape [N, 25], where N is the number of nodes.

B) First Graph Convolution Layer (GCNConv 25 → 64): The first graph convolution layer transforms node features from 25 dimensions to 64 dimensions. This operation performs neighborhood aggregation by combining each node's features with those of its connected neighbors. The resulting representation captures both local patch information and relational context within the graph. This layer contains 1,664 learnable parameters.

C) Batch Normalization and Activation: Batch normalization stabilizes training and improves convergence by normalizing the distributions of intermediate features. The ReLU activation function introduces nonlinearity, thereby enhancing the network's representational capacity. The feature dimension remains [N, 64], with 128 learnable parameters.

D) Second Graph Convolution Layer (GCNConv 64 → 64): A second graph convolution layer further refines the node embeddings while maintaining the 64-dimensional feature space. This layer enables deeper relational learning by capturing higher-order neighborhood interactions. It contains 4,60 parameters and enhances the propagation of structural features across the graph.

E) Global Mean Pooling: To perform graph-level classification, node embeddings are aggregated using global mean pooling. This operation computes the average of all node features, producing a fixed-size representation of shape [B, 64], where B denotes the batch size. This step ensures the model can handle graphs with varying numbers of nodes.

F) Fully Connected Layer (64 → 2): The pooled graph representation is passed through a linear classification layer that maps the 64-dimensional feature vector to a 2-dimensional output corresponding to the target classes. This layer contains 130 trainable parameters.

G) Output Layer (LogSoftmax): Finally, a LogSoftmax activation function is applied to produce normalized log-probabilities for each class, enabling probabilistic interpretation and compatibility with negative log-likelihood loss during training.

Layer	Type	Input Shape	Output Shape	Parameters	Description
Input	Graph Node Features	[N, 25]	[N, 25]	0	25D node features (5×5 patches)
GCN Conv 1	GCNConv	[N, 25]	[N, 64]	1664	Graph convolution: 25 → 64
BatchNorm 1	BatchNorm1d	[N, 64]	[N, 64]	128	Batch normalization + ReLU activation
GCN Conv 2	GCNConv	[N, 64]	[N, 64]	4160	Graph convolution: 64 → 64
Global Pool	GlobalMeanPool	[N, 64]	[B, 64]	0	Graph-level representation (mean pooling)
Linear	Linear	[B, 64]	[B, 2]	130	Classification layer: 64 → 2
Output	LogSoftmax	[B, 2]	[B, 2]	0	Log-softmax activation

Figure 2 Represents A Detailed Layered Architecture.

Fig 2. Layered Architecture of the model

Hyperparameter

The proposed model was implemented using PyTorch Geometric, in which deep feature embeddings from ResNet50 were converted into graph-structured representations. The architectural configuration and associated training hyperparameters are presented in Table 2. These parameters were selected empirically to achieve stable convergence and consistent performance across all evaluation metrics.

Table 2. List Of Hyperparameters Used to Train the Proposed Model.

Hyperparameter	Value
Model Name	GBLR-Net
Optimizer	AdamW
Learning Rate	9.765625e-07
Batch Size	32
Weight Decay	0.0001
Epochs	100
Dropout Rate	0.5
K (Graph Construction)	4
Number of Layers	5

Experimental Results

This section presents the experimental results obtained using the methodology described in the methodology section. The outcomes are organized to provide a clear visualization of lesion detection and the corresponding graph-based representations.

Implementation Platform

The GBLR-Net model was implemented using Python with the PyTorch framework. Graph operations were performed using PyTorch Geometric, with graphs constructed using K = 5 nearest neighbors. The model was trained using the optimizer (AdamW) with having learning rate of 9.77×10^{-7}, decay (weight) of 0.0001, and batch size of 32 over 100 epochs. The dropout rate of 0.5 was applied, and the network consisted of 4 layers. This experiment was conducted on the Google Colab platform. The runtime and computational configuration were on a T4 Graphics Processing Unit (GPU).

Figure 3 displays qualitative results, including the original fundus image, segmented lesion map, detected bounding boxes, and the associated graph overlay. In addition, this section outlines the implementation environment and evaluation procedures. Performance assessment includes confusion matrix analysis, classification reports, and feature space.

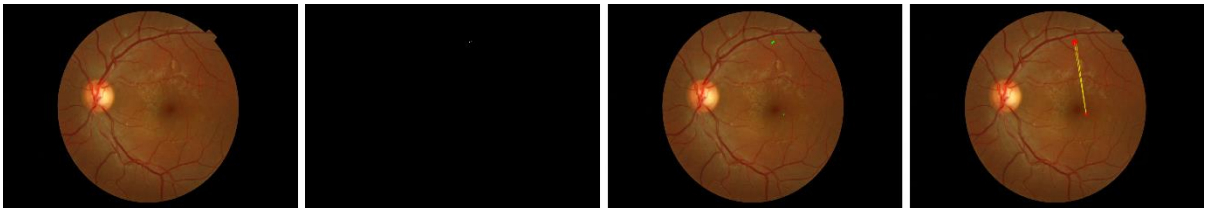

Fig (a) Fig (b) Fig (c) Fig (d)

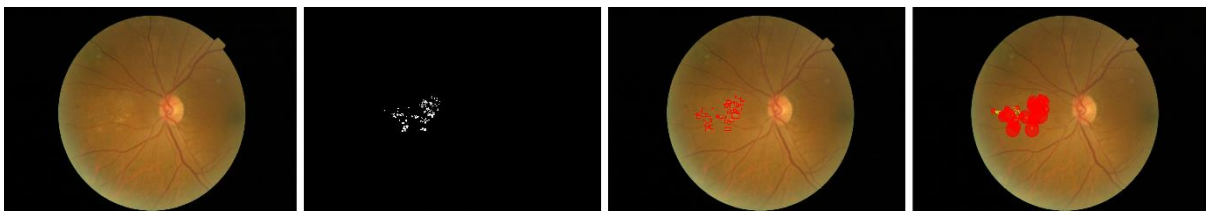

Fig (e) Fig (f) Fig (g) Fig (h)

shows moderate lesions with corresponding bounding boxes and graph overlays.

Classification Matrix

To evaluate the proposed model for early lesion classification (Normal, Mild, and Moderate), we used confusion matrices and classification reports on the Messidor-2 dataset, which consists of 1,017

Normal, 270 Mild, 347 Moderate, 75 Severe, and 35 PDR fundus images. The confusion matrices were generated by comparing the true labels of test samples with the model's predictions. The values along the diagonal correspond to correctly predicted instances, whereas the off-diagonal values denote incorrect classifications. As shown in Figure 4, the matrices were constructed for five classes: Normal, Mild, Moderate, Severe, and PDR. This visualization helps assess the model's strengths and weaknesses.

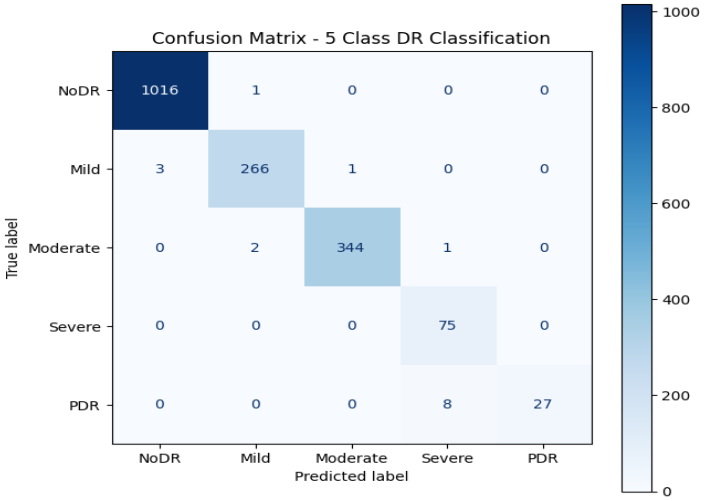

Fig 4. Confusion Matrix of 5 Classes of Messidor-2 Dataset

Model Performance Analysis

To understand the model's performance, we evaluate the classification metrics, namely recall, precision, and F1-score, across the Messidor-2 dataset. These metrics improve steadily and stabilize as the number of training epochs increases, indicating consistent learning and model convergence. Precision quantifies this model's ability to correctly identify positive samples by reducing false positives. Whereas the recall measures the model's performance in identifying true positives, thereby reducing the number of false negatives. Here, the F1 score, which is defined as the precision (harmonic mean) and recall, provides a comprehensive assessment by balancing both error types within a single metric. The sustained high F1-score observed during training suggests that the model maintains a favorable balance between sensitivity and specificity. This balance is particularly crucial in medical image analysis, where both missed detections and incorrect positive predictions can have significant clinical implications. Furthermore, the convergence of these performance indicators toward near-optimal values demonstrates the strong capability of the proposed approach. The model achieves consistent, reliable results across the Messidor-2 benchmark dataset. Figure 5 shows the model's performance analysis.

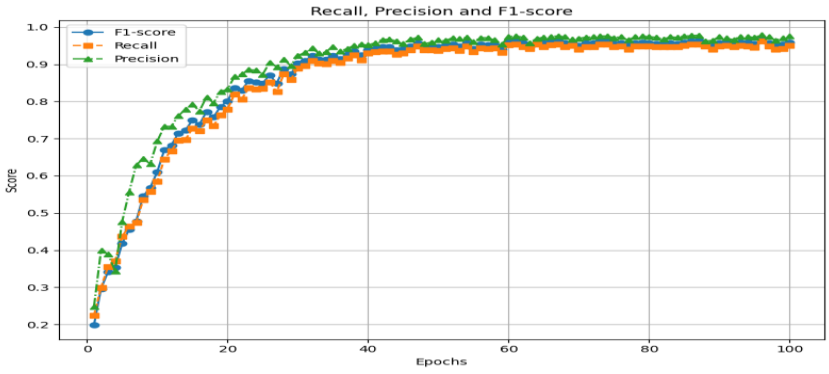

Fig5.Performance Analysis of Messidor-2

Stochastic Neighbor Embedding (t-SNE)

This study examines the quality of the learned feature representations. High-dimensional feature representations were mapped to a two-dimensional space for visualization. The projected distribution shows well-defined clusters for the five categories: Normal, Mild, Moderate, Severe, and PDR. The Normal-to-PDR categories form relatively compact and well-separated clusters, indicating strong

discriminative learning for healthy and advanced-stage cases. In contrast, some overlap is observed among the intermediate stages, reflecting the gradual progression and subtle visual differences between adjacent severity levels. Overall, the t-SNE visualization (Figure 6) confirms that the model learns structured and severity-aware feature representations, supporting its effectiveness in multi-class retinal disease classification.

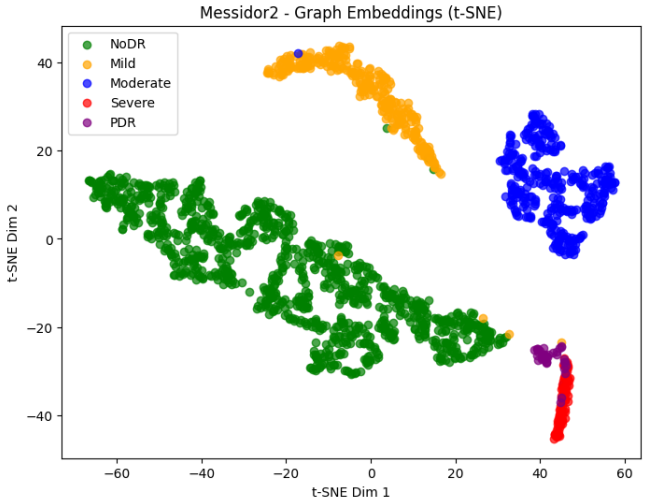

Fig 6. t-SNE Distribution of 5 Classes

Ablation Study

The ablation study was conducted to check how this model performs at different stages, and we evaluated the contribution of each component in our framework. Different versions of the model were tested by removing or modifying key components, including the lesion-level graph representation, the U-Net segmentation stage, and the graph neural network refinement. Results show that removing the graph-based representation significantly reduces classification performance, highlighting the importance of modeling inter-lesion relationships. Similarly, omitting the segmentation stage or using only raw images results in lower accuracy and F1 Scores, confirming that both precise lesion segmentation and relational feature learning are critical for optimal diabetic retinopathy assessment. The ablation study results on the Messidor-2 dataset are shown in Table 3.

Table 3. Ablation Study Based on a Model.

Model Variant	Accuracy (%)	Precision (%)	Recall (%)	F1Score (%)
No ResNet	92.20	87.30	90.71	90.28
No Dropout	93.45	90.56	92.39	91.81
One GIN Layer	90.21	91.39	91.36	90.17
Model (Macro Avg.)	99.08	99.40	99.20	99.10

Comparative Analysis based on the State-of-the-Art

In addition to evaluating preprocessing effects, we compared our GCN-enhanced framework with other state-of-the-art methods used for DR classification. Performance metrics such as precision, recall, F1-score, and accuracy were used to benchmark our model against recent convolutional and hybrid approaches, demonstrating its superior effectiveness and robustness across multiple datasets. The comparative analysis is shown in Table 4.

Table 4. Comparative Analysis Based on the State-Of-The-Art

Author	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
A. Gupta et al. [1]	-	-	-	81.80
Amina Zedadra [2]	96.00	96.50	96.80	95.80
Bhaskar Marapelli et al. [3]	94.00	95.00	93.00	94.00

Yongjia Lei et al. [4]	99.41	79.14	71.01	74.05
Yipeng Wang et al. [6]	88.70	72.27	67.84	69.69
Meiling Feng et al. [10]	78.00	77.00	98.00	84.90
J. Gnana Chandran et al. [14]	99.57	-	-	-
Guanghua Zhang et al. [17]	89.90	-	-	-
M. Kolla et al. [18]	91.00	-	-	-
Our Model for Messidor2	99.08	99.40	99.20	99.10

Discussion

The proposed GBLR-Net framework integrates ResNet50, a U-Net backbone for feature extraction, and GNN reasoning to classify fundus images into five severity categories: Normal, Mild, Moderate, NPDR, and PDR. By transforming high-dimensional convolutional embeddings into graph representations, the model captures both localized lesion characteristics and their spatial dependencies. This combination of CNN-based feature learning and GNN-driven relational modeling strengthens discrimination across progressive stages of diabetic retinopathy (DR).

Quantitative evaluation using the confusion matrix demonstrates clear separation between normal and pathological cases, indicating reliable identification of healthy samples. However, minor misclassifications were observed between adjacent severity levels, particularly between Mild, Moderate, and NPDR. This reflects the inherent clinical challenge of distinguishing visually similar and progressively evolving retinal abnormalities. The classification report further supports these findings, showing strong overall performance with relatively balanced precision, recall, and F1-scores across the five classes.

To further assess representational quality, hyperparameter analysis and t-SNE visualization were conducted. The embedding projections reveal well-defined clustering from the Normal and PDR classes, while partial overlap is observed among intermediate severity levels. These observations suggest that GBLR-Net learns meaningful and clinically relevant latent features, though fine-grained stage differentiation remains challenging.

Despite promising results, the evaluation was conducted on datasets with limited size and diversity, which may constrain generalization across varying imaging conditions and populations. Additionally, KNN-based graph construction may not fully capture global retinal context or long-range lesion interactions. Future improvements may focus on more adaptive graph modeling and multi-scale feature representations to enhance severity-level discrimination.

Conclusion and Future Work

This study proposes a GBLR-Net, a lesion-aware graph neural network for automated retinal abnormality classification, targeting five severity levels: Normal, Mild, Moderate, Severe, and PDR. The framework models lesions as graph nodes by encoding morphological, intensity, and texture features, allowing the network to capture both lesion-specific characteristics and their spatial relationships within retinal fundus images.

A K-nearest neighbor (KNN) strategy is adopted for graph construction, followed by a two-layer Graph Convolutional Network (GCN) to learn discriminative relational representations. Global mean pooling aggregates node-level embeddings into a single graph-level feature vector, which is subsequently used for five-class severity classification. Experimental evaluation demonstrates that GBLR-Net achieves strong generalization performance, yielding the best accuracy, precision, recall, and F1 score across all severity categories.

Beyond quantitative performance, the proposed framework improves interpretability by generating clinically meaningful visual outputs, including lesion segmentation overlays, severity-aware bounding boxes, and lesion graph visualizations. These components enhance transparency and support reliable clinical decision-making.

Overall, the findings confirm the effectiveness of graph-based learning for multi-class retinal disease grading and establish GBLR-Net as a robust and interpretable tool for early diabetic retinopathy screening.

Future work will investigate adaptive and attention-driven graph construction techniques and explore advanced GNN architectures, such as graph attention and transformer-based models, to further enhance fine-grained severity discrimination. Validation on larger and more diverse datasets, integration of multimodal clinical information, and lightweight optimization for real-time deployment will further strengthen its clinical applicability.

References

A. Gupta, P. Verma, and R. Srivastava, “Diabetic retinopathy detection through multiclass classification of fundus image using a convolutional neural network,” in Proceedings of the International Conference on Computer Vision and Image Processing (CVIP), pp. 210–219, Springer, 2025.

[2] Amina Zedadra, Ouarda Zedadra; Mahmoud Yassine Salah-Salah; Antonio Guerrieri, ‘‘Graph-Aware Multimodal Deep Learning for Classification of Diabetic Retinopathy Images’’, Journals & Magazines, IEEE Access, Volume: 13, 2025.

[3] Bhaskar Marapelli, D. Baburao2, Vedavyas Gurla3, Samatha Konda4, Ch.Anil Carie, Gandla Shivakanth, ‘‘Convolutional Neural Networks for Diabetic Retinopathy Detection in Retinal Images’’, International Journal of Intelligent Systems and Applications in Engineering, ISSN:2147-6799, 2024.

[4] Yongjia Lei, Shuyuan Lin, Zhiying Li, Yachao Zhang, Taotao Lai,‘‘GNNfused CapsNet with multi-head prediction for diabetic retinopathy grading’’, Engineering Applications of Artificial Intelligence 133, 107994, Elsevier 2024.

[5] Bhaumik Maan, Ashiwini T. Sapkal, Vaishali Ingale, Piyush Yadav, Abhinab Pratap Singh Chauhan, Rahul Lamba, ‘‘Classification of Diabetic Retinopathy Disease Using GNN’’, International Conference on Brain Computer Interface & Healthcare Technologies (iCon-BCIHT), 2024.

[6] Yipeng Wang, Liejun Wang, Zhiqing Guo, Shiji Song & Yanhong Li, ‘‘A graph convolutional network with dynamic weight fusion of multi-scale local features for diabetic retinopathy grading’’, Scientific Reports, Volume 14, Article number: 5791, 2024.

[7] Shahed Hossain, Md. Zahid Hasan, Risul Islam Jim, Taslima Ferdaus Shuva, Md. Tanvir Rahman; Abdullah Al-Mamun Bulbul, ‘‘GDRNet: A Novel Graph Neural Network Architecture for Diabetic Retinopathy Detection’’, IEEE International Conference on Data Mining Workshops (ICDM Workshops), 2024.

[8] I. Y. Abushawish, S. Modak, E. Abdel-Raheem, S. A. Mahmoud, and A. J. Hussain, ‘‘Deep learning in automatic diabetic retinopathy detection and grading systems: A comprehensive survey and comparison of methods," IEEE Access, Vol. 12, pp. 84785–84802, 2024.

[9] Yu Xie, Yanfeng Liang, Maoguo Gong, A. K. Qin, Yew-Soon Ong, and Tiantian He, ‘‘Semisupervised Graph Neural Networks for Graph Classification’’, IEEE Transactions on Cybernetics, Volume 53, No. 10, 2023.

[10] Meiling Feng, Jingyi Wang, Kai Wen, and Jing Sun, ‘‘Grading of Diabetic Retinopathy Images Based on Graph Neural Network’’, Journals & Magazines, IEEE Access, Volume: 11, 2023.

[11] Sumod Sundar, AND S. Sumathy, ‘‘Classification of Diabetic Retinopathy Disease Levels by Extracting Topological Features Using Graph Neural Networks’’ ,Journals & Magazines, IEEE Access, Volume: 11, 2023.

[12] Hao Yuan, Haiyang Yu, Shurui Gui, and Shuiwang Ji, ‘‘Explainability in Graph Neural Networks: A Taxonomic Survey’’, IEEE Transactions on Pattern Analysis and Machine Intelligence, Volume 45, NO. 5, MAY 2023.

[13] R. Vignesh, N. Muthukumaran, and M. P. Austin, “Detection of diabetic retinopathy image analysis using a convolutional graph neural network”, International Conference on Inventive Computation Technologies (ICICT), pp. 921–929, IEEE, 2023.

[14] J. Jasper Gnana Chandran, J. Jabez, Senduru Srinivasulu,‘‘Auto-Metric Graph Neural Network optimized with Capuchin search optimization algorithm for coinciding diabetic retinopathy and diabetic Macular edema grading”, Biomedical Signal Processing and Control, Springer 2023.

[15] Jingbo Hu, Huan Wang, Le Wang, and Ye Lu, ‘‘Graph Adversarial Transfer Learning for Diabetic Retinopathy Classification’’, Journals & Magazines, IEEE Access, Volume: 10, 2022.

[16] Decencière, E., Zhang, X., Cazuguel, G., Lay, B., Cochener, B., Trone, C., Gain, P., Ordóñez, R., Massin, P., Erginay, A., Charton, B., & Klein, J. C. (2014). Feedback on a publicly distributed image database: The Messidor database. Image Analysis & Stereology, 33(3), 231–234

[17] Guanghua Zhang, Guanghua Zhang, Bin Sun, Bin Sun, Zhixian ChenZhixian Chen, Yuxi GaoYuxi Gao, Zhaoxia Zhang, Keran Li, Keran Li, Weihua Yang Weihua Yang,‘‘Diabetic Retinopathy Grading by Deep Graph Correlation Network on Retinal Images Without Manual Annotations’’,Volume 9, 2022.

[18] M. Kolla and T. Venugopal, ‘‘Efficient classification of diabetic retinopathy using binary cnn," in 2021 International conference on computational intelligence and knowledge economy (ICCIKE). IEEE, pp. 244–247. 2021.