

Smart Diagnosis of Cauliflower Diseases Using Deep Learning and Feature Optimization

Meenalochini M.¹, Amudha P.²

¹Research Scholar, Department of Computer Science and Engineering, School of Engineering, Avinashilingam Institute for Home Science and Higher Education for Women, Coimbatore, India.

²Professor, Department of Computer Science and Engineering, School of Engineering, Avinashilingam Institute for Home Science and Higher Education for Women, Coimbatore, India.

E-mail: ¹meenalochnicse@gmail.com, ²amudha_cse@avinuty.ac.in

Abstract

Cauliflower (*Brassica oleracea* var. *botrytis*) is one of the most popular crops that are subject to a variety of diseases affecting the leaf apparatus, which impact quality and production. Despite the progress in deep learning, appropriate disease detection under real-field conditions remains a serious problem. This paper introduces an expert system GNN-PDP, which is a novel Graph Neural Network based model for the automated classification of cauliflower leaf diseases using images taken with a smartphone. A Region Growing Segmentation (RGS) is used to extract perceptual regions in this structure and statistical features are utilized as graph node features. The Salp Swarm Algorithm (SSA) finds optimal features that result in better generalization. A total of 750 images were gathered in four categories of diseases. The assessment was made based on accuracy, precision, sensitivity, specificity, and F1-score, in relation to Linear Discriminant Analysis (LDA), Random Forest (RF), Deep Neural Networks (DNN), and CNN classifiers. GNN-PDP achieved a superior classification accuracy of 89.0%, outperforming all other experiments. The model has great potential for smart agriculture in disease management.

Keywords: Cauliflower Disease Detection, Deep Learning, Graph Neural Networks, Feature Optimization, Smart Agriculture.

1. Introduction

Agriculture is an essential field in the global economy as it is the core of food supply, revenues, and rural growth. The agri-food supply chain (AFSC) encompasses the stages of production, harvesting, storage, processing and distribution whereby each stage is essential in preserving food security and stabilizing the market. The issue of AFSC management has attracted more concern over the past few decades following the increasing demand, variability in climate and issues with sustainability [1]. Cauliflower is a high-nutritional cruciferous crop, which is widely grown in different environments across the entire world. This is because many diets consist of it due to its rich content of fiber, vitamins, carotenoids, glucosinolates and phenolic compounds that are related to anticancer and cardio protective properties [2]. Nations such as India (approximately 7.8 million tons per year) and Bangladesh (73,000 tons on 9,400 acres) depend largely on cauliflower both nutritionally and economically as well as for their

agricultural efficiency [3]. Regardless of its worth, this product is immensely vulnerable to various destructive afflictions including bacterial spot rot, black leg, white rust, downy mildew, black rot, powdery mildew, and ring spot diseases. Without early detection, these diseases can result in the loss of up to 30-50 percent and threaten the food security of the region as well as the revenue of farmers [4]. Factors such as unseasonable weather conditions, inappropriate agronomic treatment, and a lack of diagnostic means contribute to the increased transmission and worsen the progress of these infections. Traditional disease detection techniques are labour-intensive, usually require expert knowledge, and have low efficiency and ergonomics. These constraints do not allow for timely action and the consequence of this is usually massive crop loss. Recent developments in machine vision and deep-learning mechanisms have introduced automation tools that can facilitate the diagnosis of plant diseases using leaf imagery. Convolutional Neural Networks (CNNs) have already been widely applied to the task of disease classification based on images [5], but these networks do not model the spatial relationships among regions in an image. Early and accurate classification of cauliflower leaf diseases in the real remains a key parameter due to a combination of visual indicators, inconsistent light, and the drawbacks of traditionally applied machine learning and deep learning models that overlook spatial correlations in leaves across imagery. Current techniques like CNNs and K-means-based classifiers cannot easily handle segmentation accuracy, class imbalance, and generic outdoor conditions in noisy data. This urgently requires a resilient, end-to-end solution that combines spatial reasoning, adaptive feature optimization and meaningful segmentation that are perceptual to enhance the performance of disease diagnosis in real-world agricultural contexts. The contributions are:

- Creation of a GNN-based model that can learn relationships between the regions of the leaves spatially by using segmented patches as the nodes in a graph and exploit the graph structure to learn the context of the patterns associated with diseases.
- RGS--designed integration with perceptually consistent segmentation of the diseased and healthy parts that proves to be more efficient than traditional K-means clustering involving structural preservation integrity.
- Implements the SSA to generate the optimal feature set to control high-dimensional feature noise and enhance the generalization of the models with reference to the selection of the most discriminative color and texture features.
- Real-field mobile-captured datasets of images taken, making the project more relevant and deployable in the wild with regard to uncontrollable agricultural conditions relevant to the proposed model.
- A full performance-based analysis of GNN-PDP versus baseline classifiers that exhibit higher performance across all disease classes.

2. Literature Review

Numerous studies have been conducted on computer vision and machine learning methods for the automatic identification of plant diseases in high-demand crops, such as cauliflower. This winter crop, crucial to Bangladesh due to its high nutritional value and economic importance, is susceptible to diseases like black rot and downy mildew, which significantly affect crop productivity. Early detection is critical for managing these diseases effectively [6]. A study surveyed control methods for black rot in Brassicaceae vegetables

caused by *Xanthomonas campestris*, categorizing them as physical (e.g., hot water seed treatment), chemical (e.g., bactericides), and biological (e.g., beneficial microbes), with a focus on integrated management [7]. Similarly, the environmental and health consequences of pesticide overuse have been addressed, advocating for eco-friendly bioremediation methods through bacteria, fungi, plants, and microalgae [8].

A study introduced "Cauli-Det," an enhanced cauliflower disease detection model using a modified YOLOv8 approach [9]. Another work utilized an agro-medical expert system based on k-means segmentation and 10 statistical/texture features for cauliflower disease diagnosis, showing improvements in feature selection and augmentation [10]. CNN models and transfer learning were used to classify four cauliflower diseases (e.g., black rot, white rust), obtaining the highest accuracy with the InceptionV3 algorithm [11]. A comparison between traditional machine learning (ML) and deep learning (DL) models for cauliflower disease recognition showed the superiority of DL models like InceptionV3, ResNet50, MobileNetV2, and VGG16 [12]. Crop-conditional CNNs were proposed to incorporate image metadata and improve classification performance, tackling sampling bias [13]. Bayesian deep learning was used to quantify uncertainty in plant disease classification, providing more accurate predictions [14]. Furthermore, improved crop disease classification was achieved using median filters, FCM clustering, and LSTM models [15]. While several models have applied ML and DL for disease detection, most fail to capture the spatial relationships between infected areas, leading to limited classification performance. Popular segmentation techniques like k-means do not preserve structural and perceptual integrity under real-world conditions. Few studies have employed advanced metaheuristic algorithms like the Salp Swarm Algorithm (SSA) for feature selection [16]. Additionally, most previous research used controlled datasets, whereas real-field environments introduce variability and noise that these methods struggle to handle. Therefore, an integrated approach combining spatial modeling, perceptual segmentation, feature optimization, and real-field data is essential to address these challenges, as explored in this paper.

3. Methodology

In the proposed study, GNN-PDP is created, a GNN-based framework capable of detecting and classifying four key diseases of cauliflower using a total of 750 mobile-captured images. The methodology also considers RGS as the technique to create regions which are considered to be significant visually, GNNs as the tool to describe spatial and relational relations, and SSA as an approach to streamline the choice of features. As described in Figure 1, the whole pipeline stream consists of image reception, preprocessing, segmentation, feature extraction, graph-based classification and enhancement, and optimization whose layout takes into consideration the issues presented by real-field data.




Figure 1. Overview of the Methodology

3.1 Image Acquisition

As described in Table 1, the white cauliflower field images (750 images) were captured during a period of 10 days by utilizing a smart phone camera in natural light settings in Hosur, Krishnagiri District, Tamil Nadu, India. The data of the healthy and diseased plants with the four target diseases is represented in the dataset. It includes 750 images of benign and pathological or sick living plants of the white cauliflower.

Table 1. Summary of the Cauliflower Disease Dataset

Attribute	Details
Vegetable	Cauliflower (white variety)
Location	Hosur, Krishnagiri District, Tamil Nadu, India
Type of the Field	Square-shaped flat land
Size of the Area	5 acres
Acquisition Duration	10 days
Collection Method	Manual capture via smartphone camera
Image Resolution	500 × 500 pixels
Classes	Bacterial spot rot, black rot, downy mildew, and healthy.
Total Images	750
Training Set	70% (525 images)
Testing Set	30% (225 images)

In order to enhance the generalization of the images and computational efficiency, the whole dataset was split into 70% training (525 images) and 30% testing (225 images). The 750 images were used to simplify the training and evaluation of the models. Examples of the acquired samples are given in Figure 2.



Figure 2. Sample Plant Pictures from the Collected Data

3.2 Pre-processing

Raw images are subjected to different light, shades and contrasts. A set of preprocessing operations was performed to improve the quality of the images and standardize the input for successive models.

3.2.1 Image Resizing

All the images obtained were rescaled to 300 pixels in height and width to ensure uniformity in the input size for the rest of the models and to reduce the computational burden. The presented resolution aims to retain the essential disease properties while balancing the processing costs. The intensity value of the $((i, j))$ pixel is denoted as (i, j) . It makes resizing possible using a bi-cubic or a bilinear interpolation-function depends on the intended scale. The approximation it shows reveals that the intensity function on a 2D surface can be approximated as:

$$f(i, j) = \sum_{x=0}^N \sum_{y=0}^N m_{xy} \times i^x \times j^y \quad (1)$$

Here, m_{xy} is the scaling factor and 2D pixels are denoted as i^x and j^y . This is done to preserve the spatial coherence, and it does not remove some vital texture features that are involved in downstream classification.

3.2.2 Histogram Equalization

The practical improvement of histogram equalization and RGB $\rightarrow$ L*a*b* color transformation was validated experimentally. Compared with RGB and HSV color models, L*a*b* achieved 7.8% higher segmentation accuracy under non-uniform lighting due to its perceptual uniformity and illumination invariance. This transformation reduced mis-segmentation in shaded regions, ensuring robust disease feature extraction in real-field conditions. The contrast of the picture is enhanced via higher traditionalization of the histogram. Suppose that S and T are respectively the row number (size) and the column (length) in pixels. S_l is the saturation of the color of the pixel number n_l , and M is the numerically calculated brightness level of the picture. Then with the help of continuity formula a mapping process between brightness S_l and every pixel brightness T_l is done as follows;

$$T_l = F(S_l) = \frac{M-1}{ST} (n_1 + n_2 + \dots + n_L) \quad (2)$$

Histogram equalization provides a mathematically grounded, computationally efficient, and practically effective means of preparing real-field cauliflower images for automated disease classification. By increasing the saliency of visual patterns associated with disease symptoms, it serves as a critical component in ensuring high accuracy in subsequent segmentation and feature extraction stages. The practical improvement of histogram

equalization and RGB→L*x*y color transformation was validated experimentally. Compared with RGB and HSV color models, L*x*y achieved 7.8% higher segmentation accuracy under non-uniform lighting due to its perceptual uniformity and illumination invariance. This transformation reduced mis-segmentation in shaded regions, ensuring robust disease feature extraction in real-field conditions.

3.2.3 RGB to L*x*y Transformation

In this study, the RGB-to-L*x*y transformation is applied as a critical preprocessing step to support effective image segmentation. RGB images captured under real-world field conditions often suffer from lighting inconsistencies and contrast variations, making them less reliable for pixel-based clustering. To enhance color uniformity and improve segmentation precision, RGB images are converted into L*x*y, a perceptually uniform and illuminant-independent color space. Unlike RGB, the L*x*y model preserves all color information while aligning pixel distances with perceptual differences. This improves the separability of healthy and diseased regions, particularly under uneven lighting, and enhances the performance of RGS. The transformation proceeds in two stages. First, the RGB values are linearly mapped to the CIE XYZ space:

$$\begin{bmatrix} i \\ j \\ k \end{bmatrix} = \begin{bmatrix} m_{11} & m_{12} & m_{13} \\ m_{21} & m_{22} & m_{23} \\ m_{31} & m_{32} & m_{33} \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} \quad (3)$$

Where, R,G and B are the red, green, and blue color values, and m_{ij} are matrix coefficients defined by the CIE standard. Next, lightness (L^*) and chromaticity coordinates (x , y) are derived using a non-linear mapping function:

$$f(x) = \begin{cases} x^{-3} & \text{if } x > 0.02 \\ 2\pi + 0.13 & \text{else} \end{cases} \quad (4)$$

This conversion improves RGS by aligning distance metrics with perceptual differences, enabling robust segmentation of cauliflower disease regions even under non-uniform lighting. This is denoted m_{ij} . It may utilize the tri-stimulus values to modulate the L*x*y color region. S_b , T_b , and R_b are the reference whites for those values in this scenario. The expression is given by Equation (6).

3.3 Image Segmentation Process

The similarity threshold (τ) in RGS was empirically tuned by analyzing intra-class variance and inter-class separation. Optimal performance was achieved at $\tau = 0.82$, balancing over-segmentation and region merging errors. Sensitivity analysis confirmed minimal accuracy deviation ($\pm 1.5\%$) for τ within 0.8–0.85. Post preprocessing, each picture is subjected to a segmentation process that isolates small images that appear intuitively significant in distinguishing diseased and healthy parts. The given study uses RGS because it maximally preserves structural contours under real-field settings. In contrast to segmenting procedures via K-means or thresholding, RGS flexibly seizes patches in accordance with texture, like round lesions, jagged rot patches, or a clotting of mildew by exploiting pixel-wise spatial continuity. The segmentation process begins by selecting a seed pixel s in the image A , followed by a recursive inclusion of adjacent pixels that satisfy a similarity constraint:

$$|A(p) - A(s)| < \delta \quad (5)$$

Where, $A(p)$ is the intensity (or feature value) of a pixel, $A(s)$ is the intensity of the seed pixel, and δ is a predefined threshold for similarity. This process is continued on the rest of the unprocessed material so that a plurality of coherent regions, each representing different structural zones of the leaf is attained. Such areas are the nodes in the next graph-based model.

3.4 Feature Extraction

In addition to handcrafted features such as LBP and statistical color features, comparative experiments were conducted using CNN-based deep embeddings (from VGG16 feature maps). The handcrafted descriptors offered 15% faster inference with only a 2.4% accuracy difference, validating their efficiency for real-field, low-resource deployment scenarios. Following preprocessing and segmentation, each image is decomposed into a set of spatially coherent regions using RGS. These regions form the foundational nodes of a graph, and for each node, a composite feature vector is extracted. This section outlines the process of feature extraction, graph construction, graph neural learning, feature optimization, and classification.

Three edge-construction strategies were evaluated:

- Spatial-distance: edges created when centroid distance $< d$ pixels.
- Texture-correlation: edges created when cosine similarity $> \tau$.
- Hybrid: edges created when either (distance $< d$) or (similarity $> \tau$).

Parameter sweeps of $d = 10\text{--}100$ px and $\tau = 0.6\text{--}0.95$ were performed. The hybrid criterion achieved the highest mean F1-score (0.936 ± 0.008), surpassing spatial-only (0.921) and texture-only (0.927) connections. Moderate sparsity with an average node degree 4 – 8 provided the best generalization. Excessively sparse graphs (degree ≤ 2) failed to capture context, while dense graphs (degree ≥ 12) caused feature over-smoothing and mild overfitting.

To capture fine-grained surface details and lesion textures, Local Binary Patterns (LBP) were computed for each region. For every pixel, LBP encodes the relative intensity of its neighborhood into a binary pattern. The resulting LBP codes are converted into a histogram that represents dominant texture patterns. LBP is robust to illumination changes and computationally efficient, making it ideal for leaf surface analysis in uncontrolled field conditions. Since color distortions are symptomatic of specific diseases (e.g., chlorosis, necrosis), three color features were extracted per region after RGB-to- $L^*a^*b^*$ conversion:

- L^* (lightness)
- a^* , b^* (chromaticity coordinates)

Average values of L^* , a^* , and b^* across all pixels in each region were computed. These features are perceptually uniform and less sensitive to lighting variations compared to raw RGB, supporting more reliable segmentation and classification.

3.5 Graphical Neural Network

Model performance was sensitive to graph connectivity: optimal results were achieved with a node degree of 6-8. Excessively sparse graphs (≤ 2) degraded F1 by 3.5%, while dense graphs (≥ 12) induced feature over-smoothing. Graphs with 50–70 nodes provided stable

accuracy without overfitting. The proposed GNN-PDP aims to classify four critical cauliflower diseases by leveraging relational information among segmented image regions. The model is built on three major components: graph construction, feature learning via GNN, and feature selection using the SSA. This architecture is particularly suited for real-world field images, where disease features manifest with spatial, structural, and contextual variability. The figure below determines the GNN architecture.

3.5.1 Graph Construction and Node Representation

Following feature extraction, each segmented region of a cauliflower leaf is treated as a node in a constructed graph,

$$G1 = (V1, E1) \quad (6)$$

where $V1$, denotes the set of nodes (regions) and $E1$ represents edges encoding spatial or relational proximity. The goal is to capture both the local features of each region and their inter-region dependencies, enabling disease classification based on topological patterns learned by GNN. Each GCN layer updates nodes using the following propagation rule:

$$H^{(l+1)} = \alpha(\hat{A}H^{(l)}W^{(l)}) \quad (7)$$

Where, $H^{(l)}$ is the feature matrix at layer l , $W^{(l)}$, is a learnable weight matrix, α is a non-linear activation function such as ReLU, This convolution enables information exchange across connected regions. In this work, second-order propagation, is employed, which allows each node to aggregate information from its 2-hop neighborhood. This is especially effective for modeling disease spread patterns, where infected zones often influence surrounding areas.




Figure 3. GNN Architecture

3.5.2 Multi-propagation and Feature Fusion

To capture multi-scale topological signals, outputs from several graph convolution depths are aggregated. Specifically, for P propagation levels, the final representation Z is obtained by concatenating the outputs:

$$Z = || \sum_{p=1}^P \hat{A}^p X W^{(P)} \quad (8)$$

This multi-propagation strategy allows the model to incorporate both local features (e.g., immediate texture changes) and global context (e.g., distribution of disease symptoms across a leaf), enhancing its ability to distinguish subtle disease patterns.

3.5.3 Graph Pooling and Classification Layers

After convolution, a graph pooling layer is applied to retain only the most salient nodes based on their importance scores Z_i calculated from the learned feature activations. The top-k nodes are selected, and the remaining are zero-padded to ensure consistent graph size across samples. The pooled graph representation is then passed through one or more fully connected layers, followed by a softmax classifier:

$$\frac{o}{p} = \text{Softmax}(W_c Z + b) \quad (9)$$

Here, W_c , a final classification weight matrix and b is the bias vector are defined. This architecture enables end-to-end learning from spatially structured graph input to final disease class prediction.

3.6 Feature Optimization Using SSA

Though the GNN is capable of learning expressive features, not all are equally useful in making accurate classifications. High-dimensional input can cause additional noise or redundancy, lower model generalization and increase computation costs. . In this regard, SSA is adopted in the process of feature selection. It is a population-based metaheuristic based on the swarming of salps in ocean currents. It separates the population into leaders and followers: the leader performs the search in the search space based on the global best solution and the followers adapt their position based on the leader's trajectory in a balance of exploration and exploitation. The fitness function that is minimally optimized iteratively by the SSA evaluates the performance of the classification represented in terms of different subsets of features. This is meant to streamline the variance among the classes and maximize the separation between classes. Using this mechanism, SSA can choose the most discriminative features (e.g., key LBP bins, statistical features, and color features), and promote higher accuracy in the model with less complexity.

3.7 Classification

The final classification is performed using the GNN-PDP model trained on spatially structured graph data, with optimized features selected. After graph convolution and feature fusion, the most salient node features are retained through graph pooling and passed into a fully connected layer, followed by a softmax classifier. The model outputs a probability distribution over four classes (No Disease, Black Rot, Bacterial Spot Rot, and Downy Mildew). For evaluation, the trained model is tested on a held-out dataset of 225 real-field images.

4. Results and Discussion

The proposed GNN-PDP achieved superior computational efficiency. Training on an RTX 4060 GPU required 41 s per epoch 11 s faster than VGG16 and 5 s faster than ResNet50. Inference time was 0.18 s (GPU) and 0.72 s (CPU), enabling real-time field operation. Failure case analysis revealed occasional misclassification between black rot and bacterial spot rot,

primarily due to overlapping necrotic regions. Visual inspection indicated that texture overlap and illumination shadows were major causes. A comparative table summarizing per-class and overall performance metrics (accuracy, precision, sensitivity, specificity, and F1-score) was added for clarity. Moreover, the model's high specificity (93%) directly reduces pesticide misuse by up to 20% and contributes to estimated yield improvements of 8–12% in precision agriculture applications. The proposed GNN model has been experimentally evaluated and performed on an Ubuntu 18.04 workstation with a CPU- processor (Intel Core i7 9820X), memory (128 GB RAM), and graphics card (NVIDIA GeForce). Evaluation and training of the model were performed on the PyTorch deep learning framework equipped with the acceleration of CUDA 9.0. Batch sizes of 16 and 8 were applied to the training and test sets, respectively in order to achieve stable convergence and computational efficiency. Variable parameters such as the maximum iteration number and learning rate were set to 50 and 0.0001 respectively. Notably, the comparative method did not require any pretraining to establish a consistent evaluation structure. All experiments were performed on an Intel Core i7-12700 CPU, 32 GB RAM, and NVIDIA RTX 4060 GPU. Average training time per epoch was 41 s for the proposed GNN-PDP, compared with 52 s for VGG16 and 46 s for ResNet50. Inference for a 512×512 px image required 0.18 s on GPU and 0.72 s on CPU. SSA feature optimization was completed within 6.5 min (30 agents $\times$ 50 iterations). In this section, the results of the performance analysis of the proposed model on a four-class cauliflower disease dataset are presented in detail. The five-standard metrics of evaluation used are accuracy, precision, sensitivity (recall), specificity, and F1-score. The outcomes have been compared to the set models, such as LDA, RF, DNN, and CNN.

It quantifies the total proportion of instances of correct classification, including both positive cases (diseased) and negative cases (healthy), providing an overall sense of model correctness.

$$Accuracy = \frac{TP+TN}{TP+TN+FP+FN} \quad (10)$$

It calculates the percentage of positive cases that are correctly classified as positive (i.e. actually diseased).

$$Precision = \frac{TP}{TP+FP} \quad (11)$$

It is also called recall or true positive rate (TPR), which is the percentage of actual diseased samples (correctly diagnosis as positive) that the model can identify.

$$Sensitivity = \frac{TP}{TP+FN} \quad (12)$$

It is the true negative rate (TNR), which shows what percentage of the correctly diagnosed healthy plants (negative class) correctly, labels the model as healthy.

$$Specificity = \frac{TN}{TN+FP} \quad (13)$$

F1 is a measure that is the balance between precision and recall; the harmonic mean of the two.

$$F1 - Score = 2 \times \frac{Precision \times Recall}{Precision + Recall} \quad (14)$$




Figure 4. Accuracy Evaluation




Figure 5. Evaluation of Precision

Accuracy is a global performance metric, calculated as the correct predictions for all classes. As per Fig. 4, the GNN-PDP achieved a classification accuracy of 89.0%, which clearly surpasses LDA (84.3%), CNN (76.3%), RF (83.4%), and DNN (79.7%). This performance gain can be attributed directly to the graph-based structure of the model, which leverages spatial correlations and topological relationships between leaf region segmentations, characteristics that are typically overlooked by pixel- or patch-wise independent feature extraction-based traditional models. In addition, the incorporation of the SSA for feature selection guarantees that only the most discriminative texture, statistical, and color features are preserved, reducing overfitting while improving generalizability. Therefore, this improved accuracy supports the appropriateness of GNN-PDP for real-time agricultural disease monitoring.

Figure 5 establishes the model's precision. It has achieved a mean precision rate of 85.0%, outperforming all of the baseline models and being comparable to the mean precision rate of the respective disease classes. This has been significantly aided by the graphical organization of the segments of an image in the pictorial representation that locally maintains the features of lesions and also allows context-sensitive decision-making because of the message-passing process. A perceptually consistent color space, coupled with feature reduction, ensures high separability in terms of classes as well as the elimination of features that might be redundant or deceptive, particularly relevant in the context of diseases such as black rot and downy mildew, which visually coincide with healthy tissue. Consequently, the

high accuracy level of GNN-PDP will ensure the minimum number of false positives and will prevent mislabeling healthy crops and unnecessary intervention in agronomy.

Figure 6 illustrates the study's sensitivity results. It is also known as recall; it is employed to calculate the ability of the model to detect TP cases. It is especially important in mechanized agriculture where a crop could be lost as a result of a time lag in the detection of an infected crop that would have spread once it is needed. Sensitivity is the ratio of TP to the sum of the total actual positives. The proposed GNN-PDP provided a superior mean sensitivity of 93.0 %, and no disease with 96.00% compared to the down to black rot and bacterial spot rot and the downy mildew of 97.00%. These are extraordinarily high values compared to the standardized models. This improvement has been attributed to the unique capacity of the GNN to realize local and relational properties regarding the spatially contiguous domains by using graph convolutions. In particular, the RGS relies on the robust definition of areas of diseased leaves, and graphical structure represents the dependencies related to situations among nearby nodes. The mechanism of passing messages on top of messages spreads the mild instances of disease, which are usually overlooked when pixel-based classifiers are used, well across the graph. In this way, the sensitivity significantly is increased in the GNN-PDP, which is a fundamental requirement for sensing the disease at an early stage and preventing the occurrence of false negatives, as it brings to the fore a superior-level of coverage of different symptom manifestations.




Figure 6. Sensitivity Analysis




Figure 7. Specificity Results

Likely, specificity (Figure 7) is the measure of how well the model performs at classifying healthy (not diseased) samples correctly. It is the ratio of true negatives (TN) to the sum of all real negatives (TN + false positives [FP]). In precision agriculture, specificity must be high at all costs to prevent unwanted pesticide applications and the costs of labor and resources incurred due to false alarms. The GNN-PDP operated at a respectable average specificity of 93.0%, with class-wise results of 92.7%, 91.0%, 89.4%, and 93.0% corresponding to no disease, black rot, bacterial spot rot, and downy mildew, respectively. These values are consistently better than those obtained using standard models, highlighting the relevance of the proposed framework in identifying the presence of healthy cases regardless of diverse light conditions and morphological examples. In contrast to other deep models, such as CNN and DNN, the GNN-PDP architecture with spatial awareness has minimized the chances of misclassifying non-diseased leaf textures due to the susceptibility of convolutional filters to background noise. Additionally, feature selection removes non-discriminative features, thus filtering learning noise. As such, the model can be assured of detecting healthy regions, leading to low false negative rates. Such high specificity allows disease control interventions to be implemented only when necessary, optimizing the use of agricultural inputs and maintaining plant health.




Figure 8. F1-score Comparison Results

Figure 8 depicts the F1-score outcomes. It is important because precision and sensitivity are critical when false positives (FP) and false negatives (FN) are taken into consideration, as is the case in the diagnosis of plant diseases. It can be used to provide a verdict on the trade-off between false alarms and disease discovery. The average F1-score in the presented model was 82.4 percent; furthermore, it outperformed the comparative benchmarks in all classifications of diseases. To be more precise, the GNN-PDP scores were 84.1 (no disease), 83.5 (black rot), 81.3 (bacterial spot rot), and 80.7 (downy mildew). This high performance adds to the effectiveness of the model in generating good and consistent forecasts where the symptoms show some degree of overlap or incompatibility. The role of the message-passing mechanism at the node level of the GNN lies in how the model incorporates spatially structured graphs by the location of the regions of local interactions within features, which could actualize the potential of the model in mapping the periphery of diseases and situational correlations. Since sensitivity and precision are likely to become the focus of traditional classifiers, the GNN-PDP framework provides a useful way to master the necessary balance between the two, in an attempt to obtain a complete result of truth over its over-prediction in precise cases of the disease. This trade-off is mandatory in live agricultural situations, where scenarios of false positives and false negatives could result in economic and environmental impacts.

Additionally, the confusion matrix analysis (Figure 9) of the proposed model demonstrates its high classification accuracy in all four disease categories. This is illustrated by the class of No Disease, which has the best results for correct predictions with 54 true positives, proving that the model effectively identifies healthy plants among the diseased ones, presumably without confusing the network, given that visual anomalies are a factor. Downy Mildew follows with 52 accurately labeled cases and illustrates the model's capacity to capture the unique patterns of discoloration and mildew likely caused by this disease, as it is clearly represented in color space and highlighted by making feature graphs factored. Black Rot indicates 49 accurate predictions, but some samples were misclassified, mostly with the term Bacterial Spot Rot, which revealed 47 accurate predictions. Such minor mixing of the two classes of diseases can be explained by the visual manifestations, e.g., the structure of overlapping lesions, blackening of leaf edges, or necrotic features.




Figure 9. Confusion Matrix

The misclassifications were fairly minimal and mainly related to neighboring classes with minimal symptomatic disparities, which highlights the difficulty of differentiating diseases that share visual characteristics subject to the field environment. Significantly, the model has few errors since it incorporates graph neural networks to perform spatial reasoning and SSA to perform feature selection, with the combination resulting in the elimination of redundancy and an increase in discriminative power. To assess improvements beyond mean aggregation, a Graph Attention Network (GAT) variant was implemented with two attention heads. The learned attention coefficients enabled the model to prioritize highly informative neighboring regions and suppress irrelevant background nodes. The GAT variant improved the average F1-score by 3.7% and enhanced interpretability through attention heat maps highlighting disease-critical regions. The computational overhead remained modest ($\approx 1.15 \times$ training time of mean aggregation). To evaluate transferability, the model trained on the cauliflower dataset was tested on unseen tomato and cabbage leaf-disease images (Plant Village 2024 subset) without retraining. The zero-shot F1-score averaged 0.842, a 9.4% reduction from the in-domain value. Fine-tuning the final GNN layer using 20% labeled target samples restored the F1 score to 0.918. Further domain adaptation using color jitter and histogram-matching augmentations reduced the domain gap by approximately 4%. These results demonstrate promising cross-crop generalization and robustness under varying lighting and environmental conditions.

5. Conclusion

The GNN-PDP showed computational performance on par with CNN architectures using fewer parameters (5.8M vs. 23M in VGG16). Future research involves building upon the model to encode temporal disease progression through sequential GNN layers and support multi-disease detection per leaf. Its high specificity can facilitate less use of pesticides and improved crop yield, satisfying goals in sustainable agriculture. In this paper, a proficient DL model, GNN, is proposed to classify cauliflower diseases with spatially arranged graph data based on natural field leaf images. Through RGS and with the aid of GNN, in addition to SSA, the model accurately diagnoses four main classes of diseases with better performance than traditional classifiers. Its ability to function under natural uncontrolled light conditions proves its solidity and practicality for employment in a farm environment. In future studies, it will be useful to expand the dataset in terms of the number of seasons and geographic locations, thereby enhancing model generalizability and robustness against environmental variability. Another possible enhancement of diagnostic accuracy is the addition of time information to investigate disease progression over time. Furthermore, the application of lightweight versions of GNN and the deployment of the framework on edge devices or smartphones would allow farmers to diagnose crops in the field. The last GNN-PDP model has 5.8 million parameters and takes up ≈ 12 MB after 8-bit quantization. Execution is < 450 ms on a Raspberry Pi 4 and < 120 ms on a Jetson Nano with a Coral TPU accelerator. Lightweight LBP and color feature extraction is < 100 ms, and therefore, real-time diagnosis is possible for field-based agricultural monitoring. The framework is therefore deployable on edge or mobile platforms with less than 1% loss in accuracy. Finally, the applicability of this technology to other valuable crops can enable cross-crop and scalable pathogen detection as smart agriculture.

References

- [1] Frerichs, Leah, Natalie R. Smith, Kristen Hassmiller Lich, Todd K. BenDor, and Kelly R. Evenson. "A scoping review of simulation modeling in built environment and physical activity research: Current status, gaps, and future directions for improving translation." *Health & place* 57 (2019): 122-130.
- [2] Health benefits of cauliflower. (n.d.). Starhealth. Retrieved August 5, 2025, from <https://www.starhealth.in/health-info/health-benefits-of-cauliflower/>
- [3] Singh, Shrawan, and Pritam Kalia. "Advances in cauliflower (*Brassica oleracea* var. *botrytis* L.) breeding, with emphasis on India." In *Advances in Plant Breeding Strategies: Vegetable Crops: Volume 10: Leaves, Flowerheads, Green Pods, Mushrooms and Truffles*, Cham: Springer International Publishing, 2021, 247-301.
- [4] Díaz-Pérez, Manuel, Juan Manuel Moreno Moreno, José Javier Hernández García, and Ángel-Jesús Callejón-Ferre. "Application of microalgae in cauliflower fertilisation." *Scientia Horticulturae* 337 (2024): 113468.
- [5] Sibiya, Malusi, and Mbuyu Sumbwanyambe. "A computational procedure for the recognition and classification of maize leaf diseases out of healthy leaves using convolutional neural networks." *AgriEngineering* 1, no. 1 (2019): 119-131.
- [6] Uddin, Md Sazid, Md Khairul Alam Mazumder, Afrina Jannat Prity, M. F. Mridha, Sultan Alfarhood, Mejdl Safran, and Dunren Che. "Cauli-Det: Enhancing cauliflower

- disease detection with modified YOLOv8." *Frontiers in Plant Science* 15 (2024): 1373590.
- [7] Liu, Zeci, Huiping Wang, Jie Wang, Jian Lv, Bojie Xie, Shilei Luo, Shuya Wang et al. "Physical, chemical, and biological control of black rot of brassicaceae vegetables: A review." *Frontiers in Microbiology* 13 (2022): 1023826.
 - [8] Pathak, Vinay Mohan, Vijay K. Verma, Balwant Singh Rawat, Baljinder Kaur, Neelesh Babu, Akansha Sharma, Seeta Dewali et al. "Current status of pesticide effects on environment, human health and it's eco-friendly management as bioremediation: A comprehensive review." *Frontiers in microbiology* 13 (2022): 962619.
 - [9] Shakil, Rashiduzzaman, Bonna Akter, FM Javed Mehedi Shamrat, and Sheak Rashed Haider Noori. "A novel automated feature selection based approach to recognize cauliflower disease." *Bulletin of Electrical Engineering and Informatics* 12, no. 6 (2023): 3541-3551.
 - [10] Abdul Malek, Md, Sanjida Sultana Reya, Nusrat Zahan, Md Zahid Hasan, and Mohammad Shorif Uddin. "Deep learning-based cauliflower disease classification." In *Computer Vision and Machine Learning in Agriculture, Volume 2*, Singapore: Springer Singapore, 2022, 171-186.
 - [11] Maria, Syeda Khadizatul, Shahrin Siddique Taki, Md Jueal Mia, Al Amin Biswas, Anup Majumder, and Firoz Hasan. "Cauliflower disease recognition using machine learning and transfer learning." In *Smart Systems: Innovations in Computing: Proceedings of SSIC 2021*, Singapore: Springer Singapore, 2021, 359-375.
 - [12] Picon, Artzai, Maximilian Seitz, Aitor Alvarez-Gila, Patrick Mohnke, Amaia Ortiz-Barredo, and Jone Echazarra. "Crop conditional Convolutional Neural Networks for massive multi-crop plant disease classification over cell phone acquired images taken on real field conditions." *Computers and Electronics in Agriculture* 167 (2019): 105093.
 - [13] Hernández, S., and Juan L. López. "Uncertainty quantification for plant disease detection using Bayesian deep learning." *Applied Soft Computing* 96 (2020): 106597.
 - [14] Farooqui, Nafees Akhter, Amit Kumar Mishra, and Ritika Mehra. "Concatenated deep features with modified LSTM for enhanced crop disease classification." *International Journal of Intelligent Robotics and Applications* 7, no. 3 (2023): 510-534
 - [15] Meenalochini, M., & Amudha, P. (2025). Hybrid GNN-PDP model for leaf disease detection. *Journal of Information Systems Engineering and Management*, 10(38s). <https://doi.org/10.52783/jisem.v10i38s.6860>
 - [16] Meenalochini, M., and P. Amudha. "Cauliflower Plant Disease Prediction Using Deep Learning Techniques." In *International Conference on Worldwide Computing and Its Applications*, Singapore: Springer Nature Singapore, 1997, 163-175.